

Amazon Quick Suite



Amazon Quick Suite: User Guide

Table of Contents

What is Amazon Quick Suite?	1
Benefits of Quick Suite	1
Pricing and availability	2
Accessing Quick Suite	2
Related services	3
Quick Suite user types	3
Are you a first-time Quick Suite user?	4
Understand key Amazon Quick Suite concepts	5
Data preparation	5
SPICE	5
Data analysis	5
Data visualization	5
Machine learning	6
Sheet	6
Dashboard	6
How Amazon Quick Suite works	7
Business intelligence (BI) with Amazon Quick Sight	9
Perform deep research with Amazon Quick Research	9
Automate workflows with Amazon Quick Flows and Amazon Quick Automate	10
Create and use Amazon Quick Suite custom agents	10
Share data and collaborate with Amazon Quick Suite spaces and folders	11
Governance and administration for Amazon Quick Suite	11
Setting up and signing in	13
Complete initial configuration tasks	13
Sign up for an Amazon Web Services account	13
Secure IAM users	14
Signing up for a subscription	14
Non-admin sign-up	15
Signing in to Quick Suite	16
Post-setup configuration	17
How to sign in to Quick Suite	17
Getting started with Amazon Quick Suite	20
Using the Amazon Quick Suite console	20
Using the Amazon Quick Suite menu and landing page	20

Searching Amazon Quick Suite	24
Choosing a language in Amazon Quick Suite	24
Getting started with Amazon Quick Sight	26
Quick start: Create an analysis using sample data	26
Create a dashboard using sample data	27
Sample data	33
Administering Amazon Quick Suite	34
Getting started	34
User types	35
Amazon Quick Suite editions	37
Supported browsers	40
Setting up	41
(Optional) Create service control policies (SCP)	42
Configure identity and access management	45
Allowlisting domains	114
Configuring VPC connections	115
VPC terminology	117
Supported VPC data sources	118
Setting up a VPC to use with Amazon Quick Suite	119
Finding information to connect to a VPC	131
Configuring Amazon Quick Sight access to Amazon data	135
Required permissions	136
Network and database configuration requirements	137
Allowing autodiscovery of Amazon resources	143
Using data in Amazon	144
Exploring your Amazon data in Amazon Quick Suite	177
Amazon service action connectors	179
Managing account	183
Manage account settings	184
Manage assets	187
Manage Amazon Q	188
Managing subscriptions	190
Configure subscriptions	190
Configure SPICE capacity	196
Managing index capacity	202
Managing identities	206

Managing user access	207
Creating and managing groups	215
Managing security	219
Managing domains	219
Multitenancy and namespaces	222
Managing mobile settings	227
Managing IP/VPC restrictions	228
Managing VPC	230
Managing Amazon KMS keys	242
Managing permissions	250
Amazon resources	250
Default access policy	259
IAM policy assignments	259
Extension access	261
Custom permissions	282
Managing customizations	296
Amazon Quick Suite brand customization	297
Managing chat customization	304
Amazon Quick Suite usage metrics	308
Analytics access and permissions	309
Analytics dashboard overview	309
Accessing analytics	312
Understanding metrics	313
Analytics limitations	314
Amazon security	314
Data protection	315
Incident response, logging, and monitoring in Amazon Quick Suite	320
Compliance validation	345
Resilience	346
Infrastructure security	347
Best practices	348
Amazon managed policies	348
Ask questions and explore data using Amazon Quick Suite chat	356
Welcome message	356
Suggested prompts	357
Expand and collapse modes	357

Agent picker	357
Scoping response generation to specific data	358
Web search	359
Upload files and chat	359
Actions	360
Flows	360
Source citations	361
Copy responses	361
Provide feedback	361
Artifact panel	362
Conversation management	362
Response events	362
Contextual awareness	363
Create and deploy AI assistants using Amazon Quick Suite chat agents	364
Amazon Quick Suite chat agent types	365
User interaction with chat agents	366
User permissions for chat agents	367
Custom permissions for chat agents	368
System chat agent	370
Custom chat agents	370
Create custom chat agent	371
Share custom chat agent	377
Update custom chat agent	377
Manager access to a chat agent	378
View chat agents	379
View chat agent details	379
Duplicate custom chat agents	380
Delete chat agents	381
Custom chat agent context configuration	381
Agent context types	381
Use a chat agent	385
Organize and collaborate using Amazon Quick Suite spaces	386
User interactions with spaces	387
User permissions for spaces	387
File upload status	388
Limitations of spaces	389

Creating spaces	389
Managing spaces	390
Viewing spaces	391
Editing spaces	391
Adding space knowledge	392
Viewing content in a space	392
Using spaces	393
Sharing spaces	393
Deleting spaces	394
Removing knowledge	395
Adding external data sources	395
Adding application actions	396
Visualize, analyze, and share data with Amazon Quick Sight	397
Connecting to data	397
Supported data sources	400
Connect to your data	405
Data source quotas	451
Supported data types and values	453
Working with datasets	466
Working with data sources	531
Refreshing data	536
Importing data into SPICE	538
Refreshing SPICE data	540
Using SPICE data in an analysis	545
View SPICE ingestion history	546
Troubleshooting skipped row errors	547
SPICE ingestion error codes	549
Updating files in a dataset	554
Preparing data	555
Describing data	556
Choosing file upload settings	557
Preparing data fields	558
Adding calculations	570
Previewing tables in a dataset	778
Joining data	778
Filtering data	786

Using SQL to customize data	818
Adding geospatial data	820
Using unsupported or custom dates	878
Adding a unique key to a Quick Sight dataset	879
Integrating SageMaker AI models	880
Preparing dataset examples	890
Analyses and reports: Visualizing data	897
Working with an analysis	897
Adding and managing sheets	911
Working with interactive sheets in Amazon Quick Sight	913
Working with paginated reports in Amazon Quick Sight	952
Working with items on sheets	968
Using themes in Amazon Quick Sight	1158
Keyboard shortcuts	1165
Dashboards and reports: Sharing data	1165
Publishing dashboards	1166
Sharing dashboards	1171
Use action connectors in dashboards	1185
Share your view of a dashboard	1188
Sending reports	1188
Subscribing to reports	1193
Threshold alerts	1194
Print a dashboard or analysis	1203
Exporting as PDFs	1204
PDF Error codes	1204
Organizing assets into folders	1206
Exploring dashboards	1219
Interacting with dashboards	1221
Interacting with paginated reports	1226
Subscribe to emails and alerts	1228
Reader generated reports	1229
Bookmarks	1233
Gaining insights with ML	1236
Understanding the ML algorithm	1238
Dataset requirements	1242
Adding insights	1243

Autonarratives	1246
Detecting outliers	1279
ML-powered forecasts and what-ifs	1291
Troubleshooting	1294
Resolving Amazon Quick Sight issues and error messages	1294
Athena issues	1294
Data source connectivity issues	1305
Login issues	1317
Visual issues	1324
Developing with Amazon Quick Sight	1327
Required knowledge	1328
Available API operations for Amazon Quick Sight	1328
Terminology and concepts	1330
Using the developer portal	1331
Developing with the Amazon Quick Sight APIs	1332
Events integration	1340
Embedded analytics	1362
Using Amazon Quick Research	1598
What you can do with Quick Research	1598
Starting new research	1598
Enter research objective	1599
Select research materials	1600
Review and finalize research plan	1604
View research report	1605
Updating research	1609
Making updates	1609
Research versions	1611
Summarizing the research	1612
Custom summary	1613
Executive summary	1613
General share out	1613
Sharing the research	1613
Download the research	1613
Sharing the research with stakeholders	1614
Admin controls	1614
Turn on/off web search	1615

Status definitions	1615
Using Amazon Quick Suite automations	1616
Types of automations	1616
Getting started	1616
Using Amazon Quick Flows	1616
What are Flows?	1617
Why Flows?	1618
Quick Flows features and capabilities	1620
Collaboration and governance	1621
Getting started	1622
Creating and managing	1630
Components and features	1659
Running and integrating	1671
Managing Quick Flows	1695
Quick Flows limits	1698
Using Amazon Quick Automate	1700
Overview	1700
Key concepts	1701
Getting started	1702
Building automations	1707
Automation actions	1714
Testing automations	1741
Deploying automations	1745
Monitoring performance	1748
Human-in-the-loop task center	1752
Integrations	1756
Supported integrations	1757
Integration capability definitions	1759
Authentication method definitions	1759
Set up integrations in the console	1760
Choose integration options	1760
View setup process examples	1760
Integration management options	1762
Integration workflows	1762
Creating a knowledge base from scratch	1762
Creating an action connector	1763

Creating a knowledge base integration from an existing data access integration	1763
Managing existing integrations	1764
Knowledge bases	1764
How knowledge bases work	1765
General workflow	1766
Common configuration settings	1766
Troubleshooting knowledge bases	1768
Action connectors	1771
How action connectors work	1771
Types of actions	1772
Bounded and unbounded agents	1774
Prerequisites	1775
Supported action connector types and available actions	1775
Action connector APIs	1779
Authentication methods	1788
Action execution methods	1790
Monitoring and maintenance	1791
Best practices	1793
Troubleshooting	1795
Integration-specific guides	1797
S3 integration	1797
Asana integration	1808
BambooHR integration	1809
Confluence integration	1813
Google Drive integration	1817
Jira integration	1821
MCP integration	1825
OneDrive integration	1831
Outlook action connector	1836
SharePoint integration	1839
Microsoft Teams integration	1847
OpenAPI integration	1850
PagerDuty integration	1857
REST API integration	1858
Salesforce integration	1862
SAP integrations	1867

ServiceNow integration	1878
Slack integration	1881
Smartsheet integration	1884
Web Crawler integration	1885
Zendesk integration	1895
Support escalation path	1897
Creating action connectors using the admin console	1897
Prerequisites	1898
Accessing the admin console	1898
Creating a new action connector	1898
Using the new connector in Amazon Quick Suite	1899
Next steps	1900
Bring Your Own Amazon Q Business Index	1900
Overview of Amazon Q Business indexes in Amazon Quick Suite	1900
Prerequisites	1901
Setting up permissions	1902
Creating knowledge bases from Amazon Q Business indexes	1902
Sharing Amazon Q Business index knowledge bases	1903
Using Amazon Q Business index knowledge bases	1903
Limitations	1905
Billing	1905
User types and capabilities	1905
Common use cases	1906
Interacting with your Amazon Q Business indexes	1907
Cross-region inference in Amazon Quick Suite integrations	1907
Supported regions for Amazon Quick Suite cross-region inference	1908
Connect Quick Suite to existing tools using extensions	1910
User interactions with extensions	1911
User capabilities and permissions	1911
Administrator responsibilities	1913
Author capabilities	1914
End user access	1915
Extension access troubleshooting	1917
Amazon Quick Suite browser extension	1917
Supported browsers	1918
Browser extension benefits	1918

Browser extension use cases	1918
Install and configure	1920
Organizational deployment	1921
Browser extension usage guidelines	1922
Browser extension known limitations	1923
Amazon Quick Suite Microsoft Outlook extension	1923
Microsoft Outlook extension author guide	1924
Microsoft Outlook extension user guide	1928
Amazon Quick Suite Slack extension	1931
Slack extension author guide	1932
Slack extension user guide	1935
Amazon Quick Suite Microsoft Word extension	1940
Microsoft Word extension author guide	1941
Microsoft Word extension user guide	1944
Amazon Quick Suite Microsoft Teams extension	1948
Microsoft Teams extension author guide	1949
Microsoft Teams extension user guide	1953
Document history	1957
Previous updates	2046
Attributions	2059

What is Amazon Quick Suite?

Amazon Quick Suite is a comprehensive, generative AI-powered business intelligence platform that makes it easy to analyze data, create visualizations, automate workflows, and collaborate across your organization. The service combines traditional business intelligence capabilities with modern AI assistance, requiring no machine learning expertise to use. You can connect to diverse data sources, create interactive dashboards, build intelligent automations, and get immediate insights through natural language conversations with AI agents.

Quick Suite includes five integrated capabilities that work together: Amazon Quick Sight for data visualization, Amazon Quick Flows for workflow automation, Amazon Quick Automate for process optimization, Amazon Quick Index for data discovery, and Amazon Quick Research for comprehensive analysis. The platform extends beyond traditional BI by bringing AI assistance directly into your existing tools through extensions for browsers, Slack, and Microsoft Office applications.

Topics

- [Benefits of Quick Suite](#)
- [Pricing and availability](#)
- [Accessing Quick Suite](#)
- [Related services](#)
- [Quick Suite user types](#)
- [Are you a first-time Quick Suite user?](#)

Benefits of Quick Suite

Some of the benefits of Quick Suite include:

AI-powered analysis and visualization

Quick Suite analyzes natural language queries across enterprise content and creates interactive dashboards from multiple data sources. You can combine diverse data types including Amazon data, third-party data, big data, spreadsheets, SaaS data, and B2B data within a unified analytical environment. Custom AI agents provide domain-specific expertise and automate analytical tasks through conversational interfaces.

Simple to deploy and manage

Quick Suite provides all machine learning infrastructure, models, and pre-built connectors with SPICE in-memory engine for analytics. As a fully managed service, it requires no infrastructure deployment or management, allowing you to focus on gaining insights from your data.

Enterprise-grade security and governance

Quick Suite supports comprehensive security through granular permissions and row-level security controls. You can access the system through federation and single sign-on capabilities. All data is protected with secure encryption at rest and in transit. The platform integrates with IAM Identity Center while ensuring all responses and visualizations respect your permissions.

Collaboration and workflow automation

Quick Suite enables seamless sharing of dashboards and insights across your organization. You can create embedded analytics for applications and websites, configure AI responses using enterprise data sources, and automate routine tasks to streamline workflows. Data and resources can be organized into dedicated project spaces for efficient knowledge sharing and task completion.

Pricing and availability

Quick Suite charges you both for user subscriptions, and for data capacity. For information about what's included in the tiers of user subscriptions and index capacity, see [Quick Suite pricing](#).

For a list of regions where Quick Suite is currently available, see [Quick Suite endpoints and quotas](#) and [Amazon Regions, websites, IP address ranges, and endpoints](#).

Accessing Quick Suite

You can access Quick Suite in the following ways in the Amazon Web Services Regions that it's available in:

[Amazon Web Services Management Console](#)

You can use the Amazon Web Services Management Console—a browser-based interface to interact with Amazon Web Services services—to access the Quick Suite console and resources. You can perform most Quick Suite tasks using the Quick Suite console.

[Amazon Quick Suite API](#)

To access Amazon Quick Sight and some Quick Suite programmatically, you can use the Amazon Q API. For more information, see the [Quick Suite API Reference](#).

[Amazon Command Line Interface](#)

The Amazon Command Line Interface (Amazon CLI) is an open source tool. You can use the Amazon CLI to interact with Amazon Web Services services using commands in your command line shell. If you want to build task-based scripts, using the command line can be faster and more convenient than using the console.

[SDKs](#)

Amazon SDKs provide language APIs for Amazon Web Services services to use programmatically.

Related services

The following are some of the other Amazon Web Services services that Quick Suite integrates with:

[Amazon Q Business](#)

Amazon Q Business is a fully managed, generative-AI powered assistant that you can configure to answer questions, provide summaries, generate content, and complete tasks based on your enterprise data. If you're already an Amazon Q Business user, you can connect your Amazon Q Business application to Quick Suite.

Quick Suite user types

There are three user personas that Quick Suite supports—readers, authors, and admins. Each persona has both standard and Pro subscription types, with Pro subscriptions providing access to advanced Amazon Quick Suite tools. For detailed information about subscription types and pricing, see [Quick Suite pricing](#).

Readers

You can use Quick Suite to access company data and find answers through chat interactions with AI agents. You can upload files, run automations, create visualizations, and share spaces

via direct links. While you cannot create datasets or agents as a reader, you are the primary consumer of prepared analytics and AI chat tools.

Authors

As a domain expert, you can build and manage Quick Suite resources. In addition to reader capabilities, you can create datasets, dashboards, automations, and agents. You have expanded sharing permissions for spaces and can use AI to create visualizations, enabling you to build the data infrastructure that supports readers.

Administrators

You can oversee the Quick Suite system by managing user access, monitoring costs, and maintaining data sources. You have full reader and author capabilities but focus primarily on system administration to ensure efficient and secure operations for all users.

Are you a first-time Quick Suite user?

If you're a first-time admin user of Quick Suite, we recommend that you read the following sections in order:

[How it works](#)

Introduces Quick Suite components and describes how they work.

[Key concepts](#)

Explains key concepts and important Quick Suite terminology.

[Setting up](#)

Outlines how to set up Quick Suite so that you can begin using it.

Amazon Quick Suite concepts

The following are some important terms that you will encounter in this guide.

Data preparation

Data *preparation* is the process of transforming data for use in an analysis. This includes making changes like the following:

- Filtering out data so that you can focus on what's important to you.
- Renaming fields to make them easier to read.
- Changing data types so that they are more useful.
- Adding calculated fields to enhance analysis.
- Creating SQL queries to refine data.

SPICE

SPICE (Super-fast, Parallel, In-memory Calculation Engine) is the robust in-memory engine that Quick Suite uses. SPICE is engineered to rapidly perform advanced calculations and serve data. The storage and processing capacity available in SPICE speeds up the analytical queries that you run against your imported data. By using SPICE, you save time because you don't need to retrieve the data every time that you change an analysis or update a visual.

Data analysis

A data *analysis* is the basic workspace for creating data visualizations, which are graphical representations of your data. Each analysis contains a collection of visualizations that you arrange and customize.

Data visualization

A data *visualization*, also known as a *visual*, is a graphical representation of data. There are many types of visualizations, including diagrams, charts, graphs, and tables. All visuals begin in AutoGraph mode, which automatically selects the best type of visualization for the fields that you

select. You can also take control and choose your own visuals. You can enhance your analytics by applying filters, changing colors, adding parameter controls, custom click actions, and more.

Machine learning

Machine learning (ML) *Insights* propose narrative add-ons that are based on an evaluation of your data. You can choose one from the list, for example forecasting or anomaly (outlier) detection. Or you can create your own. You can combine insight calculations, narrative text, colors, images, and conditions that you define.

Sheet

A *sheet* is a page that displays a set of visualizations and insights. You can imagine this as a sheet from a newspaper, except that it's filled with charts, graphs, tables, and insights. You can add more sheets, and make them work separately or together in your analysis.

Dashboard

A *dashboard* is the published version of an analysis. You can share with other users of Amazon Quick Suite for reporting purposes. You specify who has access and what they can do with the dashboard.

How Amazon Quick Suite works

Amazon Quick Suite provides five integrated API sets for comprehensive business intelligence and analytics:

- Amazon Quick Sight for interactive data visualization and business intelligence
- Amazon Quick Flows for intelligent workflow automation
- Amazon Quick Automate for streamlined business process automation
- Amazon Quick Index for data discovery and cataloging
- Amazon Quick Research for comprehensive data analysis

Data visualization and business intelligence

With Amazon Quick Sight, your applications can:

- Connect to diverse data sources including databases, data warehouses, and cloud services
- Create interactive dashboards and visualizations with automatic chart suggestions
- Perform data preparation and transformation
- Generate machine learning insights including forecasting and anomaly detection
- Embed analytics in custom applications and websites
- Share dashboards and reports across your organization

Use cases include self-service business intelligence, executive dashboards, and embedded analytics.

Intelligent workflow automation

With Amazon Quick Flows and Amazon Quick Automate, your applications can:

- Build interactive workflows that combine AI processing with structured automation
- Automate common business processes and repetitive tasks
- Integrate with external systems through action connectors
- Process natural language inputs to trigger automated workflows
- Orchestrate complex multi-step business processes

Use cases include document processing, approval workflows, and system integration automation.

AI-powered analysis and collaboration

With Quick Suite AI capabilities, your applications can:

- Create custom AI agents configured for specific organizational needs
- Analyze natural language queries across enterprise data sources
- Generate permissions-aware responses with citations and source tracing
- Organize team knowledge in collaborative spaces
- Conduct comprehensive research across large document collections
- Integrate AI assistance into existing tools through extensions

Use cases include conversational analytics, knowledge management, and research automation.

Key features

- SPICE in-memory engine for high-performance analytics
- Natural language processing for conversational data interaction
- Enterprise-grade security with granular permissions and row-level security
- Extensive third-party integrations with 25+ applications
- Scalable architecture supporting hundreds of thousands of users
- Comprehensive APIs and SDKs for programmatic access

Quick Suite lets you enhance the accuracy and relevance of AI responses by connecting custom knowledge bases and configuring agents with domain-specific expertise. For example, you can create specialized agents for financial analysis, customer support, or technical documentation that understand your organization's specific context and terminology.

The following sections cover the types of analysis that Quick Suite provides and an overview of Amazon Quick Sight, automation, and AI operations. Also covered are the differences between interactive and programmatic operations.

To demo the Quick Suite capabilities, you can see [Getting started with Amazon Quick Suite](#), which covers getting started with sample data and creating your first analysis.

Topics

- [Business intelligence \(BI\) with Amazon Quick Sight](#)

- [Perform deep research with Amazon Quick Research](#)
- [Automate workflows with Amazon Quick Flows and Amazon Quick Automate](#)
- [Create and use Amazon Quick Suite custom agents](#)
- [Share data and collaborate with Amazon Quick Suite spaces and folders](#)
- [Governance and administration for Amazon Quick Suite](#)

Business intelligence (BI) with Amazon Quick Sight

Using Amazon Quick Sight, you can access data and prepare it for use in reporting. It saves your prepared data either in SPICE memory or as a direct query. You can use a variety of data sources for analysis. When you create an analysis, the typical workflow has the following steps:

1. Create a new analysis.
2. Add new or existing datasets.
3. Choose fields to create the first chart. Quick Sight automatically suggests the best visualization.
4. Add more charts, tables, or insights to the analysis. Resize and rearrange them on one or more sheets. Use extended features to add variables, custom controls, colors, additional pages (called sheets), and more.
5. Publish the analysis as a dashboard to share it with other people.

Perform deep research with Amazon Quick Research

Amazon Quick Research capabilities in Quick Suite enable you to conduct comprehensive analysis across large volumes of data and documents using AI-powered tools. These features help you uncover insights, patterns, and relationships that might not be immediately apparent through traditional analysis methods.

Deep research functionality includes:

- Analyzing hundreds of documents, reports, or datasets simultaneously to identify trends and patterns
- Using natural language queries to explore complex data relationships across multiple sources
- Generating comprehensive summaries and insights from large document collections
- Leveraging AI to surface relevant information and connections that support decision-making

This capability is particularly valuable for tasks such as market research, competitive analysis, regulatory compliance review, and strategic planning where comprehensive data analysis is essential.

Automate workflows with Amazon Quick Flows and Amazon Quick Automate

Quick Suite automations provide intelligent workflow capabilities that combine generative AI with automated actions to streamline business processes and enhance productivity. These tools help organizations create, manage, and execute workflows that can understand natural language inputs and perform complex operations across various systems.

Automation capabilities include:

- **Amazon Quick Flows** - Interactive workflows that combine AI-powered processing with structured automation steps for complex business processes
- **Amazon Quick Automate** - Streamlined automation tools for common business processes and repetitive tasks
- Integration with external systems through action connectors to extend automation capabilities
- Natural language interfaces that allow users to trigger and interact with automated workflows

These automation tools enable organizations to reduce manual effort, improve consistency, and scale their analytical and operational processes while maintaining human oversight and control.

Create and use Amazon Quick Suite custom agents

Custom agents in Quick Suite are AI-powered assistants that you can configure to help users explore data, analyze information, and perform specific tasks within your organization's context. These agents combine natural language processing with access to your data sources, dashboards, and business knowledge to provide contextual assistance.

You can create custom agents to:

- Provide domain-specific expertise by connecting agents to relevant datasets, topics, and knowledge bases
- Automate common analytical tasks and workflows through conversational interfaces

- Enable self-service analytics for business users who prefer natural language queries over traditional BI tools
- Integrate with external systems and actions to perform tasks beyond data analysis

Custom agents can be configured with specific goals, knowledge sources, and connected tools to provide tailored assistance for different teams, use cases, or business domains within your organization.

Share data and collaborate with Amazon Quick Suite spaces and folders

Spaces and folders in Quick Suite provide organizational structures that enable teams to collaborate effectively on data analysis and share insights across your organization. Spaces serve as collaborative workspaces that aggregate data, dashboards, and knowledge for specific teams or domains, while folders provide hierarchical organization for your analytics assets.

Using spaces, you can:

- Create team-specific knowledge bases that combine datasets, dashboards, topics, and uploaded files
- Enable contextual AI conversations by grounding agent responses with relevant team data
- Reduce data silos by allowing multiple team members to contribute and access shared knowledge
- Maintain proper access controls while facilitating collaboration across teams

Folders complement spaces by providing structured organization for your analytics assets, with support for both personal and shared folders. Shared restricted folders ensure that sensitive assets remain within designated boundaries while still enabling collaboration.

Governance and administration for Amazon Quick Suite

Governance and administration features in Quick Suite provide comprehensive tools for managing users, controlling access, ensuring security, and maintaining compliance across your analytics environment. These capabilities enable administrators to establish proper controls while empowering users to work effectively with data.

Key governance and administration capabilities include:

- Identity and access management through integration with IAM Identity Center, IAM, and other identity providers
- Row-level security and column-level security to control data access at granular levels
- VPC connectivity and network security controls for secure data access
- Audit logging and monitoring capabilities to track user activities and system usage
- Resource management including SPICE capacity, user subscriptions, and feature access controls
- Brand customization and organizational settings to align with corporate standards

These features ensure that your Quick Suite deployment meets organizational security requirements while providing the flexibility needed for effective data analysis and collaboration.

Setting up and signing into Amazon Quick Suite

This section includes the essential setup tasks, such as signing up for an Amazon Web Services account, creating an administrative user, and subscribing to the Amazon Quick Suite service.

Note

For more information on setting up your Amazon Quick Suite account, including configuring identity and access management, and controlling sign-up for users, see [Administering Amazon Quick Suite](#).

Topics

- [Complete initial configuration tasks](#)
- [Signing up for an Amazon Quick Suite subscription](#)
- [Signing in to Amazon Quick Suite](#)

Complete initial configuration tasks

To use Amazon Quick Suite you must first complete the following tasks:

Topics

- [Sign up for an Amazon Web Services account](#)
- [Secure IAM users](#)

Sign up for an Amazon Web Services account

If you do not have an Amazon Web Services account, use the following procedure to create one.

To sign up for Amazon Web Services

1. Open <http://www.amazonaws.cn/> and choose **Sign Up**.
2. Follow the on-screen instructions.

Amazon sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to <http://www.amazonaws.cn/> and choosing **My Account**.

Secure IAM users

After you sign up for an Amazon Web Services account, safeguard your administrative user by turning on multi-factor authentication (MFA). For instructions, see [Enable a virtual MFA device for an IAM user \(console\)](#) in the *IAM User Guide*.

To give other users access to your Amazon Web Services account resources, create IAM users. To secure your IAM users, turn on MFA and only give the IAM users the permissions needed to perform their tasks.

For more information about creating and securing IAM users, see the following topics in the *IAM User Guide*:

- [Creating an IAM user in your Amazon Web Services account](#)
- [Access management for Amazon resources](#)
- [Example IAM identity-based policies](#)

Signing up for an Amazon Quick Suite subscription

When you first sign up for Amazon Quick Suite, you get a free trial subscription for four users for 30 days. During the process of signing up, you may set options for your identity provider.

Before you begin, make sure that you can connect to an existing Amazon Web Services account. If you don't have an Amazon Web Services account, see [Complete initial configuration tasks](#). The person who signs up for Quick Suite needs to have the correct Amazon Identity and Access Management (IAM) permissions. For more information, see [IAM policy examples for Quick Suite](#).

To test your permissions, you can use the IAM policy simulator; for more information, see [Testing IAM policies with the IAM policy simulator](#). Also, check whether your Amazon Web Services account is part of an organization based on the Amazon Organizations service. If so and you sign in as an IAM user, make sure that you didn't inherit any IAM permissions that deny access to the required permissions. For more information on Organizations, see [What is Amazon Organizations?](#)

Note

Your data is encrypted by default using Amazon-managed keys. Admins can adjust settings for custom encryption in the admin portal after signing up.

To subscribe to Quick Suite

1. Sign in to your Amazon Web Services account and open Quick Suite from the Amazon Web Services Management Console. You can find it under **Analytics** or by searching for *Quick Suite*.

Your Amazon Web Services account number is displayed for verification purposes.

2. Enter a unique account name for Quick Suite.
 - Enter a notification email address for the Quick Suite account owner or group. This email address receives service and usage notifications.
3. Choose the Amazon Web Services Region that you want to use for your initial data storage capacity, called SPICE.
4. Choose an authentication method that you want to connect to Quick Suite with. Select from one of the following:
 - **(Recommended) Password-based or Single-Sign On**
 - **IAM Identity Center**
 - **Single-Sign On Only**
 - **Active Directory**
5. Review the choices that you made, then choose **Create account**.
6. Upon completion, your Quick Suite account will be created. To open Quick Suite, choose **Go to Quick Suite**.

Non-admin sign-up process

If you are not an administrator, you can sign up for Quick Suite using a link provided by your admin. This process allows you to create your account and configure your Quick Suite environment.

To sign up as a non-admin user

1. **Getting your sign-up link**

Sign up for Amazon Quick with the link provided by your admin.

2. Account creation and QuickSight access

After you've created your account, from within QuickSight, search for Amazon Quick.

3. Account configuration

Enter your Account name, Email, region, and Authentication method.

- For authentication, Amazon Q supports Pro license integration with start URL format: <https://amzn.awsapps.com/start>
- Default region is us-east-1

4. Plan selection

Select the plan and storage options you are using for your instance of Amazon Quick.

Signing in to Amazon Quick Suite

You can sign in to Amazon Quick Suite multiple ways, depending on what your Quick Suite administrator has set up. You can sign in to Quick Suite using Amazon root, Amazon Identity and Access Management (IAM), corporate Active Directory, or your native Quick Suite credentials. If your Quick Suite account is integrated with an identity provider such as Okta, the following procedures don't apply to you.

If you're a Quick Suite administrator, make sure to allow-list the following domains within your organization's network.

User type	Domain or domains to allow-list
Users who sign in directly through Quick Suite and Active Directory users	<code>signin.aws</code> and <code>awsapps.com</code>
Amazon root user	<code>signin.aws.amazon.com</code> and <code>amazon.com</code>
IAM users	<code>signin.aws.amazon.com</code>

Post-setup configuration

After signing in to Quick Suite, you can customize your environment by configuring agent settings and other preferences.

To complete post-setup configuration

1. Agent customization access

From the left navigation menu, under Customization, select Agent customization.

2. Environment settings

Enter the settings you require for your environment. These settings will customize Amazon Q's behavior for your specific use case.

Important

We strongly recommend that you don't use the Amazon root user for your everyday tasks, even the administrative ones. Instead, adhere to the best practice of using the root user only to create your first IAM user. Then securely lock away the root user credentials and use them to perform only a few account and service management tasks. For more information, see [Amazon account root user](#) in the *IAM User Guide*.

How to sign in to Quick Suite

Use the following procedure to sign in to Quick Suite.

To sign in to Quick Suite

1. Go to <https://quicksight.aws.amazon.com/>.
2. For **Quick Suite account name**, enter your account name. This is the name that was created when the Quick Suite account was created in Amazon.

If you were invited to the Quick Suite account by email, you can find the account name inside of that email. If you don't have the email that invited you to Quick Suite, ask the Quick Suite administrator in your organization for the information that you need.

You can also find your Quick Suite account name by selecting the profile icon at the upper-right of the Quick Suite console menu. In some cases, you might not have access to your Quick Suite account or have an administrator who can provide this information, or both. If so, contact Amazon Support and open a ticket that includes your Amazon customer ID.

3. For **Username**, enter your Quick Suite user name. User names that contain a semicolon (;) aren't supported. Choose one of the following:

- For organizational users – The user name is provided by your administrator.

Your account can be based on IAM credentials or your email address if it's a root email address. Or it can be used as the user name to invite you into the Quick Suite account. If you received an invitation email from another Quick Suite user, it indicates what type of credentials to use.

- For individual users – The user name that you created for yourself.

This is usually the IAM credentials that you created.

The remaining steps vary depending on the user type you sign in as (directly through Quick Suite or as an Active Directory user, Amazon root user, or IAM user). For more information, see the following sections.

Finishing Quick Suite sign-in as a Quick Suite or Active Directory user

If you're signing in directly through Quick Suite or are using your corporate Active Directory credentials, you're redirected to `signin.aws` after you enter your account name and user name. Use the following procedure to finish signing in.

To finish signing in to Quick Suite if you sign in directly through Quick Suite or use Active Directory credentials

1. For **Password**, enter your password.

Passwords are case-sensitive and must be 8–64 characters in length. They must also contain each of the following:

- Lowercase letters (a–z)
- Uppercase letters (A–Z)
- Numbers (0–9)

- Nonalphanumeric characters (~!@#\$\$%^&* _-+= `|()\{\}[];'"<>,.?/)
2. If your account is multi-factor authentication enabled, enter the multi-factor authentication code that you receive for **MFA code**.
 3. Choose **Sign in**.

Finishing Quick Suite sign-in as an Amazon root user

If you're signing in as an Amazon root user, you're redirected to signin.aws.amazon.com (or amazon.com) to complete the sign-in process. Your user name is prefilled. Use the following procedure to finish signing in.

To finish signing in as an Amazon root user

1. Choose **Next**.
2. For **password**, enter your password. For more information about root user passwords, see [Changing the Amazon account root user password](#) in the *IAM User Guide*.
3. Choose **Sign in**.

Finishing Quick Suite sign-in as an IAM user

If you're signing in as an IAM user, you're redirected to signin.aws.amazon.com (or amazon.com) to complete the sign-in process. Your user name is prefilled. Use the following procedure to finish signing in.

To finish signing in as an IAM user

1. For **Password**, enter your password. For more information about IAM user passwords, see [Default password policy](#) in the *IAM User Guide*.
2. Choose **Sign in**.

If your sign-in process happens automatically and you need to use a different account, use a private or incognito browser window. Doing this prevents the browser from reusing cached settings.

Getting started with Amazon Quick Suite

Amazon Quick Suite is a comprehensive business intelligence platform that enables you to analyze data, create visualizations, automate workflows, and collaborate with your team. Use the topics in this section to get started with the core capabilities of Quick Suite, including Amazon Quick Sight for data analysis, Quick Flows for workflow automation, and Quick Automate for process optimization.

You can begin by using sample data to explore the platform's features, or connect to your own data sources to start building real-world solutions.

Topics

- [Using the Amazon Quick Suite console](#)
- [Getting started with Amazon Quick Sight](#)

Using the Amazon Quick Suite console

In the following topic, you can find a brief introduction to using the Amazon Quick Suite user interface.

Topics

- [Using the Amazon Quick Suite menu and landing page](#)
- [Searching Amazon Quick Suite](#)
- [Choosing a language in Amazon Quick Suite](#)

Using the Amazon Quick Suite menu and landing page

After you sign in to Amazon Quick Suite, you see the Amazon Quick Suite landing page. This page provides tabs for your analyses, dashboards, folders, data, and more. It also provides a menu bar at the top, with options for the following:

- Searching Amazon Quick Suite
- Accessing your user profile (managing Quick Suite, language selection, and help)
- Interacting with the Quick Suite chat agent
- Choosing the Amazon Region that you want to work in

- Discovering helpful resources (documentation, how-to videos, community forum, blogs, and more)

Note

Consult your administrator before changing your Amazon Web Services Region. Your default Amazon Web Services Region is configured by your Amazon Quick Suite administrator. Changing the Amazon Region changes where your work is stored.

Navigating the homepage

Using the navigation menu on the left, you can explore how to work with dashboards, build powerful analyses, manage your data sources, organize content with folders, leverage advanced analytics capabilities, and more.

The following are all the available Amazon Quick Suite capabilities:

- **Agents** – Build AI agents that intelligently analyze, plan, and work with your data to help complete tasks.
- **Spaces** – Work with Amazon Q using specific files, data, and information to complete projects or organize teams.
- **Scenarios** – Elevate your data exploration in the AI-driven workbook canvas that is enhanced with generative analysis capabilities. Allow AI to suggest relevant data to enrich your insight discovery.
- **Stories** – Use Amazon Q prompts and visuals to produce a story that incorporates the details that you provide.
- **Topics** – Create a new generative topic to introduce the generative Q&A experience for users.
- **Dashboards** – Share interactive data visualizations with your audience. You can explore and filter visualizations according to your specific needs and generate snapshots for easy sharing.
- **Analyses** – Dive deep into your data with robust visualization and analysis tools. Create charts and graphs to uncover insights before publishing to a wider audience.
- **Favorites** – Favorite the dashboards and analyses you want to easily find later.
- **Folders** – Organize your analyses and dashboards in folders.
- **Automations** – Create automated processes that reduce manual intervention and save valuable time.

- **Flows** – Create and share intelligent workflows to accomplish goals quickly and predictably.
- **Data** – Connect to data and visualize, create specialized chat assistants, or mix both.
- **Actions** – Perform tasks such as creating and updating tickets, managing incidents, and accessing project information directly.

Customizing the homepage layout

The Quick Suite homepage provides a customizable workspace where you can add widgets that enhance your analytics experience. The homepage features three key widget types:

- **Quick Suite use cases** – Offers quick access to tutorial videos and onboarding resources.
- **Recents** – Displays recently viewed resources for faster navigation.
- **Business Metrics** – Allows you to pin important visualizations and metrics directly from your dashboards.

You can easily add, remove, and rearrange these widgets to create a personalized homepage that suits your workflow.

To add a widget

1. From the Quick Suite homepage, choose **Add widget**.
2. You can choose to add either **Quick Suite use cases**, **Recents**, or **Business metrics**.
 - a. If **Quick Suite use cases** is already added to the homepage, this option will be unavailable. If you previously deleted the widget, choose **Add widget** then **Quick Suite use cases** to add it back to the homepage. **Quick Suite use cases** always appears at the top of the homepage to ensure easy access to learning resources, even if previously dismissed.
 - b. If you choose **Recents**, a widget containing recently viewed resources is added to the homepage. Choose a resource to quickly pick up where you left off.
 - c. If you choose **Business metrics**, you will be redirected to the **Dashboards** page. From there, you can open a dashboard and choose which visual you would like to pin to the homepage by selecting the home + icon.

To delete a widget

1. Navigate to the widget you would like to delete.

2. Choose the three-dot menu on the top right of the widget.
3. Choose **Delete**.
4. The selected widget should be removed from the homepage.

To reorder a widget

1. Select the widget you would like to reorder.
2. When the widget is selected, the handle icon appears on the top left of the widget.
3. Choose the handle icon and drag the widget to swap its place with another widget.

To resize a widget

1. Select the widget you would like to resize.
2. When the widget is selected, a resize icon appears on the bottom right.
3. Choose the resize icon and drag the widget up or down to reach your desired size.

Accessing user profiles

To access the user profile menu, choose the user icon at the upper right of any page in Amazon Quick Suite. Use this menu to manage Amazon Quick Suite features, choose a language, change your Amazon Web Services Region, get help from the documentation, or sign out of Amazon Quick Suite.

The following options are available from the user profile menu:

- **Manage Quick Suite** – If you have appropriate permissions, you can access administrative functions such as managing users, subscriptions, [SPICE](#) capacity, and account settings.
- **Language setting** – Choose the language you want to use in the Amazon Quick Suite user interface.
- **Region setting** – Choose the Amazon Web Services Region that you want to work in.

Note

Consult your administrator before changing your Amazon Web Services Region. Your default Amazon Web Services Region is configured by your Amazon Quick Suite administrator. Changing the Amazon Region changes where your work is stored.

- **Help** – This will open the official Quick Suite documentation, which you can view online, in Kindle, or as a PDF.
- **Sign out** – Choose this option to sign out of Amazon Quick Suite and your Amazon session.

Searching Amazon Quick Suite

From the search bar, you can search for analyses and dashboards. To use the search tool, go to the **Start Page** and choose the search box at the top-left of the page. Then enter the name, or part of the name, of the data set, analyses, or dashboard you want to find. The search is not case-sensitive.

After you locate the item that you're looking for, you can open it directly from the search results. You can modify a dataset, create an analysis from a dataset, or access an analysis or dashboard. Choose an item from the search results to open it.

Choosing a language in Amazon Quick Suite

You can choose the language that you want to use in the Amazon Quick Suite user interface. This option is set separately for each individual user. The first time a user signs in, Amazon Quick Suite detects and selects a suitable language. This choice is based on the user's browser preferences and interactions with localized Amazon websites.

Quick Suite supports the following languages:

Languages available in the Amazon Quick Suite user interface

Official name	Language code	Localized name
Dansk	da	Danish
Deutsch	de	German

Official name	Language code	Localized name
English	en	English
Español	es	Spanish
Français	fr	French
Italiano	it	Italian
Nederlands	nl	Dutch
Norsk	nb	Norwegian
Português	pt	Portuguese
Suomi	fi	Finnish
Svenska	sv	Swedish
日本語	ja	Japanese
한국어	ko	Korean
中文 (简体)	zh-CN	Simplified Chinese
中文 (繁體)	zh-TW	Traditional Chinese

Choosing a language translates only user interface elements. It doesn't translate the following:

- Amazon Quick Suite reserved keywords
- User input
- Data
- Date or number formats
- ML Insights, suggested insights, or computations in narratives (including text)

Use the following procedure to change the language in the Amazon Quick Suite interface.

1. Choose your user name at top right.
2. To open the language options menu, choose the > symbol near the current language.
3. Choose the language that you want to use.

Getting started with Amazon Quick Sight

Amazon Quick Sight is the data visualization and analysis component of Amazon Quick Suite. This section walks you through creating your first analysis using sample data, including uploading datasets, creating visualizations with AutoGraph, and sharing your results through dashboards.

Topics

- [Quick start: Create an Amazon Quick Sight analysis with a single visual using sample data](#)
- [Tutorial: Create an Amazon Quick Sight dashboard using sample data](#)
- [Sample data](#)

Quick start: Create an Amazon Quick Sight analysis with a single visual using sample data

Before you create your first analysis, make sure to complete the steps in [Setting up and signing into Amazon Quick Suite](#).

With the following procedure, you use the Web and Social Media Analytics sample dataset to create an analysis containing a line chart visual. This visual shows the count by month of people that have added themselves to the mailing list.

To create an analysis containing a line chart visual using a sample dataset

1. From the Amazon Quick Suite homepage, from Amazon Quick Sight, choose **Analyses** from the left navigation menu. If you don't have sample data, you can download it from [web-and-social-analytics.csv.zip](#). Unzip the file so you can use the .csv file.

To upload the sample data, do the following:

- a. Choose **Data** from the left navigation menu. Under the **Dataset** tab, select **New** then **Dataset**.
- b. Choose **Upload file**.

- c. Choose the sample file, `web-and-social-analytics.csv`, from your drive. If you don't see it, check that you unzipped the `web-and-social-analytics.csv.zip` file.
 - d. Confirm file upload settings by choosing **Next** on the **Confirm file upload settings** screen.
 - e. Choose **Visualize** on the **Data source details** screen.
 - f. Skip the next step. Choosing **Visualize** brings you to the same screen as the process in Step 2.
2. On the **Datasets** page, choose the **Web and Social Media Analytics** dataset, and then choose **Use in Analysis** at upper right.
3. In the **Data** pane, choose **Date**, and then choose **Mailing list adds**.

Amazon Quick Sight uses AutoGraph to create the visual, selecting the visual type that it determines is most compatible with those fields. In this case, it selects a line chart that shows mailing list adds by day, which is the date granularity default.

4. Navigate to the **Field wells** at the bottom of the **Visuals** pane.
5. Choose the **X axis** field well. Select the three-dot menu, choose **Aggregate**, and then choose **Month**.

The line chart updates to show mailing list adds by month, rather than by the default of by year.

Tutorial: Create an Amazon Quick Sight dashboard using sample data

Use the procedures in the following sections to complete these tasks:

- Create and prepare a Marketing dataset using the Web and Social Media Analytics sample data.
- Create a Marketing analysis and add several visuals to it.
- Modify the visuals in the analysis, including the following:
 - Adding another measure to an existing visual
 - Changing chart colors
 - Changing date granularity
 - Changing the size and layout of the visuals
 - Applying a filter
- Publish a dashboard based on the analysis.

Topics

- [Tutorial: Create a prepared Amazon Quick Sight dataset](#)
- [Tutorial: Create an Amazon Quick Sight analysis](#)
- [Tutorial: Modify Amazon Quick Sight visuals](#)
- [Tutorial: Create an Amazon Quick Sight dashboard](#)

Tutorial: Create a prepared Amazon Quick Sight dataset

Use the following procedure to prepare the Marketing dataset and create an analysis. If you don't see the Web and Social Media Analytics sample data already in Amazon Quick Sight, you can download it: [web-and-social-analytics.csv.zip](#).

To prepare the Marketing dataset and create an analysis

1. From the Amazon Quick Suite homepage, choose **Data** at left.
2. In the **Datasets** tab, choose **New** then **Dataset**.
3. From the existing data sources, choose the **Web and Social Media Analytics** Amazon S3 data source. Choose your table and then choose **Edit/Preview data**.

Amazon Quick Suite opens the data preparation page.

4. For **Dataset Name**, enter **Marketing Sample** to replace *Web and Social Media Analytics* for the dataset name.
5. Exclude some fields from the dataset.

In the **Fields** pane, choose the three-dot menu for the **Twitter followers cumulative** and **Mailing list cumulative** fields, and then choose **Exclude field**. To select more than one field at a time, press the Ctrl key while you select (Command key on Mac).

6. Rename a field.

In the **Dataset** preview pane, scroll to the **Website Pageviews** field and choose the pencil icon to edit.

In the **Edit field** page that opens, for **Name**, enter **Website page views**, and then choose **Apply**.

7. Add a calculated field that substitutes a text string for any 0-length string value in the **Events** field:

- a. On the data preparation page, scroll to the top of the **Fields** pane, and then choose **Add calculated field**.
- b. In the **Add calculated field** page that opens, for **Add name**, enter **populated_event**.
- c. In the **Functions** pane at right, double-click the **ifelse** function from the list of functions. This adds the function to the calculated field formula.
- d. Expand the **Field list** pane by choosing the drop-down arrow, and then double-click the **Events** field. This adds the field to the calculated field formula.
- e. In formula editor, enter the following additional functions and parameters required, in bold in the following: `ifelse(strlen({Events})=0, 'Unknown', {Events})`.

The final formula should be as follows: `ifelse(strlen({Events})=0, 'Unknown', {Events})`.

- f. Choose **Save**.

The new calculated field is created, and appears at the top of the **Fields** pane.

8. Choose **Save**.

Next steps

Create an analysis by using the procedure in [Tutorial: Create an Amazon Quick Sight analysis](#).

Tutorial: Create an Amazon Quick Sight analysis

In the following short tutorial, you create an analysis, add a visual using AutoGraph, and add another visual by choosing a specific visual type. This procedure builds on the dataset that you create and prepare in [Tutorial: Create a prepared Amazon Quick Sight dataset](#).

Create your analysis

Use the following procedure to create your analysis.

To create your analysis

1. On the Amazon Quick Sight homepage, choose **Analyses**.
2. Choose **Create analysis** and select the dataset to use.

Create a visual by using AutoGraph

Create a visual by using AutoGraph, which is selected by default.

On the analysis page, choose **Date** and **Return visitors** in the **Fields list** pane.

Amazon Quick Sight creates a line chart using this data.

Create a scatter plot visual

Create a visual by choosing a visual type and dragging fields to the field wells.

To create a scatter plot visual

1. On the analysis page, choose **Insert** and then **Add visual** on the application bar. A new, blank visual is created, and AutoGraph is selected by default.
2. In the **Visual types** pane, choose the scatter plot icon.
3. Choose fields in the **Fields list** pane to add to the **Field wells** pane:
 - Choose **Desktop Uniques** to populate the **X axis** field well.
 - Choose **Mobile Uniques** to populate the **Y axis** field well.
 - Choose **Date** to populate the **Group/Color** field well.

A scatter plot is created using these fields.

Next steps

Modify the visuals in your analysis by using the procedure in [Tutorial: Modify Amazon Quick Sight visuals](#).

Tutorial: Modify Amazon Quick Sight visuals

Use the following procedures to modify the visuals that you created using the procedures in [Tutorial: Create an Amazon Quick Sight analysis](#).

Modify the line chart visual

Modify your line chart visual by making it show an additional measure by date, and also by changing the chart color.

To modify your line chart visual

1. In your analysis, select the line chart visual.
2. Add another measure to the visual.

Select the **New visitors SEO** field in the **Fields list** pane. This measure is added to the **Value** field well, and the line chart updates with a line to represent it. The visual title also updates.

3. Change the color of the line used to represent the **Return visitors** measure.

Choose the line on the chart that represents **Return visitors**. To do this, choose the end of the line, not the middle of the line.

Choose **Color Return visitors**, and then choose the red icon from the color selector.

4. Choose the **Date** field in the **X axis** field well, choose **Aggregate**, and then choose **Month**.

Modify the scatter plot visual

Modify your scatter plot visual by changing the data granularity.

To modify your scatter chart visual

1. In the analysis, select the scatter plot visual.
2. Choose the **Group/Color** field well, choose **Aggregate**, and then choose **Month**.

The scatter plot updates to show the measures by month, rather than by the default of by year.

Modify both visuals by changing visual layout and adding a filter

Modify both visuals by changing visual size and location, and by adding a filter and applying it to both of them.

Change the visual layout

Modify both visuals by changing visual size and location.

To modify both visuals

1. In your analysis, select the line chart visual.

2. Choose the resize handle in the lower right corner of the visual and drag up and to the left, until the visual is half its former size both horizontally and vertically.
3. Repeat this procedure on the scatter plot visual.
4. Choose the move handle on the scatter plot visual, and drag it up to the right of the line chart visual so that they are side by side.

Modify both visuals by adding a filter

Modify both visuals by adding a filter and applying it to both of them.

To add a filter to both visuals

1. In the analysis, choose the scatter plot visual.
2. Choose **Insert** then **Add Filter** on the application bar.
3. Choose the **Date** field to filter on.
4. Select the new filter to expand it.
5. In the **Edit filter** pane, for **Condition**, choose the **After** comparison type.
6. Enter a start date value of 1/1/2014.

Choose **Date**, choose **2014** for the year, **January** for the month, and then choose **1** on the calendar.

7. In the **Edit filter** pane, choose **Apply** to apply the filter to the visual.

The filter is applied to the scatter plot visual. This is indicated with a filter icon on the visual drop-down menu.

8. Apply the filter to the line chart visual.

In the **Filter** pane at left, choose the **Date** filter again and choose **Single visual**, and then choose **All visuals of this dataset**.

The filter is applied to the line chart visual as well.

Next steps

Create a dashboard from your analysis by using the procedure in [Tutorial: Create an Amazon Quick Sight dashboard](#).

Tutorial: Create an Amazon Quick Sight dashboard

Use the following procedure to create a dashboard from the analysis that you created using the procedure in [Tutorial: Create an Amazon Quick Sight analysis](#).

To create a dashboard from your analysis

1. In your analysis, choose **Publish** in the application bar at upper right.
2. In the **Publish dashboard** page that opens, choose **Publish new dashboard as**, and enter the name **Marketing Dashboard**.
3. Choose **Publish dashboard**.

The dashboard is now published.

4. On the **Share dashboard** page that opens, choose the **X** icon to close it. You can share the dashboard later by using the sharing option on the dashboard page.

Sample data

To get a first look at how Amazon Quick Sight works, you can explore Amazon Quick Sight using the following sample data:

- [B2b sales data](#)
- [Business overview data](#) (revenue)
- [ML insights data](#)
- [People overview data](#) (human resources)
- [Sales pipeline data](#)
- [Web and social media analytics data](#) (marketing)

Also, a variety of datasets are available free online that you can use with Amazon Quick Sight, for example the [Amazon public datasets](#). These datasets come in a variety of formats.

Administration for Amazon Quick Suite

Use the following section to learn about Amazon Quick Suite administrative tasks. This section contains information about controlling access, managing accounts, and choosing Amazon Web Services Regions. It also includes the essential setup tasks, such as integrating your account with IAM Identity Center, assigning access to users, managing the Amazon Quick Suite service.

Topics

- [Getting started with Amazon Quick Suite](#)
- [Setting up for Quick Suite](#)
- [Configuring VPC connections in Amazon Quick Sight](#)
- [Configuring Amazon Quick Sight access to Amazon data sources](#)
- [Amazon Quick Suite account management](#)
- [Managing Quick Suite subscriptions](#)
- [Managing identities](#)
- [Managing security](#)
- [Permissions](#)
- [Customizations](#)
- [Using the Amazon Quick Suite analytics dashboard](#)
- [Amazon security in Quick Suite](#)

Getting started with Amazon Quick Suite

This chapter provides you an overview of the key Amazon Quick Suite concepts you need to understand before you dive in and start using Quick Suite.

Topics

- [Amazon Quick Suite user types](#)
- [Different editions of Amazon Quick Suite](#)
- [Supported browsers](#)

Amazon Quick Suite user types

Amazon Quick Suite has three distinct user personas based on permission levels and use cases: administrator, author, and reader. Within these personas, there are six specific user roles that can be assigned, including Pro versions that provide access to advanced Amazon Quick Suite capabilities. The following sections describe the role and task boundaries of each type.

Topics

- [Understanding Amazon Quick Suite subscriptions and roles](#)
- [Understanding reader capabilities](#)
- [Understanding author capabilities](#)
- [Understanding admin capabilities](#)

Understanding Amazon Quick Suite subscriptions and roles

The new capabilities for Amazon Quick Suite are available via the following monthly user subscriptions:

- **Amazon Quick Suite Professional** (or Amazon Quick Sight Reader Pro role)
- **Amazon Quick Suite Enterprise** (or Amazon Quick Sight Author Pro role)

To manage Amazon Quick Suite users, billing and services, you need the Admin Pro role, which is the same as Amazon Quick Suite Enterprise user subscription.

Amazon Quick Suite also provides the following monthly user subscriptions to support business only (BI) only use cases:

- **Amazon Quick Sight Reader** – For users who need to view and interact with dashboards and reports.
- **Amazon Quick Sight Author** – For users who need to create and manage BI content.

To manage Amazon Quick Sight users, billing and services, you need Admin role, which is the same as Amazon Quick Sight Author user subscription.

 Note

For pricing information, see [Amazon Quick Suite pricing](#)

Understanding reader capabilities

Readers can use Quick Suite to access company data and find answers through chat interactions with AI agents. They can upload files, run automations, create visualizations, and share spaces via direct links. Reader Pro users additionally have access to advanced Amazon Quick Suite features including AI chat agents, collaborative spaces, flows, and extensions. For detailed information about reader capabilities and subscription types, see [Amazon Quick Suite pricing](#).

Understanding author capabilities

Authors are domain experts who build and manage Quick Suite resources. In addition to reader capabilities, they can create datasets, dashboards, automations, and agents. They have expanded sharing permissions for spaces and can use AI to create visualizations. Author Pro users can additionally create content using natural language, build knowledge bases, configure actions, and access advanced automation capabilities. For detailed information about author capabilities and subscription types, see [Amazon Quick Suite pricing](#).

Understanding admin capabilities

Administrators manage user access, monitor costs, and maintain data sources with full reader and author capabilities. There are two types of administrators: system administrators oversee the broader Amazon environment including Amazon Quick Suite signup, system health, and security; Amazon Quick Suite administrators manage users and resources within Amazon Quick Suite.

When your Amazon Quick Suite account is integrated with IAM Identity Center, admin capabilities are divided between IAM permissions and Amazon Quick Suite admin role permissions. Access to some sections of the Amazon Quick Suite administration console is restricted by IAM permissions. The following table summarizes the admin actions that you can perform in Amazon Quick Suite based on the access type that you choose.

Admin action	IAM permissions	Amazon Quick Suite admin role permissions
Manage assets	Yes	No

Admin action	IAM permissions	Amazon Quick Suite admin role permissions
Security & permissions	Yes	No
Manage VPC connections	Yes	No
KMS keys	Yes	No
Account settings	Yes	No
Account customization	No	Yes
Manage users	Yes (IAM Identity Center users)	Yes (Amazon Quick Suite and IAM users)
Your subscriptions	No	Yes
Mobile settings	No	Yes
Domains and embedding	No	Yes
SPICE capacity	No	Yes

Admin and Admin Pro users have full reader and author capabilities but focus primarily on system administration to ensure efficient and secure operations for all users. For detailed information about author capabilities and subscription types, see [Amazon Quick Suite pricing](#).

Different editions of Amazon Quick Suite

Amazon Quick Suite offers Standard and Enterprise editions. Both editions offer a full set of features for creating and sharing data visualizations. Enterprise edition additionally offers encryption at rest and Microsoft Active Directory integration. In Enterprise edition, you select a Microsoft Active Directory directory in Amazon Directory Service. You use that active directory to identify and manage your Quick Suite users and administrators.

For more information about the different features offered by the Quick Suite editions and about pricing, see [Amazon Quick Suite pricing](#).

Note

New Amazon Quick Suite features are only available under the Enterprise edition.

Topics

- [Availability of editions](#)
- [User management between editions](#)
- [Permissions for the different editions](#)

Availability of editions

All editions are available in any Amazon Web Services Region that is currently supported by Amazon Quick Suite.

The capacity region in which you start your Amazon Quick Suite subscription is where your account's default [SPICE](#) capacity is allocated. However, you can purchase additional SPICE capacity and access your Amazon resources in any other supported Amazon Web Services Region.

New Amazon Quick Suite features are not available in the Standard edition. If you are an existing customer on the Standard edition, you can upgrade to Enterprise edition and provision appropriate roles to access Amazon Quick Suite features.

To manage Enterprise account settings, you must temporarily change your region for your session to US East (N. Virginia) Region. You can change it back when you have finished editing your account settings. These settings include changing your subscription's notification email, enabling IAM access requests, editing access to Amazon resources, and unsubscribing from Amazon Quick Suite.

User management between editions

User management is different between the Amazon Quick Suite Standard and Enterprise editions. However, both editions support identity federation, or Federated Single Sign-On (IAM Identity Center), through Security Assertion Markup Language 2.0 (SAML 2.0).

Topics

- [User management for standard edition](#)
- [User management for enterprise edition](#)

User management for standard edition

In the Standard edition, as a system administrator, you can invite an Amazon Identity and Access Management user and allow that user to use their credentials to access Amazon Quick Suite. Alternatively, you can invite any person with an email address to create an Amazon Quick Suite-only account. When you create a Amazon Quick Suite user account, Amazon Quick Suite sends an email to that user inviting them to activate their account.

When you create a Amazon Quick Suite user account, you also choose to assign it either an administrative or a user role. This role assignment determines the user's permissions in Amazon Quick Suite. You perform all management of users by adding, changing, and deleting accounts in Amazon Quick Suite.

User management for enterprise edition

In the Enterprise edition, as a system administrator, you can select one or more IAM Identity Center or Microsoft Active Directory groups for administrative access. All users in these groups are authorized to sign in to Amazon Quick Suite as Amazon Quick Suite administrators. You can also select one or more IAM Identity Center or Microsoft Active Directory groups in Amazon Directory Service for user access. All users in these groups are authorized to sign in to Amazon Quick Suite as users.

Important

With IAM Identity Center, share the Amazon sign in portal with end users to access Amazon Quick Suite. For more information, see [Sign in to the Amazon access portal](#).

With Active Directory, Amazon Quick Suite Administrators and users aren't automatically notified of their access to Amazon Quick Suite. You must email users with the sign-in URL, the account name, and their credentials.

You can only add or remove Enterprise edition accounts by adding or removing a person from the IAM Identity Center or Microsoft Active Directory group that you associated with Amazon Quick Suite. When you add a Amazon Quick Suite user account, its permissions depend on whether the IAM Identity Center or Microsoft Active Directory group is an administrative group or a user group in Amazon Quick Suite.

To remove a user's access to Amazon Quick Suite, remove the user from an IAM Identity Center or Microsoft Active Directory group or remove their IAM Identity Center or Microsoft Active Directory group from an associated role in Amazon Quick Suite.

Permissions for the different editions

In the Standard edition, all Amazon Quick Suite administrators can manage subscriptions and SPICE capacity. They can also add, modify, and delete accounts.

Additional IAM permissions are required to manage Amazon Quick Suite permissions to Amazon resources and to unsubscribe from Amazon Quick Suite. These tasks can only be performed by an IAM user who also has administrative permissions in Amazon Quick Suite, or by the IAM user or Amazon account (system administrator) that created the Amazon Quick Suite account.

To manage access to Amazon resources from Amazon Quick Suite, you must be logged in as one of the following:

- Any IAM user who is an Amazon Quick Suite administrator
- The IAM user or Amazon root account that created the Amazon Quick Suite account (system administrator)

All IAM Identity Center or Microsoft Active Directory users that are Amazon Quick Suite administrators can manage subscriptions and SPICE capacity.

Additional IAM permissions are required to manage access to Amazon resources or to unsubscribe from Amazon Quick Suite. Administrators need to sign in with IAM permissions to perform these tasks.

For more information on admin user permissions, see [Amazon Quick Suite user types](#).

Supported browsers

Before you start working with Amazon Quick Suite, use the following table to verify that your browser is supported for Quick Suite access.

Note

Amazon Quick Suite ended support for Microsoft Internet Explorer 11 on July 31, 2022. We can no longer ensure that the features and web pages of Quick Suite will function properly

on IE 11. Please use one of our supported browsers: Microsoft Edge (Chromium), Google Chrome, or Mozilla Firefox.

Browser	Version	Check your version
Apple Safari	13 or later	Open Safari. On the menu, choose Safari , and then choose About Safari . The version number is shown in the dialog box that displays.
Google Chrome	Last three versions	Open Chrome and type chrome://version in your address bar. The version is in the Google Chrome field at the top of the results.
Microsoft Edge (Chromium)	Latest version	Not applicable.
Mozilla Firefox	Last three versions	Open Firefox. On the menu, choose the Help icon, and then choose About Firefox . The version number is listed underneath the Firefox name.

Setting up for Quick Suite

This section helps you understand how to configure identity and access management, and control sign-up for users when you set up an Amazon Quick Suite instance.

Note

For information about how to set up your Amazon Quick Suite instance, including creating an Amazon account, signing up for an Amazon Quick Suite subscription, and signing in to Amazon Quick Suite, see [Setting up and signing into Amazon Quick Suite](#).

Topics

- [Using service control policies to restrict Amazon Quick Suite sign-up options](#)
- [Identity and access management in Quick Suite](#)
- [Allowlisting Amazon Quick Suite domains](#)

Using service control policies to restrict Amazon Quick Suite sign-up options

If you're an administrator in Amazon Organizations, you can use service control policies (SCPs) to restrict how individuals in your organization can sign up for Amazon Quick Suite. You can restrict the edition of Quick Suite they can sign up for, and also the type of user that they can sign up for.

Amazon Organizations is a user account management service that you can use to consolidate multiple Amazon accounts into an organization that you create and centrally manage. You can use SCPs in Amazon Organizations to manage the permissions in your organization. For more information, see [What is Amazon Organizations?](#) and [Service control policies](#) in the *Amazon Organizations User Guide*.

In the following topic, you can learn about two ways to restrict Quick Suite sign-up options using SCPs in Amazon Organizations. The topic includes an example SCP. To learn more about creating SCPs, see the following topics in the *Amazon Organizations User Guide*:

- [Creating, updating, and deleting service control policies](#)
- [SCP syntax](#)
- [Strategies for using SCPs](#)

Topics

- [Restricting the Quick Suite edition](#)
- [Restricting user management options](#)
- [Example SCP](#)

Restricting the Quick Suite edition

To restrict the edition of Quick Suite that your managed accounts can sign up for, use the `quicksight:Edition` condition key in your SCP. The values for this key are listed and described in the following table.

Key Name	Key Value	Description
<code>quicksight:Edition</code>	<code>standard</code>	Amazon Quick Suite Standard Edition
	<code>enterprise</code>	Amazon Quick Suite Enterprise Edition

Restricting user management options

To restrict the user management options that individuals in your organization can use to sign up for Quick Suite, use the `quicksight:DirectoryType` condition key in your SCP. The values for this key are listed and described in the following table.

Key Name	Key Value	Description
<code>quicksight:DirectoryType</code>	<code>quicksight</code>	IAM federated identities and Amazon Quick Suite-managed users
	<code>iam</code>	Only IAM federated identities
	<code>microsoft_ad</code>	Users managed in Microsoft Active Directory on Amazon Directory Service for Microsoft Active Directory
	<code>ad_connector</code>	Users managed in on-premises Active Directory and connected through AD_Connector to Amazon

Key Name	Key Value	Description
		Directory Service for Microsoft Active Directory
	iam_identity_center	Users managed in a Amazon Quick Suite account that is integrated with IAM Identity Center.

Example SCP

The following example for Quick Suite shows a service control policy that denies signing up for a Amazon Quick Suite Standard Edition and turns off the ability to sign up using Amazon Quick Suite or Active Directory credentials. This policy uses the `quicksight:subscribe` action, in addition to the condition keys previously described. For a list of Amazon Quick Suite-specific keys for use in IAM permission policies, see [Actions, resources, and condition keys for Quick Suite](#) in the *Service Authorization Reference*.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Deny",
      "Action": [
        "quicksight:Subscribe"
      ],
      "Resource": [
        "*"
      ],
      "Condition": {
        "ForAnyValue:StringEquals": {
          "quicksight:DirectoryType": [
            "iam_identity_center"
          ]
        }
      }
    }
  ],
}
```

```
    "Sid": "Statement2",
    "Effect": "Deny",
    "Action": [
        "quicksight:Subscribe"
    ],
    "Resource": [
        "*"
    ],
    "Condition": {
        "StringEquals": {
            "quicksight:Edition": "standard"
        }
    }
}
```

With this policy in effect, individuals in an organization can sign up only for Amazon Quick Suite Enterprise Edition. Additionally, they can sign up only by using the **IAM Identity Center enabled application** option. If they try to sign up for Amazon Quick Suite Standard Edition or use another form of authentication, they are restricted from signing up. They receive a message explaining that they don't have the right permissions to sign up for Amazon Quick Suite.

Identity and access management in Quick Suite

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators and Amazon Quick Suite administrators

You can use the following tools for identity and access to Quick Suite:

- [IAM Identity Center](#) (Enterprise edition only)
- [IAM federation](#) (Standard and Enterprise editions)
- [Amazon Directory Service for Microsoft Active Directory](#) (Enterprise edition only)
- [SAML-based single sign-on](#) (Standard and Enterprise edition)
- [Multifactor authentication \(MFA\)](#) (Standard and Enterprise edition)

Note

In the regions listed below, Amazon Quick Suite accounts can only use [IAM Identity Center](#) for identity and access management.

- af-south-1 Africa (Cape Town)
- ap-southeast-3 Asia Pacific (Jakarta)
- eu-south-1 Europe (Milan)
- eu-central-2 Europe (Zurich)

The following sections help you configure the identity management method of your choice for Quick Suite.

Topics

- [Using IAM](#)
- [Using IAM Identity Center](#)
- [IAM federation](#)
- [Using Active Directory with Amazon Quick Suite Enterprise edition](#)
- [Using multi-factor authentication \(MFA\) with Amazon Quick Suite](#)

Using IAM

Amazon Identity and Access Management (IAM) is an Amazon Web Services service that helps an administrator securely control access to Amazon resources. IAM administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Amazon Quick Suite resources. IAM is an Amazon Web Services service that you can use with no additional charge.

Topics

- [Introduction to IAM concepts](#)
- [Using Quick Suite with IAM](#)
- [Passing IAM roles to Quick Suite](#)
- [IAM policy examples for Quick Suite](#)
- [Provisioning users for Amazon Quick Suite](#)

- [Troubleshooting Quick Suite identity and access](#)

Introduction to IAM concepts

Amazon Identity and Access Management (IAM) is an Amazon service that helps an administrator to more securely control access to Amazon resources. Administrators control who can be *authenticated* (signed in) and *authorized* (have permissions) to use Amazon Quick Suite resources. IAM is an Amazon service that you can use with no additional charge.

IAM is used with Amazon Quick Suite in several ways, including the following:

- If your company uses IAM for their identity management, people might have IAM user names and passwords that they use to sign in to Amazon Quick Suite.
- If you want your Amazon Quick Suite users to be automatically created at first sign-in, you can use IAM to create a policy for users who are preauthorized to use Amazon Quick Suite.
- If you want to create specialized access for specific groups of Amazon Quick Suite users or to specific resources, you can use IAM policies to accomplish this.

Topics

- [Audience](#)
- [Authenticating with identities](#)
- [Managing access using policies](#)

Audience

Use the following to help understand the context of the information provided in this section, and how it applies to your role. How you use Amazon Identity and Access Management (IAM) differs depending on the work that you do in Amazon Quick Suite.

Service user – In some cases, you might use Amazon Quick Suite as an Author or Reader to interact with data, analyses, and dashboards, spaces, and agents through the Amazon Quick Suite by using the browser interface. In these cases, this section provides only background information for you. You don't directly interact with the IAM service, except if you use IAM to sign into Amazon Quick Suite.

Amazon Quick Suite administrator – If you're in charge of Amazon Quick Suite resources at your company, you probably have full access to Amazon Quick Suite. It's your job to determine

which Amazon Quick Suite features and resources your team members should access. If you have specialized requirements that you can't solve by using the Amazon Quick Suite admin panel, then you can work with your administrator to create permissions policies for your Amazon Quick Suite users. To learn more about IAM, read this page to understand the basic concepts of IAM. To learn more about how your company can use IAM with Amazon Quick Suite, see [Using Amazon Quick Suite with IAM](#).

Administrator – If you're a system administrator, you might want to learn details about how you can write policies to manage access to Amazon Quick Suite. To view examples of Amazon Quick Suite identity-based policies that you can use in IAM, see [IAM identity-based policies for Amazon Quick Suite](#).

Authenticating with identities

Authentication is how you sign in to Amazon using your identity credentials. You must be *authenticated* (signed in to Amazon) as the Amazon Web Services account root user, as an IAM user, or by assuming an IAM role.

If you access Amazon programmatically, Amazon provides a software development kit (SDK) and a command line interface (CLI) to cryptographically sign your requests by using your credentials. If you don't use Amazon tools, you must sign requests yourself. For more information about using the recommended method to sign requests yourself, see [Amazon Signature Version 4 for API requests](#) in the *IAM User Guide*.

Regardless of the authentication method that you use, you might be required to provide additional security information. For example, Amazon recommends that you use multi-factor authentication (MFA) to increase the security of your account. To learn more, see [Amazon Multi-factor authentication in IAM](#) in the *IAM User Guide*.

Topics

- [Amazon Web Services account root user](#)
- [IAM users and groups](#)
- [IAM roles](#)

Amazon Web Services account root user

When you create an Amazon Web Services account, you begin with one sign-in identity that has complete access to all Amazon Web Services services and resources in the account. This identity

is called the Amazon Web Services account *root user* and is accessed by signing in with the email address and password that you used to create the account. We strongly recommend that you don't use the root user for your everyday tasks. Safeguard your root user credentials and use them to perform the tasks that only the root user can perform. For the complete list of tasks that require you to sign in as the root user, see [Tasks that require root user credentials](#) in the *IAM User Guide*.

IAM users and groups

An [IAM user](#) is an identity within your Amazon Web Services account that has specific permissions for a single person or application. Where possible, we recommend relying on temporary credentials instead of creating IAM users who have long-term credentials such as passwords and access keys. However, if you have specific use cases that require long-term credentials with IAM users, we recommend that you rotate access keys. For more information, see [Rotate access keys regularly for use cases that require long-term credentials](#) in the *IAM User Guide*.

An [IAM group](#) is an identity that specifies a collection of IAM users. You can't sign in as a group. You can use groups to specify permissions for multiple users at a time. Groups make permissions easier to manage for large sets of users. For example, you could have a group named *IAMAdmins* and give that group permissions to administer IAM resources.

Users are different from roles. A user is uniquely associated with one person or application, but a role is intended to be assumable by anyone who needs it. Users have permanent long-term credentials, but roles provide temporary credentials. To learn more, see [Use cases for IAM users](#) in the *IAM User Guide*.

IAM roles

An [IAM role](#) is an identity within your Amazon Web Services account that has specific permissions. It is similar to an IAM user, but is not associated with a specific person. To temporarily assume an IAM role in the Amazon Web Services Management Console, you can [switch from a user to an IAM role \(console\)](#). You can assume a role by calling an Amazon CLI or Amazon API operation or by using a custom URL. For more information about methods for using roles, see [Methods to assume a role](#) in the *IAM User Guide*.

IAM roles with temporary credentials are useful in the following situations:

- **Federated user access** – To assign permissions to a federated identity, you create a role and define permissions for the role. When a federated identity authenticates, the identity is associated with the role and is granted the permissions that are defined by the role. For

information about roles for federation, see [Create a role for a third-party identity provider \(federation\)](#) in the *IAM User Guide*.

- **Temporary IAM user permissions** – An IAM user or role can assume an IAM role to temporarily take on different permissions for a specific task.
- **Cross-account access** – You can use an IAM role to allow someone (a trusted principal) in a different account to access resources in your account. Roles are the primary way to grant cross-account access. However, with some Amazon Web Services services, you can attach a policy directly to a resource (instead of using a role as a proxy). To learn the difference between roles and resource-based policies for cross-account access, see [Cross account resource access in IAM](#) in the *IAM User Guide*.
- **Cross-service access** – Some Amazon Web Services services use features in other Amazon Web Services services. For example, when you make a call in a service, it's common for that service to run applications in Amazon EC2 or store objects in Amazon S3. A service might do this using the calling principal's permissions, using a service role, or using a service-linked role.
- **Forward access sessions (FAS)** – When you use an IAM user or role to perform actions in Amazon, you are considered a principal. When you use some services, you might perform an action that then initiates another action in a different service. FAS uses the permissions of the principal calling an Amazon Web Services service, combined with the requesting Amazon Web Services service to make requests to downstream services. FAS requests are only made when a service receives a request that requires interactions with other Amazon Web Services services or resources to complete. In this case, you must have permissions to perform both actions. For policy details when making FAS requests, see [Forward access sessions](#).
- **Service role** – A service role is an [IAM role](#) that a service assumes to perform actions on your behalf. An IAM administrator can create, modify, and delete a service role from within IAM. For more information, see [Create a role to delegate permissions to an Amazon Web Services service](#) in the *IAM User Guide*.
- **Service-linked role** – A service-linked role is a type of service role that is linked to an Amazon Web Services service. The service can assume the role to perform an action on your behalf. Service-linked roles appear in your Amazon Web Services account and are owned by the service. An IAM administrator can view, but not edit the permissions for service-linked roles.
- **Applications running on Amazon EC2** – You can use an IAM role to manage temporary credentials for applications that are running on an EC2 instance and making Amazon CLI or Amazon API requests. This is preferable to storing access keys within the EC2 instance. To assign an Amazon role to an EC2 instance and make it available to all of its applications, you create an instance profile that is attached to the instance. An instance profile contains the role and

enables programs that are running on the EC2 instance to get temporary credentials. For more information, see [Use an IAM role to grant permissions to applications running on Amazon EC2 instances](#) in the *IAM User Guide*.

Managing access using policies

You control access in Amazon by creating policies and attaching them to Amazon identities or resources. A policy is an object in Amazon that, when associated with an identity or resource, defines their permissions. Amazon evaluates these policies when a principal (user, root user, or role session) makes a request. Permissions in the policies determine whether the request is allowed or denied. Most policies are stored in Amazon as JSON documents. For more information about the structure and contents of JSON policy documents, see [Overview of JSON policies](#) in the *IAM User Guide*.

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

By default, users and roles have no permissions. To grant users permission to perform actions on the resources that they need, an IAM administrator can create IAM policies. The administrator can then add the IAM policies to roles, and users can assume the roles.

IAM policies define permissions for an action regardless of the method that you use to perform the operation. For example, suppose that you have a policy that allows the `iam:GetRole` action. A user with that policy can get role information from the Amazon Web Services Management Console, the Amazon CLI, or the Amazon API.

Identity-based policies

Identity-based policies are JSON permissions policy documents that you can attach to an identity, such as an IAM user, group of users, or role. These policies control what actions users and roles can perform, on which resources, and under what conditions. To learn how to create an identity-based policy, see [Define custom IAM permissions with customer managed policies](#) in the *IAM User Guide*.

Identity-based policies can be further categorized as *inline policies* or *managed policies*. Inline policies are embedded directly into a single user, group, or role. Managed policies are standalone policies that you can attach to multiple users, groups, and roles in your Amazon Web Services account. Managed policies include Amazon managed policies and customer managed policies. To learn how to choose between a managed policy or an inline policy, see [Choose between managed policies and inline policies](#) in the *IAM User Guide*.

Resource-based policies

Resource-based policies are JSON policy documents that you attach to a resource. Examples of resource-based policies are IAM *role trust policies* and Amazon S3 *bucket policies*. In services that support resource-based policies, service administrators can use them to control access to a specific resource. For the resource where the policy is attached, the policy defines what actions a specified principal can perform on that resource and under what conditions. You must [specify a principal](#) in a resource-based policy. Principals can include accounts, users, roles, federated users, or Amazon Web Services services.

Resource-based policies are inline policies that are located in that service. You can't use Amazon managed policies from IAM in a resource-based policy.

Access control lists (ACLs)

Access control lists (ACLs) control which principals (account members, users, or roles) have permissions to access a resource. ACLs are similar to resource-based policies, although they do not use the JSON policy document format.

Amazon S3, Amazon WAF, and Amazon VPC are examples of services that support ACLs. To learn more about ACLs, see [Access control list \(ACL\) overview](#) in the *Amazon Simple Storage Service Developer Guide*.

Other policy types

Amazon supports additional, less-common policy types. These policy types can set the maximum permissions granted to you by the more common policy types.

- **Permissions boundaries** – A permissions boundary is an advanced feature in which you set the maximum permissions that an identity-based policy can grant to an IAM entity (IAM user or role). You can set a permissions boundary for an entity. The resulting permissions are the intersection of an entity's identity-based policies and its permissions boundaries. Resource-based policies that specify the user or role in the `Principal` field are not limited by the permissions boundary. An explicit deny in any of these policies overrides the allow. For more information about permissions boundaries, see [Permissions boundaries for IAM entities](#) in the *IAM User Guide*.
- **Service control policies (SCPs)** – SCPs are JSON policies that specify the maximum permissions for an organization or organizational unit (OU) in Amazon Organizations. Amazon Organizations is a service for grouping and centrally managing multiple Amazon Web Services accounts that your business owns. If you enable all features in an organization, then you can apply service control policies (SCPs) to any or all of your accounts. The SCP limits permissions for entities in

member accounts, including each Amazon Web Services account root user. For more information about Organizations and SCPs, see [Service control policies](#) in the *Amazon Organizations User Guide*.

- **Resource control policies (RCPs)** – RCPs are JSON policies that you can use to set the maximum available permissions for resources in your accounts without updating the IAM policies attached to each resource that you own. The RCP limits permissions for resources in member accounts and can impact the effective permissions for identities, including the Amazon Web Services account root user, regardless of whether they belong to your organization. For more information about Organizations and RCPs, including a list of Amazon Web Services services that support RCPs, see [Resource control policies \(RCPs\)](#) in the *Amazon Organizations User Guide*.
- **Session policies** – Session policies are advanced policies that you pass as a parameter when you programmatically create a temporary session for a role or federated user. The resulting session's permissions are the intersection of the user or role's identity-based policies and the session policies. Permissions can also come from a resource-based policy. An explicit deny in any of these policies overrides the allow. For more information, see [Session policies](#) in the *IAM User Guide*.

Multiple policy types

When multiple types of policies apply to a request, the resulting permissions are more complicated to understand. To learn how Amazon determines whether to allow a request when multiple policy types are involved, see [Policy evaluation logic](#) in the *IAM User Guide*.

Using Quick Suite with IAM

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Before you use IAM to manage access to Amazon Quick Suite, you should understand what IAM features are available to use with Amazon Quick Suite. To get a high-level view of how Amazon Quick Suite and other Amazon services work with IAM, see [Amazon Services That Work with IAM](#) in the *IAM User Guide*.

Topics

- [Amazon Quick Suite Policies \(identity-based\)](#)

- [Amazon Quick Suite policies \(resource-based\)](#)
- [Authorization based on Amazon Quick Suite tags](#)
- [Amazon Quick Suite IAM roles](#)

Amazon Quick Suite Policies (identity-based)

With IAM identity-based policies, you can specify allowed or denied actions and resources as well as the conditions under which actions are allowed or denied. Amazon Quick Suite supports specific actions, resources, and condition keys. To learn about all of the elements that you use in a JSON policy, see [IAM JSON Policy Elements Reference](#) in the *IAM User Guide*.

You can use Amazon root credentials or IAM user credentials to create an Amazon Quick Suite account. Amazon root and administrator credentials already have all of the required permissions for managing Amazon Quick Suite access to Amazon resources.

However, we recommend that you protect your root credentials, and instead use IAM user credentials. To do this, you can create a policy and attach it to the IAM user and roles that you plan to use for Amazon Quick Suite. The policy must include the appropriate statements for the Amazon Quick Suite administrative tasks you need to perform, as described in the following sections.

Important

Be aware of the following when working with Quick Suite and IAM policies:

- Avoid directly modifying a policy that was created by Quick Suite. When you modify it yourself, Quick Suite can't edit it. This inability can cause an issue with the policy. To fix this issue, delete the previously modified policy.
- If you get an error on permissions when you try to create an Amazon Quick Suite account, see [Actions Defined by Amazon Quick Suite](#) in the *IAM User Guide*.
- In some cases, you might have an Amazon Quick Suite account that you can't access even from the root account (for example, if you accidentally deleted its directory service). In this case, you can delete your old Amazon Quick Suite account, then recreate it. For more information, see [Deleting your Amazon Quick Suite subscription and closing the account](#).

Topics

- [Actions](#)

- [Resources](#)
- [Condition keys](#)
- [Examples](#)

Actions

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The `Action` element of a JSON policy describes the actions that you can use to allow or deny access in a policy. Policy actions usually have the same name as the associated Amazon API operation. There are some exceptions, such as *permission-only actions* that don't have a matching API operation. There are also some operations that require multiple actions in a policy. These additional actions are called *dependent actions*.

Include actions in a policy to grant permissions to perform the associated operation.

Policy actions in Amazon Quick Suite use the following prefix before the action: `quicksight:`. For example, to grant someone permission to run an Amazon EC2 instance with the Amazon EC2 `RunInstances` API operation, you include the `ec2:RunInstances` action in their policy. Policy statements must include either an `Action` or `NotAction` element. Amazon Quick Suite defines its own set of actions that describe tasks that you can perform with this service.

To specify multiple actions in a single statement, separate them with commas as follows:

```
"Action": [  
    "quicksight:action1",  
    "quicksight:action2"]
```

You can specify multiple actions using wildcards (*). For example, to specify all actions that begin with the word `Create`, include the following action:

```
"Action": "quicksight:Create*"
```

Amazon Quick Suite provides a number of Amazon Identity and Access Management (IAM) actions. All Amazon Quick Suite actions are prefixed with `quicksight:`, such as `quicksight:Subscribe`. For information about using Amazon Quick Suite actions in an IAM policy, see [IAM policy examples for Amazon Quick Suite](#).

To see the most up-to-date list of Amazon Quick Suite actions, see [Actions Defined by Amazon Quick Suite](#) in the *IAM User Guide*.

Resources

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Resource JSON policy element specifies the object or objects to which the action applies. Statements must include either a Resource or a NotResource element. As a best practice, specify a resource using its [Amazon Resource Name \(ARN\)](#). You can do this for actions that support a specific resource type, known as *resource-level permissions*.

For actions that don't support resource-level permissions, such as listing operations, use a wildcard (*) to indicate that the statement applies to all resources.

```
"Resource": "*"
```

Following is an example policy. It means that the caller with this policy attached, is able to invoke the CreateGroupMembership operation on any group, provided that the user name they are adding to the group is not user1.

```
{
  "Effect": "Allow",
  "Action": "quicksight:CreateGroupMembership",
  "Resource": "arn:aws-cn:quicksight:us-east-1:aws-account-id:group/default/*",
  "Condition": {
    "StringNotEquals": {
      "quicksight:UserName": "user1"
    }
  }
}
```

Some Amazon Quick Suite actions, such as those for creating resources, cannot be performed on a specific resource. In those cases, you must use the wildcard (*).

```
"Resource": "*"
```

Some API actions involve multiple resources. To specify multiple resources in a single statement, separate the ARNs with commas.

```
"Resource": [  
    "resource1",  
    "resource2"
```

To see a list of Amazon Quick Suite resource types and their Amazon Resource Names (ARNs), see [Resources Defined by Amazon Quick Suite](#) in the *IAM User Guide*. To learn with which actions you can specify the ARN of each resource, see [Actions Defined by Amazon Quick Suite](#).

Condition keys

Administrators can use Amazon JSON policies to specify who has access to what. That is, which **principal** can perform **actions** on what **resources**, and under what **conditions**.

The Condition element (or Condition *block*) lets you specify conditions in which a statement is in effect. The Condition element is optional. You can create conditional expressions that use [condition operators](#), such as equals or less than, to match the condition in the policy with values in the request.

If you specify multiple Condition elements in a statement, or multiple keys in a single Condition element, Amazon evaluates them using a logical AND operation. If you specify multiple values for a single condition key, Amazon evaluates the condition using a logical OR operation. All of the conditions must be met before the statement's permissions are granted.

You can also use placeholder variables when you specify conditions. For example, you can grant an IAM user permission to access a resource only if it is tagged with their IAM user name. For more information, see [IAM policy elements: variables and tags](#) in the *IAM User Guide*.

Amazon supports global condition keys and service-specific condition keys. To see all Amazon global condition keys, see [Amazon global condition context keys](#) in the *IAM User Guide*.

Amazon Quick Suite does not provide any service-specific condition keys, but it does support using some global condition keys. To see all Amazon global condition keys, see [Amazon Global Condition Context Keys](#) in the *IAM User Guide*.

Examples

To view examples of Amazon Quick Suite identity-based policies, see [Amazon Quick Suite Policies \(identity-based\)](#).

Amazon Quick Suite policies (resource-based)

Amazon Quick Suite doesn't support resource-based policies. However, you can use the Amazon Quick Suite console to configure access to other Amazon resources in your Amazon Web Services account.

Authorization based on Amazon Quick Suite tags

Amazon Quick Suite does not support tagging resources or controlling access based on tags.

Amazon Quick Suite IAM roles

An [IAM role](#) is an entity within your Amazon account that has specific permissions. You can use IAM roles to group permissions together to make it easier to manage user's access to Amazon Quick Suite actions.

Amazon Quick Suite doesn't support the following role features:

- Service-linked roles.
- Service roles.
- Temporary credentials (direct use): However, Amazon Quick Suite uses temporary credentials to allow users to assume an IAM role to access embedded dashboards. For more information, see [Embedded analytics for Amazon Quick Suite](#).

For more information on how Amazon Quick Suite uses IAM roles, see [Using Amazon Quick Suite with IAM](#) and [IAM policy examples for Amazon Quick Suite](#).

Passing IAM roles to Quick Suite

Applies to: Enterprise Edition

When your IAM users sign up for Quick Suite, they can choose to use the Amazon Quick Suite-managed role (this is the default role). Or they can pass an existing IAM role to Amazon Quick Suite.

Use the sections below to pass existing IAM roles to Amazon Quick Suite

Topics

- [Prerequisites](#)

- [Attaching additional policies](#)
- [Using existing IAM roles in Quick Suite](#)

Prerequisites

For your users to pass IAM roles to Amazon Quick Suite, your administrator needs to complete the following tasks:

- **Create an IAM role.** For more information about creating IAM roles, see [Creating IAM roles](#) in the *IAM User Guide*.
- **Attach a trust policy to your IAM role that allows Amazon Quick Suite to assume the role.** Use the following example to create a trust policy for the role. The following example trust policy allows the Quick Suite principal to assume the IAM role that it's attached to.

For more information about creating IAM trust policies and attaching them to roles, see [Modifying a Role \(Console\)](#) in the *IAM User Guide*.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "quicksight.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

- **Assign the following IAM permissions to your administrator (IAM users or roles):**
 - `quicksight:UpdateResourcePermissions` – This grants IAM users who are Amazon Quick Suite administrators the permission to update resource-level permissions in Amazon Quick Suite. For more information about resource types defined by Amazon Quick Suite, see [Actions, resources, and condition keys for Quick Suite](#) in the *IAM User Guide*.
 - `iam:PassRole` – This grants users permission to pass roles to Amazon Quick Suite. For more information, see [Granting a user permissions to pass a role to an Amazon service](#) in the *IAM User Guide*.

- `iam:ListRoles` – (Optional) This grants users permission to see a list of existing roles in Amazon Quick Suite. If this permission is not provided, they can use an ARN to use existing IAM roles.

Following is an example IAM permissions policy that allows managing resource-level permissions, listing IAM roles, and passing IAM roles in Quick Suite.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "iam:ListRoles",
      "Resource": "arn:aws-cn:iam::account-id:role:*"
    },
    {
      "Effect": "Allow",
      "Action": "iam:PassRole",
      "Resource": "arn:aws-cn:iam::account-id:role/path/role-name",
      "Condition": {
        "StringEquals": {
          "iam:PassedToService": [
            "quicksight.amazonaws.com"
          ]
        }
      }
    },
    {
      "Effect": "Allow",
      "Action": "quicksight:UpdateResourcePermissions",
      "Resource": "*"
    }
  ]
}
```

For more examples of IAM policies that you can use with Amazon Quick Suite, see [IAM policy examples for Amazon Quick Suite](#).

For more information about assigning permissions policies to users or user groups, see [Changing permissions for an IAM user](#) in the *IAM User Guide*.

Attaching additional policies

If you're using another Amazon service, such as Amazon Athena or Amazon S3, you can create a permissions policy that grants Amazon Quick Suite permission to perform specific actions. You can then attach the policy to the IAM roles that you later pass to Amazon Quick Suite. The following are examples of how you can set up and attach additional permissions policies to your IAM roles.

For an example managed policy for Amazon Quick Suite in Athena, see [AWSQuicksightAthenaAccess Managed Policy](#) in the *Amazon Athena User Guide*. IAM users can access this role in Amazon Quick Suite using the following ARN: `arn:aws-cn:iam::aws:policy/service-role/AWSQuicksightAthenaAccess`.

The following is an example of a permissions policy for Amazon Quick Suite in Amazon S3. For more information about using IAM with Amazon S3, see [Identity and access management in Amazon S3](#) in the *Amazon S3 User Guide*.

For information on how to create cross-account access from Amazon Quick Suite to an Amazon S3 bucket in another account, see [How do I set up cross-account access from Quick Suite to an Amazon S3 bucket in another account?](#) in the Amazon Knowledge Center.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3:ListAllMyBuckets",
      "Resource": "arn:aws-cn:s3:::*"
    },
    {
      "Action": [
        "s3:ListBucket"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws-cn:s3:::aws-athena-query-results-us-west-2-123456789"
      ]
    },
    {
      "Action": [
        "s3:GetObject",
        "s3:GetObjectVersion"
      ],
```

```

        "Effect": "Allow",
        "Resource": [
            "arn:aws-cn:s3:::aws-athena-query-results-us-west-2-123456789/*"
        ],
    },
    {
        "Action": [
            "s3:ListBucketMultipartUploads",
            "s3:GetBucketLocation"
        ],
        "Effect": "Allow",
        "Resource": [
            "arn:aws-cn:s3:::aws-athena-query-results-us-west-2-123456789"
        ],
    },
    {
        "Effect": "Allow",
        "Action": [
            "s3:PutObject",
            "s3:AbortMultipartUpload",
            "s3:ListMultipartUploadParts"
        ],
        "Resource": [
            "arn:aws-cn:s3:::aws-athena-query-results-us-west-2-123456789/*"
        ],
    }
]
}

```

Using existing IAM roles in Quick Suite

If you're a Amazon Quick Suite administrator and have permissions to update Amazon Quick Suite resources and pass IAM roles, you can use existing IAM roles in Amazon Quick Suite. To learn more about the prerequisites for passing IAM roles in Amazon Quick Suite, see the [Prerequisites](#) outlined in the previous list.

Use the following procedure to learn how to pass IAM roles in Amazon Quick Suite.

To use an existing IAM role in Amazon Quick Suite

1. In Amazon Quick Suite, choose your account name in the navigation bar at top right and choose **Manage QuickSight**.

2. On the **Manage Amazon Quick Suite** page that opens, choose **Security & Permissions** in the menu at left.
3. In the **Security & Permissions** page that opens, under **Amazon Quick Suite access to Amazon services**, choose **Manage**.
4. For **IAM role**, choose **Use an existing role**, and then do one of the following:
 - Choose the role that you want to use from the list.
 - Or, if you don't see a list of existing IAM roles, you can enter the IAM ARN for the role in the following format: `arn:aws-cn:iam::account-id:role/path/role-name`.
5. Choose **Save**.

IAM policy examples for Quick Suite

This section provides examples of IAM policies that you can use with Quick Suite.

IAM identity-based policies for Quick Suite

This section shows examples of identity-based policies to use with Quick Suite.

Topics

- [IAM identity-based policies for Amazon Quick Suite IAM console administration](#)

IAM identity-based policies for Amazon Quick Suite IAM console administration

The following example shows the IAM permissions needed for Amazon Quick Suite IAM console administration actions.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Allow",
      "Action": [
        "quicksight:*",
        "iam:AttachRolePolicy",
        "iam:DetachRolePolicy",
        "iam:ListAttachedRolePolicies",
```



```

        "iam:GetPolicy",
        "iam:CreatePolicyVersion",
        "iam:DeletePolicyVersion",
        "iam:GetPolicyVersion",
        "iam:ListPolicyVersions",
        "iam:DeleteRole",
        "iam:CreateRole",
        "iam:GetRole",
        "iam:ListRoles",
        "iam:CreatePolicy",
        "iam:ListEntitiesForPolicy",
        "iam:listPolicies",
        "s3:ListAllMyBuckets",
        "athena:ListDataCatalogs",
        "athena:GetDataCatalog"
    ],
    "Resource": [
        "*"
    ]
}
]
}

```

IAM identity-based policies for Quick Suite: dashboards

The following example shows an IAM policy that allows dashboard sharing and embedding for specific dashboards.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "quicksight:RegisterUser",
      "Resource": "*",
      "Effect": "Allow"
    },
    {
      "Action": "quicksight:GetDashboardEmbedUrl",
      "Resource": "arn:aws-cn:quicksight:us-west-2:111122223333:dashboard/1a1ac2b2-3fc3-4b44-5e5d-c6db6778df89",
      "Effect": "Allow"
    }
  ]
}

```

```
}
```

IAM identity-based policies for Quick Suite: namespaces

The following examples show IAM policies that allow a Amazon Quick Suite administrator to create or delete namespaces.

Creating namespaces

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ds:AuthorizeApplication",
        "ds:UnauthorizeApplication",
        "ds>DeleteDirectory",
        "ds>CreateIdentityPoolDirectory",
        "ds:DescribeDirectories",
        "quicksight:CreateNamespace"
      ],
      "Resource": "*"
    }
  ]
}
```

Deleting namespaces

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ds:UnauthorizeApplication",
        "ds>DeleteDirectory",
        "ds:DescribeDirectories",
        "quicksight>DeleteNamespace"
      ],
      "Resource": "*"
    }
  ]
}
```

```
    }  
  ]  
}
```

IAM identity-based policies for Quick Suite: custom permissions

The following example shows an IAM policy that allows a Amazon Quick Suite administrator or a developer to manage custom permissions.

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "quicksight:*CustomPermissions"  
      ],  
      "Resource": "*"   
    }  
  ]  
}
```

The following example shows another way to grant the same permissions as shown in the previous example.

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "quicksight:CreateCustomPermissions",  
        "quicksight:DescribeCustomPermissions",  
        "quicksight:ListCustomPermissions",  
        "quicksight:UpdateCustomPermissions",  
        "quicksight>DeleteCustomPermissions"  
      ],  
      "Resource": "*"   
    }  
  ]  
}
```

IAM identity-based policies for Quick Suite: customizing email report templates

The following example shows a policy that allows viewing, updating, and creating email report templates in Amazon Quick Suite, as well as obtaining verification attributes for an Amazon Simple Email Service identity. This policy allows a Amazon Quick Suite administrator to create and update custom email report templates, and to confirm that any custom email address they want to send email reports from is a verified identity in SES.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:DescribeAccountCustomization",
        "quicksight:CreateAccountCustomization",
        "quicksight:UpdateAccountCustomization",
        "quicksight:DescribeEmailCustomizationTemplate",
        "quicksight:CreateEmailCustomizationTemplate",
        "quicksight:UpdateEmailCustomizationTemplate",
        "ses:GetIdentityVerificationAttributes"
      ],
      "Resource": "*"
    }
  ]
}
```

IAM identity-based policies for Quick Suite: create an Enterprise account with Amazon Quick Suite managed users

The following example shows a policy that allows Amazon Quick Suite admins to create an Enterprise edition Amazon Quick Suite account with Amazon Quick Suite managed users.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Allow",
      "Action": [
        "quicksight:*",
        "iam:AttachRolePolicy",

```

```

        "iam:DetachRolePolicy",
        "iam:ListAttachedRolePolicies",
        "iam:GetPolicy",
        "iam:CreatePolicyVersion",
        "iam:DeletePolicyVersion",
        "iam:GetPolicyVersion",
        "iam:ListPolicyVersions",
        "iam:DeleteRole",
        "iam:CreateRole",
        "iam:GetRole",
        "iam:ListRoles",
        "iam:CreatePolicy",
        "iam:ListEntitiesForPolicy",
        "iam:listPolicies",
        "s3:ListAllMyBuckets",
        "athena:ListDataCatalogs",
        "athena:GetDataCatalog",
        "ds:AuthorizeApplication",
        "ds:UnauthorizeApplication",
        "ds:CheckAlias",
        "ds:CreateAlias",
        "ds:DescribeDirectories",
        "ds:DescribeTrusts",
        "ds>DeleteDirectory",
        "ds:CreateIdentityPoolDirectory"
    ],
    "Resource": [
        "*"
    ]
}
]
}

```

IAM identity-based policies for Quick Suite: creating users

The following example shows a policy that allows creating Amazon Quick Suite users only. For `quicksight:CreateReader`, `quicksight:CreateUser`, and `quicksight:CreateAdmin`, you can limit the permissions to **"Resource": "arn:aws-cn:quicksight::<YOUR_AWS_ACCOUNTID>:user/\${aws:userid}"**. For all other permissions described in this guide, use **"Resource": "*"**. The resource you specify limits the scope of the permissions to the specified resource.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "quicksight:CreateUser"
      ],
      "Effect": "Allow",
      "Resource": "arn:aws-cn:quicksight::<YOUR_AWS_ACCOUNTID>:user/
${aws:userid}"
    }
  ]
}
```

IAM identity-based policies for Quick Suite: creating and managing groups

The following example shows a policy that allows Amazon Quick Suite administrators and developers to create and manage groups.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:ListGroup",
        "quicksight:CreateGroup",
        "quicksight:SearchGroups",
        "quicksight:ListGroupMemberships",
        "quicksight:CreateGroupMembership",
        "quicksight>DeleteGroupMembership",
        "quicksight:DescribeGroupMembership",
        "quicksight:ListUsers"
      ],
      "Resource": "*"
    }
  ]
}
```

IAM identity-based policies for Quick Suite: All access for Standard edition

The following example for Amazon Quick Suite Standard edition shows a policy that allows subscribing and creating authors and readers. This example explicitly denies permission to unsubscribe from Amazon Quick Suite.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ds:AuthorizeApplication",
        "ds:UnauthorizeApplication",
        "ds:CheckAlias",
        "ds:CreateAlias",
        "ds:DescribeDirectories",
        "ds:DescribeTrusts",
        "ds>DeleteDirectory",
        "ds>CreateIdentityPoolDirectory",
        "iam:ListAccountAliases",
        "quicksight:CreateUser",
        "quicksight:DescribeAccountSubscription",
        "quicksight:Subscribe"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "Action": "quicksight:Unsubscribe",
      "Resource": "*"
    }
  ]
}
```

IAM identity-based policies for Quick Suite: All access for Enterprise edition with IAM Identity Center (Pro roles)

The following example for Amazon Quick Suite Enterprise edition shows a policy that allows a Amazon Quick Suite user to subscribe to Amazon Quick Suite, create users, and manage Active Directory in a Amazon Quick Suite account that is integrated with IAM Identity Center.

This policy also allows users to subscribe to Amazon Quick Suite Pro roles that grant access to Amazon Q in Quick Suite Generative BI capabilities. For more information about Pro roles in Amazon Quick Suite, see [Get started with Generative BI](#).

This example explicitly denies permission to unsubscribe from Amazon Quick Suite.

```
{
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Allow",
      "Action": [
        "quicksight:*",
        "iam:AttachRolePolicy",
        "iam:DetachRolePolicy",
        "iam:ListAttachedRolePolicies",
        "iam:GetPolicy",
        "iam:CreatePolicyVersion",
        "iam:DeletePolicyVersion",
        "iam:GetPolicyVersion",
        "iam:ListPolicyVersions",
        "iam>DeleteRole",
        "iam:CreateRole",
        "iam:GetRole",
        "iam:ListRoles",
        "iam:CreatePolicy",
        "iam:ListEntitiesForPolicy",
        "iam:listPolicies",
        "iam:CreateServiceLinkedRole",
        "s3:ListAllMyBuckets",
        "athena:ListDataCatalogs",
        "athena:GetDataCatalog",
        "sso:DescribeApplication",
        "sso:DescribeInstance",
        "sso:CreateApplication",
        "sso:PutApplicationAuthenticationMethod",
        "sso:PutApplicationGrant",
        "sso>DeleteApplication",
        "sso:SearchGroups",
        "sso:GetProfile",
        "sso:CreateApplicationAssignment",
        "sso>DeleteApplicationAssignment",
        "sso:ListInstances",
```



```

        "sso:DescribeRegisteredRegions",
        "organizations:DescribeOrganization",
        "user-subscriptions:CreateClaim",
        "user-subscriptions:UpdateClaim",
        "sso-directory:DescribeUser",
        "sso:ListApplicationAssignments",
        "sso-directory:DescribeGroup",
        "organizations:ListAWSServiceAccessForOrganization",
        "identitystore:DescribeUser",
        "identitystore:DescribeGroup"
    ],
    "Resource": [
        "*"
    ]
}
]
}

```

IAM identity-based policies for Quick Suite: All access for Enterprise edition with IAM Identity Center

The following example for Amazon Quick Suite Enterprise edition shows a policy that allows subscribing, creating users, and managing Active Directory in a Amazon Quick Suite account that is integrated with IAM Identity Center.

This policy does not grant permissions to create Pro roles in Amazon Quick Suite. To create a policy that grants permission to subscribe to Pro roles in Amazon Quick Suite, see [IAM identity-based policies for Amazon Quick Suite: All access for Enterprise edition with IAM Identity Center \(Pro roles\)](#).

This example explicitly denies permission to unsubscribe from Amazon Quick Suite.

```

{
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Allow",
      "Action": [
        "quicksight:*",
        "iam:AttachRolePolicy",
        "iam:DetachRolePolicy",
        "iam:ListAttachedRolePolicies",

```

```

        "iam:GetPolicy",
        "iam:CreatePolicyVersion",
        "iam:DeletePolicyVersion",
        "iam:GetPolicyVersion",
        "iam:ListPolicyVersions",
        "iam:DeleteRole",
        "iam:CreateRole",
        "iam:GetRole",
        "iam:ListRoles",
        "iam:CreatePolicy",
        "iam:ListEntitiesForPolicy",
        "iam:listPolicies",
        "s3:ListAllMyBuckets",
        "athena:ListDataCatalogs",
        "athena:GetDataCatalog",
        "sso:DescribeApplication",
        "sso:DescribeInstance",
        "sso:CreateApplication",
        "sso:PutApplicationAuthenticationMethod",
        "sso:PutApplicationGrant",
        "sso:DeleteApplication",
        "sso:SearchGroups",
        "sso:GetProfile",
        "sso:CreateApplicationAssignment",
        "sso:DeleteApplicationAssignment",
        "sso:ListInstances",
        "sso:DescribeRegisteredRegions",
        "organizations:DescribeOrganization"
    ],
    "Resource": [
        "*"
    ]
}

```

IAM identity-based policies for Quick Suite: all access for Enterprise edition with Active Directory

The following example for Amazon Quick Suite Enterprise edition shows a policy that allows subscribing, creating users, and managing Active Directory in a Amazon Quick Suite account that uses Active Directory for identity management. This example explicitly denies permission to unsubscribe from Amazon Quick Suite.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ds:AuthorizeApplication",
        "ds:UnauthorizeApplication",
        "ds:CheckAlias",
        "ds:CreateAlias",
        "ds:DescribeDirectories",
        "ds:DescribeTrusts",
        "ds>DeleteDirectory",
        "ds:CreateIdentityPoolDirectory",
        "iam:ListAccountAliases",
        "quicksight:CreateAdmin",
        "quicksight:Subscribe",
        "quicksight:GetGroupMapping",
        "quicksight:SearchDirectoryGroups",
        "quicksight:SetGroupMapping"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "Action": "quicksight:Unsubscribe",
      "Resource": "*"
    }
  ]
}
```

IAM identity-based policies for Quick Suite: active directory groups

The following example shows an IAM policy that allows Active Directory group management for an Amazon Quick Suite Enterprise edition account.

```
{
  "Statement": [
    {
      "Action": [
        "ds:DescribeTrusts",
        "quicksight:GetGroupMapping",
        "quicksight:SearchDirectoryGroups",
```

```

        "quicksight:SetGroupMapping"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ],
  "Version": "2012-10-17"
}

```

IAM identity-based policies for Quick Suite: using the admin asset management console

The following example shows an IAM policy that allows access to the admin asset management console.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:SearchGroups",
        "quicksight:SearchUsers",
        "quicksight:ListNamespaces",
        "quicksight:DescribeAnalysisPermissions",
        "quicksight:DescribeDashboardPermissions",
        "quicksight:DescribeDataSetPermissions",
        "quicksight:DescribeDataSourcePermissions",
        "quicksight:DescribeFolderPermissions",
        "quicksight:ListAnalyses",
        "quicksight:ListDashboards",
        "quicksight:ListDataSets",
        "quicksight:ListDataSources",
        "quicksight:ListFolders",
        "quicksight:SearchAnalyses",
        "quicksight:SearchDashboards",
        "quicksight:SearchFolders",
        "quicksight:SearchDatasets",
        "quicksight:SearchDatasources",
        "quicksight:UpdateAnalysisPermissions",
        "quicksight:UpdateDashboardPermissions",
        "quicksight:UpdateDataSetPermissions",
        "quicksight:UpdateDataSourcePermissions",
        "quicksight:UpdateFolderPermissions"
      ]
    }
  ]
}

```

```
    ],  
    "Resource": "*"    
  }  
]  
}
```

IAM identity-based policies for Quick Suite: using the admin key management console

The following example shows an IAM policy that allows access to the admin key management console.

```
{  
  "Version": "2012-10-17",  
  "Statement": [  
    {  
      "Effect": "Allow",  
      "Action": [  
        "quicksight:DescribeKeyRegistration",  
        "quicksight:UpdateKeyRegistration",  
        "quicksight:ListKMSKeysForUser",  
        "kms:CreateGrant",  
        "kms:ListGrants",  
        "kms:ListAliases"  
      ],  
      "Resource": "*"    
    }  
  ]  
}
```

The "quicksight:ListKMSKeysForUser" and "kms:ListAliases" permissions are required to access customer managed keys from the Amazon Quick Suite console. "quicksight:ListKMSKeysForUser" and "kms:ListAliases" are not required to use the Amazon Quick Suite key management APIs.

To specify which keys you want a user to be able to access, add the ARNs of the keys that you want the user to access to the UpdateKeyRegistration condition with the quicksight:KmsKeyArns condition key. Users can only access the keys specified in UpdateKeyRegistration. For more information about supported condition keys for Amazon Quick Suite, see [Condition keys for Amazon Quick Suite](#).

The example below grants Describe permissions for all CMKs that are registered to a Amazon Quick Suite account and Update permissions to specific CMKs that are registered to the Amazon Quick Suite account.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:DescribeKeyRegistration"
      ],
      "Resource": "arn:aws:quicksight:us-west-2:123456789012:*"
    },
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:UpdateKeyRegistration"
      ],
      "Resource": "arn:aws:quicksight:us-west-2:123456789012:*",
      "Condition": {
        "ForAllValues:StringEquals": {
          "quicksight:KmsKeyArns": [
            "arn:aws:kms:us-west-2:123456789012:key/key-id-of-key1",
            "arn:aws:kms:us-west-2:123456789012:key/key-id-of-key2",
            "..."
          ]
        }
      }
    },
    {
      "Effect": "Allow",
      "Action": [
        "kms:CreateGrant",
        "kms:ListGrants"
      ],
      "Resource": "arn:aws:kms:us-west-2:123456789012:key/*"
    }
  ]
}
```

Amazon resources Quick Suite: scoping policies in Enterprise edition

The following example for Amazon Quick Suite Enterprise edition shows a policy that allows setting default access to Amazon resources and scoping policies for permissions to Amazon resources.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "quicksight:*IAMPolicyAssignment*",
        "quicksight:AccountConfigurations"
      ],
      "Effect": "Allow",
      "Resource": "*"
    }
  ]
}
```

Provisioning users for Amazon Quick Suite

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators and Amazon Quick Suite administrators

Self-provisioning an Amazon Quick Suite administrator

Amazon Quick Suite administrators are users who can also manage Amazon Quick Suite features such as account settings and accounts. They can also purchase additional Amazon Quick Suite user subscriptions, purchase [SPICE](#), and cancel the subscription to Amazon Quick Suite for your Amazon Web Services account.

You can use an Amazon user or group policy to give users the ability to add themselves as administrators of Amazon Quick Suite. Users that have been granted this ability can only add themselves as administrators and can't use this policy to add others. Their accounts become active and billable the first time that they open Amazon Quick Suite. To set up self-provisioning, give these users permission to use the `quicksight:CreateAdmin` action.

Alternatively, you can use the following procedure to use the console to set or create the administrator for Amazon Quick Suite.

To make a user the Amazon Quick Suite administrator

1. Create the Amazon user:
 - Use IAM to create the user that you want to be the administrator of Amazon Quick Suite. Alternatively, identify an existing user in IAM for the administrator role. You can also put the user inside a new group, for manageability.
 - Grant the user (or group) sufficient permissions.
2. Sign in to your Amazon Web Services Management Console with the target user's credentials.
3. Go to <http://quicksight.aws.amazon.com/sn/console/get-user-email>, type in the target user's email address, and choose **Continue**.

On success, the target user is now an administrator in Amazon Quick Suite.

Self-provisioning an Amazon Quick Suite author

Amazon Quick Suite authors can create data sources, datasets, analyses, and dashboards. They can share analyses and dashboards with other Amazon Quick Suite users in your Amazon Quick Suite account. However, they don't have access to the **Manage Amazon Quick Suite** menu. They can't change account settings, manage accounts, purchase additional Amazon Quick Suite user subscriptions or [SPICE](#) capacity, or cancel the subscription to Amazon Quick Suite for your Amazon Web Services account. Author Pro users can additionally create content using natural language, build knowledge bases, configure actions, and access advanced automation capabilities.

You can use an Amazon user or group policy to give users the ability to create an Amazon Quick Suite author account for themselves. Their accounts become active and billable the first time they open Amazon Quick Suite. To set up self-provisioning, you need to give them permission to use the `quicksight:CreateUser` action.

Self-provisioning an Amazon Quick Suite read-only user

Amazon Quick Suite read-only users or *readers* can view and manipulate dashboards that are shared with them, but they can't make any changes or save a dashboard for further analysis. Amazon Quick Suite readers can't create data sources, datasets, analyses, or visuals. They can't do any administrative tasks. Choose this role for people who are consumers of the dashboards but

don't author their own analysis, for example, executives. Reader Pro users have access to advanced features including AI chat agents, collaborative spaces, flows, and extensions.

If you are using Microsoft Active Directory with Amazon Quick Suite, you can manage read-only permissions by using a group. Otherwise, you can bulk-invite users to use Amazon Quick Suite. You can also use an Amazon user or group policy to give people the ability to create an Amazon Quick Suite reader account for themselves.

Reader accounts become active and billable the first time they open Amazon Quick Suite. If you decide to upgrade or downgrade a user, billing for that user is prorated for the month. To set up self-provisioning, you need to give them permission to use the `quicksight:CreateReader` action.

Readers that are used to automatically or programmatically refresh dashboards for near real-time use cases must choose capacity pricing. For readers under user pricing, each reader is limited to manual use by one individual only.

Troubleshooting Quick Suite identity and access

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Use the following information to help you diagnose and fix common issues that you might encounter when working with Amazon Quick Suite and IAM.

Topics

- [I am not authorized to perform an action in Amazon Quick Suite](#)
- [I am not authorized to perform iam:PassRole](#)
- [I want to allow people outside of my Amazon account to access my Amazon Quick Suite resources](#)

I am not authorized to perform an action in Amazon Quick Suite

If the Amazon Web Services Management Console tells you that you're not authorized to perform an action, then you must contact your administrator for assistance.

The following example error occurs when the mateojackson IAM user tries to use the console to view details about a *widget* but does not have quicksight:*GetWidget* permissions.

```
User: arn:aws-cn:iam::123456789012:user/mateojackson is not authorized to perform:
quicksight:GetWidget on resource: my-example-widget
```

In this case, Mateo asks his administrator to update his policies to allow him to access the *my-example-widget* resource using the quicksight:*GetWidget* action.

I am not authorized to perform iam:PassRole

If you receive an error that you're not authorized to perform the iam:PassRole action, your policies must be updated to allow you to pass a role to Amazon Quick Suite.

Some Amazon Web Services services allow you to pass an existing role to that service instead of creating a new service role or service-linked role. To do this, you must have permissions to pass the role to the service.

The following example error occurs when an IAM user named marymajor tries to use the console to perform an action in Amazon Quick Suite. However, the action requires the service to have permissions that are granted by a service role. Mary does not have permissions to pass the role to the service.

```
User: arn:aws-cn:iam::123456789012:user/marymajor is not authorized to perform:
iam:PassRole
```

In this case, Mary's policies must be updated to allow her to perform the iam:PassRole action.

If you need help, contact your Amazon administrator. Your administrator is the person who provided you with your sign-in credentials.

I want to allow people outside of my Amazon account to access my Amazon Quick Suite resources

You can create a role that users in other accounts or people outside of your organization can use to access your resources. You can specify who is trusted to assume the role. For services that support resource-based policies or access control lists (ACLs), you can use those policies to grant people access to your resources.

To learn more, consult the following:

- To learn whether Amazon Quick Suite supports these features, see [Using Quick Suite with IAM](#).
- To learn how to provide access to your resources across Amazon Web Services accounts that you own, see [Providing access to an IAM user in another Amazon Web Services account that you own](#) in the *IAM User Guide*.
- To learn how to provide access to your resources to third-party Amazon Web Services accounts, see [Providing access to Amazon Web Services accounts owned by third parties](#) in the *IAM User Guide*.
- To learn how to provide access through identity federation, see [Providing access to externally authenticated users \(identity federation\)](#) in the *IAM User Guide*.
- To learn the difference between using roles and resource-based policies for cross-account access, see [Cross account resource access in IAM](#) in the *IAM User Guide*.

Using IAM Identity Center

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators and Amazon Quick Suite administrators

Amazon Quick Suite Enterprise edition integrates with your existing directories, using either Microsoft Active Directory or single sign-on (IAM Identity Center) using Security Assertion Markup Language (SAML). You can use Amazon Identity and Access Management (IAM) to further enhance your security, or for custom options such as embedding dashboards.

In Quick Suite Standard edition, you can manage users entirely within Quick Suite. If you prefer, you can integrate with your existing users, groups, and roles in IAM.

You can use the following tools for identity and access to Amazon Quick Suite:

- [IAM Identity Center](#) (Enterprise edition only)
- [IAM federation](#) (Standard and Enterprise editions)
- [Amazon Directory Service for Microsoft Active Directory](#) (Enterprise edition only)
- [SAML-based single sign-on](#) (Standard and Enterprise edition)
- [Multifactor authentication \(MFA\)](#) (Standard and Enterprise edition)

Note

In the regions listed below, Amazon Quick Suite accounts can only use [IAM Identity Center](#) for identity and access management.

- af-south-1 Africa (Cape Town)
- ap-southeast-3 Asia Pacific (Jakarta)
- eu-south-1 Europe (Milan)
- eu-central-2 Europe (Zurich)

IAM Identity Center helps you securely create or connect your workforce identities and manage their access across Amazon accounts and applications.

Before you integrate your Amazon Quick Suite account with IAM Identity Center, set up IAM Identity Center in your Amazon account. If you haven't set up IAM Identity Center in your Amazon organization, see [Getting started](#) in the *Amazon IAM Identity Center User Guide*.

If you want to configure an external identity provider with IAM Identity Center, see [Supported identity providers](#) to view a list of supported identity providers' configuration steps.

Topics

- [Configure your Amazon Quick Suite account with IAM Identity Center](#)

Configure your Amazon Quick Suite account with IAM Identity Center

Applies to: Enterprise Edition

Intended audience: System administrators

IAM Identity Center helps you securely create or configure your existing workforce identities and manage their access across Amazon accounts and applications. IAM Identity Center is the recommended approach for workforce authentication and authorization on Amazon for organizations of any size and type. To learn more about IAM Identity Center, see [Amazon IAM Identity Center](#).

Configure Amazon Quick Suite and IAM Identity Center so that you can sign up for a new Amazon Quick Suite account with an IAM Identity Center configured identity source. With IAM Identity Center, you can configure your external identity provider as an identity source. You can also use IAM Identity Center as an identity store if you don't want to use a third-party identity provider with Amazon Quick Suite. Identity methods can't be changed after your account is created.

When you integrate your Amazon Quick Suite account with IAM Identity Center, Amazon Quick Suite account administrators can create a new Amazon Quick Suite account that automatically has the identity provider's groups available. This simplifies asset sharing at scale in Amazon Quick Suite.

Access to some sections of the Amazon Quick Suite administration console is restricted by IAM permissions. The following table summarizes the admin actions that you can perform in Amazon Quick Suite based on the access type that you choose.

To learn more how to sign up for an Amazon Quick Suite account with IAM Identity Center, see [Signing up for an Amazon Quick Suite subscription](#).

Admin action	IAM permissions	Amazon Quick Suite admin role permissions
Manage assets	Yes	No
Security & permissions	Yes	No
Manage VPC connections	Yes	No
KMS keys	Yes	No
Account settings	Yes	No
Account customization	No	Yes
Manage users	Yes (IAM Identity Center users)	Yes (Amazon Quick Suite and IAM users)
Your subscriptions	No	Yes
Mobile settings	No	Yes

Admin action	IAM permissions	Amazon Quick Suite admin role permissions
Domains and embedding	No	Yes
SPICE capacity	No	Yes

The Amazon Quick Suite mobile app is not supported with Amazon Quick Suite accounts that are integrated with IAM Identity Center.

Considerations

The following actions permanently remove the ability for Amazon Quick Suite users to sign into Amazon Quick Suite. Amazon Quick Suite does not recommend that Amazon Quick Suite users perform these actions.

- Disabling or deleting the Amazon Quick Suite application in the IAM Identity Center console. If you want to delete your Amazon Quick Suite account, see [Closing your Amazon Quick Suite account](#).
- Migrating the Amazon Quick Suite account that contains your IAM Identity Center configuration to an Amazon Organization that does not contain the IAM Identity Center instance that your Amazon Quick Suite account is configured to.
- Deleting the IAM Identity Center instance that is configured to your Amazon Quick Suite account.
- Editing IAM Identity Center application attributes, for example the **requires assignment** attribute.

IAM federation

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

⚠ Important

Amazon Quick Suite recommends that you integrate new Amazon Quick Suite subscriptions with IAM Identity Center for identity management. This IAM identity federation user guide is provided as a reference for existing account configurations. For more information on integrating your Amazon Quick Suite account with IAM Identity Center, see [Configure your Amazon Quick Suite account with IAM Identity Center](#).

ℹ Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick Suite.

Amazon Quick Suite supports identity federation in both Standard and Enterprise editions. When you use federated users, you can manage users with your enterprise identity provider (IdP) and use Amazon Identity and Access Management (IAM) to authenticate users when they sign in to Quick Suite. You can use a third-party identity provider that supports Security Assertion Markup Language 2.0 (SAML 2.0) to provide an onboarding flow for your Amazon Quick Suite users. Such identity providers include Microsoft Active Directory Federation Services, Okta, and Ping One Federation Server. With identity federation, your users get one-click access to their Amazon Quick Suite applications using their existing identity credentials. You also have the security benefit of identity authentication by your identity provider. You can control which users have access to Amazon Quick Suite using your existing identity provider.

Topics

- [Initiating sign-on from the identity provider \(IdP\)](#)
- [Setting up IdP federation using IAM and Amazon Quick Suite](#)
- [Initiating sign-on from Quick Suite](#)
- [Setting up service provider-initiated federation with Quick Suite Enterprise edition](#)
- [Configuring email syncing for federated users in Quick Suite](#)
- [Tutorial: Amazon Quick Suite and IAM identity federation](#)

Initiating sign-on from the identity provider (IdP)

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick Suite.


In this scenario, your users initiate the sign-on process from the identity provider's portal. After the users are authenticated, they sign in to Amazon Quick Suite. After Quick Suite checks that they are authorized, your users can access Quick Suite.

Beginning with a user signing into the IdP, authentication flows through these steps:

1. The user browses to `https://applications.example.com` and signs on to the IdP. At this point, the user isn't signed in to the service provider.
2. The federation service and the IdP authenticate the user:
 - a. The federation service requests authentication from the organization's identity store.
 - b. The identity store authenticates the user and returns the authentication response to the federation service.
 - c. When authentication is successful, the federation service posts the SAML assertion to the user's browser.
3. The user opens Amazon Quick Suite:
 - a. The user's browser posts the SAML assertion to the Amazon Sign-In SAML endpoint (`https://signin.aws.amazon.com/saml`).
 - b. Amazon Sign-In receives the SAML request, processes the request, authenticates the user, and forwards the authentication token to the Amazon Quick Suite service.
4. Amazon Quick Suite accepts the authentication token from Amazon and presents Amazon Quick Suite to the user.

From the user's perspective, the process happens transparently. The user starts at your organization's internal portal and lands at an Amazon Quick Suite application portal, without ever having to supply any Amazon credentials.

In the following diagram, you can find an authentication flow between Amazon Quick Suite and a third-party identity provider (IdP). In this example, the administrator has set up a sign-in page to access Amazon Quick Suite, called `applications.example.com`. When a user signs in, the sign-in page posts a request to a federation service that complies with SAML 2.0. The end user initiates authentication from the sign-on page of the IdP.



For information from some common providers, see the following third-party documentation:

- CA – [Enabling SAML 2.0 HTTP Post Binding](#)
- Okta – [Planning a SAML deployment](#)
- Ping – [Amazon integrations](#)

Use the following topics to understand using an existing federation with Amazon:

- [Identity federation in Amazon](#) on the Amazon website

- [Providing access to externally authenticated users \(identity federation\)](#) in the *IAM User Guide*
- [Enabling SAML 2.0 federated users to access the Amazon Management Console](#) in the *IAM User Guide*

Setting up IdP federation using IAM and Amazon Quick Suite

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick Suite.

You can use an Amazon Identity and Access Management (IAM) role and a relay state URL to configure an identity provider (IdP) that is compliant with SAML 2.0. The role grants users permissions to access Amazon Quick Suite. The relay state is the portal that the user is forwarded to, after successful authentication by Amazon.

Topics

- [Prerequisites](#)
- [Step 1: Create a SAML provider in Amazon](#)
- [Step 2: Configure permissions in Amazon for your federated users](#)
- [Step 3: Configure the SAML IdP](#)
- [Step 4: Create assertions for the SAML authentication response](#)
- [Step 5: Configure the relay state of your federation](#)

Prerequisites

Before configuring your SAML 2.0 connection, do the following:

- Configure your IdP to establish a trust relationship with AWS:
 - Inside your organization's network, configure your identity store, such as Windows Active Directory, to work with a SAML-based IdP. SAML-based IdPs include Active Directory Federation Services, Shibboleth, and so on.
 - Using your IdP, generate a metadata document that describes your organization as an identity provider.
 - Set up SAML 2.0 authentication, using the same steps as for the Amazon Web Services Management Console. When this process is complete, you can configure your relay state to match the relay state of Quick Suite. For more information, see [Configure the relay state of your federation](#).
- Create an Amazon Quick Suite account and note the name to use when you configure your IAM policy and IdP. For more information on creating an Amazon Quick Suite account, see [Signing up for an Amazon Quick Suite subscription](#).

After you create the setup to federate to the Amazon Web Services Management Console as outlined in the tutorial, you can edit the relay state provided in the tutorial. You do so with the relay state of Amazon Quick Suite, described in step 5 following.

For more information, see the following resources:

- [Integrating Third-Party SAML Solution Providers with Amazon](#) in the *IAM User Guide*.
- [Troubleshooting SAML 2.0 federation with Amazon](#), also in the *IAM User Guide*.
- [Setting up trust between ADFS and Amazon and using Active Directory credentials to connect to Amazon Athena with ODBC driver](#) – This walkthrough article is helpful, although you don't need to set up Athena in order to use Amazon Quick Suite.

Step 1: Create a SAML provider in Amazon

Your SAML identity provider defines your organization's IdP to Amazon. It does so by using the metadata document that you previously generated using your IdP.

To create a SAML provider in Amazon

1. Sign in to the Amazon Web Services Management Console and open the IAM console at <https://console.amazonaws.cn/iam/>.

2. Create a new SAML provider, which is an entity in IAM that holds information about your organization's identity provider. For more information, see [Creating SAML Identity Providers](#) in the *IAM User Guide*.
3. As part of this process, upload the metadata document produced by the IdP software in your organization noted in the previous section.

Step 2: Configure permissions in Amazon for your federated users

Next, create an IAM role that establishes a trust relationship between IAM and your organization's IdP. This role identifies your IdP as a principal (trusted entity) for the purposes of federation. The role also defines which users authenticated by your organization's IdP are allowed to access Amazon Quick Suite. For more information about creating a role for a SAML IdP, see [Creating a Role for SAML 2.0 Federation](#) in the *IAM User Guide*.

After you have created the role, you can limit the role to have permissions only to Amazon Quick Suite by attaching an inline policy to the role. The following sample policy document provides access to Amazon Quick Suite. This policy allows the user access to Amazon Quick Suite and allows them to create both author accounts and reader accounts.

Note

In the following example, replace `<YOUR_AWS_ACCOUNT_ID>` with your 12-digit Amazon Web Services account ID (with no hyphens '-').

```
{
  "Statement": [
    {
      "Action": [
        "quicksight:CreateUser"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws-cn:quicksight::<YOUR_AWS_ACCOUNT_ID>:user/${aws:userid}"
      ]
    }
  ],
  "Version": "2012-10-17"
```

```
}
```

If you want to provide access to Amazon Quick Suite and also the ability to create Amazon Quick Suite admins, authors (standard users), and readers, you can use the following policy example.

```
{
  "Statement": [
    {
      "Action": [
        "quicksight:CreateAdmin"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws-cn:quicksight::<YOUR_AWS_ACCOUNT_ID>:user/${aws:userid}"
      ]
    }
  ],
  "Version": "2012-10-17"
}
```

You can view account details in the Amazon Web Services Management Console.

After you have set up SAML and the IAM policy or policies, you don't need to invite users manually. The first time that users open Amazon Quick Suite, they are provisioned automatically, using the highest level permissions in the policy. For example, if they have permissions to both `quicksight:CreateUser` and `quicksight:CreateReader`, they are provisioned as authors. If they also have permissions to `quicksight:CreateAdmin`, they are provisioned as admins. Each permission level includes the ability to create the same level user and below. For example, an author can add other authors or readers.

Users who are invited manually are created in the role assigned by the person who invited them. They don't need to have policies that grant them permissions.

Step 3: Configure the SAML IdP

After you create the IAM role, update your SAML IdP about Amazon as a service provider. To do so, install the `saml-metadata.xml` file found at <https://signin.amazonaws.cn/static/saml-metadata.xml>.

To update the IdP metadata, see the instructions provided by your IdP. Some providers give you the option to type the URL, after which the IdP gets and installs the file for you. Others require you to download the file from the URL and then provide it as a local file.

For more information, see your IdP documentation.

Step 4: Create assertions for the SAML authentication response

Next, configure the information that the IdP passes as SAML attributes to Amazon as part of the authentication response. For more information, see [Configuring SAML Assertions for the Authentication Response](#) in the *IAM User Guide*.

Step 5: Configure the relay state of your federation

Finally, configure the relay state of your federation to point to the Amazon Quick Suite relay state URL. After successful authentication by Amazon, the user is directed to Amazon Quick Suite, defined as the relay state in the SAML authentication response.

The relay state URL for Amazon Quick Suite is as follows.

```
https://quicksight.aws.amazon.com
```

Initiating sign-on from Quick Suite

Applies to: Enterprise Edition

Intended audience: System administrators

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick Suite.

In this scenario, your user initiates the sign-on process from an Amazon Quick Suite application portal without being signed on to the identity provider. In this case, the user has a federated

account managed by a third-party IdP. The user might have a user account on Quick Suite. Quick Suite sends an authentication request to the IdP. After the user is authenticated, Quick Suite opens.

Beginning with the user signing into Quick Suite, authentication flows through these steps:

1. The user opens Quick Suite. At this point, the user isn't signed in to the IdP.
2. The user attempts to sign in to Amazon Quick Suite.
3. Amazon Quick Suite redirects the user's input to the federation service and requests authentication.
4. The federation service and the IdP authenticate the user:
 - a. The federation service requests authentication from the organization's identity store.
 - b. The identity store authenticates the user and returns the authentication response to the federation service.
 - c. When authentication is successful, the federation service posts the SAML assertion to the user's browser.
 - d. The user's browser posts the SAML assertion to the Amazon Sign-In SAML endpoint (<https://signin.aws.amazon.com/saml>).
 - e. Amazon Sign-In receives the SAML request, processes the request, authenticates the user, and forwards the authentication token to the Amazon Quick Suite service.
5. Amazon Quick Suite accepts the authentication token from Amazon and presents Amazon Quick Suite to the user.

From the user's perspective, the process happens transparently. The user starts at an Amazon Quick Suite application portal. Amazon Quick Suite negotiates authentication with your organization's federation service and Amazon. Amazon Quick Suite opens, without the user needing to supply any additional credentials.

Setting up service provider–initiated federation with Quick Suite Enterprise edition

Applies to: Enterprise Edition

Intended audience: System administrators

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick Suite.

After you have finished configuring your identity provider with Amazon Identity and Access Management (IAM), you can set up service provider–initiated sign in through Amazon Quick Suite Enterprise Edition. For Quick Suite–initiated IAM federation to work, you need to authorize Quick Suite to send the authentication request to your IdP. A Quick Suite administrator can configure this by adding the following information provided by the IdP:

- The IdP URL – Quick Suite redirects users to this URL for authentication.
- The relay state parameter – This parameter relays the state that the browser session was in when it was redirected for authentication. The IdP redirects the user back to the original state after authentication. The state is provided as a URL.

The following table shows the standard authentication URL and relay state parameter for redirecting the user to the Quick Suite URL that you provide.

Identity provider	Parameter	Authentication URL
Auth0	RelayState	<code>https://<sub_domain>.auth0.com/samlp/<app_id></code>
Google accounts	RelayState	<code>https://accounts.google.com/o/saml2/initssso?idpid=<idp_id>&spid=<sp_id>&forceauthn=false</code>
Microsoft Azure	RelayState	<code>https://myapps.microsoft.com/signin/<app_name>/<app_id>?tenantId=<tenant_id></code>
Okta	RelayState	<code>https://<sub_domain>.okta.com/app/<app_name>/<app_id>/sso/saml</code>
PingFederate	TargetResource	<code>https://<host>/idp/<idp_id>/startSSO.ping?PartnerSpId=<sp_id></code>

Identity provider	Parameter	Authentication URL
PingOne	TargetResource	https://sso.connect.pingidentity.com/sso/sp/initssso?saasid=<app_id>&idpid=<idp_id>

Amazon Quick Suite supports connecting to one IdP per Amazon Web Services account. The configuration page in Amazon Quick Suite provides you with test URLs based on your entries, so you can test the settings before you turn on the feature. To make the process even more seamless, Amazon Quick Suite provides a parameter (`enable-sso=0`) to temporarily turn off Amazon Quick Suite initiated IAM federation, in case you need to disable it temporarily.

To set up Amazon Quick Suite as a service provider that can initiate IAM federation for an existing IdP

1. Make sure that you already have IAM federation set up in your IdP, in IAM, and Amazon Quick Suite. To test this setup, check if you can share a dashboard with another person in your company's domain.
2. Open Amazon Quick Suite, and choose **Manage Amazon Quick Suite** from your profile menu at upper right.

To perform this procedure, you need to be a Amazon Quick Suite administrator. If you aren't, you can't see **Manage Amazon Quick Suite** under your profile menu.

3. Choose **Single sign-on (IAM federation)** from the navigation pane.
4. For **Configuration**, **IdP URL**, enter the URL that your IdP provides to authenticate users.
5. For **IdP URL**, enter the parameter that your IdP provides to relay state, for example `RelayState`. The actual name of the parameter is provided by your IdP.

6. Test signing in:

- To test signing in with your identity provider, use the custom URL provided in **Test starting with your IdP**. You should arrive at the start page for Amazon Quick Suite, for example `https://quicksight.aws.amazon.com/sn/start`.
- To test signing in with Amazon Quick Suite first, use the custom URL provided in **Test the end-to-end experience**. The `enable-sso` parameter is appended to the URL. If `enable-sso=1`, IAM federation attempts to authenticate.

7. Choose **Save** to keep your settings.

To enable service provider–initiated IAM federation IdP

1. Make sure your IAM federation settings are configured and tested. If you're not sure about the configuration, test the connection by using the URLs from the previous procedure.
2. Open Amazon Quick Suite, and choose **Manage Amazon Quick Suite** from your profile menu.
3. Choose **Single sign-on (IAM federation)** from the navigation pane.
4. For **Status**, choose **ON**.
5. Verify that it's working by disconnecting from your IdP and opening Amazon Quick Suite.

To disable service provider initiated IAM federation

1. Open Amazon Quick Suite, and choose **Manage Amazon Quick Suite** from your profile menu.
2. Choose **Single sign-on (IAM federation)** from the navigation pane.
3. For **Status**, choose **OFF**.

Configuring email syncing for federated users in Quick Suite

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick Suite administrators

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick Suite.

In Amazon Quick Suite Enterprise edition, as an administrator you can restrict new users from using personal email addresses when provisioning through their identity provider (IdP) directly to Quick Suite. Quick Suite then uses the preconfigured email addresses passed through the IdP

when provisioning new users to your account. For example, you can make it so that only corporate-assigned email addresses are used when users are provisioned to your Amazon Quick Suite account through your IdP.

 Note

Make sure that your users are federating directly to Amazon Quick Suite through their IdP. Federating to the Amazon Web Services Management Console through their IdP and then clicking into Amazon Quick Suite results in an error and they won't be able to access Amazon Quick Suite.

When you configure email syncing for federated users in Amazon Quick Suite, users who log in to your Amazon Quick Suite account for the first time have preassigned email addresses. These are used to register their accounts. With this approach, users can manually bypass by entering an email address. Also, users can't use an email address that might differ from the email address prescribed by you, the administrator.

Amazon Quick Suite supports provisioning through an IdP that supports SAML or OpenID Connect (OIDC) authentication. To configure email addresses for new users when provisioning through an IdP, you update the trust relationship for the IAM role that they use with `AssumeRoleWithSAML` or `AssumeRoleWithWebIdentity`. Then you add a SAML attribute or OIDC token in their IdP. Last, you turn on email syncing for federated users in Amazon Quick Suite.

The following procedures describe these steps in detail.

Step 1: Update the trust relationship for the IAM role with `AssumeRoleWithSAML` or `AssumeRoleWithWebIdentity`

You can configure email addresses for your users to use when provisioning through your IdP to Amazon Quick Suite. To do this, add the `sts:TagSession` action to the trust relationship for the IAM role that you use with `AssumeRoleWithSAML` or `AssumeRoleWithWebIdentity`. By doing this, you can pass in `principal` tags when users assume the role.

The following example illustrates an updated IAM role where the IdP is Okta. To use this example, update the Federated Amazon Resource Name (ARN) with the ARN for your service provider. You can replace items in red with your Amazon and IdP service-specific information.

```
{
```

```

"Version": "2012-10-17",
"Statement": [
{
  "Effect": "Allow",
  "Principal": {
    "Federated": "arn:aws-cn:iam::account-id:saml-provider/Okta"
  },
  "Action": "sts:AssumeRoleWithSAML",
  "Condition": {
    "StringEquals": {
      "SAML:aud": "https://signin.aws.amazon.com/saml"
    }
  }
},
{
  "Effect": "Allow",
  "Principal": {
    "Federated": "arn:aws-cn:iam::account-id:saml-provider/Okta"
  },
  "Action": "sts:TagSession",
  "Condition": {
    "StringLike": {
      "aws:RequestTag/Email": "*"
    }
  }
}
]
}

```

Step 2: Add a SAML attribute or OIDC token for the IAM principal tag in your IdP

After you update the trust relationship for the IAM role as described in the preceding section, add a SAML attribute or OIDC token for the IAM Principal tag in your IdP.

The following examples illustrate a SAML attribute and an OIDC token. To use these examples, replace the email address with a variable in your IdP that points to a user's email address. You can replace items highlighted in red with your information.

- **SAML attribute:** The following example illustrates a SAML attribute.

```

<Attribute Name="https://aws.amazon.com/SAML/Attributes/
PrincipalTag:Email"><AttributeValue>john.doe@example.com</AttributeValue></Attribute>

```

Note

If you're using Okta as your IdP, make sure to turn on a feature flag in your Okta user account to use SAML. For more information, see [Okta and Amazon Partner to Simplify Access Via Session Tags](#) on the Okta blog.

- **OIDC token:** The following example illustrates an OIDC token example.

```
"https://aws.amazon.com/tags": {"principal_tags": {"Email": ["john.doe@example.com"]}}
```

Step 3: Turn on email syncing for federated users in Amazon Quick Suite

As described preceding, update the trust relationship for the IAM role and add a SAML attribute or OIDC token for the IAM Principal tag in your IdP. Then turn on email syncing for federated users in Amazon Quick Suite as described in the following procedure.

To turn on email syncing for federated users

1. From any page in Amazon Quick Suite, choose your username at top right, and then choose **Manage Amazon Quick Suite**.
2. Choose **Single sign-on (IAM federation)** in the menu at left.
3. On the **Service Provider Initiated IAM federation** page, for **Email Syncing for Federated Users**, choose **ON**.

When email syncing for federated users is on, Amazon Quick Suite uses the email addresses that you configured in steps 1 and 2 when provisioning new users to your account. Users can't enter their own email addresses.

When email syncing for federated users is off, Amazon Quick Suite asks users to input their email address manually when provisioning new users to your account. They can use any email addresses that they want.

Tutorial: Amazon Quick Suite and IAM identity federation

Applies to: Enterprise Edition and Standard Edition

Intended audience: Amazon Quick Suite Administrators and Amazon Quick Suite developers

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick Suite.

In the following tutorial, you can find a walkthrough for setting up the IdP Okta as a federation service for Amazon Quick Suite. Although this tutorial shows the integration of Amazon Identity and Access Management (IAM) and Okta, you can also replicate this solution using your choice of SAML 2.0 IdPs.

In the following procedure, you create an app in the Okta IdP using their "Amazon Account Federation" shortcut. Okta describes this integration app as follows:

"By federating Okta to Amazon Web Services (Amazon) Identity and Access Management (IAM) accounts, end users get single sign-on access to all their assigned Amazon roles with their Okta credentials. In each Amazon Web Services account, administrators set up federation and configure Amazon roles to trust Okta. When users sign in to Amazon, they get Okta single sign-in experience to see their assigned Amazon roles. They can then select a desired role, which defined their permissions for the duration of their authenticated session. Customers with large numbers of Amazon Accounts, check out the Amazon Single Sign-On app as an alternative." (<https://www.okta.com/aws/>)

To create an Okta app using Okta's "Amazon Account Federation" application shortcut

1. Sign in to your Okta dashboard. If you don't have one, create a free Okta Developer Edition account by using [this Amazon Quick Suite-branded URL](#). When you have activated your email, sign in to Okta.
2. On the Okta website, choose **<> Developer Console** at upper left, and then choose **Classic UI**.
3. Choose **Add Applications**, and choose **Add app**.
4. Enter **aws** for **Search**, and choose **Amazon Account Federation** from the search results.
5. Choose **Add** to create an instance of this application.
6. For **Application label**, enter **Amazon Account Federation - Amazon Quick Suite**.
7. Choose **Next**.

8. For **SAML 2.0, Default Relay State**, enter **https://quicksight.aws.amazon.com**.
9. Open the context (right-click) menu for **Identity Provider metadata**, and choose to save the file. Name the file `metadata.xml`. You need this file in the next procedure.

The contents of the file look similar to the following.

```
<md:EntityDescriptor xmlns:md="urn:oasis:names:tc:SAML:2.0:metadata"
  entityID="http://www.okta.com/exkffz2hATwiVft645d5">
  <md:IDPSSODescriptor WantAuthnRequestsSigned="false"
    protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protocol">
    <md:KeyDescriptor use="signing">
      <ds:KeyInfo xmlns:ds="http://www.w3.org/2000/09/xmldsig#">
        <ds:X509Data>
          <ds:X509Certificate>

MIIDpjCCAo6gAwIBAgIGAXVjA82hMA0GCSqGSIb3DQEBCwUAMIGTMQswCQYDVQQGEwJVUzETMBEG
          .
          .      (certificate content omitted)
          .
          QE/6cRdPQ6v/eaFpUL6Asd6q3sBeq+giRG4=
          </ds:X509Certificate>
        </ds:X509Data>
      </ds:KeyInfo>
    </md:KeyDescriptor>
    <md:NameIDFormat>urn:oasis:names:tc:SAML:1.1:nameid-format:emailAddress</
md:NameIDFormat>
    <md:NameIDFormat>urn:oasis:names:tc:SAML:2.0:nameid-format:unspecified</
md:NameIDFormat>
    <md:SingleSignOnService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-
POST" Location="https://dev-1054988.okta.com/app/amazon_aws/exkffz2hATwiVft645d5/
sso/saml"/>
    <md:SingleSignOnService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-
Redirect" Location="https://dev-1054988.okta.com/app/amazon_aws/
exkffz2hATwiVft645d5/sso/saml"/>
  </md:IDPSSODescriptor>
</md:EntityDescriptor>
```

10. After you have the XML file saved, scroll to the bottom of the Okta page, and choose **Done**.
11. Keep this browser window open, if possible. You need it later in the tutorial.

Next, you create an identity provider in your Amazon Web Services account.

To create a SAML provider in Amazon Identity and Access Management (IAM)

1. Sign in to the Amazon Web Services Management Console and open the IAM console at <https://console.amazonaws.cn/iam/>.
2. In the navigation pane, choose **Identity providers, Create Provider**.
3. Enter the following settings:
 - **Provider Type** – Choose **SAML** from the list.
 - **Provider Name** – Enter **Okta**.
 - **Metadata Document** – Upload the XML file `manifest.xml` from the previous procedure.
4. Choose **Next Step, Create**.
5. Locate the IdP that you created and choose it to view the settings. Note the **Provider ARN**. You need this to finish the tutorial.
6. Verify that the identity provider is created with your settings. In IAM, choose **Identity providers, Okta** (the IdP you added), **Download metadata**. The file should be the one that you recently uploaded.

Next, you create an IAM role to enable the SAML 2.0 federation to act as a trusted entity in your Amazon Web Services account. For this step, you need to choose how you want to provision users in Amazon Amazon Quick Suite. You can do one of the following:

- Grant permission to the IAM role so that first-time visitors become Amazon Quick Suite users automatically.

To create an IAM role for a SAML 2.0 federation as a trusted entity

1. Sign in to the Amazon Web Services Management Console and open the IAM console at <https://console.amazonaws.cn/iam/>.
2. In the navigation pane, choose **Roles, Create Role**.
3. For **Select type of trusted entity**, choose the card labeled **SAML 2.0 federation**.
4. For **SAML provider**, select the IdP that you created in the previous procedure, for example **Okta**.
5. Enable the option **Allow programmatic and Amazon Management Console access**.
6. Choose **Next: Permissions**.

7. Paste the following policy into the editor.

In the policy editor, update the JSON with your provider's Amazon Resource Name (ARN).

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRoleWithSAML",
      "Resource": "arn:aws-cn:iam::111111111111:saml-provider/Okta",
      "Condition": {
        "StringEquals": {
          "saml:aud": "https://signin.aws.amazon.com/saml"
        }
      }
    }
  ]
}
```

8. Choose **Review policy**.

9. For **Name**, enter **QuicksightOktaFederatedPolicy**, and then choose **Create policy**.

10. Choose **Create policy, JSON** a second time.

11. Paste the following policy into the editor.

In the policy editor, update the JSON with your Amazon Web Services account ID. It should be the same account ID that you used in the previous policy in the provider ARN.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": [
        "quicksight:CreateReader"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws-cn:quicksight::111111111111:user/${aws:userid}"
      ]
    }
  ]
}
```

```
}
```

You can omit the Amazon Web Services Region name in the ARN, as shown following.

```
arn:aws-cn:quicksight::111111111111:user/${aws:userid}
```

12. Choose **Review policy**.
13. For **Name**, enter **QuicksightCreateReader**, and then choose **Create policy**.
14. Refresh the list of policies by choosing the refresh icon at right.
15. For **Search**, enter **QuicksightOktaFederatedPolicy**. Choose the policy to enable it ().

If you don't want to use automatic provisioning, you can skip the following step.

To add a Amazon Quick Suite user, use [register-user](#). To add a Amazon Quick Suite group, use [create-group](#). To add users to the Amazon Quick Suite group, use [create-group-membership](#).

16. (Optional) For **Search**, enter **QuicksightCreateReader**. Choose the policy to enable it ().

Do this step if you want to provision Amazon Quick Suite users automatically, rather than using the Amazon Quick Suite API.

The `QuicksightCreateReader` policy activates automatic provisioning by allowing use of the `quicksight:CreateReader` action. Doing this grants dashboard subscriber (reader-level) access to first-time users. A Amazon Quick Suite administrator can later upgrade them from the Amazon Quick Suite profile menu, **Manage Amazon Quick Suite**, **Manage users**.

17. To continue attaching the IAM policy or policies, choose **Next: Tags**.
18. Choose **Next: Review**.
19. For **Role name**, enter **QuicksightOktaFederatedRole**, and choose **Create role**.
20. Verify that you completed this successfully by taking these steps:
 - a. Return to the main page of the IAM console at <https://console.amazonaws.cn/iam/>. You can use your browser's **Back** button.
 - b. Choose **Roles**.
 - c. For **Search**, enter Okta. Choose **QuicksightOktaFederatedRole** from the search results.

- d. On the **Summary** page for the policy, examine the **Permissions** tab. Verify that the role has the policy or policies that you attached to it. It should have `QuicksightOktaFederatedPolicy`. If you chose to add the ability to create users, it should also have `QuicksightCreateReader`.
- e. Use the

icon to open each policy. Verify that the text matches what is shown in this procedure. Double-check that you added your own Amazon Web Services account number in place of the example account number 111111111111.
- f. On the **Trust relationships** tab, verify that the **Trusted entities** field contains the ARN for the identity provider. You can double-check the ARN in the IAM console by opening **Identity providers, Okta**.
- g.

To create an access key for Okta

1. Sign in to the Amazon Web Services Management Console and open the IAM console at <https://console.amazonaws.cn/iam/>.
2. Add a policy that allows Okta to display a list of IAM roles to the user. To do this, choose **Policy, Create policy**.
3. Choose **JSON**, then enter the following policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:ListRoles",
        "iam:ListAccountAliases"
      ],
      "Resource": "*"
    }
  ]
}
```

4. Choose **Review Policy**.
5. For **Name**, enter **OktaListRolesPolicy**. Then choose **Create policy**.

6. Add a user so you can provide Okta with an access key.

In the navigation pane, choose **Users, Add User**.

7. Use the following settings:
 - For **User name**, enter `OktaSSOUser`.
 - For **Access type**, enable **Programmatic access**.
8. Choose **Next: Permissions**.
9. Choose **Attach existing policies directly**.
10. For **Search**, enter `OktaListRolesPolicy`, and choose `OktaListRolesPolicy` from the search results.
11. Choose **Next: Tags**, and then choose **Next: Review**.
12. Choose **Create user**. Now you can get the access key.
13. Download the key file by choosing **Download .csv**. The file contains the same access key ID and secret access key that displays on this screen. However, because Amazon doesn't display this information a second time, make sure to download the file.
14. Verify that you completed this step correctly by doing the following:
 - a. Open the IAM console, and choose **Users**. Search for `OktaSSOUser`, and open it by choosing the username from the search results.
 - b. On the **Permissions** tab, verify that the `OktaListRolesPolicy` is attached.
 - c. Use the

icon to open the policy. Verify that the text matches what is shown in this procedure.
 - d. On the **Security credentials** tab, you can check the access key, although you already downloaded it. You can return to this tab to create an access key when you need a new one.

In the following procedure, you return to Okta to provide the access key. The access key works with your new security settings to allow Amazon and the Okta IdP to work together.

To finish configuring the Okta application with Amazon settings

1. Return to your Okta dashboard. If requested to do so, sign in. If the developer console is no longer open, choose **Admin** to reopen it.

2. If you have to reopen Okta, you can return to this section by following these steps:
 - a. Sign in to Okta. Choose **Applications**.
 - b. Choose **Amazon Account Federation - Amazon Quick Suite**—the application that you created at the beginning of this tutorial.
 - c. Choose the **Sign On** tab, between **General** and **Mobile**.
3. Scroll to **Advanced Sign-On Settings**.
4. For **Identity Provider ARN (Required only for SAML IAM federation)**, enter the provider ARN from the previous procedure, for example:

```
arn:aws-cn:iam::111122223333:saml-provider/Okta
```

5. Choose **Done** or **Save**. The name of the button varies depending if you are creating or editing the application.
6. Choose the **Provisioning** tab, and at the lower part of the tab, choose **Configure API Integration**.
7. Turn on **Enable API integration** to display the settings.
8. For **Access Key** and **Secret Key**, provide the access key and secret key that you downloaded previously to a file named **OktaSSOUser_credentials.csv**.
9. Choose **Test API Credentials**. Look above the **Enable API integration** setting for a message confirming that **Amazon Account Federation was verified successfully**.
10. Choose **Save**.
11. Make sure that **To App** is highlighted at left, and choose **Edit** at right.
12. For **Create Users**, turn on the option **Enable**.
13. Choose **Save**.
14. On the **Assignments** tab, near **Provisioning** and **Import**, choose **Assign**.
15. Do one or more of the following to enable federated access:
 - To work with individual users, choose **Assign to People**.
 - To work with IAM groups, choose **Assign to Groups**. You can choose specific IAM groups or **Everyone (All users in your organization)**.
16. For each IAM user or group, do the following:
 - a. Choose **Assign, Role**.

- b. Select **QuicksightOktaFederatedRole** from the list of IAM roles.
 - c. For **SAML User Roles**, enable **QuicksightOktaFederatedRole**.
17. Choose **Save and Go Back**, and then choose **Done**.
18. Verify that you completed this step correctly by choosing the **People** or **Groups** filter at left, and checking the users or groups that you entered. If you can't complete this process because the role that you created doesn't appear in the list, return to the previous procedures to verify the settings.

To sign in to Amazon Quick Suite using Okta (IdP to service provider sign-in)

1. If you are using an Okta administrator account, switch to user mode.
2. Sign in to your Okta Applications dashboard with a user that has been granted federated access. You should see a new application with your label, for example **Amazon Account Federation - Amazon Quick Suite**.
3. Choose the application icon to launch **Amazon Account Federation - Amazon Quick Suite**.

You can now manage identities using Okta and use federated access with Quick Suite.

The following procedure is an optional part of this tutorial. If you follow its steps, you authorize Amazon Quick Suite to forward authorizations requests to the IdP on behalf of your users. Using this method, users can sign in to Amazon Quick Suite with no need to sign in using the IdP page first.

(Optional) To set up Amazon Quick Suite to send authentication requests to Okta

1. Open Amazon Quick Suite, and choose **Manage Amazon Quick Suite** from your profile menu.
2. Choose **Single sign-on (IAM federation)** from the navigation pane.
3. For **Configuration**, **IdP URL**, enter the URL that your IdP provides to authenticate users, for example `https://dev-1-----0.okta.com/home/amazon_aws/0oabababababaGQei5d5/282`. You can find this in your Okta app page, on the **General** tab, in **Embed Link**.
4. For **IdP URL**, enter `RelayState`.
5. Do one of the following:
 - To test signing in with your identity provider first, use the custom URL provided in **Test starting with your IdP**. You should arrive at the start page for Amazon Quick Suite, for example `https://quicksight.aws.amazon.com/sn/start`.

- To test signing in with Amazon Quick Suite first, use the custom URL provided in **Test the end-to-end experience**. The `enable-sso` parameter is appended to the URL. If `enable-sso=1`, IAM federation attempts to authenticate. If `enable-sso=0`, Amazon Quick Suite doesn't send the authentication request, and you sign in to Amazon Quick Suite as before.
6. For **Status**, choose **ON**.
 7. Choose **Save** to keep your settings.

You can create a deep link to a Amazon Quick Suite dashboard to allow users to use IAM federation to connect directly to specific dashboards. To do this, you append the relay state flag and dashboard URL to the Okta single sign-on URL, as described following.

To create a deep link to a Amazon Quick Suite dashboard for single sign-on

1. Locate the Okta application's single sign-on (IAM federation) URL in the `metadata.xml` file that you downloaded beginning of the tutorial. You can find the URL near the bottom of the file, in the element named `md:SingleSignOnService`. The attribute is named `Location` and the value ends with `/sso/saml`, as shown in the following example.

```
<md:SingleSignOnService Binding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-Redirect" Location="https://dev-0000001.okta.com/app/amazon_aws/abcdef2hATwiVft645d5/sso/saml"/>
```

2. Take the value of the IAM federation URL and append `?RelayState=` followed by the URL of your Amazon Quick Suite dashboard. The `RelayState` parameter relays the state (the URL) that the user was in when they were redirected to the authentication URL.
3. To the new IAM federation with the relay state added, append the URL of your Amazon Quick Suite dashboard. The resulting URL should resemble the following.

```
https://dev-1-----0.okta.com/app/amazon_aws/abcdef2hATwiVft645d5/sso/saml?RelayState=https://us-west-2.quicksight.aws.amazon.com/sn/analyses/12a12a2a-121a-212a-121a-abcd12abc1ab
```

4. If the link you create doesn't open, check that you are using the most recent IAM federation URL from the `metadata.xml`. Also check that the username you use to sign in isn't assigned in more than one IAM federation Okta app.

Using Active Directory with Amazon Quick Suite Enterprise edition

Applies to: Enterprise Edition

Intended audience: System administrators

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick Suite.

Amazon Quick Suite Enterprise edition supports both [Amazon Directory Service for Microsoft Active Directory](#) and [Active Directory Connector](#).

To create a new directory to be your identity manager for Quick Suite, use Amazon Directory Service for Microsoft Active Directory, also known as Amazon Managed Microsoft AD. This is an Active Directory host in the Amazon Cloud that offers most of the same functionality of Active Directory. Currently, you can connect to Active Directory in any Amazon Region supported by Amazon Quick Suite, except for Asia Pacific (Singapore). When you create a directory, you use it with a virtual private cloud (VPC). For more information, see [VPC](#).

If you have an existing directory that you want to use for Quick Suite, you can use Active Directory Connector. This service redirects directory requests to your Active Directory—in another Amazon Web Services Region or on-premises—without caching any information in the cloud.

For a walkthrough about creating and managing a directory with Amazon Managed Microsoft AD, see [Use an Amazon Managed Microsoft AD with Quick Suite?](#) in the Amazon Knowledge Center.

When you use Amazon Directory Service to launch a directory, Amazon creates an organizational unit (OU) with the same name as your domain. Amazon also creates an administrative account with delegated administrative rights for the OU. You can create accounts, groups, and policies within the OU by using Active Directory users and groups. For more information, see [Best Practices for Amazon Managed Microsoft AD](#) in the *Directory Service Administration Guide*.

After you establish your directory, you use it with Quick Suite by creating groups for users. Amazon Quick Suite has six specific user roles that can be assigned, including Pro versions that provide access to advanced capabilities:

- **Quick Suite admins** – Admins can change account settings, manage accounts. Admins can also purchase additional Amazon Quick Suite user subscriptions or [SPICE](#) capacity, or cancel the subscription to Amazon Quick Suite for your Amazon Web Services account. Admin Pro users have additional capabilities including creating content using natural language, building knowledge bases, configuring actions, and accessing advanced automation workflows.
- **Quick Suite authors** – Amazon Quick Suite authors can create data sources, datasets, analyses, and dashboards. They can share analyses and dashboards with other Amazon Quick Suite users. Author Pro users can additionally create content using natural language, build knowledge bases, configure actions, and access advanced automation capabilities.
- **Quick Suite readers** – Readers can view and interact with dashboards that were created by someone else. Reader Pro users have access to advanced features including AI chat agents, collaborative spaces, flows, and extensions.

You can add or refine access by applying IAM policies. For example, you can use IAM policies to allow users to subscribe themselves.

When you subscribe to Amazon Quick Suite Enterprise edition and choose Active Directory as your identity provider, you can associate your AD groups with Amazon Quick Suite. You can also add or change your AD groups later on.

Topics

- [Directory integration with Quick Suite Enterprise edition](#)

Directory integration with Quick Suite Enterprise edition

Applies to: Enterprise Edition

Intended audience: System administrators

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick Suite.

Quick Suite Enterprise supports the following options:

- Amazon Directory Service
- Amazon Directory Service with AD Connector
- On-premises Active Directory with IAM federation or AD Connector
- IAM federation using Amazon IAM Identity Center or another third-party federation service

If you want to use IAM federation with an on-premises Active Directory, you implement Amazon Directory Service as a separate Active Directory with a trust relationship to the on-premises Active Directory.

If you want to avoid using a trust relationship, you can deploy a standalone domain for authentication within Amazon. Then you can create users and groups in Active Directory. You'd then map them to users and groups in Quick Suite. In this example, users authenticate using their Active Directory login credentials. To make access to Quick Suite transparent to your users, use IAM federation in this scenario.

Using multi-factor authentication (MFA) with Amazon Quick Suite

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Important

Amazon Quick Suite recommends that you integrate new Quick Suite subscriptions with IAM Identity Center for identity management. This IAM identity federation user guide is provided as a reference for existing account configurations. For more information on

integrating your Quick Suite account with IAM Identity Center, see [Configure your Quick Suite account with IAM Identity Center](#).

Note

IAM identity federation doesn't support syncing identity provider groups with Amazon Quick Suite.

There are several ways that you can use multi-factor authentication (MFA) with Quick Suite. You can use it with Amazon Identity and Access Management (IAM). You can use it with AD Connector or your [Amazon Directory Service](#) for Microsoft Active Directory, also known as Amazon Microsoft Active Directory or Amazon Managed Microsoft Active Directory. And if you use an external identity provider (IdP), Amazon doesn't need to have any information about MFA because that is part of the authentication handled by the IdP.

For more information, see the following:

- [Using multi-factor authentication \(MFA\) in Amazon](#) in the IAM User Guide
- [Enable Multi-Factor Authentication for Amazon Managed Microsoft AD](#) in the Amazon Directory Service Administration Guide
- [Enable Multi-Factor Authentication for AD Connector](#) in the Amazon Directory Service Administration Guide

If you're a developer, see the following:

- [How do I use an MFA token to authenticate access to my Amazon resources through the Amazon CLI](#) in the [Amazon Knowledge Center](#)
- [Configuring MFA-protected API access](#) in the IAM User Guide

Allowlisting Amazon Quick Suite domains

If your end users are signing in to Amazon Quick Suite using Amazon root (not recommended), Amazon Identity and Access Management (IAM), corporate Active Directory, or native Quick Suite credentials, make sure to allow-list the following domains within your organization's network.

User type	Domain or domains to allow-list
Users who sign in directly through Amazon Quick Suite and Active Directory users	signin.aws and awsapps.com
Amazon root user	signin.aws.amazon.com and amazon.com
IAM users	signin.aws.amazon.com

Important

We strongly recommend that you don't use the Amazon root user for your everyday tasks, even the administrative ones. Instead, adhere to the best practice of using the root user only to create your first IAM user. Then securely lock away the root user credentials and use them to perform only a few account and service management tasks. For more information, see [Amazon account root user](#) in the *IAM User Guide*.

Configuring VPC connections in Amazon Quick Sight

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick Suite administrators

Note

If you're a Amazon Quick Suite system administrator configuring a VPC connection to Amazon Quick Sight, this section is for you. Amazon Quick Suite knowledge bases currently don't support VPC integrations.

Quick Suite Enterprise edition is fully integrated with the Amazon VPC service. A VPC based on this service closely resembles a traditional network that you operate in your own data center. It

enables you to secure and isolate traffic between resources. You define and control the network elements to suit your requirements, while still getting the benefit of cloud networking and the scalable infrastructure of Amazon.

By creating a VPC connection in Amazon Quick Suite, you're adding elastic network interfaces in your VPC. These network interfaces allow Amazon Quick Suite to exchange network traffic with a network instance within your VPC. You can provide all of the standard security controls for this network traffic, as you do with other traffic in your VPC. Route tables, network access control lists (ACLs), subnets, and security groups settings all apply to network traffic to and from Amazon Quick Suite in the same way that they apply to traffic between other instances in your VPC.

When you register a VPC connection with Amazon Quick Suite, you can securely connect to data that's available only in your VPC, for example:

- Data you can reach by IP address
- Data that isn't available on the public internet
- Private databases
- On-premises data

This works if you set up connectivity between the VPC and your on-premises network. For example, you might set up connectivity with Amazon Direct Connect, a virtual private network (VPN), or a proxy.

After you connect to the data, you can use it to create data analyses and publish secure data dashboards.

To further increase security, consider logging data access operations with Amazon CloudTrail, as described in [Logging Amazon Quick Suite information with CloudTrail](#). You can even create a dashboard to help you analyze your CloudTrail logs. By combining Amazon Quick Suite logs with logs from your other Amazon services, you can get a fuller view of how your data is being used.

You don't need to be an networking expert to connect and use a VPC with Amazon Quick Suite, because Amazon Quick Suite provides a user interface for adding your network information. However, the person who gathers the information that you need for setup should have some understanding of networking concepts and using VPCs. This person also needs read-only access to the services. If network changes are required, we recommend that you don't make changes to your networking configuration without expert assistance.

To use a command line interface to access your VPC, you can use the Amazon Command Line Interface (Amazon CLI). For more information on using the Amazon CLI, see the [Amazon CLI User Guide](#).

Topics

- [VPC terminology](#)
- [Supported VPC data sources](#)
- [Setting up a VPC to use with Amazon Quick Suite](#)
- [Finding information to connect to a VPC](#)

VPC terminology

The following terminology can be useful when you work with a VPC and Amazon Quick Suite.

A *VPC* is a virtual private cloud, which works like a private network to isolate the resources within it. The solution described in these topics uses an Amazon service called Amazon VPC.

A *route table* contains a set of rules, called *routes*, that are used to determine where network traffic is directed. You can view the route table in the Amazon VPC console at <https://console.amazonaws.cn/vpc/>. The VPC details display the route table that the VPC is using. You can also see **Route tables** listed in the Amazon VPC console.

A *subnet* is a defined set of network IP addresses that are used to increase the security and efficiency of network communications. You can think of them like postal codes, used for routing packages from one location to another. The **Subnets** list in the Amazon VPC console displays subnet IDs and also their associated VPC IDs, route tables, and network ACLs. You need to provide at least two subnets in different availability zones to create a VPC connection.

A *network interface* represents a virtual network card. The network interface automatically created by Amazon Quick Suite is called a *Amazon Quick Suite network interface*. Each network interface in a VPC connection is configured based on the subnet it's attached to. You can view your Amazon Quick Suite network interfaces in the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>. The network interface displays its network interface ID, subnet ID, VPC ID, security group, and the Availability Zone that it exists in. You can click on the security group name to see its group ID and its inbound and outbound rules. The term *network interface* in the following sections always means elastic network interface.

A *security group* is a set of rules that controls the network access to the resources it is associated with. Access is permitted only to and from the components defined in the security group's inbound and outbound rules. If no rules are defined, the security group prevents all access. You can view security groups from several different consoles, depending on which resource that a particular security group applies to. You can see all the security groups and their settings in one place in the VPC console. For the Amazon Quick Suite VPC connection, create a new security group.

Inbound and outbound rules define the following:

- The type of traffic to allow, for example **"All TCP"** or **"RDS"**.
- The protocol to allow (TCP, UDP, or ICMP).
- The traffic source to allow for inbound rules, or the traffic destination to allow for outbound rules. When you work with a VPC and Amazon Quick Suite, you specify the security group ID to use.
- An optional description. We recommend that you add the word **Amazon Quick Suite** to the description for Amazon Quick Suite VPC rules.

An *internet gateway* is a VPC component that allows communication between instances in your VPC and the internet. You don't need an internet gateway to use Amazon Quick Suite VPC connections.

A *VPC endpoint* enables you to privately connect your VPC to supported Amazon services without using public IP addresses. You don't need to set up a VPC endpoint to use Amazon Quick Suite VPC connections.

Supported VPC data sources

Amazon Quick Suite VPC connections work only with specific Amazon Quick Sight data sources. Use this section to know which data sources are compatible and what requirements they must meet.

The following Amazon Quick Sight data sources can connect to Amazon Quick Suite through a VPC connection:

- Amazon OpenSearch Service
- Amazon Redshift
- Amazon Relational Database Service
- Amazon Aurora

- Databricks
- Exasol
- MariaDB
- Microsoft SQL Server
- MySQL
- Oracle
- PostgreSQL
- Presto
- Snowflake
- Starburst Enterprise
- Teradata
- Trino

For a VPC data source to be accessed from Amazon Quick Sight, the following statements must be true of your configuration:

1. The Domain Name System (DNS) name of the VPC data source can be resolved from outside of your VPC.
2. The connection returns the private IP address of your instance. Databases hosted by Amazon Redshift, Amazon RDS, and Aurora automatically meet this requirement.
3. There is a clearly defined network path from the data source to Amazon Quick Sight.
4. You registered the VPC with Amazon Quick Suite by creating or using a VPC connection with the Amazon Quick Suite console.

Setting up a VPC to use with Amazon Quick Suite

Applies to: Enterprise Edition

Intended audience: System administrators

To set up a VPC to use with Amazon Quick Suite Enterprise edition, you need access to Amazon VPC and Amazon EC2. You also need access to each Amazon database service that you plan to add to Quick Suite. You can use the console, or you can use the Amazon Command Line Interface (Amazon CLI). For more information about the CLI, see the [Amazon Command Line Interface User Guide](#). To work with the CLI, go to <http://www.amazonaws.cn/cli/>.

Before you begin to set up your VPC connection in Amazon Quick Suite, make sure that you understand the components of a VPC deployment. As part of that, familiarize yourself with the VPC's subnets and security groups in relation to the destinations (databases) that you want to reach from Amazon Quick Suite. To set up a successful VPC connection, make sure that the following components work together to allow network traffic to pass between Amazon Quick Suite and your data source:

- The Amazon VPC service
- The subnets that your data source is using
- The Amazon Quick Suite elastic network interfaces and the subnets they use
- The route table
- Inbound and outbound rules for these security groups:
 - Security group for your VPC. We recommend you create a new security group to isolate the rules on the VPC security group from the rules on the Amazon Quick Suite network interface's security group).
 - Security group attached to the Amazon Quick Suite network interface.
 - Security group attached to the database server (for each database server that you want to use).
- (Optional) Amazon Route 53 Resolver inbound endpoints for private DNS resolution.

In the following topics, you can find the network components that are involved. You can also find descriptions of their roles in the network configuration of your VPC and your Amazon Quick Suite VPC connection. The network interface for Amazon Quick Suite that is automatically created during setup is called the *Amazon Quick Suite network interface (QNI)*.

If your VPC is already completely configured, skip to the next section, [Finding information to connect to a VPC](#).

Topics

- [VPC](#)

- [Subnets](#)
- [Security groups: inbound and outbound rules](#)
- [Sample rules](#)
- [Route table](#)
- [Amazon Quick Suite elastic network interface](#)
- [Inbound endpoints for Amazon Route 53 Resolver](#)

VPC

A *virtual private cloud (VPC)* is a virtual network dedicated to your Amazon account. The Amazon VPC service that provides it is a networking layer for your Amazon resources. Using Amazon VPC, you can define a virtual network in your own logically isolated area within the Amazon Cloud. A VPC closely resembles a traditional network that you might operate in your own data center, with the benefits of using the Amazon scalable infrastructure. Amazon VPC for Amazon EC2 virtual computing environments, known as *instances*, can be used for a variety of Amazon resources.

VPCs offer options that allow for flexibility in a secure environment, for example:

- To configure your VPC, you can set its IP address range, create subnets, configure route tables, network gateways, network interfaces, and security settings.
- To make the Amazon Cloud an extension of your data center, you can connect your VPC to your own corporate data center.
- You can connect your instances in the VPC to the internet, or keep your instances isolated on a private network.
- To protect the resources in each subnet, you can use multiple layers of security, including security groups and network access control lists (ACLs).

For more information, see the [Amazon VPC User Guide](#).

If you have a default VPC and don't specify a subnet when you launch an instance, the instance is launched into your default VPC. You can launch instances into your default VPC without needing to know anything about Amazon VPC.

If you don't already have a VPC or want to use a new one, you can create one by following the instructions in [Getting started with Amazon VPC](#) in the *Amazon VPC User Guide*. This section

offers guidance on how to set up your VPC. The guidance includes options for public and private subnets and for Amazon Site-to-Site VPN access for your corporate network (known as *on-premises access*). You can also use VPC peering or Amazon Direct Connect to reach an on-premises database instance.

Using the Amazon CLI

You can start to set up a VPC in Amazon EC2 by using the [aws ec2 create-vpc](#) command. To learn more about VPC settings for the Amazon CLI, see [Examples for VPC](#) in the *Amazon VPC User Guide*.

Using the Amazon EC2 console

To view your VPC or create a new one in Amazon EC2, sign in to the Amazon Web Services Management Console and open the Amazon VPC console at <https://console.amazonaws.cn/vpc/>. To create a new VPC, choose **Launch VPC Wizard** and follow the instructions. Note your new VPC ID for future reference. To view VPCs, choose **Your VPCs** on the left side.

Amazon VPC resources in VPC guides and Amazon Support articles

For general information, see [Working with VPCs and subnets](#).

For step-by-step instructions for setting up a VPC, see the following topics (choose the ones that relate to your scenario):

- [Create an IPv4 VPC and subnets using the Amazon CLI](#)
- [Sharing public subnets and private subnets](#)
- [Working with site-to-site VPN](#)
- [Amazon Site-to-Site VPN Network Administrator Guide](#) (choose your network device for specific instructions)
- [Generic Customer Gateway Device Without Border Gateway Protocol](#) (recommended for customer gateways)

If you want to migrate data source instances into the same VPC, see the following Amazon Support articles:

- [How do I change the VPC for an Amazon RDS DB instance?](#)
- [How do I move my EC2 instance to another subnet, Availability Zone, or VPC?](#)

- [How do I move my Amazon Redshift cluster from one VPC to another VPC?](#)

For troubleshooting information, see [How do I troubleshoot issues with VPC route tables?](#), an article with video created by Amazon Support.

Subnets

A *subnet* is a range of IP addresses in your VPC. You need to provide at least two subnets to create a VPC connection. Each subnet must belong a different availability zone. You can attach Amazon resources, such as Amazon EC2 instances and Amazon RDS DB instances, to subnets. You can create subnets to group instances together according to your security and operational needs.

For Amazon Quick Suite to connect to your database, the network needs to route traffic to the data sources that you want to reach from one of the subnets used by the Amazon Quick Suite network interface. Amazon Quick Suite determines which subnet to route traffic through on the backend. If the availability zone that the subnet is attached to experiences an outage, Amazon Quick Suite reroutes the traffic to one of the other subnets that are configured in the VPC connection. If the data sources are on different subnets, make sure that there is a route from the Amazon Quick Suite network interface to your database instance. By default, each subnet in a VPC is associated with one main route table and can reach the other subnets. For more information, see [VPC and Subnets](#) and [Network ACLs](#) in the *Amazon VPC User Guide*.

If you use Amazon RDS, DB instances are associated with a subnet group that you can view either in the Amazon RDS console (<https://console.amazonaws.cn/rds/>) or in the VPC console. For troubleshooting connectivity to Amazon RDS, see the Amazon Support article [How can I troubleshoot connectivity to an Amazon RDS instance that uses a public or private subnet of a VPC?](#)

Security groups: inbound and outbound rules

A *security group* acts as a virtual firewall for your instance to control inbound and outbound traffic. For each security group, you add rules that control the inbound traffic to instances, and a separate set of rules that control the outbound traffic.

For your VPC connection, create a new security group with the description QuickSight-VPC. This security group must allow all inbound TCP traffic from the security groups of the data destinations that you want to reach. The following example creates a new security group in the VPC and returns the ID of the new security group.

```
aws ec2 create-security-group \  
--name QuickSight-VPC \  
--group-name quicksight-vpc \  
--description "QuickSight-VPC" \  
--vpc-id vpc-0daeb67adda59e0cd
```

Important

Network configuration is sufficiently complex that we strongly recommend that you create a new security group for use with Amazon Quick Suite. It also makes it easier for Amazon Support to help you if you need to contact them. Creating a new group isn't absolutely required. However, the following topics are based on the assumption that you follow this recommendation.

To enable Quick Suite to successfully connect to an instance in your VPC, configure your security group rules to allow traffic between the Amazon Quick Suite network interface and the instance that contains your data. To do this, configure the security group attached to your database's instance inbound rules to allow the following traffic:

- From the port that Amazon Quick Suite is connecting to
- From one of the following options:
 - The security group ID that's associated with Amazon Quick Suite network interface (recommended)
 - or
 - The private IP address of the Amazon Quick Suite network interface

For more information, see [Security groups for your VPC](#) and [VPCs and subnets](#) in the *Amazon VPC User Guide*.

Use the topics listed below to learn more about inbound and outbound rules.

Topics

- [Inbound rules](#)
- [Outbound rules](#)

Inbound rules

Important

The following section applies to your VPC connection if the connection was created before April 27, 2023.

When you create a security group, it has no inbound rules. No inbound traffic originating from another host to your instance is allowed until you add inbound rules to the security group.

The security group attached to the Amazon Quick Suite network interface behaves differently than most security groups, because it isn't stateful. Other security groups are usually *stateful*. This means that, after they establish an outbound connection to a resource's security group, they automatically allow return traffic. In contrast, the Amazon Quick Suite network interface security group doesn't automatically allow return traffic. Because of this, adding an egress rule to the Amazon Quick Suite network interface security group doesn't work. To make it work for the Amazon Quick Suite network interface security group, make sure to add an inbound rule that explicitly authorizes the return traffic from the database host.

The inbound rule in your security group must allow traffic on all ports. It needs to do this because the destination port number of any inbound return packets is set to a randomly allocated port number.

To restrict Amazon Quick Suite to connect only to certain instances, you can specify the security group ID (recommended) or private IP address of the instances that you want to allow. In either case, your security group inbound rule still needs to allow traffic on all ports (0–65535).

To allow Amazon Quick Suite to connect to any instance in the VPC, you can configure the Amazon Quick Suite network interface security group. In this case, give it an inbound rule to allow traffic on 0.0.0.0/0 on all ports (0–65535). The security group used by the Amazon Quick Suite network interface should be different than the security groups used for your databases. We recommend that you use separate security groups for VPC connection.

Important

If you are using a long-standing Amazon RDS DB instance, check your configuration to see if you're using a DB security group. DB security groups are used with DB instances that are not in a VPC and are on the EC2-Classical platform.

If this is your configuration, and you aren't moving your DB instance into the VPC for use with Amazon Quick Suite, make sure to update your DB security group's inbound rules. Update them to allow inbound traffic from the VPC security group that you're using for Amazon Quick Suite. For more information, see [Controlling Access with Security Groups](#) in the *Amazon RDS User Guide*.

Outbound rules

Important

The following section applies to your VPC connection if the connection was created before April 27, 2023.

By default, a security group includes an outbound rule that allows all outbound traffic. We recommend that you remove this default rule and add outbound rules that allow specific outbound traffic only.

Warning

Do not configure the security group on the Amazon Quick Suite network interface with an outbound rule to allow traffic on all ports. For information on key considerations and recommendations for managing network egress traffic from VPCs, see [Security best practices for your VPC](#) in the *Amazon VPC User Guide*.

The security group attached to Amazon Quick Suite network interface should have outbound rules that allow traffic to each of the database instances in your VPC that you want Amazon Quick Suite to connect to. To restrict Amazon Quick Suite to connect only to certain instances, specify the security group ID (recommended) or the private IP address of the instances to allow. You set this up, along with the appropriate port numbers for your instances (the port that the instances are listening on), in the outbound rule.

The VPC security group must also allow outbound traffic to the security groups of the data destinations, specifically on the port or ports that the database is listening on.

Sample rules

Following, you can find some example configurations of inbound and outbound rules for Amazon RDS and Amazon Redshift.

VPC connection rules: Amazon Quick Sight: Amazon RDS for MySQL

The following tables show rule settings for connecting Amazon Quick Sight to Amazon RDS for MySQL.

Amazon Quick Sight Network interface security group: inbound rule

Type	All TCP
Protocol	TCP
Port Range	0 - 65535
Source	<i>sg-RDS11111111</i>
Description	Amazon Quick Sight - RDS MySQL

Amazon Quick Sight Network interface security group: outbound rule

Type	MYSQL/Aurora
Protocol	TCP
Port Range	3306
Source	sg-RDS11111111
Description	Amazon Quick Sight to RDS MySQL

RDS MySQL: inbound rule

Type	MYSQL/Aurora

Protocol	TCP
Port Range	3306
Source	sg-ENI33333333
Description	Amazon Quick Sight to RDS MySQL

VPC connection rules: Amazon Redshift in Amazon Quick Sight

The following tables show rule settings for connecting Amazon Quick Sight to Amazon Redshift.

Amazon Quick Sight network interface security group: inbound rule

Type	All TCP
Protocol	TCP
Port Range	0 - 65535
Source	sg-RedSh222222
Description	Amazon Quick Sight– Amazon Redshift

Amazon Quick Sight network interface security group: outbound rule

Type	Amazon Redshift
Protocol	TCP
Port Range	5439
Source	sg-RedSh222222
Description	Amazon Quick Sight– Amazon Redshift

Amazon Redshift: inbound rule

Type	Amazon Redshift
Protocol	TCP
Port Range	5439
Source	sg-ENI33333333
Description	Amazon Quick Sight– Amazon Redshift

Route table

To use VPC peering or Amazon Direct Connect to reach an on-premises database instance, update the route table that's associated with the VPC you're using with Amazon Quick Suite. For more information on route tables, see [Route tables](#) in the *Amazon VPC User Guide*.

To learn more about VPC peering and view sample scenarios and configurations, see [What is VPC peering?](#) in the *Amazon VPC Peering Guide*. For an example configuration, see [Example: Services using Amazon PrivateLink and VPC peering](#) in the *Amazon VPC User Guide*.

Using the Amazon CLI

The following example creates a route table.

```
aws ec2 create-route-table --vpc-id vpc-0daeb67adda59e0cd
```

Then you can use the `create-route` command to create a route. For more information and examples, see [create-route](#) in the *Amazon CLI Command Reference*.

For the following examples to work, make sure that you have a subnet in the VPC associated with the route table. The first example describes the route table with the specified VPC ID. The second one describes the route table with the specified route table ID.

```
aws ec2 describe-route-tables \
--filters "Name=vpc-id,Values=vpc-0daeb67adda59e0cd"
```

```
aws ec2 describe-route-tables \  
--route-table-ids rtb-45ac473a
```

The following example describes the specified associations between a specific VPC and your local gateway route tables.

```
aws ec2 describe-local-gateway-route-table-vpc-associations  
--filters "Name=vpc-id,Values=vpc-0daeb67adda59e0cd"
```

Amazon Quick Suite elastic network interface

The *Amazon Quick Suite elastic network interface* is a logical networking component in a VPC that represents a virtual network card. Quick Suite creates at least two of these network interfaces to use with a VPC connection based off of the subnets that are attached to it. Then you add the VPC connection to each Amazon Quick Sight data source you create. The Quick Suite network interface alone doesn't give Quick Suite direct access to your databases. The VPC connection works only for the Amazon Quick Sight data sources that are configured to use it.

When you use the Amazon Quick Sight data source to query a database or other instance within your VPC, all the network traffic from Amazon Quick Suite originates from this Amazon Quick Suite network interface. Because the Amazon Quick Suite network interface exists inside your VPC, traffic originating from it can reach destinations within your VPC by using their private IP addresses. Each Amazon Quick Suite network interface gets its own private IP address that comes from the subnet you configure. The private IP address is unique for each Amazon account, unlike the public IP range.

Inbound endpoints for Amazon Route 53 Resolver

Amazon Route 53 Resolver provides DNS query capabilities to your VPC. Route 53 Resolver resolves all local DNS queries and recursively looks up any DNS queries that aren't local on public DNS servers.

Amazon Quick Suite can't directly use Route 53 Resolver to query private DNS servers. However, you can set up Route 53 Resolver inbound endpoints to make these queries indirectly. For more information about inbound endpoints, see [Forwarding inbound DNS queries to your VPCs](#) in the *Route 53 Resolver Developer Guide*. To use inbound endpoints in Amazon Quick Suite, provide the IP addresses of the endpoints for **DNS resolver endpoints** when you create a VPC connection.

Finding information to connect to a VPC

Applies to: Enterprise Edition

Intended audience: System administrators

To gather the information to have ready when you create a VPC connection in Amazon Quick Suite Enterprise edition, take the steps listed following.

Steps

- [Identify the data sources to use](#)
- [Identify the Amazon Web Services Region to use](#)
- [Identify the VPC ID to use](#)
- [Identify the subnet IDs to use](#)
- [Identify the security group to use](#)

Identify the data sources to use

Start by identifying all the data sources that you want to connect to using Quick Suite. For each of these, note the database's private IP, security group, and subnets. Amazon Quick Suite connects to your data using the private IP. However, you don't need to enter this or the security group or subnet information for the VPC connection. This information helps you identify the other components you need for the Amazon Quick Suite VPC connection.

Note

For the connection to your data source to work, make sure that there's a traceable route from your data source to the VPC ID. For more details, see [Identify the data sources to use](#).

Identify the Amazon Web Services Region to use

For the connection to work, the data, the subnets, and the security group must be in the same VPC. Make sure also that you use Quick Suite in the same Amazon Web Services Region with the VPC.

You can't use Amazon Quick Suite in one Amazon Web Services Region and expect to connect to a VPC in a different Amazon Web Services Region.

If your team is already using Amazon Quick Suite, you can see your current Amazon Web Services Region displayed at the upper right of the Amazon Quick Suite home screen. You can change the Amazon Web Services Region you're using in Amazon Quick Suite by changing the Region at the upper right of the Amazon Quick Suite home screen. All the people who plan to use the data in the VPC must be using the same Amazon Web Services Region in Amazon Quick Suite.

Note

The Amazon Web Services Region that displays in the Amazon Quick Suite console doesn't have to match your Amazon CLI configuration. Take care not to mistake your current Amazon Quick Suite console settings with the settings that apply in any Amazon CLI commands that you run or the settings in other consoles. Changing the current Amazon Web Services Region in any console doesn't change the Region anywhere except for that page.

For example, let's say you have three tabs open in one browser window. You can have the Amazon Quick Suite console open in one Amazon Web Services Region, the Amazon VPC console open in a second Region, the Amazon RDS console open in a third Region, and the Amazon CLI running in a fourth Region.

Identify the VPC ID to use

The VPC ID is assigned when the VPC is created.

Using the Amazon CLI

The following `describe-vpcs` example retrieves details for all of your VPCs.

```
aws ec2 describe-vpcs
```

The following `describe-vpcs` example retrieves details for the specified VPC.

```
aws ec2 describe-vpcs \  
--vpc-ids vpc-06e4ab6c6cEXAMPLE
```

Using the Amazon VPC console

In the VPC console (<https://console.amazonaws.cn/vpc/>), choose **Your VPCs** at left. Choose the VPC-ID that you want to use. The correct one has Availability Zones in your Amazon Web Services Region and also meets the requirements described in [Finding information to connect to a VPC](#). Also note the ID of **Main Route Table**, because you need this to identify related subnets.

Tip

In the Amazon VPC console, you can filter by VPC. This option is located at the top left of the console. If you filter by your VPC ID, all the other menus display only the network elements that are in your selected VPC.

Identify the subnet IDs to use

To locate the subnet IDs for the subnets used by the VPC, open the VPC console. Locate the VPC you are using, and at least two subnets in different availability zones. Amazon Quick Suite creates its Amazon Quick Suite elastic network interface (Amazon Quick Suite network interface) for the subnets that you choose. The Amazon Quick Suite network interfaces get created after you save your VPC connection settings, described in the following section.

Your database instances can reside in different subnets. However, make sure you can trace the route from this subnet to any data destinations that you want to reach.

Using the Amazon CLI

The following example describes all existing subnets.

```
aws ec2 describe-subnets
```

The following describe-subnets example uses a filter to retrieve details for the subnets of the specified VPC.

```
aws ec2 describe-subnets \
--filters "Name=vpc-id,Values=vpc-06e4ab6c6cEXAMPLE"
```

Using the Amazon VPC console

In the VPC console (<https://console.amazonaws.cn/vpc/>), choose **Subnets** at left, and find the correct **Subnet ID**. Any subnet is correct if your database subnet has a route to the subnet that you choose at this point. In most cases, if you haven't configured the VPC network yourself, all subnets are connected.

Identify the security group to use

The security group contains rules that control the inbound and outbound network traffic on your data source instances. The security group you are using should have the description "QuickSight-VPC" to make it easier to identify.

When you locate the correct security group, copy its **Group ID** value.

Using the Amazon CLI

The following example displays the security groups in a specific Amazon Web Services Region. It displays only the group ID, name, and description. It filters the result to display only groups for a specific VPC ID that also have a description of "QuickSight-VPC".

```
aws ec2 describe-security-groups \
--region us-west-2 \
--query 'SecurityGroups[*].[GroupId, GroupName, Description]' \
--filters "Name=vpc-id,Values=vpc-06e4ab6c6cEXAMPLE"
"Name=description,Values=QuickSight-VPC"
```

The following example displays information about the security group with the ID sg-903004f8. Note that you can't reference a security group for EC2-VPC by name.

```
aws ec2 describe-security-groups
--group-ids sg-903004f8
--region us-west-2
```

The following example queries the results to describe VPC the inbound and outbound rules of a security group with a specific ID (sg-903004f8), in a specific Amazon Web Services Region (us-west-2).

```
aws ec2 describe-security-groups \
--region us-west-2 \
--group-ids sg-903004f8 \
```

```
--query 'SecurityGroups[*].[GroupId, GroupName, Description, IpPermissions, IpPermissionsEgress]'
```

The following example uses filters to describe VPC security groups that have a specific rule that allows SQL Server traffic (port 1433). The example also has a rule that allows traffic from all addresses (0.0.0.0/0). The output is filtered to display only the group IDs, names, and descriptions of the security groups. Security groups must match all filters to be returned in the results. However, a single rule doesn't have to match all filters. (EC2-VPC only)

```
aws ec2 describe-security-groups \
--filters Name=ip-permission.from-port,Values=1433 \
Name=ip-permission.to-port,Values=1433 \
Name=ip-permission.cidr,Values='0.0.0.0/0' \
--query 'SecurityGroups[*].[GroupId, GroupName, Description]'
```

Using the Amazon VPC console

In the VPC console (<https://console.amazonaws.cn/vpc/>), choose **Security groups** at left, and find the correct group ID. The correct one has your VPC ID on it. It should also have a tag or description that includes the word "QuickSight".

Configuring Amazon Quick Sight access to Amazon data sources

Use this section to help you configure access to resources in other Amazon services.

We recommend that you use SSL to secure Amazon Quick Sight connections to your data sources. To use SSL, you must have a certificate signed by a recognized certificate authority (CA). Amazon Quick Suite doesn't accept certificates that are self-signed or issued from a nonpublic CA. For more information, see [Amazon Quick Suite SSL and CA certificates](#).

Topics

- [Required permissions](#)
- [Network and database configuration requirements](#)
- [Allowing autodiscovery of Amazon resources](#)
- [Authorizing connections from Amazon Quick Sight to Amazon data stores](#)
- [Exploring your Amazon data in Amazon Quick Suite](#)

- [Amazon service action connectors](#)

Required permissions

When you connect to a Amazon Quick Sight data source that requires a user name, the user name must have SELECT permissions on some system tables. These permissions allow Amazon Quick Sight to do things such as discover table schemas and estimate table size.

The following table identifies the tables that the account must have SELECT permissions for, depending on the type of database you are connecting to. These requirements apply for all database instances you connect to, regardless of their environment. In other words, they apply whether your database instances are on-premises, in Amazon RDS, in Amazon EC2, or elsewhere.

Instance type	Tables
Amazon Aurora	INFORMATION_SCHEMA.STATISTICS INFORMATION_SCHEMA.TABLES
Amazon Redshift	pg_stats pg_class pg_namespace
MariaDB	INFORMATION_SCHEMA.STATISTICS INFORMATION_SCHEMA.TABLES
Microsoft SQL Server	DBCC SHOW_STATISTICS sp_statistics
MySQL	INFORMATION_SCHEMA.STATISTICS INFORMATION_SCHEMA.TABLES
Oracle	DBA_TAB_COLS ALL_TABLES

Instance type	Tables
	dba_segments all_segments user_segments
PostgreSQL	pg_stats pg_class pg_namespace
ServiceNow	sys_dictionary (column metadata) sys_db_object (table metadata) sys_glide_object (field type metadata)

Note

If you are using MySQL or PostgreSQL, verify that you are connecting from an allowed host or IP address. For more detail, see [Database configuration requirements for self-administered instances](#).

Network and database configuration requirements

To serve as data sources, databases need to be configured so that Amazon Quick Suite can access them. Use the following sections to make sure that your database is configured appropriately.

Important

Because a database instance on Amazon EC2 is administered by you rather than Amazon, it must meet both the [Network configuration requirements](#) as well as the [Database configuration requirements for self-administered instances](#).

Topics

- [Network configuration requirements](#)
- [Database configuration requirements for self-administered instances](#)

Network configuration requirements

Intended audience: System administrators

For you to use your database server from Amazon Quick Suite, your server must be accessible from the internet. It must also allow inbound traffic from Amazon Quick Suite servers.

If the database is on Amazon and in the same Amazon Web Services Region as your Amazon Quick Suite account, you can auto-discover the instance to make connecting to it easier. To do this, you must grant Amazon Quick Suite permissions to access it. For more information, see [Accessing data sources](#).

Topics

- [Network configuration for an Amazon instance in a default VPC](#)
- [Network configuration for an Amazon instance in a nondefault VPC](#)
- [Network configuration for an Amazon instance in a private VPC](#)
- [Network configuration for an Amazon instance that is not in a VPC](#)
- [Network configuration for a database instance other than Amazon](#)

Network configuration for an Amazon instance in a default VPC

In some cases, your database might be on an Amazon cluster or instance that you created in a default VPC. Thus, it's publicly accessible (that is, you didn't choose to make it private). In such cases, your database is already appropriately configured to be accessible from the internet. However, you still need to enable access from Amazon Quick Suite servers to your Amazon cluster or instance. For further details on how to do this, choose the appropriate topic following:

- [Authorizing connections from Amazon Quick Suite to Amazon RDS database instances](#)
- [Authorizing connections from Amazon Quick Suite to Amazon Redshift clusters](#)
- [Authorizing connections from Amazon Quick Suite to Amazon EC2 instances](#)

Network configuration for an Amazon instance in a nondefault VPC

If you are configuring an Amazon instance in a nondefault VPC, make sure that the instance is publicly accessible and that the VPC has the following:

- An internet gateway.
- A public subnet.
- A route in the route table between the internet gateway and the Amazon instance.
- Network access control lists (ACLs) in your VPC that allow traffic between the cluster or instance and Amazon Quick Suite servers. These ACLs must do the following:
 - Allow inbound traffic from the appropriate Amazon Quick Suite IP address range and all ports to the IP address and port that the database is listening on.
 - Allow outbound traffic from the database's IP address and port to the appropriate Amazon Quick Suite IP address range and all ports.

For more information about Amazon Quick Suite IP address ranges, see [IP address ranges for Amazon Quick Suite](#) following.

For more information about configuring VPC ACLs, see [Network ACLs](#).

- Security group rules that allow traffic between the cluster or instance and Amazon Quick Suite servers. For further details on how to create appropriate security group rules, see [Authorizing connections to Amazon data sources](#).

For more information about configuring a VPC in the Amazon VPC service, see [Networking in Your VPC](#).

Network configuration for an Amazon instance in a private VPC

If your database is on an Amazon cluster or instance that you created in a private VPC, you can use it with Amazon Quick Suite. For more information, see [Connecting to a Amazon VPC with Amazon Quick Suite](#).

For more information on Amazon VPC, see [Amazon VPC](#) and [Amazon VPC Documentation](#).

Network configuration for an Amazon instance that is not in a VPC

If you are configuring an Amazon instance that is not in a VPC, make sure that the instance is publicly accessible. Also, make sure that there is a security group rule that allows traffic between

the cluster or instance and Amazon Quick Suite servers. For further details on how to do this, choose the appropriate topic following:

- [Authorizing connections from Amazon Quick Suite to Amazon RDS database instances](#)
- [Authorizing connections from Amazon Quick Suite to Amazon Redshift clusters](#)
- [Authorizing connections from Amazon Quick Suite to Amazon EC2 instances](#)

Network configuration for a database instance other than Amazon

To use SSL to secure your connections to your database (*recommended*), make sure that you have a certificate signed by a recognized certificate authority (CA). Amazon Quick Suite doesn't accept certificates that are self-signed or issued from a nonpublic CA. For more information, see [Amazon Quick Suite SSL and CA certificates](#).

If your database is on a server other than Amazon, you must change that server's firewall configuration to accept traffic from the appropriate Amazon Quick Suite IP address range. For more information about Amazon Quick Suite IP address ranges, see [IP address ranges for Amazon Quick Suite](#). For any other steps that you need to take to enable internet connectivity, see your operating system documentation.

Amazon Quick Suite SSL and CA certificates

Following is a list of accepted public certificate authorities. If you are using a database instance other than Amazon, your certificate must be on this list, or it won't work.

• AAA Certificate Services	• QuoVadis Root CA 2
• AddTrust Class 1 CA Root	• QuoVadis Root CA 3
• AddTrust External CA Root	• QuoVadis Root Certification Authority
• AddTrust Qualified CA Root	• SecureTrust CA
• AffirmTrust Commercial	• Sonera Class1 CA
• AffirmTrust Networking	• Sonera Class2 CA
• AffirmTrust Premium	• Starfield Root Certificate Authority - G2
• AffirmTrust Premium ECC	• Starfield Services Root Certificate Authority - G2
• America Online Root Certification Authority 1	

- America Online Root Certification Authority 2
- Baltimore CyberTrust Code Signing Root
- Baltimore CyberTrust Root
- Buypass Class 2 Root CA
- Buypass Class 3 Root CA
- Certum CA
- Certum Trusted Network CA
- Chambers of Commerce Root
- Chambers of Commerce Root - 2008
- Class 2 Primary CA
- Class 3P Primary CA
- Deutsche Telekom Root CA 2
- DigiCert Assured ID Root CA
- DigiCert Global Root CA
- DigiCert High Assurance EV Root CA
- Entrust.net Certification Authority (2048)
- Entrust Root Certification Authority
- Entrust Root Certification Authority - G2
- Equifax Secure eBusiness CA-1
- Equifax Secure Global eBusiness CA-1
- GeoTrust Global CA
- GeoTrust Primary Certification Authority
- SwissSign Gold CA - G2
- SwissSign Platinum CA - G2
- SwissSign Silver CA - G2
- TC TrustCenter Class 2 CA II
- TC TrustCenter Class 4 CA II
- TC TrustCenter Universal CA I
- Thawte Personal Freemail CA
- Thawte Premium Server CA
- thawte Primary Root CA
- thawte Primary Root CA - G2
- thawte Primary Root CA - G3
- Thawte Server CA
- Thawte Timestamping CA
- T-TeleSec GlobalRoot Class 2
- T-TeleSec GlobalRoot Class 3
- UTN - DATACorp SGC
- UTN-USERFirst-Client Authentication and Email
- UTN-USERFirst-Hardware
- UTN-USERFirst-Object
- Valicert
- VeriSign Class 1 Public Primary Certification Authority - G3
- VeriSign Class 2 Public Primary Certification Authority - G3
- VeriSign Class 3 Public Primary Certification Authority - G3
- VeriSign Class 3 Public Primary Certification Authority - G4
- VeriSign Class 3 Public Primary Certification Authority - G5

- GeoTrust Primary Certification Authority - G2
- GeoTrust Primary Certification Authority - G3
- GeoTrust Universal CA
- Global Chambersign Root - 2008
- GlobalSign
- GlobalSign Root CA
- Go Daddy Root Certificate Authority - G2
- GTE CyberTrust Global Root
- KEYNECTIS ROOT CA
- VeriSign Universal Root Certification Authority
- XRamp Global Certification Authority

IP address ranges for Amazon Quick Suite

For more information on the IP address ranges for Amazon Quick Suite in supported Regions, see [Amazon Regions, websites, IP address ranges, and endpoints](#).

Database configuration requirements for self-administered instances

Intended audience: System administrators and Amazon Quick Suite administrators

For a database to be accessible to Amazon Quick Suite, it must meet the following criteria:

- It must be accessible from the internet. To enable internet connectivity, see your database management system documentation.
- It must be configured to accept connections and authenticate access using the user credentials that you provide as part of creating the data set.
- If you are connecting to MySQL or PostgreSQL, the database engine must be accessible from your host or IP range. This optional security limitation is specified in MySQL or PostgreSQL connection settings. If this limitation is in place, any attempt to connect from a nonspecified host or IP address is rejected, even if you have the correct username and password.

- In MySQL, the server accepts the connection only if the user and host are verified in the user table. For more information, see [Access Control, Stage 1: Connection Verification](#) in the MySQL documentation.
- In PostgreSQL, you control client authentication by using the `pg_hba.conf` file in the database cluster's data directory. However, this file might be named and located differently on your system. For more information, see [Client Authentication](#) in the PostgreSQL documentation.

Allowing autodiscovery of Amazon resources

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Each Amazon service that you access from Amazon Quick Suite needs to allow traffic from Quick Suite. Instead of opening each service console separately to add permissions, a Quick Suite administrator can do this in the administration screen. Before you begin, make sure that you have addressed the following prerequisites.

If you choose to enable autodiscovery of Amazon resources for your Quick Suite account, Quick Suite creates an Amazon Identity and Access Management (IAM) role in your Amazon Web Services account. This IAM role grants your account permission to identify and retrieve data from your Amazon data sources.

Because Amazon limits the number of IAM roles that you can create, make sure that you have at least one free role. You need this role for Amazon Quick Suite to use if you want Amazon Quick Suite to autodiscover your Amazon resources.

You can have Amazon Quick Suite autodiscover Amazon RDS DB instances or Amazon Redshift clusters that are associated with your Amazon Web Services account. These resources must be located in the same Amazon Web Services Region as your Amazon Quick Suite account.

If you choose to enable autodiscovery, choose one of the following options to make the Amazon resource accessible:

- For Amazon RDS DB instances that you created in a default VPC and didn't make private, or that aren't in a VPC (EC2-Classic instances), see [Authorizing connections from Amazon Quick Suite](#)

[to Amazon RDS instances](#). In this topic, you can find information on creating a security group to allow connections from Amazon Quick Suite servers.

- For Amazon Redshift clusters that you created in a default VPC and didn't choose to make private, or that aren't in a VPC (that is, EC2-Classic instances), see [Authorizing connections from Amazon Quick Suite to Amazon Redshift clusters](#). In this topic, you can find information on creating a security group to allow connections from Amazon Quick Suite servers.
- For an Amazon RDS DB instance or Amazon Redshift cluster that is in a nondefault VPC, see [Authorizing connections from Amazon Quick Suite to Amazon RDS instances](#) or [Authorizing connections from Amazon Quick Suite to Amazon Redshift clusters](#). In these topics, you can find information on first creating a security group to allow connections from Amazon Quick Suite servers. In addition, you can find information on then verifying that the VPC meets the requirements described in [Network configuration for an Amazon instance in a nondefault VPC](#).
- If you don't use a private VPC, set up the Amazon RDS instance to allow connections from the Amazon Quick Suite Region's public IP address.

Enabling autodiscovery is the easiest way to make this data available in Amazon Quick Suite. You can still manually create data connections whether or not you enable autodiscovery.

Authorizing connections from Amazon Quick Sight to Amazon data stores

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

For Amazon Quick Sight to access your Amazon resources, you must create security groups for them that authorize connections from the IP address ranges used by Amazon Quick Sight servers. You must have Amazon credentials that permit you to access these Amazon resources to modify their security groups.

Use the procedures in the following sections to enable Amazon Quick Sight connections.

Topics

- [Authorizing connections from Amazon Quick Sight to Amazon RDS DB instances](#)

- [Authorizing connections from Amazon Quick Sight to Amazon Redshift clusters](#)
- [Authorizing connections from Amazon Quick Suite to Amazon EC2 instances](#)
- [Authorizing connections through Amazon Lake Formation](#)
- [Authorizing connections to Amazon OpenSearch Service](#)
- [Authorizing connections to Amazon Athena](#)
- [Data access integrations](#)

Authorizing connections from Amazon Quick Sight to Amazon RDS DB instances

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

For Amazon Quick Sight to connect to an Amazon RDS DB instance, you must create a new security group for that DB instance. This security group contains an inbound rule authorizing access from the appropriate IP address range for the Quick Suite servers in that Amazon Web Services Region. To learn more about authorizing Quick Suite connections, see [Manually enabling access to an Amazon RDS instance in a VPC](#) or [Manually enabling access to an Amazon RDS instance that is not in a VPC](#).

To learn more about authorizing Amazon Quick Sight connections manually, see [Manually enabling access to an Amazon RDS instance in a VPC](#) or [Manually enabling access to an Amazon RDS instance that is not in a Amazon VPC](#).

To create and assign a security group for an Amazon RDS DB instance, you must have Amazon credentials that permit access to that DB instance.

Enabling connection from Amazon Quick Suite servers to your instance is just one of several prerequisites for creating a data set based on an Amazon database data source. For more information about what is required, see [Creating a dataset from a database](#).

Topics

- [Manually enabling Amazon Quick Sight access to an Amazon RDS instance in a VPC](#)

- [Manually enabling access from Amazon Quick Sight to an Amazon RDS instance that is not in a VPC](#)

Manually enabling Amazon Quick Sight access to an Amazon RDS instance in a VPC

Use the following procedure to enable Amazon Quick Sight access to an Amazon RDS DB instance in a VPC. If your Amazon RDS DB instance is in subnet that is private (in relation to Amazon Quick Suite) or that has Internet Gateways attached, see [Connecting to a VPC with Amazon Quick Suite](#).

To enable Amazon Quick Sight access to an Amazon RDS DB instance in a VPC

1. Sign in to the Amazon Web Services Management Console and open the Amazon RDS console at <https://console.amazonaws.cn/rds/>.
2. Choose **Databases**, locate the DB instance, and view its details. To do this, you click directly on its name (a hyperlink in the **DB identifier** column).
3. Locate **Port** and note the **Port** value. This can be a number or a range.
4. Locate **VPC** and note the **VPC** value.
5. Choose the **VPC** value to open the VPC console. In the Amazon VPC Management Console, choose **Security Groups** in the navigation pane.
6. Choose **Create Security Group**.
7. On the **Create Security Group** page, enter the security group information as follows:
 - For **Name tag** and **Group name**, enter **Amazon-QuickSight-access**.
 - For **Description**, enter **Amazon-QuickSight-access**.
 - For **VPC**, choose the VPC for your instance. This VPC is the one with the **VPC ID** that you noted previously.
8. Choose **Create**. On the confirmation page, note the **Security Group ID**. Choose **Close** to exit this screen.
9. Choose your new security group from the list, and then choose **Inbound Rules** from the tab list below.
10. Choose **Edit rules** to create a new rule.
11. On the **Edit inbound rules** page, choose **Add rule** to create a new rule.

Use the following values:

- For **Type**, choose **Custom TCP Rule**.

- For **Protocol**, choose **TCP**.
- For **Port Range**, enter the port number or range of the Amazon RDS cluster. This port number (or range) is the one that you noted previously.
- For **Source**, choose **Custom** from the list. Next to the word "Custom", enter the CIDR address block for the Amazon Web Services Region where you plan to use Amazon Quick Suite.

For example, for Europe (Ireland) you would enter Europe (Ireland)'s CIDR address block: 52.210.255.224/27. For more information on the IP address ranges for Amazon Quick Suite in supported Amazon Web Services Regions, see [Amazon Regions, websites, IP address ranges, and endpoints](#).

 Note

If you have activated Amazon Quick Suite in multiple Amazon Web Services Regions, you can create inbound rules for each Amazon Quick Suite endpoint CIDR. Doing this allows Amazon Quick Suite to have access to the Amazon RDS DB instance from any Amazon Region defined in the inbound rules.

Anyone who uses Amazon Quick Suite in multiple Amazon Web Services Regions is treated as a single user. In other words, even if you are using Amazon Quick Suite in every Amazon Web Services Region, both your Amazon Quick Suite subscription (sometimes called an 'account') and your users are global.

12. For **Description**, enter a useful description, for example "*Europe (Ireland) QuickSight*".
13. Choose **Save rules** to save your new inbound rule. Then choose **Close**.
14. Go back to the detailed view of the DB instance. Return the Amazon RDS console (<https://console.aws.amazon.com/rds/>) and choose **Databases**.
15. Choose the DB identifier for the relevant RDS instance. Choose **Modify**. The same screen displays whether you choose Modify from the databases screen or the DB instance screen: **Modify DB Instance**.
16. Locate the **Network & Security** section (the third section from the top).

The currently assigned security group or groups are already chosen for **Security Group**. Don't remove any of the existing ones unless you are sure.

Instead, choose your new security group to add it to the other groups that are selected. If you followed the name suggested previously, this group might be named something similar to **Amazon-QuickSight-access**.

17. Scroll to the bottom of the screen. Choose **Continue**, and then choose **Modify DB Instance**.
18. Choose **Apply during the next scheduled maintenance** (the screen indicates when this will occur).

Don't choose **Apply immediately**. Doing this also applies any additional changes that are in the pending modifications queue. Some of these changes might require downtime. If you bring the server down outside the maintenance window, this can cause a problem for users of this DB instance. Consult your system administrators before applying immediate changes.

19. Choose **Modify DB Instance** to confirm your changes. Then, wait for the next maintenance window to pass.

Manually enabling access from Amazon Quick Sight to an Amazon RDS instance that is not in a VPC

Use the following procedure to access an Amazon RDS DB instance that is not in a VPC. You can associate a security group with a DB instance by using **Modify** on the RDS console, the `ModifyDBInstance` Amazon RDS API, or the `modify-db-instance` Amazon CLI command.

Note

This section included for backwards compatibility purposes.

To use the console to access an Amazon RDS DB instance that is not in a VPC

1. Sign in to the Amazon Web Services Management Console and open the Amazon RDS console at <https://console.amazonaws.cn/rds/>.
2. Choose **Databases**, select the DB instance, and choose **Modify**.
3. Choose **Security Groups** in the navigation pane.
4. Choose **Create DB Security Group**.
5. Enter **Amazon-QuickSight-access** for the **Name** and **Description** values, and then choose **Create**.
6. The new security group is selected by default.

Select the details icon next to the security group, as shown following.

7. For **Connection Type**, choose **CIDR/IP**.

8. For **CIDR/IP to Authorize**, enter the appropriate CIDR address block. For more information on the IP address ranges for Amazon Quick Suite in supported Amazon Web Services Regions, see [Amazon Regions, websites, IP address ranges, and endpoints](#).
9. Choose **Authorize**.
10. Return to the **Instances** page of the Amazon RDS Management Console, choose the instance that you want to enable access to, choose **Instance Actions**, and then choose **Modify**.
11. In the **Network & Security** section, the currently assigned security group or groups already is chosen for **Security Group**. Press CTRL and choose **Amazon-QuickSight-access** in addition to the other selected groups.
12. Choose **Continue**, and then choose **Modify DB Instance**.

Authorizing connections from Amazon Quick Sight to Amazon Redshift clusters

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

You can provide access to Amazon Redshift data using three authentication methods: trusted identity propagation, run-as IAM role, or Amazon Redshift database credentials.

With trusted identity propagation, a user's identity is passed to Amazon Redshift with single sign-on that is managed by IAM Identity Center. A user that accesses a dashboard in Amazon Quick Sight has their identity propagated to Amazon Redshift. In Amazon Redshift, fine grained data permissions are applied on the data before the data is presented in a Amazon Quick Suite asset to the user. Amazon Quick Suite authors can also connect to Amazon Redshift data sources without a password input or IAM role. If Amazon Redshift Spectrum is used, all permission management is centralized in Amazon Redshift. Trusted identity propagation is supported when Amazon Quick Suite and Amazon Redshift use the same organization instance of IAM Identity Center. Trusted identity propagation is not currently supported for the following features.

- SPICE datasets
- Custom SQL on data sources
- Alerts

- Email reports
- Amazon Quick Suite Q
- CSV, Excel, and PDF exports
- Anomaly detection

For Amazon Quick Suite to connect to an Amazon Redshift instance, you must create a new security group for that instance. This security group contains an inbound rule that authorizes access from the appropriate IP address range for the Amazon Quick Suite servers in that Amazon Web Services Region. To learn more about authorizing Amazon Quick Suite connections, see [Manually enabling access to an Amazon Redshift cluster in a VPC](#).

Enabling connection from Amazon Quick Suite servers to your cluster is just one of several prerequisites for creating a data set based on an Amazon database data source. For more information about what is required, see [Creating a dataset from a database](#).

Topics

- [Enabling trusted identity propagation with Amazon Redshift](#)
- [Manually enabling access to an Amazon Redshift cluster in a VPC](#)
- [Enabling access to Amazon Redshift Spectrum](#)

Enabling trusted identity propagation with Amazon Redshift

Trusted identity propagation authenticates the end user in Amazon Redshift when they access Amazon Quick Suite assets that leverage a trusted identity propagation enabled data source. When an author creates a data source with trusted identity propagation, the identity of the data source consumers in Amazon Quick Sight is propagated and logged in CloudTrail. This allows database administrators to centrally manage data security in Amazon Redshift and automatically apply all data security rules to data consumers in Amazon Quick Suite. With other authentication methods, the data permissions of the author who created the data source are applied to all data source consumers. The data source author can choose to apply additional row and column level security to the data sources that they create in Amazon Quick Sight.

Trusted identity propagation data sources are supported only in Direct Query datasets. SPICE datasets do not currently support trusted identity propagation.

Topics

- [Prerequisites](#)
- [Enabling trusted identity propagation in Amazon Quick Sight](#)
- [Connecting to Amazon Redshift with trusted identity propagation](#)

Prerequisites

Before you get started, make sure that you have all of the required prerequisites ready.

- Trusted identity propagation is only supported for Amazon Quick Suite accounts that are integrated with IAM Identity Center. For more information, see [Configure your Amazon Quick Suite account with IAM Identity Center](#).
- An Amazon Redshift application that is integrated with IAM Identity Center. The Amazon Redshift cluster that you use must be in the same organization in Amazon Organizations as the Amazon Quick Suite account that you want to use. The cluster must also be configured with the same organization instance in IAM Identity Center that your Amazon Quick Suite account is configured to. For more information about configuring a Amazon Redshift cluster, see [Integrating IAM Identity Center](#).

Enabling trusted identity propagation in Amazon Quick Sight

To configure Amazon Quick Sight to connect to Amazon Redshift data sources with trusted identity propagation, configure Amazon Redshift OAuth scopes to your Amazon Quick Suite account.

To add a scope that allows Amazon Quick Suite to authorize identity propagation to Amazon Redshift, specify the Amazon Web Services account ID of the Amazon Quick Suite account and the service that you want to authorize identity propagation with, in this case 'REDSHIFT'.

Specify the IAM Identity Center application ARN of the Amazon Redshift cluster that you are authorizing Amazon Quick Suite to propagate user identities to. This information can be found in the Amazon Redshift console. If you don't specify authorized targets for the Amazon Redshift scope, Amazon Quick Suite authorizes users from any Amazon Redshift cluster that share the same IAM Identity Center instance. The example below configures Amazon Quick Suite to connect to Amazon Redshift data sources with trusted identity propagation.

```
aws quicksight update-identity-propagation-config --aws-account-id "AWSACCOUNTID" --  
service "REDSHIFT" --authorized-targets "arn:aws-cn:sso::XXXXXXXXXXXX:application/  
ssoins-XXXXXXXXXXXX/apl-XXXXXXXXXXXX" "arn:aws-cn:sso::XXXXXXXXXXXX:application/  
ssoins-XXXXXXXXXXXX/apl-XXXXXXXXXXXX"
```


The following example deletes OAuth scopes from a Amazon Quick Suite account.

```
aws quicksight delete-identity-propagation-config --aws-account-id "AWSACCOUNTID" --  
service "REDSHIFT" --authorized-targets "arn:aws-cn:sso::XXXXXXXXXXXX:application/  
ssoins-XXXXXXXXXXXXap1-XXXXXXXXXXXX" "arn:aws-cn:sso::XXXXXXXXXXXX:application/  
ssoins-XXXXXXXXXXXX/ap1-XXXXXXXXXXXX"
```

The following example lists all OAuth scopes that are currently on a Amazon Quick Suite account.

```
aws quicksight list-identity-propagation-configs --aws-account-id "AWSACCOUNTID"
```

Connecting to Amazon Redshift with trusted identity propagation

Use the procedure below to connect to Amazon Redshift trusted identity propagation.

To connect to Amazon Redshift with trusted identity propagation

1. Create a new dataset in Amazon Quick Suite. For more information about creating a dataset, see [Creating datasets](#).
2. Choose Amazon Redshift as the data source for the new dataset.

Note

The authentication type of an existing data source can't be changed to trusted identity propagation

3. Choose IAM Identity Center as the identity option for the data source, and then choose **Create data source**.

Manually enabling access to an Amazon Redshift cluster in a VPC

Applies to: Enterprise Edition

Use the following procedure to enable Amazon Quick Sight access to an Amazon Redshift cluster in a VPC.

To enable Amazon Quick Sight access to an Amazon Redshift cluster in a VPC

1. Sign in to the Amazon Web Services Management Console and open the Amazon Redshift console at <https://console.amazonaws.cn/redshiftv2/>.
2. Navigate to the cluster that you want to make available in Amazon Quick Suite.
3. In the **Cluster Properties** section, find **Port**. Note the **Port** value.
4. In the **Cluster Properties** section, find **VPC ID** and note the **VPC ID** value. Choose **VPC ID** to open the Amazon VPC console.
5. On the Amazon VPC console, choose **Security Groups** in the navigation pane.
6. Choose **Create Security Group**.
7. On the **Create Security Group** page, enter the security group information as follows:
 - For **Security group name**, enter **redshift-security-group**.
 - For **Description**, enter **redshift-security-group**.
 - For **VPC**, choose the VPC for your Amazon Redshift cluster. This is the VPC with the VPC ID that you noted.
8. Choose **Create security group**.

Your new security group should appear on the screen.

9. Create a second security group with the following properties.
 - For **Security group name**, enter **quicksight-security-group**.
 - For **Description**, enter **quicksight-security-group**.
 - For **VPC**, choose the VPC for your Amazon Redshift cluster. This is the VPC with the VPC ID that you noted.
10. Choose **Create security group**.
11. After you create the new security groups, create inbound rules for the new groups.

Choose the new `redshift-security-group` security group, and input the following values.

- For **Type**, choose **Amazon Redshift**.
- For **Protocol**, choose **TCP**.
- For **Port Range**, enter the port number of the Amazon Redshift cluster to which you are providing access. This is the port number that you noted in an earlier step.
- For **Source**, enter the security group ID of `quicksight-security-group`.

12. Choose **Save rules** to save your new inbound rule.
13. Repeat the previous step for quicksight-security-group and enter the following values.
 - For **Type**, choose **All traffic**.
 - For **Protocol**, choose **All**.
 - For **Port Range**, choose **All**.
 - For **Source**, enter the security group ID of redshift-security-group.
14. Choose **Save rules** to save your new inbound rule.
15. In Amazon Quick Suite, navigate to the **Manage Amazon Quick Suite** menu.
16. Choose **Manage VPC connections**, and then choose **Add VPC connection**.
17. Configure the new VPC connection with the following values.
 - For **VPC connection name**, choose a meaningful name for the VPC connection.
 - For **VPC ID**, choose the VPC in which the Amazon Redshift cluster exists.
 - For **Subnet ID**, choose the subnet for the Availability Zone (AZ) that is used for Amazon Redshift.
 - For **Security group id**, copy and paste the security group ID for quicksight-security-group.
18. Choose **Create**. It might take several minutes for the new VPC to generate.
19. In the Amazon Redshift console, navigate to the Amazon Redshift cluster that redshift-security-group is configured to. Choose **Properties**. under **Network and security settings**, enter the name of the security group.
20. In Amazon Quick Suite, choose **Datasets**, and then choose **New dataset**. Create a new dataset with the following values.
 - For **Data source**, choose **Amazon Redshift Auto-discovered**.
 - Give the data source a meaningful name.
 - The instance ID should auto populate with the VPC connection that you created in Amazon Quick Suite. If the instance ID doesn't auto populate, choose the VPC that you created from the dropdown list.
 - Enter the database credentials. If your Amazon Quick Suite account uses trusted identity propagation, choose **Single sign-on**.
21. Validate the connection, and then choose **Create data source**.

If you want to restrict the default outbound rules further, update the outbound rule of `quicksight-security-group` to allow only Amazon Redshift traffic to `redshift-security-group`. You can also delete the outbound rule that's located in the `redshift-security-group`.

Enabling access to Amazon Redshift Spectrum

Using Amazon Redshift Spectrum, you can connect Amazon Quick Suite to an external catalog with Amazon Redshift. For example, you can access the Amazon Athena catalog . You can then query unstructured data on your Amazon S3 data lake using an Amazon Redshift cluster instead of the Athena query engine.

You can also combine data sets that include data stored in Amazon Redshift and in S3. Then you can access them using the SQL syntax in Amazon Redshift.

After you've registered your data catalog (for Athena) or external schema (for a [Hive metastore](#)), you can use Amazon Quick Suite to choose the external schema and Amazon Redshift Spectrum tables. This process works just as for any other Amazon Redshift tables in your cluster. You don't need to load or transform your data.

For more information on using Amazon Redshift Spectrum, see [Using Amazon Redshift Spectrum to query external data](#) in the *Amazon Redshift Database Developer Guide*.

To connect using Redshift Spectrum, do the following:

- Create or identify an IAM role associated with the Amazon Redshift cluster.
- Add the IAM policies `AmazonS3ReadOnlyAccess` and `AmazonAthenaFullAccess` to the IAM role.
- Register an external schema or data catalog for the tables that you plan to use.

Redshift Spectrum lets you separate storage from compute, so you can scale them separately. You only pay for the queries that you run.

To connect to Redshift Spectrum tables, you don't need to grant Amazon Quick Suite access to Amazon S3 or Athena. Amazon Quick Suite needs access only to the Amazon Redshift cluster. For full details on configuring Redshift Spectrum, see [Getting started with Amazon Redshift Spectrum](#) in the *Amazon Redshift Database Developer Guide*.

Authorizing connections from Amazon Quick Suite to Amazon EC2 instances

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

For Amazon Quick Sight to connect to an Amazon EC2 instance, you must create a new security group for that instance. This security group contains an inbound rule authorizing access from the appropriate IP address range for the Quick Suite servers in that Amazon Web Services Region.

To modify the security groups for these Amazon EC2 instances, you must have Amazon credentials that permit you to access to the instances.

Enabling connection from Quick Suite servers to your instance is just one of several prerequisites for creating a data set based on an Amazon database data source. For more information about what is required, see [Creating a dataset from a database](#).

To enable Amazon Quick Suite access to an Amazon EC2 instance

1. Sign in to the Amazon Web Services Management Console and open the Amazon EC2 console at <https://console.amazonaws.cn/ec2/>.
2. If your EC2 instance is in a VPC, choose the instance to view the instance details pane. Find its VPC ID and note that ID for later use.
3. Choose **Security Groups** in the **NETWORK & SECURITY** section of the navigation pane. Then choose **Create Security Group**, as shown following.
4. Enter the security group information as follows:
 - For **Security group name**, enter **Amazon-QuickSight-access**.
 - For **Description**, enter **Amazon-QuickSight-access**.
 - For **VPC**, choose the VPC ID that you noted in step 2 if your Amazon EC2 instance is in a VPC. Otherwise, choose **No VPC**.
5. Choose **Add Rule** on the **Inbound** tab.
6. Create a new rule with the following values:
 - For **Type**, choose **Custom TCP Rule**.

- For **Protocol**, choose **TCP**.
- (Optional) For **Port Range**, enter the port number used by the instance on this Amazon EC2 instance to which you are providing access.
- For **Source**, enter the CIDR address block for the Amazon Web Services Region where you plan to use Amazon Quick Suite. For example, here is the CIDR address block for Europe (Ireland): 52.210.255.224/27. For more information on the IP address ranges for Amazon Quick Suite in supported Amazon Regions, see [Amazon Regions, websites, IP address ranges, and endpoints](#).

 Note

If you have activated Amazon Quick Suite in multiple Amazon Web Services Regions, you can create inbound rules for each Amazon Quick Suite endpoint CIDR. Doing this allows Amazon Quick Suite to have access to the Amazon RDS DB instance from any Amazon Web Services Region defined in the inbound rules.

An Amazon Quick Suite user or administrator who uses Amazon Quick Suite in multiple Amazon Regions is treated as a single user. In other words, even if you are using Amazon Quick Suite in every Amazon Web Services Region, both your Amazon Quick Suite account and your users are global.

7. Choose **Create**.
8. Choose **Instances** in the **INSTANCES** section of the navigation pane, and then choose the instance that you want to enable access to.
9. Choose **Actions**, then **Networking**, and then **Change Security Groups**.
10. In **Change Security Groups**, choose the **Amazon-QuickSight-access** security group.

Then choose **Assign Security Groups**, as shown following.

Authorizing connections through Amazon Lake Formation

Applies to: Enterprise Edition

Intended audience: System administrators

If you are querying data with Amazon Athena, you can use Amazon Lake Formation to simplify how you secure and connect to your data from Amazon Quick Sight. Lake Formation adds to the Amazon Identity and Access Management (IAM) permissions model by providing its own permissions model that is applied to Amazon analytics and machine learning services. This centrally defined permissions model controls data access at a granular level through a simple grant and revoke mechanism. You can use Lake Formation instead of, or in addition to, using scoped-down policies with IAM.

When you set up Lake Formation, you register your data sources to allow it to move the data into a new data lake in Amazon S3. Lake Formation and Athena both work seamlessly with Amazon Glue Data Catalog, making it easy to use them together. Athena databases and tables are metadata containers. These containers describe the underlying schema of the data, the data definition language (DDL) statements, and the location of the data in Amazon S3.

After Lake Formation is configured, you can use Amazon Quick Suite to access databases and tables by name or through SQL queries. Amazon Quick Suite provides a full-featured editor where you can write SQL queries. Or you can use the Athena console, the Amazon CLI, or your favorite query editor. For more information, see [Accessing Athena](#) in the *Amazon Athena User Guide*.

Use the topics below to configure a Lake Formation connection through Lake Formation or through Amazon Quick Suite.

Topics

- [Enabling connection from Lake Formation](#)
- [Enabling connection from Amazon Quick Suite](#)

Enabling connection from Lake Formation

Before you begin using this solution with Quick Suite, make sure that you can access your data using Athena with Lake Formation. After you verify that the connection is working through Athena, you need to verify only that Amazon Quick Suite can connect to Athena. Doing this means you don't have to troubleshoot connections through all three products at once. One easy way to test the connection is to use the [Athena query console](#) to run a simple SQL command, for example `SELECT 1 FROM table`.

To set up Lake Formation, the person or team who works on it needs access to create a new IAM role and to Lake Formation. They also need the information shown in the following list. For more information, see [Setting up lake formation](#) in the *Amazon Lake Formation Developer Guide*.

- Collect the Amazon Resource Names (ARNs) of the Quick Suite users and groups that need to access the data in Lake Formation. These users should be Amazon Quick Suite authors or administrators.

To find Amazon Quick Suite user and group ARNs

1. Use the Amazon CLI to find user ARNs for Amazon Quick Suite authors and admins. To do this, run the following `list-users` command in your terminal (Linux or Mac) or at your command prompt (Windows).

```
aws quicksight list-users --aws-account-id 111122223333 --namespace default --  
region us-east-1
```

The response returns information for each user. We show the Amazon Resource Name (ARN) in bold in the following example.

```
RequestId: a27a4cef-4716-48c8-8d34-7d3196e76468  
Status: 200  
UserList:  
- Active: true  
  Arn: arn:aws-cn:quicksight:us-east-1:111122223333:user/default/SaanviSarkar  
  Email: SaanviSarkar@example.com  
  PrincipalId: federated/iam/AIDAJVCZ0VSR3DESMJ7TA  
  Role: ADMIN  
  UserName: SaanviSarkar
```

To avoid using the Amazon CLI, you can construct the ARNs for each user manually.

2. (Optional) Use the Amazon CLI to find ARNs for Amazon Quick Suite groups by running the following `list-group` command in your terminal (Linux or Mac) or at your command prompt (Windows).

```
aws quicksight list-groups --aws-account-id 111122223333 --namespace default --  
region us-east-1
```

The response returns information for each group. The ARN appears in bold in the following example.

```
GroupList:
```



```
- Arn: arn:aws-cn:quicksight:us-east-1:111122223333:group/default/DataLake-Scorecard
  Description: Data Lake for CX0 Balanced Scorecard
  GroupName: DataLake-Scorecard
  PrincipalId: group/d-90671c9c12/6f9083c2-8400-4389-8477-97ef05e3f7db
  RequestId: c1000198-18fa-4277-a1e2-02163288caf6
  Status: 200
```

If you don't have any Amazon Quick Suite groups, add a group by using the Amazon CLI to run the `create-group` command. There currently isn't an option to do this from the Amazon Quick Suite console. For more information, see [Creating and managing groups in Amazon Quick Suite](#).

To avoid using the Amazon CLI, you can construct the ARNs for each group manually.

Enabling connection from Amazon Quick Suite

To work with Lake Formation and Athena, make sure that you have Amazon resource permissions configured in Amazon Quick Suite:

- Enable access to Amazon Athena.
- Enable access to the correct buckets in Amazon S3. Usually S3 access is enabled when you enable Athena. However, because you can change S3 permissions outside of that process, it's a good idea to verify them separately.

For information about how to verify or change Amazon resource permissions in Quick Suite, see [Allowing autodiscovery of Amazon resources](#) and [Accessing data sources](#).

Authorizing connections to Amazon OpenSearch Service

Applies to: Enterprise Edition

Intended audience: System administrators

Before you can use OpenSearch in a Amazon Quick Sight dataset, there are a few tasks for the Quick Suite administrator to complete with the cooperation of a person who has access to the OpenSearch console.

To get started, identify each OpenSearch domain that you want to connect to. Then gather the following information for each domain:

- The name of the OpenSearch domain.
- The OpenSearch version used by this domain.
- The Amazon Resource Name (ARN) of the OpenSearch domain.
- The HTTPS endpoint.
- The OpenSearch Dashboards URL, if you use Dashboards. You can extrapolate the Dashboards URL by appending `"/dashboards/"` to an endpoint.
- If the domain has a VPC endpoint, gather all the related information on the VPC tab of the OpenSearch Service console:
 - The VPC ID
 - The VPC security groups
 - The associated IAM role or roles
 - The associated Availability Zones
 - The associated subnets
- If the domain has a regular endpoint (not a VPC endpoint), note that it uses the public network.
- The start hour for the daily automated snapshot (if your users want to know).

Before you proceed, the Amazon Quick Suite administrator enables authorized connections from Amazon Quick Suite to OpenSearch Service. This process is required for every Amazon service that you connect to from Amazon Quick Suite. You need to do this only once per Amazon Web Services account for each Amazon service that you use as a data source.

For OpenSearch Service, the authorization process adds the Amazon managed policy `AWSQuickSightOpenSearchPolicy` to your Amazon Web Services account.

Important

Make sure that the IAM policy for your OpenSearch domain doesn't conflict with the permissions in `AWSQuickSightOpenSearchPolicy`. You can find the domain access

policy in the OpenSearch Service console. For more information, see [Configuring access policies](#) in the *Amazon OpenSearch Service Developer Guide*.

To turn on or turn off connections from Amazon Quick Suite to OpenSearch Service

1. Within Amazon Quick Suite, choose **Administrator** and **Manage Amazon Quick Suite**.
2. Choose **Security & permissions, Add or remove**.
3. To enable connections, select the **Amazon OpenSearch Service** check box.

To disable connections, clear the **Amazon OpenSearch Service** check box.

4. Choose **Update** to confirm your choices.

If needed, use the topics below to configure a OpenSearch VPC connection and permissions for Amazon Quick Suite to access OpenSearch.

Topics

- [Using a VPC connection](#)
- [Using OpenSearch permissions](#)

Using a VPC connection

In some cases, your OpenSearch domain is in a virtual private cloud (VPC) based on the Amazon VPC service. If so, make sure to determine if Amazon Quick Suite is already connected to the VPC ID that the OpenSearch domain uses. You can reuse an existing VPC connection. If you're not sure if it's working, you can test it. For more information, see [Testing the connection to your Amazon VPC data source](#).

If a connection isn't already defined in Amazon Quick Suite for the VPC that you want to use, you can create one. This task is a multistep process that you need to complete before you proceed. To learn how to add Amazon Quick Suite to a VPC and add a connection from Amazon Quick Suite to the VPC, see [Connecting to a Amazon VPC with Amazon Quick Suite](#).

Using OpenSearch permissions

After you configure Amazon Quick Suite to connect to OpenSearch Service, you might need to enable permissions in OpenSearch. For this part of the setup process, you can use the OpenSearch

Dashboards link for each OpenSearch domain. Use the following list to help determine what permissions you need:

1. For domains that use fine-grained access control, configure permissions in the form of a role. This process is similar to using scoped-down policies in Amazon Quick Suite.
2. For each domain that you create a role for, add a role mapping.

For more information, see following.

If your OpenSearch domain has [fine-grained access control](#) enabled, there are some permissions to configure so the domain is accessible from Amazon Quick Suite. Perform these steps for each domain that you want to use.

The following procedure uses OpenSearch Dashboards, which is an open-source tool that works with OpenSearch. You can find the link to Dashboards on the domain dashboard on the OpenSearch Service console.

To add permissions to a domain to allow access from Amazon Quick Suite

1. Open OpenSearch Dashboards for the OpenSearch domain that you want to work with. The URL is *opensearch-domain-endpoint*/dashboards/.
2. Choose **Security** from the navigation pane.

If you don't see the navigation pane, open it by using the menu icon at upper left. To keep the menu open, choose **Dock navigation** at lower left.

3. Choose **Roles, Create role**.
4. Name the role **quicksight_role**.

You can choose a different name, but we recommend this one because we use it in our documentation and it's thus easier to support.

5. Under **Cluster permissions**, add the following permissions:

- cluster:monitor/main
- cluster:monitor/health
- cluster:monitor/state
- indices:data/read/scroll
- indices:data/read/scroll/clear,

6. Under **Index permissions** specify ***** as the index pattern.
7. For **Index permissions**, add the following permissions:
 - `indices:admin/get`
 - `indices:admin/mappings/get`
 - `indices:admin/mappings/fields/get*`
 - `indices:data/read/search*`
 - `indices:monitor/settings/get`
8. Choose **Create**.
9. Repeat this procedure for each OpenSearch domain that you're planning to use.

Use the following procedure to add a role mapping for the permissions that you added in the previous procedure. You might find it more efficient to add the permissions and the role mapping as part of a single process. These instructions are separate for clarity.

To create a role mapping for the IAM role you added

1. Open OpenSearch Dashboards for the OpenSearch domain that you want to work with. The URL is *opensearch-domain-endpoint*/dashboards/.
2. Choose **Security** from the navigation pane.
3. Search for and open **quicksight_role** from the list.
4. On the **Mapped users** tab, choose **Manage mapping**.
5. In the **Backend roles** section, enter the ARN of the Amazon-managed IAM role for Amazon Quick Suite. Following is an example.

```
arn:aws:iam::AWS-ACCOUNT-ID:role/service-role/aws-quicksight-service-role-v0
```

6. Choose **Map**.
7. Repeat this procedure for each OpenSearch domain that you want to use.

Authorizing connections to Amazon Athena

If you need use Amazon Quick Sight with Amazon Athena or Amazon Athena Federated Query, you first need to authorize connections to Athena and the associated buckets in Amazon Simple Storage Service (Amazon S3). Amazon Athena is an interactive query service that makes it easy to

analyze data directly in Amazon S3 using standard SQL. Athena Federated Query provides access to more types of data by using Amazon Lambda. Using a connection from Quick Suite to Athena, you can write SQL queries to interrogate data that's stored in relational, non-relational, object, and custom data sources. For more information, see [Using Athena federated query](#) in the Amazon Athena User Guide.

Review the following considerations when setting up access to Athena from Quick Suite:

- Athena stores query results from Amazon Quick Sight in a bucket. By default, this bucket has a name similar to `aws-athena-query-results-AWSREGION-AWSACCOUNTID`, for example `aws-athena-query-results-us-east-2-111111111111`. Therefore, it's important to make sure Amazon Quick Sight has permissions to access the bucket Athena is currently using.
- If your data file is encrypted with an Amazon KMS key, grant permissions to the Amazon Quick Sight IAM role to decrypt the key. The easiest way to do this is to use the Amazon CLI.

You can run the KMS [create-grant](#) API operation in Amazon CLI to do this.

```
aws kms create-grant --key-id <KMS_KEY_ARN> /  
--grantee-principal <QS_ROLE_ARN> --operations Decrypt
```

The Amazon Resource Name (ARN) for the Amazon Quick Suite role has the format `arn:aws-cn:iam::<account id>:role/service-role/aws-quicksight-s3-consumers-role-v<version number>` and can be accessed from the IAM console. To find your KMS key ARN, use the S3 console. Go to the bucket that contains your data file and choose the **Overview** tab. The key is located near **KMS key ID**.

- For Amazon Athena, Amazon S3, and Athena Query Federation connections, Amazon Quick Suite uses the following IAM role by default:

```
arn:aws-cn:iam::AWS-ACCOUNT-ID:role/service-role/aws-quicksight-s3-consumers-role-v0
```

If the `aws-quicksight-s3-consumers-role-v0` is not present, then Amazon Quick Suite uses:

```
arn:aws-cn:iam::AWS-ACCOUNT-ID:role/service-role/aws-quicksight-service-role-v0
```

- If you assigned scope-down policies to your users, verify that the policies contain the `lambda:InvokeFunction` permission. Without this permission, your users can't access Athena Federated Queries. For more information about assigning IAM policies to your users in Amazon

Quick Suite, see [Setting granular access to Amazon services through IAM](#). For more information about the `lambda:InvokeFunction` permission, see [Actions, resources, and condition keys for Amazon Lambda](#) in the *IAM User Guide*.

To authorize Amazon Quick Suite to connect to Athena or Athena federated data sources

1. (Optional) If you are using Amazon Lake Formation with Athena, you also need to enable Lake Formation. For more information, see [Authorizing connections through Amazon Lake Formation](#).
2. Open your profile menu at top right and choose **Manage QuickSight**. You must be a Amazon Quick Suite administrator to do this. If you don't see **Manage QuickSight** on the profile menu, you don't have sufficient permissions.
3. Choose **Security & permissions, Add or remove**.
4. Choose the box near Amazon Athena, **Next**.

If it was already enabled, you might have to double-click it. Do this even if Amazon Athena is already enabled, so you can view the settings. No changes are saved until you choose **Update** at the end of this procedure.

5. Enable the S3 buckets you want to access.
6. (Optional) To enable Athena federated queries, select the Lambda functions you want to use.

Note

You can only see Lambda functions for the Athena catalogs in the same region of Amazon Quick Suite.

7. To confirm your changes, choose **Finish**.

To cancel, choose **Cancel**.

8. To save changes to security and permissions, choose **Update**.

To test the connection authorization settings

1. From the Amazon Quick Suite start page, choose **Datasets, New dataset**.
2. Choose the Athena card.

3. Follow the screen prompts to create a new Athena data source using the resources you need to connect to. Choose **Validate connection** to test the connection.
4. If the connection validates, you have successfully configured an Athena or Athena Federated Query connection.

If you don't have sufficient permissions to connect to an Athena dataset or run an Athena query, an error displays directing you to contact a Amazon Quick Suite administrator. This error means need to recheck your connection authorization settings to find the discrepancy.

5. After you can connect successfully, you or your Amazon Quick Suite authors can create data sources connections and share them with other Amazon Quick Suite authors. The authors can then create multiple datasets from the connections, to use in Amazon Quick Suite dashboards.

For troubleshooting information on Athena, see [Connectivity issues when using Athena with Amazon Quick Suite](#).

Using trusted identity propagation with Athena

Trusted identity propagation gives Amazon services access to Amazon resources based on the user's identity context and securely shares this user's identity with other Amazon services. These capabilities enable user access to be more easily defined, granted, and logged.

When administrators configure Quick Suite, Athena, Amazon S3 Access Grants, and Amazon Lake Formation with IAM Identity Center, they can now enable trusted identity propagation across these services and allow the user's identity to be propagated across services. When data is accessed from Quick Suite by an IAM Identity Center user, Athena or Lake Formation can make authorization decisions using the permissions defined for their user or group membership from the organization's identity provider.

Trusted identity propagation with Athena only works when permissions are managed through Lake Formation. User permissions to data reside in Lake Formation.

Prerequisites

Before you get started, make sure that you have the following required prerequisites completed.

⚠ Important

As you complete the following prerequisites, note that your IAM Identity Center instance, Athena workgroup, Lake Formation and Amazon S3 Access Grants must all be deployed in the same Amazon Region.

- Configure your Quick Suite account with IAM Identity Center. Trusted identity propagation is only supported for Quick Suite accounts that are integrated with IAM Identity Center. For more information, see [Configure your Amazon Quick Suite account with IAM Identity Center](#).

ℹ Note

To create Athena data sources, you must be an IAM Identity Center user (author) in a Quick Suite account that uses IAM Identity Center.

- An Athena workgroup that is enabled with IAM Identity Center. The Athena workgroup that you use must be using the same IAM Identity Center instance as the Quick Suite account. For more information about configuring an Athena workgroup, see [Creating an IAM Identity Center enabled Athena workgroup](#) in the *Amazon Athena User Guide*.
- Access to Athena query results bucket is managed with Amazon S3 Access Grants. For more details, see [Managing access with Amazon S3 Access Grants](#) in the *Amazon S3 User Guide*. If your query results are encrypted with an Amazon KMS key, the Amazon S3 Access Grant IAM role and the Athena workgroup role both need permissions to Amazon KMS.
 - For more information, see [Amazon S3 Access Grants and corporate directory identities](#) in the Amazon S3 User Guide.
 - The Amazon S3 Access Grant role should have the `STS:SetContext` action in its trust policy for identity propagation. To see an example, see [Register a location](#) in the Amazon S3 User Guide.
- Permissions to data must be managed with Lake Formation and Lake Formation must be configured with the same IAM Identity Center instance as Quick Suite and the Athena workgroup. For configuration information, see [Integrating IAM Identity Center](#) in the *Amazon Lake Formation Developer Guide*.
- The data lake administrator needs to grant permissions to IAM Identity Center users and groups in Lake Formation. For more details, [Granting permissions to users and groups](#) in the *Amazon Lake Formation Developer Guide*.

- The Quick Suite administrator needs to authorize connections to Athena. For details, see [Authorizing connections to Amazon Athena](#). Note, with trusted identity propagation, you do not need to give the Quick Suite role Amazon S3 bucket permissions or Amazon KMS permissions. You need to keep your users and groups that have permissions to the workgroup in Athena in sync with the Amazon S3 bucket that stores query results with Amazon S3 Access Grants permissions so that users can successfully run queries and retrieve query results in the Amazon S3 bucket using trusted identity propagation.

Configure IAM role with required permissions

To use trusted identity propagation with Athena, your Quick Suite account must have the required permissions to access your resources. To provide those permissions, you must configure your Quick Suite account to use an IAM role with the permissions.

If your Quick Suite account is already using a custom IAM role, you can modify that one. If you do not have an existing IAM role, create one by following the instructions in [Create a role for an IAM user](#) in the *IAM User Guide*.

The IAM role you create or modify must contain the following trust policy and permissions.

Required trust policy

For information about updating the trust policy of an IAM role, see [Update a role trust policy](#).

Required Athena permissions

For information about updating the trust policy of an IAM role, see [Update permissions for a role](#).

Note

The Resource uses the * wildcard. We recommend that you update it to include only the Athena resources you want to use with Quick Suite.

JSON

```
{  
  "Version": "2012-10-17",
```

```

    "Statement": [
      {
        "Effect": "Allow",
        "Action": [
          "athena:BatchGetQueryExecution",
          "athena:CancelQueryExecution",
          "athena:GetCatalogs",
          "athena:GetExecutionEngine",
          "athena:GetExecutionEngines",
          "athena:GetNamespace",
          "athena:GetNamespaces",
          "athena:GetQueryExecution",
          "athena:GetQueryExecutions",
          "athena:GetQueryResults",
          "athena:GetQueryResultsStream",
          "athena:GetTable",
          "athena:GetTables",
          "athena:ListQueryExecutions",
          "athena:RunQuery",
          "athena:StartQueryExecution",
          "athena:StopQueryExecution",
          "athena:ListWorkGroups",
          "athena:ListEngineVersions",
          "athena:GetWorkGroup",
          "athena:GetDataCatalog",
          "athena:GetDatabase",
          "athena:GetTableMetadata",
          "athena:ListDataCatalogs",
          "athena:ListDatabases",
          "athena:ListTableMetadata"
        ],
        "Resource": "*"
      }
    ]
  }
}

```

Configure your Quick Suite account to use the IAM role

After configuring the IAM role in the previous step, you must configure your Quick Suite account to use it. For information about how to do that, see [Using existing IAM roles in Quick Suite](#).

Update the identity propagation config with the Amazon CLI

To authorize Quick Suite to propagate end user identities to Athena workgroups, run the following `update-identity-propagation-config` API from the Amazon CLI, replacing the following values:

- Replace `us-west-2` with the Amazon Region that your IAM Identity Center instance is in.
- Replace `111122223333` with your Amazon account ID.

```
aws quicksight update-identity-propagation-config \  
--service ATHENA \  
--region us-west-2 \  
--aws-account-id 111122223333
```

Create an Athena dataset in Quick Suite

Now, create an Athena dataset in Quick Suite configured with the IAM Identity Center enabled Athena workgroup you want to connect to. For information about how to create an Athena dataset, see [Creating a dataset using Amazon Athena data](#).

Key callouts, considerations, and limits

The following list contains some important considerations when using trusted identity propagation with Quick Suite and Athena.

- Quick Suite Athena data sources that use trusted identity propagation have Lake Formation permissions evaluated against the IAM Identity Center end user and the IAM Identity Center groups that the user might belong to.
- When using Athena data sources that use trusted identity propagation, we recommend any fine tuned access control is done in Lake Formation. However, If you elect to use Quick Suite's scope down policy feature, scope down policies will be evaluated against the end user.
- The following features are disabled for data sources and data sets that use trusted identity propagation: SPICE datasets, Custom SQL on data sources, threshold alerts, email reports, Q Topics, stories, scenarios, CSV, Excel, and PDF exports, anomaly detection.
- If you experience high latency or timeouts, it may be because of a combination of high number of IAM Identity Center groups, Athena databases, tables, and Lake Formation rules. We recommend trying to use only the necessary number of those resources.

Data access integrations

Data access integrations in Amazon Quick Suite establish secure connections to external data sources. They serve as the foundation for creating knowledge bases. Unlike action connectors that perform actions, data access integrations focus on accessing and indexing content from third-party applications and services.

Data access integrations only configure authentication and point to the project or organization of the service. They cannot be used directly for analysis or by AI agents. You must create a connected knowledge base to make the data accessible.

How data access integrations work

Data access integrations configure authentication and establish connections to third-party service organizations or projects. You cannot use these integrations directly for analysis. You must create knowledge bases connected to the data access integration to make the data accessible for AI agents, chat interfaces, and spaces.

The relationship between data access integrations and knowledge bases is one-to-many:

- One data access integration can support multiple knowledge bases.
- Each knowledge base selects specific content from the connected data source.
- Knowledge bases inherit authentication and access permissions from their parent data access integration.

Create a data access integration

Use the following procedure to create a data access integration that establishes authentication and connection details for knowledge base creation. The following example demonstrates the process for setting up a Microsoft OneDrive data access integration, but the general steps are applicable to other data access integrations.

To create a data access integration

1. Scroll to the Set up a new integration area of the page. Find the application that you want to create an integration and knowledge base. Select "OneDrive".

Note

The Integration page defaults Knowledge bases tab and there may be existing knowledge bases that have been set up by others and shared. If you have previously set up an integration, check the data tab and use the action menu to create a knowledge base from there.

2. Select the plus (+) icon button on the application to create a new integration and knowledge base.
3. Select Bring data from Microsoft OneDrive option and click Next button.

Note

Some application integrations support data ingestion and read/write actions. The set up varies for each one. To set up actions, you'll need more information from your admin.

4. Complete the authentication process:
 - a. A Microsoft OneDrive sign in popup will appear automatically. If it doesn't, click Sign in to Microsoft OneDrive button.
 - b. Sign in using your Amazon credentials.
 - c. Wait until a success banner appears.
 - d. Click the "Next" button.
5. Select the data that should be ingested into the knowledge base using the file picker for OneDrive and click Add button.
6. Type in a knowledge base a Name and Description (optional), then click Create.
7. There will be a success toast notification and the data ingestion and sync will begin.
8. The data can take several minutes to sync, depending on the number files that are being ingested. The Status column will stay in the Syncing status until it is ready changing to Available.
9. When the knowledge base is ready, use the chat to ask questions and interact with it.

Note

By default, the chat uses 'all data and apps' that you have access to and that are set up on your behalf. If you want to chat with a single knowledge base, select the knowledge base in the chat data picker.

Note

You can also attach a knowledge base to a Space by navigating to the Space and adding it.

After successful creation, your data access integration appears in the integrations list. You can now create knowledge bases that use this integration to access and index content from the connected data source.

Note

For detailed configuration steps specific to each data source, see [the section called "Supported integrations"](#).

Supported data sources

Amazon Quick Suite supports data access integrations with the following applications and services. These integrations allow you to create knowledge bases from external data sources:

- **Amazon S3** - Access documents and files stored in S3 buckets using Amazon credentials.
- **Atlassian Confluence** - Index pages, spaces, and attachments using user authentication or service authentication.
- **Google Drive** - Connect to personal and shared drives using OAuth 2.0 authentication.
- **Microsoft OneDrive** - Access OneDrive for Business content using user authentication or service authentication.
- **Microsoft SharePoint** - Index SharePoint Online and Server content using OAuth 2.0 authentication.

- **Web Crawler** - Index content from internal and external websites using basic authentication or form/SAML authentication.

Each data source supports different authentication methods and content access capabilities. The relationship between data access integrations and knowledge bases is one-to-many - one integration can support multiple knowledge bases, each selecting specific content from the connected data source.

Data source categories

Data access integrations are organized into the following categories based on the type of content and access patterns:

Cloud storage and file systems

- Amazon S3 - Access documents and files stored in S3 buckets.
- Google Drive - Index content from personal and shared drives.
- Microsoft OneDrive - Connect to OneDrive for Business content.

Content management systems

- Atlassian Confluence - Access pages, spaces, and attachments.
- Microsoft SharePoint - Index SharePoint Online and Server content.

Web content

- Web Crawler - Index content from internal and external websites.

Authentication and security

Data access integrations use secure authentication methods to protect your data and maintain access controls. The authentication method depends on the specific data source and your organization's security requirements.

OAuth authentication

Most cloud-based integrations (Google Drive, OneDrive, Confluence Cloud) use OAuth 2.0 for secure, token-based authentication. This method allows Amazon Quick Suite to access your data without storing your credentials.

Service account authentication

Enterprise integrations may use service accounts for programmatic access. This method is common for Amazon S3 and other infrastructure-based data sources.

No authentication

Some integrations, such as web crawlers accessing public websites, may not require authentication. However, access controls are still enforced based on your Amazon Quick Suite permissions.

Note

Authentication requirements and available methods vary by user tier. Readers may have limited authentication options compared to Authors.

Access control and permissions

Data access integrations maintain security by enforcing access controls at multiple levels. When users query content through knowledge bases, Amazon Quick Suite ensures they can only access content they have permission to view.

- **Source-level permissions** - Users must have appropriate permissions in the source system (Google Drive, SharePoint, etc.).
- **Integration-level permissions** - Access to the integration itself is controlled by Amazon Quick Suite permissions.
- **Knowledge base permissions** - Individual knowledge bases can have their own access controls.
- **Entity-level access controls** - When users query content, Amazon Quick Suite verifies permissions for each document or item.

Key features and capabilities

Data access integrations provide several features to enhance your data integration experience:

- **Real-time synchronization** - Content is automatically updated when changes occur in the source system.

- **Selective indexing** - Choose specific folders, sites, or content types to include in your knowledge bases.
- **Content type support** - Index various file formats including documents, spreadsheets, presentations, and web pages.
- **Metadata preservation** - Maintain important metadata such as creation dates, authors, and tags.
- **Natural language querying** - Enable AI-powered search and question-answering across your indexed content.

Before you begin

Before creating data access integrations, ensure you have the following requirements in place:

- **Amazon Quick Suite permissions** - Author or Admin role to create and manage integrations.
- **Source system access** - Appropriate permissions in the target system (administrative access may be required for some integrations).
- **Authentication credentials** - Valid credentials or service accounts for the target system.
- **Network connectivity** - Ensure Amazon Quick Suite can access your data sources. Network requirements differ by integration type:
 - **Knowledge bases** - Do not support VPC connectivity. Data sources must be accessible over the public internet.
 - **Action connectors** - Support VPC connectivity for resource servers within your VPC. However, authentication servers must remain publicly accessible.

Exploring your Amazon data in Amazon Quick Suite

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Use this section to learn how to explore Amazon data in Amazon Quick Suite using the Amazon Web Services Management Console. Using the **Explore in Amazon Quick Sight** shortcut, you can access a customizable dashboard template showing your data. Just as with any Amazon Quick

Sight dashboard, this dashboard can be refreshed on a schedule, published, and shared with other users in your organization.

To use this feature, you must first enable Amazon S3 analytics storage class analysis for your Amazon S3 buckets. For more on enabling storage class analysis in Amazon S3, see [Amazon Amazon S3 analytics – Storage class analysis](#) in the *Amazon Amazon S3 Developer Guide*.

After you have enabled storage class analysis, you can use Amazon Quick Suite to explore your Amazon S3 analytics data.

To explore Amazon S3 analytics data in Amazon Quick Suite

1. Open the Amazon S3 console at <https://console.amazonaws.cn/s3/>.
2. Choose a bucket to explore. The bucket must have storage class analysis enabled, with at least one filter.
3. Choose the **Management** tab.
4. Then choose **Analytics**.
5. Choose **Explore in QuickSight**.

Note

If you don't have an Amazon Quick Suite account, you're prompted to create one before you can use the dashboard.

When you choose the option to explore in Amazon Quick Suite, your Amazon S3 analytics data is automatically loaded into the dashboard template. The dashboard contains multiple visualizations to help you to understand the storage access pattern of your bucket.

Use the template as is, or customize it to suit your needs. For example, one visual on the default template helps you identify infrequently accessed data. It compares the amount of data retrieved to the amount of storage consumed, for objects of different ages.

You can also add your own visualizations to the dashboard. For example, you can break down the data access patterns, using filters for storage class analysis that you already have defined in Amazon S3 analytics.

To learn more about using S3 analytics and storage class analysis, see [Amazon Amazon S3 analytics – Storage class analysis](#) in the *Amazon Amazon S3 Developer Guide*.

Amazon service action connectors

With Amazon service action connectors in Amazon Quick Suite, you can create action connectors that interact directly with Amazon services like Amazon Bedrock, Amazon Textract, and Amazon Comprehend. These connectors enable automated workflows that leverage Amazon AI and machine learning capabilities.

What you can do

Amazon service action connectors enable you to integrate powerful Amazon capabilities into your automated workflows. For example, you can use Amazon Bedrock to generate content with foundation models, Amazon Textract to extract text and data from documents, or Amazon Comprehend to analyze sentiment and extract insights from text. These action connectors allow you to build sophisticated automation workflows that combine multiple Amazon services for document processing, content generation, and data analysis—all while maintaining security through IAM role-based authentication.

Note

Amazon services action connectors can only be used with Amazon Quick Automate because they require an IAM identity for authentication. These connectors are created through the admin console and provide direct access to Amazon service APIs.

Supported Amazon services

Amazon Quick Suite supports the following Amazon services for action connectors:

- **Amazon Bedrock Agent** - Invoke Bedrock agents for complex AI workflows.
- **Amazon Bedrock Runtime** - Access foundation models for text generation and conversation.
- **Amazon Bedrock Data Automation** - Automate data processing workflows with AI.
- **Amazon Comprehend** - Analyze text for sentiment, entities, and key phrases.
- **Amazon Comprehend Medical** - Extract medical information from healthcare text.
- **Amazon Textract** - Extract text and data from documents and images.
- **Amazon S3** - Manage objects and buckets in S3.

Before you begin

Before you set up an Amazon service action connectors, make sure you have the following:

- Amazon account with access to the desired Amazon services.
- IAM role with appropriate permissions for the Amazon services you want to use.
- Amazon Quick Suite Admin access to create action connectors.
- Amazon Quick Automate access to use the action connectors in workflows.

Prepare IAM role and permissions

Before setting up the action connectors in Amazon Quick Suite, prepare your IAM role with the necessary permissions for the Amazon services you want to use.

Required IAM permissions

Configure your IAM role with the appropriate permissions based on the Amazon services you plan to use. Make sure your IAM role includes a trust policy that allows Amazon Quick Suite to assume the role for executing actions.

IAM Role for Resource Access

First, you will need to create an IAM role that will be used by Amazon Quick Suite to call the Amazon service needed in your Amazon Quick Automate workflow.

1. Login to the Amazon Console of the Amazon account where the Amazon Quick Suite subscription resides.
2. Open IAM and create a new IAM role.
3. Give it all the permissions for the Amazon service you want to invoke via action connectors. For example, you can assign a managed policy like `AmazonS3FullAccess` if you need to invoke Amazon S3.
4. In the trust relationship, give the assume role permission to `quicksight.amazonaws.com`. This allows Amazon Quick Suite to assume this role and call Amazon S3/Amazon Textract/etc on your behalf.
5. Once the Customer Role is created, take a note of the IAM role ARN.

Example trust policy:

Note

For pre-prod stages, you will also have to include "quicksight-test.amazonaws.com" in the "Service" array.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Statement1",
      "Effect": "Allow",
      "Principal": {
        "Service": [
          "quicksight.amazonaws.com"
        ]
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

Create Amazon services action connector

After preparing your IAM role and permissions, follow these steps to create your Amazon services action connector:

1. In the Amazon Quick Suite console, choose the user icon in the upper right corner.
2. From the dropdown menu, select **Manage Quick Suite**. This takes you to the Admin console.
3. In the left navigation panel, select **Amazon actions**.
4. Choose **New Action**.
5. Choose one of the supported Amazon services:
 - Bedrock Agent
 - Bedrock Runtime
 - Bedrock Data Automation
 - Comprehend
 - Comprehend Medical

- **Texttract**
 - **Amazon S3**
6. Select **Next**.
 7. Review all of the supported actions in the list for your selected service.
 8. Select **Next**.
 9. Customize the connection details:
 - **Name** - Enter a descriptive name for your action connector.
 - **Description** - Describe the purpose of this action connector.
 - **Role ARN** - Enter the ARN of the IAM role that provides the proper permissions for the Amazon service.
 10. Select **Next**.
 11. Add users you want to share the action connector with.
 12. Select **Next** to finish creating the action connector.

After creating the action connector, it becomes available for use in Amazon Quick Automate workflows. The connector will be found in Actions tab under integration.

Use action connectors in workflows

After setting up your Amazon services action connector, you can use it in Amazon Quick Automate workflows.

Access in Amazon Quick Automate

To use your Amazon action connector in Amazon Quick Automate:

1. Open Amazon Quick Automate and go to Projects.
2. Click Manage groups and select your group.
3. In the Assets sections, click **Add** → **Actions**.
4. Select the Action you created earlier and Add to the automation group.
5. Go to the automation in the group and your action will be available for use.
6. In the Actions panel, look for your Amazon service action connector.
7. Drag the desired action into your workflow.

8. Configure the action parameters as needed for your use case.
9. Test and save your workflow.

Manage existing connectors

You can modify your existing Amazon action connectors:

1. In the Admin console, choose **Amazon actions**.
2. Select your action connector from the list.
3. Choose the actions menu, then choose **Edit**.
4. Update your configuration settings as needed and choose **Save**.

Note

Amazon services action connectors use IAM role-based authentication and can only be used within Amazon Quick Automate workflows. Make sure your IAM role has the minimum required permissions for the Amazon services you plan to use.

Amazon Quick Suite account management

If you are an Amazon Quick Suite administrator, the account you use to sign in to Quick Suite is in the ADMIN Quick Suite group. There are also some permissions granted through IAM, which you might have already, or you can talk to your Amazon Web Services account administrators to learn more.

Use the following topics to manage Quick Suite.

Topics

- [Account details in Amazon Quick Suite](#)
- [Amazon Quick Suite asset management](#)
- [Manage Amazon Q permissions](#)

Account details in Amazon Quick Suite

Use this section to change the account-wide settings in Amazon Quick Suite. To open your account settings, choose the profile icon, and then select **Manage Quick Suite**. Then, click on **Account settings**.

Deleting your Amazon Quick Suite subscription and closing your account

Warning

Deleting your Amazon Quick Suite subscription permanently destroys all organizational data across all Amazon regions and cannot be reversed.

This action will immediately and permanently delete:

- Every dashboard, analysis, and dataset across your entire organization in all regions
- All user accounts and their associated work
- Custom data sources and connections
- All sharing permissions and folder structures
- Historical usage data and audit logs
- All customizations, themes, and organizational settings

Account deletion from any region affects your entire global Quick Suite deployment.

Even if you have Quick Suite resources in multiple Amazon regions, this single action deletes everything worldwide.

There is no recovery option. Once deleted, this data cannot be restored by Amazon or recovered through any means.

Before proceeding: Export critical dashboards, back up datasets, and ensure all stakeholders are informed of this permanent action.

The act of deleting your Amazon Quick Suite subscription is immediate and final. Deletion removes every Amazon Quick Suite asset on the Amazon Web Services account you are using. It doesn't delete namespaces that you added. (The **Default** namespace is deleted automatically.) You can locate and delete namespaces by using the API operations [ListNamespaces](#) and [DeleteNamespace](#).

You can terminate your Amazon Quick Suite account from the **Manage Quick Suite** menu or by using the API. To prevent someone from deleting a Amazon Quick Suite user account accidentally

or maliciously, Amazon Quick Suite uses permissions, a switch for the **Account termination protection** setting, and a required confirmation word.

After your account is deleted, you can create a new Amazon Quick Suite account. The process doesn't take more than 15 minutes. The settings for edition and user authorization method on the new account can be the same or different.

Before you can delete your Amazon Quick Suite account, make sure of the following:

- You're signed in using the IAM account or Amazon root account that was used to create your Amazon Quick Suite account.
- You understand that your Amazon Web Services account is not deleted when you terminate your Amazon Quick Suite account. To instead close your Amazon Web Services account, see [Closing an Amazon Web Services account](#).
- Terminating your account deletes all users, all uploaded data, and assets (for example, datasets, data sources, queries, dashboards, analyses, settings, and so on).

Topics

- [Permissions and access to account termination](#)
- [Deleting your account](#)

Permissions and access to account termination

You need the following special permissions to terminate a Amazon Quick Suite account. Without these permissions, you won't be able to terminate a Amazon Quick Suite user account. Contact your account administrator for help.

- You're a Amazon Quick Suite administrator and have an Admin role in Amazon Quick Suite.
- You need permissions to run the following (except if you're the root admin user (IAM) who added Amazon Quick Suite)
 - quicksight:Unsubscribe
 - ds:UnauthorizeApplication
 - ds>DeleteDirectory
 - ds:DescribeDirectories
 - quicksight:UpdateAccountSettings

- To remove custom namespaces, you need permission to run the following API operations:
 - `quicksight:ListNamespaces`
 - `quicksight>DeleteNamespace`

You don't need extra permissions to delete the default namespace.

Warning

Terminating your account is an instant action that cannot be undone by you or by Amazon.

Deleting your account

The following sections outlines the different ways you can terminate your Amazon Quick Suite instance.

To terminate your Amazon Quick Suite account without the Amazon Quick Suite UI

- Sign in to Amazon where you want to remove Amazon Quick Suite.

To terminate your account by using the Amazon Quick Suite UI

1. Choose your profile on the application bar, and then choose **Manage Quick Suite**.
2. Use one of the following methods to open the **Account termination** screen.
 - Use this [direct link](#) to the screen.
 - Choose **Account settings, Manage**.
3. On the **Account termination** page, confirm that you are viewing the correct Amazon Quick Suite account by checking the name listed for account name.
4. Toggle off **Account termination protection is on**. Doing this enables the **Delete account** section.
5. For **Type "confirm" to delete this account**, enter the word confirmation word shown on your screen.

Amazon Quick Suite asset management

Use this section to manage all of the assets in your Amazon Quick Suite account in one unified view. You can manage the following Quick Suite assets: Analyses, Dashboards, Datasets, Data sources, Shared folders, and Agents.

Here are some common reasons to use the asset manager:

- **Transfer assets** – Quickly transfer assets from one user or group to another, for example when the original owner is no longer present.
- **Onboard new employees** – Speed up onboarding of new employees by giving them access to the same assets that their teammates using.
- **Support authors** – Better support authors in tenancies by giving support engineers temporary access to the author's dashboard.
- **Revoke access** – Quickly audit and revoke permissions, for example after implementations, customer support, or unexpected events.

To manage Amazon Quick Suite assets

1. Choose the profile icon, then **Manage Quick Suite**.
2. Open the asset manager by clicking on **Manage assets**.
3. You can search for assets by name, or browse for them in a list. Choose one of the following methods:

To search by name, select the appropriate search bar, using the name as your guide. Enter your search term and press **ENTER**. Find assets a user or group has access to by using the **Search by User or Group name** bar. Find other assets by using the **Search by asset name** bar.

To browse for assets by type, select a button by its name to view a type of asset, for example: browse analyses by selecting the **Analyses** button, browse for Data sources by selecting the **Data sources** button, and so on.

4. When you are viewing a list of your search results, you can interact with the assets listed. Here are a few examples:
 - Select an asset by toggling the box at the beginning of each row. Or, select everything by clicking the box at the top left of the list.

- Change the type you are browsing for by selecting a different asset type from the **Asset type** list.
- Use the vertical dot menu at right to perform an action on the asset in that row.
- Use the **Share** button to share all of the selected assets. A popup window displays sharing options to apply to the users or groups you specify.
- Use the **Transfer** button to transfer all of the selected assets from one user or group to another. A popup window displays transferring options to apply to the users or groups you specify.

When you need to share a Amazon Quick Suite asset with 100 or more users, consider using Amazon Quick Suite groups. For more information on Amazon Quick Suite groups, see [Creating and managing groups in Amazon Quick Suite](#).

Manage Amazon Q permissions

Amazon Q is an AI-powered capability in Amazon Quick Suite that enhances your data analysis experience. Quick Suite admins can manage Amazon Q permissions in the Quick Suite administration console. These permissions control how Amazon Q interacts with your data and users.

Topics

- [Manage personalization permissions](#)
- [Manage dashboard and visual indexing for search](#)
- [Manage Dashboard Q&A](#)

Manage personalization permissions

Amazon Q can use user metadata to provide more context-aware responses. This feature allows for a more personalized experience when interacting with Amazon Q.

To manage personalization permissions:

1. Open the [Quick Suite console](#).
2. Choose the user icon at the top right, and then choose **Manage Quick Suite**.
3. Under the **Account** section, choose **Amazon Q**.

4. Locate the **Manage personalization permissions** section.
5. Use the toggle switch next to **Personalized responses** to enable or disable this feature.

When enabled, Amazon Q can read user metadata to tailor its responses to individual users. This may include considering the user's role, location, or prior documents when providing information or suggestions.

Manage dashboard and visual indexing for search

Amazon Q can index information from dashboards to make them easily searchable across applications. This feature enhances the discoverability of your dashboards and their contents.

To manage dashboard and visual indexing:

1. Open the [Quick Suite console](#).
2. Choose the user icon at the top right, and then choose **Manage Amazon Quick Suite**.
3. Under the **Account** section, choose **Amazon Q**.
4. Locate the **Manage dashboard and visual indexing for search** section.
5. Use the toggle switch next to **Dashboard & visual indexing** to enable or disable this feature.

When enabled, Amazon Q can read dashboard metadata and visual information such as keywords, data fields, and refresh times to quickly respond to queries.

Manage Dashboard Q&A

The Dashboard Q&A feature allows Amazon Q to answer questions about your data without requiring the use of Quick Suite Topics. This provides a more flexible and intuitive way for users to interact with their dashboards.

To manage Dashboard Q&A:

1. Open the [Quick Suite console](#).
2. Choose the user icon at the top right, and then choose **Manage Amazon Quick Suite**.
3. Under the **Account** section, choose **Amazon Q**.
4. Locate the **Manage Dashboard Q&A** section.
5. Use the toggle switch next to **Dashboard Q&A** to enable or disable this feature.

When enabled, Amazon Q can read dashboard metadata and visual information such as keywords, data fields, and refresh times to quickly respond to queries. This will eliminate the need for predefined Quick Suite Topics.

Managing Quick Suite subscriptions

Subscription management in Amazon Quick Suite involves configuring and optimizing your service usage to meet your organization's needs while controlling costs. This section covers essential subscription-related tasks including upgrading or downgrading your Quick Suite edition, managing SPICE memory capacity to ensure optimal performance for your in-memory datasets, and configuring data capacity limits for unstructured data storage. Proper subscription management helps you balance performance, functionality, and cost while ensuring your users have the resources they need to create and share insights effectively.

Topics

- [Configure Amazon Quick Suite subscriptions](#)
- [Configure SPICE memory capacity](#)
- [Index capacity](#)

Configure Amazon Quick Suite subscriptions

You can purchase standard user subscriptions to get discounted pricing on Amazon Quick Suite. When you invite additional users to Quick Suite, you're charged for those accounts on a month-by-month basis. If you have Enterprise edition, you have the option to take advantage of pay-per-session pricing for reader accounts. These are users who only view data dashboards, and don't need author or admin access.

To understand how Amazon Quick Suite subscription names on the pricing page map to user roles in the admin console, see [Understanding Amazon Quick Suite subscriptions and roles](#).

When you purchase an annual subscription, you pay for a Quick Suite user account on an annual rather than monthly basis. With an annual subscription, you receive a discounted price in return for the extended time commitment. You don't need to purchase an annual subscription to create or add users.

When you purchase a set of standard user subscriptions, you choose the number of accounts you want to cover. You also choose when to start the subscriptions (any time from the month following

the current month, to one year in the future) and whether to autorenew them. All subscriptions that you purchase together must use the same values for these settings.

You can edit an existing set of user subscriptions to change whether it autorenews. If the set is not yet active, you can also change the number of subscriptions it covers, or delete it entirely.

Topics

- [Viewing current subscriptions](#)
- [Purchase subscriptions](#)
- [Editing a subscription](#)
- [Delete a subscription](#)
- [Upgrading your Amazon Quick Suite subscription from Standard edition to Enterprise edition](#)

Viewing current subscriptions

Use the following procedure to view your current user subscriptions.

To view your current user subscriptions

1. Choose your user name on the application bar and then choose **Manage Quick Suite**.
2. Choose **Manage pricing**.
3. Use the subscription meter to see how many accounts you have and how they are billed. In the following example, the account has 21 users total:
 - Seven users with annual subscriptions. Only currently active subscriptions are shown here.
 - 13 month-to-month users.

Pause over any section of the meter bar to display details about that user segment.

4. Use the information in the subscriptions table to see what current and future subscriptions you have.

Purchase subscriptions

Use the following procedure to purchase subscriptions.

To purchase subscriptions

1. Choose your user name on the application bar and then choose **Manage Quick Suite**.
2. Choose **Manage pricing**.
3. Navigate to the **Authors and Admins** section, and then choose **Purchase plan**.
4. Choose or enter the number of subscriptions you want.
5. Choose the month and year when the subscriptions will start.
6. Choose whether the subscriptions autorenew.
7. Choose **Purchase**.

Editing a subscription

Use the following procedure to edit subscriptions.

To edit subscriptions

1. Choose your user name on the application bar and then choose **Manage Quick Suite**.
2. Choose **Manage pricing**.
3. Next to the set of subscriptions you want to change, choose **Manage**.
4. (Optional) If the subscriptions haven't started yet, change the number of subscriptions.
5. Choose whether the subscriptions autorenew.
6. Choose **Save changes**.

Delete a subscription

Warning

Deleting Amazon Quick Suite subscriptions affects user access. When you delete subscriptions:

- Users covered by deleted subscriptions may lose access to Quick Suite features and content
- Subscription changes affect your entire organization's access to Quick Suite capabilities
- Deleted subscriptions cannot be recovered - you must purchase new subscriptions to restore access

Before proceeding: Verify that affected users have alternative access methods and understand the implications of subscription changes.

Use the following procedure to delete subscriptions. You can only delete subscriptions that haven't started yet.

To delete subscriptions

1. Choose your user name on the application bar and then choose **Manage Quick Suite**.
2. Choose **Manage pricing**.
3. Next to the set of subscriptions that you want to delete, choose **Edit**.
4. Choose **Delete Subscription**.

Note

If you use Amazon Key Management Service or Amazon Secrets Manager with Amazon Quick Suite, you are billed for access and maintenance as described in the pricing pages for each Amazon product. For more information on how these products are billed, see the following:

- [Amazon Key Management Service Pricing page](#)
- [Amazon Secrets Manager Pricing page](#)

In your billing statement, the costs are itemized under the appropriate product and not under Amazon Quick Suite.

Upgrading your Amazon Quick Suite subscription from Standard edition to Enterprise edition

You can upgrade from Amazon Quick Suite Standard edition to Amazon Quick Suite Enterprise edition. In Enterprise edition, Amazon Quick Suite supports the following additional features:

- Reader role with pay-per-session pricing; for more pricing details, see following.

- Email reports for offline delivery of insights.
- Larger SPICE datasets with up to 500 million rows per SPICE dataset.
- Hourly refresh of SPICE data (using the Amazon Quick Suite console).
- **ML Insights** to make the most of your data, including the following:
 - Anomaly detection that can run on billions of rows of data on a schedule.
 - Contribution analysis to help you figure out key drivers.
 - One-click forecasting.
 - Customizable natural language narratives that you can use to add business context to a dashboard.
 - SageMaker AI integration.
- **Embedded analytics** in applications and portals:
 - Embed dashboards with row level security.
 - Namespaces with multitenant support for creating dashboards with embedded analytics.
 - Templates for repeatable dashboard creation and management.
 - Capacity pricing for embedding.
- **Security and governance**
 - Row-level security.
 - Private virtual private cloud (VPC) support based on Amazon VPC.
 - Folders for organization and sharing.
 - Fine-grained access control over Amazon S3, Amazon Athena, and other Amazon services and resources.
 - Amazon Lake Formation support.
- **User authentication and management options**
 - Integration with Microsoft Active Directory with support for Active Directory groups.
 - Group support for user management.

To see a full comparison of Standard edition with Enterprise edition, see [Amazon Quick Suite editions](#).

When you upgrade your account, your administrators and authors are billed at the Amazon Quick Suite Enterprise edition rates. For pay-per-session pricing, you can add additional users as readers.

Before you reprovision existing users as readers, you transfer or delete their resources, and then delete the users from your subscription.

Users who are in the reader role can view and manipulate shared dashboards, and receive emailed updates. However, readers can't add or change data sources, datasets, analyses, visuals, or administrative settings. Billing for readers is significantly lower in cost than regular user pricing. It's based on 30-minute sessions, and it's capped at a maximum amount per month for each reader. Billing for upgrades is prorated for the month of the upgrade. Upgrades to users are also prorated. If you have an annual subscription to Standard edition, it's converted to Enterprise edition and stays in place for the remaining term.

Warning

Downgrading from Enterprise edition to Standard edition isn't currently possible due to the enhanced feature set available in Enterprise edition. To perform this downgrade, unsubscribe from Amazon Quick Suite, and then start a new subscription. Also, you can't transfer users or assets between subscriptions.

Upgrading to Enterprise edition to use Active Directory connectivity isn't supported. This is because of the differences in the user identity mechanisms between Amazon Quick Suite password-based users and existing Active Directory users. However, you can upgrade to Enterprise and still use password-based users. If you want to upgrade and change how users sign in, you can unsubscribe and start a new subscription.

Use the following procedure to upgrade to Enterprise edition. To perform the upgrade, you need administrative access to Amazon Quick Suite, with security permissions to subscribe. The person performing the upgrade is usually an Amazon administrator who is also an Amazon Quick Suite administrator.

To upgrade to enterprise edition

1. Open the administrative settings page by clicking on your profile icon at top right.
2. At top left, choose **Upgrade now**.
3. Be sure that you want to upgrade.

⚠ Important

You can't undo this action.

Choose **Upgrade** to upgrade. The upgrade is instantaneous.

Billing for the upgrade to your subscription is prorated for the month of upgrade. Upgrades to Amazon Quick Suite users are also prorated.

4. (Optional) Downgrade users to readers:

- Before you start, make sure to transfer any assets your users own that you want to keep.
- Delete the users and add them back to your subscription as readers.

If you're using Active Directory, delete the authors, move them to the new reader group, then recreate them as readers in Amazon Quick Suite.

When you upgrade to Enterprise edition, your admin and author users retain their roles.

Configure SPICE memory capacity

SPICE (Super-fast, Parallel, In-memory Calculation Engine) is the robust in-memory engine that Amazon Quick Suite uses. It's engineered to rapidly perform advanced calculations and serve data. In Enterprise edition, data stored in SPICE is encrypted at rest. For more information, see [Data encryption in Quick Suite](#).

SPICE capacity is allocated separately per Amazon Web Services Region. For each Amazon Web Services account, SPICE capacity is shared by all the people using Quick Suite in a single Amazon Web Services Region. The other Amazon Web Services Regions have no SPICE capacity unless you choose to purchase some.

Quick Suite administrators can view how much [SPICE](#) capacity you have in each Amazon Web Services Region and how much of it is currently in use. Administrators can also purchase additional SPICE capacity or release unused SPICE capacity. You can only release SPICE capacity that isn't currently used by a dataset. Datasets in SPICE stay there until someone remove them from SPICE. To change that, you can either delete the datasets or change them so they aren't stored in SPICE.

Purchasing or releasing SPICE capacity only affects the capacity for the currently selected Amazon Web Services Region. Each Amazon Web Services account can have a separate Amazon Quick Suite subscription and can be used in multiple Amazon Web Services Regions.

Before you make any changes to SPICE capacity, make sure that you're using the correct Amazon Web Services account and Amazon Web Services Region. It's possible to be using different Amazon Web Services accounts or Amazon Web Services Regions at the same time in different contexts, as follows:

- If you open Amazon Quick Suite using the `http://quicksight.aws.amazon.com` URL, Amazon Quick Suite automatically selects your account and Amazon Web Services Region. You can't view your Amazon Web Services account from Amazon Quick Suite. We recommend using a different method to open Amazon Quick Suite when you want to work with SPICE capacity.
- If you open Amazon Quick Suite from the Amazon Web Services Management Console, Amazon Quick Suite opens in the account that you used to sign in to that console. However, it opens in the last Amazon Web Services Region that you selected in Amazon Quick Suite. The Amazon Web Services Management Console and the Amazon Quick Suite console each have an Amazon Web Services Region selector that works independently from the other. Changing the selected Amazon Web Services Region in the Amazon console doesn't change the Amazon Web Services Region in Amazon Quick Suite.
- If you use the Amazon Command Line Interface (Amazon CLI) to run Amazon Quick Suite commands, make sure to provide the relevant Amazon Web Services account for each Amazon Quick Suite API operation you use. The Amazon Web Services Region isn't always required, and if you don't provide it, the Amazon CLI uses your default Amazon Web Services Region from your Amazon configuration. We recommend that you always explicitly provide the Amazon Web Services Region, to make sure you apply the command to the correct Amazon Web Services Region.

You must be signed in as a Amazon Quick Suite administrator to view or manage SPICE capacity.

Topics

- [Finding your current Amazon account and Amazon Web Services Region](#)
- [Viewing SPICE capacity and usage in an Amazon Web Services Region](#)
- [Hiding SPICE capacity labels](#)
- [Purchasing SPICE capacity in an Amazon Web Services Region](#)
- [Turning on SPICE auto capacity purchasing](#)

- [Releasing SPICE capacity in an Amazon Web Services Region](#)

Finding your current Amazon account and Amazon Web Services Region

To select the correct Amazon account and Amazon Web Services Region (console)

1. Open the Amazon console, using the Amazon account that you want to view SPICE information for. If you have only one Amazon Web Services account, you can skip this step.

You can verify the account number by following these steps:

- a. On the navigation bar at the top of the page, choose the account name or number at right. If a number displays, this might be your Amazon Web Services account ID.
- b. Choose **My Security Credentials** to display your credential-related information and options. Your Amazon Web Services account ID displays near the top of the page.

To return to the original page, choose the Amazon logo at upper left.

2. Open Amazon Quick Suite by first entering "**quicksight**" into the **Find Services** search box. When the word Amazon Quick Suite appears following the search box, choose it from the list.
3. In Amazon Quick Suite, open the profile menu by choosing your profile icon at top right. The Amazon name of the Amazon Web Services Region that Amazon Quick Suite is using displays in the menu.

The same Amazon Web Services Region also displays in the URL, for example: `https://us-east-1.quicksight.aws.amazon.com/sn/admin`. If this is your URL, the profile menu displays the name N. Virginia.

To switch Amazon Web Services Regions, display the list of supported Regions by choosing the Region name from the profile menu. Then choose the Region that you want to use. Switching to a different Amazon Web Services Region changes the SPICE usage information that you can view. It also changes the Amazon Quick Suite assets that you can use, for example data sources and dashboards.

Viewing SPICE capacity and usage in an Amazon Web Services Region

To view current SPICE capacity and usage (console)

1. Open Amazon Quick Suite. Make sure that you're using the correct Amazon Web Services account and Amazon Web Services Region as described previously in [Finding your current Amazon account and Amazon Region](#).
2. Open the administration page by choosing **Manage Quick Suite** from your profile menu.
3. Choose **SPICE capacity** from the left navigation pane . The following information displays:
 - The **Total SPICE capacity** section displays the total amount of used and unused SPICE capacity. A bar graph shows how much of this storage space is in each of the following categories for this Amazon Web Services account in the Amazon Web Services Region that's currently selected in Amazon Quick Suite:
 - Purchased SPICE capacity – This is the additional SPICE capacity.
 - SPICE capacity bundled with Amazon Quick Suite – This is the total default capacity associated with your paid users.

Hover over any section of the meter to see details on that capacity type.

- The **SPICE usage** section displays the total amount of the used and unused SPICE capacity. A bar graph shows how much of this storage space is in each of the following categories for this Amazon Web Services account in the Amazon Web Services Region that's currently selected in Amazon Quick Suite:
 - Used SPICE capacity – This is the used portion of the default SPICE capacity that you get for each user.
 - Unused SPICE capacity – This is the unused portion of the default SPICE capacity that you get for each user.
 - Releasable unused capacity – This is the purchased capacity that isn't in use, and so can be released to reduce costs.

Hiding SPICE capacity labels

Amazon Quick Suite account admins can choose to hide the account-wide SPICE capacity usage and remaining size labels from Amazon Quick Suite authors. This feature is available to all enterprise accounts that use custom permissions. For more information about custom permissions in Amazon Quick Suite, see [Customizing access to Amazon Quick Suite capabilities](#)

Use the following procedure to hide SPICE capacity usage from the Amazon Quick Suite console.

1. Open the [Quick Suite console](#).
2. From any page in the Amazon Quick Suite console, choose your profile name, and then choose **Manage Quick Suite**.

The **Manage Quick Suite** menu is only available to Amazon Quick Suite admins. If you are unable to access this menu, contact your Amazon Quick Suite account admin for assistance.

3. Choose **Manage users**, and then choose **Manage permissions**.
4. Edit or create a new custom permission. For **Restrict access to**, under **Datasets**, select **Viewing account SPICE capacity**.
5. When you are finished creating or changing the custom permission, choose **Create** or **Update**.

After you create or update a custom permission to hide SPICE capacity usage, assign the new permission to users with the `UpdateUser` API.

Purchasing SPICE capacity in an Amazon Web Services Region

To purchase more SPICE capacity (console)

1. Open Amazon Quick Suite. Make sure that you're using the correct Amazon Web Services account and Amazon Web Services Region as described previously in [Finding your current Amazon account and Amazon Region](#).
2. Open the administration page by choosing **Manage Quick Suite** from your profile menu.
3. Choose **SPICE capacity** from the left navigation pane .
4. Choose the **Purchase more capacity** button.
5. Enter a number of gigabytes of SPICE capacity to purchase for the Amazon Web Services Region that is currently selected in Amazon Quick Suite.
6. To confirm your choice, choose **Purchase SPICE capacity**. To exit without making any changes, choose **Cancel**.

Turning on SPICE auto capacity purchasing

Turn on SPICE auto capacity purchasing to allow Amazon Quick Suite to automatically manage your Amazon Quick Suite account's SPICE capacity. When you turn auto capacity purchasing

on, Amazon Quick Suite evaluates how much capacity is needed based on your account's usage. As your account uses more SPICE storage, Amazon Quick Suite automatically purchases SPICE capacity as needed on your behalf. This allows users to ingest data as needed without the need to estimate usage or manually purchase SPICE data. Auto capacity purchasing makes it easier for new customers, ISVs, and larger companies to directly access SPICE without needing to understand, track, or manually purchase their account's SPICE capacity. Amazon Quick Suite admins can still purchase and release SPICE capacity manually.

Auto capacity purchasing doesn't support auto-decrement. If users want to reduce their SPICE usage, capacity release must be done manually.

By default, all new Amazon Quick Suite accounts that are created in the Amazon Quick Suite console have auto capacity purchasing turned on in the region that their capacity is located. To turn on auto capacity purchasing for other regions, Amazon Quick Suite account admins can manually turn on auto capacity from the **SPICE capacity** management page.

By default, all new Amazon Quick Suite accounts that were created with the Amazon Quick Suite API and all existing Amazon Quick Suite accounts have auto capacity purchasing turned off. To turn on auto capacity purchasing, Amazon Quick Suite account admins can manually turn on auto capacity from the **SPICE capacity** management page.

To turn SPICE capacity purchasing on or off

1. Open the [Quick Suite console](#).
2. From any page in the Amazon Quick Suite console, choose your profile name, and then choose **Manage Quick Suite**.

The **Manage Quick Suite** menu is only available to Amazon Quick Suite admins. If you are unable to access this menu, contact your Amazon Quick Suite account admin for assistance.

3. Choose **SPICE capacity**.
4. On the **SPICE Capacity** page that opens, toggle the **Auto-purchase capacity** on.

To turn auto capacity purchasing off, follow the procedure above and toggle **Auto-purchase capacity** off. When auto purchase capacity is turned off, ingestions or refreshes that exceed the account's SPICE capacity automatically fail.

Amazon Quick Suite admins can turn auto capacity pricing on or off at any time. If you turn auto capacity purchasing off after it's been in use, your account's current capacity becomes your

account's purchased capacity. If your account has no remaining capacity when you turn auto purchase off, the next ingestion or refresh will fail.

If your account already exceeds its SPICE capacity when you turn auto capacity purchasing on, Amazon Quick Suite automatically matches your account's capacity to your current usage. After Amazon Quick Suite matches your account's capacity, the auto-purchase logic starts.

Releasing SPICE capacity in an Amazon Web Services Region

To release unused SPICE capacity (console)

1. Open Amazon Quick Suite. Make sure that you're using the correct Amazon Web Services account and Amazon Web Services Region as described previously in [Finding your current Amazon account and Amazon Region](#).
2. Open the administration page by choosing **Manage Quick Suite** from your profile menu.
3. Choose **SPICE capacity** from the left navigation pane .
4. Choose **Release unused purchased capacity**.
5. Do one of the following:
 - To release all SPICE capacity from the Amazon Web Services Region that is currently selected in Amazon Quick Suite, choose **Release all**.
 - To release some of gigabytes of SPICE capacity from the Amazon Web Services Region that is currently selected in Amazon Quick Suite, enter the number of gigabytes to release.
6. To confirm your choice, choose **Release SPICE capacity**. To exit without making any changes, choose **Cancel**.

Index capacity

An index is a structured corpus of documents that allows for efficient searching, sorting, and data access. Amazon Quick Suite uses an index to store, prepare, sync any files you upload to your Quick Suite spaces. It then uses this indexed data to generate agent-driven responses during chat conversations.

You must purchase data storage capacity for your Quick Suite index. Quick Suite automatically creates an index when you sign up in your home region, and auto-provisions 50 MB default index capacity with auto-scaling enabled (a mechanism where your index capacity is adjusted

automatically based on usage). In any other region, an admin has to explicitly provision index capacity. The minimum index capacity you can purchase is 25 MB. The maximum is 10,000 MB.

The Amazon Quick Suite admin console displays both your purchased index capacity and your index capacity usage. You can switch between manual and auto-scaled index capacity modes as per your needs. When you are buying capacity manually, you can also release unused capacity. If you want to release all provisioned index capacity, you could do so after deleting corresponding data—like file uploads in spaces and knowledge bases—in your Amazon Quick Suite instance.

Topics

- [Total index capacity](#)
- [To view index capacity usage](#)
- [Release unused index capacity](#)
- [Switch between capacity modes](#)

Total index capacity

You must purchase data storage capacity for your Amazon Quick Suite index to ingest data from knowledge bases and upload files to spaces. In the index capacity page, the first horizontal bar shows your total index capacity. This is the current capacity available in your index to support file uploads and knowledge bases. When you first signup, the capacity shown here is the auto-provisioned 50 MB in the home region. When auto-scaling mode is **ON**, this capacity automatically adjusts to reflect the latest available capacity based on your ingestion needs. When manual mode is **ON**, this total capacity represents the actual capacity you have purchased for your account based on your estimate of the ingestion needs.

Note

With auto-scaling enabled, Amazon Quick Suite may temporarily provision capacity above your maximum purchase limit during periods of high concurrent usage to maintain performance. However, you are only charged for capacity up to your configured maximum purchase limit, not for any temporary over-provisioning that occurs.

To manage purchase options

1. Log in to the Amazon Quick Suite console.

2. Select **Manage Quick Suite**.
3. From the admin console left navigation menu, select **Subscriptions**, and then select **Index capacity**.
4. In **Index capacity**, from **Total index capacity** select **Manage capacity**, and do the following:
 - a. In **Manage purchase options**, do one of the following:
 - i. To automatically scale index capacity as per usage, select **Auto-scale (Recommended)**. This automatically adjusts index capacity in 25 MB increments based on your usage, purchasing or releasing as needed. Then, in **Maximum purchase limit** set the limit at which auto-purchase stops. The value of the maximum purchase limit must be 10,000. Maximum purchase limit cannot be lower than current used capacity.
 - ii. To manually purchase capacity, select **Manual**. Then, in **Index capacity**, set the total capacity you want your index to hold. The capacity must be between 25-10,000 MB and you can increase it in 25 MB increments. When you want to lower the current capacity held by the index, it can't be lowered below the used capacity. Since capacity is managed in 25 MB increments, you can only decrease to the next 25 MB unit above your current usage. For example, if your index has 200 MB capacity with 110 MB consumed, you can reduce capacity to 125 MB but not to 100 MB since part of the 100-125 MB unit is already in use.
 - b. Select **Confirm**.

To view index capacity usage

The Amazon Quick Suite admin console displays both your purchased index capacity and your index capacity usage. The following procedure shows you how to view them.

To view index data capacity usage

1. Log in to the Amazon Quick Suite console.
2. Select **Manage Quick Suite**.
3. From the left navigation menu, select **Subscriptions**, and then select **Index capacity**.
4. In **Index capacity**, you'll find your capacity usage under **Index capacity usage**. First bar shows you total purchased capacity and second bar shows you index capacity usage.

Release unused index capacity

The Amazon Quick Suite admin console displays both your purchased index capacity and your index capacity usage. If you've manually provisioned capacity, you can choose to release unused index capacity. If you've enabled auto-scaling for your index, you don't need to release index capacity as Amazon Quick Suite automatically scales index capacity to your usage. The following procedure shows you how to release unused index capacity.

Note

To release any provisioned index capacity, you must delete all datasets and files uploaded to spaces. Then, follow the steps to release unused index capacity.

To release index data capacity

1. Log in to the Amazon Quick Suite console.
2. Select **Manage Quick Suite**.
3. From the left navigation menu, select **Subscriptions**, and then select **Index capacity**.
4. In **Index capacity**, select **Manage capacity**.
5. In **Manage index capacity**, reduce the **Index capacity** to your desired value.

Switch between capacity modes

You can switch between manual and auto-scale index capacity modes at any time. When switching from auto-scale to manual, your current capacity becomes your fixed manual capacity. When switching from manual to auto-scale, you set a maximum purchase limit and Amazon Quick Suite automatically adjusts capacity based on usage.

To switch from auto-scale to manual capacity

1. Log in to the Amazon Quick Suite console.
2. Select **Manage Quick Suite**.
3. From the admin console left navigation menu, select **Subscriptions**, and then select **Index capacity**.
4. In **Index capacity**, from **Total index capacity** select **Manage capacity**.

5. In **Manage index capacity**, select **Manual**.
6. In **Index capacity**, set your desired manual capacity. The value must be between 25-10,000 MB in 25 MB increments and can't be less than current usage.
7. Select **Confirm**.

To switch from manual to auto-scale capacity

1. Log in to the Amazon Quick Suite console.
2. Select **Manage Quick Suite**.
3. From the admin console left navigation menu, select **Subscriptions**, and then select **Index capacity**.
4. In **Index capacity**, from **Total index capacity** select **Manage capacity**.
5. In **Manage index capacity**, select **Auto-scale (Recommended)**.
6. In **Maximum purchase limit**, set your maximum purchase limit. The value must be between 25-10,000 MB in 25 MB increments and can't be less than current usage.
7. Select **Confirm**.

Managing identities

Identity management in Amazon Quick Suite involves controlling who can access your Quick Suite resources and what actions they can perform. This section covers the essential tasks for managing user identities, including adding and removing users, organizing them into groups for easier permission management, and transitioning between different identity management systems. Whether you're using IAM Identity Center, federated identity providers, or other authentication methods, these topics will help you establish proper access controls and maintain security for your Quick Suite environment.

Before managing user access, it's important to understand how Amazon Quick Suite subscription names map to user roles. For information about subscription types and their capabilities, see [Understanding Amazon Quick Suite subscriptions and roles](#).

Topics

- [Managing user access inside Amazon Quick Suite](#)
- [Creating and managing groups in Amazon Quick Suite](#)

Managing user access inside Amazon Quick Suite

Intended audience: System administrators and Amazon Quick Suite administrators

Amazon Quick Suite administrators manage user access in Quick Suite. User access management in Quick Suite is determined by your Quick Suite account identity configuration. For accounts that use IAM Identity Center or Active Directory, groups are assigned to Quick Suite roles. Groups can be assigned the Admin, Author, Reader, Admin Pro, Author Pro, or Reader Pro roles. For more information about Pro roles in Quick Suite see [Get started with Generative BI](#). Note that Reader Pro maps to Amazon Quick Suite Professional subscription and Author Pro maps to Amazon Quick Suite Enterprise subscription. For more information about integrating your Amazon Quick Suite account with IAM Identity Center, see [Managing access for IAM Identity Center users](#).

To understand how subscription names relate to user roles and their capabilities, see [Understanding Amazon Quick Suite subscriptions and roles](#).

Amazon Quick Suite accounts that use Quick Suite and IAM users create users directly in Quick Suite. These users and their roles are managed at the user level. For more details, see [Managing access for Quick Suite and IAM users](#).

Topics

- [Managing access for IAM Identity Center users](#)
- [Managing access for Amazon Quick Suite and IAM users](#)

Managing access for IAM Identity Center users

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick Suite administrators

Amazon administrators can use this topic to learn more about managing accounts that are integrated with IAM Identity Center. The information in this section also applies to Quick Suite accounts that use Active Directory.

To manage Quick Suite users, you must have administrative privileges in Quick Suite and also the appropriate Amazon permissions. For more information about the necessary Amazon permissions, see [IAM policy examples for Quick Suite](#). If you are using directory groups, you need to be a network administrator.

Each Quick Suite Enterprise edition account can have an unlimited number of users. User names that contain a semicolon (;) aren't supported.

Use the following topics to add, view, and deactivate Quick Suite users.

Important

You can't remap Amazon Quick Suite users or groups from one identity store to another. For example, if you are migrating from an on-premises Active Directory to Amazon Directory Service, or the other way around, you unsubscribe and resubscribe to Amazon Quick Suite. You do this because even if the user's aliases remain the same, the underlying identity data changes. To make the transition easier, request in advance that your users document all their Amazon Quick Suite assets and settings before the migration.

Topics

- [Adding users](#)
- [Managing user access](#)
- [Deactivating user accounts](#)
- [Changing a user's role](#)
- [Deleting Enterprise accounts](#)

Adding users

With IAM Identity Center, add users to Amazon Quick Suite by associating their IAM Identity Center group to an Admin, Admin Pro, Author, Author Pro, Reader, or Reader Pro role in Quick Suite. All users in the selected groups are authorized to sign in to Quick Suite.

For more information about Pro roles in Quick Suite see [Get started with Generative BI](#).

To see which groups are integrated with your Quick Suite account, follow the procedure in [Managing user access](#).

Managing user access

Use the following procedure to view groups that are assigned to a role that grants access to Quick Suite.

1. Open the [Quick Suite console](#).
2. Choose **Manage Quick Suite**, and then choose **Manage Users**.
3. Choose **Manage role groups**.
4. In the **Manage role groups** page, use the tables to add or remove groups in IAM Identity Center or Active Directory from the Admin, User, or Reader roles in Quick Suite.

Deactivating user accounts

Deactivating a Amazon Quick Suite group or user account removes that group or user's access to Quick Suite resources, like analyses or data sets. IAM Identity Center or Active Directory users that are removed from a group that grants them access to Quick Suite lose access to Quick Suite. These users appear in the **Inactive users** list in Quick Suite until the first day of the following month. After that, the deactivated users are automatically removed from the **Inactive users** list. Before you deactivate a user, you can reassign their resources to another user with the asset management console.

If you later need to reactivate a Quick Suite user's account, put the user into a group with access to Quick Suite. Doing this restores their access to Quick Suite and to any existing resources that are still associated with that user.

Note

With IAM Identity Center integrated into your Amazon Quick Suite account or Active Directory users, you can change a user's role type by moving them to a group that is associated with a different Amazon Quick Suite role. If a user is in multiple groups that are mapped to different Amazon Quick Suite role types, the user is able to access Amazon Quick Suite with the role that offers the broadest level of access. Accounts that use other identity types can't upgrade or downgrade a user by transferring them between groups. For more information, see [Changing a user's role](#).

You can activate or deactivate multiple users at once by adding or removing one or more IAM Identity Center or Active Directory groups that are associated with a role in Amazon Quick Suite.

Changing a user's role

If you're using IAM Identity Center or Active Directory, you can change a user's role by adding or removing them from a group that's mapped to the role that you want to assign them in Quick Suite. You can also perform this task by adding a new group to a role in Quick Suite. To do this, you need both administrative privileges in Quick Suite and also appropriate Amazon permissions.

With IAM Identity Center integrated users, you can change role types for a user by moving them to a group that is associated with a different Quick Suite role. If a user belongs to multiple groups that are mapped to different role types, the user is able to access Quick Suite with the role that offers the broadest level of access.

When you make changes to users or groups in Quick Suite, it can take up to five minutes for the change to take effect. Examples of such changes are the following:

- Deleting a user
- Changing a user from an admin to an author
- Adding or removing group members

The five-minute time period allows changes to propagate throughout the system.

Deleting Enterprise accounts

If a user is deleted from IAM Identity Center or Active Directory or is removed from a group that's associated with a role in Quick Suite, the user no longer exists in Quick Suite. You do not need to delete the user in the Quick Suite application. The deleted user will appear in the **Inactive users** list in Quick Suite until the first day of the following month. After that date passes, the user is automatically removed from the list.

Managing access for Amazon Quick Suite and IAM users

Amazon Quick Suite account administrators can use this topic to learn more about managing accounts that use IAM or Quick Suite for identity federation.

To manage Quick Suite users, you must have administrative privileges in Quick Suite and also the appropriate Amazon permissions. For more information about the necessary Amazon permissions, see [IAM policy examples for Quick Suite](#).

Each Quick Suite Enterprise edition account can have an unlimited number of users. User names that contain a semicolon (;) aren't supported.

Use the topics below to learn more about managing access for Quick Suite and IAM users.

Topics

- [Inviting users to access Amazon Quick Suite](#)
- [Viewing Amazon Quick Suite account details](#)
- [Deleting a Amazon Quick Suite user account](#)

Inviting users to access Amazon Quick Suite

Applies to: Enterprise Edition and Standard Edition

Intended audience: Amazon Quick Suite administrators

Use the following procedure to invite a user to access Quick Suite.

1. Choose your user name on the application bar and then choose **Manage Quick Suite**.
2. Choose **Manage Users**. On this screen, you can manage users who already exist in your account.
3. Choose **Invite users**.
4. In the **Invite users to this account** table, enter a new user name for a person to whom you want to grant access to Quick Suite. If the user is an IAM user, enter their IAM credentials. Then press **+**. A user's IAM user name can be the same as their email address.

Repeat this step until you have entered everyone who you want to invite. Then go to the next step to enter details.

5. For **Email**, enter an email address for the account.

Note

Currently, email addresses are case-sensitive.

6. For **Role**, choose the role to assign to each person you're inviting. A *role* determines the permission level to grant to that account.

- **ADMIN roles:**

- **ADMIN** – The user is able to both use Amazon Quick Suite for authoring and for performing administrative tasks like managing users or purchasing [SPICE](#) capacity.
- **ADMIN PRO** – The user is able to perform all actions of a Amazon Quick Suite Admin and utilize applicable Amazon Quick Suite Generative BI capabilities. For more information about Pro roles in Amazon Quick Suite, see [Get started with Generative BI](#).

There are some differences in the administrative tasks that IAM users and Amazon Quick Suite administrators can perform. These differences occur because some administrative tasks require permissions in Amazon, which Amazon Quick Suite-only users lack. The differences are these:

- Amazon Quick Suite administrators can manage users, SPICE capacity, and subscriptions.
- IAM users with administrative permissions can also manage users, SPICE capacity, and subscriptions. In addition, they can manage Amazon Quick Suite permissions to Amazon resources, upgrade to Enterprise edition, and unsubscribe from Amazon Quick Suite.

If you want to create a user with administrator permissions with IAM access, check with your Amazon administrator. Make sure that the IAM user has the all necessary statements in their IAM permissions policy to work with Amazon Quick Suite resources. For more information about what statements are required, see [IAM policy examples for Amazon Quick Suite](#).

- **AUTHOR roles:**

- **AUTHOR**– The user is able to author analyses and dashboards in Amazon Quick Suite but not perform any administrative tasks in Amazon Quick Suite.
- **AUTHOR PRO**– The user is able to perform all actions of a Amazon Quick Suite Author and utilize applicable Amazon Quick Suite Generative BI capabilities. For more information about Pro roles in Amazon Quick Suite, see [Get started with Generative BI](#).

- **READER roles (Enterprise only):**

- **READER**– Users are able to interact with shared dashboards, but not author analyses or dashboards or perform any administrative tasks.
- **READER PRO**– The user is able to perform all actions of a Amazon Quick Suite Reader and utilize applicable Amazon Quick Suite Generative BI capabilities. For more information about Pro roles in Amazon Quick Suite, see [Get started with Generative BI](#).

7. For **IAM User**, verify that it says **Yes** for accounts that are associated with IAM users, and **No** for those that are Amazon Quick Suite-only.

8. (Optional) To delete a user, choose the delete icon at the end of the relevant row.
9. Choose **Invite**.

Viewing Amazon Quick Suite account details

Intended audience: Amazon Quick Suite administrators

You can view Amazon Quick Suite accounts on the **Manage Users** page. To view a Amazon Quick Suite user account, use the following procedure.

1. Choose your user name on the application bar and then choose **Manage Quick Suite**.
2. Choose **Manage Users** to view details about people who are Amazon Quick Suite users. The information that displays includes:
 - Username – The person's user name.
 - Email – The email associated with this user name.
 - Role – The security cohort that the person's user name belongs to: **ADMIN**, **ADMIN PRO**, **AUTHOR**, **AUTHOR PRO**, **READER**, or **READER PRO**.
 - Last active – The last date and time that this person accessed the Amazon Quick Suite console. Anyone who isn't an active user has a **Last active** status of User has no activity.

You can also see deleted or inactive users in this screen.

3. To find a user name, enter a part or all of a user's name or email the search box. Search is case-insensitive and wildcards aren't supported. To clear the search results and view all user names, delete your search entry.

Deleting a Amazon Quick Suite user account

Intended audience: Amazon Quick Suite administrators

 Warning

Deleting user accounts has permanent, organization-wide consequences. When you delete a user account:

- All user-owned resources are permanently removed unless explicitly transferred to another user before deletion
- Shared dashboards and analyses become inaccessible to other users if the deleted user was the owner
- Data sources and datasets owned by the user are deleted, potentially breaking dependent analyses across your organization
- **This action cannot be undone - deleted resources cannot be recovered**

Always transfer critical resources to another user before deleting an account. Review all user-owned assets using the asset management console before proceeding with any account deletion.

Accounts can be deleted by either an Amazon administrator or an Amazon Quick Suite administrator. Deleting a Amazon Quick Suite user account works the same in both the Standard and Enterprise editions of Amazon Quick Suite.

Deleting a Amazon Quick Suite user account removes or transfers their resources. In Enterprise edition, the network administrator can temporarily deactivate a Amazon Quick Suite user account by removing it from the network group that has access to Amazon Quick Suite. If a user is deleted, but not deactivated, that user can still access Amazon Quick Suite as a new user. For more information about deactivating an Enterprise account, see [Deactivating user accounts](#).

Use the following procedure to delete a Amazon Quick Suite user account.

1. Choose your user name on the application bar and then choose **Manage Quick Suite**.
2. Choose **Manage Users**.
3. Locate the account you want to delete and then choose the delete icon at the end of that row.
4. Choose to either delete or transfer any resources owned by the user and then choose **OK**.
5. Do one of the following:

- If you chose to transfer user resources, enter the user name of the account to transfer them to and then choose **Delete and transfer resources**.
- If you chose to delete user resources, choose **Delete**. You can't undo this action.

Creating and managing groups in Amazon Quick Suite

Intended audience: System administrators

Applies to: Enterprise Edition

Note

If your Amazon Quick Suite account is integrated with IAM Identity Center (recommended), groups are not managed in the Quick Suite application. Instead, groups are managed in IAM Identity Center or in the third-party identity provider that you configured in IAM Identity Center. Groups are synced automatically between Quick Suite and IAM Identity Center.

Admins with IAM credentials who have access to the Quick Suite console can organize sets of users into groups that make it easier to manage access and security. For example, you can create a group of users that you can share Amazon Quick Suite assets with all at once. You can create and manage groups using the Amazon Quick Suite console or the Amazon Command Line Interface (Amazon CLI). You can create up to 10,000 groups in a namespace. If you want to create more than 10,000 groups in a namespace, contact [Amazon Support](#).

Use the topics below to create, and modify groups with in the Amazon Quick Suite console or with the Amazon Quick Suite APIs.

Topics

- [Create a group with the Amazon Quick Suite console](#)
- [Change a group description with the Amazon Quick Suite console](#)
- [Manage group membership in the Amazon Quick Suite console](#)
- [Create and manage groups with the Amazon CLI](#)

Create a group with the Amazon Quick Suite console

Use the following procedure to create a group in the Quick Suite console.

To create a user group in the Amazon Quick Suite console:

1. On the Amazon Quick Suite start page, choose **Manage Quick Suite**, and then choose **Manage groups**.
2. Choose **NEW GROUP**.
3. On the **Create new group** page, enter the name and description of the new group in the corresponding boxes.
4. When you're finished, choose **Create** to create the new group.

Note

You can't delete a group from the Amazon Quick Suite console, but you can delete a group with the Amazon CLI. For more information on deleting a Amazon Quick Suite group with the Amazon CLI, see [Create and manage groups with the Amazon CLI](#).

Change a group description with the Amazon Quick Suite console

After you have created a new group, you can't change the group's title but you can change the group's description.

To change the description of a group:

1. On the Amazon Quick Suite start page, choose **Manage Quick Suite**, and then choose **Manage groups**.
2. Choose the group that you want to change, and then choose the **Edit** link next to the group description.
3. In the **Edit description** box that appears, enter the new description and choose **Save**.

Manage group membership in the Amazon Quick Suite console

After you create a group, you can add and remove users from the **Manage groups** page. You can't add a user to a group if you haven't added the user to your account. For more information on

adding users to your Amazon Quick Suite account, see [Managing user access inside Amazon Quick Suite](#).

To add a user to a group

1. On the Amazon Quick Suite start page, choose **Manage Quick Suite**, and then choose **Manage groups**.
2. Choose the group that you want to add a user to, and choose **ADD USER** at the page's upper right.
3. Enter the user name or email of the user that you want to add, and choose the correct user for **Search users**.

To remove a user from a group:

1. On the Amazon Quick Suite start page, choose **Manage Quick Suite**, and then choose **Manage groups**.
2. Choose the group that you want to remove a user from.
3. Find the user that you want to remove and choose **Remove**.

Choosing **remove** automatically removes the selected user from the group.

You can also search for a group member by entering the user's full user name into the search bar on the right-hand side of the group's page.

Create and manage groups with the Amazon CLI

Before you begin, make sure that you have the Amazon CLI installed. For more information, see [Installing the Amazon CLI](#) in the *Amazon CLI User Guide*.

Use the following procedure to create an Quick Suite user group.

1. Open a terminal window. If you are using Microsoft Windows, open a command prompt.
2. Enter the following command at the prompt to create a group. Substitute the correct values for your parameters.

```
aws quicksight create-group --aws-account-id=111122223333 --namespace=default --  
group-name="Sales-Management" --description="Sales Management - Forecasting"
```

You might find it easier to create the command in a text editor before entering it at the prompt. For more information on `create-group` and other available commands, see the [Amazon Quick Suite API reference](#).

3. Verify that the group exists by using a command similar to one of the following. The following command lists all groups.

```
aws quicksight list-groups --aws-account-id 111122223333 --namespace default
```

The following command describes a specific group.

```
aws quicksight describe-group --aws-account-id 111122223333 --namespace default --group-name Sales
```

The following command searches for groups in a specified Amazon Quick Suite namespace.

```
aws quicksight search-groups --region us-west-2 --aws-account-id 111122223333 --namespace default --filters "[{\"Operator\": \"StartsWith\", \"Name\": \"GROUP_NAME\", \"Value\": \"Mar\"}]"
```

4. Add a member to the new group by using a command similar to the following.

```
aws quicksight create-group-membership --aws-account-id 111122223333 --namespace default --group-name Sales --member-name Pat
```

The following command determines if a user is a member of a specified group.

```
aws quicksight describe-group-membership --region us-west-2 --aws-account-id 111122223333 --namespace default --group-name Marketing-East --member-name user
```

Enter the following command at the prompt to delete a group. Substitute the correct values for your parameters.

```
aws quicksight delete-group --aws-account-id 111122223333 --namespace default --group-name Marketing-East
```

You might find it easier to create the command in a text editor before entering it at the prompt. For more information on `delete-group` and other available commands, see the [Amazon Quick Suite API reference](#).

Managing security

Security management in Amazon Quick Suite encompasses multiple layers of protection to ensure your data and resources remain secure. This section covers key security features including domain management for controlling access to specific organizational domains, namespace configuration for multi-tenant environments, mobile device security settings, IP address and VPC restrictions, VPC connection management, and customer-managed encryption keys. These security controls help you implement defense-in-depth strategies and meet your organization's compliance requirements while maintaining the flexibility to share insights securely across your organization.

Topics

- [Managing domains](#)
- [Supporting multitenancy with isolated namespaces](#)
- [Managing mobile settings](#)
- [Managing IP/VPC restrictions](#)
- [Managing VPC connection in Amazon Quick Suite](#)
- [Encrypting your Amazon Quick Suite data with Amazon Key Management Service customer-managed keys](#)

Managing domains

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite administrators

In Amazon Quick Suite Enterprise edition, you can embed Amazon Quick Sight dashboards, visuals, consoles, and Q search bars in an app or web page. Domains that are going to host these embedded assets must be on an *allow list*, the list of approved domains for your Quick Suite subscription. This requirement protects your data by preventing unapproved domains from hosting

embedded dashboards. To embed a Amazon Quick Sight dashboard, visual, console, or Q search bar to a web page or app, add approved domains to a static allow list in the Quick Suite console. Alternatively, add them at runtime with the Quick Suite API.

Use the following sections to learn more about adding domains for embedded analytics.

Topics

- [Allow listing static domains](#)
- [Allow listing domains at runtime with the Amazon Quick Suite API](#)

Allow listing static domains

You can add static domains to your allow list through the Amazon Quick Suite console. All domains on your allow list (such as development, staging, and production) must be explicitly allowed, and they must use HTTPS. You can add up to 100 domains to the allow list.

To embed a dashboard to a static domain:

- Approve the hosting domains and subdomains for embedding.
- Publish the dashboard.
- Share the dashboard with users or groups so they can see the embedded version of it.

Use the following procedure to view or edit the list of approved domains.

To view or edit the list of approved domains

1. Choose the profile icon at top right.
2. Choose **Manage Amazon Quick Suite**. You must be an Amazon Quick Suite admin to access this screen.
3. Choose **Domains and Embedding** on the left. The domains that you can embed a dashboard in are listed at the bottom of the page.
4. (Optional) Add a new domain here by entering it in the **Domain** box. You can also choose **Include subdomains** to allow embedded dashboards on all subdomains. Choose **Add** to add the domain.

You can edit or delete existing domain by choosing the icons next to each domain in the list at the bottom of the page.

Make sure that you use a valid HTTPS URL. The following list shows examples of URLs that are valid for embedded dashboards that use a static domain:

- `https://example-1.com`
- `https://www.アマゾンドメイン.jp`
- `https://www.亚马逊域名.cn:1234`
- `https://111.222.33.44:1234`
- `https://111.222.33.44`
- `http://localhost`

The following list shows examples of URLs that are *not* valid for embedded dashboards:

- `http://example`
- `https://example.com.*.example-1.co.uk`
- `https://co.uk`
- `https://111.222.33.44.55:1234`
- `https://111.222.33.44.55`

Allow listing domains at runtime with the Amazon Quick Suite API

You can add a domain at runtime to an allow list with the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` or a `GenerateEmbedUrlForRegisteredUser` API call. The `AllowedDomains` parameter is an optional parameter. It grants you the option as a developer to override the static domains that are configured in the **Manage Amazon Quick Suite** menu.

You can list up to three domains or subdomains. Adding domains to the allow list at runtime also adds HTTP support for the domain `localhost`. The generated URL is then embedded in a developer's website. Only the domains that are listed in the parameter can access the embedded dashboard.

To embed a dashboard to a domain at runtime, see [Embedding with the Amazon Quick Suite APIs](#).

Make sure that you use a valid URL. The following list shows examples of URLs that are valid for embedded dashboards that use a runtime domain:

- <https://example-1.com>
- <http://localhost>
- <https://www.アマゾンドメイン.jp>
- https://*.sapp.amazon.com

The following list shows examples of URLs that are *not* valid for embedded dashboards:

- https://example.com.*.example-1.co.uk
- <https://co.uk>
- <https://111.222.33.44.55:1234>
- <https://111.222.33.44.55>

For more information about embedded dashboards, see [Embedding with the Amazon Quick Suite APIs](#).

Supporting multitenancy with isolated namespaces

Amazon Quick Suite Enterprise edition supports multitenancy through namespaces. An Amazon Quick Suite *namespace* is a logical container that you can use to organize clients, subsidiaries, teams, and so on. Namespaces can help you achieve the following goals:

- You can allow the users of your Amazon Quick Suite subscription to discover shared content and share with other users. At the same time, you can be sure that users in one namespace can't see or interact with users in another namespace.
- You can securely isolate data and also support diverse workloads without adding additional Amazon accounts. Access to data is still strictly controlled by Amazon security features. Users can see assets (like data and dashboards) only if they have the correct resource permissions. Also, users who have permissions can't inadvertently expose content to people who outside of their namespace. For more information, see [Supporting multitenancy with isolated namespaces](#).
- You can monitor data flows and usage reports, neatly partitioned by namespace. Categorizing data and reports by namespace can help simplify cost and security analysis.
- After you've registered users into your namespace, there's no additional administrative complexity or overhead.
- Namespaces are designed to span Amazon Web Services Regions, so the use containment doesn't change even if a person signs in to a different Amazon Web Services Region.

Namespaces currently have the following limitations:

- Custom namespaces—those that are not the default namespace—are only accessible to IAM Federated Single-Sign On users.
- Use default namespaces instead of custom namespaces if you need to support the following:
 - Integrating your Amazon Quick Suite account with IAM Identity Center. For more information on integrating your Amazon Quick Suite account with IAM Identity Center, see [Amazon security in Amazon Quick Suite](#).
 - Password-based logins.
 - Credential-based Active Directory logins.
- You can't transfer users directly from one namespace to another. You can choose to do some or all of this work programmatically. For more information, see the [Quick Suite API reference](#). At the bottom of the page of each API operation, there's a list of links to the same operation in the SDKs for other languages. To see what SDKs are available, see [SDKs and toolkits](#) in the [Amazon getting started resource center](#).
- Namespaces are useful for isolating users and permissions, but not for sharing assets. Dashboards, datasets, and analyses can be shared with users in different namespaces. By default, users can't access items that exist in the same namespace by default, but gain access to specific assets when the asset is shared with them.

If you don't have an existing Amazon Web Services account or you need to sign up for Amazon Quick Suite, read the following guidelines, then follow the applicable instructions in [Signing up for an Amazon Quick Suite subscription](#):

- Sign up for Enterprise edition.
- When asked which method you want to connect with, choose **Role Based Federation (IAM)**. Currently, namespaces support only customers who use an Amazon Identity and Access Management (IAM) role with a web identity federation. For more information, see [Creating a role for a third-party Identity Provider \(federation\)](#)
- Complete the process of signing up.
- Use the Amazon Quick Suite [CreateNamespace](#) API operation to create one or more namespaces.
- To start adding users, first follow the instructions in [Setting up IdP federation using IAM and Amazon Quick Suite](#). Then use the [RegisterUser](#) API operation to add users to the appropriate namespace.

If you already signed up for Standard edition, you can easily upgrade your subscription to Enterprise edition. The person performing the upgrade must be a Amazon Quick Suite user with administrator privileges. For more information, see [Upgrading your Amazon Quick Suite subscription](#).

If you have an Enterprise edition subscription that you've been using for some time, it's also possible to migrate your users into namespaces. When you sign up for Amazon Quick Suite and add users, all of them reside in the default namespace. All of the users can interact directly with each other and share data and dashboards with each other. To isolate your users from each other, you can create one or more additional namespaces.

⚠ Important

Amazon Quick Suite assets and resources, including datasets, data sources, dashboards, analyses, and so on, exist outside of any namespace. They're visible only to users who have resource permissions granted to them.

To implement namespaces, you use the following Amazon Quick Suite API operations:

- [CreateNamespace](#)
- [DescribeNamespace](#)
- [ListNamespaces](#)
- [DeleteNamespace](#)

Namespaces are not supported in the Regions listed below:

- af-south-1 Africa (Cape Town)
- ap-southeast-3 Asia Pacific (Jakarta)
- eu-south-1 Europe (Milan)
- eu-central-2 Europe (Zurich)

ℹ Note

If you need to install the Amazon CLI, see [Installing the Amazon CLI version 2](#) in the *Amazon Command Line Interface User Guide*.

To add users to a namespace, you use the [RegisterUser](#) API operation. Each namespace has a completely independent set of users. The user ARNs include the namespace qualifier to distinguish them, as shown in the following examples:

- Amazon Quick Suite considers these two entities to be different persons:
 - `arn:aws-cn:quicksight:us-east-1:111122223333:user/namespace-123/username123`
 - `arn:aws-cn:quicksight:us-east-1:111122223333:user/namespace-456/username123`
- Amazon Quick Suite considers these two entities to be the same person:
 - `arn:aws-cn:quicksight:us-east-1:111122223333:user/namespace-123/username123`
 - `arn:aws-cn:quicksight:us-west-2:111122223333:user/namespace-123/username123`

When you use [RegisterUser](#), you select an access level for each user. After a person's user name is assigned to one of the security cohorts, their access to the console and API is restricted. People using Amazon Quick Suite can have a single access level, as follows:

- Reader access, for read-only subscribers of a dashboard
- Author access, for analysts and dashboard designers
- Admin access, for Amazon Quick Suite administrators

To migrate existing users in one namespace to a different namespace

Follow the procedure below to migrate existing users from one namespace to a different namespace.

1. Identify the users that you want to transfer to a different namespace by using the Amazon Quick Suite user and group API operations. For more information, see [API operations for controlling access](#) in the [Quick Suite API reference](#).
2. Create users in the new namespace by using the [RegisterUser](#) API operation. Within a namespace, user names are unique.

If a namespace user starts using the Amazon Quick Suite console or API in a new Amazon Web Services Region, that user is still constrained to the namespace that you added them to. Each

namespace represents a user directory of an identity provider. As such, it originates in the primary Amazon Web Services Region where Amazon Quick Suite is set up. However, because the user directory is propagated globally in your Amazon account, the namespace is accessible from any Amazon Web Services Region where your users are using Amazon Quick Suite.

3. To identify the asset and resource permissions that the new namespace users need, use the Amazon Quick Suite API operations associated with each type of asset (dashboards, datasets, and so on). For more information, see [QuickSight API operations to control assets](#) in the [Quick Suite API reference](#).

For example, let's say you are focusing on dashboards. You can use `ListDashboards` to list all the dashboard IDs in your Amazon account. Then, to determine which users or groups can access these dashboards, you can use `DescribeDashboardPermissions` on the result set generated by `ListDashboards`. If you need to identify specific versions of a dashboard, you can `ListDashboardVersions` for that. You can also collect information about the location of the data that's used in the dashboard with the data source and dataset API operations. For more information, see [QuickSight API operations to control data resources](#) in the [Quick Suite API reference](#).

For more information about filtering API response output, see the SDK documentation for the language you're using. For information relating to the Amazon Command Line Interface (Amazon CLI), see [Controlling command output from the Amazon CLI](#) in the [Amazon Command Line Interface User Guide](#).

4. For Amazon Quick Suite assets and resources, copy the permissions that the source namespace user has for each asset. Then use, for example, `UpdateDashboardPermissions` to apply the same permissions to the target namespace user. Each asset type has its own separate set of API operations for controlling the permissions that users have to use it. For more information, see [QuickSight API operations for asset and resource permissions](#) in the [Quick Suite API reference](#).
5. When you are finished adding users and permissions, it's a good practice to allow some time for user acceptance testing. Doing this ensures that everyone is successfully using the new namespace. It also ensures that all assets and resources are accessible in the new namespace.

After you're certain that you no longer need the original user names, you can begin to deprecate their permissions in the original namespace. Finally, when the users are ready, you can remove the unused group and user names in the source namespace. Do this in each Amazon Web Services Region where your users were previously active.

Managing mobile settings

Mobile settings allow Amazon Quick Suite admins to enhance data security and maintain compliance standards for users accessing Quick Suite through mobile devices. These settings apply to all users accessing Quick Suite through the mobile application.

To access mobile settings:

1. Open the [Quick Suite console](#).
2. Choose the user icon at the top right, and then choose **Manage Amazon Quick Suite**.
3. Under the **Security** section, choose **Mobile settings**.
4. Choose settings to enhance your data security on mobile devices:
 - **Require use of biometrics or PIN to unlock device** - Check the box to enable this security requirement. When enabled, users must use fingerprint recognition, Face ID, device PINs, or other biometric authentication methods supported by their device to access Amazon Quick Suite.
 - **Require devices to use the latest version of the mobile operating system** - Check the box to enable this security requirement. When enabled, users must keep their mobile devices' operating systems up to date to access Amazon Quick Suite.
5. To manage how long users can remain logged in on mobile devices:
 - a. Locate the **Current session lifetime for all devices** setting.
 - b. Select the desired duration from the dropdown menu (default is 30 days).
 - c. Choose **Update** to apply your changes.

Best practices

- Regularly review and update mobile security settings to align with your organization's security policies
- Set an appropriate session lifetime that balances security with user convenience
- Enable biometric/PIN requirements for additional security, especially when dealing with sensitive data
- Keep operating system requirements enabled to ensure devices have the latest security updates
- Communicate changes in mobile settings to your users before implementation

Managing IP/VPC restrictions

Turning on Internet Protocol (IP) and VPC endpoint restrictions in Amazon Quick Suite

You can limit access to your organization's Amazon Quick Suite account to a predefined list of IP ranges, VPC IDs, and VPC endpoint IDs. For example, you can create an IP rule that allows users to access your Quick Suite account only from IP addresses associated with your company's office or remote virtual private network (VPN). You can also create a VPC endpoint rule that allows users access to your Quick Suite account only from the VPC that is used for Amazon Direct Connect.

For more information about setting up VPC endpoints in Quick Suite, see [Quick Suite Interface VPC endpoints \(Amazon PrivateLink\)](#) for more information about how to setup VPC endpoints.

Only admins with Amazon Identity and Access Management (IAM) credentials who have access to the Quick Suite console pages can access the IP and VPC endpoint restrictions table.

Topics

- [Adding an IP or VPC endpoint rule](#)
- [Update an existing rule](#)
- [Delete a rule](#)
- [Turning on your IP and VPC endpoint rules](#)

Adding an IP or VPC endpoint rule

An *IP rule* is created when you add a CIDR address with a public IP version 4 address to the restrictions table. A *VPC endpoint rule* is created when you add either a VPC ID or a VPC endpoint ID to the restrictions table. You can add up to IP or VPC endpoint rules to the restrictions table. You can only add rules from the Amazon Web Services Region where your account is. All traffic that is not allowed by either the IP rule or the VPC endpoint rule is blocked when the restriction is turned on.

A *CIDR address* is composed of two parts: the prefix and the suffix. The prefix is the CIDR's network address and is written like a normal IP address. The suffix shows how many bits are in the address. An example of a complete CIDR address is `10.24.34.0/23`.

IP and VPC endpoint rules apply only to Quick Suite web, embedded, and mobile access and don't restrict access to the public API. Your users can still call all API operations from restricted IP ranges.

For information on restricting calls to the public API from specific IP addresses, see [Amazon: Denies access to Amazon based on the source IP](#) in the *IAM User Guide*.

Before you save any rule changes or turn on other rules, make sure that you have a rule that includes your IP address or VPC endpoint ID. If there isn't a rule that includes allows your traffic, you can't save your changes.

To add an IP or VPC endpoint rule

1. On the Quick Suite start page, choose **Manage QuickSight**, and then choose **Security and Permissions**.
2. Choose **IP and VPC endpoint restrictions**.
3. Perform one of the following actions.
 - a. For **IP restriction**, enter the CIDR address that defines the IP range that you want to create a rule for.
 - b. For **VPC endpoint restriction**, enter either the VPC ID or the VPC endpoint ID of the endpoint that you want to create a rule for.
4. (Optional) For **Description**, enter a description for the rule. Doing this can help you differentiate your rules.
5. Choose **Add**.
6. Choose **Save changes** in the box that appears to apply the rule.

It can take up to 10 minutes for a rule to be fully implemented.

Update an existing rule

After you create an IP or VPC endpoint rule, use the IP and VPC restrictions table to make changes to the rule. Use the procedure below to update an existing IP or VPC endpoint rule in the IP and VPC restrictions table.

To update an existing IP or VPC endpoint rule

1. On the Quick Suite start page, choose **Manage QuickSight**, and then choose **Security and Permissions**.
2. Choose **IP and VPC endpoint restrictions**.
3. Choose the edit icon to the right of the rule that you want to change.

4. Make your changes and choose **Update**.
5. Choose **Save changes** in the box that appears to update the rule.

It can take up to 10 minutes for an updated rule to be fully implemented.

Delete a rule

Use the procedure below to delete an IP or VPC endpoint rule from the IP and VPC endpoint restrictions table.

To delete an IP rule

1. On the Quick Suite start page, choose **Manage QuickSight**, and then choose **Security and Permissions**.
2. Choose **IP and VPC endpoint restrictions**.
3. Make your changes and choose **Update**. A rule marked for deletion appears with a strike through it.
4. Choose **Save changes** in the box that appears to delete the rule.

It can take up to 10 minutes for an updated rule to be deleted.

Turning on your IP and VPC endpoint rules

You can turn on or turn off your account's IP and VPC endpoint restrictions by using the **Rules** option at the top of the IP and VPC restrictions page. When rules are turned on, users from sources that are not on the restrictions table can't access Quick Suite mobile, embedded, and website pages. IP and VPC endpoint rules are global and apply to all Amazon Web Services Regions.

If a user is accessing the Quick Suite account from a source that is not on the rule list when you turn on restrictions, they lose access to the account.

Account holders can audit users who make changes to the IP and VPC endpoint restrictions table by using Amazon CloudTrail. For more information, see the [Amazon CloudTrail User Guide](#).

Managing VPC connection in Amazon Quick Suite

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick Suite administrators

With Amazon Quick Suite Enterprise Edition, account admins can configure a secure, private VPC connection to a Amazon Quick Suite account from the Amazon Quick Suite console or from the Amazon Quick Suite CLI. Read the following walkthroughs on how to create, edit, and delete a VPC connection from a Amazon Quick Suite account.

Topics

- [Configuring the VPC connection in the Amazon Quick Suite console](#)
- [Configuring the VPC connection with the Amazon Quick Suite CLI](#)
- [Testing the connection to your VPC data source](#)

Configuring the VPC connection in the Amazon Quick Suite console

To create a secure private connection to the Amazon VPC service from the Amazon Quick Suite console, use the following procedure.

Prerequisites

- Sign in to Amazon Quick Suite as a Amazon Quick Suite admin to set up a VPC connection in Amazon Quick Suite. To verify that you're a Amazon Quick Suite administrator, choose your profile icon in the upper-right. If your profile menu contains the option **Manage Amazon Quick Suite**, then you're a Amazon Quick Suite administrator. Make sure your admin role in IAM includes the following permissions. The "iam:PassRole" permission needs to be applied only to the execution role that is created in the procedures below.
 - "quicksight:ListVPCConnections"
 - "quicksight:CreateVPCConnection"
 - "quicksight:DescribeVPCConnection"
 - "quicksight>DeleteVPCConnection"
 - "quicksight:UpdateVPCConnection"
 - "ec2:describeSubnets"
 - "ec2:describeVpcs"
 - "ec2:describeSecurityGroups"
 - "iam:ListRoles"

- "iam:PassRole"

The following example shows an IAM policy that applies "iam:PassRole" only to the execution role.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "iam:PassRole"
      ],
      "Resource": "arn:aws-cn:iam::111122223333:role/vpc-role-for-qs"
    }
  ]
}
```

- Before you begin, make sure that you have the following information available to copy and paste into the **VPC Connection** screen. For more information, see [Finding information to connect to a VPC](#).
 - Amazon Web Services Region – The Amazon Web Services Region where you plan to create a connection to your data source.
 - VPC ID – The ID of the VPC that contains the data, the subnets, and the security groups that you plan to use.
 - Execution role– An IAM role that contains a trust policy that allows Amazon Quick Suite to create, update, and delete network infrastructure in your account. This policy is required for all VPC connections. At minimum, the IAM policy needs the following Amazon EC2 permissions:
 - DescribeSecurityGroups
 - DescribeSubnets
 - CreateNetworkInterface
 - DeleteNetworkInterface
 - ModifyNetworkInterfaceAttribute

The following example shows an IAM policy that you can add to an existing IAM role to create, delete, or modify a VPC connection:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:CreateNetworkInterface",
        "ec2:ModifyNetworkInterfaceAttribute",
        "ec2>DeleteNetworkInterface",
        "ec2:DescribeSubnets",
        "ec2:DescribeSecurityGroups"
      ],
      "Resource": "*"
    }
  ]
}
```

After you add the necessary permissions to an IAM role, attach a trust policy to allow Amazon Quick Suite to configure the VPC connection to your account. The following example shows a trust policy that you can add to an existing IAM role to allow Amazon Quick Suite access to the role:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "quicksight.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

- Subnet IDs – The IDs of the subnets that the Amazon Quick Suite network interface is using. Each VPC connection needs at least two subnets.

- Security group IDs – The IDs of the security groups. Each VPC connection needs at least one security group.

To create a secure private connection to the Amazon VPC service from Quick Suite Enterprise edition

1. In Amazon Quick Suite, choose your profile icon in the upper-right, then choose **Manage Amazon Quick Suite**.

Only Amazon Quick Suite administrators can view the **Manage Amazon Quick Suite** option. If you don't see this option on your profile menu, you're not an administrator. In this case, contact your Amazon Quick Suite account administrators for assistance.

2. In the left navigation pane, choose **Manage VPC connections**.
3. On the **Manage VPC connections** page that opens, choose **Add VPC connection**.
4. For **VPC connection name**, enter a unique descriptive name of your choice. This name doesn't need to be an actual VPC ID or name.
5. In the **VPC ID** dropdown menu, choose the ID of the VPC in Amazon EC2 that you want to connect to your Amazon Quick Suite account. This field can't be changed later.
6. In the **Execution role** dropdown menu, choose the appropriate IAM role to use for the VPC connection. The **Execution role** dropdown only shows IAM policies that contain a trust policy that allows Amazon Quick Suite to configure the VPC connection to your account.
7. In the **Subnets** table, choose a subnet ID from the **Subnet ID** dropdown menu of at least two of the listed **Availability zones**. The Availability Zones listed in the **Subnets** table are determined based on how you configured the VPC connection in the Amazon EC2 console.
8. (Optional) If you aren't using DNS resolver endpoints, skip to the next step.

If your database host IP address must be resolved through private DNS servers in your Amazon account, enter IP addresses for Route 53 Resolver inbound endpoints (one per line).

Make sure that you are entering an endpoint, rather than a database address like the one you plan to use in Amazon Quick Suite. Most databases that are hosted by Amazon don't need to resolve DNS queries between VPCs and a customer's network. For more information, see [Resolving DNS queries between VPCs and your network](#) in the *Amazon Route 53 Developer Guide*. You only need this endpoint if you can't resolve the IP address that connects to your database by using the public DNS server system.

9. Review your choices, then choose **ADD**.

When you finish creating a VPC connection, the new connection appears in the **Manage VPC connections** table. In some cases, the status of the new VPC might be **UNAVAILABLE** until the connection is configured on the backend. After Amazon Quick Suite is finished configuring the new connection, the status of the connection switches to **AVAILABLE**, which indicates that the connection has been established. The following table describes the different **Status** values for a VPC connection.

Status	Description
AVAILABLE	The VPC connection is established and can be used.
PARTIALLY AVAILABLE	One of the network interfaces that is configured to the VPC connection is unavailable. The VPC connection can still be used.
UNAVAILABLE	The VPC connection is not established and can't be used.

To see a summary of a VPC connection, choose a VPC connection from the **VPC connection name** row of the **Manage VPC connections** table. The pop-up box that appears shows information about the network interfaces associated with the VPC connection.

The following table describes the different **Status** values for a network interface.

Status	Description
CREATING	The network interface creation is in progress.
AVAILABLE	The network interface is available for use.
CREATION_FAILURE	The network interface couldn't be created.
UPDATING	The security group associated with the network interface is updating.
UPDATE_FAILED	The security group associated with the network interface did not update successfully.

Status	Description
DELETING	The network interface is in the process of being deleted.
DELETED	The network interface is deleted and can no longer be used.
DELETION_FAILED	The network interface deletion failed and can still be used.
DELETION_SCHEDULED	This network interface is scheduled for deletion.
ATTACHMENT_FAILED_ROLLBACK_FAILED	The elastic interface failed to attach and Amazon Quick Suite was unable to delete the elastic network interface that was created within your account.

When you delete a network interface from a VPC connection, the status of the connection changes to **PARTIALLY AVAILABLE** to indicate the loss of a network interface.

To make changes to an existing VPC connection, choose the more actions (three-dots) button to the right of the connection that you want to modify, and choose **Edit**. In the **Edit VPC connection** window that appears, make your changes, and then choose **SAVE**.

To delete a VPC connection, choose the more actions (three-dots) button to the right of the connection that you want to delete and choose **Delete**. In the **Delete Amazon Quick Suite VPC Connection** pop-up that appears, confirm that you want to delete the connection, and then choose **Delete**.

Configuring the VPC connection with the Amazon Quick Suite CLI

To create a secure private connection to the Amazon VPC service from Quick Suite with the Amazon Quick Suite CLI, use the following procedure:

Prerequisites

- Before you begin, make sure that you have the following information available to copy and paste into the **VPC Connection** page. For more information, see [Finding information to connect to a VPC](#).
 - Amazon Web Services Region – The Amazon Web Services Region where you plan to create a connection to your data source.
 - VPC ID – The ID of the VPC that contains the data, the subnets, and the security groups that you plan to use.
 - Execution role– An IAM role that contains a trust policy that allows Amazon Quick Suite to create, update, and delete network infrastructure in your account. This policy is required for all VPC connections. At minimum, the IAM policy needs the following Amazon EC2 permissions:
 - DescribeSecurityGroups
 - DescribeSubnets
 - CreateNetworkInterface
 - DeleteNetworkInterface
 - ModifyNetworkInterfaceAttribute

The following example shows an IAM policy that you can add to an existing IAM role to create, delete, or modify a VPC connection:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:CreateNetworkInterface",
        "ec2:ModifyNetworkInterfaceAttribute",
        "ec2>DeleteNetworkInterface",
        "ec2:DescribeSubnets",
        "ec2:DescribeSecurityGroups"
      ],
      "Resource": "*"
    }
  ]
}
```

```
}
```

After you add the necessary permissions to an IAM role, attach a trust policy to allow Amazon Quick Suite to configure the VPC connection to your account. The following shows an example trust policy that you can add to an existing IAM role to allow Amazon Quick Suite access to the role:

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "quicksight.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

- Subnet IDs – The IDs of the subnets that the Amazon Quick Suite network interface is using. Each VPC connection needs at least two subnets.
- Security group IDs – The IDs of the security groups. Each VPC connection needs at least one security group.

Using the Amazon CLI

The following example creates a VPC connection.

```
aws quicksight create-vpc-connection \
--aws-account-id 123456789012 \
--vpc-connection-id test \
--name test \
--subnet-ids '["subnet-12345678", "subnet-12345678"]' \
--security-group-ids '["sg-12345678"]' \
--role-arn arn:aws-cn:iam::123456789012:role/test-role \
--region us-west-2
```

After you create a VPC connection, you can update, delete, or request a summary of the VPC connection.

The following example updates a VPC connection.

```
aws quicksight update-vpc-connection \  
--aws-account-id 123456789012 \  
--vpc-connection-id test \  
--name test \  
--subnet-ids '["subnet-12345678", "subnet-12345678"]' \  
--security-group-ids '["sg-12345678"]' \  
--role-arn arn:aws-cn:iam::123456789012:role/test-role \  
--region us-west-2
```

The following example deletes a VPC connection.

```
aws quicksight delete-vpc-connection \  
--aws-account-id 123456789012 \  
--vpc-connection-id test \  
--region us-west-2
```

The following example describes a VPC connection.

```
aws quicksight describe-vpc-connection \  
--aws-account-id 123456789012 \  
--vpc-connection-id test \  
--region us-west-2
```

The following table describes the different **Status** values for a network interface that `describe-vpc-connection` returns.

Status	Description
CREATING	The network interface creation is in progress.
AVAILABLE	The network interface is available for use.

Status	Description
CREATION_FAILURE	The network interface couldn't be created.
UPDATING	The security group associated with the network interface is updating.
UPDATE_FAILED	The security group associated with the network interface did not update successfully.
DELETING	The network interface is in the process of being deleted.
DELETED	The network interface is deleted and can no longer be used.
DELETION_FAILED	The network interface deletion failed and can still be used.
DELETION_SCHEDULED	This network interface is scheduled for deletion.
ATTACHMENT_FAILED_ROLLBACK_FAILED	The elastic interface failed to attach and Amazon Quick Suite was unable to delete the elastic network interface that was created within your account.

You can also use the Amazon CLI to generate a list of all VPC connections in your Amazon Quick Suite account.

```
aws quicksight list-vpc-connections \  
--aws-account-id 123456789012 \  
--region us-west-2
```

Testing the connection to your VPC data source

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite administrators and authors

To test whether you can connect to your data source through an existing Quick Suite VPC connection, use the following procedure.

Before you begin, collect the information you need to connect. If you plan to copy and paste settings from a file, make sure that it doesn't contain any of the following: formatting (list bullets or numbers), blank space (spaces, tabs), or invisible "gremlin" (non-ASCII, null (ASCII 0), or control) characters.

1. On the Amazon Quick Suite start page, choose **Manage data**.
2. On the **Datasets** page, choose **New data set**.
3. In the **FROM NEW DATA SOURCES** section of the **Create a data set** page, choose a supported data source that you want to connect to. For a list of data sources that support VPC, see [Identify the data sources to use](#).

Your data source instance must use the same VPC that you used to create the VPC connection. Also, the associated security group must be properly configured. For more information, see [Setting up a VPC to use with Amazon Quick Suite](#).

4. Enter the connection information for the data source. The fields for the data source are sometimes displayed in different order depending on which data source you choose. For more information, see [Creating a data source](#).
 - For **Data source name**, enter a descriptive name for the new data source. This name appears beside the data source logo on a tile on the **Create a data set** page. For testing purposes, name it "**VPC test -**" followed by the database name or location, whichever is unique.
 - For **Connection Type**, choose the name of the VPC connection that has a route to your data source. If the correct VPC is missing from the list, ask a Amazon Quick Suite administrator to verify that the VPC connection is correct in Amazon Quick Suite. If it looks correct, ask a system administrator to verify that the data source and VPC are set up for this purpose.
 - The name or other identifier for the server or instance to connect to. The descriptors vary depending on which one you're connecting to, but it's usually one or more of the following: hostname, IP address, cluster ID, instance ID, connector, or site based URL.
 - **Database name** shows the default database for the **Instance ID** cluster or instance. If you want to use a different database on that cluster or instance, enter its name.

- The name of the collection of data that you want to use.

The descriptor varies depending on the provider, but it's usually one of the following: database, warehouse, or catalog. In this topic, we use the word "database" as a generic term.

- For **Credentials**, enter a username and password to use for everyone who connects from Amazon Quick Suite using this data source. The username must have permissions to do the following:
 - Access the target database.
 - Read (perform a SELECT statement on) all of the tables that you want to use in that database.
5. Choose **Validate connection** to verify your connection information is correct. If your connection doesn't validate, correct the connection information and try again. If the information looks correct but doesn't validate, do one or all of the following:
- Contact your data source administrator to verify your connection settings.
 - Contact your Amazon Quick Suite administrator to verify the settings in the Amazon Quick Suite VPC connection.
 - Contact your Amazon administrator to verify that the VPC is correctly configured for use with Amazon Quick Suite.
6. After the connection validates, choose **Create data source** to save the connection profile. Or, choose **Cancel** if you don't need to save it (recommended) after testing is complete.

Encrypting your Amazon Quick Suite data with Amazon Key Management Service customer-managed keys

Amazon Quick Suite enables you to encrypt your Amazon Quick Suite data with the keys you have stored in Amazon Key Management Service. This provides you with the tools to audit access to data and satisfy regulatory security requirements. If you need to do so, you have the option to immediately lock down access to your data by revoking access to Amazon KMS keys. All data access to encrypted resources in Amazon Quick Suite is logged in Amazon CloudTrail. Administrators or auditors can trace data access in CloudTrail to identify when and where data was accessed.

To create customer-managed keys (CMKs), you use Amazon Key Management Service (Amazon KMS) in the same Amazon account and Amazon Region as the Amazon Quick Suite resource. A

Amazon Quick Suite administrator can then use a CMK to encrypt your Amazon Quick Suite data and control access.

You can create and manage CMKs in the Amazon Quick Suite console or with the Amazon Quick Suite APIs. For more information about creating and managing CMKs with the Amazon Quick Suite APIs, see [Key management operations](#).

The following rules apply to using CMKs with Amazon Quick Suite resources:

- Amazon Quick Suite doesn't support asymmetric Amazon KMS keys.
- You can have multiple CMKs and one default CMK per Amazon Web Services account per Amazon Web Services Region.
- By default, Amazon Quick Suite resources are encrypted with Amazon Quick Suite–native encryption strategies.
- Data currently encrypted by a CMK key will stay encrypted by the key.

Note

If you use Amazon Key Management Service with Amazon Quick Suite, you are billed for access and maintenance as described in the [Amazon Key Management Service Pricing page](#). In your billing statement, the costs are itemized under Amazon KMS and not under Amazon Quick Suite.

Note

Amazon Q data is encrypted by an Amazon managed key, not the default Amazon KMS key.

The key that is currently the default CMK is automatically used to encrypt the following:

- New SPICE datasets. Existing datasets need to be fully refreshed to be encrypted by the new default key.
- New report artifacts generated through the dashboard snapshot API, scheduled reports and exports, or dashboards.

All non-customer managed keys associated with Amazon Quick Suite are managed by Amazon.

Database server certificates that are not managed by Amazon are the responsibility of the customer and should be signed by a trusted CA. For more information, see [Network and database configuration requirements](#).

Use the following topics to learn more about using CMKs with Amazon Quick Suite. To learn more about data encryption in Amazon Quick Suite see [Data protection in Amazon Quick Suite](#).

Topics

- [Add a CMK to your account](#)
- [Verify the key used by Amazon Quick Suite](#)
- [Changing the default CMK](#)
- [Removing CMK encryption on your Amazon Quick Suite account](#)
- [Auditing CMK usage in CloudTrail](#)
- [Revoking access to a CMK](#)
- [Recovering encrypted Amazon Quick Suite data](#)

Add a CMK to your account

Before you begin, make sure that you have an IAM role that grants the admin user access to the Amazon Quick Suite admin key management console. For more information on the required permissions, see [IAM identity-based policies for Amazon Quick Suite: using the admin key management console](#).

You can add keys that already exist in Amazon KMS to your Amazon Quick Suite account, so that you can encrypt your Amazon Quick Suite data.

To learn more about how you can create a key to use in Amazon Quick Suite, see the [Amazon Key Management Service Developer Guide](#).

To add a new CMK to your Amazon Quick Suite account.

1. On the Amazon Quick Suite start page, choose **Manage Amazon Quick Suite**, and then choose **KMS keys**.
2. On the **KMS keys** page, choose **Manage**. The **KMS keys** dashboard opens.
3. On the **KMS Keys** dashboard, choose **Select key**.
4. On the **Select key** pop-up box, choose **Key** to open the list. Then, select the key that you want to add.

If your key isn't in the list, you can manually enter the key's ARN.

5. (Optional) Select the **Use as default encryption key for all new data in the current region of this Amazon Quick Suite account** to set the selected key as your default key. A badge appears next to the default key to indicate its status.

When you choose a default key, all new data that is created in the region that hosts your Amazon Quick Suite account is encrypted with the default key.

6. (Optional) Add more keys by repeating the previous steps in this procedure. While you can add as many keys as you want, you can only have one default key at one time.

Verify the key used by Amazon Quick Suite

When a key is used, an audit log is created in Amazon CloudTrail. You can use the log to track the key's usage. If you need to know which key the Amazon Quick Suite data is encrypted by, you can find this information in CloudTrail.

To learn more about which data can be managed with the key, see [Encrypting your Amazon Quick Suite data with Amazon Key Management Service customer-managed keys](#).

Verify the CMK that's currently used by a SPICE dataset

1. Navigate to your CloudTrail log. For more information, see [Logging Amazon Quick Suite information with CloudTrail](#).
2. Locate the most recent grant events for the SPICE dataset, using the following search arguments:
 - The event name (eventName) contains Grant.
 - The request parameters requestParameters contain the Amazon Quick Suite ARN for the dataset.

```
{
  "eventVersion": "1.08",
  "userIdentity": {
    "type": "AWSService",
    "invokedBy": "quicksight.amazonaws.com"
  },
  "eventTime": "2022-10-26T00:11:08Z",
```

```

"eventSource": "kms.amazonaws.com",
"eventName": "CreateGrant",
"awsRegion": "us-west-2",
"sourceIPAddress": "quicksight.amazonaws.com",
"userAgent": "quicksight.amazonaws.com",
"requestParameters": {
  "constraints": {
    "encryptionContextSubset": {
      "aws:quicksight:arn": "arn:aws-cn:quicksight:us-
west-2:111122223333:dataset/12345678-1234-1234-1234-123456789012"
    }
  },
  "retiringPrincipal": "quicksight.amazonaws.com",
  "keyId": "arn:aws-cn:kms:us-
west-2:111122223333:key/87654321-4321-4321-4321-210987654321",
  "granteePrincipal": "quicksight.amazonaws.com",
  "operations": [
    "Encrypt",
    "Decrypt",
    "DescribeKey",
    "GenerateDataKey"
  ]
},
....
}

```

3. Depending on the event type, one of the following applies:

CreateGrant – You can find the most recently used CMK in the key ID (keyID) for the last CreateGrant event for the SPICE dataset.

RetireGrant – If latest CloudTrail event of the SPICE datasets is RetireGrant, there is no key ID and the resource is no longer CMK encrypted.

Verify the CMK that's currently used when generating report artifacts

1. Navigate to your CloudTrail log. For more information, see [Logging Amazon Quick Sight information with Amazon CloudTrail](#).
2. Locate the most recent GenerateDataKey events for the report execution, using the following search arguments:
 - The event name (eventName) contains GenerateDataKey or Decrypt.

- The request parameters (requestParameters) contain the Amazon Quick Suite ARN for the analysis or dashboard the report was generated for.

```
{
  "eventVersion": "1.11",
  "userIdentity": {
    "type": "AWSService",
    "invokedBy": "quicksight.amazonaws.com"
  },
  "eventTime": "2025-07-23T23:33:46Z",
  "eventSource": "kms.amazonaws.com",
  "eventName": "GenerateDataKey",
  "awsRegion": "us-west-2",
  "sourceIPAddress": "quicksight.amazonaws.com",
  "userAgent": "quicksight.amazonaws.com",
  "requestParameters": {
    "keyId": "arn:aws:kms:us-west-2:111122223333:key/87654321-4321-4321-4321-210987654321",
    "keySpec": "AES_256",
    "encryptionContext": {
      "aws:quicksight:arn": "arn:aws-cn:quicksight:us-west-2:111122223333:dashboard/1ca456fe-eb34-4250-805c-b1b9350bd164",
      "aws:s3:arn": "arn:aws-cn:s3:::sn-imagegen.prod.us-west-2"
    }
  },
  ...
}
```

3. `aws:s3:arn` is the Amazon Quick Suite owned S3 bucket where your report artifacts are stored.
4. If you no longer see `GenerateDataKey`, then new report executions are no longer CMK encrypted. Existing report artifacts will remain encrypted.

Changing the default CMK

You can change the default key to another key that already exists in the **KMS keys** dashboard. When you change the default key, all new Amazon Quick Suite data is encrypted on the new key. The new default key changes how new Amazon Quick Suite data is encrypted. However, existing Amazon Quick Suite data will continue to use the previous default key.

To learn more about which data can be managed with the key, see [Encrypting your Amazon Quick Suite data with Amazon Key Management Service customer-managed keys](#).

To change the default key to an existing key

1. On the Amazon Quick Suite start page, choose **Manage Amazon Quick Suite**, and then choose **KMS keys**.
2. Choose **MANAGE** to open the **KMS keys** dashboard.
3. Navigate to the key that you want to set as your new default. Choose **Actions** (three dots) on the row of the key that you want to open the key's menu.
4. Choose **Set as default**, and then choose **Set**.

Note

The Q data key cannot be changed. Q data will remain encrypted with the current default key. In the event that this key is compromised, you can [revoke access to it](#).

The selected key is now your default key.

Removing CMK encryption on your Amazon Quick Suite account

You can remove the default key to disable data encryption in your Amazon Quick Suite account. Removing the key prevents new resources from encrypting on a CMK.

To remove CMK encryption for new Amazon Quick Suite data

1. On the Amazon Quick Suite start page, choose **Manage Amazon Quick Suite**, and then choose **KMS keys**.
2. On the **KMS keys** page, choose **Manage** to open the **KMS keys** dashboard.
3. Choose **Actions** (three dots) on the row of the default key, and then choose **Delete**.
4. In the pop-up box that appears, choose **Remove**.

After you delete the default key from your account, Amazon Quick Suite stops encrypting new Amazon Quick Suite data. Any existing encrypted data will remain encrypted. Q data remains encrypted because the Q data key cannot be changed. In the event that the deleted key is compromised, you can [revoke access to it](#).

Auditing CMK usage in CloudTrail

You can audit your account's CMK usage in Amazon CloudTrail. To audit your key usage, log in to your Amazon account, open CloudTrail, and choose **Event history**.

Revoking access to a CMK

You can revoke access to your CMKs. When you revoke access to a key that is used to encrypt your Amazon Quick Suite data, access to it is denied until you undo the revoke. The following methods are examples of how you can revoke access:

- Turn off the key in Amazon KMS.
- Add a Deny policy to your Amazon Quick Suite Amazon KMS policy in IAM.

To learn more about which data can be managed with the key, see [Encrypting your Amazon Quick Suite data with Amazon Key Management Service customer-managed keys](#).

Use the following procedure to revoke access to your CMKs in Amazon KMS.

To turn off a CMK in Amazon Key Management Service

1. Log in to your Amazon account, open Amazon KMS, and choose **Customer managed keys**.
2. Select the key that you want to turn off.
3. Open the **Key actions** menu and choose **Disable**.

To prevent further use of the CMK, you could add a Deny policy in Amazon Identity and Access Management (IAM). Use "Service": "quicksight.amazonaws.com" as the principal and the ARN of the key as the resource. Deny the following actions: "kms:Encrypt", "kms:Decrypt", "kms:ReEncrypt*", "kms:GenerateDataKey*", "kms:DescribeKey".

Important

After you revoke access by using any method, it can take up to 15 minutes for the data to become inaccessible.

Recovering encrypted Amazon Quick Suite data

To recover Amazon Quick Suite data while its access is revoked

1. Restore access to the CMK. Usually, this is enough to recover the Amazon Quick Suite data.
2. Test the Amazon Quick Suite data to check if you can see it.
3. (Optional) If the data is not fully recovered, even after you restored its access to the CMK, perform a full refresh on the data.

To learn more about which data can be managed with the key, see [Encrypting your Amazon Quick Suite data with Amazon Key Management Service customer-managed keys](#).

Permissions

Permission management in Quick Suite controls what actions users can perform and which resources they can access within your Amazon Quick Suite environment. This section covers essential permission-related tasks including configuring access to Amazon data sources, setting default access levels for new users, managing IAM policy assignments, creating custom permission profiles tailored to your organization's needs, configuring action connectors for third-party integrations, managing extension application access, and integrating with SageMaker AI for machine learning capabilities. Proper permission management ensures users have appropriate access to perform their roles while maintaining security and compliance requirements.

Topics

- [Accessing Amazon resources](#)
- [Default access policy](#)
- [IAM policy assignments](#)
- [Extension access](#)
- [Creating a custom permissions profile in Amazon Quick Suite](#)

Accessing Amazon resources

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators and Amazon Quick Suite administrators

You can control the Amazon resources that Amazon Quick Suite can access and scope down access to these resources at a more granular level. In Enterprise edition, you can also set up general access defaults for everyone in your account, and you can set up specific access for individual users and groups.

These access configurations are essential for Amazon Quick Sight data source connectivity, enabling secure connections to Amazon services like Amazon S3, Amazon RDS, Amazon Redshift, and Athena for data analysis and visualization. Proper resource access setup ensures that Amazon Quick Sight can retrieve and process data from your Amazon data sources while maintaining appropriate security boundaries.

Use the following sections to help you configure your Amazon resources to work with Quick Suite.

Before you begin, make sure that you have the correct permissions; your system administrator can give you these. To do so, your system administrator creates a policy that enables you to use certain IAM actions. Your system administrator then associates that policy with your user or group in IAM. The required actions are the following:

- **quicksight:AccountConfigurations** – To enable setting default access to Amazon resources
- **quicksight:ScopeDownPolicy** – Scoping policies for permissions to Amazon resources
- You can also bring your own IAM roles into Amazon Quick Suite. For more information, see [Passing IAM roles to Amazon Quick Suite](#).

To enable or disable the Amazon services that Amazon Quick Suite can access

1. Sign in to Amazon Quick Suite at <https://quicksight.aws.amazon.com/>.
2. At the upper right, choose your user name, and then choose **Manage Quick Suite**.
3. Choose **Security & permissions**.
4. Under **QuickSight access to Amazon services**, choose **Add or remove**.

A screen appears where you can enable all available Amazon services.

Note

If you see a permissions error, and you're an authorized Amazon Quick Suite administrator, contact your system administrator for assistance.

5. Select the check boxes for the services that you want to allow. Clear check boxes for services that you don't want to allow.

If you have already enabled an Amazon service, the check box for that service is already selected. If Amazon Quick Suite can't access a particular Amazon service, its check box is not selected.

In some cases, you might see a message like the following.

This policy used by Amazon Quick Suite for Amazon resource access was modified outside of Amazon Quick Suite, so you can no longer edit this policy to provide Amazon resource permission to Amazon Quick Suite. To edit this policy permissions, go to the IAM console and delete this policy permission with policy arn - `arn:aws-cn:iam::111122223333:policy/service-role/AWSQuickSightS3Policy`.

This type of message means that one of the IAM policies that Amazon Quick Suite uses was manually altered. To fix this, the system administrator needs to delete the IAM policy listed in the error message and reload the **Security & permissions** screen before you try again.

6. Choose **Update** to confirm, or **Cancel** to return to the previous screen.

Topics

- [Setting granular access to Amazon services through IAM](#)
- [Using Amazon Secrets Manager secrets instead of database credentials in Quick Suite](#)

Setting granular access to Amazon services through IAM

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick Suite administrators

In Enterprise edition, Amazon Quick Suite provides a way for you to set up detailed access to resources in Amazon services. Like every other Amazon service, Quick Suite uses IAM policies to control access for users and groups.

Before you begin, ask an administrator to set up the necessary IAM policies ahead of time. If these are set up, you can select them as part of the procedure in this section. For information about creating IAM policies to use with Quick Suite, see [Identity and access management in Quick Suite](#).

To assign an IAM policy to a user or group

1. Sign in to Quick Suite at <https://quicksight.aws.amazon.com/>.
2. At upper left, choose your user name, and then choose **Manage QuickSight**.
3. Choose **Security & permissions**.
4. Under **Resource access for individual users and groups**, choose **IAM policy assignments**.

The remaining steps at this point involve choosing an IAM policy to assign to the user or group. You can assign multiple IAM policies to one Amazon Quick Suite user or group. To determine permissions, Amazon Quick Suite performs a union and an intersection with the Amazon Web Services account-level policies.

If you already have active IAM policy assignments, they are listed on this page. You can search for existing assignments by using the search box. If you have drafts that aren't active yet, they are listed under **Assignment drafts**.

5. Choose one of the following:
 - To create an IAM policy assignment, choose **Add new assignment**.
 - To edit an existing assignment, choose the **Edit assignment** icon for that assignment.
 - To enable or disable a policy, select the check box for that policy, and then choose **Enable** or **Disable**. You can select multiple policy assignments at a time.
 - To delete an existing assignment, choose the **Remove assignment** icon near the name of the assignment. To confirm your choice, choose **Delete** on the confirmation screen. Or choose **Back** to cancel deletion.

If you are creating or editing an assignment, continue to the next step. Otherwise, skip to the end of this procedure.

6. On the next screen, you perform the policy assignment process, which is divided into steps. As you work through the steps, you can go forward or backward to make changes. When you exit the screen, your changes from all of the steps are saved.
 - a. **Step 1: Name assignment** – If this is a new assignment, enter a name for the assignment, and then choose **Next** to continue. If you want to change the name, choose **Step 1** at left.
 - b. **Step 2: Select an IAM policy** – Choose an IAM policy that you want to use. From this screen, you can interact with the policies as follows:
 - Choose a policy that you want to use.
 - Search for a policy name.
 - Filter the list to see all IAM policies, Amazon-managed policies, or customer-managed policies.
 - View a policy, by choosing **View policy**.

To choose a policy, choose the button beside it, and then choose **Next** to continue.

- c. **Step 3: Assign users and groups** – Choose specific users or groups. Or choose to use the selected IAM policy for all users and groups.

Choose one of the following.

- For **Assign to all users and groups**, select the check box to assign the IAM policy to all Amazon Quick Suite users and groups. Choosing this option assigns the policy to all current and future users and groups.
- Choose the users and groups you want to assign to this IAM policy. You can search for them by name, email address, or group name.

When you are finished selecting users and groups, choose **Next** to continue.

- d. **Step 4: Review and enable changes** – Save your changes.

Choose one of the following.

- To edit any of your choices, choose that step to edit it.

- To save this policy assignment as a draft, choose **Save as draft**. You can enable the draft later.
- To immediately enable this policy, choose **Save and enable**. This option overwrites any existing policy assignment with the same name.

Using Amazon Secrets Manager secrets instead of database credentials in Quick Suite

Intended audience: Amazon Quick Suite Administrators and Amazon Quick Suite developers

Amazon Secrets Manager is a secret storage service that you can use to protect database credentials, API keys, and other secret information. Using a key helps you ensure that the secret can't be compromised by someone examining your code, because the secret isn't stored in the code. For an overview, see the [Amazon Secrets Manager User Guide](#).

Quick Suite administrators can grant Amazon Quick Suite read-only access to secrets they create in Secrets Manager. These secrets can be used in place of database credentials when creating and editing data sources using the Quick Suite API.

Quick Suite supports using secrets with data source types that support credential pair authentication. Jira and ServiceNow are not currently supported.

Note

If you use Amazon Secrets Manager with Quick Suite, you are billed for access and maintenance as described in the [Amazon Secrets Manager Pricing page](#). In your billing statement, the costs are itemized under Secrets Manager and not under Amazon Quick Suite.

Use the procedures described in the following sections to integrate Secrets Manager with Amazon Quick Suite.

Topics

- [Granting Amazon Quick Suite access to Secrets Manager and selected secrets](#)
- [Creating or updating a data source with secret credentials using the Amazon Quick Suite API](#)

- [What's in the secret](#)
- [Modify a secret](#)

Granting Amazon Quick Suite access to Secrets Manager and selected secrets

If you're an administrator and you have secrets in Secrets Manager, you can grant Amazon Quick Suite read-only access to selected secrets.

To grant Amazon Quick Suite access to Secrets Manager and selected secrets

1. In Amazon Quick Suite, choose your user icon on the upper right, and then choose **Manage Quick Suite**.
2. Choose **Security & permissions** on the left.
3. Choose **Manage in Amazon Quick Suite access to Amazon resources**.
4. In **Allow access and autodiscovery for these resources**, choose **Amazon Secrets Manager, Select secrets**.

The **Amazon Secrets Manager secrets** page opens.

5. Select the secrets that you want to grant Amazon Quick Suite read-only access to.

Secrets in your Amazon Quick Suite sign-up Region are shown automatically. To select secrets outside your home Region, choose **Secrets in Other Amazon Regions**, and then enter the Amazon Resource Names (ARNs) for those secrets.

6. When you're done, choose **Finish**.

Amazon Quick Suite creates an IAM role called `aws-quicksight-secretsmanager-role-v0` in your account. It grants users in the account read-only access to the specified secrets and looks similar to the following:

When Amazon Quick Suite users create analyses from or view dashboards that use a data source with secrets, Amazon Quick Suite assumes this Secrets Manager IAM role. For more information about secret permissions policies, see [Authentication and access control for Amazon Secrets Manager](#) in the *Amazon Secrets Manager User Guide*.

The specified secret in the Amazon Quick Suite IAM role may have an additional resource policy that denies access. For more information, see [Attach a permissions policy to a secret](#) in the *Amazon Secrets Manager User Guide*.

If you're using an Amazon managed Amazon KMS key to encrypt your secret, Amazon Quick Suite doesn't require any additional permissions setup in Secrets Manager.

If you're using a customer managed key to encrypt your secret, ensure that the Amazon Quick Suite IAM role, `aws-quicksight-secretsmanager-role-v0` has `kms:Decrypt` permissions. For more information, see [Permissions for the KMS key](#) in the *Amazon Secrets Manager User Guide*.

For more information about the types of keys used in Amazon Key Management Service, see [Customer keys and Amazon keys](#) in the *Amazon Key Management Service guide*.

Creating or updating a data source with secret credentials using the Amazon Quick Suite API

After the Amazon Quick Suite administrator has granted Amazon Quick Suite read-only access to Secrets Manager, you can create and update data sources in the API using a secret the administrator selected as credentials.

Following is an example API call to create a data source in Amazon Quick Suite. This example uses the `create-data-source` API operation. You can also use the `update-data-source` operation. For more information, see [CreateDataSource](#) and [UpdateDataSource](#) in the *Amazon Quick Suite API Reference*.

The user specified in the permissions in the following API call example can delete, view, and edit data sources for the specified MySQL data source in Amazon Quick Suite. They can also view and update the data source permissions. Instead of a Amazon Quick Suite username and password, a secret ARN is used as credentials for the data source.

```
aws quicksight create-data-source
  --aws-account-id AWSACCOUNTID \
  --data-source-id DATASOURCEID \
  --name NAME \
  --type MYSQL \
  --permissions '[{"Principal": "arn:aws-cn:quicksight:region:accountID:user/namespace/username", "Actions": ["quicksight:DeleteDataSource",
"quicksight:DescribeDataSource", "quicksight:DescribeDataSourcePermissions",
"quicksight:PassDataSource", "quicksight:UpdateDataSource",
"quicksight:UpdateDataSourcePermissions"]}]' \
  --data-source-parameters='{"MySQLParameters":{"Database": "database",
"Host": "hostURL", "Port": "port"}}' \
```

```
--credentials='{"SecretArn":"arn:aws-  
cn:secretsmanager:region:accountID:secret:secretname"}' \  
--region us-west-2
```

In this call, Amazon Quick Suite authorizes `secretsmanager:GetSecretValue` access to the secret based on the API caller's IAM policy, not the IAM service role's policy. The IAM service role acts on the account level and is used when an analysis or dashboard is viewed by a user. It cannot be used to authorize secret access when a user creates or updates the data source.

When they edit a data source in the Amazon Quick Suite UI, users can view the secret ARN for data sources that use Amazon Secrets Manager as the credential type. However, they can't edit the secret, or select a different secret. If they need to make changes, for example to the database server or port, users first need to choose **Credential pair** and enter their Amazon Quick Suite account username and password.

Secrets are automatically removed from a data source when the data source is altered in the UI. To restore the secret to the data source, use the `update-data-source` API operation.

What's in the secret

Amazon Quick Suite requires the following JSON format to access your secret:

```
{  
  "username": "username",  
  "password": "password"  
}
```

The username and password fields are required for Amazon Quick Suite to access secrets. All other fields are optional and are ignored by Amazon Quick Suite.

The JSON format may vary depending on the type of database. For more information, see [JSON structure of Amazon Secrets Manager database credential secrets](#) in the *Amazon Secrets Manager User Guide*.

Modify a secret

To modify a secret, you use Secrets Manager. After you make changes to a secret, the updates become available the next time Amazon Quick Suite requests access to the secret.

Default access policy

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon Quick Suite administrators

In the Enterprise edition, you can configure specific permissions for the Amazon services that an Amazon Quick Suite user can access. If no such configuration occurs, Quick Suite uses a default set of permissions based on the user's settings. The current behavior is displayed in a blue information box.

To change the default resource access for all users (to use when no other permissions are configured)

1. Sign in to Amazon Quick Suite.
2. At upper right, choose the profile icon, and then choose **Manage Quick Suite**.
3. Under **Permissions**, choose **Default access policy**.
4. Choose one of the following:
 - Allow access to all Amazon data and resources to all users and groups.
 - Deny access to all Amazon data and resources to all users and groups.
5. Choose **Update**.

IAM policy assignments

IAM policy assignments allow you to control access to Amazon resources for users and groups within Amazon Quick Suite. By assigning IAM policies, you can define and manage permissions for various Amazon services and resources that Quick Suite may need to access. Assignments can be enabled, disabled, or deleted at any time.

To add a new IAM policy assignment:

1. Open the [Quick Suite console](#).
2. Choose the user icon at the top right, and then choose **Manage Quick Suite**.

3. Under the **Permissions** section, choose **IAM policy assignments**.
4. Choose **Add new assignment**.
5. Enter a name for this assignment and choose **Next**.
6. Select an IAM policy by either searching for it directly or scrolling through the displayed list, and then choose **Next**.
7. Assign users or groups to the IAM policy by either searching by name, email address, or group name, or scrolling through the displayed list and selecting the appropriate users or groups. You can also choose to assign the policy to all users and groups by selecting the checkbox; this will apply the policy to all current and future users and groups.
8. Choose **Next**.
9. Review your selections and choose **Save and enable**.
10. You can enable or disable the assignment at any time by selecting the checkbox next to the assignment name and choosing **Enable** or **Disable**. You can view the status of the policy assignment under the **Status** column.

To edit an IAM policy assignment:

1. Open the [Quick Suite console](#).
2. Choose the user icon at the top right, and then choose **Manage Quick Suite**.
3. Under the **Permissions** section, choose **IAM policy assignments**.
4. Choose the edit (pencil) icon.
5. Review the sections and choose **Edit** to make any changes.
6. Choose **Save and enable** to apply your updates.

To delete an IAM policy assignment:

1. Open the [Quick Suite console](#).
2. Choose the user icon at the top right, and then choose **Manage Quick Suite**.
3. Under the **Permissions** section, choose **IAM policy assignments**.
4. Choose the delete (trash) icon.
5. A confirmation modal will appear to ensure you want to proceed with the deletion. Choose **Delete** to delete the IAM policy assignment.

Extension access

Extension access connects your Amazon Quick Suite account to third party applications like Microsoft Teams, Microsoft Outlook, Microsoft Word, and Slack. Think of it as giving permission for Amazon Quick Suite to work inside the apps your team already uses, helping them accomplish their work, assisted by generative AI, and without context switching.

As an admin, you configure extension access. Then, either an admin or an author must install the extension in the Amazon Quick Suite console to make the extension available to end users. This default extension is available for use to all users on eligible subscriptions after setup. Installed extensions are available under **Connections > Extensions** in the Amazon Quick Suite menu for further configuration and management.

If you're an admin user of Amazon Quick Suite, you configure Amazon Quick Suite access to extensions, and can also deploy extensions for end user use. As an admin, you can also allow authors to edit and deploy extensions after you configure access.

User roles and extension permissions work as follows:

- **IAM Administrators:** Can enable/remove extension access, create extensions, assign owners, and control extension availability for the organization.
- **Authors:** Can create and install extensions using administrator-enabled access, and edit/delete extensions they own.
- **Readers (Reader Pro only):** Can use all enabled extensions but cannot create or modify them.

Here's what you need to know about extension access:

- **What it is:** A secure connection between your Amazon Quick Suite instance and your organization's account in third-party apps (like your Microsoft 365 or Slack workspace).
- **Why it matters:** It gives your Amazon Quick Suite instance permission to work within your organization's specific tenant or workspace in these apps.
- **Who sets it up:** Only administrators with IAM credentials can configure extension access and deploy extensions (or give authors permission to edit and deploy extensions) - it's a privileged operation that must be done before your team can start using extensions. Note that IAM admin users can access the **Extension access** pages from the **Manage Quick Suite** menu, but to configure a extension they may need to re-log in to the system and select **Connections > Extensions** from the Amazon Quick Suite menu.

Topics

- [Browser extension](#)
- [Amazon Quick Suite Microsoft Outlook extension](#)
- [Amazon Quick Suite Slack extension](#)
- [Amazon Quick Suite Microsoft Word extension](#)
- [Amazon Quick Suite Microsoft Teams extension](#)

Browser extension

The Amazon Quick Suite browser extension is available within Amazon Quick Suite to all eligible users by default and requires no administrative setup for user access. Users can find the browser extension installation link under **Connections > Extensions**. The following procedures are for IT administrators who want to automatically deploy or disable the Amazon Quick Suite browser extension across all browsers in their organization using managed device policies.

Topics

- [Deploy the browser extension](#)
- [Disable the browser extension](#)

Deploy the browser extension

Note

The following deployment guide applies only when you have a managed device fleet with managed browsers.

To deploy the browser extension to all users in your organization, you can define an **ExtensionSettings** policy and distribute this policy to your managed devices.

Topics

- [Install using ExtensionSettings policy](#)
- [Distribute ExtensionSettings policy](#)

Install using ExtensionSettings policy

The **ExtensionSettings** policy is an administrative setting for Chrome, Edge, and Firefox that allows you to manage the installation, permissions, and runtime behavior of specific extensions on managed devices. This policy provides granular control by letting administrators specify custom configurations, such as setting installation modes like force install and restricting or allowing specific site access and permissions for individual extensions or all extensions.

The format of the **ExtensionSettings** policy depends on the operating system where you want to distribute this policy. Windows, Mac, and Linux are supported.

To set the install mode in the **ExtensionSettings** policy, you need the extension identifier and the store URL:

- **Chrome/Edge**

- ID: innkphffipcmiflfibbeghfnkifiokgo
- URL: <https://clients2.google.com/service/update2/crx>

- **Firefox**

- ID: quick-browser-extension@amazon.com
- URL: <https://addons.mozilla.org/firefox/downloads/latest/amazon-quick/latest.xpi>

Distribute ExtensionSettings policy

Distributing the **ExtensionSettings** policy depends on the device and browser management solution used by your organization. You can use managed solutions like **Chrome Enterprise Core** or **Microsoft Edge** management service, or distribute the **ExtensionSettings** policy using **Group Policy**, which differs per operating system.

Disable the browser extension

Note

The following deployment guide applies only when you have a managed device fleet with managed browsers.

To disable the browser extension for all users in your organization, you can define an **ExtensionSettings** policy and distribute this policy to your managed devices.

Topics

- [Disable using ExtensionSettings policy](#)
- [Distribute ExtensionSettings policy](#)

Disable using ExtensionSettings policy

The **ExtensionSettings** policy is an administrative setting for Chrome, Edge, and Firefox that allows you to manage the installation, permissions, and runtime behavior of specific extensions on managed devices. This policy provides granular control by letting administrators specify custom configurations, such as disabling individual extensions.

The format of the **ExtensionSettings** policy depends on the operating system where you want to distribute this policy. See the documentation for the **ExtensionSettings** policy to determine the format. Windows, Mac, and Linux are supported.

To block the use of the extension in the **ExtensionSettings** policy, you need the extension identifier:

- **Chrome/Edge**

ID: innkphffipcmiflfibbeghfnkifiokgo

- **Firefox**

ID: quick-browser-extension@amazon.com

Distribute ExtensionSettings policy

Distributing the **ExtensionSettings** policy depends on the device and browser management solution used by your organization. You can use managed solutions like **Chrome Enterprise Core** or **Microsoft Edge** management service, or distribute the **ExtensionSettings** policy using **Group Policy**, which differs per operating system.

Amazon Quick Suite Microsoft Outlook extension

The Quick Suite extension for Microsoft Outlook integrates AI-powered assistance directly into your email workflows. Users can access Quick Suite capabilities within Outlook to enhance their email communication and productivity without switching between applications.

The Outlook extension enables users to:

- Summarize emails for quick understanding.
- Draft contextual responses with AI assistance.
- Use actions from action connectors configured in Amazon Quick Suite.
- Access any company knowledge sources added to your Amazon Quick Suite instance from within Microsoft Outlook.

Important

- The Amazon Quick Suite customer integrating Microsoft Outlook must have a paid Microsoft Outlook organization with an *M365 subscription for their organization*.
- Amazon Quick Suite does not use your user data for service improvement or for training its underlying large language models (LLMs).

Topics

- [Prerequisites for Microsoft Outlook extension](#)
- [Configure Microsoft Outlook extension access](#)

Prerequisites for Microsoft Outlook extension

Before configuring access to the Amazon Quick Suite extension to Microsoft Outlook, administrators must complete the following steps:

1. Have a Microsoft 365 Business subscription and be a Global Admin or have administrative permissions (specifically `AppCatalog.ReadWrite.All`).
2. Have a Amazon Quick Suite instance.

3. Your Microsoft 365 tenant ID. You can find this by going to the Azure portal > **Azure Active Directory** > **Properties**, or by using PowerShell. For detailed steps, see [How to find your tenant ID - Microsoft Entra](#) in the Microsoft Learn portal.

Note

The Microsoft Outlook extension isn't supported if you're using IAM Identity Center for end user management in Amazon Quick Suite.

Configure Microsoft Outlook extension access

As an administrator, you must allow the Amazon Quick Suite Microsoft Outlook extension to connect to your Amazon Quick Suite application environment. You can use the Amazon Quick Suite console to manage extension access configurations.

User attribute mapping

When you configure a Microsoft Outlook extension, user identity is mapped by default using the following attributes:

- **Amazon Quick Suite user attribute** - Email address is used to map Amazon Quick Suite users to their corresponding Microsoft 365 accounts. The system uses the email address to establish the connection between user identities.
- **M365 Office add-in user attribute** - User Principal Name (UPN) is used to match against Microsoft 365 user accounts. Users need to use the User Principal Name (UPN) to sign in to Microsoft 365. Email address that works on both cloud and on-premise systems.

These default mappings ensure secure and accurate user identification across both platforms without requiring additional configuration.

Note

The Microsoft Outlook extension isn't supported if you're using IAM Identity Center for end user management in Amazon Quick Suite.

Topics

- [Add Microsoft Outlook extension access](#)
- [Edit Microsoft Outlook extension access](#)
- [Delete Microsoft Outlook extension access](#)

Add Microsoft Outlook extension access

Follow these steps to create a new extension access configuration that will allow Amazon Quick Suite to integrate with your Microsoft Outlook environment.

1. Sign in to Amazon Quick Suite console.
2. In the top right, select the profile picture icon.
3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. In the top right, select **New extension access**.
6. Select **Microsoft Outlook**. Then, **Next**.
7. Configure the following fields:
 - **Name** - A name for your extension is pre-filled for you. You can edit this and enter a descriptive name for the Outlook extension (maximum 512 alphanumeric characters, hyphens allowed but no spaces).
 - **Description** (optional) - A description for your extension is pre-filled for you. You can edit this and enter a new description to provide additional context about this extension configuration (maximum 1000 characters).
 - **M365 tenant ID** - Enter your Microsoft 365 tenant identifier (must be 36 characters).
8. Select **Add** to save the new access configuration.

A success message will open up on the top right of your screen.

9. From the success message, select **View extensions** to finish installing your extension.

Note

You can also navigate to the installation screen from **Connections > Extensions** in the Amazon Quick Suite menu.

Once created, this extension access configuration enables authors and other admin in your organization to create and deploy Amazon Quick Suite extensions within your Microsoft Outlook environment.

Note

For your end users to begin using your Microsoft Outlook extension, an admin or author must finish deploying a extension after you configure extension access. Notify your authors that they can view, edit, and complete installation of this extension under **Extensions** in the left navigation once it has been shared. To learn how to do this see [Installing your Microsoft Outlook extension in the Microsoft Outlook extension author guide](#).

Edit Microsoft Outlook extension access

Use these steps to modify the configuration settings of an existing Microsoft Outlook extension access.

1. Sign in to Amazon Quick Suite console.
2. In the top right, select the profile picture icon.
3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. Select the three dot menu icon for the Microsoft Outlook extension you need to edit.
6. Select **Edit**.
7. Edit the configuration as required and select **Save** to confirm the changes.

Your changes to the Microsoft Outlook extension access configuration are saved and will take effect immediately.

Delete Microsoft Outlook extension access

Follow these steps to permanently remove a Microsoft Outlook extension access configuration. This action cannot be undone.

1. Sign in to Amazon Quick Suite console.
2. In the top right, select the profile picture icon.
3. From the drop-down menu, select **Manage account**.

4. Under **Permissions**, select **Extension access**.
5. Select the three dot menu icon for the Microsoft Outlook Extension you need to delete.
6. Select **Delete**.
7. Enter the word, "confirm", and select **DELETE**.

 Note

Deleting a extension access removes access for all users in your M365 tenant and deletes all extensions created for Outlook. If delete extension access fails, the admin must switch to the author view and delete the Outlook extensions that are using the configured extension access before returning to delete the extension access.

With Microsoft Outlook extension access configured, your team can now use Amazon Quick Suite to enhance their email workflows with AI-powered assistance for summarizing, drafting responses, and accessing organizational knowledge directly within Outlook.

Amazon Quick Suite Slack extension

The Quick Suite extension for Slack integrates AI-powered assistance directly into your team communication workflows. Users can access Quick Suite knowledge and capabilities within Slack channels and direct messages without switching between applications.

The Slack extension enables users to:

- Add Amazon Quick Suite as a collaborator using **@Amazon Quick Suite** mentions.
- Get conversation summaries and insights.
- Access organizational knowledge bases and documents directly from Slack channels.
- Generate notes and action items from discussions in Slack threads.
- Get help with data analysis and report generation using uploaded files.

 Important

- When Amazon Quick Suite is used in public Slack channels, responses are based on the invoking user's permissions. This may include content that other channel members aren't

authorized to access. Carefully evaluate using Amazon Quick Suite in public channels to prevent unintended exposure of sensitive information.

- Amazon Quick Suite does not use your user data for service improvement or for training its underlying large language models (LLMs).

Topics

- [Prerequisites for Slack extension](#)
- [Configure Slack extension access](#)

Prerequisites for Slack extension

Before adding the Amazon Quick Suite Slack Extension, administrators must complete the following requirements:

- Have a paid Slack workspace.
- Have admin access to your Slack workspace.
- Get started with Amazon Quick Suite.
- Your Slack workspace ID (must start with 'T' and be alphanumeric). One way to find your Slack workspace ID is by navigating to your Slack workspace and starting a chat with the Slack Developer Tools app running the `/sdt whoami` command. For more information, see [Locate your Slack URL or ID](#) in the Slack help center.

Configure Slack extension access

As an administrator, you must allow your Amazon Quick Suite Slack extension to connect to your Amazon Quick Suite application environment. You can use the Amazon Quick Suite console to manage extension access configurations.

User attribute mapping

When you configure a Slack extension, user identity is mapped by default using the following attributes:

- **Amazon Quick Suite user attribute** - Email address is used to map Amazon Quick Suite users to their corresponding Slack accounts. The system uses the email address to establish the connection between user identities.

- **Slack user attribute** - User Profile Email is used to match against Slack user accounts. This maps to the email address associated with the user's Slack profile.

These default mappings ensure secure and accurate user identification across both platforms without requiring additional configuration.

Topics

- [Add Slack extension access](#)
- [Edit Slack extension access](#)
- [Delete Slack extension access](#)

Add Slack extension access

Follow these steps to create a new extension access configuration that will allow Amazon Quick Suite to integrate with your Slack environment.

1. Sign in to Amazon Quick Suite console.
2. In the top right, select the profile picture icon.
3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. In the top right, select **New extension access**.
6. Select Slack, then, **Next**.
7. Configure the following fields:
 - **Name** - A name for your extension is pre-filled for you. You can edit this and enter a descriptive name for the Slack extension (maximum 512 alphanumeric characters, hyphens allowed but no spaces).
 - **Description** (optional) - A description for your extension is pre-filled for you. You can edit this and enter a new description to provide additional context about this extension configuration (maximum 1000 characters).
 - **Slack Workspace ID** - Enter your Slack workspace identifier. Workspace ID must start with T and be between 1 and 256 alphanumeric characters long.
8. Select **Add** to save the new access configuration.

A success message will open up on the top right of your screen.

9. From the success message, select **View extensions** to finish installing your extension.

 Note

You can also navigate to the installation screen from **Connections > Extensions** in the Amazon Quick Suite menu.

Once created, this extension access configuration enables authors and other admins in your organization to deploy Amazon Quick Suite Slack extensions in their workspace.

 Note

For your end users to begin using your Slack extension, an admin or author must finish deploying a extension after you configure extension access. Notify your authors that they can view, edit, and complete installation of this extension under **Extensions** in the left navigation once it has been shared. To learn how to do this see [Installing your Slack extension in the Slack extension author guide](#).

Edit Slack extension access

Use these steps to modify the configuration settings of an existing Slack extension access.

1. Sign in to Amazon Quick Suite console.
2. In the top right, select the profile picture icon.
3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. Select the three dot menu icon for the Slack Extension you need to edit.
6. Select **Edit**.
7. Edit the configuration as required and select **Save** to confirm the changes.

Your changes to the Slack extension access configuration are saved and will take effect immediately.

Delete Slack extension access

Follow these steps to permanently remove a Slack extension access configuration. This action cannot be undone.

1. Sign in to Amazon Quick Suite console.
2. In the top right, select the profile picture icon.
3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. Select the three dot menu icon for the Slack Extension you need to delete.
6. Select **Delete**.
7. Enter the word, "confirm", and select **DELETE**.

Note

Deleting a extension access removes access for all users in your Slack workspace and deletes all extensions created for Slack. If delete extension access fails, the admin must switch to the author view and delete the Slack extensions that are using the configured extension access before returning to delete the extension access.

With Slack extension access configured, your team can now use **@Amazon Quick Suite** mentions in channels and direct messages to access AI assistance and organizational knowledge directly within their Slack workspace.

Amazon Quick Suite Microsoft Word extension

The Quick Suite extension for Microsoft Word integrates AI-powered assistance directly into your document creation and editing workflows. Users can access Quick Suite capabilities within Word to enhance their document productivity and content quality without switching between applications.

The Word extension enables users to:

- Review, draft, and revise documents with AI assistance.
- Summarize documents and get insights.
- Simplify and improve writing quality.

- Use actions from action connectors configured in Amazon Quick Suite.
- Access any company knowledge sources added to your Amazon Quick Suite instance from within Microsoft Word.

Important

- The Amazon Quick Suite customer integrating Microsoft Word must have a paid Word organization with an *M365 subscription for their organization*.
- Amazon Quick Suite does not use user data for service improvement or for training its underlying large language models (LLMs).

Topics

- [Prerequisites for Microsoft Word extension](#)
- [Configure Microsoft Word extension access](#)

Prerequisites for Microsoft Word extension

Before adding the Amazon Quick Suite extension to Microsoft Word, administrators must complete the following steps:

1. Have a Microsoft 365 Business subscription and be a Global Admin or have administrative permissions (specifically `AppCatalog.ReadWrite.All`).
2. Have a Amazon Quick Suite instance.
3. Your Microsoft 365 tenant ID. You can find this by going to the Azure portal > **Azure Active Directory** > **Properties**, or by using PowerShell. For detailed steps, see [How to find your tenant ID - Microsoft Entra](#) in the Microsoft Learn portal.

Note

The Microsoft Word extension isn't supported if you're using IAM Identity Center for end user management in Amazon Quick Suite.

Configure Microsoft Word extension access

As an administrator, you must allow the Amazon Quick Suite Microsoft Word Extension to connect to your Amazon Quick Suite application environment and web experience. You can use the Amazon Quick Suite console to manage extension access configurations.

User attribute mapping

When you configure a Microsoft Word extension, user identity is mapped by default using the following attributes:

- **Amazon Quick Suite user attribute** - Email address is used to map Amazon Quick Suite users to their corresponding Microsoft 365 accounts. The system uses the email address to establish the connection between user identities.
- **M365 Office add-in user attribute** - User Principal Name (UPN) is used to match against Microsoft 365 user accounts. Users need to use the User Principal Name (UPN) to sign in to Microsoft 365. Email address that works on both cloud and on-premise systems.

These default mappings ensure secure and accurate user identification across both platforms without requiring additional configuration.

Topics

- [Add Microsoft Word extension access](#)
- [Edit Microsoft Word extension access](#)
- [Delete Microsoft Word extension access](#)

Add Microsoft Word extension access

Follow these steps to create a new extension access configuration that will allow Amazon Quick Suite to integrate with your Microsoft Word environment.

1. Sign in to Amazon Quick Suite console.
2. In the top right, select the profile picture icon.
3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. In the top right, select **New extension access**.
6. Select **Microsoft Word**. Then, select **Next**.

7. Configure the following fields:

- **Name** - A name for your extension is pre-filled for you. You can edit this and enter a descriptive name for the Word extension (maximum 512 alphanumeric characters, hyphens allowed but no spaces).
- **Description** (optional) - A description for your extension is pre-filled for you. You can edit this and enter a new description to provide additional context about this extension configuration (maximum 1000 characters).
- **M365 tenant ID** - Enter your Microsoft 365 tenant identifier (must be 36 characters).

8. Select **Add** to save the new access configuration.

A success message will open up on the top right of your screen.

9. From the success message, select **View extensions** to finish installing your extension.

Note

You can also navigate to the installation screen from **Connections > Extensions** in the Amazon Quick Suite menu.

Once created, this extension access configuration enables authors and other admin in your organization to create and deploy Amazon Quick Suite extensions within your Microsoft Word environment.

Note

For your end users to begin using your Microsoft Word extension, an admin or author must finish deploying a extension after you configure extension access. Notify your authors that they can view, edit, and complete installation of this extension under **Extensions** in the left navigation once it has been shared. To learn how to do this see [Installing your Microsoft Word extension in the Microsoft Word extension author guide](#).

Edit Microsoft Word extension access

Use these steps to modify the configuration settings of an existing Microsoft Word extension access.

1. Sign in to Amazon Quick Suite console.
2. In the top right, select the profile picture icon.
3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. Select the three dot menu icon for the **Microsoft Word** extension you need to edit.
6. Select **Edit**.
7. Edit the configuration as required and select **Save** to confirm the changes.

Your changes to the Microsoft Word extension access configuration are saved and will take effect immediately.

Delete Microsoft Word extension access

Follow these steps to permanently remove a Microsoft Word extension access configuration. This action cannot be undone.

Note

Deleting a extension access removes access for all users in your M365 tenant and deletes all extensions created for Word.

1. Sign in to Amazon Quick Suite console.
2. In the top right, select the profile picture icon.
3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. Select the three dot menu icon for the Microsoft Word Extension you need to delete.
6. Select **Delete**.
7. Enter the word, "confirm", and select **DELETE**.

Note

If the process to delete extension access fails, the admin must switch to the author view and delete the Word extensions that are using the configured extension access before returning to delete the extension access.

With Microsoft Word extension access configured, your team can now use Amazon Quick Suite to enhance their document creation workflows with AI assistance for reviewing, drafting, revising, and accessing organizational knowledge directly within Word.

Amazon Quick Suite Microsoft Teams extension

The Quick Suite extension for Microsoft Teams integrates AI-powered assistance directly into your team communication workflows. Users can mention @Amazon Quick Suite in conversations to access company knowledge, use configured action connectors, and get contextual assistance without leaving their Teams environment.

The Teams extension enables users to:

- Mention **@Amazon Quick Suite** in conversations in Teams channels to add it as a collaborator.
- Use actions from action connectors configured in Amazon Quick Suite.
- Access any company knowledge sources added to your Amazon Quick Suite instance from within Microsoft Teams.

Important

- The Amazon Quick Suite customer integrating Microsoft Teams must have a paid Microsoft Teams organization with an *M365 subscription for their organization*.
- Amazon Quick Suite does not use your user data for service improvement or for training its underlying large language models (LLMs).

Topics

- [Prerequisites for Microsoft Teams extension](#)
- [Configure Microsoft Teams extension access](#)

Prerequisites for Microsoft Teams extension

Before configuring access for the Amazon Quick Suite extension to Microsoft Teams, administrators must complete the following steps:

1. Have a Microsoft 365 Business subscription and be a Global Admin or have administrative permissions (specifically `AppCatalog.ReadWrite.All`).
2. Have a Amazon Quick Suite instance.
3. Your Microsoft 365 tenant ID. You can find this by going to the Azure portal > **Azure Active Directory** > **Properties**, or by using PowerShell. For detailed steps, see [How to find your tenant ID - Microsoft Entra](#) in the Microsoft Learn portal.

Note

The Microsoft Teams extension isn't supported if you're using IAM Identity Center for end user management in Amazon Quick Suite.

Configure Microsoft Teams extension access

As an administrator, you must allow the Amazon Quick Suite Microsoft Teams to connect to your Amazon Quick Suite application environment. You can use the Amazon Quick Suite console to manage extension access configurations.

User attribute mapping

When you configure a Microsoft Teams extension, user identity is mapped by default using the following attributes:

- **Amazon Quick Suite user attribute** - Email address is used to map Amazon Quick Suite users to their corresponding Microsoft 365 accounts. The system uses the email address to establish the connection between user identities.
- **M365 Office add-in user attribute** - User Principal Name (UPN) is used to match against Microsoft 365 user accounts. Users need to use the User Principal Name (UPN) to sign in to Microsoft 365. Email address that works on both cloud and on-premise systems.

These default mappings ensure secure and accurate user identification across both platforms without requiring additional configuration.

Topics

- [Add Microsoft Teams extension access](#)
- [Edit Microsoft Teams extension access](#)
- [Delete Microsoft Teams extension access](#)

Add Microsoft Teams extension access

Follow these steps to create a new extension access configuration that will allow Amazon Quick Suite to integrate with your Microsoft Teams environment.

1. Sign in to Amazon Quick Suite console.
2. In the top right, select the profile picture icon.
3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. In the top right, select **New extension access**.
6. Select **Microsoft Teams**, then, **Next**.
7. Configure the following fields:
 - **Name** - A name for your extension is pre-filled for you. You can edit this and enter a descriptive name for the Teams extension (maximum 512 alphanumeric characters, hyphens allowed but no spaces).
 - **Description** (optional) - A description for your extension is pre-filled for you. You can edit this and enter a new description to provide additional context about this extension configuration (maximum 1000 characters).
 - **M365 tenant ID** - Enter your Microsoft 365 tenant identifier (must be 36 characters).
8. Select **Add** to save the new access configuration.

A success message will open up on the top right of your screen.

9. From the success message, select **View extensions** to finish installing your extension.

Note

You can also navigate to the installation screen from **Connections > Extensions** in the Amazon Quick Suite menu.

Once created, this extension access configuration enables authors and other admins in your organization to deploy your Microsoft Teams Amazon Quick Suite extension within your Microsoft Teams environment.

 Note

For your end users to begin using your Microsoft Teams extension, an admin or author must finish deploying a extension after you configure extension access. Notify your authors that they can view, edit, and complete installation of this extension under **Extensions** in the left navigation once it has been shared. To learn how to do this see [Installing your Microsoft Teams extension in the Microsoft Teams extension author guide](#).

Edit Microsoft Teams extension access

Use these steps to modify the configuration settings of an existing Microsoft Teams extension access.

1. Sign in to Amazon Quick Suite console.
2. In the top right, select the profile picture icon.
3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. Select the three dot menu icon for the Microsoft Teams extension you need to edit.
6. Select **Edit**.
7. Edit the configuration as required and select **Save** to confirm the changes.

Your changes to the Microsoft Teams extension access configuration are saved and will take effect immediately.

Delete Microsoft Teams extension access

Follow these steps to permanently remove a Microsoft Teams extension access configuration. This action cannot be undone.

1. Sign in to Amazon Quick Suite console.
2. In the top right, select the profile picture icon.

3. From the drop-down menu, select **Manage account**.
4. Under **Permissions**, select **Extension access**.
5. Select the three dot menu icon for the Microsoft Teams Extension you need to delete.
6. Select **Delete**.
7. Enter the word, "confirm", and select **DELETE**.

 Note

Deleting a extension access removes access for all users in your M365 tenant and deletes all extensions created for Teams. If delete extension access fails, the admin must switch to the author view and delete the Teams extensions that are using the configured extension access before returning to delete the extension access.

With Microsoft Teams extension access configured, your team can now use **@Quick Suite** mentions in conversations to access AI assistance, company knowledge, and action connectors directly within their Teams environment.

Creating a custom permissions profile in Amazon Quick Suite

Applies to: Enterprise Edition

Intended audience: Administrators and Amazon Quick Suite developers

In Enterprise edition, you can restrict the functionality that people can access in Amazon Quick Suite. You can configure custom permissions at the account, role (admin, author, reader), and user levels for all identity types in Quick Suite. *User level custom permissions* override a role's existing default or custom role level permissions for the specified user. *User level custom permissions* and *role level custom permissions* override *account level custom permissions*.

The following limitations apply to custom permissions.

- You can't grant permissions that are above a user's default role. For example, if a user has reader access, you can't grant permissions for that user to edit dashboards.

- To customize user or role permissions, you need to be a Amazon Quick Suite administrator with the following IAM permissions:
 - `quicksight:CreateCustomPermissions`
 - `quicksight>DeleteCustomPermissions`
 - `quicksight:DescribeCustomPermissions`
 - `quicksight:ListCustomPermissions`
 - `quicksight:UpdateCustomPermissions`
 - `quicksight:DescribeAccountCustomPermissions`
 - `quicksight:UpdateAccountCustomPermissions`
 - `quicksight>DeleteAccountCustomPermissions`

You can create custom permission profiles to restrict access to any combination of the following features. Parent capabilities can be used to restrict access to an entire asset's feature sets. When parent capabilities are disabled, all associated child features will also be disabled.

Features with no parent capabilities cannot be turned off with this mechanism. Instead, they must be restricted as individual features.

Quick Suite parent capabilities

Parent capability	Functionality
Analyses	Restricts all Analysis-related features
Dashboards	Restricts all Dashboards-related features
Actions	Restricts all Actions-related features
Automate	Restricts all Automation-related features
Chat Agents	Restricts all Chat Agent-related features
Extensions	Restricts all Extensions-related features
Flows	Restricts all Flows-related features
Knowledge Base	Restricts all Knowledge Base-related features

Parent capability	Functionality
Research	Restricts all Research-related features
Spaces	Restricts all Spaces-related features

Quick Suite features

Feature	Amazon Quick Suite behavior	Parent capability
Create Chat Agents	- Can't view or access any chat agents - Agent library and navigation are hidden - Can still access and create other Quick Suite resources , such as creating spaces for file sharing with teams or flows for structured interactions (as long as those capabilities are not also restricted)	Chat Agents
Allow creators to share without approval	Flows cannot be shared by creators without approval	Flows
Use Bedrock models for output refinement	Restricts usage of Bedrock models	Flows
Enable UI agent to perform browser tasks	Restricts flows UI agent from performing browser tasks	Flows
Use internet to enhance results	Restricts usage of web-based search in Chat Agents and Research	--

Feature	Amazon Quick Suite behavior	Parent capability
Sharing analyses	Access to Share option on the File menu is disabled for analyses	Analyses
Adding or running anomaly detection	- Access to the Add anomaly to sheet option on the Insights menu is disabled for analyses - Access to the Anomaly option on the Objects menu is disabled for analyses - Users will not be able to add anomaly detection to sheets	Analyses
Print Sheet	- Access to the Print option on the File menu is disabled for analyses - Access to the Print option on the Export menu is disabled for dashboards - Users will not be able to print sheets	--
Export sheet to PDF	- Access to the Export to PDF option on the File menu is disabled for analyses - Access to the Generate PDF option on the Export menu is disabled for dashboards - Users will not be able to export sheets to a PDF file	--

Feature	Amazon Quick Suite behavior	Parent capability
Creating or updating themes	- Access to the Themes option on the Edit menu is disabled for analyses - Users will not be able to create custom themes - Users will not be able to edit or update existing themes	--
Sharing dashboards	Access to the share icon on the navigation menu is disabled for dashboards	Dashboard
Export visual to CSV	- Access to the Export to CSV option on the three-dot menu for each visual is disabled for both analyses and dashboards - Access to the Export Visual to CSV option on the Objects menu is disabled for analyses - Users will not be able to export visuals to a CSV file	--

Feature	Amazon Quick Suite behavior	Parent capability
Export visual to Excel	- Access to the Export to Excel option on the three-dot menu for each table is disabled for both analyses and dashboards - Access to the Export Table to Excel option on the Objects menu is disabled for analyses - Users will not be able to export tables to an Excel file	--
Creating or updating all datasets	Access to creating or updating all datasets will be disabled	--
Creating or updating only SPICE datasets	Access to creating or updating SPICE datasets will be disabled	--
Sharing datasets	Access to sharing datasets will be disabled	--
Viewing account SPICE capacity	Restricts retrieving the account's SPICE capacity	--
Creating or updating all data sources	Access to creating or updating all data sources will be disabled	--
Sharing data sources	Access to sharing data sources will be disabled	--

Feature	Amazon Quick Suite behavior	Parent capability
Creating shared folders	Restricts creating shared folders	--
Renaming shared folders	Restricts renaming shared folders	--
Creating or updating scheduled email reports	- Access to the Schedules option on the Schedules menu is disabled for dashboards - Access to the Recent snapshots option on the Schedules menu is disabled for dashboards - Users will not be able to create or update scheduled email reports	--
Subscribing to scheduled email reports	Users will not be able to subscribe to scheduled email reports	Dashboard
CSV attachments in scheduled email reports	- Access to the CSV option in the Content section of the Schedules menu is disabled for dashboards - Users will not be able to attach CSV files in scheduled email reports	--

Feature	Amazon Quick Suite behavior	Parent capability
Excel attachments in scheduled email reports	- Access to the Excel option in the Content section of the Schedules menu is disabled for dashboards - Users will not be able to attach Excel files in scheduled email reports	--
PDF attachments in scheduled email reports	- Access to the PDF option in the Content section of the Schedules menu is disabled for dashboards - Users will not be able to attach PDF files in scheduled email reports	--
Content within scheduled email reports	- Users will receive the content in scheduled email reports only as downloadable links that are gated behind login Include sheet in email body and File attachment options in the Schedules menu will be disregarded - Images will not be included in scheduled email reports	--
Creating or updating threshold alerts	- Access to the Alerts menu is disabled for dashboards - Users will not be able to create or update threshold alerts	--

Custom permissions profiles can be created for Amazon Quick Suite accounts that are integrated with IAM Identity Center, Active Directory, or for Amazon Quick Suite accounts that have Amazon Quick Suite managed users. The identity type that an Amazon Quick Suite account uses determines the way an Amazon Quick Suite admin configures a custom permissions profile.

The following procedure shows you how to control access to Amazon Quick Suite capabilities and respective features.

To control access to Amazon Quick Suite capabilities and features

1. Log in to the Amazon Quick Suite console.
2. Select **Manage Quick Suite**.
3. From the admin console left navigation menu, select **Permissions**, and then select **Custom permissions**.
4. In **Custom permissions**, from **Profiles**, select **New profile** or choose to edit the default profile.
5. In **New profile**, do the following:
 - In **Restrict capabilities** – Choose whether to allow specific capabilities for your system by checking or unchecking the appropriate options.
 - In **Restrict features** – Choose whether to allow specific features by checking or unchecking the appropriate options.

Creating a custom permissions profile for a Amazon Quick Suite account that is integrated with IAM Identity Center or Active Directory

Amazon Quick Suite account admins can use the following procedure to create a custom permissions profile for a Amazon Quick Suite account that is integrated with IAM Identity Center or Active Directory.

To create a custom permissions profile for a Amazon Quick Suite account that is integrated with IAM Identity Center or Active Directory

1. Sign in to the [Amazon Management Console](#).
2. Open Amazon Quick Suite.
3. The Amazon Quick Suite Admin console opens. Choose **Custom Permissions**.
4. The **Manage custom permissions** page opens. Choose one of the following options.

- To create a new custom permissions profile, choose **Create**.
 - To edit or view an existing custom permissions profile, choose the ellipsis (three dots) next to the profile that you want, and then choose **Edit**.
5. If you want to create or update a custom permissions profile, make selections for the following items.
 - For **Name**, enter a name for the custom permissions profile.
 - For **Restrictions**, choose the options that you want to deny. Any option that you don't choose is allowed. For example, if you don't want users to create or update data sources, but you want them to be able to do everything else, choose only **Creating or updating data sources**.
 6. Choose **Create** or **Update** to confirm your choices. To go back without making any changes, choose **Back**.
 7. Once you are done making changes, record the name of the custom permissions profile. Provide the name of the custom permissions profile to API users so that they can apply the custom permissions profile to roles or users.

Creating a custom permissions profile for a Amazon Quick Suite account that uses Amazon Quick Suite managed users

Amazon Quick Suite account admins can use the following procedure to create a custom permissions profile for a Amazon Quick Suite account that uses Amazon Quick Suite managed users.

To create a custom permissions profile for Amazon Quick Suite managed users

1. Open the [Quick Suite console](#).
2. From any page in the Amazon Quick Suite console, choose **Manage Quick Suite** at the top right corner.

Only Amazon Quick Suite administrators have access to the **Manage Quick Suite** menu option. If you don't have access to the **Manage Quick Suite** menu, contact your Amazon Quick Suite administrator for assistance.

3. Choose **Custom permissions**. You can also choose the **Manage users** section, and then choose **Manage custom permissions**.

4. The **Manage custom permissions** page opens. Choose one of the following options.
 - To create a new custom permissions profile, choose **Create**.
 - To edit or view an existing custom permissions profile, choose the ellipsis (three dots) next to the profile that you want, and then choose **Edit**.
5. If you want to create or update a custom permissions profile, make selections for the following items.
 - For **Name**, enter a name for the custom permissions profile.
 - For **Restrictions**, choose the options that you want to deny. Any option that you don't choose is allowed. For example, if you don't want users to create or update data sources, but you want them to be able to do everything else, choose only **Creating or updating data sources**.
6. Choose **Create** or **Update** to confirm your choices. To go back without making any changes, choose **Back**.
7. Once you are done making changes, record the name of the custom permissions profile. Provide the name of the custom permissions profile to API users so that they can apply the custom permissions profile to roles or users.

After you create a custom permissions profile, use Amazon Quick Suite APIs to add or change the custom permissions profile that is assigned to a user, role, or account. Users with sufficient permissions can also use the [AWS::QuickSight::CustomPermissions](#) Amazon CloudFormation resource to manage Amazon Quick Suite custom permissions profiles. Use the following topics to learn more about managing custom permissions profiles with the Amazon Quick Suite APIs.

- [Apply a custom permissions profile to a Amazon Quick Suite role with the Amazon Quick Suite API](#)
- [Apply a custom permissions profile to a user with the Amazon Quick Suite API](#)

Apply a custom permissions profile to a Amazon Quick Suite role with the Amazon Quick Suite API

After you create a custom permissions profile, use the Amazon Quick Suite APIs to add or change the custom permissions profile that is assigned to a role.

Before you begin, you need to set up and configure the Amazon CLI. For more information about installing the Amazon CLI, see [Install or update the latest version of the Amazon CLI](#) and [Configure the Amazon CLI](#) in the Amazon Command Line Interface User guide. You also need permissions to use the Amazon Quick Suite API.

The following example calls the UpdateRoleCustomPermission API to update the custom permissions that are assigned to a role.

```
aws quicksight update-role-custom-permission \  
--role ROLE \  
--aws-account-id AWSACCOUNTID \  
--namespace default \  
--custom-permissions-name PERMISSIONNAME \  
--region REGION
```

The following example returns the custom permissions profile that is assigned to a role.

```
aws quicksight describe-role-custom-permission \  
--role ROLE \  
--aws-account-id AWSACCOUNTID \  
--namespace default \  
--region REGION
```

The following example deletes a custom permissions profile from a role.

```
aws quicksight delete-role-custom-permission \  
--role ROLE \  
--aws-account-id AWSACCOUNTID \  
--namespace default \  
--region REGION
```

Apply a custom permissions profile to a user with the Amazon Quick Suite API

The following example applies a custom permissions profile to a user.

```
aws quicksight update-user-custom-permission \  
--aws-account-id AWSACCOUNTID \  
--namespace default \  
--user-name USER_NAME \  
--custom-permissions-name myCustomPermission
```

The following example deletes a custom permissions profile from a user.

```
aws quicksight delete-user-custom-permission \  
  --aws-account-id AWSACCOUNTID \  
  --namespace default
```

The following example adds custom permissions to a new Amazon Quick Suite IAM user.

```
aws quicksight register-user \  
  --iam-arn arn:aws-cn:iam::AWSACCOUNTID:user/USER \  
  --identity-type IAM \  
  --user-role AUTHOR \  
  --custom-permissions-name custom-permissions-profile-name \  
  --email EMAIL \  
  --aws-account-id AWSACCOUNTID \  
  --namespace default \
```

You can also associate an existing IAM user with a new permissions profile. The following example updated the custom permissions profile of an existing IAM user.

```
aws quicksight update-user \  
  --user-name USERNAME \  
  --role AUTHOR \  
  --custom-permissions-name custom-permissions-profile-name \  
  --email EMAIL \  
  --aws-account-id AWSACCOUNTID \  
  --namespace default \
```

The example below removes an existing user from a permissions profile.

```
aws quicksight update-user \  
  --user-name USERNAME \  
  --role AUTHOR \  
  --unapply-custom-permissions \  
  --email EMAIL \  
  --aws-account-id AWSACCOUNTID \  
  --namespace default
```

To test the custom permissions that are applied to a role or user, log in to the user's account. When a user logs into Amazon Quick Suite, they are granted the highest privilege role that they have access to. The highest privileged role a user can be granted is Admin. The lowest privileged role

that a user can be granted is Reader. For more information about roles in Amazon Quick Suite, see [Managing user access inside Amazon Quick Suite](#).

If you assign a custom permissions profile that restricts data source sharing to the author's role, that author is no longer able to access the controls that allow data source sharing. Instead, the affected author has view-only permissions to the data source.

Apply a custom permissions profile to an account

To apply a custom permissions profile to an account

1. Open the [Quick Suite console](#).
2. From the top right, choose the profile icon.
3. Choose **Manage Quick Suite**. Only Amazon Quick Suite administrators will be able to view this page.
4. Choose **Custom permissions**. You can also choose the **Manage users** section, and then choose **Manage Custom Permissions** if your Quick Suite account uses Quick Suite managed users.
5. Locate the desired account custom permission. In the options menu under **Actions**, choose **Set as account profile**.

Apply a custom permissions profile to an account using the Quick Suite APIs

After you have created a custom permissions profile, use the Quick Suite API to add or change the custom permissions profile that is assigned to an account.

Before you begin, you will need to set up and configure the Amazon CLI. For more information about installing the Amazon CLI, see [Install or update the latest version of the Amazon CLI](#) and [Configure the Amazon CLI](#) in the Amazon Command Line Interface user guide. You also need the following IAM permissions: `quicksight:UpdateAccountPermission`, `quicksight:DescribeAccountPermission`, and `quicksight>DeleteAccountCustomPermission`.

The following example calls the `UpdateAccountPermission` API to update the custom permissions that are assigned to an account.

```
aws quicksight update-account-custom-permission \  
--aws-account-id AWSACCOUNTID \  
--custom-permissions-name PERMISSIONNAME \  

```



```
--region REGION
```

The following example returns the custom permissions profile that is assigned to an account.

```
aws quicksight describe-account-custom-permission \  
--aws-account-id AWSACCOUNTID \  
--region REGION
```

The following example unapplies a custom permissions profile from an account.

```
aws quicksight delete-account-custom-permission \  
--aws-account-id AWSACCOUNTID \  
--region REGION
```

Customizations

Using Amazon Quick Suite, you can create a customized experience for people using either the Amazon Web Services Management Console or Quick Suite consoles embedded in your application.

Currently, different options for customizing Quick Suite are available separately in the console and the Quick Suite API. Following, you can find information about the available options.

The following customization options are currently available:

- You can customize the welcome content Quick Suite provides for new users:
 - You can accept or decline the sample assets. These assets include sample datasets and analyses that are added when a person signs in for the first time.
 - You can show or hide default introductory videos. These videos include the animation that displays for new users and also the tutorial videos shown on the Amazon Quick Suite home page.
- You can create and specify a default theme.
- You can customize dashboard report emails, paginated report emails, and alert emails by editing the email template.

⚠ Important

All customizations apply only to the Amazon Web Services Region that you are using in the API or that is selected in the Amazon Quick Suite console.

To check your Region setting, you can use one of the following procedures.

To check your Amazon Web Services Region on the Amazon Quick Suite console

1. Choose your profile icon at upper right to open the menu.
2. View your current Amazon Web Services Region, listed next to a location icon.
3. (Optional) Choose another Amazon Web Services Region from the menu to change to that Region. Remember to change back after you are finished with customizations.

To check your Amazon Web Services Region using the Amazon CLI

- On the command line, enter the following command and press Enter to view the current settings.

```
aws configure list
```

To reconfigure your default Region, use the `aws configure` command.

To keep your default Region, you can add the `--region` parameter to most CLI commands.

Topics

- [Amazon Quick Suite brand customization](#)
- [Chat agent customization in Amazon Quick Suite](#)

Amazon Quick Suite brand customization

Amazon Quick Suite allows account admins to customize their application's branding and visual theme to align with their organization's guidelines. This customization includes the following visual elements to create a cohesive look and feel across all non-administrative Amazon Quick Suite console pages, schedules, alerts, and email reports.

- Logo
- Favicon
- Associated alt text for visual assets

The following list shows the different areas customizable theme colors are grouped into.

Brand colors

- **Global navigation bar** colors are applied to the topmost bar in the Amazon Quick Suite UI and include the company logo that is displayed in the standard and embedded Amazon Quick Suite consoles.
- **Application bar** colors are applied to the secondary navigation bar that contains contextual actions.

Interaction colors

- **Accent** colors are applied to interactive elements like buttons, borders, and icons.

Surface colors

- **Primary** colors are applied to high-emphasis surfaces like the homepage background and text.
- **Secondary** colors are applied to practical surfaces like borders, backgrounds, and form fields. Secondary colors are used alongside primary colors.

Status colors

- **Success** colors are applied to success messages.
- **Danger** colors are applied to error messages.
- **Warning** colors are applied to warning messages.
- **Info** colors are applied to informational messages.

Data visualization colors

- **Dimension** colors are used to identify associations between data columns that share the same role.

- **Measure** colors are used to identify metrics or measured values.

Features

- **Visualization** colors are applied to the **Visualization** icon.
- **Insight** colors are applied to the **Insight** icon.
- **Connection** colors are applied to the **Connection** icon.
- **Automation** colors are applied to the **Automation** icon.

Use the following sections to get started with brand customization in Amazon Quick Suite.

Topics

- [Permissions for Amazon Quick Suite brand customization](#)
- [Create a custom brand in Amazon Quick Suite](#)
- [Setting a default theme for Amazon Quick Suite analyses with the Amazon Quick Suite APIs](#)

Permissions for Amazon Quick Suite brand customization

To set up a brand, you must be granted an Admin role through IAM Identity Center or IAM. Admins whose roles are granted to them within Amazon Quick Suite can't create brands. To learn more about integrating your account with IAM Identity Center, see [Configure your Amazon Quick Suite account with IAM Identity Center](#). For information about admin roles and capabilities, see [Understanding Amazon Quick Suite subscriptions and roles](#).

Admin users can only manage brands that are in the same capacity Region as their Amazon Quick Suite account.

The IAM role that you use to create a brand in Amazon Quick Suite must contain `quicksight:*` or granular action permissions to manage brands in the admin console. The following granular permissions are required for admins to work with Amazon Quick Suite brands:

- `quicksight:CreateBrand`
- `quicksight:UpdateBrand`
- `quicksight:DescribeBrand`
- `quicksight:DescribeBrandPublishedVersion`
- `quicksight:UpdateBrandPublishedVersion`

- quicksight:DeleteBrand
- quicksight:ListBrands
- quicksight:UpdateBrandAssignment
- quicksight:DescribeBrandAssignment
- quicksight:DeleteBrandAssignment

After you confirm that your Admin role contains the required permissions, you can [Create a custom brand](#) in the Amazon Quick Suite admin console.

Create a custom brand in Amazon Quick Suite

Use the following procedure to create a custom brand in Amazon Quick Suite.

1. Open the [Quick Suite console](#).
2. Choose the user icon at the top right, and then choose **Manage Quick Suite**.
3. Choose **Customize application**.
4. On the **Customize application** page that opens, choose **ADD BRAND**. The **Brand settings** page opens.
5. Navigate to the **Brand Info** section.
6. For **Brand name**, enter a name for the brand. The brand name can contain up to 512 characters.
7. (Optional) For **Brand description**, enter a description for the custom brand. The brand description can contain up to 512 characters.
8. Navigate to the **Logo** section.


Logo

Primary



Alt text

Favicon



9. For **Primary**, choose the ellipsis (three dots) next to the primary icon, and then choose **Replace image**.
10. In the **Choose image** pop up that opens, perform one of the following actions:
 - a. Drag and drop image into the **Drag an image here** box.
 - b. Choose **Select a file** to select a file from your computer.
 - c. Enter a public URL or Amazon S3 URI in the text bar.

The image that you choose must be a .jpeg, .png, or .svg format and can't exceed 1MB.

When you are finished choosing an image, choose **Apply**.

11. For **Favicon**, choose the ellipsis (three dots) next to the favicon, and then choose **Replace image**.
12. In the **Choose image** pop up that opens, perform one of the following actions:
 - a. Drag and drop image into the **Drag an image here** box.
 - b. Choose **Select a file** to select a file from your computer.
 - c. Enter a public URL or Amazon S3 URI in the text bar.

The image that you choose must be a .jpeg, .png, or .svg format and can't exceed 1MB.

When you are finished choosing an image, choose **Apply**.

13. (Optional) For **Alt text**, enter alt text to display with the logo. The alt text can contain up to 512 characters.
14. To make changes to the theme colors of the brand, navigate to the **Appearance** pane on the left and choose **Theme**.
15. The **Theme settings** page appears and displays all parts of a Amazon Quick Suite theme that can be customized. The following image shows the configuration settings of the global navigation bar.



16. To change the background color of an area, navigate to the item that you want to change and choose the **Background** color swatch.

17. In the **Custom color** pop up that appears, choose a color from the color gradient or enter a hex code value in the **HEX** bar, and then choose **APPLY**.
18. To change the foreground color of an area, navigate to the item that you want to change and choose the **Foreground** color swatch.
19. In the **Custom color** pop up that appears, choose a color from the color gradient or enter a hex code value in the **HEX** bar, and then choose **APPLY**.
20. When you are finished configuring a custom brand, choose **PUBLISH** to publish and apply the brand customization to all Amazon Quick Suite user accounts. If you don't want to publish the brand, choose **SAVE** to save the brand for later.

When you finish creating a brand in Amazon Quick Suite, the new brand appears in the brands table on the **Customize application** page of the Quick Suite admin console. The **Status** column of the brands table indicates which brand is currently published to the Quick Suite account. To make changes to a custom brand, locate the brand that you want to change in the brands table, choose the ellipsis (three dots) icon in the **Actions** column, and then choose **Publish**, **Edit**, or **Delete**.

Once you publish a brand, it can take up to 10 minutes for the new brand to propagate across all users.

Setting a default theme for Amazon Quick Suite analyses with the Amazon Quick Suite APIs

To set a default theme by using the API

1. Identify the custom theme that you want to use as the default, and locate its theme ID. If you want to use one of the QuickSight starter themes, skip this step.

To get the theme ID of a custom theme, use the [ListThemes](#) API operation for the Region where the theme is. Make sure that the theme is in the same Region with the users or groups that need to use it.

The following example shows a shell script that uses the `list-themes` command in the Amazon CLI. It sets the Amazon account ID and the Amazon Web Services Region as variables. If you previously used `aws configure` to set a default Region, adding the `--region` variable to your command overrides your default setting.

```
#declare variables
awsacct1='111122223333'
```

```
region='us-west-2'

aws quicksight list-themes \
--region $region \
--aws-account-id $awsacct1 \
--type 'CUSTOM'
```

2. Use the [ListUsers](#) or [ListGroups](#) API operation to collect the Amazon Resource Names (ARNs) for users or groups that need to use the theme as a default. You need only the top-level ARN. If all your users are part of the same group, use the group ARN.

For more information on Amazon Quick Suite ARNs, see [ARN formats](#) in the *Quick Suite API Reference*.

3. If you're using a custom theme, grant access to the theme for the ARNs that you collected in the previous step. If you're using a starter theme, skip this step because all users have access to starter themes.

The following example shows a shell script that uses the [update-theme-permissions](#) command. The grant-permissions parameter is shown using shorthand syntax. You can use JSON or YAML instead. For more information, see [Specifying parameter values](#) in the *Amazon Command Line Interface User Guide*.

```
#declare variables
awsacct1='111122223333'
namespace='default'
region='us-west-2'
theme-id='bdb844d0-0fe9-4d9d-b520-0fe602d93639' #Find this with list-themes

aws quicksight update-theme-permissions \
#Specify region if necessary: --region $region \
--aws-account-id $awsacct1 \
--theme-id $theme-id \
--grant-permissions Principal="arn:aws-
cn:quicksight:$region:$awsacct1:group/$namespace/
QuickSight_Group_Name",Actions="quicksight:DescribeTheme","quicksight:ListThemeVersions", "q
```

4. Assign the theme as the default for the same ARN or ARNs.

```
#declare variables
awsacct1='111122223333'
namespace='default'
```



```
region='us-west-2'
theme-id='bdb844d0-0fe9-4d9d-b520-0fe602d93639'

aws quicksight create-account-customization \
#Specify region if necessary: --region $region \
--aws-account-id $awsacct1 \
--namespace $namespace \
--account-customization DefaultTheme="arn:aws-
cn:quicksight:$region:$awsacct1:theme/$theme-id"
```

Currently, there are three starter themes: Classic, Midnight, and Seaside. Their ARNs are the capitalized spelling of their theme name. If you are using a starter theme instead of a custom theme, use one of the following theme ARNs:

- `arn:aws-cn:quicksight::aws:theme/CLASSIC`
- `arn:aws-cn:quicksight::aws:theme/MIDNIGHT`
- `arn:aws-cn:quicksight::aws:theme/SEASIDE`
- `arn:aws-cn:quicksight::aws:theme/RAINIER`

Chat agent customization in Amazon Quick Suite

When a user account subscribes to Quick Suite, the service automatically creates a system default chat agent. This agent powers the default chat experience, allowing Amazon Quick Suite users to leverage all chat functionalities (such as uploading files to chat, using large language model (LLM) parametric knowledge, and answering from their enterprise data) out of the box. The system default chat agent can be updated in the Amazon Quick Suite agent section by a select set of users who are designated as owners for the agent by the admin.

All chat agents (including system and custom) as well as flows are also equipped with guardrails and safety controls to ensure responsible use. Any agent or flow that you chat with will use these default guardrails powering the chat interactions:

- **Prompt leak protection** – Automatically enabled to prevent prompt injection and other LLM-breaking attacks.
- **Prompt safety** – Protects against common security threats like malicious instructions, instructions to ignore guardrails, and others.

- **Blocked words and phrases** – Protects against inappropriate content including insults, hate speech, sexual content, violence, and misconduct, for both chat requests and responses.

As an admin, you can define blocked phrases for all Amazon Quick Suite chat agents. If you do, Amazon Quick Suite ensures that chat agent and flows responses across your Amazon Quick Suite instance don't include these words or phrases. No blocked words or phrases are assigned to your chat agent or flows by default. You can choose up to 50 words or phrases to block.

Note

Admins can also control permissions for whether users can create and use chat agents and flows. For instructions on how to do that, see [Custom permissions](#).

The following sections outline how to edit the system default agent and add blocked words for chat to influence all agent responses and flows.

Topics

- [Grant user permissions to edit system default chat agent](#)
- [Edit system chat agent settings](#)
- [Adding blocked words and phrases for chat agents](#)
- [Edit blocked words and phrases for chat agents and flows](#)

Grant user permissions to edit system default chat agent

The system default chat agent can be edited by users the admin designates as owners to this agent in the admin console. The following procedure shows you how to grant admin permissions to a user so that they can edit the system default agent.

To grant permissions to edit system default agent

1. Log in to the Amazon Quick Suite console.
2. Select **Manage Quick Suite** from the Amazon Quick Suite admin console.
3. From the admin console left navigation menu, select **Account**, and then select **Manage assets**.
4. In **Manage assets**, select **Chat agents**.

5. Select **My Assistant** from the list of chat agents shown. Then, from the **Actions** menu, select **Share**.
6. In the **Share assets** modal, use the **User or Group** search bar to find the user or group you want to designate as owners for the system agent.
7. Then, from **Permissions**, choose **Owner**. Then, select **Share**.

Edit system chat agent settings

To customize your system chat agent, users designated as owners need to login to Amazon Quick Suite and visit the agent library page. Admins can directly access this page by clicking on the **Go to agent** link in admin console under **Chat agent customization**. Once in Amazon Quick Suite, follow these steps to edit your system default chat agent's properties. You can preview and test how your system agent works as you configure it.

To update system chat agent settings

1. Log in to the Amazon Quick Suite console.
2. Select **Manage Quick Suite** from the Amazon Quick Suite admin console.
3. From the admin console left navigation menu, select **Customization**, and then select **Chat agent customization**.
4. In **Chat agent customization**, in **System chat agent**, select **Go to chat agent**.
5. In the **My Assistant** page, in **Configure chat agent**, customize the following sections:
 - a. In **Chat agent persona**, configure your chat agent's personality, identity, tone, and response style. For detailed information about agent customization options, see [Working with chat agents](#) in the *Amazon Quick Suite User Guide*.

Note

This agent powers default chat interactions across all users of this account with chat feature access. Ensure that your instructions (e.g. identity and response style) work well for all users of this account.

- b. In **Reference documents**, upload files that remain active in your chat agent's memory to guide all interactions. For detailed information about file upload options, see [Working with chat agents](#) in the *Amazon Quick Suite User Guide*.

Note

Since this agent is meant for broad use, this is a place to upload enterprise-level response templates and guides.

The system default agent isn't linked to specific knowledge sources, actions, or spaces to ensure it works for all users regardless of their access permissions. This setting can't be changed.

6. In **Customization**, configure details to help recognize the chat agent. For detailed information about customization options, see [Working with chat agents](#) in the *Amazon Quick Suite User Guide*.

Adding blocked words and phrases for chat agents

Default guardrails and admin provided blocked words serve as general settings that all chat agents and flows consider when the user chats with them. Admin configured blocked words are filtered out from responses in both chat agents and flows within your Amazon Quick Suite instance.

To learn more about chat agents, see [Working with chat agents](#) in the *Amazon Quick Suite User Guide*.

To assign blocked words and phrases for all chat agents

1. Log in to the Amazon Quick Suite console.
2. Select **Manage Quick Suite**.
3. From the left navigation menu, select **Customization**, and then select **Chat agent customization**.
4. In **Chat agent customizations**, for **Guardrails and safety controls**, do the following:
 - **Add blocked words and phrases** – Select **Add** to add blocked words and phrases. You can add up to 50 words and phrases.

Edit blocked words and phrases for chat agents and flows

To edit blocked words and phrases added to chat agents and flows, use the following procedure.

To edit blocked words and phrases for all chat agents and flows

1. Log in to the Amazon Quick Suite console.
2. Select **Manage Quick Suite**.
3. From the left navigation menu, select **Customization**, and then select **Chat agent customization**.
4. In **Chat agent customization**, for **Guardrails and safety controls**, do the following:
 - **Add blocked words and phrases** – Select **Remove** to remove existing blocked words and phrases. Or, select **Add** to add new ones. You can add up to 50 words and phrases.

Using the Amazon Quick Suite analytics dashboard

Analytics in Amazon Quick Suite provides insights into how users interact with your instance across all features and capabilities. The analytics dashboard delivers the data-driven insights you need to make informed decisions about security, performance, and user experience.

Note

The analytics dashboard is only available for IAM administrators.

The Amazon Quick Suite analytics system captures detailed usage metrics from chat conversations, space interactions, flow executions, action invocations, research activities, automations, and custom agent usage. This comprehensive data collection enables you to understand not just what users are doing, but how effectively they're leveraging Quick Suite to accomplish their goals while maintaining security and compliance standards.

You can use analytics to:

- Monitor user adoption and engagement across all Quick Suite features
- Identify potential issues and monitor system health through available analytics
- Analyze user feedback and satisfaction trends to improve the user experience
- Understand resource utilization patterns and identify your most valuable content
- Make data-driven decisions about feature deployment, training needs, and resource allocation
- Monitor security-related events such as blocked messages and failed queries

Topics

- [Analytics access and permissions](#)
- [Analytics dashboard overview](#)
- [Accessing the analytics dashboard](#)
- [Understanding analytics metrics](#)
- [Analytics limitations](#)

Analytics access and permissions

Access to analytics in Amazon Quick Suite is controlled by user roles and permissions to ensure data security and privacy. Currently, only IAM administrators have access to the analytics dashboard.

To access the analytics dashboard, you need the following permission:
`quicksight:QuickSuiteUsageMetrics`

User type	Analytics capabilities
IAM administrators	• View all analytics data for the entire instance • Access detailed query and response data • Monitor team adoption and engagement metrics

Analytics dashboard overview

The Amazon Quick Suite analytics dashboard is organized into several key sections that provide different perspectives on your instance usage and security events:

Dashboard filters and controls

The analytics dashboard provides filtering options to focus on specific data:

- **Agent filter** – Select from two available options: "All Agents" to see aggregate data across all agents, or "My Assistant (System Agent)" to view system agent metrics
- **Date range picker** – Choose different time periods to analyze trends and patterns

Active user metrics

Track how users interact with your Amazon Quick Suite instance over time.

This section displays user activity metrics including:

- **Active users** – Daily, weekly, and monthly active user counts with trend comparisons
- **Usage trends** – Growth patterns and engagement changes over time
- **Feature adoption** – Usage distribution across chat, spaces, flows, and other features

Use these metrics to understand user adoption patterns and identify engagement opportunities.

Chat agent and asset metrics

Monitor chat usage patterns and user feedback to understand engagement with conversational features and identify areas for improvement.

Conversation analytics

This section provides insights into chat usage including:

- **Total conversations** – Number of chat sessions initiated
- **Total queries** – Number of messages sent by users
- **Average queries per user** – Engagement depth metrics
- **Average queries per conversation** – Conversation length patterns displayed in a time-series chart showing trends over months
- **Conversations, queries, and active users chart** – Interactive line chart displaying total conversations, total queries, and active users trends over time with monthly granularity

These metrics help optimize chat experiences and identify usage trends.

User feedback analytics

Track user satisfaction to continuously improve system quality.

This section displays user feedback data including:

- **Positive feedback** – Thumbs up interactions and trends

- **Negative feedback analysis** – Thumbs down interactions with detailed reasons and breakdown by category
- **Customer feedback trend chart** – Time-series visualization showing positive vs. negative feedback patterns over months
- **Query scope selection chart** – Analytics showing how users select different query scopes, with breakdowns by "All Resources Selected", "No Resources Selected", and "Specific Resources"
- **Feedback trends** – Satisfaction changes over time with 3-month historical data
- **Thumbs down details** – View specific queries that received negative feedback, including the full conversation context and user-provided reasons
- **Thumbs down reason breakdown** – Detailed categorization of negative feedback reasons to identify common issues and improvement areas

Leverage feedback data to enhance user satisfaction and system performance.

Query issue analysis

Identify and resolve query problems to improve system performance.

The query issue analysis section helps you identify and resolve problems:

- **Issue percentage breakdown** – View the percentage of queries that failed by issue type
- **Failed query details** – See the actual query text from users, ordered by most recent, with issue type classification
- **Unanswered queries section** – Dedicated analytics for queries that received no response, helping identify system gaps and user needs
- **Full conversation context** – Click on any failed query to see the complete conversation leading up to the issue
- **Issue trends** – Track how query issues change over time to identify patterns

Use this analysis to proactively address system issues and improve user experience.

Asset creation and usage

Monitor resource utilization to optimize content and feature adoption.

This section provides insights into resource utilization including:

- **Assets created** – New spaces, agents, flows, actions, research, and automations
- **Assets used** – Active resource engagement
- **Number of users by asset type chart** – Horizontal bar chart showing user distribution across different asset types including Actions, Automations, Chat, Custom Agents, Flows, Research, and Spaces
- **Most used assets** – Most popular content and tools (spaces, agents, actions, flows, research)

Use these insights to guide content strategy and resource allocation decisions.

Accessing the analytics dashboard

You can access the analytics dashboard from the Amazon Quick Suite console if you have the appropriate permissions.

To access the analytics dashboard

1. Log in to the Amazon Quick Suite console.
2. From the left navigation menu, select **Analytics**, and then select **Amazon Quick Suite Usage Metrics**.

The analytics dashboard opens, displaying an overview of your instance usage and available analytics.

3. Use the date picker to specify the time interval for the graphs. You can choose from the following preset intervals:
 - **Last week** – This interval starts from the previous Sunday at 00:00 and goes until the following Sunday at 00:00 in the UTC time zone.
 - **This week** – This interval starts on Monday at 00:00 and goes until 00:00 today as per the timezone relevant for the primary region of the Amazon Quick Suite account.
 - **Last month** – This interval begins on the 1st day of the last month at 00:00 and ends on the first day of this month at 00:00 in the UTC time zone.
 - **This month** – This interval begins on the 1st day of this month at 00:00 and goes until 00:00 today in the UTC time zone.

Note

If you don't see the Analytics option in the navigation menu, contact your instance administrator to request access.

Understanding analytics metrics

The analytics dashboard displays various metrics that help you understand usage patterns, performance, and security events. The following sections explain key metrics in detail.

User activity metrics

User activity metrics help you understand how many people are actively using your Amazon Quick Suite instance:

- **Daily active users** – Number of unique users who sent at least one message in a given day
- **Weekly active users** – Number of unique users who sent at least one message in the current week starting from Monday
- **Monthly active users** – Number of unique users who sent at least one message in the current month starting from the 1st day of the month

Note

Active user metrics are based on chat interactions and may not reflect all types of Quick Suite usage.

Security monitoring metrics

Security monitoring metrics help you track potential security issues and policy enforcement:

- **Blocked messages** – Messages blocked by guardrails for safety, policy, or content violations
- **Failed messages** – Messages that encountered errors, with specific error codes for troubleshooting

Analytics limitations

The following limitations apply to Amazon Quick Suite analytics:

- Analytics primarily covers chat interactions and may not capture all types of Quick Suite usage
- Data processing may have a delay of up to 24 hours
- Historical data availability depends on when analytics was first enabled
- Some metrics may be aggregated to protect individual user privacy
- Analytics data retention follows your instance's data retention policies
- Detailed query and response data is only available to instance administrators

Note

For comprehensive monitoring beyond chat interactions, consider using CloudWatch metrics and logs in conjunction with the analytics dashboard. For more information, see [Monitoring Quick Suite](#).

Amazon security in Quick Suite

Amazon Quick Suite provides a secure platform that enables you to distribute dashboards and insights to tens of thousands of users, with multiple-region availability and built-in redundancy.

Cloud security at Amazon is the highest priority. As an Amazon customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between Amazon and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – Amazon is responsible for protecting the infrastructure that runs Amazon services in the Amazon Cloud. Amazon also provides you with services that you can use securely. The effectiveness of our security is regularly tested and verified by third-party auditors as part of the [Amazon compliance programs](#). To learn about the compliance programs that apply to Quick Suite, see [Amazon Services in Scope by Compliance Program](#).

- **Security in the cloud** – Your responsibility is determined by the Amazon service that you use. You are also responsible for other factors, including the sensitivity of your data, your organization's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using Amazon Quick Suite. The following topics show you how to configure Amazon Quick Suite to meet your security and compliance objectives. You also learn how to use other Amazon services that can help you to monitor and secure your Amazon Quick Suite resources.

Amazon Quick Suite enables you to manage your users and content using a comprehensive set of security features. These include role-based access control, Microsoft Active Directory integration, Amazon CloudTrail auditing, single sign-on using Amazon Identity and Access Management (IAM) and third-party solutions, private VPC subnets, and data backup. Amazon Quick Suite can also support FedRAMP, HIPAA, PCI DSS, ISO, and SOC compliance to help you meet industry-specific or regulatory requirements.

Topics

- [Data protection in Amazon Quick Suite](#)
- [Incident response, logging, and monitoring in Amazon Quick Suite](#)
- [Compliance validation for Amazon Quick Suite](#)
- [Resilience in Amazon Quick Suite](#)
- [Infrastructure security in Amazon Quick Suite](#)
- [Best practices for security in Amazon Quick Suite](#)
- [Amazon managed policies for Amazon Quick Suite](#)

Data protection in Amazon Quick Suite

The Amazon [shared responsibility model](#) applies to data protection in Amazon Quick Suite. As described in this model, Amazon is responsible for protecting the global infrastructure that runs all of the Amazon Web Services Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. You are also responsible for the security configuration and management tasks for the Amazon Web Services services that you use. For more information about data privacy, see the [Data Privacy FAQ](#).

For data protection purposes, we recommend that you protect Amazon Web Services account credentials and set up individual users with Amazon IAM Identity Center or Amazon Identity and

Access Management (IAM). That way, each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with Amazon resources. We require TLS 1.2 and recommend TLS 1.3.
- Set up API and user activity logging with Amazon CloudTrail. For information about using CloudTrail trails to capture Amazon activities, see [Working with CloudTrail trails](#) in the *Amazon CloudTrail User Guide*.
- Use Amazon encryption solutions, along with all default security controls within Amazon Web Services services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing sensitive data that is stored in Amazon S3.
- If you require FIPS 140-3 validated cryptographic modules when accessing Amazon through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-3](#).

We strongly recommend that you never put confidential or sensitive information, such as your customers' email addresses, into tags or free-form text fields such as a **Name** field. This includes when you work with Amazon Quick Suite or other Amazon Web Services services using the console, API, Amazon CLI, or Amazon SDKs. Any data that you enter into tags or free-form text fields used for names may be used for billing or diagnostic logs. If you provide a URL to an external server, we strongly recommend that you do not include credentials information in the URL to validate your request to that server.

Topics

- [Data encryption in Amazon Quick Suite](#)
- [Inter-network traffic privacy in Amazon Quick Suite](#)

Data encryption in Amazon Quick Suite

Amazon Quick Suite uses the following data encryption features:

- Encryption at rest
- Encryption in transit

- Key management

You can find more details about data encryption at rest and data encryption in transit in the following topics. For more information about key management in Amazon Quick Suite see [Encrypting Amazon Quick Suite SPICE datasets with Amazon KMS customer-managed keys](#).

Topics

- [Encryption at rest](#)
- [Encryption in transit](#)

Encryption at rest

Amazon Quick Suite securely stores your Amazon Quick Suite metadata. This includes the following:

- Amazon Quick Suite user data, including Amazon Quick Suite user names, email addresses, and passwords. Amazon Quick Suite administrators can view user names and emails, but each user's password is completely private to each user.
- Minimal data necessary to coordinate user identification with your Microsoft Active Directory or identity federation implementation (Federated Single Sign-On (IAM Identity Center) through Security Assertion Markup Language 2.0 (SAML 2.0)).
- Data source connection data.
- Amazon Quick Suite data source credentials (username and password) or OAuth tokens to establish a data source connection are encrypted with the customer's default CMK when the customer registers a CMK with Amazon Quick Suite. If the customer does not register a CMK with Amazon Quick Suite, we will continue to encrypt the information using an Amazon Quick Suite owned Amazon KMS key.
- Names of your uploaded files, data source names, and data set names.
- Statistics that Amazon Quick Suite uses to populate machine learning (ML) insights.
- Data indexed to support Amazon Q in Quick Suite. This includes the following:
 - Topics
 - Metadata related to your dashboards
 - Your first index capacity purchase
 - Your first chat

- Your first space creation
- Your first knowledge base creation

Note

Configure a CMK prior to creating the above. Otherwise, Q data will be encrypted by an Amazon-owned key and cannot be changed later.

Amazon Quick Suite securely stores your Amazon Quick Suite data. This includes the following:

- Data-at-rest in SPICE is encrypted using hardware block-level encryption with Amazon-managed keys.
- Data-at-rest other than SPICE is encrypted using Amazon-managed KMS keys. This includes the following:
 - Email reports
 - Sample value for filters

When you delete a user, all of that user's metadata is permanently deleted. If you don't transfer that user's Amazon Quick Suite objects to another user, all of the deleted user's Amazon Quick Suite objects (data sources, datasets, analyses, and so on) are also deleted. When you unsubscribe from Amazon Quick Suite, all metadata and any data you have in SPICE is completely and permanently deleted.

Encryption in transit

Amazon Quick Suite supports encryption for all data transfers. This includes transfers from the data source to SPICE, or from SPICE to the user interface. However, encryption isn't mandatory. For some databases, you can choose whether transfers from the data source are encrypted or not. Amazon Quick Suite secures all encrypted transfers by using Secure Sockets Layer (SSL).

Inter-network traffic privacy in Amazon Quick Suite

To use Amazon Quick Suite, users need access to the internet. They also need access to a compatible browser or a mobile device with the Amazon Quick Suite mobile app installed. They don't need access to the data sources they want to analyze. This access is handled inside Amazon

Quick Suite. User connections to Amazon Quick Suite are protected through the use of SSL. So that users can access Amazon Quick Suite, allow access to HTTPS and Web Sockets Secure (wss://) protocol.

You can use a Microsoft AD connector and single sign-on (IAM Identity Center) in a corporate network environment. You can further restrict access through the identity provider. Optionally, you can also use MFA.

Amazon Quick Suite accesses data sources by using connection information supplied by the data source owner in Amazon Quick Suite. Connections are protected both between Amazon Quick Suite and on-premises applications and between Amazon Quick Suite and other Amazon resources within the same Amazon Web Services Region. For connections to any source, the data source must allow connections from Amazon Quick Suite.

Traffic between service and on-premises clients and applications

You have two connectivity options between your private network and Amazon:

- An Amazon Site-to-Site VPN connection. For more information, see [What is Amazon site-to-site VPN?](#)
- An Amazon Direct Connect connection. For more information, see [What is Amazon direct connect?](#)

If you are using Amazon API operations to interact with Amazon Quick Suite through the network, clients must support Transport Layer Security (TLS) 1.0. We recommend TLS 1.2. Clients must also support cipher suites with Perfect Forward Secrecy (PFS), such as Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Diffie-Hellman Ephemeral (ECDHE). Most modern systems such as Java 7 and later support these modes. You must sign requests using an access key ID and a secret access key that are associated with an IAM principal, or you can use the [Amazon Security Token Service \(STS\)](#) to generate temporary security credentials to sign requests.

Traffic between Amazon resources in the same region

An Amazon Virtual Private Cloud (Amazon VPC) endpoint for Amazon Quick Suite is a logical entity within a VPC that allows connectivity only to Amazon Quick Suite. The VPC routes requests to Amazon Quick Suite and routes responses back to the VPC. For more information, see the following:

- [VPC endpoints](#) in the *Amazon VPC User Guide*

- [Connecting to a Amazon VPC with Amazon Quick Suite](#)

Incident response, logging, and monitoring in Amazon Quick Suite

Intended audience: System administrators and Amazon Quick Suite administrators

Effective incident response, logging, and monitoring are essential for maintaining the security, performance, and reliability of your Amazon Quick Suite instance. This monitoring framework provides multiple layers of visibility into user activities, system performance, security events, and operational metrics across all Amazon Quick Suite features including chat, spaces, flows, actions, research, dashboards, and custom agents.

Amazon Quick Suite integrates with Amazon native monitoring and logging services to provide both real-time insights and historical analysis capabilities. The monitoring system captures detailed analytics on user engagement, conversation patterns, resource utilization, and security-related events, while CloudTrail logging ensures complete audit trails for compliance and forensic analysis.

This section covers:

- **Analytics and monitoring** - Comprehensive dashboards and metrics for tracking user adoption, performance, feedback, and security events across all Amazon Quick Suite capabilities
- **CloudTrail logging** - Complete audit trails of API calls and administrative actions for compliance and security monitoring in Amazon Quick Sight
- **Non-API event logging** - Monitoring of user interactions, content access, and system events that don't generate API calls
- **Log analysis and interpretation** - Understanding log entries, identifying security incidents, and responding to operational issues in Amazon Quick Sight

Whether you're investigating a security incident, analyzing user behavior patterns, measuring system performance, or ensuring regulatory compliance, these monitoring and logging capabilities provide the visibility and data you need to maintain a secure and well-functioning Amazon Quick Suite environment.

Topics

- [Incident response, logging, and monitoring in Amazon Quick Sight using CloudTrail](#)

- [Monitoring data in Amazon Quick Sight using CloudWatch](#)

Incident response, logging, and monitoring in Amazon Quick Sight using CloudTrail

Amazon Quick Sight is integrated with Amazon CloudTrail. This service provides a record of actions taken by a user, role, or an Amazon service in Amazon Quick Sight. CloudTrail captures all API calls for Amazon Quick Sight as events. The calls captured include some calls from the Amazon Quick Sight console and all code calls to Amazon Quick Sight API operations. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Amazon Quick Sight. If you don't configure a trail, you can still view the most recent events in the CloudTrail console in **Event history**. Using the information collected by CloudTrail, you can determine the request that was made to Amazon Quick Sight, the IP address from which the request was made, who made the request, when it was made, and additional details.

Amazon Quick Sight doesn't natively support alerting with Amazon CloudWatch or other external systems. However, it's possible to develop a custom solution to process CloudTrail logs.

Amazon Quick Sight service status can be viewed on the [Service Health Dashboard](#).

By default, the log files delivered by CloudTrail to your bucket are encrypted by Amazon [server-side encryption with Amazon S3-managed encryption keys \(SSE-S3\)](#). To provide a security layer that is directly manageable, you can instead use [server-side encryption with Amazon KMS-managed keys \(SSE-KMS\)](#) for your CloudTrail log files. Enabling server-side encryption encrypts the log files but not the digest files with SSE-KMS. Digest files are encrypted with [Amazon S3-managed encryption keys \(SSE-S3\)](#).

To learn more about CloudTrail, including how to configure and enable it, see the [Amazon CloudTrail User Guide](#).

Topics

- [Logging Amazon Quick Sight information with Amazon CloudTrail](#)
- [Tracking non-API events by using CloudTrail logs](#)
- [Example: Amazon Quick Sight log file entries](#)

Logging Amazon Quick Sight information with Amazon CloudTrail

Intended audience: System administrators

CloudTrail is enabled on your Amazon account when you create the account. When supported event activity occurs in Amazon Quick Sight, that activity is recorded in a CloudTrail event along with other Amazon service events in **Event history**. You can view, search, and download recent events in your Amazon account. For more information, see [Viewing Events with CloudTrail Event History](#).

For an ongoing record of events in your Amazon account, including events for Amazon Quick Sight, create a trail. A *trail* enables CloudTrail to deliver log files to an Amazon S3 bucket. By default, when you create a trail in the console, the trail applies to all Amazon Web Services Regions. The trail logs events from all Regions in the Amazon partition and delivers the log files to the Amazon S3 bucket that you specify. Additionally, you can configure other Amazon services to further analyze and act upon the event data collected in CloudTrail logs. For more information, see the following:

- [Overview for Creating a Trail](#)
- [CloudTrail Supported Services and Integrations](#)
- [Configuring Amazon SNS Notifications for CloudTrail](#)
- [Receiving CloudTrail Log Files from Multiple Regions](#) and [Receiving CloudTrail Log Files from Multiple Accounts](#)
- [Cross-Account CloudTrail Logging](#) in the Amazon Lake Formation Developer Guide – This topic includes instructions for including principal identities in cross-account CloudTrail logs.

Amazon Quick Sight supports logging the following actions as events in CloudTrail log files:

- Whether the request was made with root or Amazon Identity and Access Management user credentials
- Whether the request was made with temporary security credentials for an IAM role or federated user
- Whether the request was made by another Amazon service

For more information on user identity, see the [CloudTrail userIdentity Element](#).

By default, each Amazon Quick Sight log entry contains the following information:

- **userIdentity** – User identity
- **eventTime** – Event time
- **eventId** – Event Id
- **readOnly** – Read only
- **awsRegion** – Amazon Web Services Region
- **eventSource (quicksight)** – Source of the event (Amazon Quick Sight)
- **eventType (AwsServiceEvent)** – Event type (Amazon service event)
- **recipientAccountId (customer Amazon account)** – Recipient account ID (Customer Amazon account)

 Note

CloudTrail displays users as unknown if they were provisioned by Amazon Quick Sight. This display is because these users aren't a known IAM identity type.

Tracking non-API events by using CloudTrail logs

Following is a list of the non-API events you can track.

User management

- **CreateAccount** – Create Account
- **BatchCreateUser** – Create User
- **BatchResendUserInvite** – Invite User
- **UpdateGroups** – Update Groups

This event works with Enterprise edition only.

- **UpdateSpiceCapacity** – Update SPICE Capacity
- **DeleteUser** – Delete User
- **Unsubscribe** – Unsubscribe User

Subscription

- **CreateSubscription** – Create Subscription
- **UpdateSubscription** – Update Subscription
- **DeleteSubscription** – Delete Subscription

Dashboard

- **GetDashboard** – Get Dashboard
- **CreateDashboard** – Create Dashboard
- **UpdateDashboard** – Update Dashboard
- **UpdateDashboardAccess** – Update Dashboard Access
- **DeleteDashboard** – Delete Dashboard

Analysis

- **GetAnalysis** – Get Analysis
- **CreateAnalysis** – Create Analysis
- **UpdateAnalysisAccess** – Update Analysis Access

- **UpdateAnalysis** – Update Analysis
 - **RenameAnalysis** – Rename Analysis
 - **CreateVisual** – Create Visual
 - **RenameVisual** – Rename Visual
 - **DeleteVisual** – Delete Visual
 - **DeleteAnalysis** – Delete Analysis

Data source

- **CreateDataSource** – Create Data Source
 - **FlatFile** – Flat file
 - **External** – External
 - **S3** – S3
 - **ImportS3ManifestFile** – S3 Manifest File
 - **Presto** – Presto
 - **RDS** – RDS
 - **Redshift** – Redshift (manual)
- **UpdateDataSource** – Update Data Source
- **DeleteDataSource** – Delete Data Source

- **CreateDataSet** – Create Data Set

- **CustomSQL** – Custom SQL
- **SQLTable** – SQL Table
- **File** – CSV or XLSX
- **UpdateDataSet** – Update SQL Join Dataset
- **UpdateDatasetAccess** – Update Dataset Access
- **DeleteDataSet** – Delete Dataset
- **Querydatabase** – During a dataset refresh, query data source.

Example: Amazon Quick Sight log file entries

A trail is a configuration that enables delivery of events as log files to an Amazon S3 bucket that you specify. CloudTrail log files contain one or more log entries. An event represents a single request from any source and includes information about the requested action, the date and time of the action, request parameters, and so on. CloudTrail log files aren't an ordered stack trace of the public API calls, so they don't appear in any specific order.

The following example shows a CloudTrail log entry that demonstrates the BatchCreateUser action.

```
{
  "eventVersion": "1.05",
  "userIdentity": {
    "type": "Root",
    "principalId": "123456789012",
    "arn": "arn:aws-cn:iam::123456789012:root",
    "accountId": "123456789012",
    "userName": "test-username"
  },
  "eventTime": "2017-04-19T03:16:13Z",
  "eventSource": "quicksight.amazonaws.com",
  "eventName": "BatchCreateUser",
  "awsRegion": "us-west-2",
  "requestParameters": null,
  "responseElements": null,
```

```
"eventID":"e7d2382e-70a0-3fb7-9d41-a7a913422240",
"readOnly":false,
"eventType":"AwsServiceEvent",
"recipientAccountId":"123456789012",
"serviceEventDetails":
{
  "eventRequestDetails":
  {
    "users":
    {
      "test-user-11":
      {
        "role":"USER"
      },
      "test-user-22":
      {
        "role":"ADMIN"
      }
    }
  },
  "eventResponseDetails":
  {
    "validUsers":[
    ],
    "InvalidUsers":[
      "test-user-11",
      "test-user-22"
    ]
  }
}
```

Monitoring data in Amazon Quick Sight using CloudWatch

Amazon Quick Suite sends metrics to Amazon CloudWatch that you can use to observe and respond to the availability and performance of your Amazon Quick Suite environment in near real time. Currently, you can monitor metrics for Amazon Quick Sight dashboards, visuals, and dataset ingestions, as well as unstructured datasets and Quick Action Connectors, to provide your readers with a consistent, high-performing, and uninterrupted experience on Amazon Quick Suite.

For more information about using Amazon CloudWatch, see the [Amazon CloudWatch User Guide](#).

Accessing Quick Suite metrics in Amazon CloudWatch

Use the following procedure to access Amazon Quick Suite metrics in Amazon CloudWatch.

To access Amazon Quick Suite metrics in CloudWatch

1. Sign in to the Amazon Web Services account that's associated with your Amazon Quick Suite account.
2. In the upper-left corner of the Amazon Web Services Management Console home page, choose **Services**, and then choose **CloudWatch**.
3. In the navigation pane, choose **Metrics**, **All metrics**, **QuickSight**.

Topics

- [Graph metrics with the Amazon CloudWatch console](#)
- [Creating alarms with the Amazon CloudWatch console](#)
- [Metrics](#)
- [Aggregate metrics](#)
- [Aggregate SPICE metrics](#)
- [Dimensions](#)

Graph metrics with the Amazon CloudWatch console

You can also use the Amazon CloudWatch console to graph metric data generated by Quick Suite. For more information, see [Graphing metrics](#) in the *Amazon CloudWatch User Guide*.

Creating alarms with the Amazon CloudWatch console

You can create a Amazon CloudWatch alarm that monitors CloudWatch metrics for your Quick Suite assets. When the metric reaches a threshold that you specify, CloudWatch automatically sends you a notification. For examples, see [Creating Amazon CloudWatch alarms](#) in the *Amazon CloudWatch User Guide*.

Metrics

The Amazon/QuickSight namespace includes the following metrics for monitoring traffic and latency of your Amazon Quick Suite assets.

Topics

- [Per-dashboard metrics](#)
- [Per-dataset ingestion metrics](#)
- [Per-visual metrics](#)
- [Per-unstructured dataset metrics](#)
- [Per-action connector metrics](#)

Per-dashboard metrics

The following metrics track dashboard view counts and load times. You can find these metrics under the Amazon/QuickSight/Dashboard Metrics group in CloudWatch.

Metric	Description	Dimension	Unit
DashboardViewCount	The number of times that a dashboard has been viewed. This number includes all access patterns such as web, mobile, and embedded. The most useful statistic for this metric is SUM, which represents the total number of dashboard views during a set period of time.	DashboardId	Count
DashboardViewLoadTime	The amount of time that it takes for a Amazon Quick Sight dashboard to load. The measurement begins when a user reaches the Amazon Quick Sight	DashboardId	Millisecond

Metric	Description	Dimension	Unit
	dashboard and ends when all of the dashboard's visuals finish rendering. The most useful statistic for this metric is AVERAGE, which represents the average load time of a Amazon Quick Sight dashboard during a set period of time.		

Per-dataset ingestion metrics

The following metrics track ingestions for specific [SPICE](#) datasets. You can find these metrics under the Amazon/QuickSight/Ingestion Metrics group in CloudWatch.

Metric	Description	Dimension	Unit
IngestionErrorCount	The number of failed ingestions. The most useful statistic for this metric is SUM, which represents the total number of failed ingestions during a set period of time.	DatasetId	Count
IngestionInvocationCount	The number of ingestions that have been initiated. This	DatasetId	Count

Metric	Description	Dimension	Unit
	includes scheduled and manual ingestions that are initiated through the console and the Amazon Quick Sight API operations. The most useful statistic for this metric is SUM, which represents the total number of ingestions initiated during a set period of time.		
IngestionLatency	The time period between the initiation of an ingestion to the completion of the ingestion. The most useful statistic for this metric is AVERAGE, which represents the average runtime of ingestions during a set period of time.	DatasetId	Second

Metric	Description	Dimension	Unit
IngestionRowCount	The number of successful row ingestions. The most useful statistic for this metric is SUM, which represents the total amount of data ingested during a set period of time.	DatasetId	Count

Per-visual metrics

The following metrics track the load times and error counts of individual visuals on a Amazon Quick Sight dashboard. You can find these metrics under the Amazon/QuickSight/Visual Metrics group in CloudWatch.

Metric	Description	Dimension	Unit
VisualLoadTime	The time that it takes for a Amazon Quick Sight visual to receive the necessary query data for an initial paint of the visual. This includes the round-trip query time from the client, to the Amazon Quick Sight service, and then back to client. The most useful statistic for this	• DashboardId • SheetId • VisualId	Millisecond

Metric	Description	Dimension	Unit
	metric is AVERAGE, which represents the average load time of a visual during a set period of time.		
VisualLoadErrorCount	The number of times that a Amazon Quick Sight visual fails to complete a data query for the initial paint. Any error that occurs during a visual's loading period is included in this metric. The most useful statistic for this metric is SUM, which represents the total number of failed visual loads during a set period.	- DashboardId - SheetId - VisualId	Count

Per-unstructured dataset metrics

The following metrics track document statistics and indexing status for Amazon Quick Sight unstructured datasets. You can find these metrics under the Amazon/QuickSight/QuickInstanceId group in CloudWatch.

Metric	Description	Dimension	Unit
QuickIndexDocumentCount	The number of documents in the	- KnowledgeBaseId - QuickInstanceId	Count

Metric	Description	Dimension	Unit
	unstructured Quick index.		
QuickIndexExtractedTextSize	The extracted text size of the unstructured Quick index.	- KnowledgeBaseId - QuickInstanceId	Bytes
QuickIndexPurchaseInMB	The amount of storage that has been purchased for the unstructured Quick index.	- KnowledgeBaseId - QuickInstanceId	MB
DocumentsCrawled	The number of uploaded documents crawled in the unstructured Quick index.	- KnowledgeBaseId - QuickInstanceId	Count
DocumentsIndexed	The number of documents indexed in the unstructured Quick index.	- KnowledgeBaseId - QuickInstanceId	Count
DocumentsDeleted	The number of documents deleted from the unstructured Quick index.	- KnowledgeBaseId - QuickInstanceId	Count
DocumentsModified	The number of documents modified in the unstructured Quick index.	- KnowledgeBaseId - QuickInstanceId	Count

Metric	Description	Dimension	Unit
DocumentsFailedToIndex	The number of documents that failed to index in the unstructured Quick index.	- KnowledgeBaseId - QuickInstanceId	Count
ExtractedTextSize	The total text size extracted during a connector level sync in the unstructured Quick index.	- KnowledgeBaseId - QuickInstanceId	MB

Per-action connector metrics

The following metrics track the number of invocations made to Quick Action Connectors. You can find these metrics under the Amazon/QuickSight namespace in CloudWatch.

Metric	Description	Dimension	Unit
ActionInvocationCount	The number of times your Action Connector was invoked.	- ActionConnectorId - ActionConnectorType - ActionId	Count
ActionInvocationError	The number of times your Action Connector failed to invoke.	- ActionConnectorId - ActionConnectorType - ActionId - InvokeErrorCode	Count

Aggregate metrics

The Amazon/QuickSight namespace includes the following aggregate metrics for monitoring traffic and latency of your Amazon Quick Suite assets.

Topics

- [Aggregate dashboard metrics](#)
- [Aggregate ingestion metrics](#)
- [Aggregate visual metrics](#)
- [Aggregate unstructured dataset metrics](#)
- [Aggregate action connector metrics](#)

Aggregate dashboard metrics

The following metrics track view counts and load times of all dashboards in a Amazon Quick Suite account and region. You can find these metrics under the Amazon/QuickSight/Aggregate Metrics group in CloudWatch.

Metric	Description	Unit
DashboardViewCount	The number of times that all Amazon Quick Sight dashboards have been viewed across the entire Amazon Quick Suite account in the region. This number is an aggregate that includes all access patterns such as web, mobile, and embedded. The most useful statistic for this metric is SUM, which represents the total number of Amazon Quick Sight dashboard views during a set period of time.	Count
DashboardViewLoadTime	The amount of time that it takes for all Amazon Quick Sight dashboards to load. The measurement begins	Millisecond

Metric	Description	Unit
	when a user navigates to the Amazon Quick Sight dashboard and ends when all of the dashboard's visuals finish rendering. The most useful statistic for this metric is AVERAGE, which represents the average load time of all Amazon Quick Sight dashboard during a set period of time.	

Aggregate ingestion metrics

The following metrics track all ingestions associated with a Amazon Quick Suite account and Amazon Web Services Region. You can find these metrics under the Amazon/QuickSight/Aggregate Metrics group in CloudWatch.

Metric	Description	Unit
IngestionErrorCount	The number of failed ingestions. The most useful statistic for this metric is SUM, which represents the total number of failed ingestion during a set period.	Count
IngestionInvocationCount	The number of ingestion s that have been initiated. This includes scheduled and manual ingestions that are initiated through the console	Count

Metric	Description	Unit
	and the Amazon Quick Sight API operations. The most useful statistic for this metric is SUM, which represents the total number of ingestions initiated during a set period of time.	
IngestionLatency	The time period between the initiation of an ingestion to the completion of the ingestion. The most useful statistic for this metric is AVERAGE, which represents the average runtime of ingestions during a set period of time.	Second
IngestionRowCount	The number of successful row ingestions. The most useful statistic for this metric is SUM, which represents the total amount of data ingested during a set period of time.	Count

Aggregate visual metrics

The following metrics track load times and error counts of all visuals on a dashboard and in a Amazon Quick Suite account in a Region. You can find these metrics under the Amazon/QuickSight/Aggregate Metrics group for CloudWatch.

Metric	Description	Unit
VisualLoadTime	The time that it takes for all Amazon Quick Sight visuals to receive the necessary query data for an initial paint of the visuals. This includes the round-trip query time from the client, to the Amazon Quick Suite service, and then back to the client. The most useful statistic for this metric is AVERAGE, which represents the average load time of all visuals during a set period of time.	Millisecond
VisualLoadErrorCount	The number of times that all Amazon Quick Sight visuals that belong to the Amazon Quick Suite account fail to complete a data query for an initial paint. The most useful statistic for this metric is SUM, which represents the total number of failed visuals during a set period.	Count

Aggregate unstructured dataset metrics

The following metrics track all unstructured dataset metrics within a a Amazon Quick Suite account in a Region. You can find these metrics under the Amazon/QuickSight/Aggregate Metrics group for CloudWatch.

Metric	Description	Unit
QuickIndexDocumentCount	The number of documents in the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total number of documents added to your index during a set period of time.	Count
QuickIndexExtractedTextSize	The extracted text size of the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total size of all text across all documents in your index.	Bytes
QuickIndexPurchasedInMB	The amount of storage that has been purchased for the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total size of purchased storage in MB across your index.	MB
DocumentsCrawled	The number of uploaded documents crawled in the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total number	Count

Metric	Description	Unit
	of documents crawled in your index.	
DocumentsIndexed	The number of documents indexed in the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total number of documents indexed.	Count
DocumentsDeleted	The number of documents deleted from the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total number of documents deleted from your index.	Count
DocumentsModified	The number of documents modified in the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total number of documents modified in your index.	Count

Metric	Description	Unit
DocumentsFailedToIndex	The number of documents that failed to index in the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total number of documents that failed to index.	Count
ExtractedTextSize	The total text size extracted during a connector level sync in the unstructured Quick index. The most useful statistic for this metric is SUM, which represents the total size of documents extracted during a connector level sync.	MB

Aggregate action connector metrics

The following metrics track all Quick Suite action connector invocations associated with a Amazon Quick Suite account in an Amazon Region. You can find these metrics under the Amazon/QuickSight/Aggregate Metrics group for CloudWatch.

Metric	Description	Unit
ActionInvocationCount	The number of action connector invocations made. The most useful statistic for this metric is SUM, which represents the total number of action connector invocatio	Count

Metric	Description	Unit
	ns initiated during a set period of time.	
ActionInvocationError	The number of failed action connector invocations. The most useful statistic for this metric is SUM, which represents the total number of action connector invocations that failed during a set period of time.	Count

Aggregate SPICE metrics

The following metrics monitor SPICE consumption information to help you avoid reaching the SPICE consumption limit that can cause your ingestions to fail. Statistics are stored for up to 15 months so that you can access historical information to better understand the consumption trends of your Amazon Quick Suite account. You can find these metrics in the Amazon/QuickSight/Aggregate Metrics group for CloudWatch.

Metric	Description	Unit
SPICECapacityLimitInMB	This value represents the provisioned SPICE capacity at a specific point in time. This metric refreshes when an update of 1 MB or more in consumed or purchased capacity is made.	MegaBytes
SPICECapacityConsumedInMB	This value represents the consumed SPICE capacity at a specific point in time. This metric refreshes when an update of 1 MB or more	MegaBytes

Metric	Description	Unit
	in consumed or purchased capacity is made.	

Dimensions

Following is a list of Quick Suite metric dimensions that appear in Amazon CloudWatch.

Dimension	Description	
DashboardId	The public ID of a Amazon Quick Sight dashboard. You can use the <code>ListDashboards</code> API operation to see a list of every dashboard in your Amazon Quick Suite account. For more information, see ListDashboards in the <i>Amazon Quick Sight API Reference</i>.	
DatasetId	The public ID of a Amazon Quick Sight dataset. You can use the <code>ListDataSets</code> API operation to see a list of every dataset in your Amazon Quick Sight account. For more information, see ListDataSets in the <i>Amazon Quick Sight API Reference</i>.	
SheetId	The public ID of a Amazon Quick Sight sheet.	

Dimension	Description	
VisualId	The public ID of a Amazon Quick Sight visual.	
KnowledgeBaseId	The public ID of a Amazon Quick Sight knowledge base.	
QuickInstanceId	The public ID of the Quick instance.	
ActionConnectorId	The public ID of the Quick Action Connector.	
ActionConnectorType	The type of the Quick Action Connector.	
ActionId	The public ID of the Quick Action.	
InvokeErrorCode	The error code related to a failed Quick Action Connector invocation.	

Compliance validation for Amazon Quick Suite

Third-party auditors assess the security and compliance of Quick Suite as part of multiple Amazon compliance programs. These include FedRamp, HIPAA, PCI DSS, SOC, and ISO (9001, 27001, 27018, and 27019).

For the most current list of Amazon services in scope of specific compliance programs, see [Amazon services in scope by compliance program](#). For general information, see [Amazon compliance programs](#).

You can download third-party audit reports using Amazon Artifact. For more information, see [Downloading reports in Amazon Artifact](#).

Your compliance responsibility when using Amazon Quick Suite is determined by the sensitivity of your data, your company's compliance objectives, and applicable laws and regulations. Amazon provides the following resources to help with compliance:

- [Security and compliance quick start guides](#) – These deployment guides discuss architectural considerations and provide steps for deploying security- and compliance-focused baseline environments on Amazon.
- [Architecting for HIPAA security and compliance paper](#) – This paper describes how companies can use Amazon to create HIPAA-compliant applications.
- [Amazon compliance resources](#) – This collection of workbooks and guides might apply to your industry and location.
- [Amazon Config](#) – This Amazon service assesses how well your resource configurations comply with internal practices, industry guidelines, and regulations.
- [Amazon Security Hub](#) – This Amazon service provides a comprehensive view of your security state within Amazon that helps you check your compliance with security industry standards and best practices.

Resilience in Amazon Quick Suite

Quick Suite is built by Amazon and runs on Amazon-managed infrastructure. It takes full advantage of the high availability features provided by Amazon.

The Amazon global infrastructure is built around Amazon Regions and Availability Zones. Amazon Regions provide multiple physically separated and isolated Availability Zones, which are connected with low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between Availability Zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

Because Amazon Quick Suite is an Amazon-managed application, all patches and updates are applied by Amazon as needed.

For more information about Amazon Regions and Availability Zones, see [Amazon global infrastructure](#).

Infrastructure security in Amazon Quick Suite

Intended audience: Amazon Quick Suite administrators

Quick Suite is delivered as a web application, hosted on dedicated Amazon EC2 hosts, separate from Amazon virtual private clouds (VPCs). Instead of deploying Amazon Quick Suite on your own hosts, you access the Amazon Quick Suite service through Regional public endpoints. Amazon Quick Suite accesses data sources over a secured internet connection from Regional endpoints. To access data sources that are located inside a corporate network, configure the network to allow access from one of the Amazon Quick Suite public IP address blocks. We recommend that you consider using a VPC (a virtual network dedicated to your Amazon account).

For more information, see the following:

- [Global Infrastructure: The Most Extensive, Reliable, and Secure Global Cloud Infrastructure](#)
- [Amazon Regions, websites, IP address ranges, and endpoints](#)
- [Connecting to a Amazon VPC with Amazon Quick Suite](#)

As a managed service, Quick Suite is protected by the Amazon global network security procedures that are described in the [Amazon Web Services: Overview of Security Processes](#) paper.

If you use Amazon published API calls to access Amazon Quick Suite through the network, clients must support Transport Layer Security (TLS) 1.0 or later. We recommend TLS 1.2 or later. Clients must also support cipher suites with perfect forward secrecy (PFS) such as Ephemeral Diffie-Hellman (DHE) or Elliptic Curve Ephemeral Diffie-Hellman (ECDHE). Most modern systems such as Java 7 and later support these modes.

Additionally, requests must be signed by using an access key ID and a secret access key that is associated with an Amazon Identity and Access Management (IAM) principal. Or you can use the [Amazon Security Token Service](#) (Amazon STS) to generate temporary security credentials to sign requests.

You can call these API operations from any network location, but Amazon Quick Suite does support resource-based access policies, which can include restrictions based on the source IP address. You can also use Amazon Quick Suite policies to control access from specific Amazon Virtual Private Cloud (Amazon VPC) endpoints or specific VPCs. Effectively, this isolates network access to a given Amazon Quick Suite resource from only the specific VPC within the Amazon network. For more

information on using Amazon Quick Suite in a VPC, see [Connecting to a Amazon VPC with Amazon Quick Suite](#).

Best practices for security in Amazon Quick Suite

Amazon Quick Suite provides a number of security features to consider as you develop and implement your own security policies. The following best practices are general guidelines and don't represent a complete security solution. Because these best practices might not be appropriate or sufficient for your environment, treat them as helpful considerations rather than prescriptions.

Firewall – To allow users to access Amazon Quick Suite, allow access to HTTPS and WebSockets Secure (wss://) protocol. To allow Amazon Quick Suite to reach a database that is on a non-Amazon server, change that server's firewall configuration to accept traffic from the applicable Amazon Quick Suite IP address range.

SSL – Use SSL to connect to your databases, especially if you are using public networks. Using SSL with Amazon Quick Suite requires the use of certificates signed by a publicly-recognized certificate authority (CA).

Enhanced security – Use Amazon Quick Suite Enterprise edition to make use of its enhanced security capabilities, including the following.

- Store data in SPICE with encryption at rest.
- Integrate Active Directory and IAM Identity Center authentication.
- Securely access data in private VPCs and on-premises.
- Limit access to data with row level security.

VPC – (Enterprise Edition) Use a virtual private cloud (VPC) for data in Amazon data sources and for data in on-premises servers without public connectivity. For Amazon sources, VPC access for Amazon Quick Suite uses an elastic network interface for secure, private communication with data sources in a VPC. For your local data, VPC allows you to use Amazon Direct Connect to create a secure, private link with your on-premises resources.

Amazon managed policies for Amazon Quick Suite

To add permissions to users, groups, and roles, it is easier to use Amazon managed policies than to write policies yourself. It takes time and expertise to [create IAM customer managed policies](#)

that provide your team with only the permissions they need. To get started quickly, you can use our Amazon managed policies. These policies cover common use cases and are available in your Amazon Web Services account. For more information about Amazon managed policies, see [Amazon managed policies](#) in the *IAM User Guide*.

Amazon services maintain and update Amazon managed policies. You can't change the permissions in Amazon managed policies. Services occasionally add additional permissions to an Amazon managed policy to support new features. This type of update affects all identities (users, groups, and roles) where the policy is attached. Services are most likely to update an Amazon managed policy when a new feature is launched or when new operations become available. Services do not remove permissions from an Amazon managed policy, so policy updates won't break your existing permissions.

Additionally, Amazon supports managed policies for job functions that span multiple services. For example, the **ReadOnlyAccess** Amazon managed policy provides read-only access to all Amazon services and resources. When a service launches a new feature, Amazon adds read-only permissions for new operations and resources. For a list and descriptions of job function policies, see [Amazon managed policies for job functions](#) in the *IAM User Guide*.

Topics

- [Amazon managed policy: AWSQuickSightElasticsearchPolicy](#)
- [Amazon managed policy: AWSQuickSightOpenSearchPolicy](#)
- [Amazon managed policy: AWSQuickSightSageMakerPolicy](#)
- [Amazon managed policy: AWSQuickSightAssetBundleExportPolicy](#)
- [Amazon managed policy: AWSQuickSightAssetBundleImportPolicy](#)
- [Amazon Quick Suite updates to Amazon managed policies](#)

Amazon managed policy: AWSQuickSightElasticsearchPolicy

This information is provided for backward compatibility only. The `AWSQuickSightOpenSearchPolicy` Amazon managed policy replaces the `AWSQuickSightElasticsearchPolicy` Amazon managed policy.

Previously, you used the `AWSQuickSightElasticsearchPolicy` Amazon managed policy to provide access to Amazon Elasticsearch Service resources from Amazon Quick Suite. Starting on or after September 7, 2021, Amazon Elasticsearch Service is renamed to Amazon OpenSearch Service.

Wherever you are using `AWSQuickSightElasticsearchPolicy`, you can update to the new Amazon managed policy that's called `AWSQuickSightOpenSearchPolicy`. You can attach the policy to your IAM entities. Amazon Quick Suite also attaches the policy to a service role that allows Amazon Quick Suite to perform actions on your behalf. `AWSQuickSightElasticsearchPolicy` is still available and as of August 31, 2021, had the same permissions as the new policy. However, `AWSQuickSightElasticsearchPolicy` is no longer kept up-to-date with latest changes.

This policy grants read-only permissions that allow access to OpenSearch (previously known as Elasticsearch) resources from Amazon Quick Suite.

Permissions details

This policy includes the following permissions:

- `es` – Allows principals to use `es:ESHttpGet` to access your OpenSearch (previously known as Elasticsearch) domains, cluster settings, and indices. This is required to use the search service from Amazon Quick Suite.
- `es` – Allows principals to use `es:ListDomainNames` to list your OpenSearch (previously known as Elasticsearch) domains. This is required to initiate access of the search service from Amazon Quick Suite.
- `es` – Allows principals to use `es:DescribeElasticsearchDomain` to search your OpenSearch (previously known as Elasticsearch) domains. This is required to use the search service from Amazon Quick Suite.
- `es` – Allows principals to use `es:ESHttpPost` and `es:ESHttpGet` with your OpenSearch (previously known as Elasticsearch) domains. This is required to use a SQL plugin with read-only access to the search service domains from Amazon Quick Suite.

For information on the contents of this IAM policy, see [AWSQuickSightElasticsearchPolicy](#) in the IAM console.

Amazon managed policy: AWSQuickSightOpenSearchPolicy

Use the AWSQuickSightOpenSearchPolicy Amazon managed policy to provide access to Amazon OpenSearch Service resources from Amazon Quick Suite.

AWSQuickSightOpenSearchPolicy replaces AWSQuickSightElasticsearchPolicy.

As of August 31, 2021, this policy had the same permissions as the legacy policy, AWSQuickSightElasticsearchPolicy. For now, you can use them interchangeably. For the long term, we recommend updating your policy usage to AWSQuickSightOpenSearchPolicy.

You can attach AWSQuickSightOpenSearchPolicy to your IAM entities. Amazon Quick Suite also attaches this policy to a service role that allows Amazon Quick Suite to perform actions on your behalf.

This policy grants read-only permissions that allow access to OpenSearch resources from Amazon Quick Suite.

Permissions details

This policy includes the following permissions:

- `es` – Allows principals to use `es:ESHttpGet` to access your OpenSearch domains, cluster settings, and indices. This is required to use Amazon OpenSearch Service from Amazon Quick Suite.
- `es` – Allows principals to use `es:ListDomainNames` to list your OpenSearch domains. This is required to initiate access of Amazon OpenSearch Service from Amazon Quick Suite.
- `es` – Allows principals to use `es:DescribeElasticsearchDomain` and `es:DescribeDomain` to search your OpenSearch domains. This is required to use Amazon OpenSearch Service from Amazon Quick Suite.
- `es` – Allows principals to use `es:ESHttpPost` and `es:ESHttpGet` with your OpenSearch domains. This is required to use a SQL plugin with read-only access to Amazon OpenSearch Service domains from Amazon Quick Suite.

For information on the contents of this IAM policy, see [AWSQuickSightOpenSearchPolicy](#) in the IAM console.

Amazon managed policy: AWSQuickSightSageMakerPolicy

Use the AWSQuickSightSageMakerPolicy Amazon managed policy to provide access to Amazon SageMaker AI resources from Amazon Quick Suite.

You can attach `AWSQuickSightSageMakerPolicy` to your IAM entities. Amazon Quick Suite also attaches this policy to a service role that allows Amazon Quick Suite to perform actions on your behalf.

This policy grants read-only permissions that allow access to Amazon SageMaker AI resources from Amazon Quick Suite.

To view the `AWSQuickSightSageMakerPolicy`, see [AWSQuickSightSageMakerPolicy](#) in the [Amazon Managed Policy reference](#).

Permissions details

This policy includes the following permissions:

- `sagemaker` – .
- `s3` – Allows principals to use `s3:GetObject` on all Amazon S3 buckets that start with the prefix `arn:aws-cn:s3:::sagemaker.*` to access data stored in SageMaker AI default buckets. This is required to load models shared from Amazon SageMaker AI Canvas to the default Amazon SageMaker AI Canvas Amazon S3 bucket.
- `s3` – Allows principals to use `s3:PutObject` to export objects into an Amazon S3 bucket. This is required to support existing datasets from Amazon Quick Suite to Amazon SageMaker AI Canvas to build predictive models.
- `s3` – Allows principals to use `s3:ListBucket` to allow Amazon Quick Suite to validate an existing Amazon SageMaker AI Canvas bucket in Amazon S3. This is required to allow the export of data from Amazon Quick Suite to Amazon SageMaker AI Canvas to build predictive models.
- `s3` – Allows principals to use `s3:GetObject` on all Amazon Quick Suite– owned Amazon S3 buckets that start with the prefix `arn:aws-cn:s3:::quicksight-ml`. This is required to allow Amazon Quick Suite to access the predictions that are generated by Amazon SageMaker AI Canvas. The generated predictions can be appended to a Amazon Quick Suite dataset.
- `sagemaker` – Allows principals to use `sagemaker:CreateTransformJob`, `sagemaker:DescribeTransformJob`, and `sagemaker:StopTransformJob` to perform SageMaker AI transform jobs on your behalf. This is required for Amazon Quick Suite to request predictions from SageMaker AI models that can be appended to a Amazon Quick Suite dataset.
- `sagemaker` – Allows principals to use `sagemaker:ListModels` to list your SageMaker AI models. This is required to allow generated SageMaker AI models to appear in Amazon Quick Suite.

Amazon managed policy: AWSQuickSightAssetBundleExportPolicy

Use the AWSQuickSightAssetBundleExportPolicy Amazon managed policy to perform asset bundle export operations. You can attach AWSQuickSightAssetBundleExportPolicy to your IAM entities.

This policy grants read-only permissions that allow access to Amazon Quick Suite asset resources. To view the details of this policy, see [AWSQuickSightAssetBundleExportPolicy](#) in the Amazon Managed Policy reference.

This policy includes the following permissions:

- `quicksight` – Allows principals to use `quicksight:Describe*` and `quicksight:List*` to find and fetch Amazon Quick Suite assets and their corresponding permissions.
- `quicksight` – Allows principals to use `quicksight:ListTagsForResource` to fetch tags of Amazon Quick Suite assets.
- `quicksight` – Allows principals to list, execute, and get the status of an Asset bundle export job. This policy uses the `quicksight:ListAssetBundleExportJob`, `StartAssetBundleExportJob`, and `quicksight:DescribeAssetBundleExportJob` permissions.

Amazon managed policy: AWSQuickSightAssetBundleImportPolicy

Use the AWSQuickSightAssetBundleImportPolicy Amazon managed policy to perform asset bundle import operations. This managed policy does not grant permissions for any run-as-role functionality with the `iam:passrole` that is required for some VPC connection and DataSource operations. This policy also does not grant access to retrieve objects from a users Amazon S3 bucket.

You can attach the AWSQuickSightAssetBundleImportPolicy to your IAM entities. This policy grants read and write permissions that allow access to Amazon Quick Suite resources. To view the details of this policy, see [AWSQuickSightAssetBundleImportPolicy](#) in the Amazon Managed Policy reference.

This policy includes the following permissions:

- `quicksight` – Allows principals to use `quicksight:Describe*` and `quicksight:List*` to detect changes in the Amazon Quick Suite assets and their permissions.

- `quicksight` – Allows principals to use `quicksight:Create*` and `quicksight:Update*` to make changes to the Amazon Quick Suite assets and permissions from the supplied asset bundle.
- `quicksight` – Allows principals to use `quicksight:ListTagsForResource`, `quicksight:TagResource`, and `quicksight:UntagResource` to update the tags of Amazon Quick Suite assets.
- `quicksight` – Allows principals to list, execute, and get the status of an Asset bundle import job. This policy uses the `quicksight:ListAssetBundleImportJob`, `quicksight:StartAssetBundleImportJob`, and `quicksight:DescribeAssetBundleImportJob` permissions.

Amazon Quick Suite updates to Amazon managed policies

View details about updates to Amazon managed policies for Amazon Quick Suite since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the [Amazon Quick Suite Document History](#) page.

Change	Description	Date
AWSQuickSightAssetBundleExportPolicy – New policy	Amazon Quick Suite added new permissions to simplify Asset bundle export operations.	March 27, 2024
AWSQuickSightAssetBundleImportPolicy – New policy	Amazon Quick Suite added new permissions to simplify Asset bundle import operations.	March 27, 2024
AWSQuickSageMakerPolicy – Update to an existing policy	Amazon Quick Suite added new permissions to allow integration with Amazon SageMaker AI Canvas.	July 25, 2023

Change	Description	Date
AWSQuickSightElasticsearchPolicy – Update to an existing policy	Amazon Quick Suite added new permissions to provide access to Amazon OpenSearch Service resources.	September 08, 2021
AWSQuickSightOpenSearchPolicy – New policy	Amazon Quick Suite added a new policy to allow access to Amazon OpenSearch Service resources from Quick Suite.	September 08, 2021
Amazon Quick Suite started tracking changes	Amazon Quick Suite started tracking changes for its Amazon managed policies.	August 2, 2021

Ask questions, explore data, and get insights with chat in Amazon Quick Suite

You can use Amazon Quick Suite chat to ask questions, interact with your Quick Suite resources, and accomplish tasks using multiple agentic workflows using natural language. When you ask a question, Quick Suite chat analyzes the data it has access to to generate comprehensive responses.

Quick Suite chat provides you with the following capabilities.

Topics

- [Welcome message](#)
- [Suggested prompts](#)
- [Expand and collapse modes](#)
- [Agent picker](#)
- [Scoping response generation to specific data](#)
- [Web search](#)
- [Upload files and chat](#)
- [Actions](#)
- [Flows](#)
- [Source citations](#)
- [Copy responses](#)
- [Provide feedback](#)
- [Artifact panel](#)
- [Conversation management](#)
- [Response events](#)
- [Contextual awareness](#)

Welcome message

When you log in, Amazon Quick Suite chat displays a welcome message.

Suggested prompts

To help you get started on interacting with your resources in chat, Amazon Quick Suite displays suggested prompts you can use to learn more about Amazon Quick Suite features. Use these to get started on your tasks. Once you select a prompt, Amazon Quick Suite guides you step-by-step through the information it needs from you to complete the task.

Note

Suggested prompts are system generated for the system agent; you can't customize them. For custom agents, suggested prompts can be created and customized by the agent builder or owner.

Expand and collapse modes

To help enhance your productivity and streamline your chat interactions, Amazon Quick Suite offers two separate chat window display modes:

- **Collapse mode** – The default chat window display. Use for easy access to chat from within all Amazon Quick Suite contexts.
- **Expand mode** – Access chat in an expanded view. Use this mode to work in the chat window without any distractions.

You can open up the chat interface at any time by selecting the chat bubble icon from the top right navigation menu. You can also exit out of the chat at any time by selecting the cancel icon from the top navigation menu.

Agent picker

Chat agents in Amazon Quick Suite are conversational assistants that help users explore data, analyze information, and take actions. You can pick any agent you've created to use during chat, or use the system agent **My Assistant** for your conversations.

To learn more about agents, see [Working with agents](#).

Scoping response generation to specific data

When you use Amazon Quick Suite, you can define what data you want your chat responses generated from using data filters. Data filters enable you to customize your chat's expertise area. The data available to your agent depends on whether it is bound to specific Amazon Quick Suite knowledge and integrations, or if it's unbounded (not connected to specific Amazon Quick Suite resources).

By default, an agent has knowledge of all your Amazon Quick Suite resources and its underlying large language model (LLM). While using chat, you can select between the following knowledge modes:

- **All data and apps** – Amazon Quick Suite will use all available knowledge sources (including spaces, dashboards, topics, knowledge bases, action integrations, and LLM knowledge) to generate a response to your prompts.
- **General knowledge** – Amazon Quick Suite will only use LLM knowledge to generate responses to your prompts. It won't use any knowledge from Amazon Quick Suite resources.
- **Specific data and apps** – Chat with a specific or multiple spaces, dashboards, topics, knowledge bases, or actions. To choose specific resources, select **Choose resource**.

Note

You can't add datasets to a chat filter.

An agent with pre-defined knowledge, on the other hand, only has knowledge of existing Amazon Quick Suite resources already linked to it, and its underlying LLM knowledge. While chatting with an agent with pre-defined knowledge, you get the following knowledge modes:

- **Pre-defined knowledge** – Amazon Quick Suite will use only the Amazon Quick Suite resource knowledge linked to the agent, and LLM knowledge, to generate responses. You can't change this setting.
- **Additional data and apps** – Chat with a specific or multiple spaces, dashboards, topics, or knowledge bases and actions. To choose specific resources, select **Choose resource**.

Web search

Amazon Quick Suite comes with built-in web search capability. Amazon Quick Suite chat automatically searches both your connected data sources and the web, providing comprehensive answers from either or both sources as needed. You can turn web search off for any query by clicking the globe icon in the chat panel. Any content retrieved from the web is appropriately cited in the final response.

Note

Web search is available to users only if it has been enabled by admin.

Upload files and chat

You can upload documents into Amazon Quick Suite chat and ask Amazon Quick Suite chat to summarize or analyze data based on the content of the uploaded documents.

When you start a new conversation, you can upload new files, choose from a saved list of recent documents, or drag and drop files directly into the conversation.

You can upload up to 20 files at one time with the following size limits:

- 10 MB for images (.jpeg, .png)
- 5 MB for spreadsheet files (.csv, .xls, .xlsx)
- 50 MB for other supported formats

Supported file types include:

- Word (.docx)
- Excel (.xls, .xlsx)
- PowerPoint (.ppt, .pptx)
- PDF (.pdf)
- images (.jpeg, .png)
- text files (.csv, .txt, .rtf, .md)
- Outlook (.msg)

- structured data files (.json, .yaml, .xml, .vtt, and common code files)

The total parsed content for all files combined has to be under 665,000 characters.

Note

If your files exceed the file count and character count limits, you can upload them to a [space](#) instead.

Documents uploaded through the chat interface are deleted with the associated conversation after 30 days of inactivity.

Note

Audio and video files are not supported as file uploads.

Actions

Amazon Quick Suite actions boost productivity by allowing you to perform relevant tasks in supported third-party services from your Amazon Quick Suite chat interface, like sending an email using Microsoft Outlook, or creating a PagerDuty incident. You can invoke any actions configured and available to you from your Amazon Quick Suite chat interface. To select an action, use the action application from the data scoping filter and select the third party application you want to perform actions in. Then, choose the action you want to perform from the available actions. The system checks if you have authorization to perform the action before performing it.

To learn more about actions, see [Actions in Amazon Quick Suite](#).

Flows

Use Amazon Quick Suite flows to create shortcuts for everyday, repetitive tasks you perform while chatting in Amazon Quick Suite. For example, if you find yourself using Amazon Quick Suite chat to create and copy edit product launch descriptions, you can use Amazon Quick Suite flows to generate an application that takes your draft content as input, and generates a polished draft—after content, style, and grammar review—as output.

You can run a flow directly from chat within a conversation. Within a single conversation, you can talk to a chat agent as well as run a flow.

To learn more about flows, see [Flows in Amazon Quick Suite](#).

Source citations

Amazon Quick Suite chat responses provide in-text source citations in a numbered list. To view the source for a response, choose the number at the end of the sentence. The dialog box shows the title, URL, and a snippet from the source that was used to generate the response. Choose the URL to view the source document.

To view the full list of sources, choose **Sources** at the end of the response. You can use the source list to fact-check the response, or for deeper analysis.

Copy responses

You can copy and save the responses for later review and analysis. To copy a response, choose the copy icon at the end of the response.

Provide feedback

You can provide direct feedback about the responses you received from Amazon Quick Suite chat by using the thumbs-up or thumbs-down button. Your feedback is used to help address technical issues in the web experience.

If you select the thumbs-down button, you can choose from the following feedback options:

- **Inaccurate, not factually correct**
- **Incomplete answer**
- **Didn't understand my question**
- **Issue with sources**
- **Too wordy**
- **Offensive or unsafe language**
- **Other**

You can also copy conversation details from the feedback window for further troubleshooting.

Artifact panel

An *artifact* in Amazon Quick Suite is defined as any Amazon Quick Suite generated content produced using Amazon Quick Suite's underlying large language model's (LLM) generative capabilities, for example: contracts, email drafts, social media posts, tables, interactive charts, code blocks, or any Amazon Quick Suite visuals.

Amazon Quick Suite stores any artifacts it generates as responses during chat in an artifact panel. You can access any artifacts created during your chat by selecting **View more** from below the artifact in chat.

Note

You can download artifacts as text, CSV, and code files. Word, PDF, and Excel downloads are not supported. Saving artifacts to a Space is not supported.

Conversation management

Amazon Quick Suite stores conversations for up to 30 days, and you can access them in the left navigation pane. You can perform the following tasks to manage your conversations:

- **Start new conversation** – Choose **+** from the chat controls menu to start a new conversation.
- **View conversation history** – Choose the clock icon to view all conversation history for the last 30 days. You can filter conversation history to specific agents.

Note

Searching, saving, renaming, and sharing conversations with other users is not supported.

Response events

For every natural language prompt Amazon Quick Suite chat receives, it generates real-time processing steps while the request is being fulfilled. These steps can be viewed as part of **Response events** during response generation, and are also accessible once Amazon Quick Suite has finished generating a response. Response events vary based on the chat interaction you're having (whether

you're chatting with a file you've uploaded or with a Amazon Quick Suite resource) and capture any rewrites to your response before it's finalized. A response event also includes the system generated identifiers for your conversation—you can copy and save the information for troubleshooting..

Contextual awareness

Amazon Quick Suite chat is context aware. It registers the Amazon Quick Suite resource you're interacting with and scopes your chat automatically to the resource. This way, when you start chatting, the resource you're interacting with is already included as a knowledge source for your chat.

If you navigate to a specific Amazon Quick Suite resource during an on-going chat, and that resource supports chat, Amazon Quick Suite chat will automatically register your new context and ask you whether you want to chat with the specific resource you have navigated to. If you want to continue chatting in the mode you're in, you can choose to **Ignore** the message. Or you can choose **Yes** to focus your chat on the specific resource you're interacting with. If you do so, Amazon Quick Suite automatically scopes your chat interactions to the resource in context.

Create, customize, and deploy AI-powered chat agents in Amazon Quick Suite

Chat agents in Amazon Quick Suite help users explore data, analyze information, and take actions. Users can interact with chat agents using the Quick Suite chat interface. Chat agents provide assistance through open-ended conversations supported by specific goals, sources of knowledge, and any connected tools. Chat agents can evolve from simple question-answering interfaces to more advanced functions that orchestrate complex workflows.

You can use chat agents to:

- Generate content and provide answers through natural language conversations
- Analyze and summarize information from connected spaces, dashboards, topics, datasets, and uploaded files
- Invoke actions to perform predefined steps for consistent, repeatable outcomes

Note

To learn more about chatting using chat agents, see [Using Amazon Quick Suite chat](#).

Topics

- [Chat agent types](#)
- [Amazon Quick Suite user interaction with chat agents](#)
- [Amazon Quick Suite user permissions for chat agents](#)
- [Custom permissions for chat agents](#)
- [System chat agent](#)
- [Custom chat agents](#)
- [Chat agent context sources and best practices](#)
- [Use a chat agent](#)

Chat agent types

Amazon Quick Suite supports two types of chat agents:

- **System chat agent** – This chat agent ("My assistant") is automatically available to all users by default. The system chat agent serves as a base planner with no inherent data or actions of its own—it dynamically accesses resources available to each user during chat time, allowing it to be tailored to individual user permissions and available content. Admin users control system chat agent settings by assigning specific users as owners so they can customize its persona and other settings.

The system chat agent is enabled with all chat capabilities, including file upload functionality, LLM knowledge access, toxicity and other guardrails, web search. It also includes chat data scoping mechanisms that provide access to spaces, topics, dashboards, knowledge bases, and actions based on user permissions.

- **Custom chat agents** – These chat agents can be created and customized to specific use cases by users with chat agent creation capabilities in Amazon Quick Suite, and shared with anyone. Admins can choose to restrict specific users and groups from creating custom chat agents, while still allowing users to use chat agents via chat. Chat agents return responses scoped to content that their invoking users have permissions to.

Custom chat agents can interact with Amazon Quick Suite resources in the following ways:

- **Preconfigured with resources** – These chat agents use only the configured resources (for example, spaces as knowledge sources, action connectors as tools) when looking for answers or orchestrating actions as their default behavior. While chatting, users can attach additional resources as per their need or invoke actions directly. For example, chat agents configured with only spaces containing files can't take actions by default unless users explicitly invoke them. Similarly, chat agents configured with only actions rely on LLM knowledge unless users attach a space or dashboard for enterprise-specific answers. However, users may attach additional spaces or directly invoke other actions they may have access to within the chat interface.
- **Not opinionated at build time** – These chat agents are not initially configured with resources when they are built (spaces, or action connectors, or both). Chat determines the resource boundary for the chat agent. For example, if no space is configured for the chat agent, chat will default to all spaces or LLM knowledge until the user makes a change. If no action connectors are configured for the chat agent, during chat, all actions a user has access to will be available to the agent until user makes a resource selection that doesn't include actions.

User resource selection during chat affects chat agent behavior. If a user selects a space, the chat agent will only answer with data within that space and take actions available within that space (if it is not opinionated at build time). If a user selects a dashboard, topic, or knowledge base, the chat agent will only answer from that source and not take any actions since the focus is changed to a specific data source. Users are expected to select all relevant resources within chat for comprehensive assistance.

 **Note**

The system chat agent is an unopinionated chat agent by design. If you want chat agents to access all actions irrespective of data focus, configure chat agents with action connectors.

Amazon Quick Suite user interaction with chat agents

The following table shows what you can do with chat agents in the admin console versus as a Amazon Quick Suite user. For more information on which roles can access these features, refer to the [Amazon Quick Suite pricing](#) document.

Access Level	Capabilities
Admin Console	- Assign owners for the system chat agent and custom chat chat agents using Manage assets - Control whether users can create chat agents using Custom permissions - Configure instance-wide blocked words and phrases for all chat agents
Amazon Quick Suite user	- Create and maintain custom chat agents (with appropriate permissions) - Configure chat agent personality and response styles and provide reference documents to inform its behavior - Link chat agents with spaces (with dashboards, datasets, topics) as their knowledge source to look for answers - Attach action connectors to use as tools

Access Level	Capabilities
	• Share chat agents with specific users and teams • Interact with chat agents through conversations • Analyze data with chat agent assistance • Receive permissions-aware responses from chat agents

Amazon Quick Suite user permissions for chat agents

What you can do with a chat agent also depends on the permissions you're assigned for it. There are two permission types that users can be assigned:

- **Owner** – Owners can edit, share, and delete the chat agent.
- **Viewer** – Viewers can view and use the chat agent.

Note

If you don't have access to a linked resource as either a viewer or owner, resources added to the chat agent by another owner appear as "Resource unavailable". You can delete these resources as an owner, but you cannot list or add them because resource-level sharing is required.

Quick Suite admins must give users the permission to create chat agents. For information on which roles can create chat agents, refer to the [Amazon Quick Suite pricing](#) documentation. For information on how to provide access to these features, see [Custom permissions](#) in the Quick Suite Admin Guide.

The following table outlines how user permissions determine what you can do with a Amazon Quick Suite chat agent:

Permissions type	Permissions
Owners	• Can access and configure agent behavior (agent persona settings, reference documents) and resources (space and action connectors)

Permissions type	Permissions
	• Can customize chat agent details like title, description, and suggested prompts to improve usability • Can share chat agents with users and groups • Can delete the chat agent
Viewers	• Can't customize the chat agent's details • Can receive responses based on the permissions they have to resources

Custom permissions for chat agents

By default, Amazon Quick Suite enables all new features available for the Amazon Quick Suite account, so that users can access them immediately based on their subscription. You can use [custom permissions](#) to restrict specific features. As an admin, when you create a custom permissions profile, you can create two types of restrictions for chat agents:

- You can completely disable all chat agent functionality for users, including chatting with the default agent, chatting with custom agents, and creating new agents. This can be done by restricting the **Chat agent** capability.
- You can also specifically restrict the ability to create agents without impacting chat with agents. This can be done by creating a custom permission profile and restricting only the feature **Create chat agents**.

Note

If you want your users to be able to chat using the system agent but don't want them to create chat agents, restrict their chat agent creation abilities only.

Agent permissions in Amazon Quick Suite can be configured in various combinations to meet your organization's needs. To help you understand how to configure permissions, the following section provides a use-case driven approach that groups common scenarios:

Use case scenario	What users can do
Agent with all capabilities (No restrictions on chat agents, spaces, knowledge bases, actions, or flows)	• Chat with all default and custom chat agents • Create and customize new chat agents • Connect chat agents to spaces, knowledge bases, and actions • Trigger flows from the chat window from existing conversation with a chat agent • Share chat agents with other users • Use all chat features and integrations
Chat-only access to agents (Agent capabilities enabled, but chat agent creation restricted)	• Chat with existing default and shared custom chat agents • Can't create new custom chat agents • Can't edit or update custom chat agents • Full access to chat interface and chat agent library
Limited access to other capabilities (Chat agents enabled, but specific capabilities restricted)	• Chat with system and custom agents but with reduced functionality • Can't attach knowledge bases, spaces, or actions to an agent's configuration, or select these resources during chat • If spaces restricted: Agents configured with spaces will fall back to LLM knowledge only. Users can't view or select spaces in the chat resource selector • If knowledge bases restricted: Users can't select knowledge bases to use in chat. Any knowledge bases present in spaces either attached to the agent or selected by user during chat won't be used to generate responses • If actions restricted: Users can't attach actions to agents or invoke actions from chat. Any actions preconfigured with agents won't be used to execute tasks • If flows restricted: Chat won't show flows anymore, so users can't invoke flows while chatting with an agent
No chat agent access	• Can't view or access any chat agents • Agent library and navigation are hidden

Use case scenario	What users can do
(Agent capabilities completely restricted)	• Can still access and create other Amazon Quick Suite resources, such as creating spaces for file sharing with teams or flows for structured interactions (as long as those capabilities are not also restricted)

System chat agent

The Amazon Quick Suite system chat agent is automatically created when you sign up for Quick Suite. It serves as the primary interface for all users to interact with their data and perform various tasks within the Quick Suite environment. It comes equipped with default guardrails.

Ownership of the system chat agent is managed by the Quick Suite admin. Admins can disable chatting with chat agents including the system chat agent using [custom permissions](#). If enabled, users with non-admin roles (authors and readers) can view, list, invoke, and use the Quick Suite chat chat agent. Its responses are filtered based on user permissions to assets.

The Quick Suite system chat agent is set to use all chat features out of the box with minimal customization required. While these features are also available to custom chat agents, they must be configured during chat agent creation. The system chat agent includes:

- Default persona as the Quick Suite assistant, with professional tone and response style
- File upload in chat capability enabled
- Large language model (LLM) knowledge chat enabled
- Uses default guardrails configured for chat in admin console
- Access to all spaces, topics, dashboards, knowledge bases, and actions based on user permissions
- Web search capabilities

To learn how to customize a system chat agent as an Admin owner, see [Managing chat agent customization](#). To learn how to control access to the system chat chat agent, see [Manage assets](#).

Custom chat agents

Custom chat agents enable Amazon Quick Suite users to create tailored conversational interfaces for specific business needs. Unlike the system chat agent, which is available to all users with chat

permissions by default, custom chat agents can be selectively shared and configured by authorized users.

Users with chat agent creation capabilities can customize the chat agent's personality, response style, and capabilities through a natural language interface, or through the agent builder configuration flow. These users can also configure these chat agents with specific knowledge sources (like spaces) and workflows (like actions) while maintaining control over sharing and access permissions for each of those resources separately. Or, chat agents can be left unlinked for general purpose use which allows end users to use the agent with all or some of their resources.

Quick Suite admins must give users the permission to create chat agents. For information on which roles can create chat agents, refer to the [Amazon Quick Suite pricing](#) documentation. For information on how to provide or restrict access to these features, see [Custom permissions](#) in the Quick Suite Admin Guide.

Topics

- [Create and preview custom chat agents](#)
- [Share custom chat agents](#)
- [Update custom chat agents](#)
- [Manage access to a custom chat agent](#)
- [View chat agents](#)
- [View chat agent details](#)
- [Duplicate custom chat agents](#)
- [Delete chat agents](#)

Create and preview custom chat agents

You can create custom chat agents in Amazon Quick Suite from the **Chat agents** menu on the Amazon Quick Suite home page. You can preview and test your chat agent as you build it. When you create a custom chat agent, you are assigned as its owner by default. The following procedure outlines how to create a custom chat agent.

Using natural language

To create a custom chat agent using a natural language prompt

1. Log in to the Amazon Quick Suite console.

2. From the left navigation menu, select **Chat agents**, and then select **Create chat agent**.
3. In **New chat agent**, do the following:
 - a. In the text input box, enter a natural language description for what kind of chat agent you want to create. This is the prompt Amazon Quick Suite will use you create your chat agent. You can use the suggested sample instructions on the page to create your chat agent as well.
 - b. Then, select **Generate**. Selecting **Generate** takes the instruction or goal provided and expands it into chat agent instructions and configurations. It also scans the user's available resources (spaces and action connectors) and finds the most relevant matches based on the chat agent intent provided. The expanded builder UI is opened with these configurations for the user to review, customize, test, and launch.
4. On the **Configure chat agent** page, make sure your chat agent has the right settings and customize your chat agent further. Click on **Update preview** to ensure all changes are saved before you try out the chat agent in preview.
5. When you're ready, select **Launch chat agent** to publish your custom chat agent to the chat agent library and use it in chat.

 Note

- Until you select **Launch chat agent**, your chat agent is not available in the chat agent library. If you exit the creation process without launching, the preview version will be deleted and the chat agent won't be saved.
- When you launch a chat agent, it remains private by default. The chat agent only becomes available in other users' libraries after you share it. Once shared, subsequent edits and launches will publish changes to the same chat agent for all users who have access.

Using builder view directly

To create a custom chat agent using builder mode

1. Log in to the Amazon Quick Suite console.
2. From the left navigation menu, select **Chat agents**, and then select **Create chat agent**.

3. In **Agent Creator**, select **Skip**. This opens up the chat agent configuration page from which you can build and customize your chat agent.
4. Add a name for your custom chat agent. This is the name that your chat agent will be identified by.
5. (Optional) Add a description for your custom chat agent that helps users understand the purpose of the chat agent.
6. (Optional) Select an icon for your chat agent.
7. In **Configure chat agent**, customize the following sections:
 - a. In **AGENT PERSONA**, configure your chat agent's personality, identity, tone, and response style:
 - For **Agent identity** – Define the identity of your chat agent. For example, you can give it instructions on what its name and personality are. Amazon Quick Suite will use defaults if left blank.
 - For **Persona instructions** – Add instructions on how your chat agent interacts with users during chat. For example, you can define what your chat agent's primary tasks are. Amazon Quick Suite uses these to customize your chat agent's persona.
 - b. In **Communication style**, for **Pick a response style preset** – Choose a response style. You can choose between the following response style presets or add custom instructions:
 - **Executive** – Optimized for high-level business communication and strategic insights.
 - **Technical** – Optimized for detailed technical explanations and technical problem-solving.
 - **Creative** – Standard configuration for general-purpose interactions.

Each preset has the following settings:

- For **Tone** – Add a natural language prompt to define your agent's tone. Amazon Quick Suite will use it to customize your chat agent's persona.
- For **Response format** – Add a natural language prompt to define your chat agent's response style. Amazon Quick Suite will use it to customize your chat agent's response style. Define the format of your chat agent's responses, for example: "Use bullet points for lists longer than 3 items".

- For **Length** – Specify the length of your chat agent's responses. Define the length of your chat agent's responses, for example: "Keep answers under 100 words".

When you choose a preset to use, Amazon Quick Suite auto-populates the **Tone**, **Response format**, and **Length** instructions for your chat agent's responses based on the style chosen. You can further customize the existing prompts using natural language, or create new ones.

- c. In **Reference documents**, upload files that remain active in your chat agent's memory to guide all interactions. For more information about how reference documents work with other context types, see [Types of agent context](#).

- Select **Upload files** or drag and drop your documents to attach files that will guide your chat agent's responses.
- You can attach documents of .pdf, .txt, .html, .md, .csv, .doc or .docx format. Upto 100,000 characters of text will be extracted and accepted from documents uploaded.

8. (Optional) In **Knowledge sources**, choose between the following options:

- Continue without linking knowledge sources
 - Your chat agent will generate responses from large language model (LLM) knowledge and all Amazon Quick Suite resources the interacting user has access to. You can choose to select a specific knowledge source to focus on during chat. Choose this if you want to create a general purpose chat chat agent.
- To link specific existing spaces
 - Select **Link**.
 - From the **Link spaces** modal, select the spaces you want to link to your chat agent, and then select **Link**.

The system displays a success message to denote successful linking.

If you link a chat agent to a knowledge source, your chat agent will generate responses from large language model (LLM) knowledge and data from linked resources only. You can add up to 1 more knowledge source to this chat agent as temporry context during chat.

- To create and link a new space

- Select **Create**.
- From the **Create space** window that opens, select the assets you want to add to your space and then select **Create**.

If you link a chat agent to a knowledge source, your chat agent will generate responses from large language model (LLM) knowledge and data from linked resources only. You can add up to 1 more knowledge source to this chat agent as temporary context during chat.

 Note

For more information on creating and using spaces, see [Working with spaces in Amazon Quick Suite](#).

The system displays a success message to denote successful space creation. Return to the chat agent creation window.

- In the chat agent creation window, from **Knowledge sources** select **Link**.
- From the **Link spaces** modal, select the space you just created, and then select **Link**.

The system displays a success message to denote successful linking.

9. (Optional) In **Actions**, for **Actions** – Choose between the following options:

- To link specific existing actions
 - Select **Link**.
 - From the **Link action connectors** modal, select the action connectors you want to link to your chat agent, and then select **Next**.
 - In **Actions** select the actions you want to add and then select **Link**.

The system displays a success message to denote successful linking.

- To create an action connector and link new actions
 - Select **Create**.
 - From the **Actions** home page, select **New action**.
 - From the **New action** window, for **Sources** select the action connector you want to add and then select **Next**.
 - In **Actions** review the actions available and then select **Next**.

- In **Connection details**, enter the connection details needed and select **Add**.

 Note

For more information on creating and using action connectors, see [Actions in Amazon Quick Suite](#).

The system displays a success message to denote successful action addition. Return to the chat agent creation window.

- In the chat agent creation window, from **Actions** select **Link**.
- From the **Link action connectors** modal, select the action connector you just created, and then select **Link**.

The system displays a success message to denote successful linking.

10. In **Customization**, do the following:

- a. For **Welcome message** – Add a welcome message for your chat agent to display to your end user.
- b. For **Suggested prompts** – Add example prompts as conversation starters to inform your end users of the chat agents capabilities.

11. Select **Launch chat agent** to create your custom chat agent.

 Note

- Until you select **Launch chat agent**, your chat agent is not available in the chat agent library. If you exit the creation process without launching, the preview version will be deleted and the chat agent won't be saved.
- When you launch a chat agent, it remains private by default. The chat agent only becomes available in other users' libraries after you share it. Once shared, subsequent edits and launches will publish changes to the same chat agent for all users who have access.

Share custom chat agents

When you create a chat agent, you are assigned as its owner by default. You can choose to share a chat agent you own with other Amazon Quick Suite users using the following permissions:

- *Owner* permissions – User can edit, share, use, and delete the chat agent.
- *Viewer* permissions – User can view and use the chat agent.

You can also define access permissions for a chat agent globally, granting all users access to a chat agent and then defining more granular access permissions for specific users and groups.

The following procedure shows you how to share and assign permissions to chat agents.

To share custom chat agents

1. Log in to the Amazon Quick Suite console.
2. From the left navigation menu, select **Chat agents**
3. Then, from the **Actions** column for the chat agent you want to share, select the menu icon, and then select **Share**.
4. In the **Share chat agent** modal that opens, enter the user or group name you want to share the chat agent with. When the user or group name appears, select it. Repeat the action for all groups or users you want to share the chat agent with. The modal displays the users and groups you're adding.
5. Assign permissions and access information for each user or group you're sharing your chat agent with using the dropdown next to the user or group name. You can assign one of two roles:
 - **Owner** – User can edit, share, and delete the chat agent.
 - **Viewer** – User can view and use the chat agent.
6. Select **Share**.

Update custom chat agents

You can update custom chat agents in Amazon Quick Suite from the **Agents** menu on the Amazon Quick Suite home page. The following procedure outlines how to do so.

 Note

You can preview how your chat agent customizations updates perform, as you are configuring the chat agent, in the preview chat interface. Make sure to select **Update preview** before testing a change. Once all changes are tested in preview, select **Launch** to publish the changes to the chat agent currently in use.

To update a custom chat agent

1. Log in to the Amazon Quick Suite console.
2. From the left navigation menu, select **Chat agents**
3. Then, from the **Actions** column for the chat agent you want to share, select the menu icon, and then select **Edit**. This creates a preview version of your chat agent to update and test before publishing.
4. In the **Edit agent**, update your chat agent settings, and then select **Launch**.

Manage access to a custom chat agent

If you're the owner of a chat agent, you can choose to change the user access permissions for a chat agent. You can also remove user and group access to a chat chat agent. The following procedure shows how.

To manage access to custom chat agents

1. Log in to the Amazon Quick Suite console.
2. From the left navigation menu, select **Chat agents**
3. Then, from the **Actions** column for the chat agent you want to share, select the menu icon, and then select **Share**.
4. In the **Share chat agent** modal that opens, select the user and group you want to change permissions for. When the user or group name appears, select it and change the permissions. Repeat the action for all groups or users you want to change permissions for. The modal displays the users and group permissions you're changing. You can take the following three actions:
 - Change to **Owner** – User can edit, share, and delete the chat agent.

- Change to **Viewer** – User can view and use the chat agent.
- **Remove access** – Remove user or group access to a chat agent.

Note

You can also define access permissions to a chat agent globally, granting all users access to a chat agent and then defining more granular access permissions for specific users and groups. To do so, choose **Settings** and then use **Global settings** to turn this feature on or off.

5. Select **Share**.

View chat agents

You can view a list of the chat agents you have created or have access to. The following procedure shows how.

To view custom chat agents

1. Log in to the Amazon Quick Suite console.
2. From the left navigation menu, select **Chat agents**
3. The chat agents you have created or have access to will be displayed on the **Chat agents** home page.

View chat agent details

You can view the details of chat agents that are either created by you or shared with you, including all spaces and actions attached to each chat agent.

To view chat agent details

1. Log in to the Amazon Quick Suite console.
2. From the left navigation menu, select **Chat agents**
3. Then, from the list of chat agents, from the **Name** column, click on the name of the chat agent you want to view details for. A detail modal will open.

Note

You can also view chat agent details by navigating to the **Actions** column for the chat agent you want to share, selecting the menu icon, and then selecting **View chat agent details**.

4. In the chat agent details modal that opens, you will see the following details:
 - a. **Description** – The description for the chat agent.
 - b. **Instruction summary** – A summary of the instructions defined for the chat agent.
 - c. **Created by** – Details on whether the chat agent was created by you or someone else.
 - d. **Last modified** – The time the chat agent was last modified.
 - e. Under **Capabilities**:
 - **Knowledge** – The knowledge sources the chat agent is either attached to or has access to. If your chat agent is linked to a space, you'll be able to see which space it is linked to.
 - **Actions** – The actions that the chat agent is attached or has access to. You can also view this by clicking on **info** next to the agent name in the chat UI agent selector.

Duplicate custom chat agents

You can duplicate existing custom chat agents without building a chat agent from scratch. Users can duplicate chat agents from the list page using the **Duplicate** action button. When you duplicate a chat agent, Amazon Quick Suite creates a new chat agent with name format "[Original name] (copy)", which you can edit. All fields and configurations from the original chat agent are copied to the new version of the chat agent. Once duplicated, you can assign custom user access permissions to the duplicated chat agent and share it.

The following procedure outlines how to duplicate chat agents.

To duplicate custom chat agents

1. Log in to the Amazon Quick Suite console.
2. From the left navigation menu, select **Chat agents**
3. Then, from the **Actions** column for the chat agent you want to share, select the menu icon, and then select **Duplicate**.

You will be redirected to a chat agent duplication modal with prepopulated fields.

4. Make any changes you need to, including adding the right user and group permissions, and linking it to the assets you want and then select **Launch** to launch your duplicated chat agent.

Delete chat agents

You can delete a Amazon Quick Suite chat agent you own. The following procedure shows you how.

To delete a custom chat agent

1. Log in to the Amazon Quick Suite console.
2. From the left navigation menu, select **Chat agents**
3. Then, from the **Actions** column for the chat agent you want to share, select the menu icon, and then select **Delete**.

Chat agent context sources and best practices

When creating custom chat agents in Amazon Quick Suite, you have multiple options to influence chat agent behavior and frame chat agent workflows. The way you configure chat agent context and instructions significantly impacts how your chat agent performs and responds to user interactions. Understanding when to use instructions, reference documents, and spaces is crucial for creating effective chat agents that meet your specific business needs and deliver consistent, reliable results. The following sections outline how to best use chat agent context and knowledge for your use case.

Topics

- [Types of agent context](#)

Types of agent context

Amazon Quick Suite provides three distinct methods for configuring chat agent behavior, each serving different purposes in chat agent development. Understanding how each method processes content and when to use them forms the foundation of effective chat chat agent configuration.

The three context types available in Amazon Quick Suite are:

- **Persona instructions** – High-level guidance that Amazon Quick Suite enhances using generative AI
- **Reference documents** – Documents that provide exact instructions and remain active in the chat agent's memory
- **Spaces** – Collections of searchable Amazon Quick Suite resources and documents from which chat agents retrieve information to answer questions

The following section details what these are and the best practices of using them.

Topics

- [Chat agent persona instructions](#)
- [Chat agent reference documents](#)
- [Chat agents and spaces](#)

Chat agent persona instructions

You provide chat agent persona instructions in the **Persona instructions** field when you create your custom chat agent. These persona instructions are refined using generative AI to ensure that it works with all the chat features and the resources configured for the chat agent. This approach is ideal when you want the system to build upon your foundational direction. When you use persona instructions, Amazon Quick Suite uses generative AI to interpret and expand your persona instructions in the backend, augment simple guidance with additional context and capabilities, optimize the prompt for better chat agent performance, and enhance effectiveness through AI-driven improvements.

Use persona instructions to:

- Provide high level guidance on behavior and goals, which are enhanced by Amazon Quick Suite AI to customize agent responses
- Instruct how to leverage reference documents for detailed interactions
- Provide general guidance on use of spaces and action connectors in its responses

This approach works best for general-purpose chat agents and exploratory use cases where you want to benefit from AI assistance in chat agent design.

Note

The AI enhancement process may modify your original wording and structure to improve chat agent performance. Critical guidelines that must remain unchanged should be placed in reference documents instead.

Chat agent reference documents

Reference documents provide specific process and response templates that the chat agents need to follow. These reference documents are part of the chat agent's permanent context and go hand-in-hand with persona instructions and response style to inform its overall behavior.

Reference documents provide complete control over chat agent behavior by preserving your content exactly as written. This approach is essential when precision and consistency are paramount. When you upload reference documents, Amazon Quick Suite preserves precise wording and specific instructions, maintains exact formatting and structure, provides direct control over chat agent behavior, and ensures consistency across all interactions.

Use reference documents when you need:

- Exact control where precise wording and specific instructions must be preserved
- Brand consistency with specific tone, terminology, or response formats
- Complex workflows with multi-step processes that must follow exact sequences
- Predictable behavior with consistent responses across all interactions

Supported file formats:

- .pdf, .txt, .html, .md – Instructions, templates, and guidelines
- .csv – Reference data, lookup tables, and examples
- .doc, .docx – Detailed process documentation

You can attach upto 10 files of total size 50 MB as reference documents. Amazon Quick Suite extracts text from these documents and applies a 100K character limit after processing.

File prioritization and behavior

- Behavioral frameworks and reference data from files remain active even when users attach alternative content.
- Response templates may be temporarily adjusted if users provide conflicting formatting instructions.

Reference documents serve as the chat agent's operating manual, providing the behavioral framework that remains constant across all user interactions with the chat agent.

Chat agents and spaces

Spaces provide dynamic, searchable knowledge repositories that chat agents query during conversations to retrieve specific information. When you link spaces to chat agents, Amazon Quick Suite chat agents look for answers within the space to generate responses. Spaces excel at managing organizational knowledge that evolves over time, allowing multiple team members to contribute and update information while maintaining appropriate security boundaries through Amazon Quick Suite's permission system.

Use spaces with chat agents for:

- Enabling chat agents to search and retrieve current business data
- Providing chat agents access to organizational knowledge repositories
- Allowing chat agents to respect existing permission boundaries
- Connecting chat agents to collaborative document collections
- Giving chat agents access to regularly updated information
- Enabling chat agents to query live dashboards and datasets

Spaces function as the chat agent's dynamic knowledge layer, enabling real-time access to organizational information while maintaining security boundaries and collaborative workflows.

Use a chat agent

You can use a Amazon Quick Suite chat agent to ask questions and to accomplish your tasks. When you ask a question, the chat agent analyzes the latest approved data shared with you to generate a comprehensive response.

You can ask complex questions in plain language and get a detailed response. You can also use a Quick Suite chat agent to perform tasks for you, such as draft an email message, create a ticket, or use flows.

To use custom chat agents

1. Log in to the Quick Suite console.
2. From the left navigation menu, select **Chat agents**
3. Then, from the **Actions** column for the chat agent you want to share, select the menu icon, and then select **Chat**.

This will open a chat interface that you use to interact with your chat agent.

You can interact with a chat agent in the following ways:

- View and access example prompts to demonstrate possible questions and tasks
- Get contextual responses from organizational data sources
- Continue or start new conversations
- Analyze data, including summarization and comparative analysis
- Take actions like scheduling meetings and creating tickets
- View source citations to responses with clickable numbered references
- Upload files
- View conversation history

Note

To learn more about chatting using chat agents, see [Using Amazon Quick Suite chat](#).

Organize, collaborate, and share resources with spaces in Amazon Quick Suite

A space in Amazon Quick Suite is a collection of data and Quick Suite resources scoped for a particular team or domain. You can use spaces to aggregate and organize files, dashboards, topics, knowledge bases, and application actions into a unified and customizable knowledge center for your team. Spaces integrate seamlessly with Quick Suite agents for contextual conversations and are designed to scale across personal, team, and cross-team use cases.

Spaces allow your team to get the most relevant results from conversational agents and other AI tools inside Quick Suite by grounding the results with only data relevant for your task or domain. Multiple people on the team can contribute to the knowledge inside a space; this reduces data silos and streamline information discovery.

Note

Quick Suite resources added to a space respect user access permissions. If a user can't access a resource outside a space, they won't be able to access it inside a space. Files uploaded to a space are always available to everyone with access to the space.

You can use spaces to accomplish tasks like:

- Aggregating and sharing customer feedback, allowing any team member to ask questions and get summaries
- Onboarding new team members with team documents, manuals, and processes
- Analyzing hundreds of annual reports
- Summarizing and search for action items from audio meeting recordings

Topics

- [Amazon Quick Suite user interactions with spaces](#)
- [Amazon Quick Suite user permission for spaces](#)
- [Understanding file upload status](#)
- [Limitations of spaces](#)
- [Creating a space](#)

- [Managing a space](#)

Amazon Quick Suite user interactions with spaces

The following table shows how each user type in Amazon Quick Suite interacts with spaces.

User type	Capabilities
Administrators	• Control whether users with space creation capabilities can create spaces
Author Pro and Reader Pro users	• Create and maintain spaces • Add topics, dashboards, knowledge bases, and application actions to spaces • Upload files directly into spaces • Attach an agent to a space • Interact with space data through chat • Share spaces

Amazon Quick Suite user permission for spaces

What you can do with a space also depends on the permissions you're assigned for it. There are two permission types that users can be assigned:

- **Owner** – Owners can create, edit, share, and delete a space.
- **Viewer** – Viewers can view ask questions, and download files from spaces.

The following table outlines how user permissions determine what you can do with a Amazon Quick Suite space:

Permissions type	Permissions
Owners	• Create a space • Upload files to a space

Permissions type	Permissions
	• Share spaces with others • Link and unlink Amazon Quick Suite resources (topics, dashboards, knowledge bases, and application actions) to a space • Delete a space As an owner, you can designate another user co-owner of a space. If you do so, the user can undertake all post-creation management actions an owner can.
Viewers	• Download files uploaded into a space • Ask questions from data inside space • Use a specific space as context for an agent • Search for a space by name • Access a space using a direct URL • View a preset list sample questions to help get started on using the space

Understanding file upload status

When you upload files, you should see the following status messages:

Uploading

Files are being uploaded into your computer.

Processing

Files are being processed.

Text ready

The text content of the document is ready to be queried. However, the document's images and tables are still processing.

Ready

The text and media of the document has been processed.

Deleting

The document is being deleted.

Limitations of spaces

The following list outlines the limitations of spaces:

- A space can't contain other spaces.
- Users can't access Amazon Quick Suite resources (topics, dashboards) in a space if they don't have access to the resource outside it. Sharing a space doesn't automatically give access to users to assets within it.
- If an Amazon Quick Suite resource is linked to a space is deleted outside the space, users of the space won't get notified of the deletion. Deleted assets will only display a message indicating unavailability.
- Files uploaded to spaces will be rejected if your index data storage capacity is full. Reach out to your system administrator if this happens.
- Files uploaded to a space must be one of the following formats: .html, .ppt, .docx, .txt, .pdf, .csv, .xlsx, .xml, .jpeg, .svg, .json, .md, .rtf, .xslt, .mp3, .wav, .m4a, .flac, .ogg, .mp4, .mov, .m4v.

Note

You can only upload JPG images to spaces. PNG is not supported at this time.

- File uploads are limited to 5 MB for Excel and CSV files, and 50 MB for all other supported file types. You will get a "character limit exceeded" error when you upload larger documents.
- Spaces support up to 100 files (compared to 20 files in regular chat conversations), as long as the total space storage is less than 1 GB.
- You can add at most 20 resources of a specific resource type (dashboard, topic, knowledge base, or action) to a space.

Creating a space

You can create spaces in Amazon Quick Suite from the **Spaces** menu on the Quick Suite home page. You can directly add files to your space or connect existing resources to it when you create

it. If you upload a file that exceeds the file size limits, or is of an unsupported file type, Quick Suite displays an error and asks you to try again. When you create a space, you are assigned as its owner by default. The following procedure outlines how to create a space.

To create a space

1. Log in to the Quick Suite console.
2. From the left navigation menu, select **Spaces**, and then select **Create space**.
3. On the space creation page that opens, do the following:
 - a. (Optional) Enter a name for your space.
 - b. (Optional) Enter a description for your space.
 - c. Select **Add knowledge** to begin adding content to your space.

From the dropdown menu, you can select from available resource types including **Dashboards**, **Topics**, **Knowledge bases**, **Actions**, and **File uploads**.

Note

You can preview a resource before you choose to add it by clicking on its name.

After you finish adding resources, **Space knowledge** displays a list of all assets and files added to your space.

After your space is created, you can select the sparkle icon from the top left corner of the screen and start interacting with your space using any available agent.

Managing a space

You can perform the following actions for spaces.

Topics

- [Viewing a list of spaces](#)
- [Editing a space](#)
- [Adding knowledge to a space](#)
- [Viewing a list of content inside a space](#)

- [Using a space](#)
- [Sharing a space](#)
- [Deleting a space](#)
- [Removing knowledge from a space](#)
- [Adding data from external sources into your space](#)
- [Adding application actions to a space](#)

Viewing a list of spaces

You can view a list of spaces created by—or shared with—you.

To view a list of spaces

1. Log in to the Amazon Quick Suite console.
2. From the left navigation menu, select **Spaces**.

You will see a list of spaces displayed. The list also displays the name, description, owner, and last modified date for each space. You can search for a space, and filter for recent spaces.

Editing a space

You can edit an Amazon Quick Suite space. When you edit a space, you can perform the following actions:

- Edit space description and space name
- Add and delete any files and link or unlink any resources (dashboards, topics, knowledge bases, and application actions) that you've added.

To edit a space

1. Log in to the Amazon Quick Suite console.
2. From the left navigation menu, select **Spaces**.
3. In **Spaces**, select the space you want to edit.
4. Inside the space, do any of the following:

- a. To edit **Space name** or **Space description**, double click on the name or description field. Then, select the edit icon to update the description.
- b. In the **All knowledge** section, select **Add knowledge** to add new content.
- c. In **Amazon Quick Suite resources** – Link and unlink resources as needed.

Adding knowledge to a space

You can add knowledge—files and Amazon Quick Suite resources—to a space after it's been created.

To add knowledge to an existing space

1. Log in to the Amazon Quick Suite console.
2. From the left navigation menu, select **Spaces**, and then select your space from the list of spaces displayed.
3. Select **Add knowledge** to begin adding content to your space.

From the dropdown menu, you can select from available resource types including **Dashboards**, **Topics**, **Knowledge bases**, **Actions**, and **File uploads**.

Note

You can preview a resource before you choose to add it by clicking on its name.

After you finish adding resources, **Space knowledge** displays a list of all resources and files added to your space.

Viewing a list of content inside a space

You can view a list of content inside a space.

To view a list of content within a space

1. Log in to the Amazon Quick Suite console.
2. From the left navigation menu, select **Spaces**.
3. Select the name of the space whose content you want to view from the list of spaces.

On the space detail page, you will see the **Space knowledge** section displaying a list of files uploaded and resources linked to the space.

Using a space

After you finish creating your space and adding knowledge to it, you can start interacting with it through agents. You can interact with an agent while in a space in the following ways:

- Interact with a space using any agent

Note

When you switch between spaces or open up a space during an on-going chat Amazon Quick Suite detects your space and focuses your agent to it.

- Interact with a specific knowledge source within a space using an agent

To use an agent to chat with a space

1. Log in to the Amazon Quick Suite console.
2. Then, do one of the following:
 1. From the left navigation menu, select **Spaces**.
 2. In **Spaces**, select the space name.
 3. In the space that opens, start chatting with your data using the chat interface.
1. Use the sparkle chat icon to open the chat interface.
2. From the resource selector menu, select **All resources**, and then select **Choose a resource**.
3. From the **Add a resource** menu, from **Spaces**, select the space you want to chat with. Then, start chatting.

Sharing a space

You can share a space with any other Amazon Quick Suite user. When you share a space with a user, you must also choose the permission levels that user has to the space—whether owner or viewer.

To share a space

1. Log in to the Amazon Quick Suite console.
2. From the left navigation menu, select **Spaces**, and select the name of the space you want to share. The space should open.
3. From the top right corner of your space menu, select **Share**
4. In the **Share space** modal that opens, add the users and groups you want. Choose the permissions type for each user and group you add.

Note

You can also share a space from the space home page. Locate the space you want to share, select the menu icon in the **Actions**, and then select **Share**. In the **Share space** modal that opens, add users and groups and set their permissions.

Deleting a space

You can delete a Amazon Quick Suite space. Deleting a space deletes all files associated with the space. Deleting a space unlinks assets associated with the space without deleting the assets. You can only delete a space you own.

To delete a space

1. Log in to the Amazon Quick Suite console.
2. From the left navigation menu, select **Spaces**, and select the name of the space you want to share. The space should open.
3. From the top right corner of your space menu, select **Delete**, and then select **Delete** from the modal that opens.

Note

You can also delete a space from the space home page. Locate the space you want to delete, select the menu icon in the **Actions**, and then select **Delete**. In the **Delete space** modal that opens, select **Delete**.

Removing knowledge from a space

You can remove knowledge (remove files and unlink assets) from a space. When you remove files, they are removed from the Amazon Quick Suite index in which they are stored. When you remove an asset, it unlinks it from the space without deleting it from the Amazon Quick Suite system.

To remove knowledge from an existing space

1. Log in to the Amazon Quick Suite console.
2. From the left navigation menu, select **Spaces**, and then select your space from the list of spaces displayed.
3. In the **Space knowledge** section:
 - a. In **Files uploaded** – From the **Actions** menu select **Delete**.
 - b. In **Amazon Quick Suite resources** – From the **Actions** menu select **Remove from this space**.

Adding data from external sources into your space

You can pull in data from Google Drive, OneDrive, Confluence, SharePoint, web pages, and Amazon S3 into your spaces to expand the context of your chat and research agents.

To add data from external sources

1. Create a space and give it a name.
2. Select the **Knowledge bases** section in the space knowledge list.
3. Select **Add knowledge bases**.
4. In **All knowledge bases** you will see application data integrations that were created for you by others in your organization. You can select one to pull into your space, or create a new one.

Note

To learn how to create a knowledge base, refer to [Working with integrations](#).

Adding application actions to a space

You can use application actions to read and write data to external SaaS applications and MCP servers from within a space. When you add actions to a space, only those actions will be considered when you select the space as your chat context.

Note

Only authors can create action connectors.

To use actions in a space

1. Create a space and give it a name.
2. Select the **Actions** section in the space knowledge list.
3. Select **Add action**.
4. Select **All actions**, and choose an action that others have created in your organization.

Note

To learn how to create an action connector, refer to [Working with integrations](#).

Visualize, analyze, and share data with analyses, dashboards, and reports in Amazon Quick Sight

Amazon Quick Sight is a comprehensive business intelligence service that enables you to transform raw data into meaningful insights through interactive visualizations, dashboards, and reports. Whether you're connecting to databases, preparing datasets, creating analyses, or sharing dashboards with stakeholders, Amazon Quick Sight provides the tools you need to make data-driven decisions.

This section covers the complete Amazon Quick Sight workflow, from initial data connection through final report sharing. You'll learn how to connect to various data sources, prepare and transform your data, create compelling visualizations, build interactive dashboards, and leverage generative BI capabilities to accelerate your analytics workflow. Each topic builds upon the previous one, providing a comprehensive guide to maximizing your use of Amazon Quick Sight's powerful features.

Topics

- [Connecting to data in Amazon Quick Sight](#)
- [Refreshing data in Amazon Quick Sight](#)
- [Preparing data in Amazon Quick Sight](#)
- [Analyses and reports: Visualizing data in Amazon Quick Sight](#)
- [Sharing and subscribing to data in Amazon Quick Sight with dashboards and reports](#)
- [Exploring interactive dashboards in Amazon Quick Sight](#)
- [Gaining insights with machine learning \(ML\) in Amazon Quick Sight](#)
- [Troubleshooting Amazon Quick Sight](#)
- [Developing with Amazon Quick Sight](#)

Connecting to data in Amazon Quick Sight

People in many different roles use Amazon Quick Sight to help them do analysis and advanced calculations, design data dashboards, embed analytics, and make better-informed decisions. Before any of that can happen, someone who understands your data needs to add it to a [Quick Sight dataset](#). Quick Sight supports direct connections and uploads from a variety of [data sources](#).

Capabilities and use cases

Amazon Quick Suite Standard edition capabilities

After your data is available in Quick Suite Standard edition, you can do the following:

- Transform the dataset with field formatting, hierarchies, data type conversions, and calculations.
- Create one or more data analyses based on your newly created dataset.
- Share your analysis with other people so they can help design it.
- Add charts, graphs, more datasets, and multiple pages (called sheets) to your data analysis.
- Create visual appeal with customized formatting and themes.
- Make them interactive by using parameters, controls, filters, and custom actions.
- Combine data from multiple data sources, and then build new hierarchies for drilling down and calculations only available during analytics, like aggregations, window functions, and more.
- Publish your analysis as an interactive data dashboard.
- Share the dashboard so other people can use the dashboard, even if they don't use the analysis that it's based on.
- Add more data to create more analyses and dashboards.

Amazon Quick Suite Enterprise edition capabilities

After your data is available in Quick Suite Enterprise edition, you can do different things depending on your role. If you can build datasets, design analyses, and publish dashboards, you can do all of the things people using Standard edition can do.

In addition, these are some examples of additional tasks that you can do:

- Create analyses that use Quick Sight insights, including machine learning (ML) powered insights for forecasting, anomaly and outlier detection, and key driver identification.
- Design narrative insights with text, colors, images, and calculations.
- Add data from virtual private clouds (VPCs) and on-premises data sources, with data encryption at rest.
- Control access in datasets by adding row and column level security.
- Refresh imported datasets every hour.

- Share emailed reports.

Application development

If you develop applications or use the Amazon SDKs and Amazon Command Line Interface (Amazon CLI), you can do the following and more:

- Add embedded analytics and embedded interactive dashboards to websites and applications.
- Use API operations to manage data sources and datasets.
- Refresh imported data more frequently by using the data ingestion API operations.
- Script, transfer, and make templates from analyses and dashboards by using API operations.
- Programmatically assign people to security roles based on settings managed by system administrators.

Administrative functions in Quick Suite

If you perform administrative functions in Quick Suite, you can do the following and more:

- Manage security with shared folders to organize your teams' work and help them collaborate using dashboards, analytics, and datasets.
- Add Quick Sight to your VPC to enable access to data in VPC and on-premises data sources.
- Protect sensitive data with finely grained access control to Amazon data sources.
- Manually assign people to the Quick Suite author security role so they can prepare datasets, design analytics, and publish data dashboards at a fixed cost per month.
- Manually assign people to the Quick Suite reader security role so they can securely interact with published data dashboards on a pay-per-session basis.

Dashboard subscription

If you subscribe to dashboards, you can do the following:

- Use and subscribe to interactive dashboards designed by your team of experts.
- Enjoy a simplified uncluttered interface.
- View dashboard snapshots in email.
- Focus on making decisions with the data at your fingertips.

After you connect to or import data, you create a dataset to shape and prepare data to share and reuse. You can view your available datasets on the **Data** page, which you reach by choosing **Data** on the Amazon Quick Sight start page. You can view available data sources and create a new dataset

on the **Create a Data Set** page, which you reach by choosing **Create** then **New dataset** on the **Data** page.

Topics

- [Supported data sources](#)
- [Connect to your data with integrations and datasets](#)
- [Data source quotas](#)
- [Supported data types and values](#)
- [Working with datasets](#)
- [Working with data sources in Amazon Quick Sight](#)

Supported data sources

Amazon Quick Sight supports a variety of data sources that you can use to provide data for analyses. The following data sources are supported.

Connecting to relational data

You can use any of the following relational data stores as data sources for Amazon Quick Sight:

- Amazon Athena
- Amazon Aurora
- Amazon OpenSearch Service
- Amazon Redshift
- Amazon Redshift Spectrum
- Amazon S3
- Amazon S3 Analytics
- Apache Impala
- Apache Spark 2.0 or later
- Amazon IoT Analytics
- Databricks (E2 Platform only) on Spark 1.6 or later, up to version 3.0
- Exasol 7.1.2 or later
- Google BigQuery
- MariaDB 10.0 or later

- Microsoft SQL Server 2012 or later
- MySQL 5.7 or later

Note

Effective October 2023, the MySQL community has deprecated support for MySQL version 5.7. This means that Amazon Quick Sight will no longer support new features, enhancements, bug fixes, or security patches for MySQL 5.7. Support for existing query workload will take place at a best effort basis. Quick Sight customers can still use MySQL 5.7 datasets with Quick Sight, but we encourage customers to upgrade their MySQL databases (DB) to major version 8.0 or higher. To see the statement provided by Amazon RDS, see [Amazon RDS Extended Support opt-in behavior is changing. Upgrade your Amazon RDS for MySQL 5.7 database instances before February 29, 2024 to avoid potential increase in charges.](#)

Amazon RDS has updated their security settings for Amazon RDS MySQL 8.3. Any connections from Quick Sight to Amazon RDS MySQL 8.3 are SSL-enabled by default. This is the only option available for MySQL 8.3. connections.

- Oracle 12c or later
- PostgreSQL 9.3.1 or later

Note

SCRAM based authentication to PostgreSQL from Amazon Quick Sight is supported for the following connectors: RDS hosted PostgreSQL, Aurora PostgreSQL, and Vanilla PostgreSQL. If the appropriate PostgreSQL engine version is used, and the correct configurations in PostgreSQL for SCRAM are set up, no additional configurations are needed in Quick Sight. If you are still experiencing issues establishing a SCRAM authentication to PostgreSQL from Quick Sight, please create a support ticket.

- Presto 0.167 or later
- Snowflake
- Starburst
- Trino
- Teradata 14.0 or later
- Timestream

Note

You can access additional data sources not listed here by linking or importing them through supported data sources.

Amazon Redshift clusters, Amazon Athena databases, and Amazon RDS instances must be in Amazon. Other database instances must be in one of the following environments to be accessible from Amazon Quick Sight:

- Amazon EC2
- Local (on-premises) databases
- Data in a data center or some other internet-accessible environment

For more information, see [Infrastructure security in Amazon Quick Suite](#).

Importing file data

You can use files in Amazon S3 or on your local (on-premises) network as data sources. Quick Sight supports files in the following formats:

- CSV and TSV – Comma-delimited and tab-delimited text files
- ELF and CLF – Extended and common log format files
- JSON – Flat or semistructured data files
- XLSX – Microsoft Excel files

Quick Sight supports UTF-8 file encoding, but not UTF-8 (with BOM).

Files in Amazon S3 that have been compressed with zip, or gzip (www.gzip.org), can be imported as-is. If you used another compression program for files in Amazon S3, or if the files are on your local network, remove compression before importing them.

JSON data

Amazon Quick Sight natively supports JSON flat files and JSON semistructured data files.

You can either upload a JSON file or connect to your Amazon S3 bucket that contains JSON data. Amazon Quick Sight automatically performs schema and type inference on JSON files and

embedded JSON objects. Then it flattens the JSON, so you can analyze and visualize application-generated data.

Basic support for JSON flat-file data includes the following:

- Inferring the schema
- Determining data types
- Flattening the data
- Parsing JSON (JSON embedded objects) from flat files

Support for JSON file structures (.json) includes the following:

- JSON records with structures
- JSON records with root elements as arrays

You can also use the `parseJson` function to extract values from JSON objects in a text file. For example, if your CSV file has a JSON object embedded in one of the fields, you can extract a value from a specified key-value pair (KVP). For more information on how to do this, see [parseJson](#).

The following JSON features aren't supported:

- Reading JSON with a structure containing a list of records
- List attributes and list objects within a JSON record; these are skipped during import
- Customizing upload or configuration settings
- `parseJSON` functions for SQL and analyses
- Error messaging for invalid JSON
- Extracting a JSON object from a JSON structure
- Reading delimited JSON records

You can use the `parseJson` function to parse flat files during data preparation. This function extracts elements from valid JSON structures and lists.

The following JSON values are supported:

- JSON object
- String (double quoted)

- Number (integer and float)
- Boolean
- NULL

Software as a service (SaaS) data

Quick Sight can connect to a variety of Software as a Service (SaaS) data sources either by connecting directly or by using Open Authorization (OAuth).

SaaS sources that support direct connection include the following:

- Jira
- ServiceNow

SaaS sources that use OAuth require that you authorize the connection on the SaaS website. For this to work, Quick Sight must be able to access the SaaS data source over the network. These sources include the following:

- Adobe Analytics
- GitHub
- Salesforce

You can use reports or objects in the following editions of Salesforce as data sources for Amazon Quick Sight:

- Enterprise Edition
- Unlimited Edition
- Developer Edition

Local data sources

To connect to on premises data sources, you need to add your data sources and a Quick Suite-specific network interface to Amazon Virtual Private Cloud (Amazon VPC). When configured properly, a VPC based on Amazon VPC resembles a traditional network that you operate in your own data center. It enables you to secure and isolate traffic between resources. You define and control the network elements to suit your requirements, while still getting the benefit of cloud networking and the scalable infrastructure of Amazon.

For detailed information, see [Infrastructure security in Amazon Quick Suite](#).

Connect to your data with integrations and datasets

You can connect Amazon Quick Sight to different types of data sources. This includes data residing in Software-as-a-Service (SaaS) applications, flat files stored in Amazon S3 buckets, data from third-party services like Salesforce, and query results from Athena. Use the following examples to learn more about the requirements for connecting to specific data sources.

Topics

- [Creating a dataset using Amazon Athena data](#)
- [Using Amazon OpenSearch Service with Amazon Quick Sight](#)
- [Creating a dataset using Amazon S3 files](#)
- [Creating a data source using Apache Spark](#)
- [Using Databricks in Quick Sight](#)
- [Creating a dataset using Google BigQuery](#)
- [Creating a dataset using a Google Sheets data source](#)
- [Creating a dataset using an Apache Impala data source](#)
- [Creating a dataset using a Microsoft Excel file](#)
- [Creating a data source using Presto](#)
- [Using Snowflake with Amazon Quick Sight](#)
- [Using Starburst with Amazon Quick Sight](#)
- [Creating a data source and data set from SaaS sources](#)
- [Creating a dataset from Salesforce](#)
- [Using Trino with Amazon Quick Sight](#)
- [Creating a dataset using a local text file](#)
- [Using Amazon Timestream data with Amazon Quick Sight](#)

Creating a dataset using Amazon Athena data

Use the following procedure to create a new dataset that connects to Amazon Athena data or to Athena Federated Query data.

To connect to Amazon Athena

1. Begin by creating a new dataset. Choose **Data** from the navigation pane at left.
2. Choose **Create**, then choose **New dataset**.
3. a. To use an existing Athena connection profile (common), choose the card for the existing data source that you want to use. Choose **Select**.

Cards are labeled with the Athena data source icon and the name provided by the person who created the connection.

- b. To create a new Athena connection profile (less common), use the following steps:
 1. Choose **New data source**, then choose the **Athena** data source card.
 2. Choose **Next**.
 3. For **Data source name**, enter a descriptive name.
 4. For **Athena workgroup**, choose your workgroup.
 5. Choose **Validate connection** to test the connection.
 6. Choose **Create data source**.
 7. (Optional) Select an IAM role ARN for queries to run as.
4. On the **Choose your table** screen, do the following:
 - a. For **Catalog**, choose one of the following:
 - If you are using Athena Federated Query, choose the catalog you want to use.
 - Otherwise, choose **AwsDataCatalog**.
 - b. Choose one of the following:
 - To write a SQL query, choose **Use custom SQL**.
 - To choose a database and table, choose your catalog that contains your databases from the dropdown under **Catalog**. Then, choose a database from the dropdown under **Database** and choose a table from the **Tables** list that appears for your database.

If you don't have the right permissions, you receive the following error message: "You don't have sufficient permissions to connect to this dataset or run this query." Contact your Quick Suite administrator for assistance. For more information, see [Authorizing connections to Amazon Athena](#).

5. Choose **Edit/preview data**.
6. Create a dataset and analyze the data using the table by choosing **Visualize**. For more information, see [Analyses and reports: Visualizing data in Amazon Quick Sight](#).

Using Amazon OpenSearch Service with Amazon Quick Sight

Following, you can find how to connect to your Amazon OpenSearch Service data using Amazon Quick Sight.

Creating a new Quick Sight data source connection for OpenSearch Service

Following, you can find how to connect to OpenSearch Service

Before you can proceed, Amazon Quick Sight needs to be authorized to connect to Amazon OpenSearch Service. If connections aren't enabled, you get an error when you try to connect. A Quick Sight administrator can authorize connections to Amazon resources.

To authorize Quick Sight to initiate a connection to OpenSearch Service

1. Open the menu by clicking on your profile icon at top right, then choose **Manage Quick Suite**. If you don't see the **Manage Quick Suite** option on your profile menu, ask your Amazon Quick Suite administrator for assistance.
2. Choose **Security & permissions, Add or remove**.
3. Enable the option for **OpenSearch**.
4. Choose **Update**.

After OpenSearch Service is accessible, you create a data source so people can use the specified domains.

To connect to OpenSearch Service

1. Begin by creating a new dataset. Choose **Data** from the navigation pane at left, then choose **Create and New Dataset**.
2. Choose the **Amazon OpenSearch** data source card.
3. For **Data source name**, enter a descriptive name for your OpenSearch Service data source connection, for example **OpenSearch Service ML Data**. Because you can create many datasets from a connection to OpenSearch Service, it's best to keep the name simple.

4. For **Connection type**, choose the network you want to use. This can be a virtual private cloud (VPC) based on Amazon VPC or a public network. The list of VPCs contains the names of VPC connections, rather than VPC IDs. These names are defined by the Quick Suite administrator.
5. For **Domain**, choose the OpenSearch Service domain that you want to connect to.
6. Choose **Validate connection** to check that you can successfully connect to OpenSearch Service.
7. Choose **Create data source** to proceed.
8. For **Tables**, choose the one you want to use, then choose **Select** to continue.
9. Do one of the following:
 - To import your data into the Quick Sight in-memory engine (called SPICE), choose **Import to SPICE for quicker analytics**. For information about how to enable importing OpenSearch data, see [Authorizing connections to Amazon OpenSearch Service](#).
 - To allow Quick Sight to run a query against your data each time you refresh the dataset or use the analysis or dashboard, choose **Directly query your data**.

To enable autorefresh on a published dashboard that uses OpenSearch Service data, the OpenSearch Service dataset needs to use a direct query.

10. Choose **Edit/Preview** and then **Save** to save your dataset and close it.

Managing permissions for OpenSearch Service data

The following procedure describes how to view, add, and revoke permissions to allow access to the same OpenSearch Service data source. The people that you add need to be active users in Quick Sight before you can add them.

To edit permissions on a data source

1. Choose **Data** at left, then scroll down to find the data source card for your Amazon OpenSearch Service connection. An example might be US Amazon OpenSearch Service Data.
2. Choose the **Amazon OpenSearch** dataset.
3. On the dataset details page that opens, choose the **Permissions** tab.

A list of current permissions appears.

4. To add permissions, choose **Add users & groups**, then follow these steps:

- a. Add users or groups to allow them to use the same dataset.
 - b. When you're finished adding everyone that you want to add, choose the **Permissions** that you want to apply to them.
5. (Optional) To edit permissions, you can choose **Viewer** or **Owner**.
 - Choose **Viewer** to allow read access.
 - Choose **Owner** to allow that user to edit, share, or delete this Quick Sight dataset.
6. (Optional) To revoke permissions, choose **Revoke access**. After you revoke someone's access, they can't create new datasets from this data source. However, their existing datasets still have access to this data source.
7. When you are finished, choose **Close**.

Adding a new Quick Sight dataset for OpenSearch Service

After you have an existing data source connection for OpenSearch Service, you can create OpenSearch Service datasets to use for analysis.

To create a dataset using OpenSearch Service

1. From the start page, choose **Data, Create, New dataset**.
2. Scroll down to the data source card for your OpenSearch Service connection. If you have many data sources, you can use the search bar at the top of the page to find your data source with a partial match on the name.
3. Choose the **Amazon OpenSearch** data source card, and then choose **Create data set**.
4. For **Tables**, choose the OpenSearch Service index that you want to use.
5. Choose **Edit/Preview**.
6. Choose **Save** to save and close the dataset.

Adding OpenSearch Service data to an analysis

After you have an OpenSearch Service dataset available, you can add it to a Quick Sight analysis. Before you begin, make sure that you have an existing dataset that contains the OpenSearch Service data that you want to use.

To add OpenSearch Service data to an analysis

1. Choose **Analyses** at left.
2. Do one of the following:
 - To create a new analysis, choose **New analysis** at right.
 - To add to an existing analysis, open the analysis that you want to edit.
 - Choose the pencil icon near at top left.
 - Choose **Add data set**.
3. Choose the OpenSearch Service dataset that you want to add.

For information on using OpenSearch Service in visualizations, see [Limitations for using OpenSearch Service](#).

4. For more information, see [Working with analyses](#).

Limitations for using OpenSearch Service

The following limitations apply to using OpenSearch Service datasets:

- OpenSearch Service datasets support a subset of the visual types, sort options, and filter options.
- To enable autorefresh on a published dashboard that uses OpenSearch Service data, the OpenSearch Service dataset needs to use a direct query.
- Multiple subquery operations aren't supported. To avoid errors during visualization, don't add multiple fields to a field well, use one or two fields per visualization, and avoid using the **Color** field well.
- Custom SQL isn't supported.
- Crossdataset joins and self joins aren't supported.
- Calculated fields aren't supported.
- Text fields aren't supported.
- The "other" category isn't supported. If you use an OpenSearch Service dataset with a visualization that supports the "other" category, disable the "other" category by using the menu on the visual.

Creating a dataset using Amazon S3 files

To create a dataset using one or more text files (.csv, .tsv, .clf, or .elf) from Amazon S3, create a manifest for Quick Sight. Quick Sight uses this manifest to identify the files that you want to use and to the upload settings needed to import them. When you create a dataset using Amazon S3, the file data is automatically imported into [SPICE](#).

You must grant Quick Sight access to any Amazon S3 buckets that you want to read files from. For information about granting Quick Sight access to Amazon resources, see [Configuring Amazon Quick Sight access to Amazon data sources](#).

Topics

- [Supported formats for Amazon S3 manifest files](#)
- [Creating Amazon S3 datasets](#)
- [Datasets using S3 files in another Amazon account](#)

Supported formats for Amazon S3 manifest files

You use JSON manifest files to specify files in Amazon S3 to import into Quick Sight. These JSON manifest files can use either the Quick Sight format described following or the Amazon Redshift format described in [Using a manifest to specify data files](#) in the *Amazon Redshift Database Developer Guide*. You don't have to use Amazon Redshift to use the Amazon Redshift manifest file format.

If you use an Quick Sight manifest file, it must have a .json extension, for example `my_manifest.json`. If you use an Amazon Redshift manifest file, it can have any extension.

If you use an Amazon Redshift manifest file, Quick Sight processes the optional mandatory option as Amazon Redshift does. If the associated file isn't found, Quick Sight ends the import process and returns an error.

Files that you select for import must be delimited text (for example, .csv or .tsv), log (.clf), or extended log (.elf) format, or JSON (.json). All files identified in one manifest file must use the same file format. Plus, they must have the same number and type of columns. Quick Sight supports UTF-8 file encoding, but not UTF-8 with byte-order mark (BOM). If you are importing JSON files, then for `globalUploadSettings` specify `format`, but not `delimiter`, `textqualifier`, or `containsHeader`.

Make sure that any files that you specify are in Amazon S3 buckets that you have granted Quick Sight access to. For information about granting Quick Sight access to Amazon resources, see [Configuring Amazon Quick Sight access to Amazon data sources](#).

Manifest file format for Quick Sight

Quick Sight manifest files use the following JSON format.

```
{
  "fileLocations": [
    {
      "URIs": [
        "uri1",
        "uri2",
        "uri3"
      ]
    },
    {
      "URIPrefixes": [
        "prefix1",
        "prefix2",
        "prefix3"
      ]
    }
  ],
  "globalUploadSettings": {
    "format": "JSON",
    "delimiter": ",",
    "textqualifier": "'",
    "containsHeader": "true"
  }
}
```

Use the fields in the `fileLocations` element to specify the files to import, and the fields in the `globalUploadSettings` element to specify import settings for those files, such as field delimiters.

The manifest file elements are described following:

- **fileLocations** – Use this element to specify the files to import. You can use either or both of the URIs and URIPrefixes arrays to do this. You must specify at least one value in one or the other of them.
- **URIs** – Use this array to list URIs for specific files to import.

Quick Sight can access Amazon S3 files that are in any Amazon Web Services Region. However, you must use a URI format that identifies the Amazon Region of the Amazon S3 bucket if it's different from that used by your Quick Suite account.

URIs in the following formats are supported.

URI format	Example	Comments
<code>https://s3.amazonaws.com/<bucket name>/<file name></code>	<code>https://s3.amazonaws.com/amzn-s3-demo-bucket/data.csv</code>	
<code>s3://<bucket name>/<file name></code>	<code>s3://amzn-s3-demo-bucket/data.csv</code>	
<code>https://<bucket name>.s3.amazonaws.com/<file name></code>	<code>https://<i>amzn-s3-demo-bucket</i>.s3.amazonaws.com/<i>data.csv</i></code>	
<code>https://s3-<region name>.amazonaws.com/<bucket name>/<file name></code>	<code>https://<i>s3-us-east-1.amazonaws.com</i> /<i>amzn-s3-demo-bucket</i> /<i>data.csv</i></code>	This URI type identifies the Amazon Web Services Region for the Amazon S3 bucket.

URI format	Example	Comments
<code>https://<bucket name>.s3-<region name>.amazonaws.com/<file name></code>	<code>https://<i>amzn-s3-demo-bucket</i>.<i>s3-us-east-1</i>.<i>amazonaws.com</i>/<i>data.csv</i></code>	This URI type identifies the Amazon Web Services Region for the Amazon S3 bucket.

- **URIPrefixes** – Use this array to list URI prefixes for S3 buckets and folders. All files in a specified bucket or folder are imported. Quick Sight recursively retrieves files from child folders.

Quick Sight can access Amazon S3 buckets or folders that are in any Amazon Web Services Region. Make sure to use a URI prefix format that identifies the S3 bucket's Amazon Web Services Region if it's different from that used by your Quick Suite account.

URI prefixes in the following formats are supported.

URI prefix format	Example	Comments
<code>https://s3.amazonaws.com/<bucket name>/</code>	<code>https://s3.amazonaws.com/amzn-s3-demo-bucket/</code>	
<code>https://s3.amazonaws.com/<bucket name>/<folder name1>/(<folder name2>/etc.)</code>	<code>https://s3.amazonaws.com/amzn-s3-demo-bucket/folder1/</code>	
<code>s3://<bucket name></code>	<code>s3://amzn-s3-demo-bucket/</code>	
<code>s3://<bucket name>/<folder name1>/(<folder name2>/etc.)</code>	<code>s3://amzn-s3-demo-bucket/folder1/</code>	
<code>https://<bucket name>.s3.amazonaws.com</code>	<code>https://<i>amzn-s3-demo-bucket</i>.<i>s3.amazonaws.com</i></code>	

URI prefix format	Example	Comments
<code>https://s3-<region name>.amazonaws.com/<bucket name>/</code>	<code>https://s3-<i>your-region-for-example-us-east-2</i>.amazonaws.com /amzn-s3-demo-bucket /</code>	This URIPrefix type identifies the Amazon Web Services Region for the Amazon S3 bucket.
<code>https://s3-<region name>.amazonaws.com/<bucket name>/<folder name1>/(<folder name2>/etc.)</code>	<code>https://s3-<i>us-east-1</i>.amazonaws.com /amzn-s3-demo-bucket /folder1/</code>	This URIPrefix type identifies the Amazon Web Services Region for the Amazon S3 bucket.
<code>https://<bucket name>.s3-<region name>.amazonaws.com</code>	<code>https://<i>amzn-s3-demo-bucket</i>.s3-<i>us-east-1</i>.amazonaws.com</code>	This URIPrefix type identifies the Amazon Web Services Region for the Amazon S3 bucket.

- **globalUploadSettings** – (Optional) Use this element to specify import settings for the Amazon S3 files, such as field delimiters. If this element is not specified, Quick Sight uses the default values for the fields in this section.

⚠ Important

For log (.clf) and extended log (.elf) files, only the **format** field in this section is applicable, so you can skip the other fields. If you choose to include them, their values are ignored.

- **format** – (Optional) Specify the format of the files to be imported. Valid formats are **CSV**, **TSV**, **CLF**, **ELF**, and **JSON**. The default value is **CSV**.
- **delimiter** – (Optional) Specify the file field delimiter. This must map to the file type specified in the **format** field. Valid formats are commas (,) for .csv files and tabs (\t) for .tsv files. The default value is comma (,).
- **textqualifier** – (Optional) Specify the file text qualifier. Valid formats are single quote ('), double quotes (\"). The leading backslash is a required escape character for a double quote in JSON. The default value is double quotes (\"). If your text doesn't need a text qualifier, don't include this property.
- **containsHeader** – (Optional) Specify whether the file has a header row. Valid formats are **true** or **false**. The default value is **true**.

Manifest file examples for Quick Sight

The following are some examples of completed Quick Sight manifest files.

The following example shows a manifest file that identifies two specific .csv files for import. These files use double quotes for text qualifiers. The **format**, **delimiter**, and **containsHeader** fields are skipped because the default values are acceptable.

```
{
  "fileLocations": [
    {
      "URIs": [
        "https://yourBucket.s3.amazonaws.com/data-file.csv",
        "https://yourBucket.s3.amazonaws.com/data-file-2.csv"
      ]
    }
  ],
  "globalUploadSettings": {
    "textqualifier": "\""
  }
}
```

```
}
}
```

The following example shows a manifest file that identifies one specific .tsv file for import. This file also includes a bucket in another Amazon Region that contains additional .tsv files for import. The `textqualifier` and `containsHeader` fields are skipped because the default values are acceptable.

```
{
  "fileLocations": [
    {
      "URIs": [
        "https://s3.amazonaws.com/amzn-s3-demo-bucket/data.tsv"
      ]
    },
    {
      "URIPrefixes": [
        "https://s3-us-east-1.amazonaws.com/amzn-s3-demo-bucket/"
      ]
    }
  ],
  "globalUploadSettings": {
    "format": "TSV",
    "delimiter": "\t"
  }
}
```

The following example identifies two buckets that contain .clf files for import. One is in the same Amazon Web Services Region as the Quick Suite account, and one in a different Amazon Web Services Region. The `delimiter`, `textqualifier`, and `containsHeader` fields are skipped because they are not applicable to log files.

```
{
  "fileLocations": [
    {
      "URIPrefixes": [
        "https://amzn-s3-demo-bucket1.your-s3-url.com",
        "s3://amzn-s3-demo-bucket2/"
      ]
    }
  ],
  "globalUploadSettings": {
```

```
    "format": "CLF"
  }
}
```

The following example uses the Amazon Redshift format to identify a .csv file for import.

```
{
  "entries": [
    {
      "url": "https://amzn-s3-demo-bucket.your-s3-url.com/myalias-test/file-to-import.csv",
      "mandatory": true
    }
  ]
}
```

The following example uses the Amazon Redshift format to identify two JSON files for import.

```
{
  "fileLocations": [
    {
      "URIs": [
        "https://yourBucket.s3.amazonaws.com/data-file.json",
        "https://yourBucket.s3.amazonaws.com/data-file-2.json"
      ]
    }
  ],
  "globalUploadSettings": {
    "format": "JSON"
  }
}
```

Creating Amazon S3 datasets

To create an Amazon S3 dataset

1. Check [Data source quotas](#) to make sure that your target file set doesn't exceed data source quotas.
2. Create a manifest file to identify the text files that you want to import, using one of the formats specified in [Supported formats for Amazon S3 manifest files](#).
3. Save the manifest file to a local directory, or upload it into Amazon S3.

4. On the Quick Suite start page, choose **Data**.
5. On the **Data** page, choose **Create** then **New dataset**.
6. Choose the Amazon S3 icon and then choose **Next**.
7. For **Data source name**, enter a description of the data source. This name should be something that helps you distinguish this data source from others.
8. For **Upload a manifest file**, do one of the following:
 - To use a local manifest file, choose **Upload**, and then choose **Upload a JSON manifest file**. For **Open**, choose a file, and then choose **Open**.
 - To use a manifest file from Amazon S3, choose **URL**, and enter the URL for the manifest file. To find the URL of a pre-existing manifest file in the Amazon S3 console, navigate to the appropriate file and choose it. A properties panel displays, including the link URL. You can copy the URL and paste it into Quick Sight.
9. Choose **Connect**.
10. To make sure that the connection is complete, choose **Edit/Preview data**. Otherwise, choose **Visualize** to create an analysis using the data as-is.

If you choose **Edit/Preview data**, you can specify a dataset name as part of preparing the data. Otherwise, the dataset name matches the name of the manifest file.

To learn more about data preparation, see [Preparing data in Amazon Quick Sight](#).

Creating datasets based on multiple Amazon S3 files

You can use one of several methods to merge or combine files from Amazon S3 buckets inside Quick Sight:

- **Combine files by using a manifest** – In this case, the files must have the same number of fields (columns). The data types must match between fields in the same position in the file. For example, the first field must have the same data type in each file. The same goes for the second field, and the third field, and so on. Quick Sight takes field names from the first file.

The files must be listed explicitly in the manifest. However, they don't have to be inside the same Amazon S3 bucket.

In addition, the files must follow the rules described in [Supported formats for Amazon S3 manifest files](#).

For more details about combining files using a manifest, see [Creating a dataset using Amazon S3 files](#).

- **Merge files without using a manifest** – To merge multiple files into one without having to list them individually in the manifest, you can use Athena. With this method, you can simply query your text files, like they are in a table in a database. For more information, see the post [Analyzing data in Amazon S3 using Athena](#) in the Big Data blog.
- **Use a script to append files before importing** – You can use a script designed to combine your files before uploading.

Datasets using S3 files in another Amazon account

Use this section to learn how to set up security so you can use Quick Sight to access Amazon S3 files in another Amazon account.

For you to access files in another account, the owner of the other account must first set Amazon S3 to grant you permissions to read the file. Then, in Quick Sight, you must set up access to the buckets that were shared with you. After both of these steps are finished, you can use a manifest to create a dataset.

Note

To access files that are shared with the public, you don't need to set up any special security. However, you still need a manifest file.

Topics

- [Setting up Amazon S3 to allow access from a different Quick Suite account](#)
- [Setting up Quick Sight to access Amazon S3 files in another Amazon account](#)

Setting up Amazon S3 to allow access from a different Quick Suite account

Use this section to learn how to set permissions in Amazon S3 files so they can be accessed by Quick Sight in another Amazon account.

For information on accessing another account's Amazon S3 files from your Quick Sight account, see [Setting up Quick Sight to access Amazon S3 files in another Amazon account](#). For more information

about S3 permissions, see [Managing access permissions to your Amazon S3 resources](#) and [How do I set permissions on an object?](#)

You can use the following procedure to set this access from the S3 console. Or you can grant permissions by using the Amazon CLI or by writing a script. If you have a lot of files to share, you can instead create an S3 bucket policy on the `s3:GetObject` action. To use a bucket policy, add it to the bucket permissions, not to the file permissions. For information on bucket policies, see [Bucket policy examples](#) in the *Amazon S3 Developer Guide*.

To set access from a different Quick Suite account from the S3 console

1. Get the email address of the Amazon account email that you want to share with. Or you can get and use the canonical user ID. For more information on canonical user IDs, see [Amazon account identifiers](#) in the *Amazon General Reference*.
2. Sign in to the Amazon Web Services Management Console and open the Amazon S3 console at <https://console.amazonaws.cn/s3/>.
3. Find the Amazon S3 bucket that you want to share with Quick Sight. Choose **Permissions**.
4. Choose **Add Account**, and then enter an email address, or paste in a canonical user ID, for the Amazon account that you want to share with. This email address should be the primary one for the Amazon account.
5. Choose **Yes** for both **Read bucket permissions** and **List objects**.

Choose **Save** to confirm.

6. Find the file that you want to share, and open the file's permission settings.
7. Enter an email address or the canonical user ID for the Amazon account that you want to share with. This email address should be the primary one for the Amazon account.
8. Enable **Read object** permissions for each file that Quick Sight needs access to.
9. Notify the Quick Suite user that the files are now available for use.

Setting up Quick Sight to access Amazon S3 files in another Amazon account

Use this section to learn how to set up Quick Sight so you can access Amazon S3 files in another Amazon account. For information on allowing someone else to access your Amazon S3 files from their Quick Suite account, see [Setting up Amazon S3 to allow access from a different Quick Suite account](#).

Use the following procedure to access another account's Amazon S3 files from Quick Sight. Before you can use this procedure, the users in the other Amazon account must share the files in their Amazon S3 bucket with you.

To access another account's Amazon S3 files from Quick Sight

1. Verify that the user or users in the other Amazon account gave your account read and write permission to the S3 bucket in question.
2. Choose your profile icon, and then choose **Manage Quick Sight**.
3. Choose **Security & permissions**.
4. Under **Quick Sight access to Amazon services**, choose **Manage**.
5. Choose **Select S3 buckets**.
6. On the **Select Amazon S3 buckets** screen, choose the **S3 buckets you can access across Amazon** tab.

The default tab is named **S3 buckets linked to Quick Sight account**. It shows all the buckets your Quick Suite account has access to.

7. Do one of the following:
 - To add all the buckets that you have permission to use, choose **Choose accessible buckets from other Amazon accounts**.
 - If you have one or more Amazon S3 buckets that you want to add, enter their names. Each must exactly match the unique name of the Amazon S3 bucket.

If you don't have the appropriate permissions, you see the error message "We can't connect to this S3 bucket. Make sure that any S3 buckets you specify are associated with the Amazon account used to create this Quick Suite account." This error message appears if you don't have either account permissions or Quick Sight permissions.

Note

To use Amazon Athena, Quick Sight needs to access the Amazon S3 buckets that Athena uses.

You can add them here one by one, or use the **Choose accessible buckets from other Amazon accounts** option.

8. Choose **Select buckets** to confirm your selection.

9. Create a new dataset based on Amazon S3, and upload your manifest file. For more information Amazon S3 datasets, see [Creating a dataset using Amazon S3 files](#).

Creating a data source using Apache Spark

You can connect directly to Apache Spark using Quick Sight, or you can connect to Spark through Spark SQL. Using the results of queries, or direct links to tables or views, you create data sources in Quick Sight. You can either directly query your data through Spark, or you can import the results of your query into [SPICE](#).

Before you use Quick Sight with Spark products, you must configure Spark for Quick Sight.

Quick Sight requires your Spark server to be secured and authenticated using LDAP, which is available to Spark version 2.0 or later. If Spark is configured to allow unauthenticated access, Quick Sight refuses the connection to the server. To use Quick Sight as a Spark client, you must configure LDAP authentication to work with Spark.

The Spark documentation contains information on how to set this up. To start, you need to configure it to enable front-end LDAP authentication over HTTPS. For general information on Spark, see [the Apache spark website](#). For information specifically on Spark and security, see [Spark security documentation](#).

To make sure that you have configured your server for Quick Sight access, follow the instructions in [Network and database configuration requirements](#).

Using Databricks in Quick Sight

Use this section to learn how to connect from Quick Sight to Databricks.

To connect to Databricks

1. Begin by creating a new dataset. Choose **Data** from the navigation pane at left.
2. Choose **Create** then **New Dataset**.
3. Choose the **Databricks** data source card.
4. For **Data source name**, enter a descriptive name for your Databricks data source connection, for example Databricks CS. Because you can create many datasets from a connection to Databricks, it's best to keep the name simple.
5. For **Connection type**, select the type of network you're using.

- **Public network** – if your data is shared publicly.
- **VPC** – if your data is inside a VPC.

 Note

If you're using VPC, and you don't see it listed, check with your administrator.

6. For **Database server**, enter the **Hostname of workspace** specified in your Databricks connection details.
7. For **HTTP Path**, enter the **Partial URL for the spark instance** specified in your Databricks connection details.
8. For **Port**, enter the **port** specified in your Databricks connection details.
9. For **Username** and **Password**, enter your connection credentials.
10. To verify the connection is working, click **Validate connection**.
11. To finish and create the data source, click **Create data source**.

Adding a new Quick Sight dataset for Databricks

After you have an existing data source connection for Databricks data, you can create Databricks datasets to use for analysis.

To create a dataset using Databricks

1. Choose **Data** at left, then scroll down to find the data source card for your Databricks connection. If you have many data sources, you can use the search bar at the top of the page to find your data source with a partial match on the name.
2. Choose the **Databricks** data source card, and then choose **Create data set**.
3. To specify the table you want to connect to, first select the Catalog and Schema you want to use. Then, for **Tables**, select the table that you want to use. If you prefer to use your own SQL statement, select **Use custom SQL**.
4. Choose **Edit/Preview**.
5. (Optional) To add more data, use the following steps:
 - a. Choose **Add data** at top right.

- b. To connect to different data, choose **Switch data source**, and choose a different dataset.
 - c. Follow the UI prompts to finish adding data.
 - d. After adding new data to the same dataset, choose **Configure this join** (the two red dots). Set up a join for each additional table.
 - e. If you want to add calculated fields, choose **Add calculated field**.
 - f. To add a model from SageMaker AI, choose **Augment with SageMaker**. This option is only available in Quick Suite Enterprise edition.
 - g. Clear the check box for any fields that you want to omit.
 - h. Update any data types that you want to change.
6. When you are done, choose **Save** to save and close the dataset.

Quick Sight Administrator's guide to connecting Databricks

You can use Amazon Quick Sight to connect to Databricks on Amazon. You can connect to Databricks on Amazon whether you signed up for through Amazon Marketplace or through the Databricks website.

Before you can connect to Databricks, you create or identify existing resources that the connection requires. Use this section to help you gather the resources you need to connect from Quick Sight to Databricks.

- To learn how to obtain your Databricks connection details, see [Databricks ODBC and JDBC connections](#).
- To learn how to obtain your Databricks credentials—personal access token or user name and password—for authentication, see [Authentication requirements](#) in the [Databricks documentation](#).

To connect to a Databricks cluster, you need Can Attach To and Can Restart permissions. These permissions are managed in Databricks. For more information, see [Permission Requirements](#) in the [Databricks documentation](#).

- If you are setting up a private connection for Databricks, you can learn more about how to configure a VPC for use with Quick Sight, see [Connecting to a VPC with Amazon Quick Sight](#) in the Quick Sight documentation. If the connection isn't visible, verify with a system administrator that the network has open [inbound endpoints for Amazon Route 53](#). the hostname of a Databricks workspace uses a public IP, there needs to be DNS TCP and DNS UDP inbound and outbound rules to allow traffic on DNS port 53, for the Route 53 security group. An administrator

needs to create a security group with 2 inbound rules: one for DNS(TCP) on port 53 to the VPC CIDR and one for DNS(UDP) for port 53 to the VPC CIDR.

For Databricks-related details if you are using PrivateLink instead of a public connection, see [Enable Amazon PrivateLink](#) in the [Databricks documentation](#).

Creating a dataset using Google BigQuery

Note

When Quick Sight uses and transfers information that is received from Google APIs, it adheres to the [Google API Services User Data Policy](#).

Google BigQuery is a fully managed serverless data warehouse that customers use to manage and analyze their data. Google BigQuery customers use SQL to query their data without any infrastructure management.

Creating a data source connection with Google BigQuery

Prerequisites

Before you start, make sure that you have the following. These are all required to create a data source connection with Google BigQuery:

- **Project ID** – The project ID that is associated with your Google account. To find this, navigate to the Google Cloud console and choose the name of the project that you want to connect to Quick Sight. Copy the project ID that appears in the new window and record it for later use.
- **Dataset Region** – The Google region that the Google BigQuery project exists in. To find the dataset region, navigate to the Google BigQuery console and choose **Explorer**. Locate and expand the project that you want to connect to, then choose the dataset that you want to use. The dataset region appears in the pop-up that opens.
- **Google account login credentials** – The login credentials for your Google account. If you don't have this information, contact your Google account administrator.
- **Google BigQuery Permissions** – To connect your Google account with Quick Sight, make sure that your Google account has the following permissions:
 - BigQuery Job User at the Project level.

- BigQuery Data Viewer at the Dataset or Table level.
- BigQuery Metadata Viewer at the Project level.

For information about how to retrieve the previous prerequisite information, see [Unlock the power of unified business intelligence with Google Cloud BigQuery and Quick Sight](#).

Use the following procedure to connect your Quick Suite account with your Google BigQuery data source.

To create a new connection to a Google BigQuery data source from Quick Sight

1. Open the [Quick Suite console](#).
2. From the left navigation pane, choose **Data**.
3. Choose **Create** then choose **New Dataset**
4. Choose the **Google BigQuery** tile.
5. Add the data source details that you recorded in the prerequisites section earlier:
 - **Data source name** – A name for the data source.
 - **Project ID** – A Google Platform project ID. This field is case sensitive.
 - **Dataset Region** – The Google cloud platform dataset region of the project that you want to connect to.
6. Choose **Sign in**.
7. In the new window that opens, enter the login credentials for the Google account that you want to connect to.
8. Choose **Continue** to grant Quick Sight access to Google BigQuery.
9. After you create the new data source connection, continue to [Step 4](#) in the following procedure.

Adding a new Quick Sight dataset for Google BigQuery

After you create a data source connection with Google BigQuery, you can create Google BigQuery datasets for analysis. Datasets that use Google BigQuery can only be stored in SPICE.

To create a dataset using Google BigQuery

1. Open the [Quick Suite console](#).

2. From the start page, choose **Data**.
3. Choose **Create**, then **New Dataset**
4. Choose the **Google BigQuery** tile, and then choose **Create dataset**.
5. For **Tables**, do one of the following:
 - Choose the table that you want to use.
 - Choose **Use custom SQL** to use your own personal SQL statement. For more information about using custom SQL in Quick Sight, see [Using SQL to customize data](#).
6. Choose **Edit/Preview**.
7. (Optional) In the **Data prep** page that opens, you can add customizations to your data with calculated fields, filters, and joins.
8. When you are finished making changes, choose **Save** to save and close the dataset.

Creating a dataset using a Google Sheets data source

Google Sheets is a web-based spreadsheet application that enables users to create, edit, and collaborate on data in real time. With its comprehensive set of functions and formulas, it serves as a powerful data source for business intelligence and analytics. Users can organize, analyze, and share insights efficiently, while its seamless collaboration features make it an ideal platform for teams working on data-driven projects.

Admin configuration in Amazon Quick Suite

Amazon Quick Suite administrators need to perform a one-time setup to enable Google Sheets as a data source. For detailed instructions and important considerations, see [the blog](#).

Creating a dataset using a Google Sheets data source

Use the following procedure to create a dataset using a Google Sheets data source.

To create a dataset using a Google Sheets data source

1. From the Quick Suite start page, choose **Datasets**.
2. On the **Datasets** page, choose **New Dataset**.
3. Choose **Google Sheets**.
4. Enter a name for the data source, and then choose **Connect**.
5. When redirected to Google's sign-in page, do the following:

- a. Enter your Google account credentials, and then choose **Next**.
 - b. Review the permissions to authorize your Amazon account to connect with Google Sheets, and then choose **Continue**.
6. In the **Choose your table** menu, locate your data. The menu displays all folders, subfolders, sheets, and tabs from your Google account. To display the tabs, select a sheet from the displayed list.
7. Select the tab you want to work with.
8. Choose **Edit/Preview data** to navigate to the Data preparation page. Choose **Add data** to include any additional tabs.
9. Configure the join, and then select **Publish & visualize** to analyze your Google Sheets data with Quick Sight.

Note

- This connector supports only SPICE functionality.
- If your OAuth token expires (visible in the ingestion error report or when creating a new dataset), reauthorize by choosing **Edit** on the data source and updating it.

Creating a dataset using an Apache Impala data source

Apache Impala is a high-performance massively parallel processing (MPP) SQL query engine designed to run natively on Apache Hadoop. Use the procedure below to establish a secure connection between Quick Sight and Apache Impala.

All traffic between Quick Sight and Apache Impala is encrypted using SSL. Quick Sight supports standard username and password authentication for Impala connections.

To establish a connection, you'll need to configure SSL settings in your Impala instance, prepare your authentication credentials, set up the connection in Quick Sight using your Impala server details, and validate the connection to ensure secure data access.

To create a dataset using an Apache Impala data source

1. On the Quick Suite start page, choose **Data**.
2. On the **Data** page, choose **Create**.

3. Choose **Data source**.
4. Choose **Impala**, then choose **Next**.
5. Enter a name for the data source.
6. For public connections:
 - a. Enter connection details for **Database server**, **HTTP Path**, **Port**, **Username**, and **Password**.
 - b. Once the validation is successful, choose **Create data source**.
7. For private connections:
 - Coordinate with your administrator to set up a VPC connection before entering connection details.

You or your administrator can [configure the VPC connection in Quick Suite](#). SSL is enabled by default to ensure secure data transmission. If you encounter connection validation errors, please verify your connection and VPC details.

If issues persist, consult your administrator to confirm that your Certificate Authority is included in Quick Sight's [approved list of certificates](#).
8. In the **Choose your table** menu, you can either:
 - a. Choose a specific schema or table, then choose **Select**.
 - b. Choose **Use custom SQL** to write your own SQL query.
9. After completing your selection, you will be redirected to the data preparation page. Make any adjustments to your data, then choose **Publish & visualize** to analyze your Impala data in Quick Sight.

Note

This connector supports:

- Username and password authentication
- Public and private connections
- Table discovery and custom SQL queries
- Full data refresh during ingestion
- SPICE storage only

Creating a dataset using a Microsoft Excel file

To create a dataset using a Microsoft Excel file data source, upload an .xlsx file from a local or networked drive. The data is imported into [SPICE](#).

For more information about creating new Amazon S3 datasets using Amazon S3 data sources, see [Creating a dataset using an existing Amazon S3 data source](#) or [Creating a dataset using Amazon S3 files](#).

To create a dataset based on an excel file

1. Check [Data source quotas](#) to make sure that your target file doesn't exceed data source quotas.
2. On the Quick Suite start page, choose **Data**.
3. On the **Data** page, choose **Create** then **New dataset**.
4. Choose **Upload a file**.
5. In the **Open** dialog box, choose a file, and then choose **Open**.

A file must be 1 GB or less to be uploaded to Quick Sight.

6. If the Excel file contains multiple sheets, choose the sheet to import. You can change this later by preparing the data.

7.

Note

On the following screens, you have multiple chances to prepare the data. Each of these takes you to the **Prepare Data** screen. This screen is the same one where you can access after the data import is complete. It enables you to change the upload settings even after the upload is complete.

Choose **Select** to confirm your settings. Or you can choose **Edit/Preview data** to prepare the data immediately.

A preview of the data appears on the next screen. You can't make changes directly to the data preview.

8. If the data headings and content don't look correct, choose **Edit settings and prepare data** to correct the file upload settings.

Otherwise, choose **Next**.

9. On the **Data Source Details** screen, you can choose **Edit/Preview data**. You can specify a dataset name in the **Prepare Data** screen.

If you don't need to prepare the data, you can choose to create an analysis using the data as-is. Choose **Visualize**. Doing this names the dataset the same as the source file, and takes you to the **Analysis** screen. To learn more about data preparation and excel upload settings, see [Preparing data in Amazon Quick Sight](#).

Note

If at anytime you want to make changes to the file, such as adding a new field, you must make the change in Microsoft Excel and create a new dataset using the updated version in Quick Sight. For more information about possible implications of changing datasets, see [Things to consider when editing datasets](#).

Creating a data source using Presto

Presto (or PrestoDB) is an open-source, distributed SQL query engine, designed for fast analytic queries against data of any size. It supports both nonrelational and relational data sources. Supported nonrelational data sources include the Hadoop Distributed File System (HDFS), Amazon S3, Cassandra, MongoDB, and HBase. Supported relational data sources include MySQL, PostgreSQL, Amazon Redshift, Microsoft SQL Server, and Teradata.

For more information about Presto, see the following:

- [Introduction to presto](#), a description of Presto on the Amazon website.
- [Creating a presto cluster with Amazon elastic MapReduce \(EMR\)](#) in the *Amazon EMR Release Guide*.
- For general information on Presto, see the [Presto documentation](#).

The results of the queries that you run through the Presto query engine can be turned into Quick Sight datasets. Presto processes the analytic queries on the backend databases. Then it returns results to the Quick Sight client. You can directly query your data through Presto, or you can import the results of your query into SPICE.

Before you use Quick Sight as a Presto client to run queries, make sure that you configure data source profiles. You need a data source profile in Quick Sight for each Presto data source that you want to access. Use the following procedure to create a connection to Presto.

To create a new connection to a presto data source from Amazon Quick Sight (console)

1. On the Amazon Quick Sight start page, choose **Data** at left.
2. Choose **Create** then **New dataset**.
3. Choose the **Presto** tile.

Note

In most browsers, you can use Ctrl-F or Cmd-F to open a search box and enter **presto** to locate it.

4. Add the settings for the new data source:
 - **Data source name** – Enter a descriptive name for your data source connection. This name appears in the **Existing data sources** section at the bottom of the **Data sets** screen.
 - **Connection type** – Choose the connection type that you need to use to connect to Presto.

To connect through the public network, choose **Public network**.

If you use a public network, your Presto server must be secured and authenticated using Lightweight Directory Access Protocol (LDAP). For information on configuring Presto to use LDAP, see [LDAP authentication](#) in the Presto documentation.

To connect through a virtual private connection, choose the appropriate VPC name from the **VPC connections** list.

If your Presto server allows unauthenticated access, Amazon requires that you connect to it securely by using a private VPC connection. For information on configuring a new VPC, see [Configuring VPC connections in Amazon Quick Sight](#).

- **Database server** – The name of the database server.
- **Port** – The port that the server using to accept incoming connections from Amazon Quick Sight
- **Catalog** – The name of the catalog that you want to use.

- **Authentication required** – (Optional) This option only appears if you choose a VPC connection type. If the Presto data source that you're connecting to doesn't require authentication, choose **No**. Otherwise, keep the default setting (**Yes**).
- **Username** – Enter a user name to use to connect to Presto. Quick Sight applies the same user name and password to all connections that use this data source profile. If you want to monitor Quick Sight separately from other accounts, create a Presto account for each Quick Sight data source profile.

The Presto account that you use needs be able to access to the database and run SELECT statements on at least one table.

- **Password** – The password to use with the Presto user name. Amazon Quick Sight encrypts all credentials that you use in data source profile. For more information, see [Data encryption in Amazon Quick Suite](#).
 - **Enable SSL** – SSL is enabled by default.
5. Choose **Validate connection** to test your settings.
 6. After you validate your settings, choose **Create data source** to complete the connection.

Using Snowflake with Amazon Quick Sight

Snowflake is an AI data cloud platform that provides data solutions from data warehousing and collaboration to data science and generative AI. Snowflake is an [Amazon Partner](#) with multiple Amazon accreditations that include Amazon ISV Competencies in Generative AI, Machine Learning, Data and Analytics, and Retail.

Amazon Quick Sight offers two ways to connect to Snowflake: with your Snowflake login credentials or with OAuth client credentials. Use the following sections to learn about both methods of connection.

Topics

- [Creating an Quick Sight data source connection to Snowflake with login credentials](#)
- [Creating an Quick Sight data source connection to Snowflake with OAuth client credentials](#)

Creating an Quick Sight data source connection to Snowflake with login credentials

Use this section to learn how to create a connection between Quick Sight and Snowflake with your Snowflake login credentials. All traffic between Quick Sight and Snowflake is enabled by SSL.

To create a connection between Quick Sight and Snowflake

1. Open the [Quick Suite console](#).
2. From the left navigation pane, choose **Data**, then choose **Create**, then choose **New Dataset**.
3. Choose the **Snowflake** data source card.
4. In the pop up that appears, enter the following information:
 - a. For **Data source name**, enter a descriptive name for your Snowflake data source connection. Because you can create many datasets from a connection to Snowflake, it's best to keep the name simple.
 - b. For **Connection type**, choose the type of network that you're using. Choose **Public network** if your data is shared publicly. Choose **VPC** if your data is located inside a VPC. To configure a VPC connection in Quick Sight, see [Managing VPC connection in Amazon Quick Suite](#).
 - c. For **Database server** enter the hostname specified in your Snowflake connection details.
5. For **Database name and Warehouse**, enter the respective Snowflake database and warehouse that you want to connect.
6. For **Username** and **Password**, enter your Snowflake credentials.

After you have successfully created a data source connection between your Quick Sight and Snowflake accounts, you can begin [Creating datasets](#) that contain Snowflake data.

Creating an Quick Sight data source connection to Snowflake with OAuth client credentials

You can use OAuth client credentials to connect your Quick Sight account with Snowflake through the [Quick Sight APIs](#). *OAuth* is a standard authorization protocol that is often utilized for applications that have advanced security requirements. When you connect to Snowflake with OAuth client credentials, you can create datasets that contain Snowflake data with the Quick Sight APIs and in the Quick Sight UI. For more information about configuring OAuth in Snowflake, see [Snowflake OAuth overview](#).

Quick Sight supports the `client_credentials` OAuth grant type. OAuth client credentials is used to obtain an access token for machine-to-machine communication. This method is suitable for scenarios where a client needs to access resources that are hosted on a server without the involvement of a user.

In the client credentials flow of OAuth 2.0, there are several client authentication mechanisms that can be used to authenticate the client application with the authorization server. Quick Sight supports client credentials based OAuth for Snowflake for the following two mechanisms:

- **Token (Client secrets-based OAuth):** The secret-based client authentication mechanism is used with the client credentials to grant flow in order to authenticate with authorization server. This authentication scheme requires the `client_id` and `client_secret` of the OAuth client app to be stored in Secrets Manager.
- **X509 (Client private key JWT-based OAuth):** The X509 certificate key-based solution provides an additional security layer to the OAuth mechanism with client certificates that are used to authenticate instead of client secrets. This method is primarily used by private clients who use this method to authenticate with the authorization server with strong trust between the two services.

Quick Sight has validated OAuth connections with the following Identity providers:

- OKTA
- PingFederate

Storing OAuth credentials in Secrets Manager

OAuth client credentials are meant for machine-to-machine use cases and are not designed to be interactive. To create a datasource connection between Quick Sight and Snowflake, create a new secret in Secrets Manager that contains your credentials for the OAuth client app. The secret ARN that is created with the new secret can be used to create datasets that contain Snowflake data in Quick Sight. For more information about using Secrets Manager keys in Quick Sight, see [Using Amazon Secrets Manager secrets instead of database credentials in Quick Suite](#).

The credentials that you need to store in Secrets Manager are determined by the OAuth mechanism that you use. The following key/value pairs are required for X509-based OAuth secrets:

- `username`: The Snowflake account username to be used when connecting to Snowflake
- `client_id`: The OAuth client ID
- `client_private_key`: The OAuth client private key
- `client_public_key`: The OAuth client certificate public key and its encrypted algorithm (for example, `{"alg": "RS256", "kid", "cert_kid"}`)

The following key/value pairs are required for token-based OAuth secrets:

- **username:** The Snowflake account username to be used when connecting to Snowflake
- **client_id:** The OAuth client ID
- **client_secret:** the OAuth client secret

Creating a Snowflake OAuth connection with the Quick Sight APIs

After you create a secret in Secrets Manager that contains your Snowflake OAuth credentials and have connected your Quick Suite account to Secrets Manager, you can establish a data source connection between Quick Sight and Snowflake with the Quick Sight APIs and SDK. The following example creates a Snowflake data source connection using token OAuth client credentials.

```
{
  "AwsAccountId": "AWSACCOUNTID",
  "DataSourceId": "UNIQUEDATASOURCEID",
  "Name": "NAME",
  "Type": "SNOWFLAKE",
  "DataSourceParameters": {
    "SnowflakeParameters": {
      "Host": "HOSTNAME",
      "Database": "DATABASENAME",
      "Warehouse": "WAREHOUSENAME",
      "AuthenticationType": "TOKEN",
      "DatabaseAccessControlRole": "snowflake-db-access-role-name",
      "OAuthParameters": {
        "TokenProviderUrl": "oauth-access-token-endpoint",
        "OAuthScope": "oauth-scope",
        "IdentityProviderResourceUri" : "resource-uri",
        "IdentityProviderVpcConnectionProperties" : {
          "VpcConnectionArn": "IdP-VPC-connection-ARN"
        }
      }
    }
  },
  "VpcConnectionProperties": {
    "VpcConnectionArn": "VPC-connection-ARN-for-Snowflake"
  },
  "Credentials": {
    "SecretArn": "oauth-client-secret-ARN"
  }
}
```

```
}
```

For more information about the CreateDatasource API operation, see [CreateDataSource](#).

Once the connection between Quick Sight and Snowflake is established and a data source is created with the Quick Sight APIs or SDK, the new data source is displayed in Quick Sight. Quick Sight authors can use this data source to create datasets that contain Snowflake data. Tables are displayed based on the role used in the DatabaseAccessControlRole parameter that is passed in a CreateDataSource API call. If this parameter is not defined when the data source connection is created, the default Snowflake role is used.

After you have successfully created a data source connection between your Quick Sight and Snowflake accounts, you can begin [Creating datasets](#) that contain Snowflake data.

Using Starburst with Amazon Quick Sight

Starburst is a full-featured data lake analytics service built on top of a massively parallel processing (MPP) query engine, Trino. Use this section to learn how to connect from Amazon Quick Sight to Starburst. All traffic between Quick Sight and Starburst is enabled by SSL. If you're connecting to Starburst Galaxy, you can get the necessary connection details by logging in to your Starburst Galaxy account, then choose **Partner Connect** and then **Quick Sight**. You should be able to see information, such as hostname and port. Amazon Quick Sight supports basic username and password authentication to Starburst.

Quick Sight offers two ways to connect to Starburst: with your Starburst login credentials or with OAuth client credentials. Use the following sections to learn about both methods of connection.

Topics

- [Creating an Quick Sight data source connection to Starburst with login credentials](#)
- [Creating an Quick Sight data source connection to Starburst with OAuth client credentials](#)

Creating an Quick Sight data source connection to Starburst with login credentials

1. Begin by creating a new dataset. From the left navigation pane, choose **Data**, then choose **Create**, then choose **New Dataset**.
2. Choose the **Starburst** data source card.
3. Select the Starburst product type. Choose **Starburst Enterprise** for on-prem Starburst instances. Choose **Starburst Galaxy** for managed instances.

4. For **Data source name**, enter a descriptive name for your Starburst data source connection. Because you can create many datasets from a connection to Starburst, it's best to keep the name simple.
5. For **Connection type**, select the type of network you're using. Choose **Public network** if your data is shared publicly. Choose **VPC** if your data is inside a VPC. To configure a VPC connection in Amazon Quick Sight, see [Configuring the VPC connection in Amazon Quick Sight](#). This connection type is not available for Starburst Galaxy.
6. For **Database server** enter the hostname specified in your Starburst connection details.
7. For **Catalog**, enter the catalog specified in your Starburst connection details.
8. For **Port**, enter the port specified in your Starburst connection details. Defaults to 443 for Starburst Galaxy.
9. For **Username** and **Password**, enter your Starburst connection credentials.
10. To verify the connection is working, choose **Validate connection**.
11. To finish and create the data source, choose **Create data source**.

 Note

Connectivity between Amazon Quick Sight and Starburst was validated using Starburst version 420.

After you have successfully created a data source connection between your Quick Sight and Starburst accounts, you can begin [Creating datasets](#) that contain Starburst data.

Creating an Quick Sight data source connection to Starburst with OAuth client credentials

You can use OAuth client credentials to connect your Quick Sight account with Starburst through the [Quick Sight APIs](#). *OAuth* is a standard authorization protocol that is often utilized for applications that have advanced security requirements. When you connect to Starburst with OAuth client credentials, you can create datasets that contain Starburst data with the Quick Sight APIs and in the Quick Sight UI. For more information about configuring OAuth in Starburst, see [OAuth 2.0 authentication](#).

Quick Sight supports the `client_credentials` OAuth grant type. OAuth client credentials is used to obtain an access token for machine-to-machine communication. This method is suitable

for scenarios where a client needs to access resources that are hosted on a server without the involvement of a user.

In the client credentials flow of OAuth 2.0, there are several client authentication mechanisms that can be used to authenticate the client application with the authorization server. Quick Sight supports client credentials based OAuth for Starburst for the following two mechanisms:

- **Token (Client secrets-based OAuth):** The secret-based client authentication mechanism is used with the client credentials to grant flow in order to authenticate with authorization server. This authentication scheme requires the `client_id` and `client_secret` of the OAuth client app to be stored in Secrets Manager.
- **X509 (Client private key JWT-based OAuth):** The X509 certificate key-based solution provides an additional security layer to the OAuth mechanism with client certificates that are used to authenticate instead of client secrets. This method is primarily used by private clients who use this method to authenticate with the authorization server with strong trust between the two services.

Quick Sight has validated OAuth connections with the following Identity providers:

- OKTA
- PingFederate

Storing OAuth credentials in Secrets Manager

OAuth client credentials are meant for machine-to-machine use cases and are not designed to be interactive. To create a datasource connection between Quick Sight and Starburst, create a new secret in Secrets Manager that contains your credentials for the OAuth client app. The secret ARN that is created with the new secret can be used to create datasets that contain Starburst data in Quick Sight. For more information about using Secrets Manager keys in Quick Sight, see [Using Amazon Secrets Manager secrets instead of database credentials in Quick Suite](#).

The credentials that you need to store in Secrets Manager are determined by the OAuth mechanism that you use. The following key/value pairs are required for X509-based OAuth secrets:

- `username`: The Starburst account username to be used when connecting to Starburst
- `client_id`: The OAuth client ID
- `client_private_key`: The OAuth client private key

- `client_public_key`: The OAuth client certificate public key and its encrypted algorithm (for example, `{"alg": "RS256", "kid", "cert_kid"}`)

The following key/value pairs are required for token-based OAuth secrets:

- `username`: The Starburst account username to be used when connecting to Starburst
- `client_id`: The OAuth client ID
- `client_secret`: the OAuth client secret

Creating a Starburst OAuth connection with the Quick Sight APIs

After you create a secret in Secrets Manager that contains your Starburst OAuth credentials and have connected your Quick Suite account to Secrets Manager, you can establish a data source connection between Quick Sight and Starburst with the Quick Sight APIs and SDK. The following example creates a Starburst data source connection using token OAuth client credentials.

```
{
  "AwsAccountId": "AWSACCOUNTID",
  "DataSourceId": "DATASOURCEID",
  "Name": "NAME",
  "Type": "STARBURST",
  "DataSourceParameters": {
    "StarburstParameters": {
      "Host": "STARBURST_HOST_NAME",
      "Port": "STARBURST_PORT",
      "Catalog": "STARBURST_CATALOG",
      "ProductType": "STARBURST_PRODUCT_TYPE",
      "AuthenticationType": "TOKEN",
      "DatabaseAccessControlRole": "starburst-db-access-role-name",
      "OAuthParameters": {
        "TokenProviderUrl": "oauth-access-token-endpoint",
        "OAuthScope": "oauth-scope",
        "IdentityProviderResourceUri" : "resource-uri",
        "IdentityProviderVpcConnectionProperties" : {
          "VpcConnectionArn": "IdP-VPC-connection-ARN"
        }
      }
    }
  },
  "VpcConnectionProperties": {
    "VpcConnectionArn": "VPC-connection-ARN-for-Starburst"
  }
}
```

```
    },  
    "Credentials": {  
        "SecretArn": "oauth-client-secret-ARN"  
    }  
}
```

For more information about the CreateDataSource API operation, see [CreateDataSource](#).

Once the connection between Quick Sight and Starburst is established and a data source is created with the Quick Sight APIs or SDK, the new data source is displayed in Quick Sight. Quick Sight authors can use this data source to create datasets that contain Starburst data. Tables are displayed based on the role used in the DatabaseAccessControlRole parameter that is passed in a CreateDataSource API call. If this parameter is not defined when the data source connection is created, the default Starburst role is used.

After you have successfully created a data source connection between your Quick Sight and Starburst accounts, you can begin [Creating datasets](#) that contain Starburst data.

Creating a data source and data set from SaaS sources

To analyze and report on data from software as a service (SaaS) applications, you can use SaaS connectors to access your data directly from Quick Sight. The SaaS connectors simplify accessing third-party application sources using OAuth, without any need to export the data to an intermediate data store.

You can use either a cloud-based or server-based instance of a SaaS application. To connect to an SaaS application that is running on your corporate network, make sure that Quick Sight can access the application's Domain Name System (DNS) name over the network. If Quick Sight can't access the SaaS application, it generates an unknown host error.

Here are examples of some ways that you can use SaaS data:

- Engineering teams who use Jira to track issues and bugs can report on developer efficiency and bug burndown.
- Marketing organizations can integrate Quick Sight with Adobe Analytics to build consolidated dashboards to visualize their online and web marketing data.

Use the following procedure to create a data source and dataset by connecting to sources available through Software as a Service (SaaS). In this procedure, we use a connection to GitHub as an

example. Other SaaS data sources follow the same process, although the screens—especially the SaaS screens—might look different.

To create a data source and dataset by connecting to sources through SaaS

1. On the Quick Suite start page, choose **Data**.
2. On the **Data** page, choose **Create** then choose **New dataset**.
3. Choose the icon that represents the SaaS source that you want to use. For example, you might choose Adobe Analytics or GitHub.

For sources using OAuth, the connector takes you to the SaaS site to authorize the connection before you can create the data source.

4. Choose a name for the data source, and enter that. If there are more screen prompts, enter the appropriate information. Then choose **Create data source**.
5. If you are prompted to do so, enter your credentials on the SaaS login page.
6. When prompted, authorize the connection between your SaaS data source and Quick Sight.

The following example shows the authorization for Quick Sight to access the GitHub account for the Quick Sight documentation.

Note

Quick Sight documentation is now available on GitHub. If you want to make changes to this user guide, you can use GitHub to edit it directly.

(Optional) If your SaaS account is part of an organizational account, you might be asked to request organization access as part of authorizing Quick Sight. If you want to do this, follow the prompts on your SaaS screen, then choose to authorize Quick Sight.

7. After authorization is complete, choose a table or object to connect to. Then choose **Select**.
8. On the **Finish data set creation** screen, choose one of these options:
 - To save the data source and dataset, choose **Edit/Preview data**. Then choose **Save** from the top menu bar.
 - To create a dataset and an analysis using the data as-is, choose **Visualize**. This option automatically saves the data source and the dataset.

You can also choose **Edit/Preview data** to prepare the data before creating an analysis. This opens the data preparation screen. For more information about data preparation, see [Preparing dataset examples](#).

The following constraints apply:

- The SaaS source must support REST API operations for Quick Sight to connect to it.
- If you are connecting to Jira, the URL must be public address.
- If you don't have enough [SPICE](#) capacity, choose **Edit/Preview data**. In the data preparation screen, you can remove fields from the dataset to decrease its size or apply a filter that reduces the number of rows returned. For more information about data preparation, see [Preparing dataset examples](#).

Creating a dataset from Salesforce

Use the following procedure to create a dataset by connecting to Salesforce and selecting a report or object to provide data.

To create a dataset using Salesforce from a report or object

1. Check [Data source quotas](#) to make sure that your target report or object doesn't exceed data source quotas.
2. On the Quick Suite start page, choose **Data**.
3. On the **Data** page, choose **Create** then **New dataset**.
4. Choose the **Salesforce** icon.
5. Enter a name for the data source and then choose **Create data source**.
6. On the Salesforce login page, enter your Salesforce credentials.
7. For **Data elements: contain your data**, choose **Select** and then choose either **REPORT** or **OBJECT**.

Note

Joined reports aren't supported as Quick Sight data sources.

8. Choose one of the following options:

- To prepare the data before creating an analysis, choose **Edit/Preview data** to open data preparation. For more information about data preparation, see [Preparing dataset examples](#).
 - Otherwise, choose a report or object and then choose **Select**.
9. Choose one of the following options:
- To create a dataset and an analysis using the data as-is, choose **Visualize**.

 Note

If you don't have enough [SPICE](#) capacity, choose **Edit/Preview data**. In data preparation, you can remove fields from the dataset to decrease its size or apply a filter that reduces the number of rows returned. For more information about data preparation, see [Preparing dataset examples](#).

- To prepare the data before creating an analysis, choose **Edit/Preview data** to open data preparation for the selected report or object. For more information about data preparation, see [Preparing dataset examples](#).

Using Trino with Amazon Quick Sight

Trino is a massively parallel processing (MPP) query engine built to quickly query data lakes with petabytes of data. Use this section to learn how to connect from Amazon Quick Sight to Trino. All traffic between Amazon Quick Sight and Trino is enabled by SSL. Amazon Quick Sight supports basic username and password authentication to Trino.

Creating a data source connection for Trino

1. Begin by creating a new dataset. From the left navigation pane, choose **Data**. Choose **Create** then **New Dataset**.
2. Choose the **Trino** data source card.
3. For **Data source name**, enter a descriptive name for your Trino data source connection. Because you can create many datasets from a connection to Trino, it's best to keep the name simple.
4. For **Connection type**, select the type of network you're using. Choose **Public network** if your data is shared publicly. Choose **VPC** if your data is inside a VPC. To configure a VPC connection in Amazon Quick Sight, see [Configuring the VPC connection in Amazon Quick Sight](#).

5. For **Database server**, enter the hostname specified in your Trino connection details.
6. For **Catalog**, enter the catalog specified in your Trino connection details.
7. For **Port**, enter the port specified in your Trino connection details.
8. For **Username** and **Password**, enter your Trino connection credentials.
9. To verify the connection is working, choose **Validate connection**.
10. To finish and create the data source, choose **Create data source**.

Adding a new Amazon Quick Sight dataset for Trino

After you go through the [data source creation process](#) for Trino, you can create Trino datasets to use for analysis. You can create new datasets from a new or an existing Trino data source. When you are creating a new data source, Amazon Quick Sight immediately takes you to creating a dataset, which is step 3 below. If you're using an existing data source to create a new dataset, start from step 1 below.

To create a dataset using a Trino data source, see the following steps.

1. From the start page, choose **Data**. Choose **Create** then **New dataset**.
2. Choose the Trino data source you created.
3. Choose **Create data set**.
4. To specify the table you want to connect to, choose a schema. If you don't want to choose a schema, you can also use your own SQL statement.
5. To specify the table you want to connect to, first select the **Schema** you want to use. For **Tables**, choose the table that you want to use. If you prefer to use your own SQL statement, select **Use custom SQL**.
6. Choose **Edit/Preview**.
7. (Optional) To add more data, use the following steps:
8. Choose **Add data** in the top right.
9. To connect to different data, choose **Switch data source**, and choose a different dataset.
10. Follow the prompts to finish adding data.
11. After adding new data to the same dataset, choose **Configure this join** (the two red dots). Set up a join for each additional table.
12. If you want to add calculated fields, choose **Add calculated field**.

13. Clear the check box for any fields that you want to omit.
14. Update any data types that you want to change.
15. When you are done, choose **Save** to save and close the dataset.

 Note

Connectivity between Quick Sight and Trino was validated using Trino version 410.

Creating a dataset using a local text file

To create a dataset using a local text file data source, identify the location of the file, and then upload it. The file data is automatically imported into [SPICE](#) as part of creating a dataset.

To create a dataset based on a local text file

1. Check [Data source quotas](#) to make sure that your target file doesn't exceed data source quotas.

Supported file types include .csv, .tsv, .json, .clf, or .elf files.

2. On the Quick Suite start page, choose **Data**.
3. Choose **Create** then **New dataset**.
4. Choose **Upload a file**.
5. In the **Open** dialog box, browse to a file, select it, and then choose **Open**.

A file must be 1 GB or less to be uploaded to Quick Sight.

6. To prepare the data before creating the dataset, choose **Edit/Preview data**. Otherwise, choose **Visualize** to create an analysis using the data as-is.

If you choose the former, you can specify a dataset name as part of preparing the data. If you choose the latter, a dataset with the same name as the source file is created. To learn more about data preparation, see [Preparing data in Amazon Quick Sight](#).

Using Amazon Timestream data with Amazon Quick Sight

Following, you can find how to connect to your Amazon Timestream data using Amazon Quick Sight. For a brief overview, see the [Getting started with Amazon Timestream and Amazon QuickSight](#) video tutorial on YouTube.

Creating a new Amazon Quick Sight data source connection for a Timestream database

Following, you can find how to connect to Amazon Timestream from Amazon Quick Sight.

Before you can proceed, Amazon Quick Sight needs to be authorized to connect to Amazon Timestream. If connections aren't enabled, you get an error when you try to connect. A Quick Sight administrator can authorize connections to Amazon resources. To authorize, open the menu by clicking on your profile icon at top right. Choose **Manage QuickSight, Security & permissions, Add or remove**. Then enable the check box for Amazon Timestream, then choose **Update** to confirm. For more information, see [Configuring Amazon Quick Sight access to Amazon data sources](#).

To connect to Amazon Timestream

1. Begin by creating a new dataset. Choose **Data** from the navigation pane at left.
2. Choose **Create** then **New Dataset**.
3. Choose the Timestream data source card.
4. For **Data source name**, enter a descriptive name for your Timestream data source connection, for example `US Timestream Data`. Because you can create many datasets from a connection to Timestream, it's best to keep the name simple.
5. Choose **Validate connection** to check that you can successfully connect to Timestream.
6. Choose **Create data source** to proceed.
7. For **Database**, choose **Select** to view the list of available options.
8. Choose the one you want to use, then choose **Select** to continue.
9. Do one of the following:
 - To import your data into Quick Sight's in-memory engine (called SPICE), choose **Import to SPICE for quicker analytics**.
 - To allow Quick Sight to run a query against your data each time you refresh the dataset or use the analysis or dashboard, choose **Directly query your data**.

If you want to enable autorefresh on a published dashboard that uses Timestream data, the Timestream dataset needs to use a direct query.

10. Choose **Edit/Preview** and then **Save** to save your dataset and close it.
11. Repeat these steps for the number of concurrent direct connections to Timestream that you want to open in a dataset. For example, let's say you want to use four tables in a Quick Sight dataset. Currently, Quick Sight datasets connect to only one table at a time from a

Timestream data source. To use four tables in the same dataset, you need to add four data source connections in Quick Sight.

Managing permissions for Timestream data

The following procedure describes how to view, add, and revoke permissions to allow access to the same Timestream data source. The people that you add need to be active users in Quick Sight before you can add them.

To edit permissions on a dataset

1. Choose **Data** at left, then scroll down to find the dataset for your Timestream connection. An example might be US Timestream Data.
2. Choose the **Timestream** dataset to open it.
3. On the dataset details page that opens, choose the **Permissionstab**.

A list of current permissions appears.

4. To add permissions, choose **Add users & groups**, then follow these steps:
 - a. Add users or groups to allow them to use the same dataset.
 - b. When you're finished adding everyone that you want to add, choose the **Permissions** that you want to apply to them.
5. (Optional) To edit permissions, you can choose **Viewer** or **Owner**.
 - Choose **Viewer** to allow read access.
 - Choose **Owner** to allow that user to edit, share, or delete this Quick Sight data source.
6. (Optional) To revoke permissions, choose **Revoke access**. After you revoke someone's access, they can't create edit, share, or delete the dataset.
7. When you are finished, choose **Close**.

Adding a new Quick Sight dataset for Timestream

After you have an existing data source connection for Timestream data, you can create Timestream datasets to use for analysis.

Currently, you can use a Timestream connection only for a single table in a dataset. To add data from multiple Timestream tables in a single dataset, create an additional Quick Sight data source connection for each table.

To create a dataset using Amazon Timestream

1. Choose **Data** at left, then scroll down to find the data source card for your Timestream connection. If you have many data sources, you can use the search bar at the top of the page to find your data source with a partial match on the name.
2. Choose the **Timestream** data source card, and then choose **Create data set**.
3. For **Database**, choose **Select** to view a list of available databases and choose the one that you want to use.
4. For **Tables**, choose the table that you want to use.
5. Choose **Edit/Preview**.
6. (Optional) To add more data, use the following steps:
 - a. Choose **Add data** at top right.
 - b. To connect to different data, choose **Switch data source**, and choose a different dataset.
 - c. Follow the UI prompts to finish adding data.
 - d. After adding new data to the same dataset, choose **Configure this join** (the two red dots). Set up a join for each additional table.
 - e. If you want to add calculated fields, choose **Add calculated field**.
 - f. To add a model from SageMaker AI, choose **Augment with SageMaker**. This option is only available in Amazon Quick Suite Enterprise edition.
 - g. Clear the check box for any fields that you want to omit.
 - h. Update any data types that you want to change.
7. When you are done, choose **Save** to save and close the dataset.

Adding Timestream data to an analysis

Following, you can find how to add an Amazon Timestream dataset to a Quick Sight analysis. Before you begin, make sure that you have an existing dataset that contains the Timestream data that you want to use.

To add Amazon Timestream data to an analysis

1. Choose **Analyses** at left.
2. Do one of the following:
 - To create a new analysis, choose **New analysis** at right.
 - To add to an existing analysis, open the analysis that you want to edit.
 - Choose the pencil icon near at top left.
 - Choose **Add data set**.
3. Choose the Timestream dataset that you want to add.

For more information, see [Working with analyses](#).

Data source quotas

Data sources that you use with Amazon Quick Sight must conform to the following quotas.

Topics

- [SPICE quotas for imported data](#)
- [Quotas for direct SQL queries](#)

SPICE quotas for imported data

When you create a new dataset in Amazon Quick Sight, [SPICE](#) limits the number of rows you can add to a dataset. You can ingest data into SPICE from a query or from a file. Each file can have up to 2,000 columns. Each column name can have up to 127 Unicode characters. Each field can have up to 2,047 Unicode characters.

To retrieve a subset of data from a larger set, you can deselect columns or apply filters to reduce the size of the data. If you are importing from Amazon S3, each manifest can specify up to 1,000 files.

Quotas for SPICE are as follows:

- 2,047 Unicode characters for each field
- 127 Unicode characters for each column name
- 2,000 columns for each file

- 1,000 files for each manifest
- For Standard edition, 25 million (25,000,000) rows or 25 GB for each dataset
- For Enterprise edition, 1 billion (1,000,000,000) rows or 1 TB for each dataset

All quotas apply to SPICE datasets with row-level security, as well.

In rare cases, if you're ingesting large rows into SPICE, you might reach the quota for gigabytes per dataset before you reach the quota on rows. The size is based on the SPICE capacity the data occupies after ingestion into SPICE.

Quotas for direct SQL queries

If you aren't importing data into SPICE, different quotas apply for space and time. For operations such as connecting, sampling data for a dataset, and generating visuals, timeouts can occur. In some cases, these are timeout quotas set by the source database engine. In other cases, such as visualizing, Amazon Quick Sight generates a timeout after 2 minutes.

However, not all database drivers react to the 2-minute timeout, for example Amazon Redshift. In these cases, the query runs for as long as it takes for the response to return, which can result in long-running queries on your database. When this happens, you can cancel the query from the database server to free up database resources. Follow the instructions for your database server about how to do this. For example, for more information on how to cancel queries in Amazon Redshift, see [Canceling a query in Amazon Redshift](#), and [Implementing workload management in Amazon Redshift](#) in the *Amazon Redshift Database Developer Guide*.

Each result set from a direct query can have up to 2,000 columns. Each column name can have up to 127 Unicode characters. If you want to retrieve data from a larger table, you can use one of several methods to reduce the size of the data. You can deselect columns, or apply filters. In a SQL query, you can also use predicates, such as WHERE, HAVING. If your visuals time out during a direct query, you can simplify your query to optimize execution time or you can import the data into SPICE.

Quotas for queries are as follows:

- 127 Unicode characters for each column name.
- 2,000 columns for each dataset.
- 2-minute quota for generating a visual, or an optional dataset sample.
- Data source timeout quotas apply (varies for each database engine).

Supported data types and values

Amazon Quick Sight currently supports the following primitive data types: Date, Decimal, Integer, and String. The following data types are supported in SPICE: Date, Decimal-fixed, Decimal-float, Integer, and String. Quick Sight accepts Boolean values by promoting them to integers. It can also derive geospatial data types. Geospatial data types use metadata to interpret the physical data type. Latitude and longitude are numeric. All other geospatial categories are strings.

Make sure that any table or file that you use as a data source contains only fields that can be implicitly converted to these data types. Amazon Quick Sight skips any fields or columns that can't be converted. If you get an error that says "fields were skipped because they use unsupported data types", alter your query or table to remove or recast unsupported data types.

String and text data

Fields or columns that contain characters are called *strings*. A field with the data type of STRING can initially contain almost any type of data. Examples include names, descriptions, phone numbers, account numbers, JSON data, cities, post codes, dates, and numbers that can be used to calculate. These types are sometimes called textual data in a general sense, but not in a technical sense. Quick Sight doesn't support binary and character large objects (BLOBs) in dataset columns. In the Quick Sight documentation, the term "text" always means "string data".

The first time you query or import the data, Quick Sight tries to interpret the data that it identifies as other types, for example dates and numbers. It's a good idea to verify that the data types assigned to your fields or columns are correct.

For each string field in imported data, Quick Sight uses a field length of 8 bytes plus the UTF-8 encoded character length. Amazon Quick Sight supports UTF-8 file encoding, but not UTF-8 (with BOM).

Date and time data

Fields with a data type of Date also include time data, and are also known as Datetime fields. Quick Sight supports dates and times that use [supported date formats](#).

Quick Sight uses UTC time for querying, filtering, and displaying date data. When date data doesn't specify a time zone, Quick Sight assumes UTC values. When date data does specify a time zone, Quick Sight converts it to display in UTC time. For example, a date field with a time zone offset

like **2015-11-01T03:00:00-08:00** is converted to UTC and displayed in Amazon Quick Sight as **2015-11-01T15:30:00**.

For each DATE field in imported data, Quick Sight uses a field length of 8 bytes. Quick Sight supports UTF-8 file encoding, but not UTF-8 (with BOM).

Numeric data

Numeric data includes integers and decimals. Integers with a data type of INT are negative or positive numbers that don't have a decimal place. Quick Sight doesn't distinguish between large and small integers. Integers over a value of 9007199254740991 or $2^{53} - 1$ might not display exactly or correctly in a visual.

Decimals with the data type of Decimal are negative or positive numbers that contain at least one decimal place before or after the decimal point. When you choose Direct Query mode, all non-integer decimal types are marked as Decimal and the underlying engine handles the precision of the datapoint based on the data source's supported behaviors. For more information on supported data source types, see [Supported data types and values](#).

When you store your dataset in SPICE, you can choose to store your decimal values as fixed or float decimal types. Decimal-fixed data types use the format of decimal (18, 4) that allow 18 digits total and up to 4 digits after the decimal point. Decimal-fixed data types are a good choice to conduct exact mathematical operations, but Quick Sight rounds the value to the nearest ten thousandth place when the value is ingested into SPICE.

Decimal-float data types provide approximately 16 significant digits of accuracy to a value. The significant digits can be on either side of the decimal point to support numbers with many decimal places and higher numbers at the same time. For example, the Decimal-float data type supports the number 12345.1234567890 or the number 1234567890.12345. If you work with very small numbers that are close to 0, the Decimal-float data type supports up to 15 digits to the right of the decimal point, for example 0.123451234512345. The maximum value that this data type supports is $1.8 * 10^{308}$ to minimize the probability of an overflow error with your data set.

The Decimal-float data type is inexact and some values are stored as approximations instead of the real value. This may result in slight discrepancies when you store and return some specific values. The following considerations apply to the Decimal-float data type.

- If the dataset that you're using comes from an Amazon S3 data source, SPICE assigns the Decimal-float decimal type to all numeric decimal values.

- If the dataset that you're using comes from a database, SPICE uses the decimal type that the value is assigned in the database. For example, if the value is assigned a fixed-point numeric value in the database, the value will be a Decimal-fixed type in SPICE.

For existing SPICE datasets that contain fields that can be converted to the Decimal-float data type, a pop-up appears in the **Edit dataset** page. To convert fields of an existing dataset to the Decimal-float data type, choose **UPDATE FIELDS**. If you don't want to opt in, choose **DO NOT UPDATE FIELDS**. The **Update fields** pop up appears every time you open the **Edit dataset** page until the dataset is saved and published.

Supported data types from external data sources

The following table lists data types that are supported when using the following data sources with Amazon Quick Sight.

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
Amazon Athena, Presto, Starburst, Trino	• bigint • decimal • double • integer • real • smallint • tinyint	• char • varchar	• date • timestamp	• boolean
Amazon Aurora, MariaDB, and MySQL	• bigint •	• char •	• date •	

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
	decimal - double - int - integer - mediumint - numeric - smallint - tinyint	enum - set - text - varchar	datetime - timestamp - year	
Amazon OpenSearch Service	- byte - integer - long - float - double	- string (keyword string field type in OpenSearch Service) - ip	timestamp	- boolean - binary

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
Oracle	• bigint • decimal • decimal • int • money • numeric • real • smallint • smallmoney • tinyint	• char • nchar • nvarchar • text • varchar	• date • datetime • datetime2 • datetimeoffset • smalldatetime	bit

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
PostgreSQL	• bigint • decimal • double • integer • numeric • precision • real • smallint	• char • character • text • varchar • varying character	• date • timestamp	• boolean
Apache Spark	• bigint • decimal • double • integer • real • smallint • tinyint	• varchar	• date • timestamp	• boolean

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
Snowflake	• bigint • byteint • decimal • double • doubleprecision • float • float4 • float8 • int • integer • number • numeric • real • smallint • tinyint	• char • character • string • text • varchar	• date • datetime • time • timestamp • timestamp_*	• boolean

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
Microsoft SQL Server	• bigint • bit • decimal • int • money • numeric • real • smallint • smallmoney • tinyint	• char • nchar • nvarchar • text • varchar	• date • datetime • datetime2 • smalldatetime	• bit

Supported date formats

Amazon Quick Sight supports the date and time formats described in this section. Before you add data to Amazon Quick Sight, check if your date format is compatible. If you need to use an unsupported format, see [Using unsupported or custom dates](#).

The supported formats vary depending on the data source type, as follows:

Data source	Clocks	Date formats
File uploads	Both 24-hour and	Supported date and time formats are described in the Joda API documentation.

Data source	Clocks	Date formats
Amazon S3 sources	12-hour clocks	For a complete list of Joda date formats, see Class DateTimeFormat on the Joda website.
Athena		
Salesforce		For datasets stored in memory (SPICE), Amazon Quick Sight supports dates in the following range: Jan 1, 1400 00:00:00 UTC through Dec 31, 9999, 23:59:59 UTC .

Data source	Clocks	Date formats
Relational databases sources	24-hour clock only	The following data and time formats: 1. <code>dd/MM/yyyy HH:mm:ss</code> , for example <code>31/12/2016 15:30:00</code>. 2. <code>dd/MM/yyyy</code> , for example <code>31/12/2016</code>. 3. <code>dd/MMM/yyyy HH:mm:ss</code> , for example <code>31/DEC/2016 15:30:00</code>. 4. <code>dd/MMM/yyyy</code> , for example <code>31/DEC/2016</code>. 5. <code>dd-MMM-yyyy HH:mm:ss</code> , for example <code>31-DEC-2016 15:30:00</code>. 6. <code>dd-MMM-yyyy</code> , for example <code>31-DEC-2016</code>. 7. <code>dd-MM-yyyy HH:mm:ss</code> , for example <code>31-12-2016 15:30:00</code>. 8. <code>dd-MM-yyyy</code> , for example <code>31-12-2016</code>. 9. <code>MM/dd/yyyy HH:mm:ss</code> , for example <code>12/31/2016 15:30:00</code>. 10. <code>MM/dd/yyyy</code> , for example <code>12/31/2016</code>. 11. <code>MM-dd-yyyy HH:mm:ss</code> , for example <code>12-31-2016 15:30:00</code>. 12. <code>MM-dd-yyyy</code> , for example <code>12-31-2016</code>. 13. <code>MMM/dd/yyyy HH:mm:ss</code> , for example <code>DEC/31/2016 15:30:00</code>. 14.

Data source	Clocks	Date formats
		MMM/dd/yyyy , for example DEC/31/2016. 15. MMM-dd-yyyy HH:mm:ss , for example DEC-31-2016 15:30:00. 16. MMM-dd-yyyy , for example DEC-31-2016. 17. yyyy/MM/dd HH:mm:ss , for example 2016/12/31 15:30:00. 18. yyyy/MM/dd , for example 2016/12/31. 19. yyyy/MMM/dd HH:mm:ss , for example 2016/DEC/31 15:30:00. 20. yyyy/MMM/dd , for example 2016/DEC/31. 21. yyyy-MM-dd HH:mm:ss , for example 2016-12-31 15:30:00. 22. yyyy-MM-dd , for example 2016-12-31. 23. yyyy-MMM-dd HH:mm:ss , for example 2016-DEC-31 15:30:00. 24. yyyy-MMM-dd , for example 2016-DEC-31. 25. yyyyMMdd 'T' HHmmss , for example 20161231T153000. 26. yyyy-MM-dd 'T' HH:mm:ss , for example 2016-12-31T15:30:00. 27. yyyyMMdd 'T' HHmmss .SSS , for example 20161231T153000.123.

Data source	Clocks	Date formats
		28. MM/dd/yyyy HH:mm:ss.SSS , for example 12/31/2016 15:30:00.123. 29. dd/MM/yyyy HH:mm:ss.SSS , for example 31/12/2016 15:30:00.123. 30. yyyy/MM/dd HH:mm:ss.SSS , for example 2016/12/31 15:30:00.123. 31. MMM/dd/yyyy HH:mm:ss.SSS , for example DEC/31/2016 15:30:00.123. 32. dd/MMM/yyyy HH:mm:ss.SSS , for example 31/DEC/2016 15:30:00.123. 33. yyyy/MMM/dd HH:mm:ss.SSS , for example 2016/DEC/31 15:30:00.123. 34. yyyy-MM-dd'T'HH:mm:ss.SSS , for example 2016-12-31T15:30:00.123. 35. MM-dd-yyyy HH:mm:ss.SSS , for example 12-31-2016 15:30:00.123. 36. dd-MM-yyyy HH:mm:ss.SSS , for example 31-12-2016 15:30:00.123. 37. yyyy-MM-dd HH:mm:ss.SSS , for example 2016-12-31 15:30:00.123. 38. MMM-dd-yyyy HH:mm:ss.SSS , for example DEC-31-2016 15:30:00.123. 39. dd-MMM-yyyy HH:mm:ss.SSS , for example 31-DEC-2016 15:30:00.123.

Data source	Clocks	Date formats
		40. yyyy-MMM-dd HH:mm:ss.SSS , for example 2016-DEC-31 15:30:00.123.

Unsupported values in data

If a field contains values that don't conform with the data type that Amazon Quick Sight assigns to the field, the rows containing those values are skipped. For example, take the following source data.

Sales ID	Sales Date	Sales Amount
001	10/14/2015	12.43
002	5/3/2012	25.00
003	Unknown	18.17
004	3/8/2009	86.02

Amazon Quick Sight interprets **Sales Date** as a date field and drops the row containing a nondate value, so only the following rows are imported.

Sales ID	Sales Date	Sales Amount
001	10/14/2015	12.43
002	5/3/2012	25.00
004	3/8/2009	86.02

In some cases, a database field might contain values that the JDBC driver can't interpret for the source database engine. In such cases, the uninterpretable values are replaced by null so that the rows can be imported. The only known occurrence of this issue is with MySQL date, datetime, and timestamp fields that have all-zero values, for example **0000-00-00 00:00:00**. For example, take the following source data.

Sales ID	Sales Date	Sales Amount
001	2004-10-12 09:14:27	12.43
002	2012-04-07 12:59:03	25.00

003	0000-00-00 00:00:00	18.17
004	2015-09-30 01:41:19	86.02

In this case, the following data is imported.

Sales ID	Sales Date	Sales Amount

001	2004-10-12 09:14:27	12.43
002	2012-04-07 12:59:03	25.00
003	(null)	18.17
004	2015-09-30 01:41:19	86.02

Working with datasets

Datasets are the foundation of your Quick Sight analytics, serving as the prepared and structured data sources that power your analyses and dashboards. Once you've created datasets from your data sources, you need to manage them effectively throughout their lifecycle to ensure reliable, secure, and collaborative analytics.

This section covers the complete dataset management workflow, from editing and versioning datasets to sharing them with team members and implementing security controls. You'll learn how to maintain dataset integrity while supporting collaborative analytics, track which analyses depend on your datasets, and implement both row-level and column-level security to protect sensitive information. Whether you're preparing datasets for team use, troubleshooting analysis issues, or implementing data governance policies, these topics provide the essential knowledge for effective dataset management in Quick Sight.

Topics

- [Creating datasets](#)
- [Editing datasets](#)
- [Reverting datasets back to previous published versions](#)
- [Duplicating datasets](#)
- [Sharing datasets](#)
- [Tracking dashboards and analyses that use a dataset](#)
- [Using dataset parameters in Amazon Quick Suite](#)
- [Using row-level security in Amazon Quick Suite](#)
- [Using column-level security to restrict access to a dataset](#)

- [Running queries as an IAM role in Amazon Quick Suite](#)
- [Deleting datasets](#)
- [Adding a dataset to an analysis](#)

Creating datasets

You can create datasets from new or existing data sources in Amazon Quick Suite. You can use a variety of database data sources to provide data to Amazon Quick Suite. This includes Amazon RDS instances and Amazon Redshift clusters. It also includes MariaDB, Microsoft SQL Server, MySQL, Oracle, and PostgreSQL instances in your organization, Amazon EC2, or similar environments.

Topics

- [Creating datasets using new data sources](#)
- [Creating a dataset using an existing data source](#)
- [Creating a dataset using an existing dataset in Amazon Quick Suite](#)

Creating datasets using new data sources

When you create a dataset based on an Amazon service like Amazon RDS, Amazon Redshift, or Amazon EC2, data transfer charges might apply when consuming data from that source. Those charges might also vary depending on whether that Amazon resource is in the home Amazon Web Services Region that you chose for your Amazon Quick Suite account. For details on pricing, see the pricing page for the service in question.

When creating a new database dataset, you can select one table, join several tables, or create a SQL query to retrieve the data that you want. You can also change whether the dataset uses a direct query or instead stores data in [SPICE](#).

To create a new dataset

1. To create a dataset, choose **New data set** on the **Data** page. You can then create a dataset based on an existing dataset or data source, or connect to a new data source and base the dataset on that.
2. Provide connection information to the data source:
 - For local text or Microsoft Excel files, you can simply identify the file location and upload the file.

- For Amazon S3, provide a manifest identifying the files or buckets that you want to use, and also the import settings for the target files.
- For Amazon Athena, all Athena databases for your Amazon account are returned. No additional credentials are required.
- For Salesforce, provide credentials to connect with.
- For Amazon Redshift, Amazon RDS, Amazon EC2, or other database data sources, provide information about the server and database that host the data. Also provide valid credentials for that database instance.

Creating a dataset from a database

The following procedures walk you through connecting to database data sources and creating datasets. To create datasets from Amazon data sources that your Amazon Quick Suite account autodiscovered, use [Creating a dataset from an autodiscovered Amazon Redshift cluster or Amazon RDS instance](#). To create datasets from any other database data sources, use [Creating a dataset using a database that's not autodiscovered](#).

Creating a dataset from an autodiscovered Amazon Redshift cluster or Amazon RDS instance

Use the following procedure to create a connection to an autodiscovered Amazon data source.

To create a connection to an autodiscovered Amazon data source

1. Check [Data source quotas](#) to make sure that your target table or query doesn't exceed data source quotas.
2. Confirm that the database credentials you plan to use have appropriate permissions as described in [Required permissions](#).
3. Make sure that you have configured the cluster or instance for Amazon Quick Suite access by following the instructions in [Network and database configuration requirements](#).
4. On the Amazon Quick Suite start page, choose **Data**.
5. Choose **Create** then choose **New dataset**.
6. Choose either the **RDS** or the **Redshift Auto-discovered** icon, depending on the Amazon service that you want to connect to.
7. Enter the connection information for the data source, as follows:
 - For **Data source name**, enter a name for the data source.

- For **Instance ID**, choose the name of the instance or cluster that you want to connect to.
 - **Database name** shows the default database for the **Instance ID** cluster or instance. To use a different database on that cluster or instance, enter its name.
 - For **UserName**, enter the user name of a user account that has permissions to do the following:
 - Access the target database.
 - Read (perform a SELECT statement on) any tables in that database that you want to use.
 - For **Password**, enter the password for the account that you entered.
8. Choose **Validate connection** to verify your connection information is correct.
 9. If the connection validates, choose **Create data source**. If not, correct the connection information and try validating again.

 Note

Amazon Quick Suite automatically secures connections to Amazon RDS instances and Amazon Redshift clusters by using Secure Sockets Layer (SSL). You don't need to do anything to enable this.

10. Choose one of the following:

- **Custom SQL**

On the next screen, you can choose to write a query with the **Use custom SQL** option. Doing this opens a screen named **Enter custom SQL query**, where you can enter a name for your query, and then enter the SQL. For best results, compose the query in a SQL editor, and then paste it into this window. After you name and enter the query, you can choose **Edit/Preview data** or **Confirm query**. Choose **Edit/Preview data** to immediately go to data preparation. Choose **Confirm query** to validate the SQL and make sure that there are no errors.

- **Choose tables**

To connect to specific tables, for **Schema: contain sets of tables**, choose **Select** and then choose a schema. In some cases where there is only a single schema in the database, that schema is automatically chosen, and the schema selection option isn't displayed.

To prepare the data before creating an analysis, choose **Edit/Preview data** to open data preparation. Use this option if you want to join to more tables.

Otherwise, after choosing a table, choose **Select**.

11. Choose one of the following options:

- Prepare the data before creating an analysis. To do this, choose **Edit/Preview data** to open data preparation for the selected table. For more information about data preparation, see [Preparing dataset examples](#).
- Create a dataset and analysis using the table data as-is and to import the dataset data into SPICE for improved performance (recommended). To do this, check the table size and the SPICE indicator to see if you have enough capacity.

If you have enough SPICE capacity, choose **Import to SPICE for quicker analytics**, and then create an analysis by choosing **Visualize**.

 Note

If you want to use SPICE and you don't have enough space, choose **Edit/Preview data**. In data preparation, you can remove fields from the dataset to decrease its size. You can also apply a filter or write a SQL query that reduces the number of rows or columns returned. For more information about data preparation, see [Preparing dataset examples](#).

- To create a dataset and an analysis using the table data as-is, and to have the data queried directly from the database, choose the **Directly query your data** option. Then create an analysis by choosing **Visualize**.

Creating a dataset using a database that's not autodiscovered

Use the following procedure to create a connection to any database other than an autodiscovered Amazon Redshift cluster or Amazon RDS instance. Such databases include Amazon Redshift clusters and Amazon RDS instances that are in a different Amazon Web Services Region or are associated with a different Amazon account. They also include MariaDB, Microsoft SQL Server, MySQL, Oracle, and PostgreSQL instances that are on-premises, in Amazon EC2, or in some other accessible environment.

To create a connection to a database that isn't an autodiscovered Amazon Redshift cluster or RDS instance

1. Check [Data source quotas](#) to make sure that your target table or query doesn't exceed data source quotas.
2. Confirm that the database credentials that you plan to use have appropriate permissions as described in [Required permissions](#).
3. Make sure that you have configured the cluster or instance for Amazon Quick Suite access by following the instructions in [Network and database configuration requirements](#).
4. On the Amazon Quick Suite start page, choose **Manage data**.
5. Choose **Create** then choose **New data set**.
6. Choose the **Redshift Manual connect** icon if you want to connect to an Amazon Redshift cluster in another Amazon Web Services Region or associated with a different Amazon account. Or choose the appropriate database management system icon to connect to an instance of Amazon Aurora, MariaDB, Microsoft SQL Server, MySQL, Oracle, or PostgreSQL.
7. Enter the connection information for the data source, as follows:
 - For **Data source name**, enter a name for the data source.
 - For **Database server**, enter one of the following values:
 - For an Amazon Redshift cluster or Amazon RDS instance, enter the endpoint of the cluster or instance without the port number. For example, if the endpoint value is `clustername.1234abcd.us-west-2.redshift.amazonaws.com:1234`, then enter `clustername.1234abcd.us-west-2.redshift.amazonaws.com`. You can get the endpoint value from the **Endpoint** field on the cluster or instance detail page in the Amazon console.
 - For an Amazon EC2 instance of MariaDB, Microsoft SQL Server, MySQL, Oracle, or PostgreSQL, enter the public DNS address. You can get the public DNS value from the **Public DNS** field on the instance detail pane in the Amazon EC2 console.
 - For a non-Amazon EC2 instance of MariaDB, Microsoft SQL Server, MySQL, Oracle, or PostgreSQL, enter the hostname or public IP address of the database server. If you are using Secure Sockets Layer (SSL) for a secured connection (recommended), you likely need to provide the hostname to match the information required by the SSL certificate. For a list of accepted certificates see [Amazon Quick Suite SSL and CA certificates](#).
 - For **Port**, enter the port that the cluster or instance uses for connections.
 - For **Database name**, enter the name of the database that you want to use.

- For **UserName**, enter the user name of a user account that has permissions to do the following:
 - Access the target database.
 - Read (perform a SELECT statement on) any tables in that database that you want to use.
 - For **Password**, enter the password associated with the account you entered.
8. (Optional) If you are connecting to anything other than an Amazon Redshift cluster and you *don't* want a secured connection, make sure that **Enable SSL** is clear. *We strongly recommend leaving this checked*, because an unsecured connection can be open to tampering.

For more information on how the target instance uses SSL to secure connections, see the documentation for the target database management system. Amazon Quick Suite doesn't accept self-signed SSL certificates as valid. For a list of accepted certificates, see [Amazon Quick Suite SSL and CA certificates](#).

Amazon Quick Suite automatically secures connections to Amazon Redshift clusters by using SSL. You don't need to do anything to enable this.

Some databases, such as Presto and Apache Spark, must meet additional requirements before Amazon Quick Suite can connect. For more information, see [Creating a data source using Presto](#), or [Creating a data source using Apache Spark](#).

9. (Optional) Choose **Validate connection** to verify your connection information is correct.
10. If the connection validates, choose **Create data source**. If not, correct the connection information and try validating again.
11. Choose one of the following:
- **Custom SQL**

On the next screen, you can choose to write a query with the **Use custom SQL** option. Doing this opens a screen named **Enter custom SQL query**, where you can enter a name for your query, and then enter the SQL. For best results, compose the query in a SQL editor, and then paste it into this window. After you name and enter the query, you can choose **Edit/Preview data** or **Confirm query**. Choose **Edit/Preview data** to immediately go to data preparation. Choose **Confirm query** to validate the SQL and make sure that there are no errors.

- **Choose tables**

To connect to specific tables, for **Schema: contain sets of tables**, choose **Select** and then choose a schema. In some cases where there is only a single schema in the database, that schema is automatically chosen, and the schema selection option isn't displayed.

To prepare the data before creating an analysis, choose **Edit/Preview data** to open data preparation. Use this option if you want to join to more tables.

Otherwise, after choosing a table, choose **Select**.

12. Choose one of the following options:

- Prepare the data before creating an analysis. To do this, choose **Edit/Preview data** to open data preparation for the selected table. For more information about data preparation, see [Preparing dataset examples](#).
- Create a dataset and an analysis using the table data as-is and import the dataset data into SPICE for improved performance (recommended). To do this, check the table size and the SPICE indicator to see if you have enough space.

If you have enough SPICE capacity, choose **Import to SPICE for quicker analytics**, and then create an analysis by choosing **Visualize**.

 Note

If you want to use SPICE and you don't have enough space, choose **Edit/Preview data**. In data preparation, you can remove fields from the dataset to decrease its size. You can also apply a filter or write a SQL query that reduces the number of rows or columns returned. For more information about data preparation, see [Preparing dataset examples](#).

- Create a dataset and an analysis using the table data as-is and have the data queried directly from the database. To do this, choose the **Directly query your data** option. Then create an analysis by choosing **Visualize**.

Creating a dataset using an existing data source

After you make an initial connection to a Salesforce, Amazon data store, or other database data source, Amazon Quick Suite saves the connection information. It adds the data source to the **FROM**

EXISTING DATA SOURCES section of the **Create a Data Set** page. You can use these existing data sources to create new datasets without respecifying connection information.

Creating a dataset using an existing Amazon S3 data source

Use the following procedure to create a dataset using an existing Amazon S3 data source.

To create a dataset using an existing S3 data source

1. On the Amazon Quick Suite start page, choose **Data**.
2. Choose **Create** then choose **New dataset**.
3. Choose the Amazon S3 data source to use.
4. To prepare the data before creating the dataset, choose **Edit/Preview data**. To create an analysis using the data as-is, choose **Visualize**.

Creating a dataset using an existing Amazon Athena data source

To create a dataset using an existing Amazon Athena data source, use the following procedure.

To create a dataset from an existing Athena connection profile

1. On the Amazon Quick Suite start page, choose **Data**.
2. Choose **Create** then choose **New data set**.

Choose the connection profile icon for the existing data source that you want to use.

Connection profiles are labeled with the data source icon and the name provided by the person who created the connection.

3. Choose **Create data set**.

Amazon Quick Suite creates a connection profile for this data source based only on the Athena workgroup. The database and table aren't saved.

4. On the **Choose your table** screen, do one of the following:
 - To write a SQL query, choose **Use custom SQL**.
 - To choose a database and table, first select your database from the **Database** list. Next, choose a table from the list that appears for your database.

Create a dataset using an existing Salesforce data source

Use the following procedure to create a dataset using an existing Salesforce data source.

To create a dataset using an existing Salesforce data source

1. On the Amazon Quick Suite start page, choose **Data**.
2. Choose **Create** then choose **New data set**.
3. Choose the Salesforce data source to use.
4. Choose **Create Data Set**.
5. Choose one of the following:

- **Custom SQL**

On the next screen, you can choose to write a query with the **Use custom SQL** option. Doing this opens a screen named **Enter custom SQL query**, where you can enter a name for your query, and then enter the SQL. For best results, compose the query in a SQL editor, and then paste it into this window. After you name and enter the query, you can choose **Edit/Preview data** or **Confirm query**. Choose **Edit/Preview data** to immediately go to data preparation. Choose **Confirm query** to validate the SQL and make sure that there are no errors.

- **Choose tables**

To connect to specific tables, for **Data elements: contain your data**, choose **Select** and then choose either **REPORT** or **OBJECT**.

To prepare the data before creating an analysis, choose **Edit/Preview data** to open data preparation. Use this option if you want to join to more tables.

Otherwise, after choosing a table, choose **Select**.

6. On the next screen, choose one of the following options:
 - To create a dataset and an analysis using the data as-is, choose **Visualize**.

Note

If you don't have enough [SPICE](#) capacity, choose **Edit/Preview data**. In data preparation, you can remove fields from the dataset to decrease its size or apply a

filter that reduces the number of rows returned. For more information about data preparation, see [Preparing dataset examples](#).

- To prepare the data before creating an analysis, choose **Edit/Preview data** to open data preparation for the selected report or object. For more information about data preparation, see [Preparing dataset examples](#).

Creating a dataset using an existing database data source

Use the following procedure to create a dataset using an existing database data source.

To create a dataset using an existing database data source

1. On the Amazon Quick Suite start page, choose **Data**.
2. Choose **Create** then choose **New data set**.
3. Choose the database data source to use, and then choose **Create Data Set**.
4. Choose one of the following:

- **Custom SQL**

On the next screen, you can choose to write a query with the **Use custom SQL** option. Doing this opens a screen named **Enter custom SQL query**, where you can enter a name for your query, and then enter the SQL. For best results, compose the query in a SQL editor, and then paste it into this window. After you name and enter the query, you can choose **Edit/Preview data** or **Confirm query**. Choose **Edit/Preview data** to immediately go to data preparation. Choose **Confirm query** to validate the SQL and make sure that there are no errors.

- **Choose tables**

To connect to specific tables, for **Schema: contain sets of tables**, choose **Select** and then choose a schema. In some cases where there is only a single schema in the database, that schema is automatically chosen, and the schema selection option isn't displayed.

To prepare the data before creating an analysis, choose **Edit/Preview data** to open data preparation. Use this option if you want to join to more tables.

Otherwise, after choosing a table, choose **Select**.

5. Choose one of the following options:

- Prepare the data before creating an analysis. To do this, choose **Edit/Preview data** to open data preparation for the selected table. For more information about data preparation, see [Preparing dataset examples](#).
- Create a dataset and an analysis using the table data as-is and import the dataset data into [SPICE](#) for improved performance (recommended). To do this, check the SPICE indicator to see if you have enough space.

If you have enough SPICE capacity, choose **Import to SPICE for quicker analytics**, and then create an analysis by choosing **Visualize**.

 Note

If you want to use SPICE and you don't have enough space, choose **Edit/Preview data**. In data preparation, you can remove fields from the dataset to decrease its size. You can also apply a filter or write a SQL query that reduces the number of rows or columns returned. For more information about data preparation, see [Preparing dataset examples](#).

- Create a dataset and an analysis using the table data as-is and have the data queried directly from the database. To do this, choose the **Directly query your data** option. Then create an analysis by choosing **Visualize**.

Creating a dataset using an existing dataset in Amazon Quick Suite

After you create a dataset in Amazon Quick Suite, you can create additional datasets using it as a source. When you do this, any data preparation that the parent dataset contains, such as any joins or calculated fields, is kept. You can add additional preparation to the data in the new child datasets, such as joining new data and filtering data. You can also set up your own data refresh schedule for the child dataset and track the dashboards and analyses that use it.

Child datasets that are created using a dataset with RLS rules active as a source inherit the parent dataset's RLS rules. Users who are creating a child dataset from a larger parent dataset can only see the data that they have access to in the parent dataset. Then, you can add more RLS rules to the new child dataset in addition to the inherited RLS rules to further manage who can access the data that is in the new dataset. You can only create child datasets from datasets with RLS rules active in Direct Query.

Creating datasets from existing Quick Suite datasets has the following advantages:

- **Central management of datasets** – Data engineers can easily scale to the needs of multiple teams within their organization. To do this, they can develop and maintain a few general-purpose datasets that describe the organization's main data models.
- **Reduction of data source management** – Business analysts (BAs) often spend lots of time and effort requesting access to databases, managing database credentials, finding the right tables, and managing Quick Suite data refresh schedules. Building new datasets from existing datasets means that BAs don't have to start from scratch with raw data from databases. They can start with curated data.
- **Predefined key metrics** – By creating datasets from existing datasets, data engineers can centrally define and maintain critical data definitions across their company's many organizations. Examples might be sales growth and net marginal return. With this feature, data engineers can also distribute changes to those definitions. This approach means that their business analysts can get started with visualizing the right data more quickly and reliably.
- **Flexibility to customize data** – By creating datasets from existing datasets, business analysts get more flexibility to customize datasets for their own business needs. They can avoid worry about disrupting data for other teams.

For example, let's say that you're part of an ecommerce central team of five data engineers. You and your team has access to sales, orders, cancellations, and returns data in a database. You have created a Quick Suite dataset by joining 18 other dimension tables through a schema. A key metric that your team has created is the calculated field, order product sales (OPS). Its definition is: $OPS = \text{product quantity} \times \text{price}$.

Your team serves over 100 business analysts across 10 different teams in eight countries. These are the Coupons team, the Outbound Marketing team, the Mobile Platform team, and the Recommendations team. All of these teams use the OPS metric as a base to analyze their own business line.

Rather than manually creating and maintaining hundreds of unconnected datasets, your team reuses datasets to create multiple levels of datasets for teams across the organization. Doing this centralizes data management and allows each team to customize the data for their own needs. At the same time, this syncs updates to the data, such as updates to metric definitions, and maintains row-level and column-level security. For example, individual teams in your organization can use the centralized datasets. They can then combine them with the data specific to their team to create new datasets and build analyses on top of them.

Along with using the key OPS metric, other teams in your organization can reuse column metadata from the centralized datasets that you created. For example, the Data Engineering team can define metadata, such as *name*, *description*, *data type*, and *folders*, in a centralized dataset. All subsequent teams can use it.

Note

Amazon Quick Suite supports creating up to two additional levels of datasets from a single dataset.

For example, from a parent dataset, you can create a child dataset and then a grandchild dataset for a total of three dataset levels.

Creating a dataset from an existing dataset

Use the following procedure to create a dataset from an existing dataset.

To create a dataset from an existing dataset

1. From the Quick Suite start page, choose **Data** in the pane at left.
2. Choose **Create** then choose the dataset that you want to use to create a new dataset.
3. On the page that opens for that dataset, choose the drop-down menu for **Use in analysis**, and then choose **Use in dataset**.

The data preparation page opens and preloads everything from the parent dataset, including calculated fields, joins, and security settings.

4. On the data preparation page that opens, for **Query mode** at bottom left, choose how you want the dataset to pull in changes and updates from the original, parent dataset. You can choose the following options:
 - **Direct query** – This is the default query mode. If you choose this option, the data for this dataset automatically refreshes when you open an associated dataset, analysis, or dashboard. However, the following limitations apply:
 - If the parent dataset allows direct querying, you can use direct query mode in the child dataset.
 - If you have multiple parent datasets in a join, you can choose direct query mode for your child dataset only if all the parents are from the same underlying data source. For example, the same Amazon Redshift connection.

- Direct query is supported for a single SPICE parent dataset. It is not supported for multiple SPICE parent datasets in a join.
 - **SPICE** – If you choose this option, you can set up a schedule for your new dataset to sync with the parent dataset. For more information about creating SPICE refresh schedules for datasets, see [Refreshing SPICE data](#).
5. (Optional) Prepare your data for analysis. For more information about preparing data, see [Preparing data in Amazon Quick Sight](#).
 6. (Optional) Set up row-level or column-level security (RLS/CLS) to restrict access to the dataset. For more information about setting up RLS, see [Using row-level security with user-based rules to restrict access to a dataset](#). For more information about setting up CLS, see [Using column-level security to restrict access to a dataset](#).
- 
- Note**
- You can set up RLS/CLS on child datasets only. RLS/CLS on parent datasets is not supported.
7. When you're finished, choose **Save & publish** to save your changes and publish the new child dataset. Or choose **Publish & visualize** to publish the new child dataset and begin visualizing your data.

Restricting others from creating new datasets from your dataset

When you create a dataset in Amazon Quick Suite, you can prevent others from using it as a source for other datasets. You can specify if others can use it to create any datasets at all. Or you can specify the type of datasets others can or can't create from your dataset, such as direct query datasets or SPICE datasets.

Use the following procedure to learn how to restrict others from creating new datasets from your dataset.

To restrict others from creating new datasets from your dataset

1. From the Quick Suite start page, choose **Data** in the pane at left.
2. Choose **Create** then choose the dataset that you want to restrict creating new datasets from.
3. On the page that opens for that dataset, choose **Edit dataset**.

4. On the data preparation page that opens, choose **Manage** at upper right, and then choose **Properties**.
5. In the **Dataset properties** pane that opens, choose from the following options:
 - To restrict anyone from creating any type of new datasets from this dataset, turn off **Allow new datasets to be created from this one**.

The toggle is blue when creating new datasets is allowed. It's gray when creating new datasets isn't allowed.

- To restrict others from creating direct query datasets, clear **Allow direct query**.
- To restrict others from creating SPICE copies of your dataset, clear **Allow SPICE copies**.

For more information about SPICE datasets, see [Importing data into SPICE](#).

6. Close the pane.

Editing datasets

You can edit an existing dataset to perform data preparation. For more information about Quick Sight data preparation functionality, see [Preparing data in Amazon Quick Sight](#).

You can open a dataset for editing from the **Datasets** page, or from the analysis page. Editing a dataset from either location modifies the dataset for all analyses that use it.

Things to consider when editing datasets

In two situations, changes to a dataset might cause concern. One is if you deliberately edit the dataset. The other is if your data source has changed so much that it affects the analyses based on it.

Important

Analyses that are in production usage should be protected so they continue to function correctly.

We recommend the following when you're dealing with data changes:

- Document your data sources and datasets, and the visuals that rely upon them. Documentation should include screenshots, fields used, placement in field wells, filters, sorts, calculations, colors,

formatting, and so on. Record everything that you need to recreate the visual. You can also track which Quick Sight resources use a dataset in the dataset management options. For more information, see [Tracking dashboards and analyses that use a dataset](#).

- When you edit a dataset, try not to make changes that might break existing visuals. For example, don't remove columns that are being used in a visual. If you must remove a column, create a calculated column in its place. The replacement column should have the same name and data type as the original.
- If your data source or dataset changes in your source database, adapt your visual to accommodate the change, as described previously. Or you can try to adapt the source database. For example, you might create a view of the source table (document). Then if the table changes, you can adjust the view to include or exclude columns (attributes), change data types, fill null values, and so on. Or, in another circumstance, if your dataset is based on a slow SQL query, you might create a table to hold the results of the query.

If you can't sufficiently adapt the source of the data, recreate the visuals based on your documentation of the analysis.

- If you no longer have access to a data source, your analyses based on that source are empty. The visuals that you created still exist, but they can't display until they have some data to show. This result can happen if permissions are changed by your administrator.
- If you remove the dataset a visual is based on, you might need to recreate it from your documentation. You can edit the visual and select a new dataset to use with it. If you need to consistently use a new file to replace an older one, store your data in a location that is consistently available. For example, you might store your .csv file in Amazon S3 and create an S3 dataset to use for your visuals. For more information on access files stored in S3, see [Creating a dataset using Amazon S3 files](#).

Or you can import the data into a table, and base your visual on a query. This way, the data structures don't change, even if the data contained in them changes.

- To centralize data management, consider creating general, multiple-purpose datasets that others can use to create their own datasets from. For more information, see [Creating a dataset using an existing dataset in Amazon Quick Suite](#).

Editing a dataset from the Datasets page

1. From the Quick Suite start page, choose **Data** at left.

2. On the **Data** page that opens, choose the dataset that you want to edit, and then choose **Edit dataset** at upper right.

The data preparation page opens. For more information about the types of edits you can make to datasets, see [Preparing data in Amazon Quick Sight](#).

Editing a dataset in an analysis

Use the following procedure to edit a dataset from the analysis page.

To edit a dataset from the analysis page

1. In your analysis, choose the pencil icon at the top of the **Fields list** pane.
2. In **Data sets in this analysis** page that opens, choose the three dots at right of the dataset that you want to edit, and then choose **Edit**.

The dataset opens in the data preparation page. For more information about the types of edits you can make to datasets, see [Preparing data in Amazon Quick Sight](#).

Reverting datasets back to previous published versions

When you save and publish changes to a dataset in Amazon Quick Sight, a new version of the dataset is created. At any time, you can see a list of all the previous published versions of that dataset. You can also preview a specific version in that history, or even revert the dataset back to a previous version, if needed.

The following limitations apply to dataset versioning:

- Only the most recent 1,000 versions of a dataset are shown in the publishing history, and are available for versioning.
- After you exceed 1,000 published versions, the oldest versions are automatically removed from the publishing history, and the dataset can no longer be reverted back to them.

Use the following procedure to revert a dataset to a previous published version.

To revert a dataset to a previous published version

1. From the Quick Suite start page, choose **Data**.

2. On the **Data** page, choose a dataset, and then choose **Edit dataset** at upper right.

For more information about editing datasets, see [Editing datasets](#).

3. On the dataset preparation page that opens, choose the **Manage** icon in the blue toolbar at upper right, and then choose **Publishing history**.

A list of previous published versions appears at right.

4. In the **Publishing history** pane, find the version that you want and choose **Revert**.

To preview the version before reverting, choose **Preview**.

The dataset is reverted and a confirmation message appears. The **Publishing history** pane also updates to show the active version of the dataset.

Troubleshooting reverting versions

Sometimes, the dataset can't be reverted to a specific version for one the following reasons:

- The dataset uses one or more data sources that were deleted.

If this error occurs, you can't revert the dataset to a previous version.

- Reverting would make a calculated field not valid.

If this error occurs, you can edit or remove the calculated field, and then save the dataset. Doing this creates a new version of the dataset.

- One or more columns are missing in the data source.

If this error occurs, Quick Sight shows the latest schema from the data source in the preview to reconcile differences between versions. Any calculated field, field name, field type, and filter changes shown in the schema preview are from the version that you want to revert to. You can save this reconciled schema as a new version of the dataset. Or you can return to the active (latest) version by choosing **Preview** on the top (latest) version in the publishing history.

Duplicating datasets

You can duplicate an existing dataset to save a copy of it with a new name. The new dataset is a completely separate copy.

The **Duplicate dataset** option is available if both of the following are true: you own the dataset and you have permission to the data source.

To duplicate a dataset

1. From the Quick Suite start page, choose **Data** at left.
2. Choose the dataset that you want to duplicate.
3. On the dataset details page that opens, choose the drop-down for **Edit dataset**, and then choose **Duplicate**.
4. On the Duplicate dataset page that opens, give the duplicated dataset a name, and then choose **Duplicate**.

The duplicated dataset details page opens. From this page, you can edit the dataset, set up a refresh schedule, and more.

Sharing datasets

You can give other Quick Sight users and groups access to a dataset by sharing it with them. Then they can create analyses from it. If you make them co-owners, they can also refresh, edit, delete, or reshare the dataset.

Sharing a dataset

If you have owner permissions on a dataset, use the following procedure to share it.

To share a dataset

1. From the Quick Suite start page, choose **Data** at left.
2. On the **Data** page, choose the dataset that you want to share.
3. On the dataset details page that opens, choose the **Permissions** tab, and then choose **Add users & groups**.
4. Enter the user or group that you want to share this dataset with, and then choose **Add**. You can only invite users who belong to the same Quick Suite account.

Repeat this step until you have entered information for everyone you want to share the dataset with.

5. For the **Permissions** column, choose a role for each user or group to give them permissions on the dataset.

Choose **Viewer** to allow the user to create analyses and datasets from the dataset. Choose **Owner** to allow the user to do that and also refresh, edit, delete, and reshare the dataset.

Users receive emails with a link to the dataset. Groups don't receive invitation emails.

Viewing and editing the permissions of users that a dataset is shared with

If you have owner permissions on a dataset, you can use the following procedure to view, edit, or change user access to it.

To view, edit, or change user access to a dataset if you have owner permissions for it

1. From the Quick Suite start page, choose **Data** at left.
2. On the **Data** page, choose the dataset that you want to share.
3. On the dataset details page that opens, choose the **Permissions** tab.

A list of all users and groups with access to the dataset is displayed.

4. (Optional) To change permission roles for a user or group, choose the drop-down menu in the **Permissions** column for the user or group. Then choose either **Viewer** or **Owner**.

Revoking access to a dataset

If you have owner permissions on a dataset, you can use the following procedure to revoke user access to a dataset.

To revoke user access to a dataset if you have owner permissions for it

1. From the Quick Suite start page, choose **Data** at left.
2. On the **Data** page, choose the dataset that you want to share.
3. On the dataset details page that opens, choose the **Permissions** tab.

A list of all users and groups with access to the dataset is displayed.

4. In the **Actions** column for the user or group, choose **Revoke access**.

Tracking dashboards and analyses that use a dataset

When you create a dataset in Quick Sight, you can track which dashboards and analyses use that dataset. This approach is useful when you want to see which resources will be affected when you make changes to a dataset, or want to delete a dataset.

Use the following procedure to see which dashboards and analyses use a dataset.

To track resources that use a dataset

1. From the Quick Suite start page, choose **Data** in the pane at left.
2. On the **Data** page, choose the dataset that you want to track resources for.
3. In the page that opens for that dataset, choose **Edit dataset**.
4. In the data preparation page that opens, choose **Manage** at upper right, and then choose **Usage**.
5. The dashboards and analyses that use the dataset are listed in the pane that opens.

Using dataset parameters in Amazon Quick Suite

In Amazon Quick Suite, authors can use dataset parameters in direct query to dynamically customize their datasets and apply reusable logic to their datasets. A *dataset parameter* is a parameter created at the dataset level. It's consumed by an analysis parameter through controls, calculated fields, filters, actions, URLs, titles, and descriptions. For more information on analysis parameters, see [Parameters in Amazon Quick Suite](#). The following list describes three actions that can be performed with dataset parameters:

- **Custom SQL in direct query** – Dataset owners can insert dataset parameters into the custom SQL of a direct query dataset. When these parameters are applied to a filter control in a Quick Suite analysis, users can filter their custom data faster and more efficiently.
- **Repeatable variables** – Static values that appear in multiple locations in the dataset page can be modified in one action using custom dataset parameters.
- **Move calculated fields to datasets** – Quick Suite authors can copy calculated fields with parameters in an analysis and migrate them to the dataset level. This protects calculated fields at the analysis level from being accidentally modified and calculated fields be shared across multiple analyses.

In some situations, dataset parameters improve filter control performance for direct query datasets that require complex custom SQL and simplify business logic at the dataset level.

Topics

- [Dataset parameter limitations](#)
- [Creating dataset parameters in Amazon Quick Suite](#)
- [Inserting dataset parameters into custom SQL](#)
- [Adding dataset parameters to calculated fields](#)
- [Adding dataset parameters to filters](#)
- [Using dataset parameters in Quick Suite analyses](#)
- [Advanced use cases of dataset parameters](#)

Dataset parameter limitations

This section covers known limitations that you might encounter when working with dataset parameters in Amazon Quick Suite.

- When dashboard readers schedule emailed reports, selected controls don't propagate to the dataset parameters that are included in the report that's attached to the email. Instead, the default values of the parameters are used.
- Dataset parameters can't be inserted into custom SQL of datasets stored in SPICE.
- Dynamic defaults can only be configured on the analysis page of the analysis that is using the dataset. You can't configure a dynamic default at the dataset level.
- The **Select all** option is not supported on multivalue controls of analysis parameters that are mapped to dataset parameters.
- Cascading controls are not supported for dataset parameters.
- Dataset parameters can only be used by dataset filters when the dataset is using direct query.
- In a custom SQL query, only 128 dataset parameters can be used.

Creating dataset parameters in Amazon Quick Suite

Use the following procedures to get started using dataset parameters.

To create a new dataset parameter

1. From the Quick Suite start page, choose **Data** on the left, choose the ellipsis (three dots) next to the dataset that you want to change, and then choose **Edit**.
2. On the **Dataset** page that opens, choose **Parameters** on the left, and then choose the (+) icon to create a new dataset parameter.
3. In the **Create new parameter** pop-up that appears, enter a parameter name in the **Name** box.
4. In the **Data type** dropdown, choose the parameter data type that you want. Supported data types are `String`, `Integer`, `Number`, and `Datetime`. This option can't be changed after the parameter is created.
5. For **Default value**, enter the default value that you want the parameter to have.

Note

When you map a dataset parameter to an analysis parameter, a different default value can be chosen. When this happens, the default value configured here is overridden by the new default value.

6. For **Values**, choose the value type that you want the parameter to have. **Single value** parameters support single-select dropdowns, text field, and list controls. **Multiple values** parameters support multi-select dropdown controls. This option can't be changed after the parameter is created.
7. When you are finished configuring the new parameter, choose **Create** to create the parameter.

Inserting dataset parameters into custom SQL

You can insert dataset parameters into the custom SQL of a dataset in direct query mode by referencing it with `<<$parameter_name>>` in the SQL statement. At runtime, dashboard users can enter filter control values that are associated with a dataset parameter. Then, they can see the results in the dashboard visuals after the values propagate to the SQL query. You can use parameters to create basic filters based on customer input in where clauses. Alternatively, you could add `case when` or `if else` clauses to dynamically change the logic of the SQL query based on a parameter's input.

For example, say you want to add a `WHERE` clause to your custom SQL that filters data based on an end user's Region name. In this case, you create a single value parameter called `RegionName`:

```
SELECT *
FROM transactions
WHERE region = <<$RegionName>>
```

You can also let users provide multiple values to the parameter:

```
SELECT *
FROM transactions
WHERE region in (<<$RegionNames>>)
```

In the following more complex example, a dataset author refers to two dataset parameters twice based on a user's first and last names that can be selected in a dashboard filter control:

```
SELECT Region, Country, OrderDate, Sales
FROM transactions
WHERE region=
(Case
WHEN <<$UserFIRSTNAME>> In
    (select firstname from user where region='region1')
    and <<$UserLASTNAME>> In
    (select lastname from user where region='region1')
    THEN 'region1'
WHEN <<$UserFIRSTNAME>> In
    (select firstname from user where region='region2')
    and <<$UserLASTNAME>> In
    (select lastname from user where region='region2')
    THEN 'region2'
ELSE 'region3'
END)
```

You can also use parameters in SELECT clauses to create new columns in a dataset from user input:

```
SELECT Region, Country, date,
    (case
    WHEN <<$RegionName>>='EU'
    THEN sum(sales) * 0.93    --convert US dollar to euro
    WHEN <<$RegionName>>='CAN'
    THEN sum(sales) * 0.78    --convert US dollar to Canadian Dollar
    ELSE sum(sales) -- US dollar
    END
    ) as "Sales"
```

```
FROM transactions
WHERE region = <<$RegionName>>
```

To create a custom SQL query or to edit an existing query before adding a dataset parameter, see [Using SQL to customize data](#).

When you apply custom SQL with a dataset parameter, <<\$parameter_name>> is used as a placeholder value. When a user chooses one of the parameter values from a control, Quick Suite replaces the placeholder with the values that the user selects on the dashboard.

In the following example, the user enters a new custom SQL query that filters data by state:

```
select * from all_flights
where origin_state_abr = <<$State>>
```

The default value of the parameter is applied to the SQL query and the results appear in the **Preview pane**.

Adding dataset parameters to calculated fields

You can also add dataset parameters to calculated field expressions using the format `${parameter_name}`.

When you create a calculation, you can choose from the existing parameters from the list of parameters under the **Parameters** list. You can't create a calculated field that contains a multivalued parameter.

For more information on adding calculated fields, see [Using calculated fields with parameters in Amazon Quick Suite](#).

Adding dataset parameters to filters

For datasets in direct query mode, dataset authors can use dataset parameters in filters without custom SQL. Dataset parameters can't be added to filters if the dataset is in SPICE.

To add a dataset parameter to a filter

1. Open the dataset page of the dataset that you want to create a filter for. Choose **Filters** on the left, and then choose **Add filter**.
2. Enter the name that you want the filter to have and choose the field that you want filtered in the dropdown.

3. After you create the new filter, navigate to the filter in the **Filters** pane, choose the ellipsis (three dots) next to the filter, and then choose **Edit**.
4. For **Filter type**, choose **Custom filter**.
5. For **Filter condition**, choose the condition that you want.
6. Select the **Use parameter** box and choose the dataset parameter that you want the filter to use.
7. When you are finished making changes, choose **Apply**.

Using dataset parameters in Quick Suite analyses

Once you create a dataset parameter, after you add the dataset to an analysis, map the dataset parameter to a new or existing analysis parameter. After you map a dataset parameter to an analysis parameter, you can use them with filters, controls, and any other analysis parameter feature.

You can manage your dataset parameters in the **Parameters** pane of the analysis that is using the dataset that the parameters belong to. In the **Dataset Parameters** section of the **Parameters** pane, you can choose to see only the unmapped dataset parameters (default). Alternatively, you can choose to see all mapped and unmapped dataset parameters by choosing **ALL** from the **Viewing** dropdown.

Mapping dataset parameters in new Quick Suite analyses

When you create a new analysis from a dataset that contains parameters, you need to map the dataset parameters to the analysis before you can use them. This is also true when you add a dataset with parameters to an analysis. You can view all unmapped parameters in an analysis in the **Parameters** pane of the analysis. Alternatively, choose **VIEW** in the notification message that appears in the top right of the page when you create the analysis or add the dataset.

To map a dataset parameter to an analysis parameter

1. Open the [Quick Suite console](#).
2. Choose the analysis that you want to change.
3. Choose the **Parameters** icon to open the **Parameters** pane.
4. Choose the ellipsis (three dots) next to the dataset parameter that you want to map, choose **Map Parameter**, and then choose the analysis parameter that you want to map your dataset parameter to.

If your analysis doesn't have any analysis parameters, you can choose **Map parameter** and **Create new** to create an analysis parameter that is automatically mapped to the dataset parameter upon creation.

- a. (Optional) In the **Create new parameter** pop-up that appears, for **Name**, enter a name for the new analysis parameter.
- b. (Optional) For **Static default value**, choose the static default value that you want the parameter to have.
- c. (Optional) Choose **Set a dynamic default** to set a dynamic default for the new parameter.
- d. In the **Mapped dataset parameters** table, you will see the dataset parameter that you are mapping to the new analysis parameter. You can add other dataset parameters to this analysis parameter by choosing the **ADD DATASET PARAMETER** dropdown and then choosing the parameter that you want to map. You can unmap a dataset parameter by choosing the **Remove** button next to the dataset parameter that you want to remove.

For more information on creating analysis parameters, see [Setting up parameters in Amazon Quick Suite](#).

When you map a dataset parameter to an analysis parameter, the analysis parameter represents the dataset parameter wherever it is used in the analysis.

You can also map and unmap dataset parameters to analysis parameters in the **Edit parameter** window. To open the **Edit parameter** window, navigate to the **Parameters** pane, choose the ellipsis (three dots) next to the analysis parameter that you want to change, and then choose **Edit parameter**. You can add other dataset parameters to this analysis parameter by choosing the **ADD DATASET PARAMETER** dropdown and then choosing the parameter that you want to map. You can unmap a dataset parameter by choosing the **Remove** button next to the dataset parameter that you want to remove. You can also remove all mapped dataset parameters by choosing **REMOVE ALL**. When you are done making changes, choose **Update**.

When you delete an analysis parameter, all dataset parameters are unmapped from the analysis and appear in the **UNMAPPED** section of the **Parameters** pane. You can only map a dataset parameter to one analysis parameter at a time. To map a dataset parameter to a different analysis parameter, unmap the dataset parameter and then map it to the new analysis parameter.

Adding filter controls to mapped analysis parameters

After you map a dataset parameter to an analysis parameter in Quick Suite, you can create filter controls for filters, actions, calculated fields, titles, descriptions, and URLs.

To add a control to a mapped parameter

1. In the **Parameters** pane of the analysis page, choose the ellipsis (three dots) next to the mapped analysis parameter that you want, and then choose **Add control**.
2. In the **Add control** window that appears, enter the **Name** that you want and choose the **Style** that you want the control to have. For single value controls, choose between Dropdown, List, and Text field. For multivalue controls, choose Dropdown.
3. Choose **Add** to create the control.

Advanced use cases of dataset parameters

This section covers more advanced options and use cases working with dataset parameters and dropdown controls. Use the following walkthroughs to create dynamic dropdown values with dataset parameters.

Using multivalue controls with dataset parameters

When you use dataset parameters that are inserted into the custom SQL of a dataset, the dataset parameters commonly filter data by values from a specific column. If you create a dropdown control and assign the parameter as the value, the dropdown only shows the value that the parameter filtered. The following procedure shows how you can create a control that is mapped to a dataset parameter and shows all unfiltered values.

To populate all assigned values in a dropdown control

1. Create a new single-column dataset in SPICE or direct query that includes all unique values from the original dataset. For example, let's say that your original dataset is using the following custom SQL:

```
select * from all_flights
      where origin_state_abr = <<$State>>
```

To create a single-column table with all unique origin states, apply the following custom SQL to the new dataset:

```
SELECT distinct origin_state_abr FROM all_flights
order by origin_state_abr asc
```

The SQL expression returns all unique states in alphabetic order. The new dataset does not have any dataset parameters.

2. Enter a **Name** for the new dataset, and then save and publish the dataset. In our example, the new dataset is called `State Codes`.
3. Open the analysis that contains the original dataset, and add the new dataset to the analysis. For information on adding datasets to an existing analysis, see [Adding a dataset to an analysis](#).
4. Navigate to the **Controls** pane and find the dropdown control that you want to edit. Choose the ellipsis (three dots) next to the control, and then choose **Edit**.
5. In the **Format control** that appears on the left, and choose **Link to a dataset field** in the **Values** section.
6. For the **Dataset** dropdown that appears, choose the new dataset that you created. In our example, the `State Codes` dataset is chosen.
7. For the **Field** dropdown that appears, choose the appropriate field. In our example, the `origin_state_abr` field is chosen.

After you finish linking the control to the new dataset, all unique values appear in the control's dropdown. These include the values that are filtered out by the dataset parameter.

Using controls with Select all options

By default, when one or more dataset parameters are mapped to an analysis parameter and added to a control, the `Select all` option is not available. The following procedure shows a workaround that uses the same example scenario from the previous section.

Note

This walkthrough is for datasets that are small enough to load in direct query. If you have a large dataset and want to use the `Select All` option, it is recommended that you load the dataset into SPICE. However, if you want to use the `Select All` option with dataset parameters, this walkthrough describes a way to do so.

To begin, let's say you have a direct query dataset with custom SQL that contains a multivalue parameter called `States`:

```
select * from all_flights
where origin_state_abr in (<<$States>>)
```

To use the **Select all** option in a control that uses dataset parameters

1. In the **Parameters** pane of the analysis, find the dataset parameter that you want to use and choose **Edit** from the ellipsis (three dots) next to the parameter.
2. In the **Edit parameter** window that appears, enter a new default value in the **Static multiple default values** section. In our example, the default value is `All States`. Note that the example uses a leading space character so that the default value appears as the first item in the control.
3. Choose **Update** to update the parameter.
4. Navigate to the dataset that contains the dataset parameter that you're using in the analysis-by-analysis. Edit the custom SQL of the dataset to include a default use case for your new static multiple default values. Using the `All States` example, the SQL expression appears as follows:

```
select * from public.all_flights
where
  ' All States' in (<<$States>>) or
  origin_state_abr in (<<$States>>)
```

If the user chooses `All States` in the control, the new SQL expression returns all unique records. If the user chooses a different value from the control, the query returns values that were filtered by the dataset parameter.

Using controls with **Select all** and multivalue options

You can combine the previous **Select all** procedure with the multivalue control method discussed earlier to create dropdown controls that contain a **Select all** value in addition to multiple values that the user can select. This walkthrough assumes that you have followed the previous procedures, that you know how to map dataset parameters to analysis parameters, and that you can create controls in an analysis. For more information on mapping analysis parameters, see [Mapping dataset parameters in new Quick Suite analyses](#). For more information on creating

controls in an analysis that is using dataset parameters, see [Adding filter controls to mapped analysis parameters](#).

To add multiple values to a control with a Select all option and a mapped dataset parameter

1. Open the analysis that has the original dataset with a `Select all` custom SQL expression and a second dataset that includes all possible values of the filtered column that exists in the original dataset.
2. Navigate to the secondary dataset that was created earlier to return all values of a filtered column. Add a custom SQL expression that adds your previously configured `Select all` option to the query. The following example adds the `All States` record to the top of the list of returned values of the dataset:

```
(Select ' All States' as origin_state_abr)
  Union All
  (SELECT distinct origin_state_abr FROM all_flights
   order by origin_state_abr asc)
```

3. Go back to the analysis that the datasets belong to and map the dataset parameter that you are using to the analysis parameter that you created in step 3 of the previous procedure. The analysis parameter and dataset parameter can have the same name. In our example, the analysis parameter is called `States`.
4. Create a new filter control or edit an existing filter control and choose **Hide Select All** to hide the disabled **Select All** option that appears in multivalue controls.

Once you create the control, users can use the same control to select all or multiple values of a filtered column in a dataset.

Using row-level security in Amazon Quick Suite

Applies to: Enterprise Edition

In the Enterprise edition of Amazon Quick Suite, you can restrict access to a dataset by configuring row-level security (RLS) on it. You can do this before or after you have shared the dataset. When you share a dataset with RLS with dataset owners, they can still see all the data. When you share it with readers, however, they can only see the data restricted by the permission dataset rules.

Also, when you embed Amazon Quick Suite dashboards in your application for unregistered users of Quick Suite, you can use row-level security (RLS) with tags. In this case, you use tags to specify which data your users can see in the dashboard depending on who they are.

You can restrict access to a dataset using username or group-based rules, tag-based rules, or both.

Choose user-based rules if you want to secure data for users or groups provisioned (registered) in Quick Suite. To do so, select a permissions dataset that contains rules set by columns for each user or group accessing the data. Only users or groups identified in the rules have access to data.

Choose tag-based rules only if you are using embedded dashboards and want to secure data for users not provisioned (unregistered users) in Quick Suite. To do so, define tags on columns to secure data. Values to tags must be passed when embedding dashboards.

Topics

- [Using row-level security with user-based rules to restrict access to a dataset](#)
- [Using row-level security with tag-based rules to restrict access to a dataset when embedding dashboards for anonymous users](#)

Using row-level security with user-based rules to restrict access to a dataset

Applies to: Enterprise Edition

In the Enterprise edition of Amazon Quick Suite, you can restrict access to a dataset by configuring row-level security (RLS) on it. You can do this before or after you have shared the dataset. When you share a dataset with RLS with dataset owners, they can still see all the data. When you share it with readers, however, they can only see the data restricted by the permission dataset rules. By adding row-level security, you can further control their access.

Note

When applying SPICE datasets to row-level security, each field in the dataset can contain up to 2,047 Unicode characters. Fields that contain more than this quota are truncated during ingestion. To learn more about SPICE data quotas, see [SPICE quotas for imported data](#).

To do this, you create a query or file that has one column named `UserName`, `GroupName`, or both. Or you can create a query or file that has one column named `UserARN`, `GroupARN`, or both. You can think of this as *adding a rule* for that user or group. Then you can add one column to the query or file for each field that you want to grant or restrict access to. For each user or group name that you add, you add the values for each field. You can use `NULL` (no value) to mean all values. To see examples of dataset rules, see [Creating dataset rules for row-level security](#).

To apply the dataset rules, you add the rules as a permissions dataset to your dataset. Keep in mind the following points:

- The permissions dataset can't contain duplicate values. Duplicates are ignored when evaluating how to apply the rules.
- Each user or group specified can see only the rows that *match* the field values in the dataset rules.
- If you add a rule for a user or group and leave all other columns with no value (`NULL`), you grant them access to all the data.
- If you don't add a rule for a user or group, that user or group can't see any of the data.
- The full set of rule records that are applied per user must not exceed 999. This limitation applies to the total number of rules that are directly assigned to a username, plus any rules that are assigned to the user through group names.
- If a field includes a comma (,) Amazon Quick Suite treats each word separated from another by a comma as an individual value in the filter. For example, in ('AWS' , 'INC'), `AWS, INC` is considered as two strings: `AWS` and `INC`. To filter with `AWS, INC`, wrap the string with double quotation marks in the permissions dataset.

If the restricted dataset is a SPICE dataset, the number of filter values applied per user can't exceed 192,000 for each restricted field. This applies to the total number of filter values that are directly assigned to a username, plus any filter values that are assigned to the user through group names.

If the restricted dataset is a direct query dataset, the number of filter values applied per user varies from data sources.

Exceeding the filter value limit may cause the visual rendering to fail. We recommend adding an additional column to your restricted dataset to divide the rows into groups based on the original restricted column so that the filter list can be shortened.

Amazon Quick Suite treats spaces as literal values. If you have a space in a field that you are restricting, the dataset rule applies to those rows. Amazon Quick Suite treats both NULLs and blanks (empty strings "") as "no value". A NULL is an empty field value.

Depending on what data source your dataset is coming from, you can configure a direct query to access a table of permissions. Terms with spaces inside them don't need to be delimited with quotes. If you use a direct query, you can easily change the query in the original data source.

Or you can upload dataset rules from a text file or spreadsheet. If you are using a comma-separated value (CSV) file, don't include any spaces on the given line. Terms with spaces inside them need to be delimited with quotation marks. If you use dataset rules that are file-based, apply any changes by overwriting the existing rules in the dataset's permissions settings.

Datasets that are restricted are marked with the word **RESTRICTED** in the **Data** screen.

Child datasets that are created from a parent dataset that has RLS rules active retain the same RLS rules that the parent dataset has. You can add more RLS rules to the child dataset, but you can't remove the RLS rules that the dataset inherits from the parent dataset.

Child datasets that are created from a parent dataset that has RLS rules active can only be created with Direct Query. Child datasets that inherit the parent dataset's RLS rules aren't supported in SPICE.

Row-level security works only for fields containing textual data (string, char, varchar, and so on). It doesn't currently work for dates or numeric fields. Anomaly detection is not supported for datasets that use row-level security (RLS).

Creating dataset rules for row-level security

Use the following procedure to create a permissions file or query to use as dataset rules.

To create a permissions files or query to use as dataset rules

1. Create a file or a query that contains the dataset rules (permissions) for row-level security.

It doesn't matter what order the fields are in. However, all the fields are case-sensitive. Make sure that they exactly match the field names and values.

The structure should look similar to one of the following. Make sure that you have at least one field that identifies either users or groups. You can include both, but only one is required, and only one is used at a time. The field that you use for users or groups can have any name you choose.

Note

If you are specifying groups, use only Amazon Quick Suite groups or Microsoft AD groups.

The following example shows a table with groups.

GroupName	Region	Segment	
EMEA-Sales	EMEA	Enterprise, SMB, Startup	
US-Sales	US	Enterprise	
US-Sales	US	SMB, Startup	
US-Sales	US	Startup	
APAC-Sales	APAC	Enterprise, SMB	
Corporate-Reporting			
APAC-Sales	APAC	Enterprise, Startup	

The following example shows a table with usernames.

UserName	Region	Segment	
AlejandroRosalez	EMEA	Enterprise, SMB, Startup	
MarthaRivera	US	Enterprise	
NikhilJayashankar	US	SMB, Startup	
PauloSantos	US	Startup	

UserName	Region	Segment	
SaanviSarkar	APAC	Enterprise, SMB	
sales-tps@example.com			
ZhangWei	APAC	Enterprise, Startup	

The following example shows a table with user and group Amazon Resource Names (ARNs).

UserARN	GroupARN	Region
arn:aws-cn:quicksight:us-east-1:123456789012:user/default/Bob	arn:aws-cn:quicksight:us-east-1:123456789012:group/default/group-1	APAC
arn:aws-cn:quicksight:us-east-1:123456789012:user/default/Sam	arn:aws-cn:quicksight:us-east-1:123456789012:group/default/group-2	US

Or if you use a .csv file, the structure should look similar to one of the following.

```

UserName,Region,Segment
AlejandroRosalez,EMEA,"Enterprise,SMB,Startup"
MarthaRivera,US,Enterprise
NikhilJayashankars,US,SMB
PauloSantos,US,Startup
SaanviSarkar,APAC,"SMB,Startup"
sales-tps@example.com,"",""
ZhangWei,APAC-Sales,"Enterprise,Startup"

```

```

GroupName,Region,Segment
EMEA-Sales,EMEA,"Enterprise,SMB,Startup"
US-Sales,US,Enterprise

```

```
US-Sales,US,SMB
US-Sales,US,Startup
APAC-Sales,APAC,"SMB,Startup"
Corporate-Reporting,"",""
APAC-Sales,APAC,"Enterprise,Startup"
```

```
UserARN,GroupARN,Region
arn:aws-cn:quicksight:us-east-1:123456789012:user/Bob,arn:aws-cn:quicksight:us-east-1:123456789012:group/group-1,APAC
arn:aws-cn:quicksight:us-east-1:123456789012:user/Sam,arn:aws-cn:quicksight:us-east-1:123456789012:group/group-2,US
```

Following is a SQL example.

```
/* for users*/
select User as UserName, Region, Segment
from tps-permissions;

/* for groups*/
select Group as GroupName, Region, Segment
from tps-permissions;
```

2. Create a dataset for the dataset rules. To make sure that you can easily find it, give it a meaningful name, for example **Permissions-Sales-Pipeline**.

Rules Dataset flagging for row-level security

Use the following procedure to appropriately flag a dataset as a rules dataset.

Rules Dataset is a flag that distinguishes permission datasets used for row-level security from regular datasets. If a permissions dataset was applied to a regular dataset before March 31, 2025, it will have a Rules Dataset flag in the **Dataset** landing page.

If a permissions dataset was not applied to a regular dataset by March 31, 2025, it will be categorized as a regular dataset. To use it as a rules dataset, duplicate the permissions dataset and flag it as a rules dataset on the console when creating the dataset. Select EDIT DATASET and under the options, choose DUPLICATE AS RULES DATASET.

To successfully duplicate it as a rules dataset, ensure the original dataset has: 1. Required user metadata or group metadata column(s) and 2. Only string type columns.

To create a new rules dataset on the console, select **NEW RULES DATASET** under the **NEW DATASET** dropdown. When creating a rules dataset programmatically, add the following parameter: [UseAs: RLS_RULES](#). This is an optional parameter that is only used to create a rules dataset. Once a dataset has been created, either through the console or programmatically, and flagged as either a rules dataset or a regular dataset, it cannot be changed.

Once datasets are flagged as rules datasets, Amazon Quick Suite will apply strict SPICE ingestion rules on them. To ensure data integrity, SPICE ingestions for rules datasets will fail if there are invalid rows or cells exceeding length limits. You must fix the ingestion issues in order to re-initiate a successful ingestion. Strict ingestion rules are only applicable to rules datasets. Regular datasets will not have dataset ingestion failures when there are skipped rows or string truncations.

Applying row-level security

Use the following procedure to apply row-level security (RLS) by using a file or query as a dataset that contains the rules for permissions.

To apply row-level security by using a file or query

1. Confirm that you have added your rules as a new dataset. If you added them, but don't see them under the list of datasets, refresh the screen.
2. On the **Data** page, choose the dataset
3. On the dataset details page that opens, for **Row-level security**, choose **Set up**.
4. On the **Set up row-level security** page that opens, choose **User-based rules**.
5. From the list of datasets that appears, choose your permissions dataset.

If your permissions dataset doesn't appear on this screen, return to your datasets, and refresh the page.

6. For **Permissions policy** choose **Grant access to dataset**. Each dataset has only one active permissions dataset. If you try to add a second permissions dataset, it overwrites the existing one.

Important

Some restrictions apply to NULL and empty string values when working with row-level security:

- If your dataset has NULL values or empty strings ("") in the restricted fields, these rows are ignored when the restrictions are applied.

- Inside the permissions dataset, NULL values and empty strings are treated the same. For more information, see the following table.
- To prevent accidentally exposing sensitive information, Amazon Quick Suite skips empty RLS rules that grant access to everyone. An *empty RLS rule* occurs when all columns of a row have no value. Quick Suite RLS treats NULL, empty strings (""), or empty comma separated strings (for example ",",) as no value.
- After skipping empty rules, other nonempty RLS rules still apply.
- If a permission dataset has only empty rules and all of them were skipped, no one will have access to any data restricted by this permission dataset.

Rules for UserName, GroupName, Region, Segment	Granted access	
AlejandroRosalez,EMEA-Sales,EMEA,"Enterprise,SMB,Startup"	Sees all EMEA Enterprise, SMB, and Startup	
sales-tps@example.com,Corporate-Reporting,"", ""	Sees all rows	
User or group has no entry	Sees no rows	
","", "", ""	Skipped; sees no rows if all other rules are empty.	
NULL,"","",NULL	Skipped; sees no rows if all other rules are empty.	

Anyone whom you shared your dashboard with can see all the data in it, unless the dataset is restricted by dataset rules.

7. Choose **Apply dataset** to save your changes. Then, on the **Save data set rules?** page, choose **Apply and activate**. Changes in permissions apply immediately to existing users.
8. (Optional) To remove permissions, first remove the dataset rules from the dataset.

Make certain that the dataset rules are removed. Then, choose the permissions dataset and choose **Remove data set**.

To overwrite permissions, choose a new permissions dataset and apply it. You can reuse the same dataset name. However, make sure to apply the new permissions in the **Permissions** screen to make these permissions active. SQL queries dynamically update, so these can be managed outside of Amazon Quick Suite. For queries, the permissions are updated when the direct query cache is automatically refreshed.

If you delete a file-based permissions dataset before you remove it from the target dataset, restricted users can't access the dataset. While the dataset is in this state, it remains marked as **RESTRICTED**. However, when you view **Permissions** for that dataset, you can see that it has no selected dataset rules.

To fix this, specify new dataset rules. Creating a dataset with the same name is not enough to fix this. You must choose the new permissions dataset on the **Permissions** screen. This restriction doesn't apply to direct SQL queries.

Using row-level security with tag-based rules to restrict access to a dataset when embedding dashboards for anonymous users

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite Administrators and Amazon Quick Suite developers

When you embed Amazon Quick Suite dashboards in your application for users who are not provisioned (registered) in Quick Suite, you can use row-level security (RLS) with tags. In this case, you use tags to specify which data your users can see in the dashboard depending on who they are.

For example, let's say you're a logistics company that has a customer-facing application for various retailers. Thousands of users from these retailers access your application to see metrics related to how their orders are getting shipped from your warehouse.

You don't want to manage thousands of users in Quick Suite, so you use anonymous embedding to embed the selected dashboards in your application that your authenticated and authorized users

can see. However, you want to make sure retailers see only data that is for their business and not for others. You can use RLS with tags to make sure your customers only see data that's relevant to them.

To do so, complete the following steps:

1. Add RLS tags to a dataset.
2. Assign values to those tags at runtime using the `GenerateEmbedUrlForAnonymousUser` API operation.

For more information about embedding dashboards for anonymous users using the `GenerateEmbedUrlForAnonymousUser` API operation, see [Embedding Amazon Quick Sight dashboards for anonymous \(unregistered\) users](#).

Before you can use RLS with tags, keep in mind the following points:

- Using RLS with tags is currently only supported for anonymous embedding, specifically for embedded dashboards that use the `GenerateEmbedUrlForAnonymousUser` API operation.
- Using RLS with tags isn't supported for embedded dashboards that use the `GenerateEmbedUrlForRegisteredUser` API operation or the old `GetDashboardEmbedUrl` API operation.
- RLS tags aren't supported with Amazon Identity and Access Management (IAM) or the Quick Suite identity type.
- When applying SPICE datasets to row-level security, each field in the dataset can contain up to 2,047 Unicode characters. Fields that contain more than this quota are truncated during ingestion. To learn more about SPICE data quotas, see [SPICE quotas for imported data](#).

Step 1: Add RLS tags to a dataset

You can add tag-based rules to a dataset in Amazon Quick Suite. Alternatively, you can call the `CreateDataSet` or `UpdateDataSet` API operation and add tag-based rules that way. For more information, see [Add RLS tags to a dataset using the API](#).

Use the following procedure to add RLS tags to a dataset in Quick Suite.

To add RLS tags to a dataset

1. From the Quick Suite start page, choose **Data** at left.

2. Choose the dataset that you want to add RLS to.
3. On the dataset details page that opens, for **Row-level security**, choose **Set up**.
4. On the **Set up row-level security** page that opens, choose **Tag-based rules**.
5. For **Column**, choose a column that you want to add tag rules to.

For example, in the case for the logistics company, the `retailer_id` column is used.

Only columns with a string data type are listed.

6. For **Tag**, enter a tag key. You can enter any tag name that you want.

For example, in the case for the logistics company, the tag key `tag_retailer_id` is used. Doing this sets row-level security based on the retailer that's accessing the application.

7. (Optional) For **Delimiter**, choose a delimiter from the list, or enter your own.

You can use delimiters to separate text strings when assigning more than one value to a tag. The value for a delimiter can be 10 characters long, at most.

8. (Optional) For **Match all**, choose the `*`, or enter your own character or characters.

This option can be any character that you want to use when you want to filter by all the values in that column in the dataset. Instead of listing the values one by one, you can use the character. If this value is specified, it can be at least one character, or at most 256 characters long.

9. Choose **Add**.

The tag rule is added to the dataset and is listed at the bottom, but it isn't applied yet. To add another tag rule to the dataset, repeat steps 5–9. To edit a tag rule, choose the pencil icon that follows the rule. To delete a tag rule, choose the delete icon that follows the rule. You can add up to 50 tags to a dataset.

10. When you're ready to apply the tag rules to the dataset, choose **Apply rules**.
11. On the **Turn on tag-based security?** page that opens, choose **Apply and activate**.

The tag-based rules are now active. On the **Set up row-level security** page, a toggle appears for you to turn tag rules on and off for the dataset.

To turn off all tag-based rules for the dataset, switch the **Tag-Based rules** toggle off, and then enter "confirm" in the text box that appears.

On the **Data** page, a lock icon appears in the dataset row to indicate that tag rules are enabled.

You can now use tag rules to set tag values at runtime, described in [Step 2: Assign values to RLS tags at runtime](#). The rules only affect Quick Suite readers when active.

 Important

After tags are assigned and enabled on the dataset, make sure to give Quick Suite authors permissions to see any of the data in the dataset when authoring a dashboard. To give Quick Suite authors permission to see data in the dataset, create a permissions file or query to use as dataset rules. For more information, see [Creating dataset rules for row-level security](#).

After you create a tag-based rule, a new **Manage rules** table appears that shows how your tag-based rules relate to each other. To make changes to the rules listed in the **Manage rules** table, choose the pencil icon that follows the rule. Then add or remove tags, and choose **Update**. To apply your updated rule to the dataset, choose **Apply**.

(Optional) Add the OR condition to RLS tags

You can also add the OR condition to your tag-based rules to further customize the way data is presented to your Quick Suite account users. When you use the OR condition with your tag-based rules, visuals in Quick Suite appear if at least one tag defined in the rule is valid.

To add the OR condition to your tag-based rules

1. In the **Manage rules** table, choose **Add OR condition**.
2. In the **Select tag** dropdown list that appears, choose the tag that you want to create an OR condition for. You can add up to 50 OR conditions to the **Manage rules** table. You can add multiple tags to a single column in a dataset, but at least one column tag needs to be included in a rule.
3. Choose **Update** to add the condition to your rule, then choose **Apply** to apply the updated rule to your dataset.

Add RLS tags to a dataset using the API

Alternatively, you can configure and enable tag-based row-level security on your dataset by calling the `CreateDataSet` or `UpdateDataSet` API operation. Use the following examples to learn how.

CreateDataSet

The following is an example for creating a dataset that uses RLS with tags. It assumes the scenario of the logistics company described previously. The tags are defined in the `row-level-permission-tag-configuration` element. The tags are defined on the columns that you want to secure the data for. For more information about this optional element, see [RowLevelPermissionTagConfiguration](#) in the *Amazon Quick Suite API Reference*.

```
create-data-set
  --aws-account-id <value>
  --data-set-id <value>
  --name <value>
  --physical-table-map <value>
  [--logical-table-map <value>]
  --import-mode <value>
  [--column-groups <value>]
  [--field-folders <value>]
  [--permissions <value>]
  [--row-level-permission-data-set <value>]
  [--column-level-permission-rules <value>]
  [--tags <value>]
  [--cli-input-json <value>]
  [--generate-cli-skeleton <value>]
  [--row-level-permission-tag-configuration
  '{
    "Status": "ENABLED",
    "TagRules":
      [
        {
          "TagKey": "tag_retailer_id",
          "ColumnName": "retailer_id",
          "TagMultiValueDelimiter": ",",
          "MatchAllValue": "*"
        },
        {
          "TagKey": "tag_role",
          "ColumnName": "role"
        }
      ],
    "TagRuleConfigurations":
      [
        tag_retailer_id
      ],
```

```
[
  tag_role
]
}'
]
```

The tags in this example are defined in the TagRules part of the element. In this example, two tags are defined based on two columns:

- The tag_retailer_id tag key is defined for the retailer_id column. In this case for the logistics company, this sets row-level security based on the retailer that's accessing the application.
- The tag_role tag key is defined for the role column. In this case for the logistics company, this sets an additional layer of row-level security based on the role of the user accessing your application from a specific retailer. An example is store_supervisor or manager.

For each tag, you can define TagMultiValueDelimiter and MatchAllValue. These are optional.

- TagMultiValueDelimiter – This option can be any string that you want to use to delimit the values when you pass them at runtime. The value can be 10 characters long, at most. In this case, a comma is used as the delimiter value.
- MatchAllValue – This option can be any character that you want to use when you want to filter by all the values in that column in the dataset. Instead of listing the values one by one, you can use the character. If specified, this value can be at least one character, or at most 256 characters long. In this case, an asterisk is used as the match all value.

While configuring the tags for dataset columns, turn them on or off using the mandatory property Status. For enabling the tag rules use the value ENABLED for this property. By turning on tag rules, you can use them to set tag values at runtime, described in [Step 2: Assign values to RLS tags at runtime](#).

The following is an example of the response definition.

```
{
  "Status": 201,
  "Arn": "arn:aws-cn:quicksight:us-west-2:11112222333:dataset/RLS-Dataset",
  "DataSetId": "RLS-Dataset",
```



```
"RequestId": "aa4f3c00-b937-4175-859a-543f250f8bb2"
}
```

UpdateDataSet

UpdateDataSet

You can use the UpdateDataSet API operation to add or update RLS tags for an existing dataset.

The following is an example of updating a dataset with RLS tags. It assumes the scenario of the logistics company described previously.

```
update-data-set
  --aws-account-id <value>
  --data-set-id <value>
  --name <value>
  --physical-table-map <value>
  [--logical-table-map <value>]
  --import-mode <value>
  [--column-groups <value>]
  [--field-folders <value>]
  [--row-level-permission-data-set <value>]
  [--column-level-permission-rules <value>]
  [--cli-input-json <value>]
  [--generate-cli-skeleton <value>]
  [--row-level-permission-tag-configuration
' {
  "Status": "ENABLED",
  "TagRules":
  [
    {
      "TagKey": "tag_retailer_id",
      "ColumnName": "retailer_id",
      "TagMultiValueDelimiter": ",",
      "MatchAllValue": "*"
    },
    {
      "TagKey": "tag_role",
      "ColumnName": "role"
    }
  ],
  "TagRuleConfigurations":
```

```
[
  tag_retailer_id
],
[
  tag_role
]
}'
]
```

The following is an example of the response definition.

```
{
  "Status": 201,
  "Arn": "arn:aws-cn:quicksight:us-west-2:11112222333:dataset/RLS-Dataset",
  "DataSetId": "RLS-Dataset",
  "RequestId": "aa4f3c00-b937-4175-859a-543f250f8bb2"
}
```

Important

After tags are assigned and enabled on the dataset, make sure to give Quick Suite authors permissions to see any of the data in the dataset when authoring a dashboard.

To give Quick Suite authors permission to see data in the dataset, create a permissions file or query to use as dataset rules. For more information, see [Creating dataset rules for row-level security](#).

For more information about the `RowLevelPermissionTagConfiguration` element, see [RowLevelPermissionTagConfiguration](#) in the *Amazon Quick Suite API Reference*.

Step 2: Assign values to RLS tags at runtime

You can use tags for RLS only for anonymous embedding. You can set values for tags using the `GenerateEmbedUrlForAnonymousUser` API operation.

The following example shows how to assign values to RLS tags that were defined in the dataset in the previous step.

```
POST /accounts/AwsAccountId/embed-url/anonymous-user
```

```

HTTP/1.1
Content-type: application/json
{
  "AwsAccountId": "string",
  "SessionLifetimeInMinutes": integer,
  "Namespace": "string", // The namespace to which the anonymous end user virtually
belongs
  "SessionTags": // Optional: Can be used for row-level security
  [
    {
      "Key": "tag_retailer_id",
      "Value": "West,Central,South"
    }
    {
      "Key": "tag_role",
      "Value": "shift_manager"
    }
  ],
  "AuthorizedResourceArns":
  [
    "string"
  ],
  "ExperienceConfiguration":
  {
    "Dashboard":
    {
      "InitialDashboardId": "string"
      // This is the initial dashboard ID the customer wants the user to land on. This
ID goes in the output URL.
    }
  }
}

```

The following is an example of the response definition.

```

HTTP/1.1 Status
Content-type: application/json

{
  "EmbedUrl": "string",
  "RequestId": "string"
}

```

RLS support without registering users in Quick Suite is supported only in the `GenerateEmbedUrlForAnonymousUser` API operation. In this operation, under `SessionTags`, you can define the values for the tags associated with the dataset columns.

In this case, the following assignments are defined:

- Values `West`, `Central`, and `South` are assigned to the `tag_retailer_id` tag at runtime. A comma is used for the delimiter, which was defined in `TagMultipleValueDelimiter` in the dataset. To use call values in the column, you can set the value to `*`, which was defined as the `MatchAllValue` when creating the tag.
- The value `shift_manager` is assigned to the `tag_role` tag.

The user using the generated URL can view only the rows having the `shift_manager` value in the `role` column. That user can view only the value `West`, `Central`, or `South` in the `retailer_id` column.

For more information about embedding dashboards for anonymous users using the `GenerateEmbedUrlForAnonymousUser` API operation, see [Embedding Amazon Quick Sight dashboards for anonymous \(unregistered\) users](#), or [GenerateEmbedUrlForAnonymousUser](#) in the *Amazon Quick Suite API Reference*

Using column-level security to restrict access to a dataset

In the Enterprise edition of Quick Suite, you can restrict access to a dataset by configuring column-level security (CLS) on it. A dataset or analysis with CLS enabled has the restricted

symbol next to it. By default, all users and groups have access to the data. By using CLS, you can manage access to specific columns in your dataset.

If you use an analysis or dashboard that contains datasets with CLS restrictions that you don't have access to, you can't create, view, or edit visuals that use the restricted fields. For most visual types, if a visual has restricted columns that you don't have access to, you can't see the visual in your analysis or dashboard.

Tables and pivot tables behave differently. If a table or pivot table uses restricted columns in the **Rows** or **Columns** field wells, and you don't have access to these restricted columns, you can't see the visual in an analysis or dashboard. If a table or pivot table has restricted columns in the **Values** field well, you can see the table in an analysis or dashboard with only the values that you have access to. The values for restricted columns show as Not Authorized.

To enable column-level security on an analysis or dashboard, you need administrator access.

To create a new analysis with CLS

1. On the Quick Suite start page, choose the **Analyses** tab.
2. At upper right, choose **New analysis**.
3. Choose a dataset, and choose **Column-level security**.
4. Select the columns that you want to restrict, and then choose **Next**. By default, all groups and users have access to all columns.
5. Choose who can access each column, and then choose **Apply** to save your changes.

To use an existing analysis for CLS

1. On the Quick Suite start page, choose the **Data** tab.
2. On the Data page, open your dataset
3. On the dataset details page that opens, for **Column-level security**, choose **Set up**.
4. Select the columns that you want to restrict, and then choose **Next**. By default, all groups and users have access to all columns.
5. Choose who can access each column, and then choose **Apply** to save your changes.

To create a dashboard with CLS

1. On the Quick Suite navigation pane, choose the **Analyses** tab.
2. Choose the analysis that you want to create a dashboard of.
3. At upper right, choose **Publish**.
4. Choose one of the following:
 - To create a new dashboard, choose **Publish new dashboard as** and enter a name for the new dashboard.
 - To replace an existing dashboard, choose **Replace an existing dashboard** and choose the dashboard from the list.

Additionally, you can choose **Advanced publish options**. For more information, see [Publishing dashboards](#).

5. Choose **Publish dashboard**.
6. (Optional) Do one of the following:
 - To publish a dashboard without sharing, choose **x** at the upper right of the **Share dashboard with users** screen when it appears. You can share the dashboard later by choosing **Share** from the application bar.
 - To share the dashboard, follow the procedure in [Sharing Amazon Quick Sight dashboards](#).

Running queries as an IAM role in Amazon Quick Suite

You can enhance data security by using fine-grained access policies rather than broader permissions for data sources connected to Amazon Athena, Amazon Redshift or Amazon S3. You start by creating an Amazon Identity and Access Management (IAM) role with permissions to be activated when a person or an API starts a query. Then, an Quick Suite administrator or a developer assigns the IAM Role to an Athena or Amazon S3 data source. With the role in place, any person or API that runs the query has the exact permissions necessary to run the query.

Here are some things to consider before you commit to implementing run-as roles to enhance data security:

- Articulate how the additional security works to your advantage.
- Work with your Quick Suite administrator to learn if adding roles to data sources helps you to better meet your security goals or requirements.
- Ask if this type of security, for the number of data sources and people and applications involved, can be feasibly documented and maintained by your team? If not, then who will undertake that part of the work?
- In a structured organization, locate stakeholders in parallel teams in Operations, Development, and IT Support. Ask for their experience, advice, and willingness to support your plan.
- Before you launch your project, consider doing a proof of concept that involves the people who need access to the data.

The following rules apply to using run-as roles with Athena, Amazon Redshift, and Amazon S3:

- Each data source can have only one associated RoleArn. Consumers of the data source, who typically access datasets and visuals, can generate many different types of queries. The role places boundaries on which queries work and which don't work.

- The ARN must correspond to an IAM role in the same Amazon Web Services account as the Quick Suite instance that uses it.
- The IAM role must have a trust relationship allowing Quick Suite to assume the role.
- The identity that calls Quick Suite's APIs must have permission to pass the role before they can update the `RoleArn` property. You only need to pass the role when creating or updating the role ARN. The permissions aren't re-evaluated later on. Similarly, the permission isn't required when the role ARN is omitted.
- When the role ARN is omitted, the Athena or Amazon S3 data source uses the account-wide role and scope-down policies.
- When the role ARN is present, the account-wide role and any scope-down policies are both ignored. For Athena data sources, Lake Formation permissions are not ignored.
- For Amazon S3 data sources, both the manifest file and the data specified by the manifest file must be accessible using the IAM role.
- The ARN string needs to match an existing IAM role in the Amazon Web Services account and Amazon Web Services Region where the data is located and queried.

When Quick Suite connects to another service in Amazon, it uses an IAM role. By default, this less granular version of the role is created by Quick Suite for each service it uses, and the role is managed by Amazon Web Services account administrators. When you add an IAM role ARN with a custom permissions policy, you override the broader role for your data sources that need extra protection. For more information about policies, see [Create a customer managed policy](#) in the IAM User Guide.

Run queries with Athena data sources

Use the API to attach the ARN to the Athena data source. To do so, add the role ARN in the [RoleArn](#) property of [AthenaParameters](#). For verification, you can see the role ARN on the **Edit Athena data source** dialog box. However, **Role ARN** is a read-only field.

To get started, you need a custom IAM role, which we demonstrate in the following example.

Keep in mind that the following code example is for learning purposes only. Use this example in a temporary development and testing environment only, and not in a production environment. The policy in this example doesn't secure any specific resource, which must be in a deployable policy. Also, even for development, you need to add your own Amazon account information.

The following commands create a simple new role and attach a few policies that grant permissions to Quick Suite.

```
aws iam create-role \  
  --role-name TestAthenaRoleForQuickSight \  
  --description "Test Athena Role For QuickSight" \  
  --assume-role-policy-document '{  
    "Version": "2012-10-17",  
    "Statement": [  
      {  
        "Effect": "Allow",  
        "Principal": {  
          "Service": "quicksight.amazonaws.com"  
        },  
        "Action": "sts:AssumeRole"  
      }  
    ]  
  }'
```

After you've identified or created an IAM role to use with each data source, attach the policies by using the `attach-role-policy`.

```
aws iam attach-role-policy \  
  --role-name TestAthenaRoleForQuickSight \  
  --policy-arn arn:aws-cn:iam::222222222222:policy/service-role/  
AWSQuickSightS3Policy1  
  
aws iam attach-role-policy \  
  --role-name TestAthenaRoleForQuickSight \  
  --policy-arn arn:aws-cn:iam::aws:policy/service-role/AWSQuicksightAthenaAccess1  
  
aws iam attach-role-policy \  
  --role-name TestAthenaRoleForQuickSight \  
  --policy-arn arn:aws-cn:iam::aws:policy/AmazonS3Access1
```

After you verify your permissions, you can use the role in Quick Suite data sources by creating a new role or updating an existing role. When using these commands, update the Amazon Web Services account ID and Amazon Web Services Region to match your own.

Remember, these example code snippets are not for production environments. Amazon strongly recommends that you identify and use a set of least privilege policies for your production cases.


```
aws quicksight create-data-source
  --aws-account-id 222222222222 \
  --region us-east-1 \
  --data-source-id "athena-with-custom-role" \
  --cli-input-json '{
    "Name": "Athena with a custom Role",
    "Type": "ATHENA",
    "data sourceParameters": {
      "AthenaParameters": {
        "RoleArn": "arn:aws-cn:iam::222222222222:role/
TestAthenaRoleForQuickSight"
      }
    }
  }'
```

Run queries with Amazon Redshift data sources

Connect your Amazon Redshift data with the run-as role to enhance your data security with fine-grained access policies. You can create a run-as role for Amazon Redshift data sources that use a public network or a VPC connection. You specify the connection type that you want to use in the **Edit Amazon Redshift data source** dialog box. The run-as role is not supported for Amazon Redshift Serverless data sources.

To get started, you need a custom IAM role, which we demonstrate in the following example. The following commands create a sample new role and attach policies that grant permissions to Quick Suite.

```
aws iam create-role \
  --role-name TestRedshiftRoleForQuickSight \
  --description "Test Redshift Role For QuickSight" \
  --assume-role-policy-document '{
    "Version": "2012-10-17",
    "Statement": [
      {
        "Effect": "Allow",
        "Principal": {
          "Service": "quicksight.amazonaws.com"
        },
        "Action": "sts:AssumeRole"
      }
    ]
  }
```

```
}'
```

After you identify or create an IAM role to use with each data source, attach the policies with an `attach-role-policy`. If the `redshift:GetClusterCredentialsWithIAM` permission is attached to the role that you want to use, the values for `DatabaseUser` and `DatabaseGroups` are optional.

```
aws iam attach-role-policy \
--role-name TestRedshiftRoleForQuickSight \
--policy-arn arn:aws-cn:iam:111122223333:policy/service-role/
AWSQuickSightRedshiftPolicy

aws iam create-policy --policy-name RedshiftGetClusterCredentialsPolicy1 \
--policy-document file://redshift-get-cluster-credentials-policy.json

aws iam attach-role-policy \
--role-name TestRedshiftRoleForQuickSight \
--policy-arn arn:aws-cn:iam:111122223333:policy/RedshiftGetClusterCredentialsPolicy1
// redshift-get-cluster-credentials-policy.json
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "RedshiftGetClusterCredentialsPolicy",
      "Effect": "Allow",
      "Action": [
        "redshift:GetClusterCredentials"
      ],
      "Resource": [
        "*"
      ]
    }
  ]
}
```

The example above creates a data source that uses the `RoleARN`, `DatabaseUser`, and `DatabaseGroups` IAM parameters. If you want to establish the connection only through the IAM `RoleARN` parameter, attach the `redshift:GetClusterCredentialsWithIAM` permission to your role, shown in the example below.

```
aws iam attach-role-policy \
--role-name TestRedshiftRoleForQuickSight \
--policy-arn arn:aws:iam:111122223333:policy/RedshiftGetClusterCredentialsPolicy1 //
redshift-get-cluster-credentials-policy.json {
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "RedshiftGetClusterCredentialsPolicy",
      "Effect": "Allow",
      "Action": [ "redshift:GetClusterCredentialsWithIAM" ],
      "Resource": [ "*" ]
    }
  ]
}
```

After you verify your permissions, you can use the role in Quick Suite data sources by creating a new role or updating an existing role. When using these commands, update the Amazon account ID and Amazon Region to match your own.

```
aws quicksight create-data-source \
--region us-west-2 \
--endpoint https://quicksight.us-west-2.quicksight.aws.com/ \
--cli-input-json file://redshift-data-source-iam.json \
redshift-data-source-iam.json is shown as below
{
  "AwsAccountId": "AWSACCOUNTID",
  "DataSourceId": "DATSOURCEID",
  "Name": "Test redshift demo iam",
  "Type": "REDSHIFT",
  "DataSourceParameters": {
    "RedshiftParameters": {
      "Database": "integ",
      "Host": "redshiftdemocluster.us-west-2.redshift.amazonaws.com",
      "Port": 8192,
      "ClusterId": "redshiftdemocluster",
      "IAMParameters": {
        "RoleArn": "arn:aws-
cn:iam::222222222222:role/TestRedshiftRoleForQuickSight",
        "DatabaseUser": "user",
        "DatabaseGroups": ["admin_group", "guest_group", "guest_group_1"]
      }
    }
  }
```

```

    },
    "Permissions": [
      {
        "Principal": "arn:aws-cn:quicksight:us-east-1:AWSACCOUNTID:user/default/demoname",
        "Actions": [
          "quicksight:DescribeDataSource",
          "quicksight:DescribeDataSourcePermissions",
          "quicksight:PassDataSource",
          "quicksight:UpdateDataSource",
          "quicksight>DeleteDataSource",
          "quicksight:UpdateDataSourcePermissions"
        ]
      }
    ]
  }
}

```

If your data source uses the VPC connection type, use the following VPC configuration.

```

{
  "AwsAccountId": "AWSACCOUNTID",
  "DataSourceId": "DATSOURCEID",
  "Name": "Test redshift demo iam vpc",
  "Type": "REDSHIFT",
  "DataSourceParameters": {
    "RedshiftParameters": {
      "Database": "mydb",
      "Host": "vpcdemo.us-west-2.redshift.amazonaws.com",
      "Port": 8192,
      "ClusterId": "vpcdemo",
      "IAMParameters": {
        "RoleArn": "arn:aws-cn:iam::222222222222:role/TestRedshiftRoleForQuickSight",
        "DatabaseUser": "user",
        "AutoCreateDatabaseUser": true
      }
    }
  },
  "VpcConnectionProperties": {
    "VpcConnectionArn": "arn:aws-cn:quicksight:us-west-2:222222222222:vpcConnection/VPC Name"
  },
  "Permissions": [

```

```
{
  "Principal": "arn:aws-cn:quicksight:us-east-1:222222222222:user/
default/demoname",
  "Actions": [
    "quicksight:DescribeDataSource",
    "quicksight:DescribeDataSourcePermissions",
    "quicksight:PassDataSource",
    "quicksight:UpdateDataSource",
    "quicksight>DeleteDataSource",
    "quicksight:UpdateDataSourcePermissions"
  ]
}
```

If your data source uses the `redshift:GetClusterCredentialsWithIAM` permission and doesn't use the `DatabaseUser` or `DatabaseGroups` parameters, grant the role access to some or all tables in the schema. To see if a role has been granted `SELECT` permissions to a specific table, input the following command into the Amazon Redshift Query Editor.

```
SELECT
u.username,
t.schemaname||'.'||t.tablename,
has_table_privilege(u.username,t.tablename,'select') AS user_has_select_permission
FROM
pg_user u
CROSS JOIN
pg_tables t
WHERE
u.username = 'IAMR:RoleName'
AND t.tablename = tableName
```

For more information about the `SELECT` action in the Amazon Redshift Query Editor, see [SELECT](#).

To grant `SELECT` permissions to the role, input the following command in the Amazon Redshift Query Editor.

```
GRANT SELECT ON { [ TABLE ] table_name [, ...] | ALL TABLES IN SCHEMA
schema_name [, ...] } TO "IAMR:RoleName";
```

For more information about the `GRANT` action in the Amazon Redshift Query Editor, see [GRANT](#).

Run queries with Amazon S3 data sources

Amazon S3 data sources contain a manifest file that Quick Suite uses to find and parse your data. You can upload a JSON manifest file through the Quick Suite console, or you can provide a URL that points to a JSON file in an S3 bucket. If you choose to provide a URL, Quick Suite must be granted permission to access the file in Amazon S3. Use the Quick Suite administration console to control access to the manifest file and the data that it references.

With the **RoleArn** property, you can grant access to the manifest file and the data that it references through a custom IAM role that overrides the account-wide role. Use the API to attach the ARN to the manifest file of the Amazon S3 data source. To do so, include the role ARN in the [RoleArn](#) property of [S3Parameters](#). For verification, you can see the role ARN in the **Edit S3 data source** dialog box. However, **Role ARN** is a read-only field, as shown in the following screenshot.

To get started, create an Amazon S3 manifest file. Then, you can either upload it to Amazon Quick Suite when you create a new Amazon S3 dataset or place the file into the Amazon S3 bucket that contains your data files. View the following example to see what a manifest file might look like:

```
{
  "fileLocations": [
    {
      "URIPrefixes": [
        "s3://quicksightUser-run-as-role/data/"
      ]
    }
  ],
  "globalUploadSettings": {
    "format": "CSV",
    "delimiter": ",",
    "textqualifier": "'",
    "containsHeader": "true"
  }
}
```

For instructions on how to create a manifest file, see [Supported formats for Amazon S3 manifest files](#).

After you have created a manifest file and added it to your Amazon S3 bucket or uploaded it to Quick Suite, create or update an existing role in IAM that grants `s3:GetObject` access. The following example illustrates how to update an existing IAM role with the Amazon API:

```
aws iam put-role-policy \
  --role-name QuickSightAccessToS3RunAsRoleBucket \
  --policy-name GrantS3RunAsRoleAccess \
  --policy-document '{
    "Version": "2012-10-17",
    "Statement": [
      {
        "Effect": "Allow",
        "Action": "s3:ListBucket",
        "Resource": "arn:aws-cn:s3::s3-bucket-name"
      },
      {
        "Effect": "Allow",
        "Action": "s3:GetObject",
        "Resource": "arn:aws-cn:s3::s3-bucket-name/manifest.json"
      },
      {
        "Effect": "Allow",
        "Action": "s3:GetObject",
        "Resource": "arn:aws-cn:s3::s3-bucket-name/*"
      }
    ]
  }'
```

After your policy grants `s3:GetObject` access, you can begin creating data sources that apply the updated `put-role-policy` to the Amazon S3 data source's manifest file.

```
aws quicksight create-data-source --aws-account-id 111222333444 --region us-west-2 --
endpoint https://quicksight.us-west-2.quicksight.aws.com/ \
  --data-source-id "s3-run-as-role-demo-source" \
  --cli-input-json '{
    "Name": "S3 with a custom Role",
    "Type": "S3",
    "DataSourceParameters": {
      "S3Parameters": {
        "RoleArn": "arn:aws-cn:iam::111222333444:role/
QuickSightAccessRunAsRoleBucket",
        "ManifestFileLocation": {
          "Bucket": "s3-bucket-name",
          "Key": "manifest.json"
        }
      }
    }
  }'
```

```
}'
```

After you verify your permissions, you can use the role in Quick Suite data sources, either by creating a new role or updating an existing role. When using these commands, be sure to update the Amazon Web Services account ID and Amazon Web Services Region to match your own.

Deleting datasets

Important

Currently, deleting a dataset is irreversible and can cause irreversible loss of work. Deletes don't cascade to delete dependent objects. Instead, dependent objects stop working, even if you replace the deleted dataset with an identical dataset.

Before you delete a dataset, we strongly recommend that you first point each dependent analysis or dashboard to a new dataset.

Currently, when you delete a dataset while dependent visuals still exist, the analyses and dashboards that contain those visuals have no way to assimilate new metadata. They remain visible, but they can't function. They can't be repaired by adding an identical dataset.

This is because datasets include metadata that is integral to the analyses and dashboards that depend on that dataset. This metadata is uniquely generated for each dataset. Although the Quick Sight engine can read the metadata, it isn't readable by humans (for example, it doesn't contain field names). So, an exact replica of the dataset has different metadata. Each dataset's metadata is unique, even for multiple datasets that share the same name and the same fields.

To delete a dataset

1. Make sure that the dataset isn't being used by any analysis or dashboard that someone wants to keep using.

On the **Data** page, choose the dataset that you no longer need. Then choose **Delete Dataset** at upper-right.

2. If you receive a warning if this dataset is in use, track down all dependent analyses and dashboards and point them at a different dataset. If this isn't feasible, try one or more of these best practices instead of deleting it:
 - Rename the dataset, so that the dataset is clearly deprecated.

- Filter the data, so that the dataset has no rows.
- Remove everyone else's access to the dataset.

We recommend that you use whatever means you can to inform owners of dependent objects that this dataset is being deprecated. Also, make sure that you provide sufficient time for them to take action.

3. After you make sure that there are no dependent objects that will stop functioning after the dataset is deleted, choose the dataset and choose **Delete Data Set**. Confirm your choice, or choose **Cancel**.

Important

Currently, deleting a dataset is irreversible and can cause irreversible loss of work. Deletes don't cascade to delete dependent objects. Instead, dependent objects stop working, even if you replace the deleted dataset with an identical dataset.

Adding a dataset to an analysis

After you have created an analysis, you can add more datasets to the analysis. Then, you can use them to create more visuals.

From within the analysis, you can open any dataset for editing, for example to add or remove fields, or perform other data preparation. You can also remove or replace data sets.

The currently selected dataset displays at the top of the **Data** pane. This is the dataset that is used by the currently selected visual. Each visual can use only one dataset. Choosing a different visual changes the selected dataset to the one used by that visual.

To change the selected dataset manually, choose the dataset list at the top of the **Data** pane and then choose a different dataset. This deselects the currently selected visual if it doesn't use this dataset. Then, choose a visual that uses the selected dataset. Or choose **Add** in the **Visuals** pane to create a new visual using the selected dataset.

If you choose **Suggested** on the tool bar to see suggested visuals, you'll see visuals based on the currently selected dataset.

Only filters for the currently selected dataset are shown in the **Filter** pane, and you can only create filters on the currently selected dataset.

Topics

- [Replacing datasets](#)
- [Remove a dataset from an analysis](#)

Use the following procedure to add a dataset to an analysis or edit a dataset used by an analysis.

To add a dataset to an analysis

1. On the analysis page, navigate to the **Data** pane and expand the **Dataset** dropdown.
2. Choose **Add a new dataset** to add a dataset. Or, choose **Manage datasets** to edit a dataset. For more information about editing a dataset, see [Editing datasets](#).
3. A list of your datasets appears. Choose a dataset and then choose **Select**. To cancel, choose **Cancel**.

Replacing datasets

In an analysis, you can add, edit, replace, or remove datasets. Use this section to learn how to replace your dataset.

When you replace a dataset, the new dataset should have similar columns, if you expect the visual to work the way you designed it. Replacing the dataset also clears the undo and redo history for the analysis. This means you can't use the undo and redo buttons on the application bar to navigate your changes. So, when you decide to change the dataset, your analysis design should be somewhat stable—not in the middle of an editing phase.

To replace a dataset

1. On the analysis page, navigate to the **Data** pane and expand the **Dataset** dropdown.
2. Choose **Manage datasets**.
3. Choose the ellipsis (three dots) next to the dataset that you want to replace, and then choose **Replace**.
4. In the **Select replacement dataset** page, choose a dataset from the list, and then choose **Select**.

Note

Replacing a dataset clears the undo and redo history for this analysis.

The dataset is replaced with the new one. The field list and visuals are updated with the new dataset.

At this point, you can choose to add a new dataset, edit the new dataset, or replace it with a different one. Choose **Close** to exit.

If your new dataset doesn't match

In some cases, the selected replacement dataset doesn't contain all of the fields and hierarchies used by the visuals, filters, parameters, and calculated fields in your analysis. If so, you receive a warning from Quick Sight that shows a list of mismatched or missing columns.

If this happens, you can update the field mapping between the two datasets.

To update the field mapping

1. In the **Mismatch in replacement dataset** page, choose **Update field mapping**.
2. In the **Update field mapping** page, choose the drop-down menu for the field(s) you want to map and choose a field from the list to map it to.

If the field is missing from the new dataset, choose **Ignore this field**.

3. Choose **Confirm** to confirm your updates.
4. Choose **Close** to close the page and return to your analysis.

The dataset is replaced with the new one. The fields list and visuals are updated with the new dataset.

Any visuals that were using a field that's now missing from the new dataset update to blank. You can readd fields to the visual or remove the visual from your analysis.

If you change your mind after replacing the dataset, you can still recover. Let's say you replace the dataset and then find that it's too difficult to change your analysis to match the new dataset. You can undo any changes you made to your analysis. You can then replace the new dataset with the original one, or with a dataset that more closely matches the requirements of the analysis.

Remove a dataset from an analysis

Use the following procedure to delete a dataset from an analysis.

To delete a dataset from an analysis

1. On the analysis page, navigate to the **Data** pane and expand the **Dataset** dropdown.
2. Choose **Manage datasets**.
3. Choose the ellipsis (three dots) next to the dataset that you want to replace, and then choose **Remove**. You can't delete a dataset if it's the only one in the analysis.
4. Choose **Close** to close the dialog box.

Working with data sources in Amazon Quick Sight

Use a data source to access an external data store. Amazon S3 data sources save the manifest file information. In contrast, Salesforce and database data sources save connection information like credentials. In such cases, you can easily create multiple datasets from the data store without having to re-enter information. Connection information isn't saved for text or Microsoft Excel files.

Topics

- [Creating a data source](#)
- [Editing a data source](#)
- [Deleting a data source](#)

Creating a data source

Intended audience: Amazon Quick Suite authors

As an analysis author in Amazon Quick Suite, you don't need to understand anything about the infrastructure that you use to connect to your data. You set up a new data source only once.

After a data source is set up, you can access it from its tile in the Quick Suite console. You can use it to create one or more datasets. After a dataset is set up, you can also access the dataset from its tile. By abstracting away the technical details, Amazon Quick Sight simplifies data connections.

Note

You don't need to store connection settings for files that you plan to upload manually. For more information about file uploads, see [Creating datasets](#).

Before you begin adding a new data-source connection profile to Amazon Quick Suite, first collect the information that you need to connect to the data source. In some cases, you might plan to copy and paste settings from a file. If so, make sure that the file doesn't contain formatting characters (list bullets or numbers) or blank space characters (spaces, tabs). Also make sure that the file doesn't contain nontext "gremlin" characters such as non-ASCII, null (ASCII 0), and control characters.

The following list includes the information to collect the most commonly used settings:

- The data source to connect to.

Make sure that you know which source that you need to connect to for reporting. This source might be different than the source that stores, processes, or provides access to the data.

For example, let's say that you're a new analyst in a large company. You want to analyze data from your ordering system, which you know uses Oracle. However, you can't directly query the online transaction processing (OLTP) data. A subset of data is extracted and stored in a bucket on Amazon S3, but you don't have access to that either. Your new co-workers explain that they use Amazon Glue crawlers to read the files and Amazon Lake Formation to access them. With more research, you learn that you need to use an Amazon Athena query as your data source in Amazon Quick Sight. The point here is that it isn't always obvious which type of data source to choose.

- A descriptive name for the new data source tile.

Each new data source connection needs a unique and descriptive name. This name displays on the Amazon Quick Sight list of existing data sources, which is at the bottom of the **Create a Data Set** screen. Use a name that makes it easy to distinguish your data sources from other similar data sources. Your new Amazon Quick Sight data source profile displays both the database software logo and the custom name that you assign.

- The name of the server or instance to connect to.

A unique name or other identifier identifies the server connector of the data source on your network. The descriptors vary depending on which one you're connecting to, but it's usually one or more of the following:

- Hostname
- IP address
- Cluster ID
- Instance ID
- Connector
- Site-based URL
- The name of the collection of data that you want to use.

The descriptor varies depending on the data source, but it's usually one of the following:

- Database
- Warehouse
- S3 bucket
- Catalog
- Schema

In some cases, you might need to include a manifest file or a query.

- The user name that you want Amazon Quick Sight to use.

Every time Amazon Quick Sight connects using this data source profile (tile), it uses the user name from the connection settings. In some cases, this might be your personal login. But if you're going to share this with other people, ask the system administrator about creating credentials to use for Amazon Quick Sight connections.

- What type of connection to use. You can choose a public network or a VPC connection. If you have more than one VPC connection available, identify which one to use to reach your source of data.
- Additional settings, such as Secure Sockets Layer (SSL) or API tokens, are required by some data sources.

After you save the connection settings as a data source profile, you can create a dataset by selecting its tile. The connections are stored as data source connection profiles in Amazon Quick Sight.

To view your existing connection profiles, open the Quick Suite start page, choose **Data**, choose **Create**, and then choose **New Dataset**.

For a list of supported data source connections and examples, see [Connect to your data with integrations and datasets](#).

After you create a data source in Quick Sight, you can [create a dataset](#) in Quick Sight that contains data from the connected data source. You can also [update data source connection](#) information at any time.

Editing a data source

You can edit an existing database data source to update the connection information, such as the server name or the user credentials. You can also edit an existing Amazon Athena data source to update the data source name. You can't edit Amazon S3 or Salesforce data sources.

Editing a database data source

Use the following procedure to edit a database data source.

1. From the Quick Suite start page, choose **Data** at left. Choose **Create** and then choose **New dataset**.
2. Choose a database data source.
3. Choose **Edit Data Source**.
4. Modify the data source information:
 - If you are editing an autodiscovered database data source, you can modify any of the following settings:
 - For **Data source name**, enter a name for the data source.
 - For **Instance ID**, choose the name of the instance or cluster that you want to connect to from the list provided.
 - **Database name** shows the default database for the **Instance ID** cluster or instance. If you want to use a different database on that cluster or instance, enter its name.
 - For **UserName**, enter the user name of a user account that has permissions to do the following:

- Access the target database.
 - Read (perform a SELECT statement on) any tables in that database that you want to use.
 - For **Password**, enter the password for the account that you entered.
- If you are editing an external database data source, you can modify any of the following settings:
 - For **Data source name**, enter a name for the data source.
 - For **Database server**, enter one of the following values:
 - For an Amazon Redshift cluster, enter the endpoint of the cluster without the port number. For example, if the endpoint value is `clustername.1234abcd.us-west-2.redshift.amazonaws.com:1234`, then enter `clustername.1234abcd.us-west-2.redshift.amazonaws.com`. You can get the endpoint value from the **Endpoint** field on the cluster detail page on the Amazon Redshift console.
 - For an Amazon EC2 instance of PostgreSQL, MySQL, or SQL Server, enter the public DNS address. You can get the public DNS value from the **Public DNS** field on the instance detail pane in the EC2 console.
 - For a non–Amazon EC2 instance of PostgreSQL, MySQL, or SQL Server, enter the hostname or public IP address of the database server.
 - For **Port**, enter the port that the cluster or instance uses for connections.
 - For **Database name**, enter the name of the database that you want to use.
 - For **UserName**, enter the user name of a user account that has permissions to do the following:
 - Access the target database.
 - Read (perform a SELECT statement on) any tables in that database that you want to use.
 - For **Password**, enter the password for the account that you entered.
5. Choose **Validate connection**.
 6. If the connection validates, choose **Update data source**. If not, correct the connection information and try validating again.
 7. If you want to create a new dataset using the updated data source, proceed with the instructions at [Creating a dataset from a database](#). Otherwise, close the **Choose your table** dialog box.

Editing an Athena data source

Use the following procedure to edit an Athena data source.

1. From the Quick Suite start page, choose **Data** at left. Choose **Create** and then choose **New dataset**.
2. Choose an Athena data source.
3. Choose **Edit Data Source**.
4. For **Data source name**, enter a new name.
5. The **Manage data source sharing** screen appears. On the **Users** tab, locate the user that you want to remove.
6. If you want to create a new dataset using the updated data source, proceed with the instructions at [Creating a dataset using Amazon Athena data](#). Otherwise, close the **Choose your table** dialog box.

Deleting a data source

You can delete a data source if you no longer need it. Deleting a query-based database data source makes any associated datasets unusable. Deleting an Amazon S3, Salesforce, or SPICE-based database data source doesn't affect your ability to use any associated datasets. This is because the data is stored in [SPICE](#). However, you can no longer refresh those datasets.

To delete a data source

1. Choose the data source that you want to delete.
2. Choose **Delete**.

Refreshing data in Amazon Quick Sight

When refreshing data, Amazon Quick Sight handles datasets differently depending on the connection properties and the storage location of the data.

If Quick Sight connects to the data store by using a direct query, the data automatically refreshes when you open an associated dataset, analysis, or dashboard. Filter controls are refreshed automatically every 24 hours.

To refresh SPICE datasets, Quick Sight must independently authenticate using stored credentials to connect to the data. Quick Sight can't refresh manually uploaded data—even from S3 buckets, even though it's stored in SPICE—because Quick Sight doesn't store its connection and location metadata. If you want to automatically refresh data that's stored in an S3 bucket, create a dataset by using the **S3** data source card.

For files that you manually uploaded to SPICE, you refresh these manually by importing the file again. If you want to reuse the name of the original dataset for the new file, first rename or delete the original dataset. Then give the preferred name to the new dataset. Also, check that the field names are the same name and data type. Open your analysis, and replace the original dataset with the new dataset. For more information, see [Replacing datasets](#).

You can refresh your [SPICE](#) datasets at any time. Refreshing imports the data into SPICE again, so the data includes any changes since the last import.

For Amazon Quick Sight Standard Edition, you can do a full refresh of your SPICE data at any time. For Amazon Quick Sight Enterprise Edition, you can do a full refresh or an incremental refresh (SQL-based data sources only) at any time.

 Note

If your dataset uses CustomSQL, refreshing incrementally might not benefit you. If the SQL query is complex, your database may not be able to optimize the filter with the look-back window. This can cause the query that pulls in the data to take longer than a full refresh. We recommend that you try reducing query execution time by refactoring the custom SQL. Note that results might vary depending on the type of optimization you make.

You can refresh SPICE data by using any of the following approaches:

- You can use the options on **Datasets** page.
- You can refresh a dataset while editing a dataset.
- You can schedule refreshes in the dataset settings.
- You can use the [CreateIngestion](#) API operation to refresh the data.

When you create or edit a SPICE dataset, you can enable email notifications about data loading status. This option notifies the owners of the dataset if the data fails to load or refresh. To turn on notifications, select the **Email owners when a refresh fails** option that appears on the **Finish data**

set creation screen. This option isn't available for datasets that you create by using **Upload a File** on the datasets page.

In the following topics, you can find an explanation of different approaches to refreshing and working with SPICE data.

Topics

- [Importing data into SPICE](#)
- [Refreshing SPICE data](#)
- [Using SPICE data in an analysis](#)
- [View SPICE ingestion history](#)
- [Troubleshooting skipped row errors](#)
- [SPICE ingestion error codes](#)
- [Updating files in a dataset](#)

Importing data into SPICE

When you import data into a dataset rather than using a direct SQL query, it becomes *SPICE data* because of how it's stored. *SPICE (Super-fast, Parallel, In-memory Calculation Engine)* is the robust in-memory engine that Amazon Quick Sight uses. It's engineered to rapidly perform advanced calculations and serve data. In Enterprise edition, data stored in SPICE is encrypted at rest.

When you create or edit a dataset, you choose to use either SPICE or a direct query, unless the dataset contains uploaded files. Importing (also called *ingesting*) your data into SPICE can save time and money:

- Your analytical queries process faster.
- You don't need to wait for a direct query to process.
- Data stored in SPICE can be reused multiple times without incurring additional costs. If you use a data source that charges per query, you're charged for querying the data when you first create the dataset and later when you refresh the dataset.

SPICE capacity is allocated separately for each Amazon Web Services Region. Default SPICE capacity is automatically allocated to your home Amazon Web Services Region. For each Amazon account, SPICE capacity is shared by all the people using Quick Sight in a single Amazon Web

Services Region. The other Amazon Web Services Regions have no SPICE capacity unless you choose to purchase some. Quick Sight administrators can view how much [SPICE](#) capacity you have in each Amazon Web Services Region and how much of it is currently in use. A Quick Sight administrator can purchase more SPICE capacity or release unused SPICE capacity as needed. For more information, see [Configure SPICE memory capacity](#).

Topics

- [Estimating the size of SPICE datasets](#)

Estimating the size of SPICE datasets

The size of a dataset in SPICE relative to your Quick Suite account's SPICE capacity is called *logical size*. A dataset's logical size isn't the same as the size of the dataset's source file or table. The computation of a dataset's logical size occurs after all the data type transformations and calculated columns are defined during data preparation. These fields are materialized in SPICE in a way that enhances query performance. Any changes you make in an analysis have no effect on the logical size of the data in SPICE. Only changes that are saved in the dataset apply to SPICE capacity.

The logical size of a SPICE dataset depends on the data types of the dataset fields and the number of rows in the dataset. The three types of SPICE data are decimals, dates, and strings. You can transform a field's data type during the data preparation phase to fit your data visualization needs. For example, the file you want to import might contain all strings (text). But for these to be used in a meaningful way in an analysis, you prepare the data by changing the data types to their proper form. Fields containing prices can be changed from strings to decimals, and fields containing dates can be changed from strings to dates. You can also create calculated fields and exclude fields that you don't need from the source table. When you are finished preparing your dataset and all transformations are complete, you can estimate the logical size of the final schema.

Note

Geospatial data types use metadata to interpret the physical data type. Latitude and longitude are numeric. All other geospatial categories are strings.

In the formula below, decimals and dates are calculated as 8 bytes per cell with 4 extra bytes for auxillary. Strings are calculated based on the text's length in UTF-8 encoding plus 24 bytes for auxillary. String data types require more space because of the extra indexing required by SPICE to provide high query performance.

```
Logical dataset size in bytes =  
(Number of Numeric cells * (12 bytes per cell))  
+ (Number of Date cells * (12 bytes per cell))  
+ SUM ((24 bytes + UTF-8 encoded length) per Text cell)
```

The formula above should only be used to estimate the size of a single dataset in SPICE. The SPICE capacity usage is the total size of all datasets in an account in a specific region. Quick Sight does not recommend that you use this formula to estimate the total SPICE capacity that your Quick Sight account is using.

Refreshing SPICE data

Refreshing a dataset

Use the following procedure to refresh a [SPICE](#) dataset based on an Amazon S3 or database data source in the **Data** tab. If there is a schema change in a database, Quick Sight will not be able to auto-detect it, resulting in an ingestion failure. Edit and save the dataset to update the schema and avoid ingestion failures.

To refresh SPICE data from the Data tab

1. Select **Data** from the left navigation menu. In the **Datasets** tab, choose the dataset to open it.
2. On the dataset details page that opens, choose the **Refresh** tab and then choose **Refresh now**.
3. Keep the refresh type as **Full refresh**.
4. If you are refreshing an Amazon S3 dataset, choose one of the following options for **S3 Manifest**:
 - To use the same manifest file you last provided to Amazon Quick Sight, choose **Existing Manifest**. If you have changed the manifest file at the file location or URL that you last provided, the data returned reflects those changes.
 - To specify a new manifest file by uploading it from your local network, choose **Upload Manifest**, and then choose **Upload manifest file**. For **Open**, choose a file, and then choose **Open**.
 - To specify a new manifest file by providing a URL, enter the URL of the manifest in **Input manifest URL**. You can find the manifest file URL in the Amazon S3 console by opening the context (right-click) menu for the manifest file, choosing **Properties**, and looking at the **Link** box.

5. Choose **Refresh**.
6. If you are refreshing an Amazon S3 dataset, choose **OK**, then **OK** again.
If you are refreshing a database dataset, choose **OK**.

Incrementally refreshing a dataset

Applies to: Enterprise Edition

For SQL-based data sources, such as Amazon Redshift, Amazon Athena, PostgreSQL, or Snowflake, you can refresh your data incrementally within a look-back window of time.

An *incremental refresh* queries only data defined by the dataset within a specified look-back window. It transfers all insertions, deletions, and modifications to the dataset, within that window's timeframe, from its source to the dataset. The data currently in SPICE that's within that window is deleted and replaced with the updates.

With incremental refreshes, less data is queried and transferred for each refresh. For example, let's say you have a dataset with 180,000 records that contains data from January 1 to June 30. On July 1, you run an incremental refresh on the data with a look-back window of seven days. Quick Sight queries the database asking for all data since June 24 (7 days ago), which is 7,000 records. Quick Sight then deletes the data currently in SPICE from June 24 and after, and appends the newly queried data. The next day (July 2), Quick Sight does the same thing, but queries from June 25 (7,000 records again), and then deletes from the existing dataset from the same date. Rather than having to ingest 180,000 records every day, it only has to ingest 7,000 records.

Use the following procedure to incrementally refresh a [SPICE](#) dataset based on a SQL data source from the **Datasets** tab.

To incrementally refresh a SQL-based SPICE dataset

1. Choose **Data** from the left navigation menu. On the **Datasets** tab, choose the dataset to open it.
2. On the dataset details page that opens, choose the **Refresh** tab and then choose **Refresh now**.
3. For **Refresh type**, choose **Incremental refresh**.
4. If this is your first incremental refresh on the dataset, choose **Configure**.
5. On the **Configure incremental refresh** page, do the following:

- a. For **Date column**, choose a date column that you want to base the look-back window on.
- b. For **Window size**, enter a number for **size**, and then choose an amount of time that you want to look back for changes.

You can choose to refresh changes to the data that occurred within a specified number of hours, days, or weeks from now. For example, you can choose to refresh changes to the data that occurred within two weeks of the current date.

6. Choose **Submit**.

Refreshing a dataset during data preparation

Use the following procedure to refresh a [SPICE](#) dataset based on an Amazon S3 or database data source during data preparation.

To refresh SPICE data during data preparation

1. Choose **Data** from the left navigation menu. On the **Datasets** tab, choose the dataset, and then choose **Edit Data Set**.
2. On the dataset screen, choose **Refresh now**.
3. Keep the refresh type set to **Full refresh**.
4. (Optional) If you are refreshing an Amazon S3 dataset, choose one of the following options for **S3 Manifest**:
 - To use the same manifest file that you last provided to Amazon Quick Sight, choose **Existing Manifest**. If you have changed the manifest file at the file location or URL that you last provided, the data returned reflects those changes.
 - To specify a new manifest file by uploading it from your local network, choose **Upload Manifest**, and then choose **Upload manifest file**. For **Open**, choose a file, and then choose **Open**.
 - To specify a new manifest file by providing a URL, enter the URL of the manifest in **Input manifest URL**. You can find the manifest file URL in the Amazon S3 console by opening the context (right-click) menu for the manifest file, choosing **Properties**, and looking at the **Link** box.
5. Choose **Refresh**.
6. If you are refreshing an Amazon S3 dataset, choose **OK**, then **OK** again.

If you are refreshing a database dataset, choose **OK**.

Refreshing a dataset on a schedule

Use the following procedure to schedule refreshing the data. If your dataset is based on a direct query and not stored in [SPICE](#), you can refresh your data by opening the dataset. You can also refresh your data by refreshing the page in an analysis or dashboard.

To refresh [SPICE](#) data on a schedule

1. Choose **Data** from the left navigation menu. On the **Datasets** tab, choose the dataset to open it.
2. On the dataset details page that opens, choose the **Refresh** tab and then choose **Add new schedule**.
3. On the **Create a refresh schedule** screen, choose settings for your schedule:
 - a. For **Time zone**, choose the time zone that applies to the data refresh.
 - b. For **Starting time**, choose a date and time for the refresh to start. Use HH:MM and 24-hour format, for example 13:30.
 - c. For **Frequency**, choose one of the following:
 - For Standard or Enterprise editions, you can choose **Daily**, **Weekly**, or **Monthly**.
 - **Daily**: Repeats every day.
 - **Weekly**: Repeats on the same day of each week.
 - **Monthly**: Repeats on the same day number of each month. To refresh data on the 29th, 30th or 31st day of the month, choose **Last day of month** from the list.
 - For Enterprise edition only, you can choose **Hourly**. This setting refreshes your dataset every hour, beginning at the time that you choose. So, if you select 1:05 as the starting time, the data refreshes at five minutes after the hour, every hour.

If you decide to use an hourly refresh, you can't also use additional refresh schedules. To create an hourly schedule, remove any other existing schedules for that dataset. Also, remove any existing hourly schedule before you create a daily, weekly, or monthly schedule.

4. Choose **Save**.

Scheduled dataset ingestions take place within 10 minutes of the scheduled date and time.

Using the Quick Suite console, you can create five schedules for each dataset. When you have created five, the **Create** button is turned off.

Incrementally refreshing a dataset on a schedule

Applies to: Enterprise Edition

For SQL-based data sources, such as Amazon Redshift, Athena, PostgreSQL, or Snowflake, you can schedule incremental refreshes. Use the following procedure to incrementally refresh a [SPICE](#) dataset based on a SQL data source in the **Datasets** tab.

To set an incremental refresh schedule for a SQL-based SPICE dataset

1. Choose **Data** from the left navigation menu. On the **Datasets** tab, choose the dataset to open it.
2. On the dataset details page that opens, choose the **Refresh** tab and then choose **Add new schedule**.
3. On the **Create a schedule** page, for **Refresh type**, choose **Incremental refresh**.
4. If this is your first incremental refresh for this dataset, choose **Configure**, and then do the following:
 - a. For **Date column**, choose a date column that you want to base the look-back window on.
 - b. For **Window size**, enter a number for **size**, and then choose an amount of time that you want to look back for changes.

You can choose to refresh changes to the data that occurred within a specified number of hours, days, or weeks from now. For example, you can choose to refresh changes to the data that occurred within two weeks of the current date.

- c. Choose **Submit**.
5. For **Time zone**, choose the time zone that applies to the data refresh.
 6. For **Repeats**, choose one of the following:
 - You can choose **Every 15 minutes**, **Every 30 minutes**, **Hourly**, **Daily**, **Weekly**, or **Monthly**.

- **Every 15 minutes:** Repeats every 15 minutes, beginning at the time you choose. So, if you select 1:05 as the starting time, the data refreshes at 1:20, then again at 1:35, and so on.
 - **Every 30 minutes:** Repeats every 30 minutes, beginning at the time you choose. So, if you select 1:05 as the starting time, the data refreshes at 1:35, then again at 2:05, and so on.
 - **Hourly:** Repeats every hour, beginning at the time you choose. So, if you select 1:05 as the starting time, the data refreshes at five minutes after the hour, every hour.
 - **Daily:** Repeats every day.
 - **Weekly:** Repeats on the same day of each week.
 - **Monthly:** Repeats on the same day number of each month. To refresh data on the 29th, 30th or 31st day of the month, choose **Last day of month** from the list.
 - If you decide to use refresh every 15 or 30 minutes, or hourly, you can't also use additional refresh schedules. To create a refresh schedule every 15 minutes, 30 minutes, or hourly, remove any other existing schedules for that dataset. Also, remove any existing minute or hourly schedule before you create a daily, weekly, or monthly schedule.
7. For **Starting**, choose a date for the refresh to start.
 8. For **At**, specify the time that the refresh should start. Use HH:MM and 24-hour format, for example 13:30.

Scheduled dataset ingestions take place within 10 minutes of the scheduled date and time.

In some cases, something might go wrong with the incremental refresh dataset that makes you want to roll back your dataset. Or you might no longer want to refresh the dataset incrementally. If so, you can delete the scheduled refresh.

To do so, choose the dataset on the **Datasets** page, choose **Schedule a refresh**, and then choose the x icon to the right of the scheduled refresh. Deleting an incremental refresh configuration starts a full refresh. As part of this full refresh, all the configurations prepared for incremental refreshes are removed.

Using SPICE data in an analysis

When you use stored data to create an analysis, a data import indicator appears next to the dataset list at the top of the **Fields list** pane. When you first open the analysis and the dataset is importing, a spinner icon appears.

After the SPICE import is complete, the indicator displays the percentage of rows that were successfully imported. A message also appears at the top of the visualization pane to provide counts of the rows imported and skipped.

If any rows were skipped, you can choose **View summary** in this message bar to see details about why those rows failed to import. To edit the dataset and resolve the issues that led to skipped rows, choose **Edit data set**. For more information about common causes for skipped rows, see [Troubleshooting skipped row errors](#).

If an import fails altogether, the data import indicator appears as an exclamation point icon, and an **Import failed** message is displayed.

View SPICE ingestion history

You can view the ingestion history for SPICE datasets to find out, for example, when the latest ingestion started and what its status is.

The SPICE ingestion history page includes the following information:

- Date and time that the ingestion started (UTC)
- Status of the ingestion
- Amount of time that the ingestion took
- The number of aggregated rows in the dataset.
- The number of rows ingested during a refresh.
- Rows skipped and rows ingested (imported) successfully
- The job type for the refresh: scheduled, full refresh, and so on

Use the following procedure to view a dataset's SPICE ingestion history.

To view a dataset's SPICE ingestion history

1. From the homepage, choose **Data** at left.
2. On the **Datasets** tab, choose the dataset that you want to examine.
3. On the dataset details page that opens, choose the **Refresh** tab.

SPICE ingestion history is shown at bottom.

4. (Optional) Choose a time frame to filter the entries from the last hour to the last 90 days.
5. (Optional) Choose a specific job status to filter the entries, for example **Running** or **Completed**. Otherwise, you can view all entries by choosing **All**.

Troubleshooting skipped row errors

When you import data, Amazon Quick Sight previews a portion of your data. If it can't interpret a row for any reason, Quick Sight skips the row. In some cases, the import will fail. When this happens, Quick Sight returns an error message that explains the failure.

Fortunately, there's a limited number of things that can go wrong. Some issues can be avoided by being aware of examples like the following:

- Make sure that there is no inconsistency between the field data type and the field data, for example occasional string data in a field with a numeric data type. Here are a few examples that can be difficult to detect when scanning the contents of a table:
 - ' ' – Using an empty string to indicate a missing value
 - 'NULL ' – Using the word "null" to indicate a missing value
 - \$1000 – Including a dollar sign in a currency value turns it into a string
 - 'O'Brien' – Using punctuation to mark a string that itself contains the same punctuation.

However, this type of error isn't always this easy to find, especially if you have a lot of data, or if your data is typed in by hand. For example, some customer service or sales applications involve entering information verbally provided by customers. The person who originally typed in the data might have put it in the wrong field. They might add, or forget to add, a character or digit. For example, they might enter a date of "0/10/12020" or enter someone's gender in a field meant for age.

- Make sure that your imported file is correctly processed with or without a header. If there is a header row, make sure that you choose the **Contains header** upload option.
- Make sure that the data doesn't exceed one or more of the [Data source quotas](#).
- Make sure that the data is compatible with the [Supported data types and values](#).
- Make sure that your calculated fields contain data that works with the calculation, rather than being incompatible with or excluded by the function in the calculated field. For example, if you have a calculated field in your dataset that uses [parseDate](#), Quick Sight skips rows where that field doesn't contain a valid date.

Quick Sight provides a detailed list of the errors that occur when the SPICE engine attempts to ingest data. When a saved dataset reports skipped rows, you can view the errors so you can take action to fix the issues.

To view errors for rows that were skipped during SPICE ingestion (data import)

1. Choose **Data** at left. In the **Datasets** tab, choose the problematic dataset to open it.
2. On the dataset details page that opens, choose the **Refresh** tab.

SPICE ingestion history is shown at bottom.

3. For the ingestion with the error, choose **View error summary**. This link is located under the **Status** column.
4. Examine the **File import log** that opens. It displays the following sections:
 - **Summary** – Provides a percentage score of how many rows were skipped out of the total number of rows in the import. For example, if there are 864 rows skipped out of a total of 1,728, the score is 50.00%.
 - **Skipped Rows** – Provides the row count, field name, and error message for each set of similar skipped rows.
 - **Troubleshooting** – Provides a link to download a file that contains error information.
5. Under **Troubleshooting**, choose **Download error rows file**.

The error file has a row for each error. The file is named `error-report_123_fe8.csv`, where `123_fe8` is replaced with a unique identifying string. The file contains the following columns:

- **ERROR_TYPE** – The type or error code for the error that occurred when importing this row. You can look up this error in the [SPICE ingestion error codes](#) section that follows this procedure.
 - **COLUMN_NAME** – The name of the column in your data that caused the error.
 - All the columns from your imported row – The remaining columns duplicate the entire row of data. If a row has more than one error, it can appear multiple times in this file.
6. Choose **Edit data set** to make changes to your dataset. You can filter the data, omit fields, change data types, adjust existing calculated fields, and add calculated fields that validate the data.
 7. After you've made changes indicated by the error codes, import the data again. If more SPICE ingestion errors appear in the log, step through this procedure again to fix all remaining errors.

Tip

If you can't solve the data issues in a reasonable amount of time by using the dataset editor, consult the administrators or developers who own the data. In the long run, it's more cost-effective to cleanse the data closer to its source, rather than adding exception processing while you're preparing the data for analysis. By fixing it at the source, you avoid a situation where multiple people fix the errors in different ways, resulting in different reporting results later on.

To practice troubleshooting skipped rows

1. Download [CSV files for troubleshooting skipped rows.zip](#).
2. Extract the files into a folder that you can use to upload the sample .csv file into Quick Sight.

The zip file contains the following two text files:

- `sample dataset - data ingestion error.csv` – A sample .csv file that contains issues that cause skipped rows. You can try to import the file yourself to see how the error process works.
 - `sample data ingestion error file` – A sample error file generated during SPICE ingestion while importing the sample .csv file into Quick Sight.
3. Import the data by following these steps:
 - a. Choose **Data, Datasets** tab, **New, Dataset**.
 - b. Choose **Upload a file**.
 - c. Find and choose the file named `sample dataset - data ingestion error.csv`.
 - d. Choose **Upload a file, Edit settings and prepare data**.
 - e. Choose **Save** to exit.
 4. Choose your dataset to view its information, then choose **View error summary**. Examine the errors and the data to help you resolve the issues.

SPICE ingestion error codes

The following list of errors codes and descriptions can help you understand and troubleshoot issues with data ingestion into SPICE.

Error codes for skipped rows

The following list of errors codes and descriptions can help you understand and troubleshoot issues with skipped rows.

ARITHMETIC_EXCEPTION – An arithmetic exception occurred while processing a value.

ENCODING_EXCEPTION – An unknown exception occurred while converting and encoding data to SPICE.

OPENSEARCH_CURSOR_NOT_ENABLED – The OpenSearch domain doesn't have SQL cursors enabled ("opendistro.sql.cursor.enabled" : "true"). For more information, see [Authorizing connections to Amazon OpenSearch Service](#).

INCORRECT_FIELD_COUNT – One or more rows have too many fields. Make sure that the number of fields in each row matches the number of fields defined in the schema.

INCORRECT_SAGEMAKER_OUTPUT_FIELD_COUNT – The SageMaker AI output has an unexpected number of fields.

INDEX_OUT_OF_BOUNDS – The system requested an index that isn't valid for the array or list being processed.

MALFORMED_DATE – A value in a field can't be transformed to a valid date. For example, if you try to convert a field that contains a value like "sale date" or "month-1", the action generates a malformed date error. To fix this error, remove nondatetime values from your data source. Check that you aren't importing a file with a column header mixed into the data. If your string contains a date or time that doesn't convert, see [Using unsupported or custom dates](#).

MISSING_SAGEMAKER_OUTPUT_FIELD – A field in the SageMaker AI output is unexpectedly empty.

NUMBER_BITWIDTH_TOO_LARGE – A numeric value exceeds the length supported in SPICE. For example, your numeric value has more than 19 digits, which is the length of a bigint data type. For a long numeric sequence that isn't a mathematical value, use a string data type.

NUMBER_PARSE_FAILURE – A value in a numeric field is not a number. For example, a field with a data type of int contains a string or a float.

SAGEMAKER_OUTPUT_COLUMN_TYPE_MISMATCH – The data type defined in the SageMaker AI schema doesn't match the data type received from SageMaker AI.

STRING_TRUNCATION – A string is being truncated by SPICE. Strings are truncated where the length of the string exceeds the SPICE quota. For more information about SPICE, see [Importing data into SPICE](#). For more information about quotas, see [Service Quotas](#).

UNDEFINED – An unknown error occurred while ingesting data.

UNSUPPORTED_DATE_VALUE – A date field contains a date that is in a supported format but is not in the supported range of dates, for example "12/31/1399" or "01/01/10000". For more information, see [Using unsupported or custom dates](#).

Error codes during data import

For imports and data refresh jobs that fail, Quick Sight provides an error code indicating what caused the failure. The following list of errors codes and descriptions can help you understand and troubleshoot issues with data ingestion into SPICE.

ACCOUNT_CAPACITY_LIMIT_EXCEEDED – This data exceeds your current SPICE capacity. Purchase more SPICE capacity or clean up existing SPICE data and then retry this ingestion.

CONNECTION_FAILURE – Amazon Quick Sight can't connect to your data source. Check the data source connection settings and try again.

CUSTOMER_ERROR – There was a problem parsing the data. If this persists, contact Amazon Quick Sight technical support.

DATA_SET_DELETED – The data source or dataset was deleted or became unavailable during ingestion.

DATA_SET_SIZE_LIMIT_EXCEEDED – This dataset exceeds the maximum allowable SPICE dataset size. Use filters to reduce the dataset size and try again. For information on SPICE quotas, see [Data source quotas](#).

DATA_SOURCE_AUTH_FAILED – Data source authentication failed. Check your credentials and use the **Edit data source** option to replace expired credentials.

DATA_SOURCE_CONNECTION_FAILED – Data source connection failed. Check the URL and try again. If this error persists, contact your data source administrator for assistance.

DATA_SOURCE_NOT_FOUND – No data source found. Check your Amazon Quick Sight data sources.

DATA_TOLERANCE_EXCEPTION – There are too many invalid rows. Amazon Quick Sight has reached the quota of rows that it can skip and still continue ingesting. Check your data and try again.

FAILURE_TO_ASSUME_ROLE – Amazon Quick Sight couldn't assume the correct Amazon Identity and Access Management (IAM) role. Verify the policies for Amazon Quick Sight-service-role in the IAM console.

FAILURE_TO_PROCESS_JSON_FILE – Amazon Quick Sight couldn't parse a manifest file as valid JSON.

IAM_ROLE_NOT_AVAILABLE – Amazon Quick Sight doesn't have permission to access the data source. To manage Amazon Quick Sight permissions on Amazon resources, go to the **Security and Permissions** page under the **Manage Amazon Quick Sight** option as an administrator.

INGESTION_CANCELED – The ingestion was canceled by the user.

INGESTION_SUPERSEDED – This ingestion has been superseded by another workflow. This happens when a new ingestion is created while another one is still in progress. Avoid manually editing the dataset multiple times in a short period, because each manual edit creates a new ingestion which will supersede and end the previous ingestion.

INTERNAL_SERVICE_ERROR – An internal service error occurred.

INVALID_DATA_SOURCE_CONFIG – Invalid values appeared in connection settings. Check your connection details and try again.

INVALID_DATAPREP_SYNTAX – Your calculated field expression contains invalid syntax. Correct the syntax and try again.

INVALID_DATE_FORMAT – An invalid date format appeared.

IOT_DATA_SET_FILE_EMPTY – No Amazon IoT Analytics data was found. Check your account and try again.

IOT_FILE_NOT_FOUND – An indicated Amazon IoT Analytics file wasn't found. Check your account and try again.

OAUTH_TOKEN_FAILURE – Credentials to the data source have expired. Renew your credentials and retry this ingestion.

PASSWORD_AUTHENTICATION_FAILURE – Incorrect credentials appeared for a data source. Update your data source credentials and retry this ingestion.

PERMISSION_DENIED – Access to the requested resources was denied by the data source. Request permissions from your database administrator or ensure proper permission has been granted to Amazon Quick Sight before retrying.

QUERY_TIMEOUT – A query to the data source timed out waiting for a response. Check your data source logs and try again.

ROW_SIZE_LIMIT_EXCEEDED – The row size quota exceeded the maximum.

S3_FILE_INACCESSIBLE – Couldn't connect to an S3 bucket. Make sure that you grant Amazon Quick Sight and users necessary permissions before you connect to the S3 bucket.

S3_MANIFEST_ERROR – Couldn't connect to S3 data. Make sure that your S3 manifest file is valid. Also verify access to the S3 data. Both Amazon Quick Sight and the Amazon Quick Sight user need permissions to connect to the S3 data.

S3_UPLOADED_FILE_DELETED – The file or files for the ingestion were deleted (between ingestions). Check your S3 bucket and try again.

SOURCE_API_LIMIT_EXCEEDED_FAILURE – This ingestion exceeds the API quota for this data source. Contact your data source administrator for assistance.

SOURCE_RESOURCE_LIMIT_EXCEEDED – A SQL query exceeds the resource quota of the data source. Examples of resources involved can include the concurrent query quota, the connection quota, and physical server resources. Contact your data source administrator for assistance.

SPICE_TABLE_NOT_FOUND – An Amazon Quick Sight data source or dataset was deleted or became unavailable during ingestion. Check your dataset in Amazon Quick Sight and try again. For more information, see [Troubleshooting skipped row errors](#).

SQL_EXCEPTION – A general SQL error occurred. This error can be caused by query timeouts, resource constraints, unexpected data definition language (DDL) changes before or during a query, and other database errors. Check your database settings and your query, and try again.

SQL_INVALID_PARAMETER_VALUE – An invalid SQL parameter appeared. Check your SQL and try again.

SQL_NUMERIC_OVERFLOW – Amazon Quick Sight encountered an out-of-range numeric exception. Check related values and calculated columns for overflows, and try again.

SQL_SCHEMA_MISMATCH_ERROR – The data source schema doesn't match the Amazon Quick Sight dataset. Update your Amazon Quick Sight dataset definition.

SQL_TABLE_NOT_FOUND – Amazon Quick Sight can't find the table in the data source. Verify the table specified in the dataset or custom SQL and try again.

SSL_CERTIFICATE_VALIDATION_FAILURE – Amazon Quick Sight can't validate the Secure Sockets Layer (SSL) certificate on your database server. Check the SSL status on that server with your database administrator and try again.

UNRESOLVABLE_HOST – Amazon Quick Sight can't resolve the host name of the data source. Verify the host name of the data source and try again.

UNROUTABLE_HOST – Amazon Quick Sight can't reach your data source because it's inside a private network. Ensure that your private VPC connection is configured correctly in Enterprise Edition, or allow Amazon Quick Sight IP address ranges to allow connectivity for Standard Edition.

Updating files in a dataset

To get the latest version of files, you can update files in your dataset. You can update these types of files:

- Comma-delimited (CSV) and tab-delimited (TSV) text files
- Extended and common log format files (ELF and CLF)
- Flat or semistructured data files (JSON)
- Microsoft Excel (XLSX) files

Before updating a file, make sure that the new file has the same fields in the same order as the original file currently in the dataset. If there are field (column) discrepancies between the two files, an error occurs and you need to fix the discrepancies before attempting to update again. You can do this by editing the new file to match the original. Note that if you want to add new fields, you can append them after the original fields in the file. For example, in a Microsoft Excel spreadsheet, you can append new fields to the right of the original fields.

To update a file in a dataset

1. In Quick Sight, choose **Data** at left.
2. In the **Datasets** tab, choose the dataset that you want to update, and then choose **Edit dataset**.

3. On the data preparation page that opens, choose the drop-down list for the file that you want to update, and then choose **Update file**.
4. On the **Update file** page that opens, choose **Upload file**, and then navigate to a file.

Quick Sight scans the file.

5. If the file is a Microsoft Excel file, choose the sheet that you want on the **Choose your sheet** page that opens, and then choose **Select**.
6. Choose **Confirm file update** on the following page. A preview of some of the sheet columns is shown for your reference.

A message that the file updated successfully appears at top right and the table preview updates to show the new file data.

Preparing data in Amazon Quick Sight

Datasets store any data preparation you have done on that data, so that you can reuse that prepared data in multiple analyses. Data preparation provides options such as adding calculated fields, applying filters, and changing field names or data types. If you are basing the data source on a SQL database, you can also use data preparation to join tables. Or you can enter a SQL query if you want to work with data from more than a single table.

If you want to transform the data from a data source before using it in Amazon Quick Sight, you can prepare it to suit your needs. You then save this preparation as part of the dataset.

You can prepare a dataset when you create it, or by editing it later. For more information about creating a new dataset and preparing it, see [Creating datasets](#). For more information about opening an existing dataset for data preparation, see [Editing datasets](#).

Use the following topics to learn more about data preparation.

Topics

- [Describing data](#)
- [Choosing file upload settings](#)
- [Preparing data fields for analysis in Amazon Quick Sight](#)
- [Adding calculations](#)
- [Previewing tables in a dataset](#)
- [Joining data](#)

- [Filtering data in Amazon Quick Sight](#)
- [Using SQL to customize data](#)
- [Adding geospatial data](#)
- [Using unsupported or custom dates](#)
- [Adding a unique key to an Amazon Quick Sight dataset](#)
- [Integrating Amazon SageMaker AI models with Amazon Quick Sight](#)
- [Preparing dataset examples](#)

Describing data

Using Amazon Quick Sight, you can add information, or *metadata*, about the columns (fields) in your datasets. By adding metadata, you make the dataset self-explanatory and easier to reuse. Doing this can help data curators and their customers know where the data came from and what it means. It's a way of communicating to the people who use your dataset or combine it with other datasets to build dashboards. Metadata is especially important for information that is shared between organizations.

After you add metadata to a dataset, field descriptions become available to anyone who is using the dataset. A column description appears when someone who is actively browsing the **Fields** list pauses on a field name. Column descriptions are visible to people who are editing a dataset or an analysis, but not to someone who is viewing a dashboard. Descriptions aren't formatted. You are able to enter line feeds and formatting marks and these are preserved by the editor. However, the description tooltip that displays is only able to show words, numbers, and symbols—but not formatting.

To edit a description to a column or field

1. From the Quick Suite homepage, choose **Data** at left.
2. In the **Data** tab, choose the dataset that you want to work on.
3. On the dataset details page that opens, choose **Edit dataset** at upper right.
4. On the dataset page that opens, choose a column in the table preview at bottom or in the field list at left.
5. To add or change the description, do one of the following:
 - At the bottom of the screen, open the settings for the field from the pencil icon next to the field's name.

- In the field list, open the settings for the field from the menu next to the field's name. Then choose **Edit name & description** from the context menu.
6. Add or change the description for the field.

To delete an existing description, delete all the text in the Description box.
 7. (Optional) For **Name**, if you want to change the name of the field, you can enter a new one here.
 8. Choose **Apply** to save your changes. Choose cancel to exit.

Choosing file upload settings

If you are using a file data source, confirm the upload settings, and correct them if necessary.

Important

If it's necessary to change upload settings, make these changes before you make any other changes to the dataset. Changing upload settings causes Amazon Quick Sight to reimport the file. This process overwrites any changes you have made so far.

Changing text file upload settings

Text file upload settings include the file header indicator, file format, text delimiter, text qualifier, and start row. If you are working with an Amazon S3 data source, the upload settings you select are applied to all files you choose to use in this dataset.

Use the following procedure to change text file upload settings.

1. On the data preparation page, open the **Upload Settings** pane by choosing the expand icon.
2. In **File format**, choose the file format type.
3. If you chose the **custom separated (CUSTOM)** format, specify the separating character in **Delimiter**.
4. If the file doesn't contain a header row, deselect the **Files include headers** check box.
5. If you want to start from a row other than the first row, specify the row number in **Start from row**. If the **Files include headers** check box is selected, the new starting row is treated as the header row. If the **Files include headers** check box is not selected, the new starting row is treated as the first data row.

6. In **Text qualifier**, choose the text qualifier, either single quotes (') or double quotes (").

Changing Microsoft Excel file upload settings

Microsoft Excel file upload settings include the range header indicator and whole worksheet selector.

Use the following procedure to change Microsoft Excel file upload settings.

1. On the data preparation page, open the **Upload Settings** pane by choosing the expand icon.
2. Leave **Upload whole sheet** selected.
3. If the file doesn't contain a header row, deselect the **Range contains headers** check box.

Preparing data fields for analysis in Amazon Quick Sight

Before you start analyzing and visualizing your data, you can prepare the fields (columns) in your dataset for analysis. You can edit field names and descriptions, change the data type for fields, set up drill-down hierarchies for fields, and more.

Use the following topics to prepare fields in your dataset.

Topics

- [Editing field names and descriptions](#)
- [Setting fields as a dimensions or measures](#)
- [Changing a field data type](#)
- [Adding drill-downs to visual data in Quick Sight](#)
- [Selecting fields](#)
- [Organizing fields into folders in Amazon QuickSight](#)
- [Mapping and joining fields](#)

Editing field names and descriptions

You can change any field name and description from what is provided by the data source. If you change the name of a field used in a calculated field, make sure also to change it in the calculated field function. Otherwise, the function fails.

To change a field name or description

1. In the **Fields** pane of the data prep page, choose the three-dot icon on the field that you want to change. Then choose **Edit name & description**.
2. Enter the new name or description that you want to change, and choose **Apply**.

You can also change the name and description of a field on the data prep page. To do this, select the column header of the field that you want to change in the **Dataset** table in that page's lower half. Then make any changes there.

Setting fields as a dimensions or measures

In the **Field list** pane, dimension fields have blue icons and measure fields have green icons. *Dimensions* are text or date fields that can be items, like products, or attributes that are related to measures. You can use dimensions to partition these items or attributes, like sales date for sales figures. *Measures* are numeric values that you use for measurement, comparison, and aggregation.

In some cases, Quick Sight interprets a field as a measure that you want to use it as a dimension (or the other way around). If so, you can change the setting for that field.

Changing a field's measure or dimension setting changes it for all visuals in the analysis that use that dataset. However, it doesn't change it in the dataset itself.

Changing a field's dimension or measure setting

Use the following procedure to change a field's dimension or measure setting

To change a field's dimension or measure setting

1. In the **Field list** pane, hover over the field that you want to change.
2. Choose the selector icon to the right of the field name, and then choose **Convert to dimension** or **Convert to measure** as appropriate.

Changing a field data type

When Quick Sight retrieves data, it assigns each field a data type based on the data in the field. The possible data types are as follows:

- **Date** – The date data type is used for date data in a supported format. For information about the date formats Quick Sight supports, see [Data source quotas](#).

- **Decimal** – The decimal data type is used for numeric data that requires one or more decimal places of precision, for example 18.23. The decimal data type supports values with up to four decimal places to the right of the decimal point. Values that have a higher scale than this are truncated to the fourth decimal place in two cases. One is when these values are displayed in data preparation or analyses, and one is when these values are imported into Quick Sight. For example, 13.00049 is truncated to 13.0004.
- **Geospatial** – The geospatial data type is used for geospatial data, for example longitude and latitude, or cities and countries.
- **Integer** – The int data type is used for numeric data that only contains integers, for example 39.
- **String** – The string data type is used for nondate alphanumeric data.

Quick Sight reads a small sample of rows in the column to determine the data type. The data type that occurs most in the small sample size is the suggested type. In some cases, there might be blank values (treated as strings by Quick Sight) in a column that contains mostly numbers. In these cases, it might be that the String data type is the most frequent type in the sample set of rows. You can manually modify the data type of the column to make it integer. Use the following procedures to learn how.

Changing a field data type during data prep

During data preparation, you can change the data type of any field from the data source. On the **Change data type** menu, you can change calculated fields that don't include aggregations to geospatial types. You can make other changes to the data type of a calculated field by modifying its expression directly. Quick Sight converts the field data according to the data type that you choose. Rows that contain data that is incompatible with that data type are skipped. For example, suppose that you convert the following field from String to Integer.

```
10020
36803
14267a
98457
78216b
```

All records containing alphabetic characters in that field are skipped, as shown following.

```
10020
36803
```

98457

If you have a database dataset with fields whose data types aren't supported by Quick Sight, use a SQL query during data preparation. Then use CAST or CONVERT commands (depending on what is supported by the source database) to change the field data types. For more information about adding a SQL query during data preparation, see [Using SQL to customize data](#). For more information about how different source data types are interpreted by Quick Sight, see [Supported data types from external data sources](#).

You might have numeric fields that act as dimensions rather than metrics, for example ZIP codes and most ID numbers. In these cases, it's helpful to give them a string data type during data preparation. Doing this lets Quick Sight understand that they are not useful for performing mathematical calculations and can only be aggregated with the Count function. For more information about how Quick Sight uses dimensions and measures, see [Setting fields as a dimensions or measures](#).

In [SPICE](#), numbers converted from numeric into an integer are truncated by default. If you want to round your numbers instead, you can create a calculated field using the [round](#) function. To see whether numbers are rounded or truncated before they are ingested into SPICE, check your database engine.

To change a field data type during data prep

1. From the Quick Sight homepage, choose **Data** at left. In the **Data** tab, choose the dataset that you want, and then choose **Edit dataset**.
2. In the data preview pane, choose the data type icon under the field you want to change.
3. Choose the target data type. Only data types other than the one currently in use are listed.

Changing a field data type in an analysis

You can use the **Field list** pane, visual field wells, or on-visual editors to change numeric field data types within the context of an analysis. Numeric fields default to displaying as numbers, but you can choose to have them display as currency or as a percentage instead. You can't change the data types for string or date fields.

Changing a field's data type in an analysis changes it for all visuals in the analysis that use that dataset. However, it doesn't change it in the dataset itself.

Note

If you are working in a pivot table visual, applying a table calculation changes the data type of the cell values in some cases. This type of change occurs if the data type doesn't make sense with the applied calculation.

For example, suppose that you apply the Rank function to a numeric field that you modified to use a currency data type. In this case, the cell values display as numbers rather than currency. Similarly, if you apply the Percent difference function instead, the cell values display as percentages rather than currency.

To change a field's data type

1. Choose one of the following options:
 - In the **Field list** pane, hover over the numeric field that you want to change. Then choose the selector icon to the right of the field name.
 - On any visual that contains an on-visual editor associated with the numeric field that you want to change, choose that on-visual editor.
 - Expand the **Field wells** pane, and then choose the field well associated with the numeric field that you want to change.
2. Choose **Show as**, and then choose **Number**, **Currency**, or **Percent**.

Adding drill-downs to visual data in Quick Sight

All visual types except pivot tables offer the ability to create a hierarchy of fields for a visual element. The hierarchy lets you drill down to see data at different levels of the hierarchy. For example, you can associate the country, state, and city fields with the x-axis on a bar chart. Then, you can drill down or up to see data at each of those levels. As you drill down each level, the data displayed is refined by the value in the field you drill down on. For example, if you drill down on the state of California, you see data on all of the cities in California.

The field wells you can use to create drill-downs varies by visual type. Refer to the topic on each visual type to learn more about its drill-down support.

Drill-down functionality is added automatically for dates when you associate a date field with the drill-down field well of a visual. In this case, you can always drill up and down through the levels

of date granularity. Drill-down functionality is also added automatically for geospatial groupings, after you define these in the dataset.

Use the following table to identify the field wells/on-visual editors that support drill-down for each visual type.

Visual type	Field well or on-visual editor
Bar charts (all horizontal)	Y axis and Group/Color
Bar charts (all vertical)	X axis and Group/Color
Combo charts (all)	X axis and Group/Color
Geospatial charts	Geospatial and Color
Heat map	Rows and Columns
KPIs	Trend Group
Line charts (all)	X axis and Color
Pie chart	Group/Color
Scatter plot	Group/Color
Tree map	Group by

 Important

Drill-downs are not supported for tables or pivot tables.

Adding a drill-down

Use the following procedure to add drill-down levels to a visual.

To add drill-down levels to a visual

1. On the analysis page, choose the visual that you want to add drill-downs to.

2. Drag a field item into a **Field well**.
3. If your dataset has a defined hierarchy, you can drag the entire hierarchy into the field well as one. An example is geospatial or coordinate data. In this case, you don't need to follow the remaining steps.

If you don't have a predefined hierarchy, you can create one in your analysis, as described in the remaining steps.

4. Drag a field that you want to use in the drill-down hierarchy to an appropriate field well, depending on the visual type. Make sure that the label for the dragged field says **Add drill-down layer**. Position the dragged field above or below the existing field based on where you want it to be in the hierarchy you're creating.
5. Continue until you have added all of the levels of hierarchy that you want. To remove a field from the hierarchy, choose the field, and then choose **Remove**.
6. To drill down or up to see data at a different level of the hierarchy, choose an element on the visual (like a line or bar), and then choose **Drill down to <lower level>** or **Drill up to <higher level>**. In this example, from the `car-make` level you can drill down to `car-model` to see data at that level. If you drill down to `car-model` from the **Ford** `car-make`, you see only `car-models` in that `car-make`.

After you drill down to the `car-model` level, you can then drill down further to see `make-year` data, or go back up to `car-make`. If you drill down to `make-year` from the bar representing **Ranger**, you see only years for that model of car.

Selecting fields

When you prepare data, you can select one or more fields to perform an action on them, such as excluding them or adding them to a folder.

To select one or more fields in the data preparation pane, click or tap the field or fields in the **Fields** pane at left. You can then choose the field menu (the three dots) to the right of the field name and choose an action to take. The action is performed on all selected fields.

You can select or deselect all fields at once by choosing either **All** or **None** at the top of the **Fields** pane.

If you edit a dataset and exclude a field that is used in a visual, that visual breaks. You can fix it the next time you open that analysis.

Searching for fields

If you have a long field list in the **Fields** pane, you can search to locate a specific field by entering a search term for **Search fields**. Any field whose name contains the search term is shown.

Search is case-insensitive and wildcards are not supported. Choose the cancel icon (X) to the right of the search box to return to viewing all fields.

Organizing fields into folders in Amazon QuickSight

When prepping your data in Quick Sight, you can use folders to organize your fields for multiple authors across your enterprise. Arranging fields into folders and subfolders can make it easier for authors to find and understand fields in your dataset.

You can create folders while preparing your dataset, or when editing a dataset. For more information about creating a new dataset and preparing it, see [Creating datasets](#). For more information about opening an existing dataset for data preparation, see [Editing datasets](#).

While performing an analysis, authors can expand and collapse folders, search for specific fields within folders, and see your descriptions of folders on the folder menu. Folders appear at the top of the **Fields** pane in alphabetical order.

Creating a folder

Use the following procedure to create a new folder in the **Fields** pane.

To create a new folder

1. On the data preparation page, in the **Fields** pane, select the three-dot icon, and choose **Add to folder**.

To select more than one field at a time, press the Ctrl key while you select (Command key on Mac).

2. On the **Add to folder** page that appears, choose **Create a new folder** and enter a name for the new folder.
3. Choose **Apply**.

The folder appears at the top of the **Fields** pane with the fields that you chose inside it. Fields inside folders are arranged in alphabetical order.

Creating a subfolder

To further organize your data fields in the **Fields** pane, you can create subfolders within parent folders.

To create a subfolder

1. On the data preparation page, in the **Fields** pane, select the field menu for a field already in a folder and choose **Move to folder**.
2. On the **Move to folder** page that appears, choose **Create a new folder** and enter a name for the new folder.
3. Choose **Apply**.

The subfolder appears within the parent folder at the top of the list of fields. Subfolders are arranged in alphabetical order.

Adding fields to an existing Folder

Use the following procedure to add fields to an existing folder in the **Fields** pane.

To add one or more fields to a folder

1. On the data preparation page, in the **Fields** pane, select the fields that you want to add to a folder.

To select more than one field at a time, press the Ctrl key while you select (Command key on Mac).
2. On the field menu, choose **Add to folder**.
3. On the **Add to folder** page that appears, choose a folder for **Existing folder**.
4. Choose **Apply**.

The field or fields are added to the folder.

Moving fields between folders

Use the following procedure to move fields between folders in the **Fields** pane.

To move fields between folders

1. On the data preparation page, in the **Fields** pane, select the fields that you want to move to another folder.

To select more than one field at a time, press the Ctrl key while you select (Command key on Mac).

2. On the field menu, choose **Move to folder**.
3. On the **Move to folder** page that appears, choose a folder for **Existing folder**.
4. Choose **Apply**.

Removing fields from a folder

Use the following procedure to remove fields from a folder in the **Fields** pane. Removing a field from a folder doesn't delete the field.

To remove fields from a folder

1. On the data preparation page, in the **Fields** pane, select the fields that you want to remove.
2. On the field menu, choose **Remove from folder**.

The fields that you selected are removed from the folder and placed back in the list of fields in alphabetical order.

Editing a folder name and adding a folder description

You can edit the name or add a description of a folder to provide context about the data fields inside it. The folder name appears in the **Fields** pane. While performing an analysis, authors can read your folder's description when they select the folder menu in the **Fields** pane.

To edit a folder name or edit or add a description for a folder

1. On the data preparation page, in the **Fields** pane, select the folder menu for the folder that you want to edit and choose **Edit name & description**.
2. On the **Edit folder** page that appears, do the following:
 - For **Name**, enter a name for the folder.
 - For **Description**, enter a description of the folder.

3. Choose **Apply**.

Moving folders

You can move folders and subfolders to new or existing folders in the **Fields** pane.

To move a folder

1. On the data preparation page, in the **Fields** pane, choose **Move folder** on the folder menu.
2. On the **Move folder** page that appears, do one of the following:
 - Choose **Create a new folder** and enter a name for the folder.
 - For **Existing folder**, choose a folder.
3. Choose **Apply**.

The folder appears within the folder that you chose in the **Fields** pane.

Removing folders from the fields pane

Use the following procedure to remove a folder from the **Fields** pane.

To remove a folder

1. On the data preparation page, in the **Fields** pane, choose **Remove folder** on the folder menu.
2. On the **Remove folder?** page that appears, choose **Remove**.

The folder is removed from the **Fields** pane. Any fields that were in the folder are placed back in the list of fields in alphabetical order. Removing folders doesn't exclude fields from view or delete fields from the dataset.

Mapping and joining fields

When you are using different datasets together in Quick Sight, you can simplify the process of mapping fields or joining tables during the data preparation stage. You should already be verifying that your fields have the correct data type and an appropriate field name. However, if you already know which datasets are going to be used together, you can take a couple of extra steps to make your work easier later on.

Mapping fields

Quick Sight can automatically map fields between datasets in the same analysis. The following tips can help make it easier for Quick Sight to automatically map fields between datasets, for example if you are creating a filter action across datasets:

- **Matching field names** – Field names must match exactly, with no differences in case, spacing, or punctuation. You can rename fields that describe the same data, so an automatic mapping is accurate.
- **Matching data types** – Fields must have the same data type for automatic mapping. You can change the data types while you are preparing the data. This step also gives you the opportunity to discover whether you need to filter out any data that isn't the correct data type.
- **Using calculated fields** – You can use calculated fields to create a matching field, and give it the correct name and data type for automatic mapping.

Note

After an automatic mapping exists, you can rename a field without breaking the field mapping. However, if you change the data type, the mapping is broken.

For more information on field mapping for filter actions across datasets, see [Creating and editing custom actions in Amazon Quick Sight](#).

Joining fields

You can create joins between data from different data sources, including files or databases. The following tips can help make it easier for you to join data from different files or data sources:

- **Similar field names** – It is simpler to join fields when you can see what should match; for example, **Order ID** and **order-id** seem as if they should be the same. But if one is a work order, and the other is a purchase order, then the fields are probably different data. If possible, make sure that the files and tables that you want to join have field names making it clear what data they contain.
- **Matching data types** – Fields must have the same data type before you can join on them. Make sure that the files and tables that you want to join having matching data types in join fields. You can't use a calculated field for a join. Also, you can't join two existing datasets. You create the joined dataset by directly accessing the source data.

For more information on joining data across data sources, see [Joining data](#).

Adding calculations

Create calculated fields to transform your data by using one or more of the following:

- [Operators](#)
- [Functions](#)
- Fields that contain data
- Other calculated fields

You can add calculated fields to a dataset during data preparation or from the analysis page. When you add a calculated field to a dataset during data preparation, it's available to all analyses that use that dataset. When you add a calculated field to a dataset in an analysis, it's available only in that analysis. For more information about adding calculated fields, see the following topics.

Topics

- [Adding calculated fields](#)
- [Order of evaluation in Amazon Quick Sight](#)
- [Using level-aware calculations in Quick Sight](#)
- [Calculated field function and operator reference for Amazon Quick Suite](#)

Adding calculated fields

Create calculated fields to transform your data by using one or more of the following:

- [Operators](#)
- [Functions](#)
- Aggregate functions (you can only add these to an analysis)
- Fields that contain data
- Other calculated fields

You can add calculated fields to a dataset during data preparation or from the analysis page. When you add a calculated field to a dataset during data preparation, it's available to all analyses that use

that dataset. When you add a calculated field to a dataset in an analysis, it's available only in that analysis.

Analyses support both single-row operations and aggregate operations. Single-row operations are those that supply a (potentially) different result for every row. Aggregate operations supply results that are always the same for entire sets of rows. For example, if you use a simple string function with no conditions, it changes every row. If you use an aggregate function, it applies to all the rows in a group. If you ask for the total sales amount for the US, the same number applies to the entire set. If you ask for data on a particular state, the total sales amount changes to reflect your new grouping. It still provides one result for the entire set.

By creating the aggregated calculated field within the analysis, you can then drill down into the data. The value of that aggregated field is recalculated appropriately for each level. This type of aggregation isn't possible during dataset preparation.

For example, let's say that you want to figure out the percentage of profit for each country, region, and state. You can add a calculated field to your analysis, $(\text{sum}(\text{salesAmount} - \text{cost})) / \text{sum}(\text{salesAmount})$. This field is then calculated for each country, region, and state, at the time your analyst drills down into the geography.

Topics

- [Adding calculated fields to an analysis](#)
- [Adding calculated fields to a dataset](#)
- [Handling decimal values in calculated fields](#)

Adding calculated fields to an analysis

When you add a dataset to an analysis, every calculated field that exists in the dataset is added to the analysis. You can add additional calculated fields at the analysis level to create calculated fields that are available only in that analysis.

To add a calculated field to an analysis

1. Open the [Quick Suite console](#).
2. Open the analysis that you want to change.
3. In the **Data** pane, choose **Add** at top left, and then choose **+ CALCULATED FIELD**.
 - a. In the calculations editor that opens, do the following:

- b. Enter a name for the calculated field.
 - c. Enter a formula using fields from your dataset, functions, and operators.
4. When finished, choose **Save**.

For more information about how to create formulas using the available functions in Quick Sight, see [Calculated field function and operator reference for Amazon Quick Suite](#).

Adding calculated fields to a dataset

Amazon Quick Sight authors can generate calculated fields during the data preparation phase of a dataset's creation. When you create a calculated field for a dataset, the field becomes a new column in the dataset. All analyses that use the dataset inherit the dataset's calculated fields.

If the calculated field operates at the row level and the dataset is stored in SPICE, Quick Sight computes and materializes the result in SPICE. If the calculated field relies on an aggregation function, Quick Sight retains the formula and performs the calculation when the analysis is generated. This type of calculated field is called an unmaterialized calculated field.

To add or edit a calculated field for a dataset

1. Open the dataset that you want to work with. For more information, see [Editing datasets](#).
2. On the data prep page, do one of the following:
 - To create a new field, choose **Add calculated field** at left.
 - To edit an existing calculated field, choose it from **Calculated fields** at left, then choose **Edit** from the context (right-click) menu.
3. In the calculation editor, enter a descriptive name for **Add title** to name the new calculated field. This name appears in the field list in the dataset, so it should look similar to the other fields. For this example, we name the field `Total Sales This Year`.
4. (Optional) Add a comment, for example to explain what the expression does, by enclosing text in slashes and asterisks.

```
/* Calculates sales per year for this year*/
```

5. Identify the metrics, functions, and other items to use. For this example, we need to identify the following:
 - The metric to use

- Functions: `ifelse` and `dateDiff`

We want to build a statement like "If the sale happened during this year, show the total sales, and otherwise show 0."

To add the `ifelse` function, open the **Functions** list. Choose **All** to close the list of all functions. Now you should see the function groups: **Aggregate**, **Conditional**, **Date**, and so on.

Choose **Conditional**, and then double-click on `ifelse` to add it to the workspace.

```
ifelse()
```

6. Place your cursor inside the parenthesis in the workspace, and add three blank lines.

```
ifelse(  
  
  
  
)
```

7. With your cursor on the first blank line, find the `dateDiff` function. It's listed for **Functions** under **Dates**. You can also find it by entering **date** for **Search functions**. The `dateDiff` function returns all functions that have *date* as part of their name. It doesn't return all functions listed under **Dates**; for example, the `now` function is missing from the search results.

Double-click on `dateDiff` to add it to the first blank line of the `ifelse` statement.

```
ifelse(  
  dateDiff()  
  
  
)
```

Add the parameters that `dateDiff` uses. Place your cursor inside the `dateDiff` parentheses to begin to add `date1`, `date2`, and `period`:

1. For `date1`: The first parameter is the field that has the date in it. Find it under **Fields**, and add it to the workspace by double-clicking it or entering its name.

2. For date2, add a comma, then choose `truncate()` for **Functions**. Inside its parenthesis, add period and date, like this: `truncate( "YYYY", now() )`
3. For period: Add a comma after date2 and enter `YYYY`. This is the period for the year. To see a list of all the supported periods, find `dateDiff` in the **Functions** list, and open the documentation by choosing **Learn more**. If you're already viewing the documentation, as you are now, see [dateDiff](#).

Add a few spaces for readability, if you like. Your expression should look like the following.

```
ifelse(
  dateDiff( {Date}, truncate( "YYYY", now() ) ,"YYYY" )

)
```

8. Specify the return value. For our example, the first parameter in `ifelse` needs to return a value of `TRUE` or `FALSE`. Because we want the current year, and we're comparing it to this year, we specify that the `dateDiff` statement should return `0`. The `if` part of the `ifelse` evaluates as true for rows where there is no difference between the year of the sale and the current year.

```
dateDiff( {Date}, truncate( "YYYY", now() ) ,"YYYY" ) = 0
```

To create a field for `TotalSales` for last year, you can change `0` to `1`.

Another way to do the same thing is to use `addDateTime` instead of `truncate`. Then for each previous year, you change the first parameter for `addDateTime` to represent each year. For this, you use `-1` for last year, `-2` for the year before that, and so on. If you use `addDateTime`, you leave the `dateDiff` function = `0` for each year.

```
dateDiff( {Discharge Date}, addDateTime(-1, "YYYY", now() ) ,"YYYY" ) = 0 /*
Last year */
```

9. Move your cursor to the first blank line, just under `dateDiff`. Add a comma.

For the `then` part of the `ifelse` statement, we need to choose the measure (metric) that contains the sales amount, `TotalSales`.

To choose a field, open the **Fields** list and double-click a field to add it to the screen. Or you can enter the name. Add curly braces { } around names that contain spaces. It's likely that your metric has a different name. You can know which field is a metric by the number sign in front of it (#).

Your expression should look like the following now.

```
ifelse(  
  dateDiff( {Date}, truncDate( "YYYY", now() ) , "YYYY" ) = 0  
  , {TotalSales}  
)
```

10. Add an else clause. The `ifelse` function doesn't require one, but we want to add it. For reporting purposes, you usually don't want to have any null values, because sometimes rows with nulls are omitted.

We set the else part of the `ifelse` to 0. The result is that this field is 0 for rows that contain sales from previous years.

To do this, on the blank line add a comma and then a 0. If you added the comment at the beginning, your finished `ifelse` expression should look like the following.

```
/* Calculates sales per year for this year*/  
ifelse(  
  dateDiff( {Date}, truncDate( "YYYY", now() ) , "YYYY" ) = 0  
  , {TotalSales}  
  , 0  
)
```

11. Save your work by choosing **Save** at upper right.

If there are errors in your expression, the editor displays an error message at the bottom. Check your expression for a red squiggly line, then hover your cursor over that line to see what the error message is. Common errors include missing punctuation, missing parameters, misspellings, and invalid data types.

To avoid making any changes, choose **Cancel**.

To add a parameter value to a calculated field

1. You can reference parameters in calculated fields. By adding the parameter to your expression, you add the current value of that parameter.
2. To add a parameter, open the **Parameters** list, and select the parameter whose value you want to include.
3. (Optional) To manually add a parameter to the expression, type the name of the parameter. Then enclosed it in curly braces {}, and prefix it with a \$, for example `${parameterName}`.

You can change the data type of any field in your dataset, including the types of calculated fields. You can only choose data types that match the data that's in the field.

To change the data type of a calculated field

- For **Calculated fields** (at left), choose the field that you want to change, then choose **Change data type** from the context (right-click) menu.

Unlike the other fields in the dataset, calculated fields can't be disabled. Instead, delete them.

To delete a calculated field

- For **Calculated fields** (at left), choose the field that you want to change, then choose **Delete** from the context (right-click) menu.

Handling decimal values in calculated fields

When your dataset uses Direct Query mode, the calculation of the decimal data type is determined by the behavior of the source engine that the dataset originates from. In some particular cases, Quick Sight applies special handlings to determine the output calculation's data type.

When your dataset uses SPICE query mode and a calculated field is materialized, the data type of the result is contingent on the specific function operators and the data type of the input. The tables below show the expected behavior for some numeric calculated fields.

Unary operators

The following table shows which data type is output based on the operator you use and the data type of the value that you input. For example, if you input an integer to an abs calculation, the output value's data type is integer.

Operator	Input type	Output type
abs	Decimal-fixed	Decimal-fixed
	Int	Int
	Decimal-float	Decimal-float
ceil	Decimal-fixed	Int
	Int	Int
	Decimal-float	Int
exp	Decimal-fixed	Decimal-float
	Int	Decimal-float
	Decimal-float	Decimal-float
floor	Decimal-fixed	Int
	Int	Int
	Decimal-float	Int
ln	Decimal-fixed	Decimal-float
	Int	Decimal-float
	Decimal-float	Decimal-float
log	Decimal-fixed	Decimal-float
	Int	Decimal-float
	Decimal-float	Decimal-float
round	Decimal-fixed	Decimal-fixed
	Int	Decimal-fixed

Operator	Input type	Output type
	Decimal-float	Decimal-fixed
sqrt	Decimal-fixed	Decimal-float
	Int	Decimal-float
	Decimal-float	Decimal-float

Binary operators

The following tables show which data type is output based on the data types of the two values that you input. For example, for an arithmetic operator, if you provide two integer data types, the result of the calculation output as an integer.

For basic operators (+, -, *):

	Integer	Decimal-fixed	Decimal-float
Integer	Integer	Decimal-fixed	Decimal-float
Decimal-fixed	Decimal-fixed	Decimal-fixed	Decimal-float
Decimal-float	Decimal-float	Decimal-float	Decimal-float

For division operators (/):

	Integer	Decimal-fixed	Decimal-float
Integer	Decimal-float	Decimal-float	Decimal-float
Decimal-fixed	Decimal-float	Decimal-fixed	Decimal-float

	Integer	Decimal-fixed	Decimal-float
Decimal-float	Decimal-float	Decimal-float	Decimal-float

For exponential and mod operators (^, %):

	Integer	Decimal-fixed	Decimal-float
Integer	Decimal-float	Decimal-float	Decimal-float
Decimal-fixed	Decimal-float	Decimal-float	Decimal-float
Decimal-float	Decimal-float	Decimal-float	Decimal-float

Order of evaluation in Amazon Quick Sight

When you open or update an analysis, before displaying it Amazon Quick Sight evaluates everything that is configured in the analysis in a specific sequence. Amazon Quick Sight translates the configuration into a query that a database engine can run. The query returns the data in a similar way whether you connect to a database, a software as a service (SaaS) source, or the Amazon Quick Sight analytics engine ([SPICE](#)).

If you understand the order that the configuration is evaluated in, you know the sequence that dictates when a specific filter or calculation is applied to your data.

The following illustration shows the order of evaluation. The column on the left shows the order of evaluation when no level aware calculation window (LAC-W) nor aggregate (LAC-A) function is involved. The second column shows the order of evaluation for analyses that contain calculated fields to compute LAC-W expressions at the prefilter (PRE_FILTER) level. The third column shows the order of evaluation for analyses that contain calculated fields to compute LAC-W expressions at the preaggregate (PRE_AGG) level. The last column shows the order of evaluation for analyses that contain calculated fields to compute LAC-A expressions. Following the illustration, there is a more detailed explanation of the order of evaluation. For more information about level aware calculations, see [Using level-aware calculations in Quick Sight](#).

The following list shows the sequence in which Amazon Quick Sight applies the configuration in your analysis. Anything that's set up in your data set happens outside your analysis, for example calculations at the data set level, filters, and security settings. These all apply to the underlying data. The following list only covers what happens inside the analysis.

1. **LAC-W Prefilter level:** Evaluates the data at the original table cardinality before analysis filters
 - a. **Simple calculations:** Calculations at scalar level without any aggregations or window calculations. For example, `date_metric/60`, `parseDate(date, 'yyyy/MM/dd')`, `ifelse(metric > 0, metric, 0)`, `split(string_column, '|' 0)`.
 - b. **LAC-W function PRE_FILTER:** If any LAC-W PRE_FILTER expression is involved in the visual, Amazon Quick Sight firstly computes the window function at the original table level, before any filters. If the LAC-W PRE_FILTER expression is used in filters, it is applied at this point. For example, `maxOver(Population, [State, County], PRE_FILTER) > 1000`.
2. **LAC-W PRE_AGG:** Evaluates the data at the original table cardinality before aggregations
 - a. **Filters added during analysis:** Filters created for un-aggregated fields in the visuals are applied at this point, which are similar to WHERE clauses. For example, `year > 2020`.
 - b. **LAC-W function PRE_AGG:** If any LAC-W PRE_AGG expression is involved in the visual, Amazon Quick Sight computes the window function before any aggregation applied. If the LAC-W PRE_AGG expression is used in filters, it is applied at this point. For example, `maxOver(Population, [State, County], PRE_AGG) > 1000`.
 - c. **Top/bottom N filters:** Filters that are configured on dimensions to display top/bottom N items.
3. **LAC-A level:** Evaluate aggregations at customized level, before visual aggregations
 - a. **Custom-level aggregations:** If any LAC-A expression is involved in the visual, it is calculated at this point. Based on the table after the filters mentioned above, Amazon QuickSight computes the aggregation, grouped by the dimensions that are specified in the calculated fields. For example, `max(Sales, [Region])`.
4. **Visual level:** Evaluates aggregations at visual level, and post-aggregation table calculations, with the remaining configurations applied in the visuals
 - a. **Visual-level aggregations:** Visual aggregations should always be applied except for tabular tables (where dimension is empty). With this setting, aggregations based on the fields in the field wells are calculated, grouped by the dimensions that put into the visuals. If any

filter is built on top of aggregations, it is applied at this point, similar to HAVING clauses. For example, `min(distance) > 100`.

- b. **Table calculations:** If there is any post-aggregation table calculation (it should take aggregated expression as operand) referenced in the visual, it is calculated at this point. Amazon Quick Sight performs window calculations after visual aggregations. Similarly, filters built on such calculations are applied.
- c. **Other category calculations:** This type of calculation only exists in line/bar/pie/donut charts. For more information, see [Display limits](#).
- d. **Totals and subtotals:** Totals and Subtotals are calculated in donut charts (only totals), tables (only totals) and pivot tables, if requested.

Using level-aware calculations in Quick Sight

Applies to: Enterprise Edition and Standard Edition

With *Level-aware calculations* (LAC) you can specify the level of granularity that you want to compute window functions or aggregate functions. There are two types of LAC functions: level-aware calculation - aggregate (LAC-A) functions, and level-aware calculation - window (LAC-W) functions.

Topics

- [Level-aware calculation - aggregate \(LAC-A\) functions](#)
- [Level-aware calculation - window \(LAC-W\) functions](#)

Level-aware calculation - aggregate (LAC-A) functions

With LAC-A functions, you can specify at what level to group the computation. By adding one argument into an existing aggregate function, such as `sum()` , `max()` , `count()`, you can define any group-by level that you want for the aggregation. The level added can be any dimension independent of the dimensions added to the visual. For example:

```
sum(measure, [group_field_A])
```

To use LAC-A functions, type them directly in the calculation editor by adding the intended aggregation levels as the second argument between brackets. Following is an example of an aggregate function and a LAC-A function, for comparison.

- Aggregate function: `sum({sales})`
- LAC-A function: `sum({sales}, [{Country},{Product}])`

The LAC-A results are computed with the specified level in the brackets [], can be used as operand of an aggregate function. The group-by level of the aggregate function is visual level, with **Group by** fields added to the field well of the visual.

In addition to creating a static LAC group key in the bracket [], you can make it dynamically adapted to visual group-by fields, by putting a parameter `$visualDimensions` in the bracket. This is a system-provided parameter, in contrast to user-defined parameter. The `[$visualDimensions]` parameter represents the fields added to the **Group by** field well in current visual. The following examples show how to dynamically add group keys to the visual dimensions or remove group keys from visual dimensions

- LAC-A with dynamic-added group key: `sum({sales}, [${visualDimensions}, {Country},{Products}])`

It calculates, before the visual level aggregation is calculated, the sum of sales, grouping by country, products, and any other fields in the **Group by** field well .

- LAC-A with dynamic-removed group key: `sum({sales}, [${visualDimensions}, !{Country},!{Products}])`

It calculates, before visual level aggregation is calculated, the sum of sales, grouping by the fields in the visual's **Group by** field well, except country and product.

You can specify added group key or removed group key in on LAC expression, but not both.

LAC-A functions are supported for the following aggregate functions:

- [avg](#)
- [count](#)
- [distinct_count](#)
- [max](#)

- [median](#)
- [min](#)
- [percentile](#)
- [percentileCont](#)
- [percentileDisc \(percentile\)](#)
- [stdev](#)
- [stdevp](#)
- [sum](#)
- [var](#)
- [varp](#)

LAC-A examples

You can do the following with LAC-A functions:

- Run calculations that are independent of the levels in the visual. For example, if you have the following calculation, the sales numbers are aggregated only at the country level, but not across other dimensions (Region or Product) in the visual.

```
sum({Sales}, [{Country}])
```

- Run calculations for the dimensions that are not in the visual. For example, if you have the following function, you can calculate the average total country sales by region.

```
sum({Sales}, [{Country}])
```

Though Country is not included in the visual, the LAC-A function first aggregates the sales at the Country level and then the visual level calculation generates the average number for each region. If the LAC-A function isn't used to specify the level, the average sales are calculated at the lowest granular level (the base level of the dataset) for each region (showing in the sales column).

- Use LAC-A combined with other aggregate functions and LAC-W functions. There are two ways you can nest LAC-A functions with other functions.

- You can write a nested syntax when you create a calculation. For example, the LAC-A function can be nested with a LAC-W function to calculate the total sales by country of each product's average price:

```
sum(avgOver({Sales}, [{Product}], PRE_AGG), [{Country}])
```

- When adding a LAC-A function into a visual, the calculation can be further nested with visual-level aggregate functions that you selected in the fields well. For more information about changing the aggregation of fields in the visual, see [Changing or adding aggregation to a field by using a field well](#).

LAC-A limitations

The following limitations apply to LAC-A functions:

- LAC-A functions are supported for all additive and non-additive aggregate functions, such as `sum()`, `count()`, and `percentile()`. LAC-A functions are not supported for conditional aggregate functions that end with "if", such as `sumif()` and `countif()`, nor for period aggregate functions that start with "periodToDate", such as `periodToDateSum()` and `periodToDateMax()`.
- Row-level and column-level totals are not currently supported for LAC-A functions in tables and pivot tables. When you add row-level or column-level totals to the chart, the total number will show as blank. Other non-LAC dimensions are not impacted.
- Nested LAC-A functions are not currently supported. A limited capability of LAC-A functions nested with regular aggregate functions and LAC-W functions are supported.

For example, the following functions are valid:

- Aggregation(LAC-A()). For example: `max(sum({sales}, [{country}]))`
- LAC-A(LAC-W()). For example: `sum(sumOver({Sales}, [{Product}], PRE_AGG), [{Country}])`

The following functions are not valid:

- LAC-A(Aggregation()). For example: `sum(max({sales}), [{country}])`
- LAC-A(LAC-A()). For example: `sum(max({sales}, [{country}]), [category])`
- LAC-W(LAC-A()). For example: `sumOver(sum({Sales}, [{Product}]), [{Country}], PRE_AGG)`

Level-aware calculation - window (LAC-W) functions

With LAC-W functions, you can specify the window or partition to compute the calculation. LAC-W functions are a group of window functions, such as `sumover()`, `(maxover)`, `denseRank`, that you can run at the prefilter or preaggregate level. For example: `sumOver(measure, [partition_field_A], pre_agg)`.

LAC-W functions used to be called level aware aggregations (LAA).

LAC-W functions help you to answer the following types of questions:

- How many of my customers made only 1 purchase order? Or 10? Or 50? We want the visual to use the count as a dimension rather than a metric in the visual.
- What are the total sales per market segment for customers whose lifetime spend is greater than \$100,000? The visual should only show the market segment and the total sales for each.
- How much is the contribution of each industry to the entire company's profit (percent of total)? We want to be able to filter the visual to show some of the industries, and how they contribute to the total sales for the displayed industries. However, we also want to see each industry's percent of total sales for the entire company (including industries that are filtered out).
- What are the total sales in each category as compared to the industry average? The industry average should include all of the categories, even after filtering.
- How are my customers grouped into cumulative spending ranges? We want to use the grouping as a dimension rather than a metric.

For more complex questions, you can inject a calculation or filter before Quick Sight gets to a specific point in its evaluation of your settings. To directly influence your results, you add a calculation level keyword to a table calculation. For more information on how Quick Sight evaluates queries, see [Order of evaluation in Amazon Quick Sight](#).

The following calculation levels are supported for LAC-W functions:

- **PRE_FILTER** – Before applying filters from the analysis, Quick Sight evaluates prefilter calculations. Then it applies any filters that are configured on these prefilter calculations.
- **PRE_AGG** – Before computing display-level aggregations, Quick Sight performs preaggregate calculations. Then it applies any filters that are configured on these preaggregate calculations. This work happens before applying top and bottom *N* filters.

You can use the PRE_FILTER or PRE_AGG keyword as a parameter in the following table calculation functions. When you specify a calculation level, you use an unaggregated measure in the function. For example, you can use `countOver({ORDER ID}, [{Customer ID}], PRE_AGG)`. By using PRE_AGG, you specify that the `countOver` executes at the preaggregate level.

- [avgOver](#)
- [countOver](#)
- [denseRank](#)
- [distinctCountOver](#)
- [minOver](#)
- [maxOver](#)
- [percentileRank](#)
- [rank](#)
- [stdevOver](#)
- [stdevpOver](#)
- [sumOver](#)
- [varOver](#)
- [varpOver](#)

By default, the first parameter for each function must be an aggregated measure. If you use either PRE_FILTER or PRE_AGG, you use a nonaggregated measure for the first parameter.

For LAC-W functions, the visual aggregation defaults to MIN to eliminate duplicates. To change the aggregation, open the field's context (right-click) menu, and then choose a different aggregation.

For examples of when and how to use LAC-W functions in real life scenarios, see the following post in the Amazon Big Data Blog: [Create advanced insights using Level Aware Aggregations in Amazon QuickSight](#).

Calculated field function and operator reference for Amazon Quick Suite

You can add calculated fields to a dataset during data preparation or from the analysis page. When you add a calculated field to a dataset during data preparation, it's available to all analyses that use that dataset. When you add a calculated field to a dataset in an analysis, it's available only in that analysis.

You can create calculated fields to transform your data by using the following functions and operators.

Topics

- [Operators](#)
- [Functions by category](#)
- [Functions](#)
- [Aggregate functions](#)
- [Table calculation functions](#)

Operators

You can use the following operators in calculated fields. Quick Suite uses the standard order of operations: parentheses, exponents, multiplication, division, addition, subtraction (PEMDAS). Equal (=) and not equal (<>) comparisons are case-sensitive.

- Addition (+)
- Subtraction (-)
- Multiplication (*)
- Division (/)
- Modulo (%) – See also `mod( )` in the following list.
- Power (^) – See also `exp( )` in the following list.
- Equal (=)
- Not equal (<>)
- Greater than (>)
- Greater than or equal to (>=)
- Less than (<)
- Less than or equal to (<=)
- AND
- OR
- NOT

Amazon Quick Suite supports applying the following mathematical functions to an expression.

- [Mod](#)(*number*, *divisor*) – Finds the remainder after dividing a number by a divisor.
- [Log](#)(*expression*) – Returns the base 10 logarithm of a given expression.
- [Ln](#)(*expression*) – Returns the natural logarithm of a given expression.
- [Abs](#)(*expression*) – Returns the absolute value of a given expression.
- [Sqrt](#)(*expression*) – Returns the square root of a given expression.
- [Exp](#)(*expression*) – Returns the base of natural log e raised to the power of a given expression.

To make lengthy calculations easier to read, you can use parenthesis to clarify groupings and precedence in calculations. In the following statement, you don't need parentheses. The multiplication statement is processed first, and then the result is added to five, returning a value of 26. However, parentheses make the statement easier to read and thus maintain.

```
5 + (7 * 3)
```

Because parenthesis are first in the order of operations, you can change the order in which other operators are applied. For example, in the following statement the addition statement is processed first, and then the result is multiplied by three, returning a value of 36.

```
(5 + 7) * 3
```

Example: Arithmetic operators

The following example uses multiple arithmetic operators to determine a sales total after discount.

```
(Quantity * Amount) - Discount
```

Example: (/) Division

The following example uses division to divide 3 by 2. A value of 1.5 is returned. Amazon Quick Suite uses floating point divisions.

```
3/2
```

Example: (=) equal

Using = performs a case-sensitive comparison of values. Rows where the comparison is TRUE are included in the result set.

In the following example, rows where the Region field is **South** are included in the results. If the Region is **south**, these rows are excluded.

```
Region = 'South'
```

In the following example, the comparison evaluates to FALSE.

```
Region = 'south'
```

The following example shows a comparison that converts Region to all uppercase (**SOUTH**), and compares it to **SOUTH**. This returns rows where the region is **south**, **South**, or **SOUTH**.

```
toUpper(Region) = 'SOUTH'
```

Example: (<>)

The not equal symbol <> means *less than or greater than*.

So, if we say **x<>1**, then we are saying *if x is less than 1 OR if x is greater than 1*. Both < and > are evaluated together. In other words, *if x is any value except 1*. Or, *x is not equal to 1*.

Note

Use <>, not !=.

The following example compares Status Code to a numeric value. This returns rows where the Status Code is not equal to **1**.

```
statusCode <> 1
```

The following example compares multiple statusCode values. In this case, active records have activeFlag = 1. This example returns rows where one of the following applies:

- For active records, show rows where the status isn't 1 or 2
- For inactive records, show rows where the status is 99 or -1

```
( activeFlag = 1 AND (statusCode <> 1 AND statusCode <> 2) )  
OR  
( activeFlag = 0 AND (statusCode= 99 OR statusCode= -1) )
```

Example: (^)

The power symbol ^ means *to the power of*. You can use the power operator with any numeric field, with any valid exponent.

The following example is a simple expression of 2 to the power of 4 or (2 * 2 * 2 * 2). This returns a value of 16.

```
2^4
```

The following example computes the square root of the revenue field.

```
revenue^0.5
```

Example: AND, OR, and NOT

The following example uses AND, OR, and NOT to compare multiple expressions. It does so using conditional operators to tag top customers NOT in Washington or Oregon with a special promotion, who made more than 10 orders. If no values are returned, the value 'n/a' is used.

```
ifelse(( NOT (State = 'WA' OR State = 'OR')) AND Orders > 10), 'Special Promotion  
XYZ', 'n/a')
```

Example: Creating comparison lists like "in" or "not in"

This example uses operators to create a comparison to find values that exist, or don't exist, in a specified list of values.

The following example compares promoCode a specified list of values. This example returns rows where the promoCode is in the list (1, 2, 3).

```
promoCode = 1  
OR promoCode = 2  
OR promoCode = 3
```

The following example compares promoCode a specified list of values. This example returns rows where the promoCode is NOT in the list (**1, 2, 3**).

```
NOT(promoCode = 1  
OR promoCode = 2  
OR promoCode = 3  
)
```

Another way to express this is to provide a list where the promoCode is not equal to any items in the list.

```
promoCode <> 1  
AND promoCode <> 2  
AND promoCode <> 3
```

Example: Creating a "between" comparison

This example uses comparison operators to create a comparison showing values that exist between one value and another.

The following example examines OrderDate and returns rows where the OrderDate is between the first day and last day of 2016. In this case, we want the first and last day included, so we use "or equal to" on the comparison operators.

```
OrderDate >= "1/1/2016" AND OrderDate <= "12/31/2016"
```

Functions by category

In this section, you can find a list of the functions available in Amazon Quick Suite, sorted by category.

Topics

- [Aggregate functions](#)
- [Conditional functions](#)
- [Date functions](#)
- [Numeric functions](#)
- [Mathematical functions](#)

- [String functions](#)
- [Table calculations](#)

Aggregate functions

The aggregate functions for calculated fields in Amazon Quick Suite include the following. These are only available during analysis and visualization. Each of these functions returns values grouped by the chosen dimension or dimensions. For each aggregation, there is also a conditional aggregation. These perform the same type of aggregation, based on a condition.

- [avg](#) averages the set of numbers in the specified measure, grouped by the chosen dimension or dimensions.
- [avgIf](#) calculates the average based on a conditional statement.
- [count](#) calculates the number of values in a dimension or measure, grouped by the chosen dimension or dimensions.
- [countIf](#) calculates the count based on a conditional statement.
- [distinct_count](#) calculates the number of distinct values in a dimension or measure, grouped by the chosen dimension or dimensions.
- [distinct_countIf](#) calculates the distinct count based on a conditional statement.
- [max](#) returns the maximum value of the specified measure, grouped by the chosen dimension or dimensions.
- [maxIf](#) calculates the maximum based on a conditional statement.
- [median](#) returns the median value of the specified measure, grouped by the chosen dimension or dimensions.
- [medianIf](#) calculates the median based on a conditional statement.
- [min](#) returns the minimum value of the specified measure, grouped by the chosen dimension or dimensions.
- [minIf](#) calculates the minimum based on a conditional statement.
- [percentile](#) (alias of `percentileDisc`) computes the n th percentile of the specified measure, grouped by the chosen dimension or dimensions.
- [percentileCont](#) calculates the n th percentile based on a continuous distribution of the numbers of the specified measure, grouped by the chosen dimension or dimensions.
- [percentileDisc \(percentile\)](#) calculates the n th percentile based on the actual numbers of the specified measure, grouped by the chosen dimension or dimensions.

- [periodToDateAvg](#) averages the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateCount](#) calculates the number of values in a dimension or measure for a given time granularity (for instance, Quarter) up to a point in time including duplicates.
- [periodToDateMax](#) returns the maximum value of the specified measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateMedian](#) returns the median value of the specified measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateMin](#) returns the minimum value of the specified measure or date for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDatePercentile](#) calculates the percentile based on the actual numbers in measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDatePercentileCont](#) calculates percentile based on a continuous distribution of the numbers in the measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateStDev](#) calculates the standard deviation of the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time based on a sample.
- [periodToDateStDevP](#) calculates the population standard deviation of the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time based on a sample.
- [periodToDateSum](#) adds the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateVar](#) calculates the sample variance of the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateVarP](#) calculates the population variance of the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time.
- [stdev](#)) calculates the standard deviation of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a sample.
- [stdevIf](#) calculates the sample standard deviation based on a conditional statement.
- [stdevp](#) calculates the standard deviation of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a biased population.
- [stdevpIf](#) calculates the population deviation based on a conditional statement.

- [var](#)) calculates the variance of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a sample.
- [varif](#) calculates the sample variance based on a conditional statement.
- [varp](#)) calculates the variance of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a biased population.
- [varplf](#) calculates the population variance based on a conditional statement.
- [sum](#)) adds the set of numbers in the specified measure, grouped by the chosen dimension or dimensions.
- [sumif](#)) calculates the sum based on a conditional statement.

Conditional functions

The conditional functions for calculated fields in Amazon Quick Suite include the following:

- [Coalesce](#) returns the value of the first argument that is not null.
- [Ifelse](#) evaluates a set of *if, then* expression pairings, and returns the value of the *then* argument for the first *if* argument that evaluates to true.
- [in](#) evaluates an expression to see if it is in a given list of values.
- [isNotNull](#) evaluates an expression to see if it is not null.
- [isNull](#) evaluates an expression to see if it is null. If the expression is null, `isNull` returns true, and otherwise it returns false.
- [notin](#) evaluates an expression to see if it is not in a given list of values.
- [nullif](#) compares two expressions. If they are equal, the function returns null. If they are not equal, the function returns the first expression.
- [switch](#) returns an expression that matches the first label equal to the condition expression.

Date functions

The date functions for calculated fields in Amazon Quick Suite include the following:

- [addDateTime](#) adds or subtracts a unit of time to the date or time provided.
- [addWorkDays](#) adds or subtracts the given number of work days to the date or time provided.
- [dateDiff](#) returns the difference in days between two date fields.
- [epochDate](#) converts an epoch date into a standard date.

- [Extract](#) returns a specified portion of a date value.
- [formatDate](#) formats a date using a pattern you specify.
- [isWorkDay](#) returns TRUE if a given date-time value is a work or business day.
- [netWorkDays](#) returns the number of working days between the provided two date values.
- [Now](#) returns the current date and time, using either settings for a database, or UTC for file and Salesforce.
- [truncDate](#) returns a date value that represents a specified portion of a date.

Numeric functions

The numeric functions for calculated fields in Amazon Quick Suite include the following:

- [Ceil](#) rounds a decimal value to the next highest integer.
- [decimalToInt](#) converts a decimal value to an integer.
- [Floor](#) decrements a decimal value to the next lowest integer.
- [intToDecimal](#) converts an integer value to a decimal.
- [Round](#) rounds a decimal value to the closest integer or, if scale is specified, to the closest decimal place.

Mathematical functions

The mathematical functions for calculated fields in Amazon Quick Suite include the following:

- [Mod](#)(*number*, *divisor*) – Finds the remainder after dividing a number by a divisor.
- [Log](#)(*expression*) – Returns the base 10 logarithm of a given expression.
- [Ln](#)(*expression*) – Returns the natural logarithm of a given expression.
- [Abs](#)(*expression*) – Returns the absolute value of a given expression.
- [Sqrt](#)(*expression*) – Returns the square root of a given expression.
- [Exp](#)(*expression*) – Returns the base of natural log e raised to the power of a given expression.

String functions

The string (text) functions for calculated fields in Amazon Quick Suite include the following:

- [Concat](#) concatenates two or more strings.
- [contains](#) checks if an expression contains a substring.
- [endsWith](#) checks if the expression ends with the substring specified.
- [Left](#) returns the specified number of leftmost characters from a string.
- [Locate](#) locates a substring within another string, and returns the number of characters before the substring.
- [Ltrim](#) removes preceding blank space from a string.
- [parseDate](#) parses a string to determine if it contains a date value, and returns the date if found.
- [parseDecimal](#) parses a string to determine if it contains a decimal value.
- [parseInt](#) parses a string to determine if it contains an integer value.
- [parseJson](#) parses values from a native JSON or from a JSON object in a text field.
- [Replace](#) replaces part of a string with a new string.
- [Right](#) returns the specified number of rightmost characters from a string.
- [Rtrim](#) removes following blank space from a string.
- [Split](#) splits a string into an array of substrings, based on a delimiter that you choose, and returns the item specified by the position.
- [startsWith](#) checks if the expression starts with the substring specified.
- [Strlen](#) returns the number of characters in a string.
- [Substring](#) returns the specified number of characters in a string, starting at the specified location.
- [toLowerCase](#) formats a string in all lowercase.
- [toString](#) formats the input expression as a string.
- [toUpperCase](#) formats a string in all uppercase.
- [trim](#) removes both preceding and following blank space from a string.

Table calculations

Table calculations form a group of functions that provide context in an analysis. They provide support for enriched aggregated analysis. By using these calculations, you can address common business scenarios such as calculating percentage of total, running sum, difference, common baseline, and rank.

When you are analyzing data in a specific visual, you can apply table calculations to the current set of data to discover how dimensions influence measures or each other. Visualized data is your result set based on your current dataset, with all the filters, field selections, and customizations applied. To see exactly what this result set is, you can export your visual to a file. A table calculation function performs operations on the data to reveal relationships between fields.

Lookup-based functions

- [difference](#) calculates the difference between a measure based on one set of partitions and sorts, and a measure based on another.
- [lag](#) calculates the lag (previous) value for a measure.
- [lead](#) calculates the lead (following) value for a measure.
- [percentDifference](#) calculates the percentage difference between the current value and a comparison value.

Over functions

- [avgOver](#) calculates the average of a measure over one or more dimensions.
- [countOver](#) calculates the count of a field over one or more dimensions.
- [distinctCountOver](#) calculates the distinct count of the operand partitioned by the specified attributes at a specified level.
- [maxOver](#) calculates the maximum of a measure over one or more dimensions.
- [minOver](#) the minimum of a measure over one or more dimensions.
- [percentileOver](#) (alias of `percentileDiscOver`) calculates the n th percentile of a measure partitioned by a list of dimensions.
- [percentileContOver](#) calculates the n th percentile based on a continuous distribution of the numbers of a measure partitioned by a list of dimensions.
- [percentileDiscOver](#) calculates the n th percentile based on the actual numbers of a measure partitioned by a list of dimensions.
- [percentOfTotal](#) calculates the percentage that a measure contributes to the total.
- [periodOverPeriodDifference](#) calculates the difference of a measure over two different time periods as specified by period granularity and offset.
- [periodOverPeriodLastValue](#) calculates the last (previous) value of a measure from a previous time period as specified by period granularity and offset.

- [periodOverPeriodPercentDifference](#) calculates the percent difference of a measure over two different time periods as specified by period granularity and offset.
- [periodToDateAvgOverTime](#) calculates the average of a measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateCountOverTime](#) calculates the count of a dimension or measure for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateMaxOverTime](#) calculates the maximum of a measure or date for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateMinOverTime](#) calculates the minimum of a measure or date for a given time granularity (for instance, a quarter) up to a point in time.
- [periodToDateSumOverTime](#) calculates the sum of a measure for a given time granularity (for instance, a quarter) up to a point in time.
- [sumOver](#) calculates the sum of a measure over one or more dimensions.
- [stdevOver](#) calculates the standard deviation of the specified measure, partitioned by the chosen attribute or attributes, based on a sample.
- [stdevpOver](#) calculates the standard deviation of the specified measure, partitioned by the chosen attribute or attributes, based on a biased population.
- [varOver](#) calculates the variance of the specified measure, partitioned by the chosen attribute or attributes, based on a sample.
- [varpOver](#) calculates the variance of the specified measure, partitioned by the chosen attribute or attributes, based on a biased population.

Ranking functions

- [rank](#) calculates the rank of a measure or a dimension.
- [denseRank](#) calculates the rank of a measure or a dimension, ignoring duplicates.
- [percentileRank](#) calculates the rank of a measure or a dimension, based on percentile.

Running functions

- [runningAvg](#) calculates a running average for a measure.
- [runningCount](#) calculates a running count for a measure.
- [runningMax](#) calculates a running maximum for a measure.

- [runningMin](#) calculates a running minimum for a measure.
- [runningSum](#) calculates a running sum for a measure.

Window functions

- [firstValue](#) calculates the first value of the aggregated measure or dimension partitioned and sorted by specified attributes.
- [lastValue](#) calculates the last value of the aggregated measure or dimension partitioned and sorted by specified attributes.
- [windowAvg](#) calculates the average of the aggregated measure in a custom window that is partitioned and sorted by specified attributes.
- [windowCount](#) calculates the count of the aggregated measure in a custom window that is partitioned and sorted by specified attributes.
- [windowMax](#) calculates the maximum of the aggregated measure in a custom window that is partitioned and sorted by specified attributes.
- [windowMin](#) calculates the minimum of the aggregated measure in a custom window that is partitioned and sorted by specified attributes.
- [windowSum](#) calculates the sum of the aggregated measure in a custom window that is partitioned and sorted by specified attributes.

Functions

In this section, you can find a list of functions available in Amazon Quick Suite. To view a list of functions sorted by category, with brief definitions, see [Functions by category](#).

Topics

- [addDateTime](#)
- [addWorkDays](#)
- [Abs](#)
- [Ceil](#)
- [Coalesce](#)
- [Concat](#)
- [contains](#)
- [decimalToInt](#)

- [dateDiff](#)
- [endsWith](#)
- [epochDate](#)
- [Exp](#)
- [Extract](#)
- [Floor](#)
- [formatDate](#)
- [Ifelse](#)
- [in](#)
- [intToDecimal](#)
- [isNotNull](#)
- [isNull](#)
- [isWorkDay](#)
- [Left](#)
- [Locate](#)
- [Log](#)
- [Ln](#)
- [Ltrim](#)
- [Mod](#)
- [netWorkDays](#)
- [Now](#)
- [notIn](#)
- [nullIf](#)
- [parseDate](#)
- [parseDecimal](#)
- [parseInt](#)
- [parseJson](#)
- [Replace](#)
- [Right](#)

- [Round](#)
- [Rtrim](#)
- [Split](#)
- [Sqrt](#)
- [startsWith](#)
- [Strlen](#)
- [Substring](#)
- [switch](#)
- [toLowerCase](#)
- [toString](#)
- [toUpperCase](#)
- [trim](#)
- [truncDate](#)

addDateTime

addDateTime adds or subtracts a unit of time from a datetime value. For example, addDateTime(2, 'YYYY', parseDate('02-JUL-2018', 'dd-MMM-yyyy')) returns 02-JUL-2020. You can use this function to perform date math on your date and time data.

Syntax

```
addDateTime(amount, period, datetime)
```

Arguments

amount

A positive or negative integer value that represents the amount of time that you want to add or subtract from the provided datetime field.

period

A positive or negative value that represents the amount of time that you want to add or subtract from the provided datetime field. Valid periods are as follows:

- **YYYY:** This returns the year portion of the date.
- **Q:** This returns the quarter that the date belongs to (1–4).
- **MM:** This returns the month portion of the date.
- **DD:** This returns the day portion of the date.
- **WK:** This returns the week portion of the date. The week starts on Sunday in Amazon Quick Suite.
- **HH:** This returns the hour portion of the date.
- **MI:** This returns the minute portion of the date.
- **SS:** This returns the second portion of the date.
- **MS:** This returns the millisecond portion of the date.

datetime

The date or time that you want to perform date math on.

Return type

Datetime

Example

Let's say you have a field called `purchase_date` that has the following values.

```
2018 May 13 13:24
2017 Jan 31 23:06
2016 Dec 28 06:45
```

Using the following calculations, `addDateTime` modifies the values as shown following.

```
addDateTime(-2, 'YYYY', purchaseDate)
```

```
2016 May 13 13:24
2015 Jan 31 23:06
2014 Dec 28 06:45
```

```
addDateTime(4, 'DD', purchaseDate)
```

```
2018 May 17 13:24
2017 Feb 4 23:06
2017 Jan 1 06:45
```

```
addDateTime(20, 'MI', purchaseDate)
```

```
2018 May 13 13:44
2017 Jan 31 23:26
2016 Dec 28 07:05
```

addWorkDays

addWorkDays Adds or subtracts a designated number of work days to a given date value. The function returns a date for a work day, that falls a designated work days after or before a given input date value.

Syntax

```
addWorkDays(initDate, numWorkDays)
```

Arguments

initDate

A valid non-NULL date that acts as the start date for the calculation.

- **Dataset field** – Any date field from the dataset that you are adding this function to.
- **Date function** – Any date output from another date function, for example `parseDate`, `epochDate`, `addDateTime`., and so on.

Example

```
addWorkDays(epochDate(1659484800), numWorkDays)
```

- **Calculated fields** – Any Quick Suite calculated field that returns a date value.

Example

```
calcFieldStartDate = addDateTime(10, "DD", startDate)
addWorkDays(calcFieldStartDate, numWorkDays)
```

- **Parameters** – Any Quick Suite datetime parameter.

Example

```
addWorkDays($paramStartDate, numWorkDays)
```

- Any combination of the above stated argument values.

numWorkDays

A non-NULL integer that acts as the end date for the calculation.

- **Literal** – An integer literal directly typed in the expression editor.

Example

- **Dataset field** – Any date field from the dataset

Example

- **Scalar function or calculation** – Any scalar Quick Suite function that returns an integer output from another, for example `decimalToInt`, `abs`, and so on.

Example

```
addWorkDays(initDate, decimalToInt(sqrt (abs(numWorkDays)) ) )
```

- **Calculated field** – Any Quick Suite calculated field that returns a date value.

Example

```
someOtherIntegerCalcField = (num_days * 2) + 12
addWorkDays(initDate, someOtherIntegerCalcField)
```

- **Parameter** – Any Quick Suite datetime parameter.

Example

```
addWorkDays(initDate, $param_numWorkDays)
```

- Any combination of the above stated argument values.

Return type

Integer

Output values

Expected output values include:

- Positive integer (when `start_date < end_date`)
- Negative integer (when `start_date > end_date`)
- NULL when one or both of the arguments get a null value from the `dataset` field.

Input errors

Disallowed argument values cause errors, as shown in the following examples.

- Using a literal NULL as an argument in the expression is disallowed.

Example

```
addWorkDays(NULL, numWorkDays)
```

Example

Error

At least one of the arguments in this function does not have correct type.
Correct the expression and choose Create again.

- Using a string literal as an argument, or any other data type other than a date, in the expression is disallowed. In the following example, the string **"2022-08-10"** looks like a date, but it is actually a string. To use it, you would have to use a function that converts to a date data type.

Example

```
addWorkDays("2022-08-10", 10)
```

Example

Error

Expression `addWorkDays("2022-08-10", numWorkDays)` for function `addWorkDays` has

```
incorrect argument type addWorkDays(String, Number).  
Function syntax expects Date, Integer.
```

Example

A positive integer as numWorkDays argument will yield a date in the future of the input date. A negative integer as numWorkDays argument will yield a resultant date in the past of the input date. A zero value for the numWorkDays argument yields the same value as input date whether or not it falls on a work day or a weekend.

The addWorkDays function operates at the granularity: DAY. Accuracy cannot be preserved at any granularity which is lower or higher than DAY level.

```
addWorkDays(startDate, endDate)
```

Let's assume there is a field named employmentStartDate with the following values:

```
2022-08-10 2022-08-06 2022-08-07
```

Using the above field and following calculations, addWorkDays returns the modified values as shown below:

```
addWorkDays(employmentStartDate, 7)
```

```
2022-08-19  
2022-08-16  
2022-08-16
```

```
addWorkDays(employmentStartDate, -5)
```

```
2022-08-02  
2022-08-01  
2022-08-03
```

```
addWorkDays(employmentStartDate, 0)
```

```
2022-08-10  
2022-08-06  
2022-08-07
```

The following example calculates the total pro-rated bonus to be paid to each employee for 2 years based on how many days each employee has actually worked.

```
last_day_of_work = addWorkDays(employment_start_date, 730)
total_days_worked = netWorkDays(employment_start_date, last_day_of_work)
total_bonus = total_days_worked * bonus_per_day
```

Abs

abs returns the absolute value of a given expression.

Syntax

```
abs(expression)
```

Arguments

expression

The expression must be numeric. It can be a field name, a literal value, or another function.

Ceil

ceil rounds a decimal value to the next highest integer. For example, ceil(29.02) returns 30.

Syntax

```
ceil(decimal)
```

Arguments

decimal

A field that uses the decimal data type, a literal value like **17.62**, or a call to another function that outputs a decimal.

Return type

Integer

Example

The following example rounds a decimal field to the next highest integer.

```
ceil(salesAmount)
```

The following are the given field values.

```
20.13  
892.03  
57.54
```

For these field values, the following values are returned.

```
21  
893  
58
```

Coalesce

`coalesce` returns the value of the first argument that is not null. When a non-null value is found, the remaining arguments in the list are not evaluated. If all arguments are null, the result is null. 0-length strings are valid values and are not considered equivalent to null.

Syntax

```
coalesce(expression1, expression2 [, expression3, ...])
```

Arguments

`coalesce` takes two or more expressions as arguments. All of the expressions must have the same data type or be able to be implicitly cast to the same data type.

expression

The expression can be numeric, datetime, or string. It can be a field name, a literal value, or another function.

Return type

`coalesce` returns a value of the same data type as the input arguments.

Example

The following example retrieves a customer's billing address if it exists, her street address if there is no billing address, or returns "No address listed" if neither address is available.

```
coalesce(billingAddress, streetAddress, 'No address listed')
```

Concat

concat concatenates two or more strings.

Syntax

```
concat(expression1, expression2 [, expression3 ...])
```

Arguments

concat takes two or more string expressions as arguments.

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

Return type

String

Examples

The following example concatenates three string fields and adds appropriate spacing.

```
concat(salutation, ' ', firstName, ' ', lastName)
```

The following are the given field values.

salutation	firstName	lastName
Ms.	Li	Juan
Dr.	Ana Carolina	Silva
Mr.	Nikhil	Jayashankar

For these field values, the following values are returned.

```
Ms. Li Juan  
Dr. Ana Carolina Silva  
Mr. Nikhil Jayashankar
```

The following example concatenates two string literals.

```
concat('Hello', 'world')
```

The following value is returned.

```
Helloworld
```

contains

`contains` evaluates if the substring that you specify exists within an expression. If the expression contains the substring, `contains` returns true, and otherwise it returns false.

Syntax

```
contains(expression, substring, string-comparison-mode)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

substring

The set of characters to check against the *expression*. The substring can occur one or more times in the *expression*.

string-comparison-mode

(Optional) Specifies the string comparison mode to use:

- `CASE_SENSITIVE` – String comparisons are case-sensitive.
- `CASE_INSENSITIVE` – String comparisons are case-insensitive.

This value defaults to CASE_SENSITIVE when blank.

Return type

Boolean

Examples

Default case sensitive example

The following case sensitive example evaluates if state_nm contains **New**.

```
contains(state_nm, "New")
```

The following are the given field values.

```
New York  
new york
```

For these field values, the following values are returned.

```
true  
false
```

Case insensitive example

The following case insensitive example evaluates if state_nm contains **new**.

```
contains(state_nm, "new", CASE_INSENSITIVE)
```

The following are the given field values.

```
New York  
new york
```

For these field values, the following values are returned.

```
true  
true
```

Example with conditional statements

The contains function can be used as the conditional statement within the following If functions: [avgIf](#), [minIf](#), [distinct_countIf](#), [countIf](#), [maxIf](#), [medianIf](#), [stdevIf](#), [stdevplf](#), [sumIf](#), [varIf](#), and [varplf](#).

The following example sums Sales only if state_nm contains **New**.

```
sumIf(Sales, contains(state_nm, "New"))
```

Does NOT contain example

The conditional NOT operator can be used to evaluate if the expression does not contain the specified substring.

```
NOT(contains(state_nm, "New"))
```

Example using numeric values

Numeric values can be used in the expression or substring arguments by applying the toString function.

```
contains(state_nm, toString(5) )
```

decimalToInt

decimalToInt converts a decimal value to the integer data type by stripping off the decimal point and any numbers after it. decimalToInt does not round up. For example, decimalToInt(29.99) returns 29.

Syntax

```
decimalToInt(decimal)
```

Arguments

decimal

A field that uses the decimal data type, a literal value like **17.62**, or a call to another function that outputs a decimal.

Return type

Integer

Example

The following example converts a decimal field to an integer.

```
decimalToInt(salesAmount)
```

The following are the given field values.

```
20.13  
892.03  
57.54
```

For these field values, the following values are returned.

```
20  
892  
57
```

dateDiff

`dateDiff` returns the difference in days between two date fields. If you include a value for the period, `dateDiff` returns the difference in the period interval, rather than in days.

Syntax

```
dateDiff(date1, date2, [period])
```

Arguments

`dateDiff` takes two dates as arguments. Specifying a period is optional.

date 1

The first date in the comparison. A date field or a call to another function that outputs a date.

date 2

The second date in the comparison. A date field or a call to another function that outputs a date.

period

The period of difference that you want returned, enclosed in quotes. Valid periods are as follows:

- YYYY: This returns the year portion of the date.
- Q: This returns the date of the first day of the quarter that the date belongs to.
- MM: This returns the month portion of the date.
- DD: This returns the day portion of the date.
- WK: This returns the week portion of the date. The week starts on Sunday in Amazon Quick Suite.
- HH: This returns the hour portion of the date.
- MI: This returns the minute portion of the date.
- SS: This returns the second portion of the date.
- MS: This returns the millisecond portion of the date.

Return type

Integer

Example

The following example returns the difference between two dates.

```
dateDiff(orderDate, shipDate, "MM")
```

The following are the given field values.

orderDate	shipdate
01/01/18	03/05/18
09/13/17	10/20/17

For these field values, the following values are returned.

```
2
1
```

endsWith

endsWith evaluates if the expression ends with a substring that you specify. If the expression ends with the substring, endsWith returns true, and otherwise it returns false.

Syntax

```
endsWith(expression, substring, string-comparison-mode)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

substring

The set of characters to check against the *expression*. The substring can occur one or more times in the *expression*.

string-comparison-mode

(Optional) Specifies the string comparison mode to use:

- CASE_SENSITIVE – String comparisons are case-sensitive.
- CASE_INSENSITIVE – String comparisons are case-insensitive.

This value defaults to CASE_SENSITIVE when blank.

Return type

Boolean

Examples

Default case sensitive example

The following case sensitive example evaluates if `state_nm` endsWith **"York"**.

```
endsWith(state_nm, "York")
```

The following are the given field values.


```
New York  
new york
```

For these field values, the following values are returned.

```
true  
false
```

Case insensitive example

The following case insensitive example evaluates if `state_nm` endsWith **"york"**.

```
endsWith(state_nm, "york", CASE_INSENSITIVE)
```

The following are the given field values.

```
New York  
new york
```

For these field values, the following values are returned.

```
true  
true
```

Example with conditional statements

The `endsWith` function can be used as the conditional statement within the following If functions: [avgIf](#), [minIf](#), [distinct_countIf](#), [countIf](#), [maxIf](#), [medianIf](#), [stdevIf](#), [stdevpIf](#), [sumIf](#), [varIf](#), and [varpIf](#).

The following example sums Sales only if `state_nm` ends with **"York"**.

```
sumIf(Sales, endsWith(state_nm, "York"))
```

Does NOT contain example

The conditional NOT operator can be used to evaluate if the expression does not start with the specified substring.

```
NOT(endsWith(state_nm, "York"))
```

Example using numeric values

Numeric values can be used in the expression or substring arguments by applying the `toString` function.

```
endsWith(state_nm, toString(5) )
```

epochDate

`epochDate` converts an epoch date into a standard date in the format `yyyy-MM-ddTkk:mm:ss.SSSZ`, using the format pattern syntax specified in [Class DateTimeFormat](#) in the Joda project documentation. An example is `2015-10-15T19:11:51.003Z`.

`epochDate` is supported for use with analyses based on datasets stored in Quick Suite (SPICE).

Syntax

```
epochDate(epochdate)
```

Arguments

epochdate

An epoch date, which is an integer representation of a date as the number of seconds since 00:00:00 UTC on January 1, 1970.

epochdate must be an integer. It can be the name of a field that uses the integer data type, a literal integer value, or a call to another function that outputs an integer. If the integer value is longer than 10 digits, the digits after the 10th place are discarded.

Return type

Date

Example

The following example converts an epoch date to a standard date.

```
epochDate(3100768000)
```

The following value is returned.

```
2068-04-04T12:26:40.000Z
```

Exp

exp returns the base of natural log e raised to the power of a given expression.

Syntax

```
exp(expression)
```

Arguments

expression

The expression must be numeric. It can be a field name, a literal value, or another function.

Extract

extract returns a specified portion of a date value. Requesting a time-related portion of a date that doesn't contain time information returns 0.

Syntax

```
extract(period, date)
```

Arguments

period

The period that you want extracted from the date value. Valid periods are as follows:

- YYYY: This returns the year portion of the date.
- Q: This returns the quarter that the date belongs to (1–4).
- MM: This returns the month portion of the date.
- DD: This returns the day portion of the date.
- WD: This returns the day of the week as an integer, with Sunday as 1.
- HH: This returns the hour portion of the date.
- MI: This returns the minute portion of the date.
- SS: This returns the second portion of the date.

- MS: This returns the millisecond portion of the date.

**Note**

Extracting milliseconds is not supported in Presto databases below version 0.216.

date

A date field or a call to another function that outputs a date.

Return type

Integer

Example

The following example extracts the day from a date value.

```
extract('DD', orderDate)
```

The following are the given field values.

```
orderDate
=====
01/01/14
09/13/16
```

For these field values, the following values are returned.

```
01
13
```

Floor

`floor` decrements a decimal value to the next lowest integer. For example, `floor(29.08)` returns 29.

Syntax

```
floor(decimal)
```

Arguments

decimal

A field that uses the decimal data type, a literal value like **17.62**, or a call to another function that outputs a decimal.

Return type

Integer

Example

The following example decrements a decimal field to the next lowest integer.

```
floor(salesAmount)
```

The following are the given field values.

```
20.13  
892.03  
57.54
```

For these field values, the following values are returned.

```
20  
892  
57
```

formatDate

`formatDate` formats a date using a pattern you specify. When you are preparing data, you can use `formatDate` to reformat the date. To reformat a date in an analysis, you choose the format option from the context menu on the date field.

Syntax

```
formatDate(date, ['format'])
```

Arguments

date

A date field or a call to another function that outputs a date.

format

(Optional) A string containing the format pattern to apply. This argument accepts the format patterns specified in [Supported date formats](#).

If you don't specify a format, this string defaults to yyyy-MM-ddTkk:mm:ss:SSS.

Return type

String

Example

The following example formats a UTC date.

```
formatDate(orderDate, 'dd-MMM-yyyy')
```

The following are the given field values.

```
order date
=====
2012-12-14T00:00:00.000Z
2013-12-29T00:00:00.000Z
2012-11-15T00:00:00.000Z
```

For these field values, the following values are returned.

```
13 Dec 2012
28 Dec 2013
14 Nov 2012
```

Example

If the date contains single quotes or apostrophes, for example yyyyMMdd 'T' HHmmss, you can handle this date format by using one of the following methods.

- Enclose the entire date in double quotes, as shown in the following example:

```
formatDate({myDateField}, "yyyyMMdd'T'HHmmss")
```

- Escape the single quotes or apostrophes by adding a backslash (\) to the left of them, as shown in the following example:

```
formatDate({myDateField}, 'yyyyMMdd\'T\'HHmmss')
```

Ifelse

ifelse evaluates a set of *if, then* expression pairings, and returns the value of the *then* argument for the first *if* argument that evaluates to true. If none of the *if* arguments evaluate to true, then the value of the *else* argument is returned.

Syntax

```
ifelse(if-expression-1, then-expression-1 [, if-expression-n, then-expression-n ...], else-expression)
```

Arguments

ifelse requires one or more *if,then* expression pairings, and requires exactly one expression for the *else* argument.

if-expression

The expression to be evaluated as true or not. It can be a field name like **address1**, a literal value like **'Unknown'**, or another function like `toString(salesAmount)`. An example is `isNull(Fieldname)`.

If you use multiple AND and OR operators in the *if* argument, enclose statements in parentheses to identify processing order. For example, the following *if* argument returns records with a month of 1, 2, or 5 and a year of 2000.

```
ifelse((month = 5 OR month < 3) AND year = 2000, 'yes', 'no')
```

The next *if* argument uses the same operators, but returns records with a month of 5 and any year, or with a month of 1 or 2 and a year of 2000.

```
ifelse(month = 5 OR (month < 3 AND year = 2000), 'yes', 'no')
```

then-expression

The expression to return if its *if* argument is evaluated as true. It can be a field name like **address1**, a literal value like **'Unknown'**, or a call to another function. The expression must have the same data type as the other then arguments and the else argument.

else-expression

The expression to return if none of the *if* arguments evaluate as true. It can be a field name like **address1**, a literal value like **'Unknown'**, or another function like `toString(salesAmount)`. The expression must have the same data type as all of the then arguments.

Return type

`ifelse` returns a value of the same data type as the values in *then-expression*. All data returned *then* and *else* expressions must be of the same data type or be converted to the same data type.

Examples

The following example generates a column of aliases for field `country`.

```
ifelse(country = "United States", "US", country = "China", "CN", country = "India",  
       "IN", "Others")
```

For such use cases evaluating each value in a field against a list of literals, and returns the result corresponding to the first matching value., function `switch` is recommended to simplify your work. The previous example can be rewritten to the following statement using [switch](#):

```
switch(country, "United States", "US", "China", "CN", "India", "IN", "Others")
```

The following example categorizes sales per customer into human-readable levels.

```
ifelse(salesPerCustomer < 1000, "VERY_LOW", salesPerCustomer < 10000, "LOW",  
       salesPerCustomer < 100000, "MEDIUM", "HIGH")
```

The following example uses AND, OR, and NOT to compare multiple expressions using conditional operators to tag top customers NOT in Washington or Oregon with a special promotion, who made more than 10 orders. If no values are returned, the value `'n/a'` is used.


```
ifelse(( NOT (State = 'WA' OR State = 'OR')) AND Orders > 10), 'Special Promotion XYZ', 'n/a')
```

The following examples use only OR to generate a new column that contains the name of continent that corresponds to each country.

```
ifelse(country = "United States" OR country = "Canada", "North America", country = "China" OR country = "India" OR country = "Japan", "Asia", "Others")
```

The previous example can be simplified as shown in the next example. The following example uses `ifelse` and [in](#) to create a value in a new column for any row where the tested value is in a literal list. You could use `ifelse` with [notin](#) as well.

```
ifelse(in(country,["United States", "Canada"]), "North America", in(country, ["China","Japan","India"]), "Asia", "Others")
```

Authors are able to save a literal list in a multivalue parameter and use it in the [in](#) or [notin](#) functions. The following example is an equivalent of the previous example, except that the literal lists are stored in two multivalue parameters.

```
ifelse(in(country,${NorthAmericaCountryParam}), "North America", in(country, ${AsiaCountryParam}), "Asia", "Others")
```

The following example assigns a group to a sales record based on the sales total. The structure of each `if-then` phrase mimics the behavior of *between*, a keyword that doesn't currently work in calculated field expressions. For example, the result of the comparison `salesTotal >= 0 AND salesTotal < 500` returns the same values as the SQL comparison `salesTotal between 0 and 499`.

```
ifelse(salesTotal >= 0 AND salesTotal < 500, 'Group 1', salesTotal >= 500 AND salesTotal < 1000, 'Group 2', 'Group 3')
```

The following example tests for a NULL value by using `coalesce` to return the first non-NULL value. Instead of needing to remember the meaning of a NULL in a date field, you can use a readable description instead. If the disconnect date is NULL, the example returns the suspend date, unless both of those are NULL. Then `coalesce(DiscoDate, SuspendDate, '12/31/2491')` returns '12/31/2491'. The return value must match the other data types. This date might seem

like an unusual value, but a date in the 25th century reasonably simulates the "end of time," defined as the highest date in a data mart.

```
ifelse ( (coalesce(DiscoDate, SuspendDate, '12/31/2491') = '12/31/2491'), 'Active
subscriber', 'Inactive subscriber')
```

The following shows a more complex example in a more readable format, just to show that you don't need to compress your code all into one long line. This example provides for multiple comparisons of the value a survey result. It handles potential NULL values for this field and categorizes two acceptable ranges. It also labels one range that needs more testing and another that's not valid (out of range). For all remaining values, it applies the `else` condition, and labels the row as needing a retest three years after the date on that row.

```
ifelse
(
  isNull({SurveyResult}), 'Untested',
  {SurveyResult}=1, 'Range 1',
  {SurveyResult}=2, 'Range 2',
  {SurveyResult}=3, 'Need more testing',
  {SurveyResult}=99, 'Out of Range',
  concat
  (
    'Retest by ',
    toString
    (
      addDateTime(3, "YYYY", {Date})
    )
  )
)
```

The following example assigns a "manually" created region name to a group of states. It also uses spacing and comments, wrapped in `/* */`, to make it easier to maintain the code.

```
ifelse
(
  /* NE REGION*/
  locate('New York, New Jersey, Connecticut, Vermont, Maine, Rhode Island, New
Hampshire',{State})) > 0,
  'Northeast',

  /* SE REGION*/
  locate('Georgia, Alabama, South Carolina, Louisiana',{State})) > 0,
```

```
'Southeast',
'Other Region'
)
```

The logic for the region tagging breaks down as follows:

1. We list the states that we want for each region, enclosing each list in quotation marks to make each list a string, as follows:
 - 'New York, New Jersey, Connecticut, Vermont, Maine, Rhode Island, New Hampshire'
 - 'Georgia, Alabama, South Carolina, Louisiana'
 - You can add more sets, or use countries, cities, provinces, or What3Words if you want.
2. We ask if the value for State (for each row) is found in the list, by using the `locate` function to return a nonzero value if the state is found in the list, as follows.

```
locate('New York, New Jersey, Connecticut, Vermont, Maine, Rhode Island, New
Hampshire',{State})
```

and

```
locate('Georgia, Alabama, South Carolina, Louisiana',{State})
```

3. The `locate` function returns a number instead of a TRUE or FALSE, but `ifelse` requires the TRUE/FALSE Boolean value. To get around this, we can compare the result of `locate` to a number. If the state is in the list, the return value is greater than zero.
 - a. Ask if the state is present.

```
locate('New York, New Jersey, Connecticut, Vermont, Maine, Rhode Island, New
Hampshire',{State}) > 0
```

- b. If it's present the region, label it as the specific region, in this case a Northeast region.

```
/*The if expression:*/      locate('New York, New Jersey, Connecticut, Vermont,
Maine, Rhode Island, New Hampshire',{State}) > 0,
/*The then expression:*/    'Northeast',
```

4. Because we have states that aren't in a list, and because `ifelse` requires a single `else` expression, we provide 'Other Region' as the label for the leftover states.

```
/*The if expression:*/      locate('New York, New Jersey, Connecticut, Vermont, Maine,
Rhode Island, New Hampshire',{State}) > 0,
/*The then expression:*/    'Northeast',
/*The else expression:*/    'Other Region'
```

5. We wrap all that in the `ifelse( )` function to get the final version. The following example leaves out the Southeast region states that were in the original. You can add them back in place of the *<insert more regions here>* tag.

If you want to add more regions, you can construct more copies of those two lines and alter the list of states to suit your purpose. You can change the region name to something that suits you, and change the field name from `State` to anything that you need.

```
ifelse
(
/*The if expression:*/      locate('New York, New Jersey, Connecticut, Vermont, Maine,
Rhode Island, New Hampshire',{State}) > 0,
/*The then expression:*/    'Northeast',

/*<insert more regions here>*/

/*The else expression:*/    'Other Region'
)
```

Note

There are other ways to do the initial comparison for the if expression. For example, suppose that you pose the question "What states are not missing from this list?" rather than "Which states are on the list?" If you do, you might phrase it differently. You might compare the `locate` statement to zero to find values that are missing from the list, and then use the `NOT` operator to classify them as "not missing," as follows.

```
/*The if expression:*/      NOT (locate('New York, New Jersey, Connecticut,
Vermont, Maine, Rhode Island, New Hampshire',{State}) = 0),
```

Both versions are correct. The version that you choose should make the most sense to you and your team, so you can maintain it easily. If all the options seem equal, choose the simplest.

in

`in` evaluates if an expression exists within a literal list. If the list contains the expression, `in` returns true, and otherwise it returns false. `in` is case sensitive for string type inputs.

`in` accepts two kinds of literal list, one is manually entered list and the other is a [multivalue parameter](#).

Syntax

Using a manually entered list:

```
in(expression, [literal-1, ...])
```

Using a multivalue parameter:

```
in(expression, $multivalue_parameter)
```

Arguments

expression

The expression to be compared with the elements in literal list. It can be a field name like address, a literal value like ' **Unknown**', a single value parameter, or a call to another scalar function—provided this function is not an aggregate function or a table calculation.

literal list

(required) This can be a manually entered list or a multivalue parameter. This argument accepts up to 5,000 elements. However, in a direct query to a third party data source, for example Oracle or Teradata, the restriction can be smaller.

- **manually entered list** – One or more literal values in a list to be compared with the expression. The list should be enclosed in square brackets. All the literals to compare must have the same datatype as the expression.
- **multivalue parameter** – A pre-defined multivalue parameter passed in as a literal list. The multivalue parameter must have the same datatype as the expression.

Return type

Boolean: TRUE/FALSE

Example with a static list

The following example evaluates the `origin_state_name` field for values in a list of string. When comparing string type input, `in` only supports case sensitive comparison.

```
in(origin_state_name, ["Georgia", "Ohio", "Texas"])
```

The following are the given field values.

```
"Washington"  
  "ohio"  
  "Texas"
```

For these field values the following values are returned.

```
false  
  false  
  true
```

The third return value is true because only "Texas" is one of the included values.

The following example evaluates the `fl_date` field for values in a list of string. In order to match the type, `toString` is used to cast the date type to string type.

```
in(toString(fl_date), ["2015-05-14", "2015-05-15", "2015-05-16"])
```

Literals and NULL values are supported in expression argument to be compared with the literals in list. Both of the following two examples will generate a new column of TRUE values.

```
in("Washington", ["Washington", "Ohio"])
```

```
in(NULL, [NULL, "Ohio"])
```

Example with multivalue parameter

Let's say an author creates a [multivalue parameter](#) that contains a list of all the state names. Then the author adds a control to allow the reader to select values from the list.

Next, the reader selects three values—"Georgia", "Ohio", and "Texas"—from the parameter's drop down list control. In this case, the following expression is equivalent to the first example, where those three state names are passed as the literal list to be compared with the `original_state_name` field.

```
in (origin_state_name, ${stateName MultivalueParameter})
```

Example with `ifelse`

`in` can be nested in other functions as a boolean value. One example is that authors can evaluate any expression in a list and return the value they want by using `in` and `ifelse`. The following example evaluates if the `dest_state_name` of a flight is in a particular list of US states and returns different categories of the states based on the comparison.

```
ifelse(in(dest_state_name, ["Washington", "Oregon", "California"]), "WestCoastUSState",  
      "Other US State")
```

`intToDecimal`

`intToDecimal` converts an integer value to the decimal data type.

Syntax

```
intToDecimal(integer)
```

Arguments

int

A field that uses the integer data type, a literal value like **14**, or a call to another function that outputs an integer.

Return type

Decimal

Example

The following example converts an integer field to a decimal.

```
intToDecimal(price)
```

The following are the given field values.

```
20  
892  
57
```

For these field values, the following values are returned.

```
20.0  
892.0  
58.0
```

You can apply formatting inside an analysis, for example to format price as currency.

isNotNull

`isNotNull` evaluates an expression to see if it is not null. If the expression is not null, `isNotNull` returns true, and otherwise it returns false.

Syntax

```
isNotNull(expression)
```

Arguments

expression

The expression to be evaluated as null or not. It can be a field name like **address1** or a call to another function that outputs a string.

Return type

Boolean

Example

The following example evaluates the `sales_amount` field for null values.

```
isNotNull(salesAmount)
```


The following are the given field values.

```
20.13  
(null)  
57.54
```

For these field values, the following values are returned.

```
true  
false  
true
```

isNull

`isNull` evaluates an expression to see if it is null. If the expression is null, `isNull` returns true, and otherwise it returns false.

Syntax

```
isNull(expression)
```

Arguments

expression

The expression to be evaluated as null or not. It can be a field name like **address1** or a call to another function that outputs a string.

Return type

Boolean

Example

The following example evaluates the `sales_amount` field for null values.

```
isNull(salesAmount)
```

The following are the given field values.

```
20.13
```

```
(null)
57.54
```

For these field values, the following values are returned.

```
false
true
false
```

The following example tests for a NULL value in an `ifelse` statement, and returns a human-readable value instead.

```
ifelse( isNull({ActiveFlag}) , 'Inactive', 'Active')
```

isWorkDay

`isWorkDay` evaluates a given date-time value to determine if the value is a workday or not.

`isWorkDay` assumes a standard 5-day work week starting from Monday and ending on Friday. Saturday and Sunday are assumed to be weekends. The function always calculates its result at the DAY granularity and is exclusive of the given input date.

Syntax

```
isWorkDay(inputDate)
```

Arguments

inputDate

The date-time value that you want to evaluate. Valid values are as follows:

- Dataset fields: Any date field from the dataset that you are adding this function to.
- Date Functions: Any date output from another date function, for example, `parseDate`.
- Calculated fields: Any Quick Suite calculated field that returns a date value.
- Parameters: Any Quick Suite `DateTime` parameter.

Return type

Integer (0 or 1)

Example

The following example determines whether or not the `application_date` field is a work day.

Let's assume that there's a field named `application_date` with the following values:

```
2022-08-10  
2022-08-06  
2022-08-07
```

When you use these fields and add the following calculations, `isWorkDay` returns the below values:

```
isWorkDay({application_date})
```

```
1  
0  
0
```

The following example filters employees whose employment ends on a work day and determines whether their employment began on work day or a weekend using conditional formatting:

```
is_start_date_work_day = isWorkDay(employment_start_date)  
is_end_date_work_day = isWorkDay(employment_end_date)
```

Left

`left` returns the leftmost characters from a string, including spaces. You specify the number of characters to be returned.

Syntax

```
left(expression, limit)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

limit

The number of characters to be returned from *expression*, starting from the first character in the string.

Return type

String

Example

The following example returns the first 3 characters from a string.

```
left('Seattle Store #14', 3)
```

The following value is returned.

```
Sea
```

Locate

locate locates a substring that you specify within another string, and returns the number of characters until the first character in the substring. The function returns 0 if it doesn't find the substring. The function is 1-based.

Syntax

```
locate(expression, substring, start)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

substring

The set of characters in *expression* that you want to locate. The substring can occur one or more times in *expression*.

start

(Optional) If *substring* occurs more than once, use *start* to identify where in the string the function should start looking for the substring. For example, suppose that you want to find the second example of a substring and you think it typically occurs after the first 10 characters. You specify a *start* value of 10. It should start from 1.

Return type

Integer

Examples

The following example returns information about where the first occurrence of the substring 'and' appears in a string.

```
locate('1 and 2 and 3 and 4', 'and')
```

The following value is returned.

```
3
```

The following example returns information about where the first occurrence of the substring 'and' appears in a string after the fourth character.

```
locate('1 and 2 and 3 and 4', 'and', 4)
```

The following value is returned.

```
9
```

Log

log returns the base 10 logarithm of a given expression.

Syntax

```
log(expression)
```

Arguments

expression

The expression must be numeric. It can be a field name, a literal value, or another function.

Ln

Ln returns the natural logarithm of a given expression.

Syntax

```
ln(expression)
```

Arguments

expression

The expression must be numeric. It can be a field name, a literal value, or another function.

Ltrim

Ltrim removes preceding blank space from a string.

Syntax

```
ltrim(expression)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like '**12 Main Street**', or a call to another function that outputs a string.

Return type

String

Example

The following example removes the preceding spaces from a string.

```
ltrim('   Seattle Store #14')
```

The following value is returned.

```
Seattle Store #14
```

Mod

Use the mod function to find the remainder after dividing the number by the divisor. You can use the mod function or the modulo operator (%) interchangeably.

Syntax

```
mod(number, divisor)
```

```
number%divisor
```

Arguments

number

The number is the positive integer that you want to divide and find the remainder for.

divisor

The divisor is the positive integer that you are dividing by. If the divisor is zero, this function returns an error on dividing by 0.

Example

The following examples return the modulo of 17 when dividing by 6. The first example uses the % operator, and the second example uses the mod function.

```
17%6
```

```
mod( 17, 6 )
```

The following value is returned.

netWorkDays

netWorkDays returns the number of working days between the provided two date fields or even custom date values generated using other Quick Suite date functions such as `parseDate` or `epochDate` as an integer.

netWorkDays assumes a standard 5-day work week starting from Monday and ending on Friday. Saturday and Sunday are assumed to be weekends. The calculation is inclusive of both `startDate` and `endDate`. The function operates on and shows results for DAY granularity.

Syntax

```
netWorkDays(startDate, endDate)
```

Arguments

startDate

A valid non-NULL date that acts as the start date for the calculation.

- Dataset fields: Any date field from the dataset that you are adding this function to.
- Date Functions: Any date output from another date function, for example, `parseDate`.
- Calculated fields: Any Quick Suite calculated field that returns a date value.
- Parameters: Any Quick Suite `DateTime` parameter.
- Any combination of the above stated argument values.

endDate

A valid non-NULL date that acts as the end date for the calculation.

- Dataset fields: Any date field from the dataset that you are adding this function to.
- Date Functions: Any date output from another date function, for example, `parseDate`.
- Calculated fields: Any Quick Suite calculated field that returns a date value.
- Parameters: Any Quick Suite `DateTime` parameter.
- Any combination of the above stated argument values.

Return type

Integer

Output values

Expected output values include:

- Positive integer (when `start_date < end_date`)
- Negative integer (when `start_date > end_date`)
- NULL when one or both of the arguments get a null value from the `dataset` field.

Example

The following example returns the number of work days falling between two dates.

Let's assume that there's a field named `application_date` with the following values:

```
netWorkDays({startDate}, {endDate})
```

The following are the given field values.

```
startDate endDate netWorkDays
9/4/2022 9/11/2022 5
9/9/2022 9/2/2022 -6
9/10/2022 9/11/2022 0
9/12/2022 9/12/2022 1
```

The following example calculates the number of days worked by each employee and the salary expended per day for each employee:

```
days_worked = netWorkDays({employment_start_date}, {employment_end_date})
salary_per_day = {salary}/{days_worked}
```

The following example filters employees whose employment ends on a work day and determines whether their employment began on work day or a weekend using conditional formatting:

```
is_start_date_work_day = netWorkDays(employment_start_date)
is_end_date_work_day = netWorkDays(employment_end_date)
```

Now

For database datasets that directly query the database, now returns the current date and time using the settings and format specified by the database server. For SPICE and Salesforce data sets, now returns the UTC date and time, in the format yyyy-MM-ddTkk:mm:ss:SSSZ (for example, 2015-10-15T19:11:51:003Z).

Syntax

```
now()
```

Return type

Date

notIn

notIn evaluates if an expression exists within a literal list. If the list doesn't contain the expression, notIn returns true, and otherwise it returns false. notIn is case sensitive for string type inputs.

notIn accepts two kinds of literal list, one is manually entered list and the other is a [multivalue parameter](#).

Syntax

Using a manually entered list:

```
notIn(expression, [literal-1, ...])
```

Using a multivalue parameter:

```
notIn(expression, $multivalue_parameter)
```

Arguments

expression

The expression to be compared with the elements in literal list. It can be a field name like address, a literal value like 'Unknown', a single value parameter, or a call to another scalar function—provided this function is not an aggregate function or a table calculation.

literal list

(required) This can be a manually entered list or a multivalue parameter. This argument accepts up to 5,000 elements. However, in a direct query to a third party data source, for example Oracle or Teradata, the restriction can be smaller.

- ***manually entered list*** – One or more literal values in a list to be compared with the expression. The list should be enclosed in square brackets. All the literals to compare must have the same datatype as the expression.
- ***multivalue parameter*** – A pre-defined multivalue parameter passed in as a literal list. The multivalue parameter must have the same datatype as the expression.

Return type

Boolean: TRUE/FALSE

Example with a manually entered list

The following example evaluates the `origin_state_name` field for values in a list of string. When comparing string type input, `notIn` only supports case sensitive comparison.

```
notIn(origin_state_name, ["Georgia", "Ohio", "Texas"])
```

The following are the given field values.

```
"Washington"  
  "ohio"  
  "Texas"
```

For these field values the following values are returned.

```
true  
  true  
  false
```

The third return value is false because only "Texas" is one of the excluded values.

The following example evaluates the `fl_date` field for values in a list of string. In order to match the type, `toString` is used to cast the date type to string type.

```
notIn(toString(fl_date),["2015-05-14","2015-05-15","2015-05-16"])
```

Literals and NULL values are supported in expression argument to be compared with the literals in list. Both of the following two examples will generate a new column of FALSE values.

```
notIn("Washington",["Washington","Ohio"])
```

```
notIn(NULL,[NULL,"Ohio"])
```

Example with multivalue parameter

Let's say an author creates a [multivalue parameter](#) that contains a list of all the state names. Then the author adds a control to allow the reader to select values from the list.

Next, the reader selects three values—"Georgia", "Ohio", and "Texas"—from the parameter's drop down list control. In this case, the following expression is equivalent to the first example, where those three state names are passed as the literal list to be compared with the `original_state_name` field.

```
notIn (origin_state_name, ${stateName MultivalueParameter})
```

Example with ifelse

`notIn` can be nested in other functions as a boolean value. One example is that authors can evaluate any expression in a list and return the value they want by using `notIn` and `ifelse`. The following example evaluates if the `dest_state_name` of a flight is in a particular list of US states and returns different categories of the states based on the comparison.

```
ifelse(notIn(dest_state_name,["Washington", "Oregon", "California"]),  
      "notWestCoastUSState", "WestCoastUSState")
```

nullIf

`nullIf` compares two expressions. If they are equal, the function returns null. If they are not equal, the function returns the first expression.

Syntax

```
nullIf(expression1, expression2)
```

Arguments

`nullIf` takes two expressions as arguments.

expression

The expression can be numeric, datetime, or string. It can be a field name, a literal value, or another function.

Return type

String

Example

The following example returns nulls if the reason for a shipment delay is unknown.

```
nullIf(delayReason, 'unknown')
```

The following are the given field values.

```
delayReason  
=====  
unknown  
back ordered  
weather delay
```

For these field values, the following values are returned.

```
(null)  
back ordered  
weather delay
```

parseDate

`parseDate` parses a string to determine if it contains a date value, and returns a standard date in the format `yyyy-MM-ddTkk:mm:ss.SSSZ` (using the format pattern syntax specified in [Class](#)

[DateTimeFormat](#) in the Joda project documentation), for example 2015-10-15T19:11:51.003Z. This function returns all rows that contain a date in a valid format and skips any rows that don't, including rows that contain null values.

Quick Suite supports dates in the range from Jan 1, 1900 00:00:00 UTC to Dec 31, 2037 23:59:59 UTC. For more information, see [Supported date formats](#).

Syntax

```
parseDate(expression, ['format'])
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'1/1/2016'**, or a call to another function that outputs a string.

format

(Optional) A string containing the format pattern that *date_string* must match. For example, if you are using a field with data like **01/03/2016**, you specify the format 'MM/dd/yyyy'. If you don't specify a format, it defaults to yyyy-MM-dd. Rows whose data doesn't conform to *format* are skipped.

Different date formats are supported based on the type of dataset used. Use the following table to see details of supported date formats.

Date source type	Supported date formats
File, Amazon Athena, and Salesforce datasets	All date format patterns specified in Supported date formats .
Direct query of Amazon Aurora, MariaDB, and MySQL databases	• MM/dd/yyyy • dd/MM/yyyy • yyyy/MM/dd •

Date source type	Supported date formats
	MMM/dd/yyyy • dd/MMM/yyyy • yyyy/MMM/dd • MM/dd/yyyy HH:mm:ss • dd/MM/yyyy HH:mm:ss • yyyy/MM/dd HH:mm:ss • MMM/dd/yyyy HH:mm:ss • dd/MMM/yyyy HH:mm:ss • yyyy/MMM/dd HH:mm:ss • MM-dd-yyyy • dd-MM-yyyy • yyyy-MM-dd • MMM-dd-yyyy • dd-MMM-yyyy • yyyy-MMM-dd • MM-dd-yyyy HH:mm:ss • dd-MM-yyyy HH:mm:ss • yyyy-MM-dd HH:mm:ss •

Date source type	Supported date formats
	MMM-dd-yyyy HH:mm:ss • dd-MMM-yyyy HH:mm:ss • yyyy-MMM-dd HH:mm:ss • MM/dd/yyyy HH:mm:ss.SSS • dd/MM/yyyy HH:mm:ss.SSS • yyyy/MM/dd HH:mm:ss.SSS • MMM/dd/yyyy HH:mm:ss.SSS • dd/MMM/yyyy HH:mm:ss.SSS • yyyy/MMM/dd HH:mm:ss.SSS • MM-dd-yyyy HH:mm:ss.SSS • dd-MM-yyyy HH:mm:ss.SSS • yyyy-MM-dd HH:mm:ss.SSS • MMM-dd-yyyy HH:mm:ss.SSS • dd-MMM-yyyy HH:mm:ss.SSS • yyyy-MMM-dd HH:mm:ss.SSS

Date source type	Supported date formats
Direct query of Snowflake	- dd/MM/yyyy - dd/MM/yyyy HH:mm:ss - dd-MM-yyyy - dd-MM-yyyy HH:mm:ss - MM/dd/yyyy - MM/dd/yyyy HH:mm:ss - MM-dd-yyyy - MM-dd-yyyy HH:mm:ss - yyyy/MM/dd - yyyy/MM/dd HH:mm:ss - yyyy-MM-dd - yyyy-MM-dd HH:mm:ss - MM/dd/yyyy HH:mm:ss.SSS - dd/MM/yyyy HH:mm:ss.SSS - yyyy/MM/dd HH:mm:ss.SSS - MMM/dd/yyyy HH:mm:ss.SSS - dd/MMM/yyyy HH:mm:ss.SSS - yyyy/MMM/dd HH:mm:ss.SSS

Date source type	Supported date formats
	MM-dd-yyyy HH:mm:ss.SSS • dd-MM-yyyy HH:mm:ss.SSS • yyyy-MM-dd HH:mm:ss.SSS • MMM-dd-yyyy HH:mm:ss.SSS • dd-MMM-yyyy HH:mm:ss.SSS • yyyy-MMM-dd HH:mm:ss.SSS

Date source type	Supported date formats
Direct query of Microsoft SQL Server databases	• dd-MM-yyyy • MM/dd/yyyy • dd/MM/yyyy • yyyy/MM/dd • MMM/dd/yyyy • dd/MMM/yyyy • yyyy/MMM/dd • dd/MM/yyyy HH:mm:ss • yyyy/MM/dd HH:mm:ss • MMM/dd/yyyy HH:mm:ss • dd/MMM/yyyy HH:mm:ss • yyyy/MMM/dd HH:mm:ss • MM-dd-yyyy • yyyy-MM-dd • MMM-dd-yyyy • yyyy-MMM-dd • MM-dd-yyyy HH:mm:ss • dd-MM-yyyy HH:mm:ss •

Date source type	Supported date formats
	yyyy-MM-dd HH:mm:ss • MMM-dd-yyyy HH:mm:ss • dd-MMM-yyyy HH:mm:ss • yyyy-MMM-dd HH:mm:ss • MM/dd/yyyy HH:mm:ss.SSS • dd/MM/yyyy HH:mm:ss.SSS • yyyy/MM/dd HH:mm:ss.SSS • MMM/dd/yyyy HH:mm:ss.SSS • dd/MMM/yyyy HH:mm:ss.SSS • yyyy/MMM/dd HH:mm:ss.SSS • MM-dd-yyyy HH:mm:ss.SSS • dd-MM-yyyy HH:mm:ss.SSS • yyyy-MM-dd HH:mm:ss.SSS • MMM-dd-yyyy HH:mm:ss.SSS • dd-MMM-yyyy HH:mm:ss.SSS • yyyy-MMM-dd HH:mm:ss.SSS

Date source type	Supported date formats
Direct query of Amazon Redshift or PostgreSQL databases Also, datasets from any DBMS that are stored in Quick Suite SPICE	• MM/dd/yyyy • dd/MM/yyyy • yyyy/MM/dd • MMM/dd/yyyy • dd/MMM/yyyy • yyyy/MMM/dd • MM/dd/yyyy HH:mm:ss • dd/MM/yyyy HH:mm:ss • yyyy/MM/dd HH:mm:ss • MMM/dd/yyyy HH:mm:ss • dd/MMM/yyyy HH:mm:ss • yyyy/MMM/dd HH:mm:ss • MM-dd-yyyy • dd-MM-yyyy • yyyy-MM-dd • MMM-dd-yyyy • dd-MMM-yyyy • yyyy-MMM-dd •

Date source type	Supported date formats
	MM-dd-yyyy HH:mm:ss • dd-MM-yyyy HH:mm:ss • yyyy-MM-dd HH:mm:ss • MMM-dd-yyyy HH:mm:ss • dd-MMM-yyyy HH:mm:ss • yyyy-MMM-dd HH:mm:ss • yyyyMMdd'T'HHmmss • yyyy-MM-dd'T'HH:mm:ss • MM/dd/yyyy HH:mm:ss.SSS • dd/MM/yyyy HH:mm:ss.SSS • yyyy/MM/dd HH:mm:ss.SSS • MMM/dd/yyyy HH:mm:ss.SSS • dd/MMM/yyyy HH:mm:ss.SSS • yyyy/MMM/dd HH:mm:ss.SSS • MM-dd-yyyy HH:mm:ss.SSS • dd-MM-yyyy HH:mm:ss.SSS • yyyy-MM-dd HH:mm:ss.SSS • MMM-dd-yyyy HH:mm:ss.SSS •

Date source type	Supported date formats
	dd-MMM-yyyy HH:mm:ss.SSS
	yyyy-MMM-dd HH:mm:ss.SSS

Return type

Date

Example

The following example evaluates prodDate to determine if it contains date values.

```
parseDate(prodDate, 'MM/dd/yyyy')
```

The following are the given field values.

```
prodDate
-----
01-01-1999
12/31/2006
1/18/1982
7/4/2010
```

For these field values, the following rows are returned.

```
12-31-2006T00:00:00.000Z
01-18-1982T00:00:00.000Z
07-04-2010T00:00:00.000Z
```

parseDecimal

`parseDecimal` parses a string to determine if it contains a decimal value. This function returns all rows that contain a decimal, integer, or null value, and skips any rows that don't. If the row contains an integer value, it is returned as a decimal with up to 4 decimal places. For example, a value of '2' is returned as '2.0'.

Syntax

```
parseDecimal(expression)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like '9.62', or a call to another function that outputs a string.

Return type

Decimal

Example

The following example evaluates fee to determine if it contains decimal values.

```
parseDecimal(fee)
```

The following are the given field values.

```
fee
-----
2
2a
12.13
3b
3.9
(null)
198.353398
```

For these field values, the following rows are returned.

```
2.0
12.13
3.9
(null)
```



```
198.3533
```

parseInt

`parseInt` parses a string to determine if it contains an integer value. This function returns all rows that contain a decimal, integer, or null value, and skips any rows that don't. If the row contains a decimal value, it is returned as the nearest integer, rounded down. For example, a value of '2.99' is returned as '2'.

Syntax

```
parseInt(expression)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like '3', or a call to another function that outputs a string.

Return type

Integer

Example

The following example evaluates `feeType` to determine if it contains integer values.

```
parseInt(feeType)
```

The following are the given field values.

```
feeType
-----
2
2.1
2a
3
3b
(null)
```

5

For these field values, the following rows are returned.

2
2
3
(null)
5

parseJson

Use `parseJson` to extract values from a JSON object.

If your dataset is stored in Quick Suite SPICE, you can use `parseJson` when you are preparing a data set, but not in calculated fields during analysis.

For direct query, you can use `parseJson` both during data preparation and analysis. The `parseJson` function applies to either strings or to JSON native data types, depending on the dialect, as shown in the following table.

Dialect	Type
PostgreSQL	JSON
Amazon Redshift	String
Microsoft SQL Server	String
MySQL	JSON
Teradata	JSON
Oracle	String
Presto	String
Snowflake	Semistructured data type object and array
Hive	String

Syntax

```
parseJson(fieldName, path)
```

Arguments

fieldName

The field containing the JSON object that you want to parse.

path

The path to the data element you want to parse from the JSON object. Only letters, numbers, and blank spaces are supported in the path argument. Valid path syntax includes:

- `$` – Root object
- `.` – Child operator
- `[]` – Subscript operator for array

Return type

String

Example

The following example evaluates incoming JSON to retrieve a value for item quantity. By using this during data preparation, you can create a table out of the JSON.

```
parseJson({jsonField}, "$.items.qty")
```

The following shows the JSON.

```
{
  "customer": "John Doe",
  "items": {
    "product": "Beer",
    "qty": 6
  },
  "list1": [
    "val1",
```

```
    "val2"  
  ],  
  "list2": [  
    {  
      "list21key1": "list1value1"  
    }  
  ]  
}
```

For this example, the following value is returned.

```
6
```

Example

The following example evaluates `JSONObject1` to extract the first key value pair (KVP), labeled "State", and assign the value to the calculated field that you are creating.

```
parseJson(JSONObject1, "$.state")
```

The following are the given field values.

```
JSONObject1  
-----  
{ "State": "New York", "Product": "Produce", "Date Sold": "1/16/2018", "Sales  
  Amount": "$3423.39" }  
{ "State": "North Carolina", "Product": "Bakery Products", "Date Sold": "2/1/2018", "Sales  
  Amount": "$3226.42" }  
{ "State": "Utah", "Product": "Water", "Date Sold": "4/24/2018", "Sales Amount": "$7001.52" }
```

For these field values, the following rows are returned.

```
New York  
North Carolina  
Utah
```

Replace

`replace` replaces part of a string with another string that you specify.

Syntax

```
replace(expression, substring, replacement)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

substring

The set of characters in *expression* that you want to replace. The substring can occur one or more times in *expression*.

replacement

The string you want to have substituted for *substring*.

Return type

String

Example

The following example replaces the substring 'and' with 'or'.

```
replace('1 and 2 and 3', 'and', 'or')
```

The following string is returned.

```
1 or 2 or 3
```

Right

right returns the rightmost characters from a string, including spaces. You specify the number of characters to be returned.

Syntax

```
right(expression, limit)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

limit

The number of characters to be returned from *expression*, starting from the last character in the string.

Return type

String

Example

The following example returns the last five characters from a string.

```
right('Seattle Store#14', 12)
```

The following value is returned.

```
tle Store#14
```

Round

round rounds a decimal value to the closest integer if no scale is specified, or to the closest decimal place if scale is specified.

Syntax

```
round(decimal, scale)
```

Arguments

decimal

A field that uses the decimal data type, a literal value like **17.62**, or a call to another function that outputs a decimal.

scale

The number of decimal places to use for the return values.

Return type

Decimal

Example

The following example rounds a decimal field to the closest second decimal place.

```
round(salesAmount, 2)
```

The following are the given field values.

```
20.1307  
892.0388  
57.5447
```

For these field values, the following values are returned.

```
20.13  
892.04  
58.54
```

Rtrim

`rtrim` removes following blank space from a string.

Syntax

```
rtrim(expression)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

Return type

String

Example

The following example removes the following spaces from a string.

```
rtrim('Seattle Store #14  ')
```

For these field values, the following values are returned.

```
Seattle Store #14
```

Split

`split` splits a string into an array of substrings, based on a delimiter that you choose, and returns the item specified by the position.

You can only add `split` to a calculated field during data preparation, not to an analysis. This function is not supported in direct queries to Microsoft SQL Server.

Syntax

```
split(expression, delimiter , position)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street;1402 35th Ave;1818 Elm Ct;11 Janes Lane'**, or a call to another function that outputs a string.

delimiter

The character that delimits where the string is broken into substrings. For example, `split('one|two|three', '|', 2)` becomes the following.

```
one
```



```
two  
three
```

If you choose `position = 2`, `split` returns `'two'`.

position

(Required) The position of the item to return from the array. The position of the first item in the array is 1.

Return type

String array

Example

The following example splits a string into an array, using the semicolon character (;) as the delimiter, and returns the third element of the array.

```
split('123 Test St;1402 35th Ave;1818 Elm Ct;11 Janes Lane', ';', 3)
```

The following item is returned.

```
1818 Elm Ct
```

This function skips items containing null values or empty strings.

Sqrt

`sqrt` returns the square root of a given expression.

Syntax

```
sqrt(expression)
```

Arguments

expression

The expression must be numeric. It can be a field name, a literal value, or another function.

startsWith

`startsWith` evaluates if the expression starts with a substring that you specify. If the expression starts with the substring, `startsWith` returns true, and otherwise it returns false.

Syntax

```
startsWith(expression, substring, string-comparison-mode)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

substring

The set of characters to check against the *expression*. The substring can occur one or more times in the *expression*.

string-comparison-mode

(Optional) Specifies the string comparison mode to use:

- `CASE_SENSITIVE` – String comparisons are case-sensitive.
- `CASE_INSENSITIVE` – String comparisons are case-insensitive.

This value defaults to `CASE_SENSITIVE` when blank.

Return type

Boolean

Examples

Default case sensitive example

The following case sensitive example evaluates if `state_nm` startsWith **New**.

```
startsWith(state_nm, "New")
```

The following are the given field values.

```
New York  
new york
```

For these field values, the following values are returned.

```
true  
false
```

Case insensitive example

The following case insensitive example evaluates if `state_nm` startsWith **new**.

```
startsWith(state_nm, "new", CASE_INSENSITIVE)
```

The following are the given field values.

```
New York  
new york
```

For these field values, the following values are returned.

```
true  
true
```

Example with conditional statements

The `startsWith` function can be used as the conditional statement within the following If functions: [avgIf](#), [minIf](#), [distinct_countIf](#), [countIf](#), [maxIf](#), [medianIf](#), [stdevIf](#), [stdevplf](#), [sumIf](#), [varIf](#), and [varplf](#).

The following example sums Sales only if `state_nm` starts with **New**.

```
sumIf(Sales,startsWith(state_nm, "New"))
```

Does NOT contain example

The conditional NOT operator can be used to evaluate if the expression does not start with the specified substring.

```
NOT(startsWith(state_nm, "New"))
```

Example using numeric values

Numeric values can be used in the expression or substring arguments by applying the `toString` function.

```
startsWith(state_nm, toString(5) )
```

Strlen

`strlen` returns the number of characters in a string, including spaces.

Syntax

```
strlen(expression)
```

Arguments

expression

An expression can be the name of a field that uses the string data type like **address1**, a literal value like '**Unknown**', or another function like `substring(field_name,0,5)`.

Return type

Integer

Example

The following example returns the length of the specified string.

```
strlen('1421 Main Street')
```

The following value is returned.

```
16
```

Substring

`substring` returns the characters in a string, starting at the location specified by the *start* argument and proceeding for the number of characters specified by the *length* arguments.

Syntax

```
substring(expression, start, length)
```

Arguments

expression

An expression can be the name of a field that uses the string data type like **address1**, a literal value like **'Unknown'**, or another function like `substring(field_name,1,5)`.

start

The character location to start from. *start* is inclusive, so the character at the starting position is the first character in the returned value. The minimum value for *start* is 1.

length

The number of additional characters to include after *start*. *length* is inclusive of *start*, so the last character returned is (*length* - 1) after the starting character.

Return type

String

Example

The following example returns the 13th through 19th characters in a string. The beginning of the string is index 1, so you begin counting at the first character.

```
substring('Fantasy and Science Fiction',13,7)
```

The following value is returned.

```
Science
```

switch

`switch` compares a *condition-expression* with the literal labels, within a set of literal label and *return-expression* pairings. It then returns the *return-expression* corresponding to the first literal label that's equal to the *condition-expression*. If no label equals to the *condition-expression*, `switch` returns the *default-expression*. Every *return-expression* and *default-expression* must have the same datatype.

Syntax

```
switch(condition-expression, label-1, return-expression-1 [, label-n, return-expression-n ...],  
      default-expression)
```

Arguments

`switch` requires one or more *if,then* expression pairings, and requires exactly one expression for the *else* argument.

condition-expression

The expression to be compared with the label-literals. It can be a field name like `address`, a literal value like `'Unknown'`, or another scalar function like `toString(salesAmount)`.

label

The literal to be compared with the *condition-expression* argument, all of the literals must have the same data type as *condition-expression* argument. `switch` accepts up to 5000 labels.

return-expression

The expression to return if the value of its label equals to the value of the *condition-expression*. It can be a field name like `address`, a literal value like `'Unknown'`, or another scalar function like `toString(salesAmount)`. All of the *return-expression* arguments must have the same data type as the *default-expression*.

default-expression

The expression to return if no value of any label arguments equals to the value of *condition-expression*. It can be a field name like `address`, a literal value like `'Unknown'`, or another scalar function like `toString(salesAmount)`. The *default-expression* must have the same data type as all of the *return-expression* arguments.

Return type

`switch` returns a value of the same data type as the values in *return-expression*. All data returned *return-expression* and *default-expression* must be of the same data type or be converted to the same data type.

General Examples

The following example returns the Amazon Web Services Region code of input region name.

```
switch(region_name,  
      "US East (N. Virginia)", "us-east-1",  
      "Europe (Ireland)", "eu-west-1",  
      "US West (N. California)", "us-west-1",  
      "other regions")
```

The following are the given field values.

```
"US East (N. Virginia)"  
"US West (N. California)"  
"Asia Pacific (Tokyo)"
```

For these field values the following values are returned.

```
"us-east-1"  
"us-west-1"  
"other regions"
```

Use `switch` to replace `ifelse`

The following `ifelse` use case is an equivalent of the previous example, for `ifelse` evaluating whether values of one field equals to different literal values, using `switch` instead is a better choice.

```
ifelse(region_name = "US East (N. Virginia)", "us-east-1",  
      region_name = "Europe (Ireland)", "eu-west-1",  
      region_name = "US West (N. California)", "us-west-1",  
      "other regions")
```

Expression as return value

The following example uses expressions in *return-expressions*:

```
switch({origin_city_name},  
      "Albany, NY", {arr_delay} + 20,  
      "Alexandria, LA", {arr_delay} - 10,  
      "New York, NY", {arr_delay} * 2,  
      {arr_delay})
```

The preceding example changes the expected delay time for each flight from a particular city.

toLowerCase

toLowerCase formats a string in all lowercase. toLowerCase skips rows containing null values.

Syntax

```
toLowerCase(expression)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

Return type

String

Example

The following example converts a string value into lowercase.

```
toLowerCase('Seattle Store #14')
```

The following value is returned.

```
seattle store #14
```

toString

toString formats the input expression as a string. toString skips rows containing null values.

Syntax

```
toString(expression)
```

Arguments

expression

An expression can be a field of any data type, a literal value like **14.62**, or a call to another function that returns any data type.

Return type

String

Example

The following example returns the values from `payDate` (which uses the date data type) as strings.

```
toString(payDate)
```

The following are the given field values.

```
payDate
-----
1992-11-14T00:00:00.000Z
2012-10-12T00:00:00.000Z
1973-04-08T00:00:00.000Z
```

For these field values, the following rows are returned.

```
1992-11-14T00:00:00.000Z
2012-10-12T00:00:00.000Z
1973-04-08T00:00:00.000Z
```

toUpper

`toUpper` formats a string in all uppercase. `toUpper` skips rows containing null values.

Syntax

```
toUpper(expression)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

Return type

String

Example

The following example converts a string value into uppercase.

```
toUpper('Seattle Store #14')
```

The following value is returned.

```
SEATTLE STORE #14
```

trim

`trim` removes both preceding and following blank space from a string.

Syntax

```
trim(expression)
```

Arguments

expression

The expression must be a string. It can be the name of a field that uses the string data type, a literal value like **'12 Main Street'**, or a call to another function that outputs a string.

Return type

String

Example

The following example removes the following spaces from a string.

```
trim('   Seattle Store #14   ')
```

For these field values, the following values are returned.

```
Seattle Store #14
```

truncDate

`truncDate` returns a date value that represents a specified portion of a date. For example, requesting the year portion of the value 2012-09-02T00:00:00.000Z returns 2012-01-01T00:00:00.000Z. Specifying a time-related period for a date that doesn't contain time information returns the initial date value unchanged.

Syntax

```
truncDate('period', date)
```

Arguments

period

The period of the date that you want returned. Valid periods are as follows:

- YYYY: This returns the year portion of the date.
- Q: This returns the date of the first day of the quarter that the date belongs to.
- MM: This returns the month portion of the date.
- DD: This returns the day portion of the date.
- WK: This returns the week portion of the date. The week starts on Sunday in Amazon Quick Suite.
- HH: This returns the hour portion of the date.

- MI: This returns the minute portion of the date.
- SS: This returns the second portion of the date.
- MS: This returns the millisecond portion of the date.

date

A date field or a call to another function that outputs a date.

Return type

Date

Example

The following example returns a date representing the month of the order date.

```
truncDate('MM', orderDate)
```

The following are the given field values.

```
orderDate
=====
2012-12-14T00:00:00.000Z
2013-12-29T00:00:00.000Z
2012-11-15T00:00:00.000Z
```

For these field values, the following values are returned.

```
2012-12-01T00:00:00.000Z
2013-12-01T00:00:00.000Z
2012-11-01T00:00:00.000Z
```

Aggregate functions

Aggregate functions are only available during analysis and visualization. Each of these functions returns values grouped by the chosen dimension or dimensions. For each aggregation, there is also a conditional aggregation. These perform the same type of aggregation, based on a condition.

When a calculated field formula contains an aggregation, it becomes a custom aggregation. To make sure that your data is accurately displayed, Amazon Quick Suite applies the following rules:

- Custom aggregations can't contain nested aggregate functions. For example, this formula doesn't work: `sum( avg(x)/avg(y) )`. However, nesting nonaggregated functions inside or outside aggregate functions does work. For example, `ceil( avg(x) )` works. So does `avg( ceil(x) )`.
- Custom aggregations can't contain both aggregated and nonaggregated fields, in any combination. For example, this formula doesn't work: `Sum(sales)+quantity`.
- Filter groups can't contain both aggregated and nonaggregated fields.
- Custom aggregations can't be converted to a dimension. They also can't be dropped into the field well as a dimension.
- In a pivot table, custom aggregations can't be added to table calculations.
- Scatter plots with custom aggregations need at least one dimension under **Group/Color** in the field wells.

For more information about supported functions and operators, see [Calculated field function and operator reference for Amazon Quick Suite](#).

The aggregate functions for calculated fields in Quick Suite include the following.

Topics

- [avg](#)
- [avgIf](#)
- [count](#)
- [countIf](#)
- [distinct_count](#)
- [distinct_countIf](#)
- [max](#)
- [maxIf](#)
- [median](#)
- [medianIf](#)
- [min](#)
- [minIf](#)
- [percentile](#)

- [percentileCont](#)
- [percentileDisc \(percentile\)](#)
- [periodToDateAvg](#)
- [periodToDateCount](#)
- [periodToDateMax](#)
- [periodToDateMedian](#)
- [periodToDateMin](#)
- [periodToDatePercentile](#)
- [periodToDatePercentileCont](#)
- [periodToDateStDev](#)
- [periodToDateStDevP](#)
- [periodToDateSum](#)
- [periodToDateVar](#)
- [periodToDateVarP](#)
- [stdev](#)
- [stdevp](#)
- [stdevlf](#)
- [stdevplf](#)
- [sum](#)
- [sumlf](#)
- [var](#)
- [varlf](#)
- [varp](#)
- [varplf](#)

avg

The avg function averages the set of numbers in the specified measure, grouped by the chosen dimension or dimensions. For example, `avg(salesAmount)` returns the average for that measure grouped by the (optional) chosen dimension.

Syntax

```
avg(decimal, [group-by level])
```

Arguments

decimal

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example calculates the average sales.

```
avg({Sales})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the average sales at the Country level, but not across other dimensions (Region or Product) in the visual.

```
avg({Sales}, [{Country}])
```

avgIf

Based on a conditional statement, the avgIf function averages the set of numbers in the specified measure, grouped by the chosen dimension or dimensions. For example, avgIf(ProdRev, CalendarDay >= \${BasePeriodStartDate} AND CalendarDay <=

`${BasePeriodEndDate} AND SourcingType <> 'Indirect')` returns the average for that measure grouped by the (optional) chosen dimension, if the condition evaluates to true.

Syntax

```
avgIf(dimension or measure, condition)
```

Arguments

decimal

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

count

The count function calculates the number of values in a dimension or measure, grouped by the chosen dimension or dimensions. For example, `count(product type)` returns the total number of product types grouped by the (optional) chosen dimension, including any duplicates. The `count(sales)` function returns the total number of sales completed grouped by the (optional) chosen dimension, for example salesperson.

Syntax

```
count(dimension or measure, [group-by level])
```

Arguments

dimension or measure

The argument must be a measure or a dimension. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example calculates the count of sales by a specified dimension in the visual. In this example, the count of sales by month are shown.

```
count({Sales})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the count of sales at the Country level, but not across other dimensions (Region or Product) in the visual.

```
count({Sales}, [{Country}])
```

countIf

Based on a conditional statement, the countIf function calculates the number of values in a dimension or measure, grouped by the chosen dimension or dimensions.

Syntax

```
countIf(dimension or measure, condition)
```

Arguments

dimension or measure

The argument must be a measure or a dimension. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

Return type

Integer

Example

The following function returns a count of the sales transactions (Revenue) that meet the conditions, including any duplicates.

```
countIf (  
    Revenue,  
    # Conditions  
        CalendarDay >= ${BasePeriodStartDate} AND  
        CalendarDay <= ${BasePeriodEndDate} AND  
        SourcingType <> 'Indirect'  
)
```

distinct_count

The `distinct_count` function calculates the number of distinct values in a dimension or measure, grouped by the chosen dimension or dimensions. For example, `distinct_count(product type)` returns the total number of unique product types grouped by the (optional) chosen dimension, without any duplicates. The `distinct_count(ship date)` function returns the total number of dates when products were shipped grouped by the (optional) chosen dimension, for example region.

Syntax

```
distinct_count(dimension or measure, [group-by level])
```

Arguments

dimension or measure

The argument must be a measure or a dimension. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Example

The following example calculates the total number of dates when products were ordered grouped by the (optional) chosen dimension in the visual, for example region.

```
distinct_count({Order Date})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the average sales at the Country level, but not across other dimensions (Region) in the visual.

```
distinct_count({Order Date}, [Country])
```

distinct_countIf

Based on a conditional statement, the `distinct_countIf` function calculates the number of distinct values in a dimension or measure, grouped by the chosen dimension or dimensions. For example, `distinct_countIf(product type)` returns the total number of unique product types grouped by the (optional) chosen dimension, without any duplicates. The `distinct_countIf(ProdRev, CalendarDay >= ${BasePeriodStartDate} AND CalendarDay <= ${BasePeriodEndDate} AND SourcingType <> 'Indirect')` function returns the total number of dates when products were shipped grouped by the (optional) chosen dimension, for example region, if the condition evaluates to true.

Syntax

```
distinct_countIf(dimension or measure, condition)
```

Arguments

dimension or measure

The argument must be a measure or a dimension. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

max

The max function returns the maximum value of the specified measure or date, grouped by the chosen dimension or dimensions. For example, `max(sales goal)` returns the maximum sales goals grouped by the (optional) chosen dimension.

Syntax

```
max(measure, [group-by level])
```

Arguments

measure

The argument must be a measure or a date. Null values are omitted from the results. Literal values don't work. The argument must be a field.

Maximum dates work only in the **Value** field well of tables and pivot tables.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example returns the max sales value for each region. It is compared to the total, minimum, and median sales values.

```
max({Sales})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the max sales at the Country level, but not across other dimensions (Region) in the visual.

```
max({Sales}, [Country])
```

maxIf

Based on a conditional statement, the maxIf function returns the maximum value of the specified measure, grouped by the chosen dimension or dimensions. For example, `maxIf(ProdRev, CalendarDay >= ${BasePeriodStartDate} AND CalendarDay <= ${BasePeriodEndDate} AND SourcingType <> 'Indirect')` returns the maximum sales goals grouped by the (optional) chosen dimension, if the condition evaluates to true.

Syntax

```
maxIf(measure, condition)
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

median

The median aggregation returns the median value of the specified measure, grouped by the chosen dimension or dimensions. For example, `median(revenue)` returns the median revenue grouped by the (optional) chosen dimension.

Syntax

```
median(measure, [group-by level])
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example returns the median sales value for each region. It is compared to the total, maximum, and minimum sales.

```
median({Sales})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the median sales at the Country level, but not across other dimensions (Region) in the visual.

```
median({Sales}, [Country])
```

medianIf

Based on a conditional statement, the `medianIf` aggregation returns the median value of the specified measure, grouped by the chosen dimension or dimensions. For example, `medianIf(Revenue, SaleDate >= ${BasePeriodStartDate} AND SaleDate <= ${BasePeriodEndDate})` returns the median revenue grouped by the (optional) chosen dimension, if the condition evaluates to true.

Syntax

```
medianIf(measure, condition)
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

min

The `min` function returns the minimum value of the specified measure or date, grouped by the chosen dimension or dimensions. For example, `min(return_rate)` returns the minimum rate of returns grouped by the (optional) chosen dimension.

Syntax

```
min(measure, [group-by level])
```

Arguments

measure

The argument must be a measure or a date. Null values are omitted from the results. Literal values don't work. The argument must be a field.

Minimum dates work only in the **Value** field well of tables and pivot tables.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example returns the minimum sales value for each region. It is compared to the total, max, and median sales.

```
min({Sales})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the minimum sales at the Country level, but not across other dimensions (Region) in the visual.

```
min({Sales}, [Country])
```

minIf

Based on a conditional statement, the minIf function returns the minimum value of the specified measure, grouped by the chosen dimension or dimensions. For example, minIf(ProdRev, CalendarDay >= \${BasePeriodStartDate} AND CalendarDay <= \${BasePeriodEndDate} AND SourcingType <> 'Indirect') returns the minimum rate of returns grouped by the (optional) chosen dimension, if the condition evaluates to true.

Syntax

```
minIf(measure, condition)
```


Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

percentile

The percentile function calculates the percentile of the values in measure, grouped by the dimension that's in the field well. There are two varieties of percentile calculation available in Quick Suite:

- [percentileCont](#) uses linear interpolation to determine result.
- [percentileDisc \(percentile\)](#) uses actual values to determine result.

The percentile function is an alias of percentileDisc.

percentileCont

The percentileCont function calculates percentile based on a continuous distribution of the numbers in the measure. It uses the grouping and sorting that are applied in the field wells. It answers questions like: What values are representative of this percentile? To return an exact percentile value that might not be present in your dataset, use percentileCont. To return the nearest percentile value that is present in your dataset, use percentileDisc instead.

Syntax

```
percentileCont(expression, percentile, [group-by level])
```

Arguments

measure

Specifies a numeric value to use to compute the percentile. The argument must be a measure or metric. Nulls are ignored in the calculation.

percentile

The percentile value can be any numeric constant 0–100. A percentile value of 50 computes the median value of the measure.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Returns

The result of the function is a number.

Usage notes

The `percentileCont` function calculates a result based on a continuous distribution of the values from a specified measure. The result is computed by linear interpolation between the values after ordering them based on settings in the visual. It's different from `percentileDisc`, which simply returns a value from the set of values that are aggregated over. The result from `percentileCont` might or might not exist in the values from the specified measure.

Examples of `percentileCont`

The following examples help explain how `percentileCont` works.

Example Comparing median, `percentileCont`, and `percentileDisc`

The following example shows the median for a dimension (category) by using the median, `percentileCont`, and `percentileDisc` functions. The median value is the same as the `percentileCont` value. `percentileCont` interpolates a value, which might or might not be in the data set. However, because `percentileDisc` always displays a value that exists in the dataset, the two results might not match. The last column in this example shows the difference between the two values. The code for each calculated field is as follows:

- `50%Cont = percentileCont( example , 50 )`
- `median = median( example )`
- `50%Disc = percentileDisc( example , 50 )`

- `Cont-Disc = percentileCont( example , 50 ) - percentileDisc( example , 50 )`
- `example = left( category, 1 )` (To make a simpler example, we used this expression to shorten the names of categories down to their first letter.)

example	median	50%Cont	50%Disc	Cont-Disc
A	22.48	22.48	22.24	0.24
B	20.96	20.96	20.95	0.01
C	24.92	24.92	24.92	0
D	24.935	24.935	24.92	0.015
E	14.48	14.48	13.99	0.49

Example 100th percentile as maximum

The following example shows a variety of `percentileCont` values for the `example` field. The calculated fields `n%Cont` are defined as `percentileCont( {example} , n )`. The interpolated values in each column represent the numbers that fall into that percentile bucket. In some cases, the actual data values match the interpolated values. For example, the column `100%Cont` shows the same value for every row because 6783.02 is the highest number.

example	50%Cont	75%Cont	99%Cont	100%Cont
A	20.97	84.307	699.99	6783.02
B	20.99	88.84	880.98	6783.02
C	20.99	90.48	842.925	6783.02
D	21.38	85.99	808.49	6783.02

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the 30th percentile based on a continuous distribution of the numbers at the Country level, but not across other dimensions (Region) in the visual.

```
percentileCont({Sales}, 30, [Country])
```

percentileDisc (percentile)

The percentileDisc function calculates the percentile based on the actual numbers in measure. It uses the grouping and sorting that are applied in the field wells. The percentile function is an alias of percentileDisc.

Use this function to answer the following question: Which actual data points are present in this percentile? To return the nearest percentile value that is present in your dataset, use percentileDisc. To return an exact percentile value that might not be present in your dataset, use percentileCont instead.

Syntax

```
percentileDisc(expression, percentile, [group-by level])
```

Arguments

measure

Specifies a numeric value to use to compute the percentile. The argument must be a measure or metric. Nulls are ignored in the calculation.

percentile

The percentile value can be any numeric constant 0–100. A percentile value of 50 computes the median value of the measure.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Returns

The result of the function is a number.

Usage notes

percentileDisc is an inverse distribution function that assumes a discrete distribution model. It takes a percentile value and a sort specification and returns an element from the given set.

For a given percentile value P, `percentileDisc` uses the sorted values in the visual and returns the value with the smallest cumulative distribution value that is greater than or equal to P.

Examples of `percentileDisc`

The following examples help explain how `percentileDisc` works.

Example Comparing median, `percentileDisc`, and `percentileCont`

The following example shows the median for a dimension (category) by using the `percentileCont`, and `percentileDisc`, and `median` functions. The median value is the same as the `percentileCont` value. `percentileCont` interpolates a value, which might or might not be in the data set. However, because `percentileDisc` always displays the closest value that exists in the dataset, the two results might not match. The last column in this example shows the difference between the two values. The code for each calculated field is as follows:

- `50%Cont = percentileCont( example , 50 )`
- `median = median( example )`
- `50%Disc = percentileDisc( example , 50 )`
- `Cont-Disc = percentileCont( example , 50 ) - percentileDisc( example , 50 )`
- `example = left( category, 1 )` (To make a simpler example, we used this expression to shorten the names of categories down to their first letter.)

example	median	50%Cont	50%Disc	Cont-Disc
A	22.48	22.48	22.24	0.24
B	20.96	20.96	20.95	0.01
C	24.92	24.92	24.92	0
D	24.935	24.935	24.92	0.015
E	14.48	14.48	13.99	0.49

Example 100th percentile as maximum

The following example shows a variety of `percentileDisc` values for the `example` field. The calculated fields `n%Disc` are defined as `percentileDisc( {example} , n)`. The values in each column are actual numbers that come from the dataset.

example	50%Disc	75%Disc	99%Disc	100%Disc
A	20.97	73.98	699.99	6783.02
B	42.19	88.84	820.08	6783.02
C	30.52	90.48	733.44	6783.02
D	41.38	85.99	901.29	6783.0

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the 30th percentile based on a continuous distribution of the numbers at the Country level, but not across other dimensions (Region) in the visual.

```
percentile({Sales}, 30, [Country])
```

periodToDateAvg

The periodToDateAvg function averages the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time, relative to that period.

Syntax

```
periodToDateAvg(
  measure,
  dateTime,
  period,
  endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick Suite begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateAvg(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```

periodToDateCount

The periodToDateCount function calculates the number of values in a dimension or measure, including duplicates, for a given time granularity (for instance, a quarter) up to a point in time, relative to that period.

Syntax

```
periodToDateCount(  
  measure,  
  dateTime,  
  period,  
  endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick Suite begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateCount(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```

periodToDateMax

The periodToDateMax function returns the maximum value of the specified measure for a given time granularity (for instance, a quarter) up to a point in time, relative to that point.

Syntax

```
periodToDateMax(  
  measure,  
  dateTime,  
  period,  
  endDate (optional))
```


Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick Suite begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateMax(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```

periodToDateMedian

The periodToDateMedian function returns the median value of the specified measure for a given time granularity (for instance, a quarter) up to a point in time, relative to that period.

Syntax

```
periodToDateMedian(  
  measure,  
  dateTime,  
  period,
```

```
endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick Suite begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateMedian(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```

periodToDateMin

The periodToDateMin function returns the minimum value of the specified measure or date, or a given time granularity (for instance, a quarter) up to a point in time, relative to that period.

Syntax

```
periodToDateMin(
```

```
measure,  
dateTime,  
period,  
endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick Suite begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateMin(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-  
yyyy"))
```

periodToDatePercentile

The periodToDatePercentile function calculates the percentile based on the actual numbers in measure for a given time granularity (for instance, a quarter) up to a point in time, relative to that period. It uses the grouping and sorting that are applied in the field wells.

To return the nearest percentile value that is present in your dataset, use `periodToDatePercentile`. To return an exact percentile value that might not be present in your dataset, use `periodToDatePercentileCont` instead.

Syntax

```
periodToDatePercentile(  
  measure,  
  percentile,  
  dateTime,  
  period,  
  endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

percentile

The percentile value can be any numeric constant 0-100. A percentile of 50 computes the median value of the measure.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to `now()` if omitted.

Example

The following example calculates the week-to-date, 90th percentile of fare amount per payment type for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick Suite begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDatePercentile(fare_amount, 90, pickupDatetime, WEEK, parseDate("06-30-2021",  
"MM-dd-yyyy"))
```

periodToDatePercentileCont

The `periodToDatePercentileCont` function calculates percentile based on a continuous distribution of the numbers in the measure for a given time granularity (for instance, a quarter) up to a point in time in that period. It uses the grouping and sorting that are applied in the field wells.

To return an exact percentile value that might not be present in your dataset, use `periodToDatePercentileCont`. To return the nearest percentile value that is present in your dataset, use `periodToDatePercentile` instead.

Syntax

```
periodToDatePercentileCont(  
  measure,  
  percentile,  
  dateTime,  
  period,  
  endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

percentile

The percentile value can be any numeric constant 0-100. A percentile of 50 computes the median value of the measure.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date, 90th percentile of fare amount per payment type for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick Suite begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDatePercentileCont(fare_amount, 90, pickupDatetime, WEEK,  
    parseDate("06-30-2021", "MM-dd-yyyy"))
```

periodToDateStDev

The periodToDateStDev function calculates the standard deviation of the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time, based on a sample and relative to that period.

Syntax

```
periodToDateStDev(  
    measure,  
    dateTime,  
    period,  
    endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick Suite begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateStDev(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```

periodToDateStDevP

The periodToDateStDevP function calculates the population standard deviation of the set of numbers in the specified measure, for a given time granularity (for instance, a quarter) up to a point in time, based on a sample in that period.

Syntax

```
periodToDateStDevP(  
  measure,  
  dateTime,
```

```
period,  
endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick Suite begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateStDevP(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```

periodToDateSum

The periodToDateSum function adds the specified measure for a given time granularity (for instance, a quarter) up to a point in time, relative to that period.

Syntax

```
periodToDateSum(
```



```
measure,  
dateTime,  
period,  
endDate)
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following function calculates the week to date sum of fare amount per payment, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick Suite begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateSum(fare_amount, pickUpDateTime, WEEK, parseDate("06-30-2021", "MM-dd-  
yyyy"))
```

periodToDateVar

The periodToDateVar function calculates the sample variance of the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time in that period.

Syntax

```
periodToDateVar(  
    measure,  
    dateTime,  
    period,  
    endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to now() if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick Suite begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateVar(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```

periodToDateVarP

The `periodToDateVarP` function calculates the population variance of the set of numbers in the specified measure for a given time granularity (for instance, a quarter) up to a point in time, relevant to that period.

Syntax

```
periodToDateVarP(  
  measure,  
  dateTime,  
  period,  
  endDate (optional))
```

Arguments

measure

The argument must be a field. Null values are omitted from the results. Literal values don't work.

dateTime

The Date dimension over which you're computing PeriodToDate aggregations.

period

The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

endDate

(Optional) The date dimension that you're ending computing periodToDate aggregations. It defaults to `now()` if omitted.

Example

The following example calculates the week-to-date minimum fare amount per payment type, for the week of 06-30-21. For simplicity in the example, we filtered out only a single payment. 06-30-21 is Wednesday. Quick Suite begins the week on Sundays. In our example, that is 06-27-21.

```
periodToDateVarP(fare_amount, pickUpDatetime, WEEK, parseDate("06-30-2021", "MM-dd-yyyy"))
```

stdev

The `stdev` function calculates the standard deviation of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a sample.

Syntax

```
stdev(measure, [group-by level])
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example returns the standard deviation of test scores for a class, using a sample of the test scores recorded.

```
stdev({Score})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the standard deviation of test scores at the subject level, but not across other dimensions (Class) in the visual.

```
stdev({Score}, [Subject])
```

stdevp

The stdevp function calculates the population standard deviation of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions.

Syntax

```
stdevp(measure, [group-by level])
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example returns the standard deviation of test scores for a class using all the scores recorded.

```
stdevp({Score})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the standard deviation of test scores at the subject level, but not across other dimensions (Class) in the visual using all the scores recorded.

```
stdevp({Score}, [Subject])
```

stdevIf

Based on a conditional statement, the `stdevIf` function calculates the standard deviation of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a sample.

Syntax

```
stdevIf(measure, conditions)
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

stdevpIf

Based on a conditional statement, the `stdevpIf` function calculates the standard deviation of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a biased population.

Syntax

```
stdevpIf(measure, conditions)
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

sum

The sum function adds the set of numbers in the specified measure, grouped by the chosen dimension or dimensions. For example, `sum(profit amount)` returns the total profit amount grouped by the (optional) chosen dimension.

Syntax

```
sum(measure, [group-by level])
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example returns the sum of sales.

```
sum({Sales})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example calculates the sum of sales at the Country level, but not across other dimensions (Region and Product) in the visual.

```
sum(Sales, [Country])
```

sumIf

Based on a conditional statement, the sumIf function adds the set of numbers in the specified measure, grouped by the chosen dimension or dimensions. For example, `sumIf(ProdRev, CalendarDay >= ${BasePeriodStartDate} AND CalendarDay <= ${BasePeriodEndDate} AND SourcingType <> 'Indirect')` returns the total profit amount grouped by the (optional) chosen dimension, if the condition evaluates to true.

Syntax

```
sumIf(measure, conditions)
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

Examples

The following example uses a calculated field with sumIf to display the sales amount if Segment is equal to SMB.

```
sumIf(Sales, Segment='SMB')
```

The following example uses a calculated field with sumIf to display the sales amount if Segment is equal to SMB and Order Date greater than year 2022.

```
sumIf(Sales, Segment='SMB' AND {Order Date} >='2022-01-01')
```

var

The var function calculates the sample variance of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions.

Syntax

```
var(measure, [group-by level])
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example returns the variance of a sample of test scores.

```
var({Scores})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example returns the variance of a sample of test scores at the subject level, but not across other dimensions (Class) in the visual.

```
var({Scores}, [Subject])
```

varIf

Based on a conditional statement, the `varIf` function calculates the variance of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a sample.

Syntax

```
varIf(measure, conditions)
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

varp

The `varp` function calculates the population variance of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions.

Syntax

```
varp(measure, [group-by level])
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

group-by level

(Optional) Specifies the level to group the aggregation by. The level added can be any dimension or dimensions independent of the dimensions added to the visual.

The argument must be a dimension field. The group-by level must be enclosed in square brackets []. For more information, see [Level-aware calculation - aggregate \(LAC-A\) functions](#).

Examples

The following example returns the variance of a population of test scores.

```
varp({Scores})
```

You can also specify at what level to group the computation using one or more dimensions in the view or in your dataset. This is called a LAC-A function. For more information about LAC-A functions, see [Level-aware calculation - aggregate \(LAC-A\) functions](#). The following example returns the variance of a population test scores at the subject level, but not across other dimensions (Class) in the visual.

```
varp({Scores}, [Subject])
```

varpIf

Based on a conditional statement, the `varpIf` function calculates the variance of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a biased population.

Syntax

```
varpIf(measure, conditions)
```

Arguments

measure

The argument must be a measure. Null values are omitted from the results. Literal values don't work. The argument must be a field.

condition

One or more conditions in a single statement.

Table calculation functions

When you are analyzing data in a specific visual, you can apply table calculations to the current set of data to discover how dimensions influence measures or each other. *Visualized data* is your result set based on your current dataset, with all the filters, field selections, and customizations applied. To see exactly what this result set is, you can export your visual to a file. A *table calculation function* performs operations on the data to reveal relationships between fields.

In this section, you can find a list of the functions available in table calculations that you can perform on visualized data in Amazon Quick Suite.

To view a list of functions sorted by category, with brief definitions, see [Functions by category](#).

Topics

- [difference](#)
- [distinctCountOver](#)
- [lag](#)
- [lead](#)
- [percentDifference](#)
- [avgOver](#)
- [countOver](#)
- [maxOver](#)
- [minOver](#)
- [percentileOver](#)
- [percentileContOver](#)
- [percentileDiscOver](#)
- [percentOfTotal](#)
- [periodOverPeriodDifference](#)
- [periodOverPeriodLastValue](#)
- [periodOverPeriodPercentDifference](#)
- [periodToDateAvgOverTime](#)
- [periodToDateCountOverTime](#)
- [periodToDateMaxOverTime](#)
- [periodToDateMinOverTime](#)
- [periodToDateSumOverTime](#)
- [stdevOver](#)
- [stdevpOver](#)
- [varOver](#)
- [varpOver](#)
- [sumOver](#)

- [denseRank](#)
- [rank](#)
- [percentileRank](#)
- [runningAvg](#)
- [runningCount](#)
- [runningMax](#)
- [runningMin](#)
- [runningSum](#)
- [firstValue](#)
- [lastValue](#)
- [windowAvg](#)
- [windowCount](#)
- [windowMax](#)
- [windowMin](#)
- [windowSum](#)

difference

The difference function calculates the difference between a measure based on one set of partitions and sorts, and a measure based on another.

The difference function is supported for use with analyses based on SPICE and direct query data sets.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
difference
(
    measure
    ,[ sortorder_field ASC_or_DESC, ... ]
    ,lookup_index,
    ,[ partition field, ... ]
)
```

Arguments

measure

An aggregated measure that you want to see the difference for.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

lookup index

The lookup index can be positive or negative, indicating a following row in the sort (positive) or a previous row in the sort (negative). The lookup index can be 1–2,147,483,647. For the engines MySQL, MariaDB and Aurora MySQL-Compatible Edition, the lookup index is limited to just 1.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the difference between of `sum({Billed Amount})`, sorted by `Customer Region` ascending, compared to the next row, and partitioned by `Service Line`.

```
difference(  
    sum( {Billed Amount} ),  
    [{Customer Region} ASC],  
    1,  
    [{Service Line}]  
)
```

The following example calculates the difference between `Billed Amount` compared to the next line, partitioned by (`[{Customer Region}]`). The fields in the table calculation are in the field wells of the visual.

```
difference(  
    sum( {Billed Amount} ),  
    [{Customer Region} ASC],  
    1  
)
```

The red highlights show how each amount is added ($a + b = c$) to show the difference between amounts a and c.

distinctCountOver

The `distinctCountOver` function calculates the distinct count of the operand partitioned by the specified attributes at a specified level. Supported levels are `PRE_FILTER` and `PRE_AGG`. The operand must be unaggregated.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
distinctCountOver  
(  
    measure or dimension field  
    ,[ partition_field, ... ]  
    ,calculation level  
)
```

Arguments

measure or dimension field

The measure or dimension that you want to do the calculation for, for example `{Sales Amt}`. Valid values are `PRE_FILTER` and `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.

This value defaults to POST_AGG_FILTER when blank. POST_AGG_FILTER is not a valid level for this operation and will result in an error message. For more information, see [Using level-aware calculations in Amazon Quick Suite](#).

Example

The following example gets the distinct count of Sales partitioned over City and State at the PRE_AGG level.

```
distinctCountOver
(
  Sales,
  [City, State], PRE_AGG
)
```

lag

The lag function calculates the lag (previous) value for a measure based on specified partitions and sorts.

lag is supported for use with analyses based on SPICE and direct query data sets.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
lag
(
  lag
  (
    measure
    , [ sortorder_field ASC_or_DESC, ... ]
    , lookup_index
    , [ partition_field, ... ]
  )
)
```


Arguments

measure

The measure that you want to get the lag for. This can include an aggregate, for example `sum({Sales Amt})`.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

lookup index

The lookup index can be positive or negative, indicating a following row in the sort (positive) or a previous row in the sort (negative). The lookup index can be 1–2,147,483,647. For the engines MySQL, MariaDB, and Amazon Aurora MySQL-Compatible Edition, the lookup index is limited to just 1.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the previous `sum(sales)`, partitioned by the state of origin, in the ascending sort order on `cancellation_code`.

```
lag
(
    sum(Sales),
    [cancellation_code ASC],
    1,
    [origin_state_nm]
)
```

The following example uses a calculated field with `lag` to display sales amount for the previous row next to the amount for the current row, sorted by `Order Date`. The fields in the table calculation are in the field wells of the visual.

```
lag(  
    sum({Sales}),  
    [{Order Date} ASC],  
    1  
)
```

The following screenshot shows the results of the example.

The following example uses a calculated field with `lag` to display the sales amount for the previous row next to the amount for the current row, sorted by `Order Date` partitioned by `Segment`.

```
lag  
(  
    sum(Sales),  
    [Order Date ASC],  
    1, [Segment]  
)
```

The following screenshot shows the results of the example.

lead

The `lead` function calculates the lead (following) value for a measure based on specified partitions and sorts.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
lead  
(  
    measure  
    ,[ sortorder_field ASC_or_DESC, ... ]  
    ,lookup_index,  
    ,[ partition_field, ... ]
```

```
)
```

Arguments

measure

The measure that you want to get the lead for. This can include an aggregate, for example `sum({Sales Amt})`.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

lookup index

The lookup index can be positive or negative, indicating a following row in the sort (positive) or a previous row in the sort (negative). The lookup index can be 1–2,147,483,647. For the engines MySQL, MariaDB, and Amazon Aurora MySQL-Compatible Edition, the lookup index is limited to just 1.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the next `sum(sales)`, partitioned by the state of origin, in the ascending sort order on `cancellation_code`.

```
lead
(
  sum(sales),
  [cancellation_code ASC],
  1,
  [origin_state_nm]
)
```

The following example uses a calculated field with lead to display the amount for the next row beside the amount for the current row, sorted by Customer Segment. The fields in the table calculation are in the field wells of the visual.

```
lead(  
    sum({Billed Amount}),  
    [{Customer Segment} ASC],  
    1  
)
```

The following screenshot shows the results of the example.

percentDifference

The percentDifference function calculates the percentage difference between the current value and a comparison value, based on partitions, sorts, and lookup index.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
percentDifference  
(  
    measure  
    ,[ sortorder_field ASC_or_DESC, ... ]  
    ,lookup index  
    ,[ partition_field, ... ]  
)
```

Arguments

measure

An aggregated measure that you want to see the percent difference for.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

lookup index

The lookup index can be positive or negative, indicating a following row in the sort (positive) or a previous row in the sort (negative). The lookup index can be 1–2,147,483,647. For the engines MySQL, MariaDB and Aurora MySQL-Compatible Edition, the lookup index is limited to just 1.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the percentage of difference between the sum(Sales) for the current and the previous State, sorted by Sales.

```
percentDifference
(
  sum(amount),
  [sum(amount) ASC],
  -1,
  [State]
)
```

The following example calculates the percent that a specific Billed Amount is of another Billed Amount, sorted by ([{Customer Region} ASC]). The fields in the table calculation are in the field wells of the visual.

```
percentDifference
(
  sum( {Billed Amount} ),
  [{Customer Region} ASC],
  1
)
```

The following screenshot shows the results of the example. The red letters show that the total Billed Amount for the Customer Region **APAC** is 24 percent less than the amount for the **EMEA** region.

avgOver

The `avgOver` function calculates the average of a measure partitioned by a list of dimensions.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
avgOver
(  
    measure  
    , [ partition_field, ... ]  
    , calculation level  
)
```

The following example shows the average Billed Amount over Customer Region. The fields in the table calculation are in the field wells of the visual.

```
avgOver
(  
    sum({Billed Amount}),  
    [{Customer Region}]  
)
```

The following screenshot shows the results of the example. With the addition of Service Line, the total amount billed for each is displayed, and the average of these three values displays in the calculated field.

Arguments

measure

The measure that you want to do the calculation for, for example `sum({Sales Amt})`. Use an aggregation if the calculation level is set to `NULL` or `POST_AGG_FILTER`. Don't use an aggregation if the calculation level is set to `PRE_FILTER` or `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to POST_AGG_FILTER when blank. For more information, see [Using level-aware calculations in Quick Suite](#).

Example

The following example gets the average sum(Sales) partitioned over City and State.

```
avgOver
(
    sum(Sales),
    [City, State]
)
```

countOver

The countOver function calculates the count of a dimension or measure partitioned by a list of dimensions.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
countOver
(
    measure or dimension field
    ,[ partition_field, ... ]
    ,calculation level
)
```

Arguments

measure or dimension field

The measure or dimension that you want to do the calculation for, for example `sum( {Sales Amt} )`. Use an aggregation if the calculation level is set to `NULL` or `POST_AGG_FILTER`. Don't use an aggregation if the calculation level is set to `PRE_FILTER` or `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to `POST_AGG_FILTER` when blank. For more information, see [Using level-aware calculations in Quick Suite](#).

Example

The following example gets the count of Sales partitioned over City and State.

```
countOver
(
  Sales,
  [City, State]
)
```

The following example gets the count of {County} partitioned over City and State.

```
countOver
(
  {County},
```



```
[City, State]
)
```

The following example shows the count of `Billed Amount` over `Customer Region`. The fields in the table calculation are in the field wells of the visual.

```
countOver
(
  sum({Billed Amount}),
  [{Customer Region}]
)
```

The following screenshot shows the results of the example. Because there are no other fields involved, the count is one for each region.

If you add additional fields, the count changes. In the following screenshot, we add `Customer Segment` and `Service Line`. Each of those fields contains three unique values. With 3 segments, 3 service lines, and 3 regions, the calculated field shows 9.

If you add the two additional fields to the partitioning fields in the calculated field, `countOver( sum({Billed Amount}), [{Customer Region}, {Customer Segment}, {Service Line}]`, then the count is again 1 for each row.

maxOver

The `maxOver` function calculates the maximum of a measure or date partitioned by a list of dimensions.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
maxOver
(
  measure
  , [ partition_field, ... ]
  , calculation level
```

```
)
```

Arguments

measure

The measure that you want to do the calculation for, for example `sum({Sales Amt})`. Use an aggregation if the calculation level is set to `NULL` or `POST_AGG_FILTER`. Don't use an aggregation if the calculation level is set to `PRE_FILTER` or `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to `POST_AGG_FILTER` when blank. For more information, see [Using level-aware calculations in Quick Suite](#).

Example

The following example calculates the maximum `sum(Sales)`, partitioned by `City` and `State`.

```
maxOver
(
    sum(Sales),
    [City, State]
)
```

The following example shows the maximum `Billed Amount` over `Customer Region`. The fields in the table calculation are in the field wells of the visual.

```
maxOver  
(  
    sum({Billed Amount}),  
    [{Customer Region}]  
)
```

The following screenshot shows the results of the example. With the addition of Service Line, the total amount billed for each is displayed, and the maximum of these three values displays in the calculated field.

minOver

The `minOver` function calculates the minimum of a measure or date partitioned by a list of dimensions.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
minOver  
(  
    measure  
    ,[ partition_field, ... ]  
    ,calculation level  
)
```

Arguments

measure

The measure that you want to do the calculation for, for example `sum({Sales Amt})`. Use an aggregation if the calculation level is set to `NULL` or `POST_AGG_FILTER`. Don't use an aggregation if the calculation level is set to `PRE_FILTER` or `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to **POST_AGG_FILTER** when blank. For more information, see [Using level-aware calculations in Quick Suite](#).

Example

The following example calculates the min sum(*Sales*), partitioned by *City* and *State*.

```
minOver
(
    sum(Sales),
    [City, State]
)
```

The following example shows the minimum Billed Amount over Customer Region. The fields in the table calculation are in the field wells of the visual.

```
minOver
(
    sum({Billed Amount}),
    [{Customer Region}]
)
```

The following screenshot shows the results of the example. With the addition of *Service Line*, the total amount billed for each is displayed, and the minimum of these three values displays in the calculated field.

percentileOver

The `percentileOver` function calculates the *n*th percentile of a measure partitioned by a list of dimensions. There are two varieties of the `percentileOver` calculation available in Quick Suite:

- [percentileContOver](#) uses linear interpolation to determine result.
- [percentileDiscOver](#) uses actual values to determine result.

The `percentileOver` function is an alias of `percentileDiscOver`.

percentileContOver

The `percentileContOver` function calculates the percentile based on the actual numbers in measure. It uses the grouping and sorting that are applied in the field wells. The result is partitioned by the specified dimension at the specified calculation level.

Use this function to answer the following question: Which actual data points are present in this percentile? To return the nearest percentile value that is present in your dataset, use `percentileDiscOver`. To return an exact percentile value that might not be present in your dataset, use `percentileContOver` instead.

Syntax

```
percentileContOver (  
    measure  
    , percentile-n  
    , [partition-by, ...]  
    , calculation-level  
)
```

Arguments

measure

Specifies a numeric value to use to compute the percentile. The argument must be a measure or metric. Nulls are ignored in the calculation.

percentile-n

The percentile value can be any numeric constant 0–100. A percentile value of 50 computes the median value of the measure.

partition-by

(Optional) One or more dimensions that you want to partition by, separated by commas. Each field in the list is enclosed in { } (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

calculation-level

Specifies where to perform the calculation in relation to the order of evaluation. There are three supported calculation levels:

- PRE_FILTER
- PRE_AGG
- POST_AGG_FILTER (default) – To use this calculation level, specify an aggregation on measure, for example `sum(measure)`.

PRE_FILTER and PRE_AGG are applied before the aggregation occurs in a visualization. For these two calculation levels, you can't specify an aggregation on measure in the calculated field expression. To learn more about calculation levels and when they apply, see [Order of evaluation in Amazon Quick Suite](#) and [Using level-aware calculations in Quick Suite](#).

Returns

The result of the function is a number.

Example of percentileContOver

The following example helps explain how percentileContOver works.

Example Comparing calculation levels for the median

The following example shows the median for a dimension (category) by using different calculation levels with the percentileContOver function. The percentile is 50. The dataset is filtered by a region field. The code for each calculated field is as follows:

- `example = left( category, 1 )` (A simplified example.)
- `pre_agg = percentileContOver ( {Revenue} , 50 , [ example ] , PRE_AGG)`
- `pre_filter = percentileContOver ( {Revenue} , 50 , [ example ] , PRE_FILTER)`
- `post_agg_filter = percentileContOver ( sum ( {Revenue} ) , 50 , [ example ] , POST_AGG_FILTER )`

example	pre_filter	pre_agg	post_agg_filter

0	106,728	119,667	4,117,579
1	102,898	95,946	2,307,547
2	97,807	93,963	554,570
3	101,043	112,585	2,709,057
4	96,533	99,214	3,598,358
5	106,293	97,296	1,875,648
6	97,118	69,159	1,320,672
7	100,201	90,557	969,807

percentileDiscOver

The `percentileDiscOver` function calculates the percentile based on the actual numbers in measure. It uses the grouping and sorting that are applied in the field wells. The result is partitioned by the specified dimension at the specified calculation level. The `percentileOver` function is an alias of `percentileDiscOver`.

Use this function to answer the following question: Which actual data points are present in this percentile? To return the nearest percentile value that is present in your dataset, use `percentileDiscOver`. To return an exact percentile value that might not be present in your dataset, use `percentileContOver` instead.

Syntax

```
percentileDiscOver (
    measure
    , percentile-n
    , [partition-by, ...]
    , calculation-level
)
```

Arguments

measure

Specifies a numeric value to use to compute the percentile. The argument must be a measure or metric. Nulls are ignored in the calculation.

percentile-n

The percentile value can be any numeric constant 0–100. A percentile value of 50 computes the median value of the measure.

partition-by

(Optional) One or more dimensions that you want to partition by, separated by commas. Each field in the list is enclosed in { } (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

calculation-level

Specifies where to perform the calculation in relation to the order of evaluation. There are three supported calculation levels:

- PRE_FILTER
- PRE_AGG
- POST_AGG_FILTER (default) – To use this calculation level, you need to specify an aggregation on measure, for example `sum(measure)`.

PRE_FILTER and PRE_AGG are applied before the aggregation occurs in a visualization. For these two calculation levels, you can't specify an aggregation on measure in the calculated field expression. To learn more about calculation levels and when they apply, see [Order of evaluation in Amazon Quick Suite](#) and [Using level-aware calculations in Quick Suite](#).

Returns

The result of the function is a number.

Example of percentileDiscOver

The following example helps explain how percentileDiscOver works.

Example Comparing calculation levels for the median

The following example shows the median for a dimension (category) by using different calculation levels with the percentileDiscOver function. The percentile is 50. The dataset is filtered by a region field. The code for each calculated field is as follows:

- `example = left( category, 1 )` (A simplified example.)
- `pre_agg = percentileDiscOver ( {Revenue} , 50 , [ example ] , PRE_AGG)`
- `pre_filter = percentileDiscOver ( {Revenue} , 50 , [ example ] , PRE_FILTER)`
- `post_agg_filter = percentileDiscOver ( sum ( {Revenue} ) , 50 , [ example ] , POST_AGG_FILTER )`

example	pre_filter	pre_agg	post_agg_filter
0	106,728	119,667	4,117,579
1	102,898	95,946	2,307,547
2	97,629	92,046	554,570
3	100,867	112,585	2,709,057
4	96,416	96,649	3,598,358
5	106,293	97,296	1,875,648
6	97,118	64,395	1,320,672
7	99,915	90,557	969,807

Example The median

The following example calculates the median (the 50th percentile) of Sales partitioned by City and State.

```
percentileDiscOver
(
  Sales,
  50,
  [City, State]
)
```

The following example calculates the 98th percentile of sum({Billed Amount}) partitioned by Customer Region. The fields in the table calculation are in the field wells of the visual.

```
percentileDiscOver
(
  sum({Billed Amount}),
  98,
  [{Customer Region}]
)
```

The following screenshot shows the how these two examples look on a chart.

percentOfTotal

The `percentOfTotal` function calculates the percentage a measure contributes to the total, based on the dimensions specified.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
percentOfTotal
(  
    measure  
    , [ partition_field, ... ]  
)
```

Arguments

measure

An aggregated measure that you want to see the percent of total for. Currently, the distinct count aggregation is not supported for `percentOfTotal`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example creates a calculation for the percent of total Sales contributed by each State.

```
percentOfTotal
(  
    sum(Sales),  
    [State]  
)
```

The following example calculates the percent that a specific Billed Amount is when compared to the total Billed Amount, partitioned by ([Service Line] ASC). The fields in the table calculation are in the field wells of the visual.

```
percentOfTotal
```

```
(  
    sum( {Billed Amount} ),  
    [{Service Line}]  
)
```

The following screenshot shows the results of the example. The red highlights show that the partition field with the value "Billing" has three entries, one for each region. The total billed amount for this service line is divided into three percentages, which total 100 percent. Percentages are rounded and might not always add up to exactly 100 percent.

periodOverPeriodDifference

The `periodOverPeriodDifference` function calculates the difference of a measure over two different time periods as specified by period granularity and offset. Unlike a difference calculation, this function uses a date-based offset instead of a fixed sized offset. This ensures that only the correct dates are compared, even if data points are missing in the dataset.

Syntax

```
periodOverPeriodDifference(  
    measure,  
    date,  
    period,  
    offset)
```

Arguments

measure

An aggregated measure that you want to perform the `periodOverPeriod` calculation on.

dateTime

The Date dimension over which we are computing Period-Over-Period calculations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

The default value is the visual date dimension granularity.

offset

(Optional) The offset can be a positive or negative integer representing the prior time period (specified by period) that you want to compare against. For instance, period of a quarter with offset 1 means comparing against the previous quarter.

The default value is 1.

Example

The following example uses a calculated field `PeriodOverPeriod` to display the sales amount difference of yesterday

```
periodOverPeriodDifference(sum(Sales), {Order Date})
```

The following example uses a calculated field `PeriodOverPeriod` to display the sales amount difference of previous 2 months. Below example is comparing sales of Mar2020 with Jan2020.

```
periodOverPeriodDifference(sum(Sales), {Order Date}, MONTH, 1)
```

periodOverPeriodLastValue

The `periodOverPeriodLastValue` function calculates the last (previous) value of a measure from the previous time period as specified by the period granularity and offset. This function uses a date-based offset instead of a fixed sized offset. This ensures only the correct dates are compared, even if data points are missing in the dataset.

Syntax

```
periodOverPeriodLastValue(  
  measure,  
  date,  
  period,  
  offset)
```

Arguments

measure

An aggregated measure that you want to see the difference for.

date

The date dimension over which you're computing periodOverPeriod calculations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

This argument defaults to the granularity of the visual aggregation

offset

(Optional) The offset can a positive or negative integer representing the prior time period (specified by period) that you want to compare against. For instance, period of a quarter with offset 1 means comparing against the previous quarter.

This argument default value is 1.

Example

The following example calculates the month over month value in sales with the visual dimension granularity and default offset of 1.

```
periodOverPeriodLastValue(sum(Sales), {Order Date})
```

The following example calculates the month over month value in sales with a fixed granularity of MONTH and fixed offset of 1.

```
periodOverPeriodLastValue(sum(Sales), {Order Date},MONTH, 1)
```

periodOverPeriodPercentDifference

The periodOverPeriodPercentDifference function calculates the percent difference of a measure over two different time periods as specified by the period granularity and offset. Unlike

percentDifference, this function uses a date-based offset instead of a fixed sized offset. This ensures only the correct dates are compared, even if data points are missing in the dataset.

Syntax

```
periodOverPeriodPercentDifference(  
  measure,  
  date,  
  period,  
  offset)
```

Arguments

measure

An aggregated measure that you want to see the difference for.

date

The date dimension over which you're computing periodOverPeriod calculations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

This argument defaults to the granularity of the visual aggregation

offset

(Optional) The offset can a positive or negative integer representing the prior time period (specified by period) that you want to compare against. For instance, period of a quarter with offset 1 means comparing against the previous quarter.

This argument default value is 1.

Example

The following example calculates the month over month percent difference in sales with the visual dimension granularity and default offset of 1.

```
periodOverPeriodPercentDifference(sum(Sales), {Order Date})
```

The following example calculates the month over month percent difference in sales with a fixed granularity of MONTH and fixed offset of 1.

```
periodOverPeriodPercentDifference(sum(Sales), {Order Date}, MONTH, 1)
```

periodToDateAvgOverTime

The `periodToDateAvgOverTime` function calculates the average of a measure for a given time granularity (for instance, a quarter) up to a point in time.

Syntax

```
periodToDateAvgOverTime(  
    measure,  
    dateTime,  
    period)
```

Arguments

measure

An aggregated measure that you want to do the calculation

dateTime

The date dimension over which you're computing PeriodOverTime calculations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

The default value is the visual's date dimension granularity.

Example

The following function calculates the average fare amount month over month.

```
periodToDateAvgOverTime(sum({fare_amount}), pickupDatetime, MONTH)
```

periodToDateCountOverTime

The `periodToDateCountOverTime` function calculates the count of a dimension or measure for a given time granularity (for instance, a quarter) up to a point in time.

Syntax

```
periodToDateCountOverTime(  
  measure,  
  dateTime,  
  period)
```

Arguments

measure

An aggregated measure that you want to do the calculation

dateTime

The date dimension over which you're computing PeriodOverTime calculations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

The default value is the visual's date dimension granularity.

Example

The following example calculates the count of vendors month over month.

```
periodToDateCountOverTime(count(vendorid), pickupDatetime, MONTH)
```


periodToDateMaxOverTime

The `periodToDateMaxOverTime` function calculates the maximum of a measure for a given time granularity (for instance, a quarter) up to a point in time.

Syntax

```
periodToDateMaxOverTime(  
    measure,  
    dateTime,  
    period)
```

Arguments

measure

An aggregated measure that you want to do the calculation

dateTime

The date dimension over which you're computing `PeriodOverTime` calculations.

period

(Optional) The time period across which you're computing the computation. Granularity of `YEAR` means `YearToDate` computation, `Quarter` means `QuarterToDate`, and so on. Valid granularities include `YEAR`, `QUARTER`, `MONTH`, `WEEK`, `DAY`, `HOUR`, `MINUTE`, and `SECONDS`.

The default value is the visual's date dimension granularity.

Example

The following example calculates the maximum fare amount month over month.

```
periodToDateMaxOverTime(max({fare_amount}), pickupDatetime, MONTH)
```

periodToDateMinOverTime

The `periodToDateMinOverTime` function calculates the minimum of a measure for a given time granularity (for instance, a quarter) up to a point in time.

Syntax

```
periodToDateMinOverTime(  
    measure,  
    dateTime,  
    period)
```

Arguments

measure

An aggregated measure that you want to do the calculation

dateTime

The date dimension over which you're computing PeriodOverTime calculations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

The default value is the visual's date dimension granularity.

Example

The following example calculates the minimum fare amount month over month.

```
periodToDateMinOverTime(min({fare_amount}), pickupDatetime, MONTH)
```

periodToDateSumOverTime

The periodToDateSumOverTime function calculates the sum of a measure for a given time granularity (for instance, a quarter) up to a point in time.

Syntax

```
periodToDateSumOverTime(  
    measure,  
    dateTime,  
    period)
```

```
measure,  
dateTime,  
period)
```

Arguments

measure

An aggregated measure that you want to do the calculation

dateTime

The date dimension over which you're computing PeriodOverTime calculations.

period

(Optional) The time period across which you're computing the computation. Granularity of YEAR means YearToDate computation, Quarter means QuarterToDate, and so on. Valid granularities include YEAR, QUARTER, MONTH, WEEK, DAY, HOUR, MINUTE, and SECONDS.

The default value is the visual's date dimension granularity.

Example

The following function returns the total fare amount month over month.

```
periodToDateSumOverTime(sum({fare_amount}), pickupDatetime, MONTH)
```

stdevOver

The `stdevOver` function calculates the standard deviation of the specified measure, partitioned by the chosen attribute or attributes, based on a sample.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
stdevOver  
(  
    measure
```

```
,[ partition_field, ... ]  
,calculation level  
)
```

Arguments

measure

The measure that you want to do the calculation for, for example `sum( {Sales Amt} )`. Use an aggregation if the calculation level is set to `NULL` or `POST_AGG_FILTER`. Don't use an aggregation if the calculation level is set to `PRE_FILTER` or `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (default) table calculations are computed when the visuals display.

This value defaults to `POST_AGG_FILTER` when blank. For more information, see [Using level-aware calculations in Quick Suite](#).

Example

The following example calculates the standard deviation of `sum(Sales)`, partitioned by `City` and `State`, based on a sample..

```
stdevOver  
(  
    sum(Sales),  
    [City, State]  
)
```

The following example calculates the standard deviation of Billed Amount over Customer Region, based on a sample. The fields in the table calculation are in the field wells of the visual.

```
stdevOver
(  
    sum({Billed Amount}),  
    [{Customer Region}]  
)
```

stdevpOver

The stdevpOver function calculates the standard deviation of the specified measure, partitioned by the chosen attribute or attributes, based on a biased population.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
stdevpOver
(  
    measure  
    ,[ partition_field, ... ]  
    ,calculation level  
)
```

Arguments

measure

The measure that you want to do the calculation for, for example sum({Sales Amt}). Use an aggregation if the calculation level is set to NULL or POST_AGG_FILTER. Don't use an aggregation if the calculation level is set to PRE_FILTER or PRE_AGG.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (default) table calculations are computed when the visuals display.

This value defaults to **POST_AGG_FILTER** when blank. For more information, see [Using level-aware calculations in Quick Suite](#).

Example

The following example calculates the standard deviation of `sum(Sales)`, partitioned by `City` and `State`, based on a biased population.

```
stdevpOver
(  
    sum(Sales),  
    [City, State]  
)
```

The following example calculates the standard deviation of `Billed Amount` over `Customer Region`, based on a biased population. The fields in the table calculation are in the field wells of the visual.

```
stdevpOver
(  
    sum({Billed Amount}),  
    [{Customer Region}]  
)
```

varOver

The `varOver` function calculates the variance of the specified measure, partitioned by the chosen attribute or attributes, based on a sample.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
varOver
```

```
(  
    measure  
    , [ partition_field, ... ]  
    , calculation_level  
)
```

Arguments

measure

The measure that you want to do the calculation for, for example `sum({Sales Amt})`. Use an aggregation if the calculation level is set to `NULL` or `POST_AGG_FILTER`. Don't use an aggregation if the calculation level is set to `PRE_FILTER` or `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to `POST_AGG_FILTER` when blank. For more information, see [Using level-aware calculations in Quick Suite](#).

Example

The following example calculates the variance of `sum(Sales)`, partitioned by `City` and `State`, based on a sample.

```
varOver  
(  
    sum(Sales),  
    [City, State]
```

```
)
```

The following example calculates the variance of Billed Amount over Customer Region, based on a sample. The fields in the table calculation are in the field wells of the visual.

```
varOver  
(  
    sum({Billed Amount}),  
    [{Customer Region}]  
)
```

varpOver

The varpOver function calculates the variance of the specified measure, partitioned by the chosen attribute or attributes, based on a biased population.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
varpOver  
(  
    measure  
    ,[ partition_field, ... ]  
    ,calculation level  
)
```

Arguments

measure

The measure that you want to do the calculation for, for example sum({Sales Amt}). Use an aggregation if the calculation level is set to NULL or POST_AGG_FILTER. Don't use an aggregation if the calculation level is set to PRE_FILTER or PRE_AGG.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to **POST_AGG_FILTER** when blank. For more information, see [Using level-aware calculations in Quick Suite](#).

Example

The following example calculates the variance of `sum(Sales)`, partitioned by `City` and `State`, based on a biased population.

```
varpOver  
(  
    sum(Sales),  
    [City, State]  
)
```

The following example calculates the variance of `Billed Amount` over `Customer Region`, based on a biased population. The fields in the table calculation are in the field wells of the visual.

```
varpOver  
(  
    sum({Billed Amount}),  
    [{Customer Region}]  
)
```

sumOver

The `sumOver` function calculates the sum of a measure partitioned by a list of dimensions.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
sumOver
```

```
(  
    measure  
    , [ partition_field, ... ]  
    , calculation level  
)
```

Arguments

measure

The measure that you want to do the calculation for, for example `sum({Sales Amt})`. Use an aggregation if the calculation level is set to `NULL` or `POST_AGG_FILTER`. Don't use an aggregation if the calculation level is set to `PRE_FILTER` or `PRE_AGG`.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it is more than one word. The entire list is enclosed in `[ ]` (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (default) table calculations are computed when the visuals display.

This value defaults to `POST_AGG_FILTER` when blank. For more information, see [Using level-aware calculations in Quick Suite](#).

Example

The following example calculates the sum of `sum(Sales)`, partitioned by `City` and `State`.

```
sumOver  
(  
    sum(Sales),  
    [City, State]  
)
```

The following example sums Billed Amount over Customer Region. The fields in the table calculation are in the field wells of the visual.

```
sumOver
(
    sum({Billed Amount}),
    [{Customer Region}]
)
```

The following screenshot shows the results of the example. With the addition of Customer Segment, the total amount billed for each is summed for the Customer Region, and displays in the calculated field.

denseRank

The denseRank function calculates the rank of a measure or a dimension in comparison to the specified partitions. It counts each item only once, ignoring duplicates, and assigns a rank "without holes" so that duplicate values share the same rank.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
denseRank
(
    [ sort_order_field ASC_or_DESC, ... ]
    ,[ partition_field, ... ]
)
```

Arguments

sort order field

One or more aggregated fields, either measures or dimensions or both, that you want to sort the data by, separated by commas. You can either specify ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to POST_AGG_FILTER when blank. For more information, see [Using level-aware calculations in Quick Suite](#).

Example

The following example densely ranks `max(Sales)`, based on a descending sort order, by State and City. Any cities with the same `max(Sales)` are assigned the same rank, and the next city is ranked consecutively after them. For example, if three cities share the same ranking, the fourth city is ranked as second.

```
denseRank
(  
  [max(Sales) DESC],  
  [State, City]  
)
```

The following example densely ranks `max(Sales)`, based on a descending sort order, by State. Any states with the same `max(Sales)` are assigned the same rank, and the next is ranked consecutively after them. For example, if three states share the same ranking, the fourth state is ranked as second.

```
denseRank
(  
  [max(Sales) DESC],
```

```
[State]
)
```

rank

The `rank` function calculates the rank of a measure or a dimension in comparison to the specified partitions. It counts each item, even duplicates, once and assigns a rank "with holes" to make up for duplicate values.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
rank
(
  [ sort_order_field ASC_or_DESC, ... ]
  ,[ partition_field, ... ]
)
```

Arguments

sort order field

One or more aggregated measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.

- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to POST_AGG_FILTER when blank. For more information, see [Using level-aware calculations in Quick Suite](#).

Example

The following example ranks `max(Sales)`, based on a descending sort order, by State and City, within the State of **WA**. Any cities with the same `max(Sales)` are assigned the same rank, but the next rank includes the count of all previously existing ranks. For example, if three cities share the same ranking, the fourth city is ranked as fourth.

```
rank
(
  [max(Sales) DESC],
  [State, City]
)
```

The following example ranks `max(Sales)`, based on an ascending sort order, by State. Any states with the same `max(Sales)` are assigned the same rank, but the next rank includes the count of all previously existing ranks. For example, if three states share the same ranking, the fourth state is ranked as fourth.

```
rank
(
  [max(Sales) ASC],
  [State]
)
```

The following example ranks Customer Region by total Billed Amount. The fields in the table calculation are in the field wells of the visual.

```
rank(
  [sum({Billed Amount}) DESC]
)
```

The following screenshot shows the results of the example, along with the total Billed Amount so you can see how each region ranks.

percentileRank

The `percentileRank` function calculates the percentile rank of a measure or a dimension in comparison to the specified partitions. The percentile rank value(*x*) indicates that the current item is above *x*% of values in the specified partition. The percentile rank value ranges from 0 (inclusive) to 100 (exclusive).

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
percentileRank
(  
    [ sort_order_field ASC_or_DESC, ... ]  
    , [ {partition_field}, ... ]  
)
```

Arguments

sort order field

One or more aggregated measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

calculation level

(Optional) Specifies the calculation level to use:

- **PRE_FILTER** – Prefilter calculations are computed before the dataset filters.
- **PRE_AGG** – Preaggregate calculations are computed before applying aggregations and top and bottom *N* filters to the visuals.
- **POST_AGG_FILTER** – (Default) Table calculations are computed when the visuals display.

This value defaults to POST_AGG_FILTER when blank. For more information, see [Using level-aware calculations in Quick Suite](#).

Example

The following example does a percentile ranking of max(Sales) in descending order, by State.

```
percentileRank
(
    [max(Sales) DESC],
    [State]
)
```

The following example does a percentile ranking of Customer Region by total Billed Amount. The fields in the table calculation are in the field wells of the visual.

```
percentileRank(
    [sum({Billed Amount}) DESC],
    [{Customer Region}]
)
```

The following screenshot shows the results of the example, along with the total Billed Amount so you can see how each region compares.

runningAvg

The runningAvg function calculates a running average for a measure based on the specified dimensions and sort orders.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
runningAvg
(
    measure
    ,[ sortorder_field ASC_or_DESC, ... ]
    ,[ partition_field, ... ]
)
```


Arguments

measure

An aggregated measure that you want to see the running average for.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates a running average of `sum(Sales)`, sorted by Sales, partitioned by City and State.

```
runningAvg
(
  sum(Sales),
  [Sales ASC],
  [City, State]
)
```

The following example calculates a running average of `Billed Amount`, sorted by month (`[truncDate("MM",Date) ASC]`). The fields in the table calculation are in the field wells of the visual.

```
runningAvg
(
  sum({Billed Amount}),
  [truncDate("MM",Date) ASC]
)
```

runningCount

The `runningCount` function calculates a running count for a measure or dimension, based on the specified dimensions and sort orders.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
runningCount
(  
  measure_or_dimension  
  , [ sortorder_field ASC_or_DESC, ... ]  
  , [ partition_field, ... ]  
)
```

Arguments

measure or dimension

An aggregated measure or dimension that you want to see the running count for.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates a running count of `sum(Sales)`, sorted by `Sales`, partitioned by `City` and `State`.

```
runningCount
```

```
(  
    sum(Sales),  
    [Sales ASC],  
    [City, State]  
)
```

The following example calculates a running count of Billed Amount, sorted by month ([truncDate("MM",Date) ASC]). The fields in the table calculation are in the field wells of the visual.

```
runningCount  
(  
    sum({Billed Amount}),  
    [truncDate("MM",Date) ASC]  
)
```

runningMax

The runningMax function calculates a running maximum for a measure based on the specified dimensions and sort orders.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
runningMax  
(  
    measure  
    ,[ sortorder_field ASC_or_DESC, ... ]  
    ,[ partition_field, ... ]  
)
```

Arguments

measure

An aggregated measure that you want to see the running maximum for.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates a running maximum of `sum(Sales)`, sorted by `Sales`, partitioned by `City` and `State`.

```
runningMax
(  
  sum(Sales),  
  [Sales ASC],  
  [City, State]  
)
```

The following example calculates a running maximum of `Billed Amount`, sorted by month (`[truncDate("MM",Date) ASC]`). The fields in the table calculation are in the field wells of the visual.

```
runningMax
(  
  sum({Billed Amount}),  
  [truncDate("MM",Date) ASC]  
)
```

runningMin

The `runningMin` function calculates a running minimum for a measure based on the specified dimensions and sort orders.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
runningMin
```

```
(  
  measure  
  , [ sortorder_field ASC_or_DESC, ... ]  
  , [ partition_field, ... ]  
)
```

Arguments

measure

An aggregated measure that you want to see the running minimum for.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates a running minimum of `sum(Sales)`, sorted by Sales, partitioned by City and State.

```
runningMin  
(  
  sum(Sales),  
  [Sales ASC],  
  [City, State]  
)
```

The following example calculates a running minimum of `Billed Amount`, sorted by month (`[truncDate("MM", Date) ASC]`). The fields in the table calculation are in the field wells of the visual.

```
runningMin
(  
  sum({Billed Amount}),  
  [truncDate("MM",Date) ASC]  
)
```

runningSum

The `runningSum` function calculates a running sum for a measure based on the specified dimensions and sort orders.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
runningSum
(  
  measure  
  ,[ sortorder_field ASC_or_DESC, ... ]  
  ,[ partition_field, ... ]  
)
```

Arguments

measure

An aggregated measure that you want to see the running sum for.

sort order field

One or more measures and dimensions that you want to sort the data by, separated by commas. You can specify either ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates a running sum of `sum(Sales)`, sorted by Sales, partitioned by City and State.

```
runningSum
(  
    sum(Sales),  
    [Sales ASC],  
    [City, State]  
)
```

The following example calculates a running sum of Billed Amount, sorted by month (`[truncDate("MM",Date) ASC]`). The fields in the table calculation are in the field wells of the visual.

```
runningSum
(  
    sum({Billed Amount}),  
    [truncDate("MM",Date) ASC]  
)
```

The following screenshot shows the results of the example. The red labels show how each amount is added ($a + b = c$) to the next amount, resulting in a new total.

firstValue

The `firstValue` function calculates the first value of the aggregated measure or dimension partitioned and sorted by specified attributes.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
firstValue
(  
    aggregated measure or dimension,  
    [ sort_attribute ASC_or_DESC, ... ],  
    [ partition_by_attribute, ... ]  
)
```

Arguments

aggregated measure or dimension

An aggregated measure or dimension that you want to see the first value for.

sort attribute

One or more aggregated fields, either measures or dimensions or both, that you want to sort the data by, separated by commas. You can either specify ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it's more than one word. The entire list is enclosed in [] (square brackets).

partition by attribute

(Optional) One or more measure or dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the first Destination Airport, sorted by Flight Date, partitioned by Flight Date ascending and Origin Airport.

```
firstValue(  
    {Destination Airport}  
    [{Flight Date} ASC],  
    [  
        {Origin Airport},  
        {Flight Date}  
    ]  
)
```

lastValue

The lastValue function calculates the last value of the aggregated measure or dimension partitioned and sorted by specified attributes.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
lastValue  
(  
    aggregated measure or dimension,  
    [ sort_attribute ASC_or_DESC, ... ],  
    [ partition_by_attribute, ... ]  
)
```

Arguments

aggregated measure or dimension

An aggregated measure or dimension that you want to see the last value for.

sort attribute

One or more aggregated fields, either measures or dimensions or both, that you want to sort the data by, separated by commas. You can either specify ascending (ASC) or descending (DESC) sort order.

Each field in the list is enclosed in {} (curly braces), if it's more than one word. The entire list is enclosed in [] (square brackets).

partition by attribute

(Optional) One or more measures or dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the last value for Destination Airport. This calculation is sorted by the Flight Date value and partitioned by the Flight Date value sorted in ascending order and the Origin Airport value.

```
lastValue(  
    [{Destination Airport}],
```

```
[{Flight Date} ASC],  
[  
    {Origin Airport},  
    truncate('DAY', {Flight Date})  
]  
)
```

windowAvg

The windowAvg function calculates the average of the aggregated measure in a custom window that is partitioned and sorted by specified attributes. Usually, you use custom window functions on a time series, where your visual shows a metric and a date field. For example, you can use windowAvg to calculate a moving average, which is often used to smooth out the noise in a line chart.

Window functions aren't supported for MySQL versions earlier than 8 and MariaDB versions earlier than 10.2.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
windowAvg  
(  
    measure  
    , [sort_order_field ASC/DESC, ...]  
    , start_index  
    , end_index  
    , [ partition_field, ... ]  
)
```

Arguments

measure

The aggregated metric that you want to get the average for, for example sum({Revenue}).

sort attribute

One or more aggregated fields, either measures or dimensions or both, that you want to sort the data by, separated by commas. You can either specify ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in { } (curly braces), if it's more than one word. The entire list is enclosed in [] (square brackets).

start index

The start index is a positive integer, indicating n rows above the current row. The start index counts the available data points above the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

end index

The end index is a positive integer, indicating n rows below the current row. The end index counts the available data points below the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in { } (curly braces), if it's more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the moving average of `sum(Revenue)`, partitioned by `SaleDate`. The calculation includes three rows above and two row below of the current row.

```
windowAvg
(
    sum(Revenue),
    [SaleDate ASC],
    3,
    2
)
```

The following screenshot shows the results of this moving average example. The `sum(Revenue)` field is added to the chart to show the difference between the revenue and the moving average of revenue.

windowCount

The windowCount function calculates the count of the aggregated measure or dimension in a custom window that is partitioned and sorted by specified attributes. Usually, you use custom window functions on a time series, where your visual shows a metric and a date field.

Window functions aren't supported for MySQL versions earlier than 8 and MariaDB versions earlier than 10.2.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
windowCount
(
    measure_or_dimension
    , [sort_order_field ASC/DESC, ...]
    , start_index
    , end_index
    , [partition_field, ... ]
)
```

Arguments

measure or dimension

The aggregated metric that you want to get the average for, for example sum({Revenue}).

sort attribute

One or more aggregated fields, either measures or dimensions or both, that you want to sort the data by, separated by commas. You can either specify ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it's more than one word. The entire list is enclosed in [] (square brackets).

start index

The start index is a positive integer, indicating *n* rows above the current row. The start index counts the available data points above the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

end index

The end index is a positive integer, indicating n rows below the current row. The end index counts the available data points below the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it's more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the moving count of `sum(Revenue)`, partitioned by `SaleDate`. The calculation includes three rows above and two row below of the current row.

```
windowCount
(
    sum(Revenue),
    [SaleDate ASC],
    3,
    2
)
```

windowMax

The `windowMax` function calculates the maximum of the aggregated measure in a custom window that is partitioned and sorted by specified attributes. Usually, you use custom window functions on a time series, where your visual shows a metric and a date field. You can use `windowMax` to help you identify the maximum of the metric over a period time.

Window functions aren't supported for MySQL versions earlier than 8 and MariaDB versions earlier than 10.2.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
windowMax
```

```
(  
    measure  
    , [ sort_order_field ASC/DESC, ... ]  
    , start_index  
    , end_index  
    , [ partition_field, ... ]  
)
```

Arguments

measure

The aggregated metric that you want to get the average for, for example `sum( {Revenue} )`.

sort attribute

One or more aggregated fields, either measures or dimensions or both, that you want to sort the data by, separated by commas. You can either specify ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it's more than one word. The entire list is enclosed in [] (square brackets).

start index

The start index is a positive integer, indicating *n* rows above the current row. The start index counts the available data points above the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

end index

The end index is a positive integer, indicating *n* rows below the current row. The end index counts the available data points below the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it is more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the trailing 12-month maximum of `sum(Revenue)`, partitioned by `SaleDate`. The calculation includes 12 rows above and 0 row below of the current row.

```
windowMax
(
    sum(Revenue),
    [SaleDate ASC],
    12,
    0
)
```

The following screenshot shows the results of this trailing 12-month example. The `sum(Revenue)` field is added to the chart to show the difference between the revenue and the trailing 12-month maximum revenue.

windowMin

The `windowMin` function calculates the minimum of the aggregated measure in a custom window that is partitioned and sorted by specified attributes. Usually, you use custom window functions on a time series, where your visual shows a metric and a date field. You can use `windowMin` to help you identify the minimum of the metric over a period time.

Window functions aren't supported for MySQL versions earlier than 8 and MariaDB versions earlier than 10.2.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
windowMin
(
    measure
    , [sort_order_field ASC/DESC, ...]
    , start_index
    , end_index
    , [ partition_field, ... ]
)
```

Arguments

measure

The aggregated metric that you want to get the average for, for example `sum( {Revenue} )`.

sort attribute

One or more aggregated fields, either measures or dimensions or both, that you want to sort the data by, separated by commas. You can either specify ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in `{ }` (curly braces), if it's more than one word. The entire list is enclosed in `[ ]` (square brackets).

start index

The start index is a positive integer, indicating n rows above the current row. The start index counts the available data points above the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

end index

The end index is a positive integer, indicating n rows below the current row. The end index counts the available data points below the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in `{ }` (curly braces), if it's more than one word. The entire list is enclosed in `[ ]` (square brackets).

Example

The following example calculates the trailing 12-month minimum of `sum(Revenue)`, partitioned by `SaleDate`. The calculation includes 12 rows above and 0 row below of the current row.

```
windowMin  
(
```



```
sum(Revenue),  
[SaleDate ASC],  
12,  
0  
)
```

The following screenshot shows the results of this trailing 12-month example. The `sum(Revenue)` field is added to the chart to show the difference between the revenue and the trailing 12-month minimum revenue.

windowSum

The `windowSum` function calculates the sum of the aggregated measure in a custom window that is partitioned and sorted by specified attributes. Usually, you use custom window functions on a time series, where your visual shows a metric and a date field.

Window functions aren't supported for MySQL versions earlier than 8 and MariaDB versions earlier than 10.2.

Syntax

The brackets are required. To see which arguments are optional, see the following descriptions.

```
windowSum  
(  
    measure  
    , [sort_order_field ASC/DESC, ...]  
    , start_index  
    , end_index  
    , [ partition_field, ... ]  
)
```

Arguments

measure

The aggregated metric that you want to get the sum for, for example `sum( {Revenue} )`.

For the engines MySQL, MariaDB, and Amazon Aurora with MySQL compatibility, the lookup index is limited to just 1. Window functions aren't supported for MySQL versions below 8 and MariaDB versions earlier than 10.2.

sort attribute

One or more aggregated fields, either measures or dimensions or both, that you want to sort the data by, separated by commas. You can either specify ascending (**ASC**) or descending (**DESC**) sort order.

Each field in the list is enclosed in {} (curly braces), if it's more than one word. The entire list is enclosed in [] (square brackets).

start index

The start index is a positive integer, indicating n rows above the current row. The start index counts the available data points above the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

end index

The end index is a positive integer, indicating n rows below the current row. The end index counts the available data points below the current row, rather than counting actual time periods. If your data is sparse (missing months or years, for example), adjust the indexes accordingly.

partition field

(Optional) One or more dimensions that you want to partition by, separated by commas.

Each field in the list is enclosed in {} (curly braces), if it's more than one word. The entire list is enclosed in [] (square brackets).

Example

The following example calculates the moving sum of `sum(Revenue)`, sorted by `SaleDate`. The calculation includes two rows above and one row ahead of the current row.

```
windowSum
(
    sum(Revenue),
    [SaleDate ASC],
    2,
    1
)
```

The following example show a trailing 12-month sum.

```
windowSum(sum(Revenue), [SaleDate ASC], 12, 0)
```

The following screenshot shows the results of this trailing 12-month sum example. The `sum(Revenue)` field is added to the chart to show the difference between the revenue and the trailing 12-month sum of revenue.

Previewing tables in a dataset

You can preview each individual data table within a dataset. When you choose a data table to preview, a read-only preview of the table appears in a new tab in the data preview section. You can have multiple table preview tabs open at once.

You can only preview tables that you have access to in a dataset. If a table doesn't appear in the top half of the data preparation space, you can't preview the table.

The **Dataset** tab contains all transformations, like new columns or filters. Table preview tabs don't show any of your transforms.

To preview a data table

1. On the Quick Suite homepage, choose **Data** at left.
2. In the **Data** tab, choose the dataset that you want, and choose **Edit dataset**.
3. Choose the data table that you want to preview, choose the down arrow to open the menu, and choose **Show table preview**.

Joining data

You can use the join interface in Amazon Quick Sight to join objects from one or more data sources. By using Amazon Quick Sight to join the data, you can merge disparate data without duplicating the data from different sources.

Types of joined datasets

A join is performed between two Quick Sight *logical tables*, where each logical table contains information about how to fetch data. When editing a dataset in Quick Sight, the join diagram at the top half of the page shows each logical table as a rectangular block.

There are two different types of joined datasets in Quick Sight: same-source and cross-source. A dataset is considered same-source when it doesn't have any joins, or when all of the following conditions are met:

- If any of the logical tables refer to a Quick Sight data source:
 - All of the logical tables in this dataset must refer to the same Quick Sight data source. This doesn't apply if two separate Quick Sight data sources refer to the same underlying database. It must be the exact same Quick Sight data source. For more information about using a single data source, see [Creating a dataset using an existing data source](#).
- If any of the logical tables refer to a Quick Sight dataset that is a parent dataset:
 - The parent dataset must use direct query.
 - The parent dataset must refer to the same Quick Sight data source.

If the above conditions aren't met, the dataset is considered a cross-source join.

Facts about joining datasets

Both same-source and cross-source dataset joins have the following limitations.

What's the maximum number of tables a joined dataset can contain?

All joined datasets can contain up to 32 tables.

How large can joined data be?

The maximum allowed size of a join is determined by the query mode and query engine that is used. The list below provides information about the different size limits for the tables to be joined. The size limit applies to all secondary tables combined. There are no join size limits for the primary table.

- **Same-source tables** – When tables originate from a single query data source, Quick Sight imposes no restrictions on the join size. This does not override join size limitations that the source query engine may have in place.
- **Cross-source datasets** – This type of join contains tables from different data sources that aren't stored in SPICE. For these types of joins, Quick Sight automatically identifies the largest table in the dataset. The combined size of all other secondary tables must be less than 1 GB.
- **Datasets stored in SPICE** – This type of join contains tables that are all ingested into SPICE. The combined size of all secondary tables in this join cannot exceed 20 GB.

For more information about SPICE dataset size calculations, see [Estimating the size of SPICE datasets](#).

Can a joined dataset use direct query?

Same-source datasets support direct query, assuming there are no other restrictions on using direct query. For example, S3 data sources don't support direct query, so a same-source S3 dataset must still use SPICE.

Cross-source datasets must use SPICE.

Can calculated fields be used in a join?

All joined datasets can use calculated fields, but calculated fields can't be used in any on-clauses.

Can geographical data be used in a join?

Same-source datasets support geographical data types, but geographical fields can't be used in any on-clauses.

Cross-source datasets don't support geographical data in any form.

For some examples of joining tables across data sources, see the [Joining across data sources on Amazon Quick Sight](#) post on the Amazon Big Data Blog.

Creating a join

Use the following procedure to join tables to use in a dataset. Before you begin, import or connect to your data. You can create a join between any of the data sources supported by Amazon Quick Sight, except Internet of Things (IoT) data. For example, you can add comma-separated value (.csv) files, tables, views, SQL queries, or JSON objects in an Amazon S3 bucket.

To add one or more joins

1. Open the dataset that you want to work with.
2. (Optional) Before you get started, decide if you want to disable the autogenerated preview based on of a sample of your data. To turn that off, choose **Auto-preview** at top right. It's turned on by default.
3. If you haven't already chosen a query mode, choose **Query mode**.

Choose **SPICE** to store your dataset in [SPICE](#), or choose **Direct query** to pull live data every time. If your dataset contains one or more manually uploaded files, your dataset is automatically stored in SPICE.

If you choose **SPICE**, the data is ingested into Quick Sight. Visuals that use the dataset run queries in SPICE, instead of on the database.

If you choose **Direct query**, the data isn't ingested into SPICE. Visuals that use the dataset run queries on the database, instead of in SPICE.

If you choose **Query mode**, make sure to set unique keys in the join, if applicable, to improve performance when loading visuals.

4. On the data preparation page, choose **Add data**.
5. In the **Add data** page that opens, choose one of the following options and complete the steps following:
 - Add data from a dataset:
 1. Choose **Dataset**.
 2. Select a dataset from the list.
 3. Choose **Select**.
 - Add data from a data source:
 1. Choose **Data source**.
 2. Select a data source from the list.
 3. Choose **Select**.
 4. Select a table from the list.
 5. Choose **Select**.
 - Create self-joins by adding a table multiple times. A counter appears after the name. An example is **Product**, **Product (2)**, and **Product (3)**. Field names in the **Fields** or **Filters** sections include the same counter so you can know which instance of the table a field came from.
 - Add a new file by choosing **Upload a file**, and then choose the file that you want to join.
6. (Optional) Choose **Use custom SQL** to open the query editor and write a query for a SQL data source.

7. (Optional) After you add data, interact with each table by choosing its menu icon. Rearrange the tables by dragging and dropping them.

An icon with red dots appears to indicate that you need to configure this join. Two red dots appear for joins that aren't yet configured. To create joins, choose the first join configuration icon.

8. (Optional) To change an existing join, reopen **Join configuration** by choosing the join icon between two tables.

The **Join Configuration** pane opens. On the join interface, specify the join type and the fields to use to join the tables.

9. At the bottom of the screen, you can see options to set a field in one table equal to a field in another table.

- In the **Join clauses** section, choose the join column for each table.

(Optional) If the tables that you selected join on multiple columns, choose **Add a new join clause**. Doing this adds another row to the join clauses, so you can specify the next set of columns to join. Repeat this process until you have identified all of the join columns for the two data objects.

10. In the **Join configuration** pane, choose the kind of join to apply. If the join fields are a unique key for one or both tables, enable the unique key setting. Unique keys only apply to direct queries, not to SPICE data.

For more information about joins, see [Join types](#).

11. Choose **Apply** to confirm your choices.

To cancel without making changes, choose **Cancel**.

12. The join icon in the workspace changes to show the new relationship.

13. (Optional) In the **Fields** section, you can use each field's menu to do one or more of the following:

- **Add a hierarchy** to a geospatial field.
- **Include** or **Exclude** the field.
- **Edit name & description** of the field.
- **Change data type**.
- **Add a calculation** (a calculated field).

- **Restrict access to only me**, so only you can see it. This can be helpful when you are adding fields to a dataset that's already in use.

14. (Optional) In the **Filters** section, you can add or edit filters. For more information, see [Filtering data in Amazon Quick Sight](#).

Join types

Amazon Quick Sight supports the following join types:

- Inner joins
- Left and right outer joins
- Full outer joins

Let's take a closer look at what these join types do with your data. For our example data, we're using the following tables named `widget` and `safety_rating`.

```
SELECT * FROM safety-rating
```

```
rating_id safety_rating
```

```
1      A+
2      A
3      A-
4      B+
5      B
```

```
SELECT * FROM WIDGET
```

```
widget_id    widget safety_rating_id
```

```
1      WidgetA  3
2      WidgetB  1
3      WidgetC  1
4      WidgetD  2
5      WidgetE
6      WidgetF  5
7      WidgetG
```


Inner joins

Use an inner join
()
when you want to see only the data where there is a match between two tables. For example, suppose that you perform an inner join on the **safety-rating** and **widget** tables.

In the following result set, widgets without safety ratings are removed, and safety ratings without associated widgets are removed. Only the rows that match perfectly are included.

```
SELECT * FROM safety-rating
INNER JOIN widget
ON safety_rating.rating_id = widget.safety_rating_id
```

rating_id	safety_rating	widget_id	widget	safety_rating_id
3	A-	1	WidgetA	3
1	A+	2	WidgetB	1
1	A+	3	WidgetC	1
2	A	4	WidgetD	2
5	B	6	WidgetF	5

Left and right outer joins

These are also known as left or right outer joins. Use a left
()
or right
()
outer join when you want to see all the data from one table, and only the matching rows from the other table.

In a graphical interface, you can see which table is on the right or the left. In a SQL statement, the first table is considered to be on the left. Therefore, choosing a left outer join as opposed to a right outer join depends only on how the tables are laid out in your query tool.

For example, suppose that you perform a left outer join
()
on **safety-rating** (the left table) and **widgets** (the right table). In this case, all **safety-rating** rows are returned, and only matching **widget** rows are returned. You can see blanks in the result set where there is no matching data.

```
SELECT * FROM safety-rating
```

```
LEFT OUTER JOIN widget
ON safety_rating.rating_id = widget.safety_rating_id
```

rating_id	safety_rating	widget_id	widget	safety_rating_id
1	A+	2	WidgetB	1
1	A+	3	WidgetC	1
2	A	4	WidgetD	2
3	A-	1	WidgetA	3
4	B+			
5	B	6	WidgetF	5

If you instead use a right outer join

(
)
call the tables in the same order so safety-rating is on the left and widgets is on the right. In this case, only matching safety-rating rows are returned, and all widget rows are returned. You can see blanks in the result set where there is no matching data.

```
SELECT * FROM safety-rating
RIGHT OUTER JOIN widget
ON safety_rating.rating_id = widget.safety_rating_id
```

rating_id	safety_rating	widget_id	widget	safety_rating_id
3	A-	1	WidgetA	3
1	A+	2	WidgetB	1
1	A+	3	WidgetC	1
2	A	4	WidgetD	2
		5	WidgetE	
5	B	6	WidgetF	5
		7	WidgetG	

Full outer joins

These are sometimes called just outer joins, but this term can refer to either a left outer, right outer, or full outer join. To define the meaning, we use the complete name: full outer join.

Use a full outer join

(
)
to see data that matches, plus data from both tables that doesn't match. This type of join includes all rows from both tables. For example, if you perform a full outer join on the safety-rating and widget tables, all rows are returned. The rows are aligned where they matched, and all extra data is included on separate rows. You can see blanks in the result set where there is no matching data.

```
SELECT * FROM safety-rating
FULL OUTER JOIN widget
ON safety_rating.rating_id = widget.safety_rating_id
```

rating_id	safety_rating	widget_id	widget	safety_rating_id
1	A+	2 WidgetB	1	
1	A+	3 WidgetC	1	
2	A	4 WidgetD	2	
3	A-	1 WidgetA	3	
4	B+			
5	B	6 WidgetF	5	
		5 WidgetE		
		7 WidgetG		

Filtering data in Amazon Quick Sight

You can use filters to refine the data in a dataset or an analysis. For example, you can create a filter on a region field that excludes data from a particular region in a dataset. You can also add a filter to an analysis, such as a filter on the range of dates that you want to include in any visuals in your analysis.

When you create a filter in a dataset, that filter applies to the entire dataset. Any analyses and subsequent dashboards created from that dataset contains the filter. If someone creates a dataset from your dataset, the filter also is in the new dataset.

When you create a filter in an analysis, that filter only applies to that analysis and any dashboards you publish from it. If someone duplicates your analysis, the filter persists in the new analysis. In analyses, you can scope filters to a single visual, some visuals, all visuals that use this dataset, or all applicable visuals.

Also, when you create filters in an analysis, you can add a filter control to your dashboard. For more information about filter controls, see [Adding filter controls to analysis sheets](#).

Each filter you create applies only to a single field. You can apply filters to both regular and calculated fields.

There are several types of filters you can add to datasets and analyses. For more information about the types of filters you can add, and some of their options, see [Filter types in Amazon Quick Suite](#).

If you create multiple filters, all top-level filters apply together using AND. If you group filters by adding them inside a top-level filter, the filters in the group apply using OR.

Amazon Quick Sight applies all of the enabled filters to the field. For example, suppose that there is one filter of `state = WA` and another filter of `sales >= 500`. Then the dataset or analysis only contains records that meet both of those criteria. If you disable one of these, only one filter applies.

Take care that multiple filters applied to the same field aren't mutually exclusive.

Use the following sections to learn how to view, add, edit, and delete filters.

Topics

- [Viewing existing filters](#)
- [Adding filters](#)
- [Cross-sheet filters and controls](#)
- [Filter types in Amazon Quick Suite](#)
- [Adding filter controls to analysis sheets](#)
- [Editing filters](#)
- [Enabling or disabling filters](#)
- [Deleting filters](#)

Viewing existing filters

When you edit a dataset or open an analysis, you can view any existing filters that were created. Use the following procedures to learn how.

Viewing filters in datasets

1. Open the [Quick Suite console](#).
2. From the Quick Suite homepage, choose **Data** at left.
3. In the **Datasets** tab, choose the dataset that you want, and then choose **Edit dataset**.
4. On the data preparation page that opens, choose **Filters** at lower left to expand the **Filters** section.

Any filters that are applied to the dataset appear here. If a single field has multiple filters, they are grouped together. They display in order of create date, with the oldest filter on top.

Viewing filters in analyses

Use the following procedure to view filters in analyses.

To view a filter in an analysis

1. From the Quick Suite homepage, choose **Analyses**.
2. On the **Analyses** page, choose the analysis that you want to work with.
3. In the analysis, choose the **Filter** icon to open the **Filters** pane.

Any filters applied to the analysis appear here.

The way that a filter is scoped is listed at the bottom of each filter. For more information about scoping filters, see [Adding filters](#).

Adding filters

You can add filters to a dataset or an analysis. Use the following procedures to learn how.

Adding filters to datasets

Use the following procedure to add filters to datasets.

To add a filter to a dataset

1. Open the [Quick Suite console](#).
2. From the Quick Suite homepage, choose **Data** at left.
3. In the **Datasets** tab, choose the dataset that you want, and then choose **Edit dataset**.
4. On the data preparation page that opens, choose **Add filter** at lower left, and then choose a field that you want to filter.

The filter is added to the **Filters** pane.

5. Choose the new filter in the pane to configure the filter. Or you can choose the three dots to the right of the new filter and choose **Edit**.

Depending on the data type of the field, your options for configuring the filter vary. For more information about the types of filters that you can create and their configurations, see [Filter types in Amazon Quick Suite](#).

6. When finished, choose **Apply**.

Note

The data preview shows you the results of your combined filters only as they apply to the first 1,000 rows. If all of the first 1,000 rows are filtered out, then no rows show in the preview. This effect occurs even when rows after the first 1,000 aren't filtered out.

Adding filters in analyses

Use the following procedure to add filters to analyses.

To add a filter to an analysis

1. Open the [Quick Suite console](#).
2. From the Quick Suite homepage, choose **Analyses**.
3. On the **Analyses** page, choose the analysis that you want to work with.
4. In the analysis, choose the **Filter** icon to open the **Filters** pane, and then choose **ADD**.
5. Choose the new filter in the pane to configure it. Or you can choose the three dots to the right of the new filter and choose **Edit**.
6. In the **Edit filter** pane that opens, for **Applied to**, choose one of the following options.
 - **Single visual** – The filter applies to the selected item only.
 - **Single sheet** – The filter applies to a single sheet.
 - **Cross sheet** – The filter applies to multiple sheets in the dataset.

Depending on the data type of the field, your remaining options for configuring the filter vary. For more information about the types of filters you can create and their configurations, see [Filter types in Amazon Quick Suite](#).

Cross-sheet filters and controls

Cross-sheet filters and controls are filters that are scoped to either your entire analysis or dashboard or multiple sheets within your analysis and dashboard.

Filters

Creating a Cross-Sheet Filter

1. Once you have [Added a filter](#), you update the scope of the filter to cross-sheet. By default, this applies to all of the the sheets in your analysis.
2. If the **Apply cross-datasets** box is checked, then the filter will be applied to all visuals from up to 100 different datasets that are applicable to all sheets in the filter scope.
3. If you want to customize the sheets that it is applied to, then choose the Cross-sheet icon. You can then view the sheets the filter is currently applied to or toggle on the custom select sheets.
4. When you enable **Custom select sheets**, you can select which sheets to apply the filter to.
5. Follow the steps at [Editing filters in analyses](#). Your changes will be applied to all of the filters for all of the sheets you have selected. This includes newly added sheets if the filter is scoped to your entire analysis.

Removing a Cross-Sheet Filter

Deleting

If you have no controls created from these filters, see [Deleting filters in analyses](#).

If you have controls created then:

1. Follow the instructions at [Deleting filters in analyses](#).
2. If you choose **Delete Filter and Controls**, the controls will be deleted from all pages. This may impact the layout of your analysis. Alternatively, you can remove these controls individually.

Downscoping

If you want to remove a cross-sheet filter, you can also do this by changing the filter scope:

1. Follow the instructions at [Editing filters in analyses](#) to get to the filter.
2. One of the edits you can make is changing the scope. You can switch to **Single sheet** or **Single visual**. You can also remove a sheet from the Cross-sheet selection.

Or the custom sheet selection:

3. If there are controls, you will see a modal to warn you that you will be bulk-removing controls from any of the sheets where the filter no longer applies and this can impact your layout. You can also remove the controls individually. For more information, see [Removing a Cross-Sheet Control](#).
4. If you add controls to the **Top of all sheets in filter scope** then new sheets will by default be added with this new control if the filter is scoped to your entire analysis.

Controls

Creating a Cross-Sheet Control

New filter control

1. Create a cross-sheet filter. For more information, see [Filters](#).
2. From the three-dot menu, you can see an option that says **Add control**. Hovering over this, you will see three options:
 - **Top of all sheets in filter scope**
 - **Top of this sheet**
 - **Inside this sheet**

If you want to add to multiple-sheets within the sheets themselves, you can do that sheet-by-sheet. Or you can add to the top and then use the option on each control to **Move to sheet**. For more information, see [Editing a Cross-Sheet Control](#).

Increasing Scope of Existing Control

1. Navigate to the existing filter in the analysis
2. Change the scope of what sheets this filter is **Applied to** to **Cross-sheet**.
3. If there is already a control created from the filter, you will see a modal, which if you check the box will bulk-add controls to the top of all the sheets in the filter scope. This will not impact the position of the already created control if it is on the sheet.

Editing a Cross-Sheet Control

1. Go to the cross-sheet control and select the three-dot menu if the control is pinned to the top or the edit pencil icon if the control is on the sheet. You will be presented with the following options:
 - **Go to filter** (which directs you to the cross-sheet filter for you to edit or review)
 - **Move to sheet** (which moves the control into the analysis pane)
 - **Reset**
 - **Refresh**
 - **Edit**
 - **Remove**
2. Choose **Edit**. This brings up the **Format Control** pane on the right side of your analysis.
3. You can then edit your control. The top section labeled **Cross-sheet settings** will apply to all controls, whereas any settings outside of this section are not applicable to all controls and only to the specific control you're editing. For instance, **Relevant value** is not a cross-sheet control setting.
4. You can also see the sheets that this control is on as well as the location (Top or Sheet) that the control is on for each sheet. You can do this by choosing **Sheets(8)**.

Removing a Cross-Sheet Control

You can remove controls in two places. First, from the control:

1. Go to the cross-sheet control and select the three-dot menu if the control is pinned to the top or the edit pencil icon if the control is on the sheet. You will be presented with the following options:
 - **Go to filter** (which directs you to the cross-sheet filter for you to edit or review)
 - **Move to sheet** (which moves the control into the analysis pane)
 - **Reset**
 - **Refresh**
 - **Edit**
 - **Remove**
2. Choose **Remove**

Second, you can remove controls from the filter:

1. Choose the three-dot menu on the cross-sheet filter that the cross-sheet controls are created from. You will see that instead of an option to **Add control** there is now an option to **Manage control**.
2. Hover over **Manage control**. You will be presented with the following options:
 - **Move inside this sheet**
 - **Top of this sheet**

These options are applicable to just the control on the sheet, depending on where the current control is. If you don't have controls on all of the sheets within the filter scope, you will get the option to **Add to top of all sheets in filter scope**. This will not move sheet controls to the top of the sheet if you have already added them to the sheet in the analysis. You will also get the option to **Remove from this sheet** or **Remove from all sheets**.

Filter types in Amazon Quick Suite

You can create several different types of filters in Quick Suite. The type of filter you create mostly depends on the data type of the field that you want to filter.

In datasets, you can create the following types of filters:

- Text filters
- Numeric filters
- Date filters

In analyses, you can create the same types of filters as you can in datasets. You can also create:

- Group filters with and/or operators
- Cascading filters
- Nested filters

Use the following sections to learn more about each type of filter you can create and some of their options.

Topics

- [Adding text filters](#)
- [Adding nested filters](#)
- [Adding numeric filters](#)
- [Adding date filters](#)
- [Adding filter conditions \(group filters\) with AND and OR operators](#)
- [Creating cascading filters](#)

Adding text filters

When you add a filter using a text field, you can create the following types of text filters:

- **Filter list** (Analyses only) – This option creates a filter that you can use to select one or more field values to include or exclude from all the available values in the field. For more information about creating this type of text filter, see [Filtering text field values by a list \(analyses only\)](#).
- **Custom filter list** – With this option, you can enter one or more field values to filter on, and whether you want to include or exclude records that contain those values. The values that you enter must match the actual field values exactly for the filter to be applied to a given record. For more information about creating this type of text filter, see [Filtering text field values by a custom list](#).
- **Custom filter** – With this option, you enter a single value that the field value must match in some way. You can specify that the field value must equal, not equal, starts with, ends with, contains, or does not contain the value you specify. If you choose an equal comparison, the specified value and actual field value must match exactly in order for the filter to be applied to a given record. For more information about creating this type of text filter, see [Filtering a single text field value](#).
- **Top and bottom filter** (Analyses only) – You can use this option to show the top or bottom *n* value of one field ranked by the values in another field. For example, you might show the top five salespeople based on revenue. You can also use a parameter to allow dashboard users to dynamically choose how many top or bottom ranking values to show. For more information about creating top and bottom filters, see [Filtering a text field by a top or bottom value \(analyses only\)](#).

Filtering text field values by a list (analyses only)

In analyses, you can filter a text field by selecting values to include or exclude from a list of all value in the field.

To filter a text field by including and excluding values

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. For **Filter type**, choose **Filter list**.
4. For **Filter condition**, choose **Include** or **Exclude**.
5. Choose the field values that you want to filter on. To do this, select the check box in front of each value.

If there are too many values to choose from, enter a search term into the box above the checklist and choose **Search**. Search terms are case-insensitive and wildcards aren't supported. Any field value that contains the search term is returned. For example, searching on L returns al, AL, la, and LA.

The values display alphabetically in the control, unless there are more than 1,000 distinct values. Then the control displays a search box instead. Each time that you search for the value that you want to use, it starts a new query. If the results contain more than 1,000 values, you can scroll through the values with pagination.

6. When finished, choose **Apply**.

Filtering text field values by a custom list

You can specify one or more field values to filter on, and whether you want to include or exclude records that contain those values. The specified value and actual field value must match exactly for the filter to be applied to a given record.

To filter text field values by a custom list

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. For **Filter type**, choose **Custom filter list**.

4. For **Filter condition**, choose **Include** or **Exclude**.
5. For **List**, enter a value in the text box. The value must match an existing field value exactly.
6. (Optional) To add additional values, enter them in the text box, one per line.
7. For **Null options** choose **Exclude nulls**, **Include nulls**, or **Nulls only**.
8. When finished, choose **Apply**.

Filtering a single text field value

With the **Custom filter** filter type, you specify a single value that the field value must equal or not equal, or must match partially. If you choose an equal comparison, the specified value and actual field value must match exactly for the filter to be applied to a given record.

To filter a text field by a single value

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. For **Filter type**, choose **Custom filter**.
4. For **Filter condition**, choose one of the following:
 - **Equals** – When you choose this option, the values included or excluded in the field must match the value that you enter exactly.
 - **Does not equal** – When you choose this option, the values included or excluded in the field must match the value that you enter exactly.
 - **Starts with** – When you choose this option, the values included or excluded in the field must start with the value that you enter.
 - **Ends with** – When you choose this option, the values included or excluded in the field must start with the value that you enter.
 - **Contains** – When you choose this option, the values included or excluded in the field must contain the whole value that you enter.
 - **Does not contain** – When you choose this option, the values included or excluded in the field must not contain any part of the value that you enter.

 Note

Comparison types are case-sensitive.

5. Do one of the following:

- For **Value**, enter a literal value.
- Select **Use parameters** to use an existing parameter, and then choose a parameter from the list.

For parameters to appear in this list, create your parameters first. Usually, you create a parameter, add a control for it, and then add a filter for it. For more information, see [Parameters in Amazon Quick Suite](#).

The values display alphabetically in the control, unless there are more than 1,000 distinct values. Then the control displays a search box instead. Each time that you search for the value that you want to use, it starts a new query. If the results contain more than 1,000 values, you can scroll through the values with pagination.

6. For **Null options** choose **Exclude nulls**, **Include nulls**, or **Nulls only**.
7. When finished, choose **Apply**.

Filtering a text field by a top or bottom value (analyses only)

You can use a **Top and bottom filter** to show the top or bottom n value of one field ranked by the values in another field. For example, you might show the top five salespeople based on revenue. You can also use a parameter to allow dashboard users to dynamically choose how many top or bottom ranking values to show.

To create a top and bottom text filter

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. For **Filter type**, choose **Top and bottom filter**.
4. Choose **Top** or **Bottom**.
5. For **Show top** integer (or **Show bottom** integer), do one of the following:

- Enter the number of top or bottom items to show.
- To use a parameter for the number of top or bottom items to show, select **Use parameters**. Then choose an existing integer parameter.

For example, let's say that you want to show the top three salespersons by default. However, you want the dashboard viewer to be able to choose whether to show 1–10 top salespersons. In this case, take the following actions:

- Create an integer parameter with a default value.
 - To link the number of displayed items to a parameter control, create a control for the integer parameter. Then you make the control a slider with a step size of 1, a minimum value of 1, and a maximum value of 10.
 - To make the control work, link it to a filter by creating a top and bottom filter on Salesperson by Weighted Revenue, enable **Use parameters**, and choose your integer parameter.
6. For **By**, choose a field to base the ranking on. If you want to show the top five salespeople per revenue, choose the revenue field. You can also set the aggregate that you want to perform on the field.
 7. (Optional) Choose **Tie breaker** and then choose another field to add one or more aggregations as tie breakers. This is useful, in the case of this example, when there are more than five results returned for the top five salespeople per revenue. This situation can happen if multiple salespeople have the same revenue amount.

To remove a tie breaker, use the delete icon.

8. When finished, choose **Apply**.

Adding nested filters

Nested filters are advanced filters that can be added to a Quick Suite analysis. A Nested filter filters a field using a subset of data defined by another field in that same dataset. This allows authors to show additional contextual data without the need to filter data out if the data point doesn't meet an initial condition.

Nested filters function similarly to a correlated subquery in SQL or a market basket analysis. For example, say you want to perform a market basket analysis on your sales data. You can use nested filters to find the sales quantity by product for customers who have or have not purchased

a specific product. You can also use nested filters to identify groups of customers that did not purchase a selected product or who only purchased a specific list of products.

Nested filters can only be added at the analysis level. You can't add a nested filter to a dataset.

Use the procedure below to add a nested filter to a Quick Suite analysis.

1. Open the [Quick Suite console](#).
2. Choose **Analyses**, and then choose the analysis that you want to add a nested filter to.
3. Create a new filter on the text field that you want to filter on. For more information about creating a filter, see [Adding filters in analyses](#).
4. After you create the new filter, locate the new filter in the **Filters** pane. Choose the ellipsis (three dots) next to the new filter, and then choose **Edit filter**. Alternatively, choose the filter entity in the **Filters** pane to open the **Edit filter** pane.
5. The **Edit filter** pane opens. Open the **Filter type** dropdown menu, navigate to the **Advanced filter** section, and then choose **Nested filter**.
6. For **Qualifying condition**, choose **Include** or **Exclude**. The *qualifying condition* allows you to run a not in the set query on the data in your analysis. In our sales example above, the qualifying condition determines if the filter returns a list of customers who did buy the specific product or a list of customers who did not buy the product.
7. For **Nested field**, choose the text field that you want to filter data with. The nested field cannot be the same as the primary field selected in step 3. Category fields are the only supported field type for the inner filter.
8. For **Nested filter type**, choose the filter type that you want. The filter type that you choose determines the final configuration steps for the nested filter. Available filter types and information about their configuration can be found in the list below.
 - [Filter list](#)
 - [Custom filter list](#)
 - [Custom filter](#)

Adding numeric filters

Fields with decimal or int data types are considered numeric fields. You create filters on numeric fields by specifying a comparison type, for example **Greater than** or **Between**, and a comparison

value or values as appropriate to the comparison type. Comparison values must be positive integers and can't contain commas.

You can use the following comparison types in numeric filters:

- Equals
- Does not equal
- Greater than
- Greater than or equal to
- Less than
- Less than or equal to
- Between

 Note

To use a top and bottom filter for numeric data (analyses only), first change the field from a measure to a dimension. Doing this converts the data to text. Then you can use a text filter. For more information, see [Adding text filters](#).

In analyses, for datasets based on database queries, you can also optionally apply an aggregate function to the comparison value or values, for example **Sum** or **Average**.

You can use the following aggregate functions in numeric filters:

- Average
- Count
- Count distinct
- Max
- Median
- Min
- Percentile
- Standard deviation
- Standard deviation - population
- Sum

- Variance
- Variance - population

Creating numeric filters

Use the following procedure to create a numeric field filter.

To create a numeric field filter

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. (Optional) For **Aggregation**, choose an aggregation. No aggregation is applied by default. This option is available only when creating numeric filters in an analysis.
4. For **Filter condition**, choose a comparison type.
5. Do one of the following:

- If you chose a comparison type other than **Between**, enter a comparison value.

If you chose a comparison type of **Between**, enter the beginning of the value range in **Minimum value** and the end of the value range in **Maximum value**.

- (Analyses only) To use an existing parameter, enable **Use parameters**, then choose your parameter from the list.

For parameters to appear in this list, create your parameters first. Usually, you create a parameter, add a control for it, and then add a filter for it. For more information, see [Parameters in Amazon Quick Suite](#). The values display alphabetically in the control, unless there are more than 1,000 distinct values. Then the control displays a search box instead. Each time you search for the value that you want to use, it initiates a new query. If the results contain more than 1,000 values, you can scroll through the values with pagination.

6. (Analyses only) For **Null options** choose **Exclude nulls**, **Include nulls**, or **Nulls only**.
7. When finished, choose **Apply**.

Adding date filters

You create filters on date fields by selecting the filter conditions and date values that you want to use. There are three filter types for dates:

- **Range** – A series of dates based on a time range and comparison type. You can filter records based on whether the date field value is before or after a specified date, or within a date range. You enter date values in the format MM/DD/YYYY. You can use the following comparison types:
 - **Between** – Between a start date and an end date
 - **After** – After a specified date
 - **Before** – Before a specified date
 - **Equals** – On a specified date

For each comparison type, you can alternatively choose a rolling date relative to a period or dataset value.

- **Relative** (analyses only) – A series of date and time elements based on the current date. You can filter records based on the current date and your selected unit of measure (UOM). Date filter units include years, quarters, months, weeks, days, hours, and minutes. You can exclude current period, add support for Next N filters similar to Last N with an added capability to allow for Anchor date. You can use the following comparison types:
 - **Previous** – The previous UOM—for example, the previous year.
 - **This** – This UOM, which includes all dates and times that fall within the select UOM, even if they occur in the future.
 - **To date or up to now** – UOM to date, or UOM up to now. The displayed phrase adapts to the UOM that you choose. However, in all cases this option filters out data that is not between the beginning of the current UOM and the current moment.
 - **Last n** – The last specified number of the given UOM, which includes all of this UOM and all of the last $n - 1$ UOM. For example, let's say today is May 10, 2017. You choose to use *years* as your UOM, and set Last n years to 3. The filtered data includes data for all of 2017, plus all of 2016, and all of 2015. If you have any data for the future dates of the current year (2017 in this example), these records are included in your dataset.
- **Top and bottom** (analyses only) – A number of date entries ranked by another field. You can show the top or bottom n for the type of date or time UOM you choose, based on values in another field. For example, you can choose to show the top 5 sales days based on revenue.

Comparisons are applied inclusive to the date specified. For example, if you apply the filter **Before 1/1/16**, the records returned include all rows with date values through 1/1/16 23:59:59. If you don't want to include the date specified, you can clear the option to **Include this date**. If you want

to omit a time range, you can use the **Exclude the last N periods** option to specify the number and type of time periods (minutes, days, and so on) to filter out.

You can also choose to include or exclude nulls, or exclusively show rows that contain nulls in this field. If you pass in a null date parameter (one without a default value), it doesn't filter the data until you provide a value.

Note

If a column or attribute has no time zone information, then the client query engine sets the default interpretation of that date-time data. For example, suppose that a column contains a timestamp, rather than a timestamptz, and you are in a different time zone than the data's origin. In this case, the engine can render the timestamp differently than you expect. Amazon Quick Suite and [SPICE](#) both use Universal Coordinated Time (UTC) times.

Use the following sections to learn how to create date filters in datasets and analyses.

Creating date filters in datasets

Use the following procedure to create a range filter for a date field in a dataset.

To create a range filter for a date field in a dataset

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. For **Condition**, choose a comparison type: **Between**, **After**, or **Before**.

To use **Between** as a comparison, choose **Start date** and **End date** and choose dates from the date picker controls that appear.

You can choose if you want to include either or both the start and end dates in the range by selecting **Include start date** or **Include end date**.

To use **Before** or **After** comparisons, enter a date or choose the date field to bring up the date picker control and choose a date instead. You can include this date (the one you chose), to exclude the last N time periods, and specify how to handle nulls.

4. For **Time granularity**, choose **Day**, **Hour**, **Minute**, or **Second**.

5. When finished, choose **Apply**.

Creating date filters in analyses

You can create date filters in analyses as described following.

Creating range date filters in analyses

Use the following procedure to create a range filter for a date field in an analysis.

To create a range filter for a date field in an analysis

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. For **Filter type**, choose **Date & time range**.
4. For **Condition**, choose a comparison type: **Between**, **After**, **Before**, or **Equals**.

To use **Between** as a comparison, choose **Start date** and **End date** and choose dates from the date picker controls that appear.

You can choose to include either or both the start and end dates in the range by selecting **Include start date** or **Include end date**.

To use a **Before**, **After**, or **Equals** comparison, enter a date or choose the date field to bring up the date picker control and choose a date instead. You can include this date (the one you chose), to exclude the last N time periods, and specify how to handle nulls.

To **Set a rolling date** for your comparison, choose **Set a rolling date**.

In the **Set a rolling date** pane that opens, choose **Relative date** and then select if you want to set the date to **Today**, **Yesterday**, or you can specify the **Filter condition** (start of or end of), **Range** (this, previous, or next), and **Period** (year, quarter, month, week, or day).

5. For **Time granularity**, choose **Day**, **Hour**, **Minute**, or **Second**.
6. (Optional) If you are filtering by using an existing parameter, instead of specific dates, choose **Use parameters**, then choose your parameter or parameters from the list. To use **Before**, **After**, or **Equals** comparisons, choose one date parameter. You can include this date in the range.

To use **Between**, enter both the start date and end date parameters separately. You can include the start date, the end date, or both in the range.

To use parameters in a filter, create them first. Usually, you create a parameter, add a control for it, and then add a filter for it. For more information, see [Parameters in Amazon Quick Suite](#).

7. For **Null options** choose **Exclude nulls**, **Include nulls**, or **Nulls only**.
8. When finished, choose **Apply**.

Creating relative date filters in analyses

Use the following procedure to create a relative filter for a date field in an analysis.

To create a relative filter for a date field in an analysis

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. For **Filter type**, choose **Relative dates**.
4. For **Time granularity**, choose a granularity of time that you want to filter by (days, hours, minutes).
5. For **Period**, choose a unit of time (years, quarters, quarters, months, weeks, days).
6. For **Range**, choose how you want the filter to relate to the time frame. For example, if you choose to report on months, your options are previous month, this month, month to date, last N months, and next N months.

If you choose Last N or Next N years, quarters, months, weeks, or days, enter a number for **Number of**. For example, last 3 years, next 5 quarters, last 5 days.

7. For **Null options** choose **Exclude nulls**, **Include nulls**, or **Nulls only**.
8. For **Set dates relative to**, choose one of the following options:
 - **Current date time** – If you choose this option, you can set it to **Exclude last**, and then specify the number and type of time periods.
 - **Date and time from a parameter** – If you choose this option, you can select an existing datetime parameter.

9. (Optional) If you are filtering by using an existing parameter, instead of specific dates, enable **Use parameters**, then choose your parameter or parameters from the list.

To use parameters in a filter, create them first. Usually, you create a parameter, add a control for it, and then add a filter for it. For more information, see [Parameters in Amazon Quick Suite](#).

10. When finished, choose **Apply**.

Creating top and bottom date filters in analyses

Use the following procedure to create a top and bottom filter for a date field in an analysis.

To create a top and bottom filter for a date field in an analysis

1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. For **Filter type**, choose **Top and bottom**.
4. Select **Top** or **Bottom**.
5. For **Show**, enter the number of top or bottom items you want to show and choose a unit of time (years, quarters, months, weeks days, hours, minutes).
6. For **By**, choose a field to base the ranking on.
7. (Optional) Add another field as a tie breaker, if the field for **By** has duplicates. Choose **Tie breaker**, and then choose another field. To remove a tie breaker, use the delete icon.
8. (Optional) If you are filtering by using an existing parameter, instead of specific dates, select **Use parameters**, then choose your parameter or parameters from the list.

To use a parameter for **Top and bottom**, choose an integer parameter for the number of top or bottom items to show.

To use parameters in a filter, create them first. Usually, you create a parameter, add a control for it, and then add a filter for it. For more information, see [Parameters in Amazon Quick Suite](#).

9. When finished, choose **Apply**.

Adding filter conditions (group filters) with AND and OR operators

In analyses, when you add multiple filters to a visual, Quick Suite uses the AND operator to combine them. You can also add filter conditions to a single filter with the OR operator. This is called a compound filter, or filter group.

To add multiple filters using the OR operator, create a filter group. Filter grouping is available for all types of filters in analyses.

When you filter on multiple measures (green fields marked with #), you can apply the filter conditions to an aggregate of that field. Filters in a group can contain either aggregated or nonaggregated fields, but not both.

To create a filter group

1. Create a new filter in an analysis. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. In the expanded filter, choose **Add filter condition** at bottom, and then choose a field to filter on.
4. Choose the conditions to filter on.

The data type of the field that you selected determines the options available here. For example, if you chose a numeric field, you can specify the aggregation, filter condition, and values. If you chose a text field, you can choose the filter type, filter condition, and values. And if you chose a date field, you can specify the filter type, condition, and time granularity. For more information about these options, see [Filter types in Amazon Quick Suite](#).

5. (Optional) You can add additional filter conditions to the filter group by choosing **Add filter condition** again at bottom.
6. (Optional) To remove a filter from the filter group, choose the trash-can icon near the field name.
7. When finished, choose **Apply**.

The filters appear as a group in the **Filters** pane.

Creating cascading filters

The idea behind cascading any action, such as a filter, is that choices in the higher levels of a hierarchy affect the lower levels of a hierarchy. The term *cascading* comes from the way that a cascade waterfall flows from one tier to the next.

To set up cascading filters, you need a trigger point where the filter is activated, and target points where the filter is applied. In Quick Suite, the trigger and target points are included in visuals.

To create a cascading filter, you set up an action, not a filter. This approach is because you need to define how the cascading filter is activated, which fields are involved, and which visuals are filtered when someone activates it. For more information, including step-by-step instructions, see [Using custom actions for filtering and navigating](#).

There are two other ways to activate a filter across multiple visuals:

- **For a filter that is activated from a widget on a dashboard** – The widget is called a *sheet control*, which is a custom menu that you can add to the top of your analysis or dashboard. The most common sheet control is a drop-down list, which displays a list of options to choose from when you open it. To add one of these to your analysis, create a parameter, add a control to the parameter, and then add a filter that uses the parameter. For more information, see [Setting up parameters in Amazon Quick Suite](#), [Using a control with a parameter in Amazon Quick Suite](#), and [Adding filter controls to analysis sheets](#).
- **For a filter that always applies to multiple visuals** – This is a regular filter, except that you set its scope to apply to multiple (or all) visuals. This type of filter doesn't really cascade, because there is no trigger point. It always filters all the visuals that it's configured to filter. To add this type of filter to your analysis, create or edit a filter and then choose its scope: **Single visual**, **Single sheet**, or **Cross sheets**. Note the option to **Apply cross-datasets**. If this box is checked, then the filter will be applied to all visuals from different datasets that are applicable on all sheets in the filter scope. For more information, see [Filters](#).

Adding filter controls to analysis sheets

When you're designing an analysis, you can add a filter to the analysis sheet near the visuals that you want to filter. It appears in the sheet as a control that dashboard viewers can use when you publish the analysis as a dashboard. The control uses the analysis theme settings so it looks like it's part of the sheet.

Filter controls share some settings with their filters. They apply to one, some, or all of the objects on the same sheet.

Use the following sections add and customize filter controls to an analysis. To learn how to add cross-sheet controls, see [Controls](#).

Topics

- [Adding filter controls](#)
- [Pinning filter controls to the top of a sheet](#)
- [Customizing filter controls](#)
- [Cascading filter controls](#)

Adding filter controls

Use the following procedure to add a filter control.

To add a filter control

1. Open the [Quick Suite console](#).
2. From the Quick Suite homepage, choose **Analyses**, and then choose the analysis that you want to work with.
3. In the analysis, choose **Filter**.
4. If you don't already have some filters available, create one. For more information about creating filters, see [Adding filters](#).
5. In the **Filters** pane, choose the three dots to the right of the filter that you want to add a control for, and choose **Add to sheet**.

The filter control is added to the sheet, usually at the bottom. You can resize it or drag it to different positions on the sheet. You can also customize how it appears and how dashboard viewers can interact with it. For more information about customizing filter controls, see the following sections.

Pinning filter controls to the top of a sheet

Use the following procedure to pin filter controls to the top of a sheet.

To pin a control to the top of a sheet

1. On the filter control that you want to move, choose the three dots next to the pencil icon and choose **Pin to top**.

The filter is pinned to the top of the sheet and is collapsed. You can click it to expand it.

2. (Optional) To unpin the control, expand it and hover over it at the top of the sheet until three dots appear. Choose the three dots, and then choose **Move to sheet**.

Customizing filter controls

Depending on the data type of the field and the type of filter, filter controls have different settings available. You can customize how they appear in the sheet and how dashboard viewers can interact with them.

To customize a filter control

1. Choose the filter control in the sheet.
2. On the filter control, choose the pencil icon.

If the filter control is pinned to the top of the sheet, expand it and hover your cursor over it until the three dots appear. Choose the three dots, and then choose **Edit**.

3. In the **Format control** pane that opens, do the following:
 - a. For **Display name**, enter a name for the filter control.
 - b. (Optional) To hide the display name from the filter control, clear the check box for **Show title**.
 - c. For **Title font size**, choose the title font size that you want to use. The options range from extra small to extra large. The default setting is medium.

The remaining steps depend on the type of field the control is referencing. For options by filter type, see the following sections.

Date filters

If your filter control is from a date filter, use the following procedure to customize the remaining options.

To customize further options for a date filter

1. In the **Format control** pane, for **Style**, choose one of the following options:
 - **Date picker – range** – Displays a set of two fields to define a time range. You can enter a date or time, or you can choose a date from the calendar control. You can also customize how you want the dates to appear in the control by entering a date token for **Date format**. For more information, see [Customizing date formats in Quick Suite](#).
 - **Date picker – relative** – Displays settings like the time period, its relation to the current date and time, and the option to exclude time periods. You can also customize how you want the dates to appear in the control by entering a date token for **Date format**. For more information, see [Customizing date formats in Quick Suite](#).
 - **Text field** – Displays a box where you can enter the top or bottom *N* date.

Helper text is included in the text field control by default, but you can choose to remove it by clearing the **Show helper text in control** option.

By default, Quick Suite visuals are reloaded whenever a change is made to a control. For Calendar and Relative date picker controls, authors can add an **Apply** button to a control that delays visual reload until the user chooses **Apply**. This allows users to make multiple changes at a time without additional queries. This setting can be configured with the **Show an apply button** checkbox in the **Control options** section of the **Format control** pane.

2. When finished, choose **Apply**.

Text filters

If your filter control is from a text filter, for example dimensions, categories, or labels, use the following procedure to customize the remaining options.

To customize further options for a text filter

1. In the **Format control** pane, for **Style**, choose one of the following options:
 - **Dropdown** – Displays a dropdown list with buttons that you can use to select a single value.

When you select this option, you can choose the following options for **Values**:

 - **Filter** – Displays all the values that are available in the filter.
 - **Specific values** – Enables you to enter the values to display, one entry per line.

You can also choose to **Hide Select all option from the control values**. This removes the option to select or clear the selection of all values in the filter control.

- **Dropdown - multiselect** – Displays a dropdown list with boxes that you can use to select multiple values.

When you select this option, you can choose the following options for **Values**:

- **Filter** – Displays all the values that are available in the filter.
- **Specific values** – Enables you to enter the values to display, one entry per line.

By default, Quick Suite visuals are reloaded whenever a change is made to a control. For Multiselect dropdown controls, authors can add an **Apply** button to a control that delays visual reload until the user chooses **Apply**. This allows users to make multiple changes at a time without additional queries. This setting can be configured with the **Show an apply button** checkbox in the **Control options** section of the **Format control** pane.

- **List** – Displays a list with buttons that you can use to select a single value.

When you select this option, you can choose the following options for **Values**:

- **Filter** – Displays all the values that are available in the filter.
- **Specific values** – Enables you to enter the values to display, one entry per line.

You can also choose the following:

- **Hide search bar when control is on sheet** – Hides the search bar in the filter control, so users can't search for specific values.
- **Hide Select all option from the control values** – Removes the option to select or clear the selection of all values in the filter control.
- **List - multiselect** – Displays a list with boxes that you can use to select multiple values.

When you select this option, you can choose the following options for **Values**:

- **Filter** – Displays all the values that are available in the filter.
- **Specific values** – Enables you to enter the values to display, one entry per line.

You can also choose the following:

- **Hide search bar when control is on sheet** – Hides the search bar in the filter control, so users can't search for specific values.

- **Hide Select all option from the control values** – Removes the option to select or clear the selection of all values in the filter control.
- **Text field** – Displays a text box where you can enter a single entry. Text fields support up to 79950 characters.

When you select this option, you can choose the following:

- **Show helper text in control** – Removes the helper text in text fields.
- **Text field - multiline** – Displays a text box where you can enter multiple entries. Multiline text fields support up to 79950 characters across all entries.

When you select this option, you can choose the following:

- For **Separate values by**, choose how you want to separate values you enter into the filter control. You can choose to separate values by a line break, comma, pipe (|), or semicolon.
- **Show helper text in control** – Removes the helper text in text fields.

2. When finished, choose **Apply**.

Numeric filters

If your filter control is from a numeric filter, use the following procedure to customize the remaining options.

To customize further options for a numeric filter

1. In the **Format control** pane, for **Style**, choose one of the following options:

- **Dropdown** – Displays a list where you can select a single value.

When you select this option, you can choose the following options for **Values**:

- **Filter** – Displays all the values that are available in the filter.
- **Specific values** – Enables you to enter the values to display, one entry per line.

You can also choose to **Hide Select all option from the control values**. This removes the option to select or clear the selection of all values in the filter control.

- **Filter** – Displays all the values that are available in the filter.
- **Specific values** – Enables you to enter the values to display, one entry per line.
- **Hide Select all option from the control values** – Removes the option to select or clear the selection of all values in the filter control.

- **List** – Displays a list with buttons that enable selecting a single value.

When you select this option, you can choose the following options for **Values**:

- **Filter** – Displays all the values that are available in the filter.
- **Specific values** – Enables you to enter the values to display, one entry per line.

You can also choose the following:

- **Hide search bar when control is on sheet** – Hides the search bar in the filter control, so users can't search for specific values.
- **Hide Select all option from the control values** – Removes the option to select or clear the selection of all values in the filter control.
- **Slider** – Displays a horizontal bar with a toggle that you can slide to change the value. If you have a ranged filter for values between a minimum and a maximum, the slider provides a toggle for each number. For sliders, you can specify the following options:
 - **Minimum value** – Displays the smaller value at the left of the slider.
 - **Maximum value** – Displays the larger value at the right of the slider.
 - **Step size** – Enables you to set the number of increments that the bar is divided into.
- **Text box** – Displays a box where you can enter the value. When you select this option, you can choose the following:
 - **Show helper text in control** – Removes the helper text in text fields.

2. When finished, choose **Apply**.

Cascading filter controls

You can limit the values displayed in the control, so they only show values that are valid for what is selected in other controls. This is called a cascading control.

When creating cascading controls, the following limitations apply:

1. Cascading controls must be tied to dataset columns from the same dataset.
2. The child control must be a dropdown or list control.
3. For parameter controls, the child control must be linked to a dataset column.
4. For filter controls, the child control must be linked to a filter (instead of showing only specific values).
5. The parent control must be one of the following:

- a. A string, integer, or numeric parameter control.
- b. A string filter control (excluding top-bottom filters).
- c. A non-aggregated numeric filter control.
- d. A date filter control (excluding top-bottom filters).

To create a cascading control

1. Choose **Show relevant values only**. Note that this option might not be available for all filter control types.
2. In the **Show relevant values only** pane that opens, choose one or more controls from the available list.
3. Choose a field to match the value to.
4. Choose **Update**.

Editing filters

You can edit filters at any time in a dataset or analysis.

You can't change the field a filter applies to. To apply a filter to a different field, create a new filter instead.

Use the following procedures to learn how to edit filters.

Editing filters in datasets

Use the following procedure to edit filters in datasets.

To edit a filter in a dataset

1. Open the [Quick Suite console](#).
2. From the Quick Suite homepage, choose **Data** at left.
3. Under the **Datasets** tab, choose the dataset that you want, and then choose **Edit dataset**.
4. On the data preparation page that opens, choose **Filters** at lower left.
5. Choose the filter that you want to edit.
6. When finished editing, choose **Apply**.

Editing filters in analyses

Use the following procedure to edit filters in analyses.

To edit a filter in an analysis

1. Open the [Quick Suite console](#).
2. From the Quick Suite homepage, choose **Analyses** at left.
3. On the **Analyses** page, choose the analysis that you want to work with.
4. In the analysis, choose the **Filter** icon shown to open the **Filters** pane.
5. Choose the filter that you want to edit.
6. When finished editing, choose **Apply**.

Enabling or disabling filters

You can use the filter menu to enable or disable a filter in a dataset or an analysis. When you create a filter, it's enabled by default. Disabling a filter removes the filter from the field, but it doesn't delete the filter from the dataset or analysis. Disabled filters are grayed out in the filters pane. If you want to re-apply the filter to the field, you can simply enable it.

Use the following procedures to learn how to enable or disable filters.

Disabling filters in datasets

Use the following procedure to disable filters in datasets.

To disable a filter in a dataset

1. From the Quick Suite homepage, choose **Data** at left.
2. Under the **Datasets** tab, choose the dataset that you want, and then choose **Edit dataset**.
3. On the data preparation page that opens, choose **Filters** at lower left.
4. In the **Filters** pane at left, choose the three dots to the right of the filter that you want to disable, and then choose **Disable**. To enable a filter that was disabled, choose **Enable**.

Disabling filters in analyses

Use the following procedure to disable filters in analyses.

To disable a filter in an analysis

1. Open the [Quick Suite console](#).
2. From the Quick Suite homepage, choose **Analyses**.
3. On the **Analyses** page, choose the analysis that you want to work with.
4. In the analysis, choose the **Filter** icon to open the **Filters** pane.
5. In the **Filters** pane that opens, choose the three dots to the right of the filter that you want to disable, and then choose **Disable**. To enable a filter that was disabled, choose **Enable**.

Deleting filters

You can delete filters at any time in a dataset or analysis. Use the following procedures to learn how.

Deleting filters in datasets

Use the following procedure to delete filters in datasets.

To delete a filter in a dataset

1. Open the [Quick Suite console](#).
2. From the Quick Suite homepage, choose **Data**.
3. Under the **Datasets** tab, choose the dataset that you want, and then choose **Edit dataset**.
4. On the data preparation page that opens, choose **Filters** at lower left.
5. Choose the filter that you want to delete, and then choose **Delete filter**.

Deleting filters in analyses

Use the following procedure to delete filters in analyses.

To delete a filter in an analysis

1. Open the [Quick Suite console](#).
2. From the Quick Suite homepage, choose **Analyses**.
3. On the **Analyses** page, choose the analysis that you want to work with.
4. In the analysis, choose the **Filter** icon to open the **Filters** pane.
5. Choose the filter that you want to delete, and then choose **Delete filter**.

Using SQL to customize data

When you create a dataset or prepare your data for use in an analysis, you can customize the data in the query editor.

The query editor is made up of multiple components, as follows:

- **Query mode** – At the top left, you can choose between direct query or SPICE query modes:
 - **Direct query** – To run the SELECT statement directly against the database
 - **SPICE** – To run the SELECT statement against data that was previously stored in memory
- **Fields** – Use this section to disable fields you want to remove from the final dataset. You can add calculated fields in this section, and augment your data with SageMaker AI
- **Query archive** – Use this section to find previous version of your SQL queries.
- **Filters** – Use this section to add, edit, or remove filters.
- **Schema explorer** – This section only appears while you are editing SQL. You can use it to explore your schemas, tables, fields, and data types.
- **SQL editor** – Use this to edit your SQL. The SQL editor, which offers syntax highlighting, basic autocomplete, autoindent, and line numbering. You can specify a SQL query only for datasets that come from data sources compatible with SQL. Your SQL must conform to the target database requirements regarding syntax, capitalization, command termination, and so on. If you prefer, you can instead paste SQL from another editor.
- **Data workspace** – When the SQL editor is closed, the data workspace displays at top right with a grid background. Here you can see a graphical representation of your data objects, including queries, tables, files, and joins created in the join editor.

To view details about each table, use the data source options menu and choose **Table details** or **Edit SQL Query**. Details display for table name and alias, schema, data source name, and data source type. For upload settings on a file, choose **Configure upload settings** from the data source options menu to view or change the following settings:

- **Format** – the file format, CSV, CUSTOM, CLF, and so on
- **The starting row** – the row to start with
- **The text qualifier** – double quote or single quote
- **Header** – indicates if the file includes a header row
- **Preview rows** – A preview of the sampled rows appear at bottom right when the join configuration editor isn't in use.

- **Join configuration editor** – The join editor opens when you have more than one data object in the data workspace. To edit a join, you select the join icon between two tables (or files). Choose a join type and the fields to join on, by using the join configuration panel at the bottom of the screen. Then choose **Apply** to create the join. You must complete all joins before you can save your work.

To add more queries, tables, or files, use the **Add data** option above the workspace.

Creating a basic SQL query

Use the following procedure to connect to a data source by using a custom SQL query.

To create a basic SQL query

1. Create a new data source and validate the connection.
2. Fill in the options necessary to connection, however you don't need to select a schema or a table.
3. Choose **Use custom SQL**.
4. (Optional) You can enter your query in the SQL editor, or continue on to the next step to use the full-screen version. To enter it now, create a name for the query. Then type or paste a SQL query into the editor. The SQL editor offers syntax highlighting, basic autocomplete, autoindent, and line numbering.

(Optional) Choose **Confirm query** to validate it and view settings for direct query, SPICE memory, and SageMaker AI settings.

5. Choose **Edit/Preview data**. The full query editor appears with the SQL editor displayed. The query is processed and a sample of the query results displays in the data preview pane. You can make changes to the SQL and confirm them by choosing **Apply**. When you are done with the SQL, choose **Close** to continue.
6. At the top, enter a name for the dataset. Then choose **Save & visualize**.

Modifying existing queries

To update a SQL query

1. Open the dataset that you want to work with.

2. In the workspace with the grid, locate the box-shaped object that represents the existing query.
3. Open the options menu on the query object and choose **Edit SQL query**. If this option doesn't appear in the list, the query object isn't based on SQL.

To view previous versions of queries, open the **Query archive** at left.

Adding geospatial data

You can flag geographic fields in your data, so that Amazon Quick Sight can display them on a map. Amazon Quick Sight can chart latitude and longitude coordinates. It also recognizes geographic components such as country, state or region, county or district, city, and ZIP code or postal code. You can also create geographic hierarchies that can disambiguate similar entities, for example the same city name in two states.

Note

Geospatial charts in Amazon Quick Sight aren't currently supported in some Amazon Web Services Regions, including in China. We are working on adding support for more Regions.

Use the following procedures to add geospatial data types and hierarchies to your dataset.

To add geospatial data types and hierarchies to your dataset

1. On the data preparation page, label the geographic components with the correct data type.

There are several ways to do this. One is to choose the field under **Fields** and use the ellipses icon (...) to open the context menu.

Then choose the correct geospatial data type.

You can also change the data type in the work area with the data sample. To do this, choose the data type listed under the field name. Then choose the data type you want to assign.

2. Verify that all geospatial fields necessary for mapping are labeled as geospatial data types. You can check this by looking for the place marker icon. This icon appears under the field names across the top of the page, and also in the **Fields** pane on the left.

Also check the name of the data type, for example latitude or country.

3. (Optional) You can set up a hierarchy or grouping for geographical components (state, city), or for latitude and longitude coordinates. For coordinates, you must add both latitude and longitude to the geospatial field wells.

To create a hierarchy or grouping, first choose one of these fields in the **Fields** pane. Each field can only belong to one hierarchy. It doesn't matter which one you choose first, or what order you add the fields in.

Choose the ellipsis icon (...) next to the field name. Then choose **Add to a hierarchy**.

4. On the **Add field to hierarchy** screen, choose one of the following:
 - Choose **Create a new geospatial hierarchy** to create a new grouping.
 - Choose **Add to existing geospatial hierarchy** to add a field to a grouping that already exists. The existing hierarchies displayed include only those of matching geospatial types.

Choose **Add** to confirm your choice.

5. On the **Create hierarchy** screen, name your hierarchy.

If you are creating a latitude and longitude grouping, the **Create hierarchy** screen appears. Depending on whether you chose latitude or longitude in the previous steps, either latitude or longitude displays on this screen. Make sure that your latitude field shows under **Field to use for latitude**. Also make sure that your longitude shows under **Field to use for longitude**.

For geographical components, the **Create hierarchy** screen has two choices:

- Choose **This hierarchy is for a single country** if your data only contains one country. Choose the specific country from the list. Your data doesn't need to contain every level of the hierarchy. You can add fields to the hierarchy in any order.
- Choose **This hierarchy is for multiple countries** if your data contains more than one country. Choose the field that contains the country names.

For either hierarchy type, choose **Update** to continue.

6. Continue by adding as many fields to the hierarchy as you need to.

Your geospatial groupings appear in the **Fields** pane.

Changing a geospatial grouping

You can change a geospatial hierarchy or grouping that exists in a dataset.

Use the following procedure to edit or disband a geospatial hierarchy.

To edit or disband a geospatial hierarchy

1. Open the dataset. In the **Fields** pane, choose the hierarchy name.
2. Choose the ellipsis icon (...), then choose one of the following options.

Choose **Disband hierarchy** to remove the hierarchy from the dataset. You can't undo this operation. However, you can recreate your hierarchy or grouping by starting again at step 1. Disbanding the hierarchy doesn't remove any fields from the dataset.

Choose **Edit hierarchy** to make changes to the hierarchy. Doing this reopens the creation screens, so you can make different choices in rebuilding your hierarchy.

Geospatial troubleshooting

Use this section to discover the Amazon Quick Sight requirements for correctly processing geospatial data. If Amazon Quick Sight doesn't recognize your geospatial data as geospatial, use this section to help troubleshoot the issue. Make sure that your data follows the guidelines listed, so that it works in geospatial visuals.

Note

Geospatial charts in Amazon Quick Sight currently aren't supported in some Amazon Web Services Regions, including in China. We are working on adding support for more Regions. If your geography follows all the guidelines listed here, and still generates errors, contact the Amazon Quick Sight team from within the Amazon Quick Sight console.

Topics

- [Geocoding issues](#)
- [Issues with latitude and longitude](#)
- [Supported administrative areas and postal codes by country](#)

Geocoding issues

Amazon Quick Sight geocodes place names into latitude and longitude coordinates. It uses these coordinates to display place names on the map. Amazon Quick Sight skips any places that it can't geocode.

For this process to work properly, your data must include at least the country. Also, there can't be duplicate place names inside of a parent place name.

A few issues prevent place names from showing up on a map chart. These issues include unsupported, ambiguous, or invalid locations, as described following.

Topics

- [Issues with unsupported areas](#)
- [Issues with ambiguous locations](#)
- [Issues with invalid geospatial data](#)
- [Issues with the default country in geocoding](#)

Issues with unsupported areas

To map unsupported locations, include latitude and longitude coordinates in your data. Use these coordinates in the geospatial field well to make locations show on a map chart.

Issues with ambiguous locations

Geospatial data can't contain ambiguous locations. For example, suppose that the data contains a city named **Springfield**, but the next level in the hierarchy is country. Because multiple states have a city named **Springfield**, it isn't possible to geocode the location to a specific point on a map.

To avoid this problem, you can add enough geographical data to indicate what location should show on a map chart. For example, you can add a state level into your data and its hierarchy. Or, you might add latitude and longitude.

Issues with invalid geospatial data

Invalid geospatial data occurs when a place name (a city, for example) is listed under an incorrect parent (a state, for example). This issue might be a simple misspelling, or data entry error.

Note

Amazon Quick Sight doesn't support regions (for example, West Coast or South) as geospatial data. However, you can use a region as a filter in a visual.

Issues with the default country in geocoding

Make sure that you are using the correct default country.

The default for each hierarchy is based on the country or country field that you choose when you create the hierarchy.

To change this default, you can return to the **Create hierarchy** screen. Then edit or create a hierarchy, and choose a different country.

If you don't create a hierarchy, your default country is based on your Amazon Web Services Region. For details, see the following table.

Region	Default country
US West (Oregon) Region	US
US East (Ohio) Region	
US East (N. Virginia) Region	
Asia Pacific (Singapore)	Singapore
Asia Pacific (Sydney)	Australia
Europe (Ireland) Region	Ireland

Issues with latitude and longitude

Amazon Quick Sight uses latitude and longitude coordinates in the background to find place names on a map. However, you can also use coordinates to create a map without using place names. This approach also works with unsupported place names.

Latitude and longitude values must be numeric. For example, the map point indicated by **28.5383355 -81.3792365** is compatible with Amazon Quick Sight. But **28° 32' 18.0096'' N 81° 22' 45.2424'' W** is not.

Topics

- [Valid ranges for latitude and longitude coordinates](#)
- [Using coordinates in degrees, minutes, and seconds \(DMS\) format](#)

Valid ranges for latitude and longitude coordinates

Amazon Quick Sight supports latitude and longitude coordinates within specific ranges.

Coordinate	Valid range
Latitude	Between -90 and 90
Longitude	Between -180 to 180

Amazon Quick Sight skips any data outside these ranges. Out-of-range points can't be mapped on a map chart.

Using coordinates in degrees, minutes, and seconds (DMS) format

You can use a calculated field with a formula to create a numeric latitude and longitude out of character strings. Use this section to find different ways that you can create calculated fields in Amazon Quick Sight, to parse GPS latitude and longitude into numeric latitude and longitude.

The following sample converts latitude and longitude to numeric format from separate fields. For example, suppose that you parse **51° 30' 26.4636'' N 0° 7' 39.9288'' W** using space as a delimiter. In this case, you can use something like the following sample to convert the resulting fields to numeric latitude and longitude.

In this example, the seconds are followed by two single quotation marks. If your data has a double quotation mark instead, then you can use `strlen(LatSec)-1` instead of `strlen(LatSec)-2`.

```
/*Latitude*/  
    ifelse(
```

```

    LatDir = "N",
    parseInt(split(LatDeg, "°", 1)) +
        (parseFloat(split(LatMin, "'", 1) ) /60) +
        (parseFloat((substring(LatSec, 1, strlen(LatSec)-2) ) ) /3600),
    (parseInt(split(LatDeg, "°", 1)) +
        (parseFloat(split(LatMin, "'", 1) ) /60) +
        (parseFloat((substring(LatSec, 1, strlen(LatSec)-2) ) ) /3600)) * -1
    )

/*Longitude*/
    ifelse(
    LongDir = "E",
    parseInt(split(LongDeg, "°", 1)) +
        (parseFloat(split(LongMin, "'", 1) ) /60) +
        (parseFloat((substring(LongSec, 1, strlen(LongSec)-2) ) ) /3600),
    (parseInt(split(LongDeg, "°", 1)) +
        (parseFloat(split(LongMin, "'", 1) ) /60) +
        (parseFloat((substring(LongSec, 1, strlen(LongSec)-2) ) ) /3600)) * -1
    )

```

If your data doesn't include the symbols for degree, minute and second, the formula looks like the following.

```

/*Latitude*/
    ifelse(
        LatDir = "N",
        (LatDeg + (LatMin / 60) + (LatSec / 3600)),
        (LatDeg + (LatMin / 60) + (LatSec / 3600)) * -1
    )

/*Longitude*/
    ifelse(
        LongDir = "E",
        (LongDeg + (LongMin / 60) + (LongSec / 3600)),
        (LongDeg + (LongMin / 60) + (LongSec / 3600)) * -1
    )

```

The following sample converts **53°21'N 06°15'W** to numeric format. However, without the seconds, this location doesn't map as accurately.

```
/*Latitude*/
```

```

ifelse(
  right(Latitude, 1) = "N",
  (parseInt(split(Latitude, '°', 1)) +
    parseDecimal(substring(Latitude, (locate(Latitude, '°',3)+1), 2) ) / 60) ,
  (parseInt(split(Latitude, '°', 1)) +
    parseDecimal(substring(Latitude, (locate(Latitude, '°',3)+1), 2) ) / 60) * -1
)

/*Longitude*/
ifelse(
  right(Longitude, 1) = "E",
  (parseInt(split(Longitude, '°', 1)) +
    parseDecimal(substring(Longitude, (locate(Longitude, '°',3)+1), 2) ) / 60) ,
  (parseInt(split(Longitude, '°', 1)) +
    parseDecimal(substring(Longitude, (locate(Longitude, '°',3)+1), 2) ) / 60) *
-1
)

```

The formats of GPS latitude and longitude can vary, so customize your formulas to match your data. For more information, see the following:

- [Degrees Minutes Seconds to Decimal Degrees](#) on LatLong.net
- [Converting Degrees/Minutes/Seconds to Decimals using SQL](#) on Stack Overflow
- [Geographic Coordinate Conversion](#) on Wikipedia

Supported administrative areas and postal codes by country

The following is a list of supported administrative areas by country.

Supported administrative areas

Country name	Country code	Country	State	County	City
Aruba	ABW	Country	Regions	Zones	
Afghanistan	AFG	Country	Wilayat	Wuleswali	Localities/Urban Areas

Country name	Country code	Country	State	County	City
Angola	AGO	Country	Provinces /Províncias	Municipios	Localities/Urban Areas
Anguilla	AIA	Country	Parishes		
Albania	ALB	Country	Qarqe/Qark	Communes/Bashki	Njësi/Localities/Urban Areas
Andorra	AND	Country	Parishes/Parròquies	Localities/Urban Areas	
United Arab Emirates	ARE	Country	Emirates	Municipalities	Cities/Localities/Urban Areas
Argentina	ARG	Country	Provinces /Provincias	Departamentos/Departments	Comunas/Barrios
Armenia	ARM	Country	Provinces /Marzpet		Localities/Urban Areas
American Samoa	ASM	Country	Districts	Counties	Villages

Country name	Country code	Country	State	County	City
Antarctica	ATA	Country			
French Southern Territories	ATF	Country	Districts		
Antigua and Barbuda	ATG	Country	Parishes		Localities/Urban Areas
Australia	AUS	Country	States	Local Government Areas	Suburbs/Urban Centers
Austria	AUT	Country	States/Bundesländer	Districts/Bezirke	Municipalities/Gemeinden/Urban Areas/Stadtteil
Azerbaijan	AZE	Country	Regions/Iqtisadi Rayonlar	Districts/Rayonlar	Localities/Urban Areas
Burundi	BDI	Country	Provinces	Communes	Localities/Urban Areas

Country name	Country code	Country	State	County	City
Belgium	BEL	Country	Regions/ Gewest	Provinces /Provincie	Districts/ Arrondissements/ Municipalities/ Communes
Benin	BEN	Country	Departments	Communes	Localities/ Urban Areas
Bonaire, Sint Eustasius, and Saba	BES	Country	Municipalities		Localities/ Urban Areas
Burkina Faso	BFA	Country	Regions	Provinces	Communes/ Localities/ Urban Areas
Bangladesh	BGD	Country	Divisions /Bibhag	Districts/ Zila	Subdistricts/ Upzila/ Localities/ Urban Areas
Bulgaria	BGR	Country	Oblasts	Obshtina	Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Bahrain	BHR	Country	Governorates	Constituencies	Localities
Bahamas	BHS	Country	Island Groups	Districts	Towns
Bosnia and Herzegovina	BIH	Country	Federation/ Republika	Kanton	Opština/ Localities /Urban Areas
Saint Barthélemy	BLM	Country			Localities/ Urban Areas
Belarus	BLR	Country	Voblasts	Rayon	Selsoviet/ Localities/ Urban Areas
Belize	BLZ	Country	Districts	Constituencies	Localities/ Urban Areas
Bermuda	BMU	Country	Parishes		Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Bolivia	BOL	Country	Provinces /Provinci as	Departame ntos/ Depa rtments	Municipal ities/ Mun icipios/L ocalities /Urban Areas
Brazil	BRA	Country	Provinces /States/ U nidades	Municipal ities/ Mun icipios	Localitie s/Urban Areas
Barbados	BRB	Country	Parishes		Localitie s/Urban Areas
Brunei	BRN	Country	Districts/ Dawaïr	Subdistri cts/ Mukim	Villages/ Kampung/ L ocalities /Urban Areas
Bhutan	BTN	Country	Districts/ Dzongkha g		Localitie s/Urban Areas
Bouvet Island	BVT	Country			

Country name	Country code	Country	State	County	City
Botswana	BWA	Country	Districts	Subdistricts	Localities/Urban Areas
Central African Republic	CAF	Country	Regions	Prefectures	Sub Prefectures/Communes
Canada	CAN	Country	Provinces/Territories	Census Divisions	Census Subdivisions/Localities/Urban Areas
Switzerland	CHE	Country	Cantons/Kanton/Cantone/Chanton	District/Bezirk/Distretto/Circulo	"Commune/Gemeinde/Comune/Cumün/Localities/Urban Areas"
Chile	CHL	Country	Regions/Regiones	Province/Provincias	Communes/Comunas/Localities/Urban Areas

Country name	Country code	Country	State	County	City
China, People's Republic of	CHN	Country	Provinces	Prefectures	Cities/Counties
Cote d'Ivoire	CIV	Country	Districts	Regions	Departments/Sub Prefectures
Cameroon	CMR	Country	Provinces /Regions	Departments	Arrondissements/Cities
Congo, Democratic Republic of	COD	Country	Provinces	Districts	Localities/Urban Areas
Congo, Republic of	COG	Country	Departments		Communes/Arrondissements
Cook Islands	COK	Country	Island Councils		
Colombia	COL	Country	Departments	Municipios	Localities/Urban Areas

Country name	Country code	Country	State	County	City
Comoros	COM	Country	Autonomo s Islands/ îles Autonome:		Villes/Vi llages
Clipperton Island	CPT	Country			
Cape Verde	CPV	Country	Ilhas	Concelhos	Localitie s/Urban Areas
Costa Rica	CRI	Country	Provincia s	Cantons	Distritos /Localiti es/ Urban Areas
Cuba	CUB	Country	Provincia s	Município s	Localitie s/Urban Areas
Curaçoa	CUW	Country			Localitie s/Urban Areas
Cayman Islands	CYM	Country	Districts		

Country name	Country code	Country	State	County	City
Cyprus	CYP	Country	Districts/ Eparchies	Municipal ities/Dim os	Localitie s/Urban Areas/ Sinikia
Czech Republic	CZE	Country	Regions/ Kraj	Municipal ities/Orp	Obec/ Mesto
Germany	DEU	Country	Bundeslan d/States	Kreis/Di stricts	Gemeinde/ Municipal ities/Sta dtteil/L ocalities/ Urban Areas
Djibouti	DJI	Country	Regions		Localitie s/Urban Areas
Dominica	DMA	Country	Parishes		Localitie s/Urban Areas
Denmark	DNK	Country	Regions	Provinces	Municipal ities/Loc alities/ Urban Areas

Country name	Country code	Country	State	County	City
Dominican Republic	DOM	Country	Regions/ Regiones	Provinces /Provincias	Municipal ities/ Mun icipios/L ocalities /Urban Areas
Algeria	DZA	Country	Provinces /Wilayas	Districts	Municipal ities/Bal adiyas/ Lo calities/ Urban Areas
Ecuador	ECU	Country	Provinces	Cantons	Parishes/ Localitie s/Urban Areas
Egypt	EGY	Country	Governora tes/ Muhaf azat	Municipal Divisions /Markaz	Towns/ Cities/ Sub Municipal Divisions
Eritrea	ERI	Country	Regions/ Zoba	Districts/ Subzobas	Localitie s/Urban Areas

Country name	Country code	Country	State	County	City
Spain	ESP	Country	Autonomos Communities/ Comunidades Autonomas	Provincias	Municipios/Localities/ Urban Areas
Estonia	EST	Country	Maakond	Omavalitsus/Linn/ Vald	Küla/ Localities/ Urban Areas
Ethiopia	ETH	Country	Regions/ Kililoch	Zones/ Zonouch	Localities/ Urban Areas
Finland	FIN	Country	Regions/ Maaakunta	Sub-Regions/ Seutukunta	Municipalities/Kunta/ Localities/ Urban Areas
Fiji	FJI	Country	Divisions	Provinces	Districts/ Villages
Falkland Islands	FLK	Country			

Country name	Country code	Country	State	County	City
France	FRA	Country	Regions	Départements	Arrondissements/ Cantons
Faroe Islands	FRO	Country	Regions/ Syslur	Municipalities/ Kommunur	Localities/ Urban Areas
Federated States of Micronesia	FSM	Country	States		
Gabon	GAB	Country	Provinces	Departments	Localities/ Urban Areas
United Kingdom	GBR	Country	Nations	Counties	Districts/ Localities/ Urban Areas
Georgia	GEO	Country	Regions/ Mkhare	Municipalities/ Municipality	Localities/ Urban Areas
Ghana	GHA	Country	Regions	Districts	Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Gibraltar	GIB	Country			Localities/Urban Areas
Guinea	GIN	Country	Regions	Prefectures	Sub Prefectures/Localities/Urban Areas
Guadeloupe	GLP	Country	Arrondissements	Communes	Localities/Urban Areas
Gambia	GMB	Country	Regions	Districts	Localities/Urban Areas
Guinea Bissau	GNB	Country	Regions	Sectors	Localities/Urban Areas
Equatorial Guinea	GNQ	Country	Regions	Provincias	Distritos/Localities/Urban Areas

Country name	Country code	Country	State	County	City
Greece	GRC	Country	Regions/ P eriphenie s	Regional Units Peri Enotities	Municipal ities/ Domo i/ Locali ties/ Urban Areas
Grenada	GRD	Country	States	Parishes/ Dependenc ies	Localitie s/Urban Areas
Greenland	GRL	Country	Municipal ities/ Kom munia		
Guatemala	GTM	Country	Departmer ts/ Depart amentos	Municipal ities/ Mun icipios	Localitie s/Urban Areas
French Guiana	GUF	Country	Arrondiss ements	Communes	Localitie s/Urban Areas
Guam		Country = USA	States	Districts	

Country name	Country code	Country	State	County	City
Guyana	GUY	Country	Regions	Neighborhood Councils	Localities/Urban Areas
Hong Kong	HKG	Country	Districts	Subdistricts	Localities/Urban Areas
Heard and McDonald Islands	HMD	Country			
Honduras	HND	Country	Departments/ Departamentos	Municipalities/ Municipios	Localities/Urban Areas
Croatia	HRV	Country	Counties	Municipalities	Localities/Urban Areas
Haiti	HTI	Country	Departments/ Départements	Districts/ Arrondissements	Communes/ Localities/Urban Areas
Hungary	HUN	Country	Regiok	Megyék	Járások/ Városok

Country name	Country code	Country	State	County	City
Indonesia	IDN	Country	Provinces /Provinsi	Regency/ Kabupaten	Districts/ Kecamatan/ Localities/ Urban Areas
India	IND	Country	States/ Territories	Districts	Subdistricts/ Towns/ Localities/ Urban Areas
British Indian Ocean Territory	IOT	Country			
Ireland	IRL	Country	Regions	Counties	Electoral Divisions/ Localities/ Urban Areas
Iran	IRN	Country	Provinces /Ostanha	Counties/ Shahrestan	Localities/ Dehestân

Country name	Country code	Country	State	County	City
Iraq	IRQ	Country	Governorates/ Muhafazat	Districts/ Qadaa/ Kaza	Urban Areas/ Localities
Iceland	ISL	Country	Regions/ Landsvæðing	Municipalities/ Sveitarfélög	Localities/ Urban Areas
Israel	ISR	Country	Districts	Cities/ Local Councils	Localities/ Urban Areas
Italy	ITA	Country	Regioni	Province	Communes/ Localities/ Urban Areas
Jamaica	JAM	Country	Counties	Parishes	Constituencies/ Localities/ Urban Areas
Jordan	JOR	Country	Governorates	Districts	Subdistricts/ Cities

Country name	Country code	Country	State	County	City
Japan	JPN	Country	Prefectur es		Cities/Di stricts/M unicipali ties
Kazakhstan	KAZ	Country	Regions/ Oblystar	Districts/ Audandar	Towns/ Kent/ Localit ies/ Urban Areas
Kenya	KEN	Country	Counties	Constitue ncies	Localitie s/Urban Areas/ Suburbs
Kyrgyzstan	KGZ	Country	Regions/ O blasttar	Districts/ Raions	Localitie s/Urban Areas
Cambodia	KHM	Country	Provinces /Khaet	Districts/ Srők	Communes/ Khum/ Loca lities/ Urban Areas
Kiribati	KIR	Country	Districts	Island Councils	

Country name	Country code	Country	State	County	City
Saint Kitts and Nevis	KNA	Country	Parishes	States	Localities/ Urban Areas
South Korea	KOR	Country	Provinces /Do	Districts/ Si/Gun	Localities/ Urban Areas
Kuwait	KWT	Country	Governorates/ Muhafazah	Areas/ Min taqah	Cities/ Communities
Laos	LAO	Country	Provinces / Khoueng	Districts/ Muang	Localities/ Urban Areas
Lebanon	LBN	Country	Governorates/ Muhafazat	Districts/ Qadaa	Municipalities/ Localities/ Urban Areas
Liberia	LBR	Country	Counties	Districts	Clans/ Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Libya	LBY	Country	Districts/ Shabiya		Cities/Localities/ Urban Areas
Saint Lucia	LCA	Country	Districts/ Quarters		Localities/ Urban Areas
Liechtenstein	LIE	Country	Districts/ Bezirk	Municipalities/ Gemeinden	Localities/ Urban Areas
Sri Lanka	LKA	Country	Provinces	Districts	Divisional Secretariats/ Localities/ Urban Areas
Lesotho	LSO	Country	Districts	Constituencies	Community Councils/ Localities
Lithuania	LTU	Country	Apskritis	Savivaldybė	Seniūnija

Country name	Country code	Country	State	County	City
Luxembourg	LUX	Country	Cantons/ Kantounen/ Kantone	Communes/ Gemengen/ Gemeinden	Localities/ Ortschaft/ Uertschaft/Cities
Latvia	LVA	Country	Regions	Municipalities/ Novadi	Pilsētas/ Pagasti/ Localities/ Urban Areas
Macao	MAC	Country	Parishes	Districts	
Saint Martin	MAF	Country			Localities/ Urban Areas
Morocco	MAR	Country	Regions	Provinces/ Prefectures	Communes/ Localities/ Urban Areas
Monaco	MCO	Country	Communes	Wards/ Quartiers	

Country name	Country code	Country	State	County	City
Moldova	MDA	Country	Raion	Comuna	Localities/Urban Areas
Madagascar	MDG	Country	Regions/ Faritra	Districts	Communes/ Localities/Urban Areas
Maldives	MDV	Country	Atolls/Cities	Islands	
Mexico	MEX	Country	Estados	Municipios/ Delegaciones	Colonias/ Localities/Urban Areas
Marshall Islands	MHL	Country	Municipalities		
Macedonia	MKD	Country	Statistical Regions	Opština	Localities/Urban Areas
Mali	MLI	Country	Regions	Communes	Localities/Urban Areas

Country name	Country code	Country	State	County	City
Malta	MLT	Country	Districts	Local Councils/ Kunsilli Lokali	Localities/ Urban Areas
Myanmar	MMR	Country	States/ Regions/ Union Territories	Districts	Townships/ Localities/ Urban Areas
Montenegro	MNE	Country	Opštine/ Municipalities		Localities/ Urban Areas
Mongolia	MNG	Country	Regions	Provinces/ Aimags	Districts/ Sums/ Localities/ Urban Areas
Northern Mariana Islands	MNP	Country	Municipalities		
Mozambique	MOZ	Country	Provinces	Districts/ Distritos	Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Mauritania	MRT	Country	Regions	Départements	Localities/Urban Areas
Montserrat	MSR	Country	Parishes	Regions	Localities/Urban Areas
Martinique	MTQ	Country	Arrondissements	Communes	Localities/Urban Areas
Mauritius	MUS	Country	Islands	Districts	Wards/ Localities/ Urban Areas
Malawi	MWI	Country	Regions	Districts	Localities/Urban Areas
Malaysia	MYS	Country	States/ Negeri	Districts/ Daïra/Daerah	Subdistricts/ Mukim/ Localities/ Urban Area/ Bahagian Kecil

Country name	Country code	Country	State	County	City
Mayotte	MYT	Country	Communes		Villages
Namibia	NAM	Country	Provinces	Constituencies	Suburbs/ Localities
New Caledonia	NCL	Country	Provinces	Communes	
Niger	NER	Country	Regions	Departments	Localities/ Urban Areas
Nigeria	NGA	Country	States	Local Government Areas	Towns/ Cities
Nicaragua	NIC	Country	Departments/ Departamentos	Municipalities/ Municipios	Localities/ Urban Areas
Niue	NIU	Country	Villages		Towns
Netherlands	NLD	Country	Counties/ Fylker	Districts/ Okonomisk	Municipalities, Kommuner, Localities, or Urban Areas

Country name	Country code	Country	State	County	City
Norway	NOR	Country	Counties/ Fylker	Districts/ Okonomisk	Municipal ities, Kommuner, Localities, or Urban Areas
Nepal	NPL	Country	Provinces / Pradeshaharu	Districts/ Jilla	Municipal ities/Loc alities/ Urban Areas
Nauru	NRU	Country	Districts		
New Zealand	NZL	Country	Regions	Territorial Authoriti es	Statistic al Areas/ Loc alities/ Urban Areas
Oman	OMN	Country	Governorates/ Muhafazah	Provinces /Wilayat	Cities/ Urban Areas/ Com munities
Pakistan	PAK	Country	Provinces	Districts	Localities/ Tehsils

Country name	Country code	Country	State	County	City
Panama	PAN	Country	Provinces /Provinci as	Districts/ Distrito	Corregimi entos/ Loc alities/ Urban Areas
Pitcairn Islands	PCN	Country	Islands		
Peru	PER	Country	Regions	Districts	Distritos /Localiti es/ Urban Areas
Philippines	PHL	Country	Regions/ Rehiyon	Provinces / Lalawiga n	Municipal ities/ Mun isipiyos/ Cities/Lu ngsod
Palau	PLW	Country	States		
Papua New Guinea	PNG	Country	Regions	Provinces	Districts /Localiti es/ Urban Areas

Country name	Country code	Country	State	County	City
Poland	POL	Country	Provinces / Voivodes hips	Counties/ Powiats	Communes/ Gminas/ Towns/ Dziel nicas
North Korea	PRK	Country	Provinces		Localitie s/Urban Areas
Portugal	PRT	Country	Districts/ Distritos	Municipal ities/Con celhos	Civil Parish/Fr eguesias/ Localitie s/Urban Areas
Paraguay	PRY	Country	Departmer ts	Distritos	Localitie s/Urban Areas
Palestine	PSE	Country	Territori es	Governora tes/ Muhaf azat	Localitie s/Urban Areas
French Polynesia	PYF	Country	Subdivisi ons/lles	Communes	

Country name	Country code	Country	State	County	City
Qatar	QAT	Country	Municipal ities/Bal adiyat	Zones	Localitie s/Urban Areas
Réunion	REU	Country	Arrondiss ements	Communes	Localitie s/Urban Areas
Romania	ROU	Country	Regions/ Judete	Communes	Towns/ Oraş
Russia	RUS	Country	Federal District/ Federal'n yy Okrug	Oblast'	Rayon/ Raion/ Urban Area/ Gorod
Rwanda	RWA	Country	Provinces	Districts	Sectors/ S ecteurs/ L ocalities /Urban Areas
Saudi Arabia	SAU	Country	Regions/ Manatiq	Governora tes/ Muhaf azat	Municipal ities/ Amanah

Country name	Country code	Country	State	County	City
Sudan	SDN	Country	States/ Wilaya'at		Localities/ Urban Areas
Senegal	SEN	Country	Regions	Departments	Arrondissements/ Localities/ Urban Areas
Singapore	SGP	Country	Districts	Constituencies	Wards
Saint Helena	SHN	Country	Islands	Districts	Localities/ Urban Areas
Solomon Islands	SLB	Country	Provinces	Constituencies	Wards
Sierra Leone	SLE	Country	Provinces	Districts	Chiefdoms/ Localities/ Urban Areas
El Salvador	SLV	Country	Departments/ Departamentos	Municipalities/ Municipios	Localities/ Urban Areas

Country name	Country code	Country	State	County	City
San Marino	SMR	Country	Municipalities/Castelli	Localities/Urban Areas	
Somalia	SOM	Country	Regions/ Gobolada		Localities/Urban Areas
Saint Pierre and Miquelon	SPM	Country	Communes		
Serbia	SRB	Country	Autonomous Pokrajina /Regions	Okrug/ Districts	Opština/ Municipalities/Localities/ Urban Areas
South Sudan	SSD	Country	States/ Wilayat	Counties	Localities/Urban Areas
São Tomé and Príncipe	STP	Country	Provinces	Districts	Localities/Urban Areas
Suriname	SUR	Country	Districts/ Distrikt	Resorts	Localities/Urban Areas

Country name	Country code	Country	State	County	City
Slovakia	SVK	Country	Regions/ Kraje	Districts/ Okresy	Municipal ities/ Obec/ Mestská cast
Slovenia	SVN	Country	Regions/ Regi	Upravne Enote	Municipal ities/ Obcine/ Local ities/Urb an Areas
Sweden	SWE	Country	Counties	Municipal ities	Localitie s/Urban Areas
Eswatini	SWZ	Country	Regions	Tinkhundl a	Towns/ Suburbs/ Loca lities
Sint Maarten	SXM	Country	Settlemen ts		
Seychelles	SYC	Country	Districts		Localitie s/Urban Areas

Country name	Country code	Country	State	County	City
Syria	SYR	Country	Governorates	Districts/ Muhafazah	Cities/Localities/ Urban Areas
Turks and Caicos Islands	TCA	Country	Districts	Localities	
Chad	TCD	Country	Regions	Départements	Arrondissements/ Localities/ Urban Areas
Togo	TGO	Country	Regions/ Provinces	Prefectures	Localities/ Urban Areas
Thailand	THA	Country	Provinces / Changwat	Districts/ Amphoe	Subdistricts/ Tambon/Localities/ Urban Areas
Tajikistan	TJK	Country	Provinces / Regions	Districts/ Raion/Raion	Localities/ Urban Areas

Country name	Country code	Country	State	County	City
Tokelau	TKL	Country	Atolls		
Turkmenistan	TKM	Country	Provinces /Welayat	Districts/ Etraplar	Towns
East Timor (Timor-Le ste)	TLS	Country	Municipal ities	Administr ative Post	Localitie s/Urban Areas
Tonga	TON	Country	Subdivisi ons		
Trinidad and Tobago	TTO	Country	Municipal ities		Localitie s/Urban Areas
Tunisia	TUN	Country	Governate s/ Wilayahs	Delegatio ns/ Mutama diyats	Municipal ities/Sh aykhats/ Lo calities/ Urban Areas

Country name	Country code	Country	State	County	City
Turkey	TUR	Country	Provinces /Il	Districts/ İlce	Urban Areas/ Belde/ Subdistricts/ Bucak/ Neighborhoods/ Mahalle
Tuvalu	TUV	Country	Islands		
Tanzania	TZA	Country	Provinces /Mkoa	Districts/ Wilaya	Localities/Urban Areas
Uganda	UGA	Country	Regions	Districts	Counties/ Localities/Urban Areas
Ukraine	UKR	Country	Oblast/ Mista/ Avtonomna Respublika	Raions	Settlement Councils/ Rural Councils/ Localities/Urban Areas
United States Minor Outlying Islands	UMI	Country	Islands/ Atolls		

Country name	Country code	Country	State	County	City
Uruguay	URY	Country	Departments/ Departamentos	Municipios/ Municipalidades/S ecciones	Segmentos/ Localities/ Urban Areas
United States of America	USA	Country	States/ Territories	Counties	MCD/ CCD/ Post Localities/ Municipalities
Uzbekistan	UZB	Country	Regions/ Viloyatlar	Districts/ Tumanlar	Localities/ Urban Areas
Vatican City	VAT	Country			Localities/ Urban Areas
Saint Vincent and the Grenadines	VCT	Country	Parishes	Divisions	Localities/ Urban Areas
Venezuela	VEN	Country	States/ Estados	Municipalities/ Municipios	Localities/ Urban Areas/ Parish/ Parroquias

Country name	Country code	Country	State	County	City
British Virgin Islands	VGB	Country	Districts		
Vietnam	VNM	Country	Provinces /Cities	Districts	Wards/ Loc alities/ Urban Areas
Vanuatu	VUT	Country	Provinces		
Wallis and Futuna Islands	WLF	Country	Districts/ Rayaumes		
Samoa	WSM	Country	Districts/ Itūmālō	Towns	Localitie s/Urban Areas
Kosovo	XKS	Country	Districts	Municipal ities	Localitie s/Urban Areas
Yemen	YEM	Country	Governora tes/ Muhaf azat	Districts/ Muderiah	Localitie s/Urban Areas
South Africa	ZAF	Country	Provinces	Districts	Municipal ities/War ds

Country name	Country code	Country	State	County	City
Zambia	ZMB	Country	Provinces	Districts	Suburbs/ L ocalities
Zimbabwe	ZWE	Country	Provinces	Districts/ Muderiah	Localitie s/Urban Areas

The following is a list of the supported postal code formats by country, including the number of digits and an example postal code.

 Note

PO BOX zipcodes are not supported postal code formats. Union territory zip codes that are used in India are also not supported.

Supported postal codes

Country	Postal format	Example
Afghanistan	4 digit	1001
Albania	4 digit	1001
Algeria	5 digit	01000
American Samoa	5 digit	96799
Andorra	5 digit	AD100

Country	Postal format	Example
Anguilla	6 digit	AI-2640
Argentina	5 digit	A4126
Armenia	2 digit	00
Australia	4 digit	0800
Austria	4 digit	1010
Azerbaijan	2 digit	01
Brunei Darussalam	6 digit	BA1111
Bahrain	4 digit	0101
Bangladesh	2 digit	10
Belarus	6 digit	202115
Belgium	4 digit	1000
Bermuda	4 digit	CR 01
Bhutan	2 digit	11
Bosnia and Herzegovina	5 digit	70101
Brazil	5 digit	01001

Country	Postal format	Example
British Indian Ocean Territory	Alphanumeric - 5 digit	BBND 1
British Virgin Islands	4 digit	1110
Bulgaria	4 digit	1000
Cabo Verde	4 digit	1101
Cambodia	2 digit	01
Canada	3 digit	A0A
Cayman Islands	Alphanumeric - 7 digit	KY1-1000
Chile	3 digit	100
China	4 digit	0100
Colombia	4 digit	0500
Costa Rica	5 digit	10101
Croatia	5 digit	10000
Cuba	1 digit	1
Cyprus	4 digit	1010

Country	Postal format	Example
Czechia	5 digit	100 00
Democratic Republic of the Congo	4 digit	1001
Denmark	4 digit	1050
Dominican Republic	5 digit	10101
Ecuador	6 digit	010101
Egypt	2 digit	11
El Salvador	4 digit	1101
Estonia	5 digit	10001
Falkland Islands	Alphanumeric - 5 digit	FIQQ 1
Faroe Islands	3 digit	100
Finland	5 digit	00100
France	5 digit	01000
French Guiana	5 digit	97300
French Polynesia	5 digit	98701

Country	Postal format	Example
Georgia	2 digit	01
Germany	5 digit	01067
Ghana	2 digit	A2
Gibraltar	Alphanumeric - 5 digit	GX11 1
Greece	5 digit	104 31
Greenland	4 digit	3900
Guadeloupe	5 digit	97100
Guam	5 digit	96910
Guatemala	5 digit	01001
Guernsey	Alphanumeric - 4 digit, 5 digit	GY1 1, GY10 1
Guinea-Bissau	4 digit	1000
Haiti	4 digit	1110
Holy See	5 digit	00120
Honduras	2 digit	11

Country	Postal format	Example
Hungary	4 digit	1007
Iceland	3 digit	101
India	6 digit	110001
Indonesia	5 digit	10110
Iran	2 digit	11
Iraq	2 digit	10
Ireland	3 digit	A41
Isle of Man	Alphanumeric - 4 digit	IM1 1
Israel	5 digit	10292
Italy	5 digit	00010
Japan	7 digit	001-0010
Jersey	Alphanumeric - 4 digit	JE2 3
Jordan	5 digit	11100
Kazakhstan	4 digit	0100

Country	Postal format	Example
Kenya	1 digit	0
Kiribati	6 digit	KI0101
Kosovo	5 digit	10000
Kuwait	2 digit	00
Kyrgyzstan	4 digit	7200
Laos	2 digit	01
Latvia	4 digit	1001
Lesotho	1 digit	1
Liberia	2 digit	10
Liechtenstein	4 digit	9485
Lithuania	5 digit	00100
Luxembourg	4 digit	1110
Macedonia	4 digit	1000
Madagascar	3 digit	101

Country	Postal format	Example
Malawi	3 digit	101
Malaysia	5 digit	01000
Maldives	2 digit	00
Malta	3 digit	ATD
Marshall Islands	3 digit	969
Martinique	5 digit	97200
Mauritius	3 digit	111
Mayotte	5 digit	97600
Mexico	5 digit	01000
Micronesia	5 digit	96941
Moldova	4 digit	2001
Monaco	5 digit	98000
Mongolia	4 digit	1200
Montenegro	5 digit	81000

Country	Postal format	Example
Montserrat	4 digit	1120
Morocco	5 digit	10000
Mozambique	4 digit	1100
Myanmar	2 digit	01
Namibia	3 digit	100
Nepal	3 digit	101
Netherlands	4 digit	1011
New Caledonia	5 digit	98800
New Zealand	4 digit	0110
Nicaragua	3 digit	110
Niger	4 digit	1000
Nigeria	4 digit	1002
Niue	4 digit	9974
Norfolk Island	4 digit	2899

Country	Postal format	Example
Northern Mariana Islands	5 digit	96950
Norway	4 digit	0010
Oman	1 digit	1
Pakistan	2 digit	10
Palau	5 digit	96939
Palestine	4 digit	P104
Papua New Guinea	3 digit	111
Paraguay	6 digit	001001
Peru	5 digit	01000
Philippines	4 digit	1000
Pitcairn	Alphanumeric - 5 digit	PCRN 1
Poland	5 digit	00-002
Portugal	4 digit	1000
Puerto Rico	5 digit	00601

Country	Postal format	Example
Romania	6 digit	010011
Russia	6 digit	101000
Réunion	5 digit	97400
Saint Barthélemy	5 digit	97133
Saint Helena, Ascension and Tristan da Cunha	Alphanumeric - 5 digit	ASCN 1
Saint Lucia	7 digit	LC01 101
Saint Martin	5 digit	97150
Saint Pierre and Miquelon	5 digit	97500
Saint Vincent and the Grenadines	4 digit	VC01
Samoa	2 digit	11
San Marino	5 digit	47890
Saudi Arabia	2 digit	12
Senegal	5 digit	10000
Serbia	5 digit	11000

Country	Postal format	Example
Singapore	6 digit	018906
Slovakia	5 digit	010 01
Slovenia	4 digit	1000
South Africa	4 digit	0001
South Georgia and the South Sandwich Islands	Alphanumeric - 5 digit	SIQQ 1
South Korea	5 digit	01000
Spain	5 digit	01001
Sri Lanka	2 digit	00
Sudan	2 digit	11
Svalbard and Jan Mayen	4 digit	8099
Swaziland	1 digit	H
Sweden	5 digit	111 15
Switzerland	4 digit	1000
Tajikistan	4 digit	7340

Country	Postal format	Example
Tanzania, United Republic of	3 digit	111
Thailand	5 digit	10100
Timor-Leste	4 digit	TL10
Trinidad and Tobago	2 digit	10
Tunisia	4 digit	1000
Turkey	5 digit	01010
Turkmenistan	3 digit	744
Turks and Caicos Islands	Alphanumeric - 5 digit	TKCA 1
U.S. Virgin Islands	5 digit	00802
Ukraine	3 digit, 5 digit	070, 01001
United Kingdom	Alphanumeric - 2 to 5 digits	B1, AL1, AB10, AB10 1
United States	5 digit	00001
Uruguay	5 digit	11000
Uzbekistan	4 digit	1000

Country	Postal format	Example
Venezuela	4 digit	0000
Vietnam	5 digit	01106
Wallis and Futuna	5 digit	98600
Zambia	5 digit	10100

Using unsupported or custom dates

Amazon Quick Sight natively supports a limited number of date formats. However, you can't always control the format of the data provided to you. When your data contains a date in an unsupported format, you can tell Amazon Quick Sight how to interpret it.

You can do this by editing the dataset, and changing the format of the column from text or numeric to date. A screen appears after you make this change, so you can enter the format. For example, if you are using a relational data source, you can specify MM-dd-yyyy for a text field containing '09-19-2017', so it is interpreted as 2017-09-19T00:00:00.000Z. If you are using a nonrelational data source, you can do the same thing starting with a numeric field or a text field.

Amazon Quick Sight only supports text to date for relational (SQL) sources.

For more information on supported date formats, see [Supported date formats](#).

Use this procedure to help Amazon Quick Sight understand dates in different formats.

1. For a dataset containing unsupported date formats, edit the data as follows. For the column containing your datetime data, change the data type from text to date. Do this by choosing the colorful data type icon beneath the column name in the data preview.

Note

Integer dates that aren't Unix epoch datetimes don't work as is. For example, these formats are not supported as integers: MMddyy, MMddyyyy, ddMMyy, ddMMyyyy, and

yyMMdd. The workaround is to first change them to text format. Make sure all your rows contain six digits (not five). Then change the text data type to datetime. For more information on Unix epoch datetimes, see [epochDate](#).

When you change the data type to date, the **Edit date format** screen appears.

2. Enter your date format, indicating which parts are month, date, year, or time. Formats are case-sensitive.
3. Choose **Validate** to make sure Amazon Quick Sight can now interpret your datetime data with the format you specified. Rows that don't validate are skipped and omitted from the dataset.
4. When you are satisfied with the results, choose **Update**. Otherwise, choose **Close**.

Adding a unique key to an Amazon Quick Sight dataset

Quick Suite authors can configure a unique key column to a Quick Sight dataset during data preparation. This unique key acts as a global sorting key for the dataset and optimizes query generation for table visuals. When a user creates a table visual in Quick Sight and adds the unique key column to the value field well, data is sorted from left to right up to the unique key column. All columns to the right of the unique key column are ignored in the sort order. Tables that don't contain a unique key are sorted based on the order that the columns appear in the dataset.

The following limitations apply to unique keys:

- Unique keys are only supported for unaggregated tables.
- If a dataset column is used for column level security (CLS), the column can't also be used as the unique key.

Use the following procedure to designate a unique key for a dataset in Amazon Quick Sight.

To set up a unique key

1. Open the [Quick Suite console](#).
2. Choose **Data**.
3. Perform one of the following actions:

- a. Navigate to the dataset that you want to add a unique key to, choose the ellipsis (three dots) next to the dataset, and then choose **Edit**.
 - b. Choose **New** then **Dataset**. Choose the dataset that you want to add and then choose **Edit data source**. For more information about creating new datasets in Amazon Quick Sight, see [Creating datasets](#).
4. The data preparation page for the dataset opens. Navigate to the **Fields** pane and locate the field that you want to set as the unique key.
 5. Choose the ellipsis (three dots) next to the field name, and then choose **Set as unique key**.

After you create a unique key, a key icon appears next to the field to show that the field is now the unique key for the dataset. When you save and publish the dataset, the unique key configuration is applied to the dataset and to all dashboards and analyses that are created with that dataset. To remove a unique key from a dataset, navigate to the data preparation page for the dataset, choose the ellipsis next to the unique key field, and then choose **Remove as unique key**. After you remove a unique key from a dataset, you can designate a different field as the unique key.

Integrating Amazon SageMaker AI models with Amazon Quick Sight

Note

You don't need any technical experience in machine learning (ML) to author analyses and dashboards that use the ML-powered features in Amazon Quick Sight.

You can augment your Amazon Quick Suite Enterprise edition data with Amazon SageMaker AI machine learning models. You can run inferences on data stored in SPICE imported from any data source supported by Quick Suite. For a full list of supported data sources, see [Supported data sources](#).

Using Quick Suite with SageMaker AI models can save the time that you might otherwise spend managing data movement and writing code. The results are useful both for evaluating the model and—when you're satisfied with the results—for sharing with decision-makers. You can begin immediately after the model is built. Doing this surfaces your data scientists' prebuilt models, and enables you to apply the data science to your datasets. Then you can share these insights in your predictive dashboards. With the Quick Suite serverless approach, the process scales seamlessly, so you don't need to worry about inference or query capacity.

Amazon Quick Suite supports SageMaker AI models that use regression and classification algorithms. You can apply this feature to get predictions for just about any business use case. Some examples include predicting the likelihood of customer churn, employee attrition, scoring sales leads, and assessing credit risks. To use Quick Suite to provide predictions, the SageMaker AI model data for both input and output must be in tabular format. In multiclass or multilabel classification use cases, each output column has to contain a single value. Quick Suite doesn't support multiple values inside a single column.

Topics

- [How SageMaker AI integration works](#)
- [Costs incurred \(no additional costs with integration itself\)](#)
- [Usage guidelines](#)
- [Defining the schema file](#)
- [Adding a SageMaker AI model to your Quick Sight dataset](#)
- [Build predictive models with SageMaker AI Canvas](#)

How SageMaker AI integration works

In general, the process works like this:

1. An Amazon Quick Suite administrator adds permissions for Quick Suite to access SageMaker AI. To do this, open **Security & Permissions** settings from the **Manage Quick Suite** page. Go to **Quick Suite access to Amazon services**, and add SageMaker AI.

When you add these permissions, Quick Suite is added to an Amazon Identity and Access Management (IAM) role that provides access to list all the SageMaker AI models in your Amazon account. It also provides permissions to run SageMaker AI jobs that have names that are prefixed with quicksight-auto-generated-.

2. We recommend that you connect to an SageMaker AI model that has an inference pipeline, because it automatically performs data preprocessing. For more information, see [Deploy an Inference Pipeline](#) in the *SageMaker AI Developer Guide*.
3. After you identify the data and the pretrained model that you want to use together, the owner of the model creates and provides a schema file. This JSON file is a contract with SageMaker AI. It provides metadata about the fields, data types, column order, output, and settings that the model expects. The optional settings component provides the instance size and count of the compute instances to use for the job.

If you're the data scientist who built the model, create this schema file using the format documented following. If you're a consumer of the model, get the schema file from the owner of the model.

4. In Quick Suite, you begin by creating a new dataset with the data that you want to make predictions on. If you're uploading a file, you can add the SageMaker AI model on the upload settings screen. Otherwise, add the model on the data preparation page.

Before you proceed, verify the mappings between the dataset and the model.

5. After the data is imported into the dataset, the output fields contain the data returned from SageMaker AI. You use these fields just as you use other fields, within the guidelines described in [Usage guidelines](#).

When you run SageMaker AI integration, Quick Suite passes a request to SageMaker AI to run batch transform jobs with inference pipelines. Quick Suite starts provisions and deployment of the instances needed in your Amazon account. When processing is complete, these instances are shut down and terminated. The compute capacity incurs costs only when it's processing models.

To make it easier for you to identify them, Quick Suite names all its SageMaker AI jobs with the prefix `quicksight-auto-generated-`.

6. The output of the inference is stored in SPICE and appended to the dataset. As soon as the inference is complete, you can use the dataset to create visualizations and dashboards using the prediction data.
7. The data refresh starts every time you save the dataset. You can start the data refresh process manually by refreshing the SPICE dataset, or you can schedule it to run at a regular interval. During each data refresh, the system automatically calls SageMaker AI batch transform to update the output fields with new data.

You can use the Amazon Quick Sight SPICE ingestion API operations to control the data refresh process. For more information about using these API operations, see the [Amazon Quick Sight API Reference](#).

Costs incurred (no additional costs with integration itself)

Using this feature doesn't require an additional fee in itself. Your costs include the following:

- The cost of model deployment through SageMaker AI, which is incurred only when the model is running. Saving a dataset—after either creating or editing it—or refreshing its data starts the

data ingestion process. This process includes calling SageMaker AI if the dataset has inferred fields. Costs are incurred in the same Amazon account where your Quick Suite subscription is.

- Your Quick Suite subscription costs are as follows:
 - The cost of storing your data in the in-memory calculation engine in Quick Suite (SPICE). If you are adding new data to SPICE, you might need to purchase enough SPICE capacity to accommodate it.
 - Quick Suite subscriptions for the authors or admins who build the datasets.
 - Pay-per-session charges for viewers (readers) to access interactive dashboards.

Usage guidelines

In Amazon Quick Suite, the following usage guidelines apply to this Enterprise edition feature:

- The processing of the model occurs in SPICE. Therefore, it can only apply to datasets that are stored in SPICE. The process currently supports up to 500 million rows per dataset.
- Only Quick Suite admins or authors can augment datasets with ML models. Readers can only view the results when they are part of a dashboard.
- Each dataset can work with one and only one ML model.
- Output fields can't be used to calculate new fields.
- Datasets can't be filtered by fields that are integrated with the model. In other words, if your dataset field is currently mapped to the ML model, you can't filter on that field.

In SageMaker AI, the following usage guidelines apply to a pretrained model that you use with Amazon Quick Sight:

- When you create the model, associate it with the Amazon Resource Name (ARN) for the appropriate IAM role. The IAM role for the SageMaker AI model needs to have access to the Amazon S3 bucket that Amazon Quick Sight uses.
- Make sure that your model supports .csv files for both input and output. Make sure that your data is in a tabular format.
- Provide a schema file that contains metadata about the model, including the list of input and output fields. Currently, you must create this schema file manually.
- Consider the amount of time that it takes to complete your inference, which depends on a number of factors. These include the complexity of the model, the amount of data, and the

compute capacity defined. Completing the inference can take several minutes to several hours. Amazon Quick Sight caps all data ingestion and inferencing jobs to a maximum of 10 hours. To reduce the time it takes to perform an inference, consider increasing the instance size or the number of instances.

- Currently, you can use only batch transforms for integration with SageMaker AI, not real-time data. You can't use an SageMaker AI endpoint.

Defining the schema file

Before you use an SageMaker AI model with Quick Sight data, create the JSON schema file that contains the metadata that Amazon Quick Sight needs to process the model. The Amazon Quick Suite author or admin uploads the schema file when configuring the dataset.

The schema fields are defined as follows. All fields are required unless specified in the following description. Attributes are case-sensitive.

inputContentType

The content type that this SageMaker AI model expects for the input data. The only supported value for this is "text/csv". Quick Sight doesn't include any of the header names that you add to the input file.

outputContentType

The content type of the output that is produced by the SageMaker AI model that you want to use. The only supported value for this is "text/csv".

input

A list of features that the model expects in the input data. Quick Sight produces the input data in exactly the same order. This list contains the following attributes:

- *name* – The name of the column. If possible, make this the same as the name of the corresponding column in the QuickSight dataset. This attribute is limited to 100 characters.
- *type* – The data type of this column. This attribute takes the values "INTEGER", "STRING", and "DECIMAL".
- *nullable* – (Optional) The nullability of the field. The default value is `true`. If you set `nullable` to `false`, Quick Sight drops rows that don't contain this value before calling SageMaker AI. Doing this helps avoid causing SageMaker AI to fail on missing required data.

output

A list of output columns that the SageMaker AI model produces. Quick Sight expects these fields in exactly the same order. This list contains the following attributes:

- *name* – This name becomes the default name for the corresponding new column that's created in Quick Sight. You can override the name specified here in Quick Sight. This attribute is limited to 100 characters.
- *type* – The data type of this column. This attribute takes the values "INTEGER", "STRING", and "DECIMAL".

instanceTypes

A list of the ML instance types that SageMaker AI can provision to run the transform job. The list is provided to the Amazon Quick Suite user to choose from. This list is limited to the types supported by SageMaker AI. For more information on supported types, see [TransformResources](#) in the *SageMaker AI Developer Guide*.

defaultInstanceType

(Optional) The instance type that is presented as the default option in the SageMaker AI wizard in Quick Sight. Include this instance type in *instanceTypes*.

instanceCount

(Optional) The instance count defines how many of the selected instances for SageMaker AI to provision to run the transform job. This value must be a positive integer.

description

This field provides a place for the person who owns the SageMaker AI model to communicate with the person who is using this model in Quick Sight. Use this field to provide hints about successfully using this model. For example, this field can contain information about selecting an effective instance type to choose from the list in *instanceTypes*, based on the size of dataset. This field is limited to 1,000 characters.

version

The version of the schema, for example "1.0".

The following example shows the structure of the JSON in the schema file.

```
{  
    "inputContentType": "CSV",
```

```

    "outputContentType": "CSV",
    "input": [
      {
        "name": "buying",
        "type": "STRING"
      },
      {
        "name": "maint",
        "type": "STRING"
      },
      {
        "name": "doors",
        "type": "INTEGER"
      },
      {
        "name": "persons",
        "type": "INTEGER"
      },
      {
        "name": "lug_boot",
        "type": "STRING"
      },
      {
        "name": "safety",
        "type": "STRING"
      }
    ],
    "output": [
      {
        "name": "Acceptability",
        "type": "STRING"
      }
    ],
    "description": "Use ml.m4.xlarge instance for small datasets, and ml.m4.4xlarge
for datasets over 10 GB",
    "version": "1.0",
    "instanceCount": 1,
    "instanceTypes": [
      "ml.m4.xlarge",
      "ml.m4.4xlarge"
    ],
    "defaultInstanceType": "ml.m4.xlarge"
  }

```

The structure of the schema file is related to the kind of model that is used in examples provided by SageMaker AI.

Adding a SageMaker AI model to your Quick Sight dataset

Using the following procedure, you can add a pretrained SageMaker AI model to your dataset, so that you can use predictive data in analyses and dashboards.

Before you begin, have the following items available:

- The data that you want to use to build the dataset.
- The name of the SageMaker AI model that you want to use to augment the dataset.
- The schema of the model. This schema includes field name mappings and data types. It's helpful if it also contains recommended settings for instance type and number of instances to use.

To augment your Amazon Quick Sight dataset with SageMaker AI

1. Create a new dataset from the start page by choosing **Datasets**, and then choose **New dataset**.

You can also edit an existing dataset.

2. Choose **Augment with SageMaker** on the data preparation screen.
3. For **Select your model**, choose the following settings:
 - **Model** – Choose the SageMaker AI model to use to infer fields.
 - **Name** – Provide a descriptive name for the model.
 - **Schema** – Upload the JSON schema file provided for the model.
 - **Advanced settings** – QuickSight recommends the selected defaults based on your dataset. You can use specific runtime settings to balance the speed and cost of your job. To do this, enter the SageMaker AI ML instance types for **Instance type** and number of instances for **Count**.

Choose **Next** to continue.

4. For **Review inputs**, review the fields that are mapped to your dataset. Quick Sight attempts to automatically map the fields in your schema to the fields in your dataset. You can make changes here if the mapping needs adjustment.

Choose **Next** to continue.

5. For **Review outputs**, view the fields that are added to your dataset.

Choose **Save and prepare data** to confirm your choices.

6. To refresh the data, choose the dataset to view details. Then either choose **Refresh Now** to manually refresh the data, or choose **Schedule refresh** to set up a regular refresh interval. During each data refresh, the system automatically runs the SageMaker AI batch transform job to update the output fields with new data.

Build predictive models with SageMaker AI Canvas

Amazon Quick Suite authors can export data into SageMaker AI Canvas to build ML models that can be sent back to Quick Suite. Authors can use these ML models to augment their datasets with predictive analytics that can be used to build analyses and dashboards.

Prerequisites

- A Quick Suite account that's integrated with IAM Identity Center. If your Quick Suite account isn't integrated with IAM Identity Center, create a new Quick Suite account and choose **Use IAM Identity Center enabled application** as the identity provider.
 - For more information on IAM Identity Center, see [Getting started](#).
 - To learn more about integrating your Quick Suite with IAM Identity Center, see [Configure your Amazon Quick Suite account with IAM Identity Center](#).
 - To import assets from an existing Quick Suite account to a new Quick Suite account that's integrated with IAM Identity Center, see [Asset bundle operations](#).
- A new SageMaker AI domain that is integrated with IAM Identity Center. For more information about onboarding to SageMaker AI Domain with IAM Identity Center, see [Onboard to SageMaker AI Domain using IAM Identity Center](#).

Topics

- [Build a predictive model in SageMaker AI Canvas from Amazon Quick Sight](#)
- [Create a dataset with a SageMaker AI Canvas model](#)
- [Considerations](#)

Build a predictive model in SageMaker AI Canvas from Amazon Quick Sight

To build a predictive model in SageMaker AI Canvas

1. Log in to Amazon Quick Suite and navigate to the tabular table or pivot table that you want to create a predictive model for.
2. Open the on-visual menu and choose **Build a predictive model**.
3. In the **Build a predictive model in SageMaker AI Canvas** pop up that appears, review the information presented and then choose **EXPORT DATA TO SAGEMAKER CANVAS**.
4. In the **Exports** pane that appears, choose **GO TO SAGEMAKER CANVAS** when the export is completed to go to the SageMaker AI Canvas console.
5. In SageMaker AI Canvas, create a predictive model with the data that you exported from Quick Sight. You can choose to follow a guided tour that helps you create the predictive model, or you can skip the tour and work at your own pace. For more information about creating a predictive model in SageMaker AI Canvas, see [Build a model](#).
6. Send the predictive model back to Quick Sight. For more information about sending a model from SageMaker AI Canvas to Amazon Quick Sight, see [Send your model to Amazon Quick Sight](#).

Create a dataset with a SageMaker AI Canvas model

After you create a predictive model in SageMaker AI Canvas and send it back to Quick Sight, use the new model to create a new dataset or apply it to an existing dataset.

To add a predictive field to a dataset

1. Open the Quick Suite console, choose **Data** at left, and choose the **Datasets** tab.
2. Upload a new dataset or choose an existing dataset.
3. Choose **Edit**.
4. On the dataset's data prep page, choose **ADD**, and then choose **Add predictive field** to open the **Augment with SageMaker AI** modal.
5. For **Model**, choose the model that you sent to Quick Sight from SageMaker AI Canvas. The schema file automatically populates in the **Advanced settings** pane. Review the inputs, and then choose **Next**.
6. On the **Review outputs** pane, enter a field name and description for a column to be targeted by the model that you created in SageMaker AI Canvas.

7. When you are finished, choose **Prepare data**.
8. After you choose **Prepare data**, you are redirected to the dataset page. To publish the new dataset, choose, **Publish & Visualize**.

When you publish a new dataset that uses a model from SageMaker AI Canvas, the data is imported into SPICE and a batch inference job begins in SageMaker AI. It can take up to 10 minutes for these processes to complete.

Considerations

The following limitations apply to the creation of SageMaker AI Canvas models with Quick Sight data.

- The **Build a predictive model** option that is used to send data to SageMaker AI Canvas is only available on table and tabular pivot table visuals. The table or pivot table visual must have between 2 and 1,000 fields and at least 500 rows.
- Datasets that contain integer or geographic data types will experience schema mapping errors when you add a predictive field to the dataset. To resolve this issue, remove the integer or geographic data types from the dataset or convert them to a new data type.

Preparing dataset examples

You can prepare data in any dataset to make it more suitable for analysis, for example changing a field name or adding a calculated field. For database datasets, you can also determine the data used by specifying a SQL query or joining two or more tables.

Use the following topics to learn how to prepare datasets.

Topics

- [Preparing a dataset based on file data](#)
- [Preparing a dataset based on Salesforce data](#)
- [Preparing a dataset based on database data](#)

Preparing a dataset based on file data

Use the following procedure to prepare a dataset based on text or Microsoft Excel files from either your local network or Amazon S3.

To prepare a dataset based on text or Microsoft Excel files from a local network or S3

1. Open a file dataset for data preparation by choosing one of the following options:
 - Create a new local file dataset, and then choose **Edit/Preview data**. For more information about creating a new dataset from a local text file, see [Creating a dataset using a local text file](#). For more information about creating a new dataset from a Microsoft Excel file, see [Creating a dataset using a Microsoft Excel file](#).
 - Create a new Amazon S3 dataset, and then choose **Edit/Preview data**. For more information about creating a new Amazon S3 dataset using a new Amazon S3 data source, see [Creating a dataset using Amazon S3 files](#). For more information about creating a new Amazon S3 dataset using an existing Amazon S3 data source, see [Creating a dataset using an existing Amazon S3 data source](#).
 - Open an existing Amazon S3, text file, or Microsoft Excel dataset for editing, from either the analysis page or the **Your Datasets** page. For more information about opening an existing dataset for data preparation, see [Editing datasets](#).
2. (Optional) On the data preparation page, enter a new name into the dataset name box on the application bar.

This name defaults to the file name for local files. For example, it defaults to **Group 1** for Amazon S3 files.

3. Review the file upload settings and correct them if necessary. For more information about file upload settings, see [Choosing file upload settings](#).

Important

If you want to change upload settings, make this change before you make any other changes to the dataset. New upload settings cause Amazon Quick Sight to reimport the file. This process overwrites all of your other changes.

4. Prepare the data by doing one or more of the following:
 - [Selecting fields](#)
 - [Editing field names and descriptions](#)
 - [Changing a field data type](#)
 - [Adding calculated fields](#)
 - [Filtering data in Amazon Quick Sight](#)

5. Check the [SPICE](#) indicator to see if you have enough capacity to import the dataset. File datasets automatically load into SPICE. The import happens when you choose either **Save & visualize** or **Save**.

If you don't have access to enough SPICE capacity, you can make the dataset smaller by using one of the following options:

- Apply a filter to limit the number of rows.
- Select fields to remove from the dataset.

 Note

The SPICE indicator doesn't update to how much space you save by removing fields or filtering the data. It continues to reflect the SPICE usage from the last import.

6. Choose **Save** to save your work, or **Cancel** to cancel it.

You might also see **Save & visualize**. This option appears based on the screen that you started from. If this option isn't there, you can create a new visualization by starting from the dataset screen.

Preparing a dataset based on a Microsoft Excel file

Use the following procedure to prepare a Microsoft Excel dataset.

To prepare a Microsoft Excel dataset

1. Open a text file dataset for preparation by choosing one of the following options:
 - Create a new Microsoft Excel dataset, and then choose **Edit/Preview data**. For more information about creating a new Excel dataset, see [Creating a dataset using a Microsoft Excel file](#).
 - Open an existing Excel dataset for editing. You can do this from the analysis page or the **Your Datasets** page. For more information about opening an existing dataset for data preparation, see [Editing datasets](#).
2. (Optional) On the data preparation page, enter a name into the dataset name box in the application bar. If you don't rename the dataset, its name defaults to the Excel file name.

3. Review the file upload settings and correct them if necessary. For more information about file upload settings, see [Choosing file upload settings](#).

 Important

If it's necessary to change upload settings, make this change before you make any other changes to the dataset. Changing upload settings causes Amazon Quick Sight to reimport the file. This process overwrites any changes you have made so far.

4. (Optional) Change the worksheet selection.
5. (Optional) Change the range selection. To do this, open **Upload Settings** from the on-dataset menu beneath the login name at upper right.
6. Prepare the data by doing one or more of the following:
 - [Selecting fields](#)
 - [Editing field names and descriptions](#)
 - [Changing a field data type](#)
 - [Adding calculated fields](#)
 - [Filtering data in Quick Sight](#)
7. Check the [SPICE](#) indicator to see if you have enough space to import the dataset. Amazon Quick Sight must import Excel datasets into SPICE. This import happens when you choose either **Save & visualize** or **Save**.

If you don't have enough SPICE capacity, you can choose to make the dataset smaller using one of the following methods:

- Apply a filter to limit the number of rows.
- Select fields to remove from the dataset.
- Define a smaller range of data to import.

 Note

The SPICE indicator doesn't update to reflect your changes until after your load them. It shows the SPICE usage from the last import.

8. Choose **Save** to save your work, or **Cancel** to cancel it.

You might also see **Save & visualize**. This option appears based on the screen that you started from. If this option isn't there, you can create a new visualization by starting from the dataset screen.

Preparing a dataset based on Salesforce data

Use the following procedure to prepare a Salesforce dataset.

To prepare a Salesforce dataset

1. Open a Salesforce dataset for preparation by choosing one of the following options:
 - Create a new Salesforce dataset and choose **Edit/Preview data**. For more information about creating a new Salesforce dataset using a new Salesforce data source, see [Creating a dataset from Salesforce](#). For more information about creating a new Salesforce dataset using an existing Salesforce data source, see [Create a dataset using an existing Salesforce data source](#).
 - Open an existing Salesforce dataset for editing from either the analysis page or the **Your Datasets** page. For more information about opening an existing dataset for data preparation, see [Editing datasets](#).
2. (Optional) On the data preparation page, enter a name into the dataset name box in the application bar if you want to change the dataset name. This name defaults to the report or object name.
3. (Optional) Change the data element selection to see either reports or objects.
4. (Optional) Change the data selection to choose a different report or object.

If you have a long list in the **Data** pane, you can search to locate a specific item by entering a search term into the **Search tables** box. Any item whose name contains the search term is shown. Search is case-insensitive and wildcards are not supported. Choose the cancel icon (X) to the right of the search box to return to viewing all items.

5. Prepare the data by doing one or more of the following:
 - [Selecting fields](#)
 - [Editing field names and descriptions](#)
 - [Changing a field data type](#)
 - [Adding calculated fields](#)
 - [Filtering data in Quick Sight](#)

6. Check the [SPICE](#) indicator to see if you have enough space to import the dataset. Importing data into SPICE is required for Salesforce datasets. Importing occurs when you choose either **Save & visualize** or **Save**.

If you don't have enough SPICE capacity, you can remove fields from the dataset or apply a filter to decrease its size. For more information about adding and removing fields from a dataset, see [Selecting fields](#).

 Note

The SPICE indicator doesn't update to reflect the potential savings of removing fields or filtering the data. It continues to reflect the size of the dataset as retrieved from the data source.

7. Choose **Save** to save your work, or **Cancel** to cancel it.

You might also see **Save & visualize**. This option appears based on the screen you started from. If this option isn't there, you can create a new visualization by starting from the dataset screen.

Preparing a dataset based on database data

Use the following procedure to prepare a dataset based on a query to a database. The data for this dataset can be from an Amazon database data source like Amazon Athena, Amazon RDS, or Amazon Redshift, or from an external database instance. You can choose whether to import a copy of your data into [SPICE](#), or to query the data directly.

To prepare a dataset based on a query to a database

1. Open a database dataset for preparation by choosing one of the following options:
 - Create a new database dataset and choose **Edit/Preview data**. For more information about creating a new dataset using a new database data source, see [Creating a dataset from a database](#). For more information about creating a new dataset using an existing database data source, see [Creating a dataset using an existing database data source](#).
 - Open an existing database dataset for editing from either the analysis page or the **Your Datasets** page. For more information about opening an existing dataset for data preparation, see [Editing datasets](#).

2. (Optional) On the data preparation page, enter a name into the dataset name box on the application bar.

This name defaults to the table name if you selected one before data preparation. Otherwise, it's **Untitled data source**.

3. Decide how your data is selected by choosing one of the following:

- To use a single table to provide data, choose a table or change the table selection.

If you have a long table list in the **Tables** pane, you can search for a specific table by typing a search term for **Search tables**.

Any table whose name contains the search term is shown. Search is case-insensitive and wildcards are not supported. Choose the cancel icon (X) to the right of the search box to return to viewing all tables.

- To use two or more joined tables to provide data, choose two tables and join them using the join pane. You must import data into Quick Sight if you choose to use joined tables. For more information about joining data using the Amazon Quick Sight interface, see [Joining data](#).
- To use a custom SQL query to provide data in a new dataset, choose **Switch to Custom SQL** tool on the **Tables** pane. For more information, see [Using SQL to customize data](#).

To change the SQL query in an existing dataset, choose **Edit SQL** on the **Fields** pane to open the SQL pane and edit the query.

4. Prepare the data by doing one or more of the following:

- [Selecting fields](#)
- [Editing field names and descriptions](#)
- [Changing a field data type](#)
- [Adding calculated fields](#)
- [Filtering data in Quick Sight](#)

5. If you aren't joining tables, choose whether to query the database directly or to import the data into SPICE by selecting either the **Query** or **SPICE** radio button. We recommend using SPICE for enhanced performance.

If you want to use SPICE, check the SPICE indicator to see if you have enough space to import the dataset. Importing occurs when you choose either **Save & visualize** or **Save**.

If you don't have enough space, you can remove fields from the dataset or apply a filter to decrease its size.

 Note

The SPICE indicator doesn't update to reflect the potential savings of removing fields or filtering the data. It continues to reflect the size of the dataset as retrieved from the data source.

6. Choose **Save** to save your work, or **Cancel** to cancel it.

You might also see an option to **Save & visualize**. This option appears based on the screen you started from. If this option isn't there, you can create a new visualization by starting from the dataset screen.

Analyses and reports: Visualizing data in Amazon Quick Sight

Following, you can find descriptions of how to create and customize Amazon Quick Sight charts, arrange charts in a dashboard, and more.

Topics

- [Working with an analysis in Amazon Quick Sight](#)
- [Adding and managing sheets](#)
- [Working with interactive sheets in Amazon Quick Sight](#)
- [Working with paginated reports in Amazon Quick Sight](#)
- [Working with items on sheets in Amazon Quick Sight analyses](#)
- [Using themes in Amazon Quick Sight](#)
- [Accessing Amazon Quick Sight using keyboard shortcuts](#)

Working with an analysis in Amazon Quick Sight

In Quick Sight, an analysis is the same thing as a dashboard, except that it can only be accessed by the authors you choose. You can keep it private, and make it as robust and detailed as you like. When and if you decide to publish it, the shared version of it is called a dashboard.

Use the following sections to learn how to interact with a Quick Sight analysis.

Topics

- [Starting an analysis in Quick Sight](#)
- [Adding a title and description to an analysis](#)
- [Sharing Quick Sight analyses](#)
- [Renaming an analysis](#)
- [Duplicating analyses](#)
- [Customize date and time values of an analysis](#)
- [The analysis menu](#)
- [Configure analysis settings](#)
- [Item limits for Amazon Quick Sight analyses in the Quick Sight APIs](#)
- [Saving changes to analyses](#)
- [Exporting data from Quick Sight analyses](#)
- [Deleting an analysis](#)

Starting an analysis in Quick Sight

In Quick Sight, you analyze and visualize your data in analyses. When you're finished, you can publish your analysis as a dashboard to share with others in your organization.

Use the following procedure to create a new analysis.

To create a new analysis

1. On the Amazon Quick Suite start page, choose **Analyses**, and then choose **New analysis**.
2. Choose the dataset that you want to include in your new analysis, and then choose **USE IN ANALYSIS** in the top right.
3. In the **New sheet** pop-up that appears, choose the sheet type that you want. You can choose between an **Interactive sheet** and a **Paginated report**. To create a paginated report, you need the paginated reports add-on for your account. For more information about paginated reports, see [Working with paginated reports in Amazon Quick Sight](#). For more information on sheets, see [Adding and managing sheets](#).
4. (Optional) If you choose **Interactive sheet**, follow these steps:

- (Optional): Choose the type layout that you want for your interactive sheet. You can choose one of the following options:
 - Free-form
 - Tiled

The default option is **Free-form**.

For more information about interactive sheet layouts, see [Types of layout](#).

- Choose the canvas size that you want your sheet optimized for. You can choose one of the following options:
 - 1024px
 - 1280px
 - 1366px
 - 1600px
 - 1920px

For more information on formatting interactive sheets, see [Working with interactive sheets in Amazon Quick Sight](#).

5. (Optional) If you choose **Paginated report**, follow these steps:

- (Optional) Choose the paper size that you want for your paginated report. You can choose from the following options:
 - US letter (8.5 x 11 in.)
 - US legal (8.5 x 14 in.)
 - A0 (841 x 1189 mm)
 - A1 (594 x 841 mm)
 - A2 (420 x 594 mm)
 - A3 (297 x 420 mm)
 - A4 (210 x 297 mm)
 - A5 (148 x 210 mm)
 - Japan B4 (257 x 364 mm)
 - Japan B5 (182 x 257 mm)

- (Optional) Choose the orientation of the sheet. You can choose between **Portrait** or **Landscape**. The default option is portrait.

Before you can create Amazon Quick Sight paginated reports, first get the paginated reporting add-on for your Quick Suite account. For more information on getting the paginated reporting add-on, see [Get the Quick Suite paginated reports add-on](#).

For more information on formatting paginated reports, see [Working with paginated reports in Amazon Quick Sight](#).

6. Choose **Add**.
7. Create a visual. For more information about creating visuals, see [Adding visuals to Quick Sight analyses](#).

After you are done creating the analysis, you can iterate on it by modifying the visual, adding more visuals, adding scenes to the default story, or adding more stories.

Adding a title and description to an analysis

In addition to the analysis name, you can add a title and description to an analysis. A useful title and description provides context about the information in the analysis.

Add a title and description

Use the following procedure to add a title and description to an analysis. Titles and descriptions can contain up to 1,024 characters. Titles and descriptions are not supported for paginated reports.

To add a title and description to an analysis

1. On the analysis page, choose **Sheets** in the application bar and then choose **Add title**.
2. For **Sheet title**, enter a title and press **Enter**. To remove a title, choose **Sheets** in the application bar and then choose **Delete title**. Or, to remove the title, you can select the title and then choose the x-shaped delete icon.

To create a dynamic sheet title, you can add existing parameters to the sheet title. For more information, see [Using parameters in titles and descriptions in Amazon Quick Suite](#).

3. Choose **Sheets** in the application bar, and then choose **Add description**.
4. In the description space that appears on the sheet, enter the description that you want and press **Enter**. To remove a description, choose **Sheets** in the application bar and then choose

Delete description. Or, to remove the description, you can select the description and then choose the x-shaped delete icon.

Sharing Quick Sight analyses

You can share an analysis with one or more other users by emailing them a link, making it easy to collaborate and disseminate findings. You can only share an analysis with other users in your Quick Suite account.

After you share an analysis, you can review the other users who have access to it, and also revoke access from any user.

Topics

- [Sharing an analysis](#)
- [Viewing the users that an analysis is shared with](#)
- [Revoking access to an analysis](#)

Sharing an analysis

Use the following procedure to share an analysis.

To share an analysis

1. Open the [Quick Suite console](#).
2. Open the analysis that you want to change.
3. On the analysis page, choose **File** on the application bar, and then choose **Share**.

You can only share analyses with users or groups who are in your Quick Suite account.

4. Add a user or group to share with. To do this, for **Type a user name or email**, enter the first user or group that you want to share this analysis with. Then choose **Share**. Repeat this step until you have entered information for everyone you want to share the analysis with.

To edit sharing permissions for this analysis, choose **Manage analysis permissions**.

The **Manage analysis permissions** screen appears. On this screen, choose **Invite user** to edit permissions and add more users or groups.

5. For **Permission**, choose the role to assign to each user or group. The role determines the permission level to grant to that user or group.

6. Choose **Share**.

The users that you have shared the analysis with get emails with a link to the analysis. Groups don't receive invitation emails.

Viewing the users that an analysis is shared with

If you have shared an analysis, you can use the following procedure to see which users or groups have access to it.

To view which users or groups have access to an analysis

1. Open the [Quick Suite console](#).
2. Open the analysis that you want to change.
3. On the analysis page, choose **File** on the application bar, and then choose **Share**.
4. Choose **Manage analysis permissions**.
5. Review who this analysis has been shared with. You can search to locate a specific account by typing a search term. The search returns any user, group, or email address that contains the search term. Searching is case-sensitive, and wildcards are not supported. Delete the search term to view all users and groups.

Revoking access to an analysis

Use the following procedure to revoke access to an analysis.

To revoke access to an analysis

1. Open the [Quick Suite console](#).
2. Open the analysis that you want to change.
3. On the analysis page, choose **File** on the application bar, and then choose **Share**.
4. Choose **Manage analysis permissions**.
5. Locate the user or group whose access you want to revoke, and then choose the trash-can icon next to the user or group.
6. Choose **Confirm**.

Renaming an analysis

Use the following procedure to rename an analysis.

To rename an analysis

1. Open the analysis that you want to rename.
2. In the **Analysis name** box in the application bar, select the current name and then enter a new name.

Duplicating analyses

You can duplicate analyses in Quick Sight. Use the following procedure to learn how.

To duplicate an analysis

1. From the Quick Suite homepage, choose **Analyses**, and then open the analysis that you want to duplicate.
2. In the analysis, choose **Save as** in the application bar at upper right.
3. In the **Save a copy** page that opens, enter a name for the analysis, and then choose **Save**.

The new analysis opens. You can find the original analysis by returning to the Quick Suite start page and selecting **Analyses**.

Customize date and time values of an analysis

In Amazon Quick Suite, authors can set custom time zones and week start days of an analysis. When you set a custom week start or time zone, all visuals in the analysis that use datetime data are formatted to reflect the time zone or week start that the analysis uses.

Setting custom time zones in an analysis

Quick Suite authors can use *custom time zones* to help manage data across multiple geographic regions. When you set a custom time zone, all visible dimensions, measures, calculated fields, and filters are converted to the chosen time zone at query run time. Daylight Savings Time (DST) adjustments are applied automatically to eliminate the need for time consuming workarounds that don't accurately handle historical dates.

Custom time zones refer to the use of IANA time zone abbreviations that represent specific geographic regions around the world. Each time zone is defined as an offset from Coordinated Universal Time (UTC). Time zones are different from simple offsets because they incorporate DST.

The default time zone for all analyses is UTC.

The following rules apply to time zones.

- **Datetime displays with a granularity that is lower than hour are converted to the selected time zone.** For example, if you set the timezone of an analysis to `America/New_York` (UTC-04:00), the datetime value `Dec .1, 2020 12:00am` in UTC+00:00 is converted and displayed as `Nov .30, 2020 7:00pm`. Daylight Savings Time (DST) is incorporated into the datetime conversion.
- **Datetime literals that are added to calculations or selected in filters honor the selected time zone of the analysis.** For example, if you manually enter a literal into a calculated field such as `01-01-2022 7:00pm`, or select a fixed filter time, Quick Sight applies the chosen timezone to the literal value.
- **Measures that are aggregated above the hour/minute granularity are aggregated based on the timezone that the analysis is set to.** When Quick Sight processes a dataset, all timestamps are initially converted at the lowest granularity level. Values are then aggregated based on the boundary of the selected time zone for the analysis. For example, a sum of hourly revenues at the day level with a UTC+00:00 time zone aggregate all hourly revenues from 12am-11pm for the UTC time zone. When you convert UTC+00:00 to `New_York` (UTC-04:00), all revenue datapoints are aggregated from 8:00pm-7:00pm(+1day) in UTC to correspond with the start and end of the day in `New_York` (UTC-04:00).
- **The `now()` function, rolling date filter, and parameters are converted to the chosen time zone.** Relative date filters, rolling date filters, and relative date parameters that use the `now()` function also honor the chosen time zone when they are applied to the visual. For example, when you select a relative filter such as `last week` or a rolling date filter such as `start of the month`, the chosen timezone is automatically applied to the filter to display the values `last week of New_York time zone` and `start of the month of New_York time zone`, respectively.

To set the custom time zone of an analysis

1. From the analysis that you want to change, navigate to the top menu and choose **Edit**.
2. Choose **Analysis settings**, and then choose **Date and time**.

3. Toggle **Convert time zone** ON, and choose the **Time zone** that you want.
4. Choose **Apply**.

When an analysis is assigned a time zone, an icon appears at the top of the analysis that indicates which time zone the analysis uses. This icon also appears on any dashboard that is published from the analysis.

Considerations

The following considerations apply to custom time zones.

- To use custom time zones, all datetime columns in a dataset must be normalized to UTC. If your datetime columns aren't normalized in your data source, you need to convert the columns in your data source before you can use this feature.
- For analyses that are not assigned a custom time zone, author and reader experiences are unaffected.
- Once a time zone is added to an analysis, the time zone is applied to all visuals and sheets in the analysis.
- Quick Suite authors can choose only one time zone for an analysis. All dashboards that are published from the analysis use the time zone that the analysis uses. To create a dashboard that uses a different time zone than the one that the analysis uses, change the time zone of the analysis and republish the dashboard.
- Quick Suite readers can't change the time zone of a dashboard.
- If you set the time zone of an analysis that uses a dataset that is stored in Direct Query and experience slow load times, consider storing the dataset in SPICE. SPICE is engineered to handle time zone conversions in a performant way.
- Custom time zones do not support the following database engines:
 - Timestream
 - OpenSearch Service
 - Teradata
 - SqlServer

Setting custom week start days in an analysis

Quick Suite authors can define the week start day of an analysis to align their data with the schedule that their company or industry follows. When you set a custom week start day, all dimensions, calculated fields, and filters that are aggregated at the week level are calculated to align with the new week start day. The default week start day is Sunday.

To set the custom week start day of an analysis

1. From the analysis that you want to change, navigate to the top menu and choose **Edit**.
2. Choose **Analysis settings**, and then choose **Date and time**.
3. For **Custom start day**, choose the start day that you want.
4. Choose **Apply**.

Considerations

The following considerations apply to custom week start days.

- Datetime fields are converted at run time. When you work with calculated fields that use datetime values, define the fields at the analysis level instead of the dataset level.
- Once you choose a new week start day, the change is applied to all visuals and sheets in the analysis.
- Quick Suite authors can choose only one week start day for an analysis. All dashboards that are published from the analysis use the week start day that the analysis uses. To create a dashboard that uses a different week start day than the one that the analysis uses, change the week start day of the analysis and republish the dashboard.
- Quick Suite readers can't change the week start day of a dashboard.

The analysis menu

While working on an analysis, Amazon Quick Suite provides several menu options. You use these menu options to efficiently perform tasks without needing to manually navigate through your analysis to find the assets that you want to change.

You can use these options to perform the following tasks.

- *File* – Perform analysis management tasks, including creating, sharing, and publishing. Authors can use this option to make changes across all sheets or visuals in an analysis.

- *Edit* – Navigate between changes that you make to the analysis. You can undo or redo changes that you make.
- *Data* – Manage datasets, data fields, and parameters. Changes that you make by using this option are applied to all sheets in the analysis.
- *Insert* – Use an ingress point where you can add visuals, text boxes, insights, reporting objects, filters, and parameters to an analysis. The content that you insert can be data or objects.
- *Sheets* – Manage the sheet settings of the analysis, including layout settings, actions to add or remove assets from a sheet, and sheet properties.
- *Objects* – Manage objects and their features, including style, canvas placement, sizing, card background, and borders. You can also manage these objects by using the **Properties** pane when working on a visual object.
- *Search* – Access the *Quick search* bar. Quick search is a search bar that will begin showing results for the asset you are searching for as you type. The suggested results continue to modify as you type until you see the result that you're looking for.

To use quick search, open the **Search** menu, and in the **Search analysis actions** box, begin typing a name or phrase associated with the asset you are trying to find.

Configure analysis settings

Amazon Quick Suite authors can use the Analysis settings menu to configure the refresh and date time settings of an analysis. To access the Analysis Settings menu, choose **Edit**, and then choose **Analysis Settings**. The following settings can be configured in the Analysis settings menu:

Refresh settings

- **Reload visuals every time I switch sheets** – Use this setting to reload every visual in a Quick Sight analysis whenever the user switches to a different sheet in the analysis.
- **Update visuals manually** – Use this setting to only update applicable visuals in an analysis when the user applies their changes. When this setting is toggled on, the analysis loads the visuals blank by default because the queries won't be fired until the user selects the **UPDATE VISUALS** button located in the toolbar or on the impacted visuals. The **UPDATE VISUALS** button confirms that the user is finished with the filter and control choices that they want to apply to the affected visuals. The image below shows the **UPDATE VISUALS** button.

When **Update visuals manually** is toggled on, authors can still add visuals, edit visuals, and edit control selections, but the affected visuals won't update until the author applies the new

changes. This allows authors to build analyses without increasing their database load and gives better control over which values are loaded in an analysis.

Date and time settings

- **Convert time zone** – Use this setting to convert all date field related visualizations, filters, and parameters to reflect the chosen time zone. All daylight savings adjustments are made automatically. For more information about time zone configuration, see [Customize date and time values of an analysis](#).
- **Start of the week** – Use this setting to choose the week start day for an analysis.

Interactivity

- Use this setting to highlight specific data points across visuals in a sheet. When you select or hover over a data point on a visual, related data across other visuals will stand out, while unrelated data is dimmed. Highlighting allows you to understand correlations, spot patterns, trends, and outliers, and facilitate stronger, more informed analyses. Select either **On selection** or **On hover** to turn highlighting on, or **No highlight** to turn it off.
- To customize highlighting on a per-sheet level see [Adding and managing sheets](#).

Item limits for Amazon Quick Sight analyses in the Quick Sight APIs

Use the following table to review the current limits or quotas for different analysis items in Amazon Quick Sight that are created and managed with the Amazon Quick Sight APIs. If your analysis contains more than the supported number of analysis items, remove items to optimize the performance of the analysis. New analysis items cannot be added to an analysis that contains more than the supported number of analysis items.

Analysis item	Limit
Sheets	20 sheets per analysis
Visuals	50 visuals per sheet
Calculated fields	500 per analysis and 200 per dataset*
Bookmarks	200 per dashboard

Analysis item	Limit
Custom actions	10 per visual
Filter groups	2000 per analysis
Filters	20 filters per filter group
Parameters	200 per analysis
Controls	200 per sheet
Text boxes	100 per sheet
Image components	10 per sheet
Layer map visuals	5 per sheet

* The per dataset limit applies to calculations that were created in the analysis. Dataset level calculations are not included in this limit. For more information about dataset level calculations, see [Adding calculated fields](#).

Saving changes to analyses

When working on an analysis, you can set Autosave either on (the default) or off. When Autosave is on, your changes are automatically saved every minute or so. When Autosave is off, your changes aren't automatically saved, which means that you can make changes and pursue different lines of inquiry without permanently altering the analysis. If you decide that you want to save your results after all, re-enable Autosave. Your changes up to that point are then saved.

In either Autosave mode, you can undo or redo up to 200 changes that you make by choosing **Undo** or **Redo** on the application bar.

Changing the Autosave mode

Changed to To change the Autosave mode for an analysis, choose **File** and then choose **Autosave On** or **Autosave OFF**.

When Autosave can't save changes

Suppose that one of the following things occurs:

- Autosave is on and another user makes a conflicting change to the analysis.
- Autosave is on and there is a service failure, such that your most recent changes can't be saved.
- Autosave is off, you turn it on, and one of the backlogged changes now being saved to the server conflicts with another user's changes.

In this case, Amazon Quick Sight gives you the option to do one of two things. You can either let Amazon Quick Sight turn Autosave off and continue working in unsaved mode, or reload the analysis from the server and then redo your most recent changes.

If your client authentication expires while you are editing an analysis, you are directed to sign in again. On successful sign-in, you are directed back to the analysis where you can continue working normally.

If your permissions on the analysis are revoked while you are editing it, you can't make any further changes.

Exporting data from Quick Sight analyses

Note

Export files can directly return information from the dataset import. This makes the files vulnerable to CSV injection if the imported data contains formulas or commands. For this reason, export files can prompt security warnings. To avoid malicious activity, turn off links and macros when reading exported files.

You can export data from an analysis to a CSV or PDF file. To export data from an analysis or dashboard to a CSV file, follow the procedure in [Exporting data from visuals](#).

Use the procedure below to export an analysis as a PDF.

1. From the analysis that you want to export, choose **File > Export to PDF**. Quick Sight begins to prepare the analysis for download.
2. Choose **VIEW EXPORTS** in the blue pop-up to open the **Exports** pane on the right.
3. Choose **DOWNLOAD** in the green pop-up.
4. To see all analyses or reports that are ready to download, choose **File** then **Exports**. The Exports panel will open on the right side of the screen. Select **Click to download** next to the file that you want to save to your preferred location.

The process for exporting to a PDF works the same way for both dashboards and analyses.

You can also attach a PDF to dashboard email reports. For more information, see [Scheduling and sending Quick Sight reports by email](#).

Deleting an analysis

If you have the permissions to do so, you can delete an analysis from the **Analyses** page. When you delete an analysis, it doesn't affect any dashboards that are based on that analysis. They continue to show the deleted analysis, but you can't make changes to the analysis once you delete it. Navigate to the Analyses page and find the analysis that you want to remove. Choose the details icon (:) on the analysis, then choose **Delete**. Confirm your choice by choosing **Delete** again. You can't undo this action.

Adding and managing sheets

A *sheet* is a set of visuals that are viewed together in a single page. When you create an analysis, you place visuals in the workspace on a sheet. You can imagine this as a sheet from a newspaper, except that it is filled with data visualizations. You can add more sheets, and make them work separately or together in your analysis.

The top sheet, also called the default sheet, is the one on the far left. This sheet displays on top in an analysis or dashboard. Each analysis can contain up to 20 sheets.

You can share analyses and publish dashboards with multiple sheets. You can also schedule email reports for any combination of sheets in an analysis.

When you create a new analysis or a new sheet in an existing analysis, you choose whether to make the new sheet an **Interactive sheet** or a **Paginated report**. This way, you can have analyses for interactive sheets only, analyses for paginated reports only, or you can have an analysis that includes both interactive sheets and paginated reports.

An *Interactive Sheet* is a collection of data expressed in visuals that users can interact with when the sheet is published to a dashboard. Amazon Quick Suite authors can add different controls and filters to their interactive sheets. Dashboard viewers can use these to gain detailed information from the published data. For more information on interactive sheets, see [Working with interactive sheets in Amazon Quick Sight](#).

A *Paginated Report* is a collection of tables, charts, and visuals that are used to convey business critical information, such as daily transaction summaries or weekly business reports. In order to

create paginated reports in Quick Sight, add the **Paginated reporting Add-on** to your Quick Suite account. To get the **Paginated reporting Add-on** and start working with paginated reports, see [Working with paginated reports in Amazon Quick Sight](#).

Use the following list of actions to work with sheets:

- To add a new sheet, choose the plus sign (+) to the right of the sheet tabs, choose the type of sheet that you want, and then choose **ADD**.
- To rename a sheet, choose the name of the sheet and start typing. **Rename** is also available from the sheet menu.
- To duplicate a sheet, choose the name of the sheet, then choose **Duplicate** from the sheet menu. You can only duplicate a sheet if **Autosave** is turned on.
- To duplicate an interactive sheet and convert it to a paginated report, choose the name of the sheet, then choose **Duplicate to report** from the sheet menu. You can't convert a paginated report to an interactive sheet.
- To delete a sheet, choose the name of the sheet, then choose **Delete** from the sheet menu. You can't delete the sheet if it's the only sheet in the analysis.
- To change the order of the sheets, choose the name of the sheet and drag it to a new position.
- To copy a visual to a new sheet, choose **Duplicate visual to** from the on-visual menu. Then choose the target sheet. Filters exist only on the sheet that you create them on. To duplicate filters, recreate them on the target sheet.
- To highlight specific data points across visuals in a sheet, go to the **Sheets** tab and select **Layout Settings**. Under the **Interactivity** section, select either **On selection** or **On hover** to turn highlighting on, or **No highlight** to turn it off. By default, sheet highlighting follows the same settings as analysis highlighting.

When you select or hover over a data point on a visual, related data across other visuals will stand out, while unrelated data is dimmed. Highlighting allows you to understand correlations, spot patterns, trends, and outliers, and facilitate stronger, more informed analyses.

You can use the parameter controls on the top sheet to control multiple sheets. To do this, open each sheet that you want to work with the parameter. Then add a filter that uses the same parameter used in the control on the top sheet. Or, if you want a new sheet to operate independently, you can add parameters and parameter controls to it that are separate from those on the top sheet.

Working with interactive sheets in Amazon Quick Sight

An *Interactive Sheet* is a collection of data expressed in visuals that users can interact with when the sheet is published to a dashboard. Amazon Quick Suite authors can add different layouts, controls, and filters to their interactive sheets that dashboard viewers can use to gain detailed information from the published data. By default, every sheet in an analysis is an interactive sheet. If your account doesn't have the **Paginated reporting Add-on**, you can only create and publish interactive sheets.

For more information on creating an interactive sheet, see [Starting an analysis in Quick Sight](#).

For more information on formatting interactive sheets, see the following topics.

Topics

- [Customizing dashboard layouts in Amazon Quick Sight](#)
- [Parameters in Amazon Quick Suite](#)
- [Using custom actions for filtering and navigating](#)

Customizing dashboard layouts in Amazon Quick Sight

You can customize a dashboard's layout to organize your data to fit your business requirements. You can choose from three dashboard layouts. You can also change the size, background color, border color, and interactions of a visual to create a fully customized dashboard.

Use the following topics to learn more about customizing dashboards and visuals.

Topics

- [Types of layout](#)
- [Choosing a layout](#)
- [Customizing visuals in a free-form layout](#)
- [Conditional rules](#)

Types of layout

There are three dashboard layout designs you can choose from: **Tiled**, **Free-form**, and **Classic**.

Tiled layout

Visuals in a **Tiled** layout snap to a grid with standard spacing and alignment. You can make visuals any size and place them wherever you want within a dashboard, but visuals can't overlap.

Dashboards are displayed as designed, with options to fit to screen or view at actual size. You can also fit an entire dashboard to your window by choosing **Fit to window** for **View** in the top-right corner. This option was previously called **Optimized**.

Note

On mobile devices, tiled layout dashboards appear as a single column in portrait mode or exactly as designed in landscape mode.

Free-form layout

Visuals in a **Free-form** layout can be placed anywhere in your dashboard using precise coordinates. You can drag a visual to the exact place you want, or you can enter the coordinates of the visual's location. Use the following procedure to enter the exact coordinates of the visual's location.

Dashboards are displayed the way that you choose to design them, with options to fit to screen or to view at its actual size. You can optimize free-form layouts for viewing at specific resolutions, with the default being 1,600 pixels. You can also fit an entire dashboard to a browser's window by choosing **Fit to window** for **View** in the top-right corner.

Note

Dashboards with optimized resolutions might appear bigger or smaller on a viewer's computer if the viewer's computer resolution doesn't equal the set resolution of the dashboard.

Switching from **Free-form** to another layout might cause some visual elements to shift. On mobile devices, **Free-form** dashboards appear as published with no changes to the layout.

Classic layout

Visuals in a **Classic** layout snap to a grid with standard spacing and alignment. Dashboards hide data or change formatting to fit smaller screen sizes. For example, if you change a visual to make it considerably smaller, the on-visual menu and editors are hidden so that the chart elements have more room to display. Bar chart visuals can also display fewer data points.

If you reduce the size of the browser window, Amazon Quick Sight resizes and if necessary reorders visuals for optimal display. For example, smaller visuals that were side by side might be displayed sequentially. The original layout is restored when the size of the browser window is increased again.

Note

On mobile devices, classic layout dashboards appear as a single column or exactly as designed in landscape mode.

Choosing a layout

To change a dashboard's layout

1. Open the [Quick Suite console](#).
2. From the Quick Suite homepage, choose **Analyses**, and then choose the analysis that you want to change.
3. On the analysis page, choose **Edit** and then choose **Analysis Settings**.
4. Expand **Sheet layout** and choose the layout that you want to use.
5. When finished, choose **Apply**.

Customizing visuals in a free-form layout

You can use the free-form layout to fully customize the color, size, location, and visibility of each visual in a dashboard.

Organizing visuals

Besides dragging a visual to its preferred location within a dashboard, there are many different ways to move a visual to the exact location it needs to be.

To enter the coordinates of the visual's location

1. Choose the visual that you want.
2. On the menu in the upper-right corner of the visual, select the **Format visual** icon.
3. In the **Properties** pane that opens, choose **Placement**.
4. Enter the **X** and **Y** coordinates of the location you want to place your visual. You can also adjust the size of the visual by entering **Width** and **Height** values.

Selected visuals can also be moved pixel-by-pixel using your keyboard's arrow keys.

You can overlay visuals on top of one another to create multi-layered visuals that show data.

Visuals can be organized into multiple layers that can be manually moved to the front and back.

To move overlaid visuals to the front and back

1. Choose the visual that you want.
2. On the three-dot menu in the upper-right hand side of the visual, choose **Menu options**.
3. For **Menu options**, choose from the following:
 - **Send to back** sends the visual to the back.
 - **Send backward** sends the visual one layer back.
 - **Bring forward** rings the visual one layer forward.
 - **Bring to front** brings a visual to the front.

Changing a visual's background color

The colors of a visual's background, border, and selection frame can be customized in the **Display settings** pane of the **Properties** pane.

To change the color of a visual's background, border, or selection frame

1. Choose the visual that you want to change.
2. On the menu in the upper-right hand side of the visual, choose the **Properties** icon.
3. In the **Properties** pane that appears on the left, choose **Display settings**.
4. Navigate to the **Card style** section and perform one or more of the available actions:

- To change the color of a visual's background, choose the **Background** color box, and then choose the color that you want.
- To change the color of a visual's border, choose the **Border** color box, and then choose the color that you want.
- To change the color of a visual's selection frame, choose the **Selection** color box, and then choose the color that you want.

If you want to use a custom color for your visual's background, border, or selection frame, choose the color box of the property that you want to change, and then choose **Custom color**. In the **Custom color** window that appears, choose your custom color or enter the color's hexadecimal code. When you are finished, choose **Apply**.

You can also reset a visual's customized background back to its default appearance.

To reset the appearance of a visual

1. Choose the visual that you want to change.
2. On the menu in the upper-right hand side of the visual, choose the **Properties** icon.
3. In the **Properties** pane that appears on the left, choose **Display settings**.
4. Choose the color that you want to reset, and then choose **Reset to default**.

Hiding visual backgrounds, borders, and selection colors

You can also choose not to show the background border, or selection color of a visual. This is useful for when you want to overlap multiple visuals. You can hide a visual's background, border, and selection colors by choosing the eye icon next to the **Border**, **Background**, or **Selection** color boxes. You can also remove a visual's loading animation by clearing the **Show loading animation** box.

Disabling visual menus

Use the **Interactions** panel of the **Properties** pane to hide the **Context** menu and **On-visual** menu from selected visuals. You can hide secondary visual menus to make the visual less crowded or to make a visual act like an overlay.

The **Context** menu opens on data-point clicks. Common actions in the **Context menu** include **Focus**, **Exclude**, and **Drill-down**.

The **On-visual** menu appears on the top-right side of a visual. The **On-visual** menu is used to access the **Properties** pane, **Maximize** the visual, access the **Menu options** panel, and review an **Anomaly insight**.

You can turn off the secondary visual menus by clearing the **Context menu** and **On-visual menu** options.

Note

You can't preview changes to the **Interactions** panel in **Analyses**. Publish the dashboard to view your changes.

Conditional rules

This feature is currently available with the **Free-form** layout. Conditional rules are used to hide or show visuals when specific conditions are met. This can be useful when you have multiple versions of the same visual overlapped with each other and want the dashboard viewer to see a version that best represents the parameter value they select.

Conditional rules use parameters and parameter controls to hide and show visuals. Parameters are named variables that can transfer a value for use by an action or an object. This feature supports string and number parameters. To make the parameters accessible to the dashboard viewer, you add a parameter control. A parameter control allows users to choose a value to use in a predefined filter or URL action. For more information about parameters and parameter controls, see [Parameters in Amazon Quick Suite](#).

Use the sections below to set up and use conditional rules.

Topics

- [Hiding a visual by default](#)
- [Setting a conditional rule](#)
- [Using conditional rules](#)

Hiding a visual by default

In the **Interactions** pane of the **Properties** pane, you can choose to hide a visual by default. Doing this can be useful if you want the viewer to only see visuals based on specific conditions.

To hide a visual by default

1. From the Quick Suite homepage, choose **Analyses**, and then choose the analysis that you want to customize.
2. Choose the visual that you want to add a rule to.
3. On the menu in the upper-right hand side of the visual, choose **Properties**.
4. In the **Properties** pane that opens, choose **Interactions** and open the **Rules** dropdown.
5. In the **Rules** menu, choose **Hide this visual by default**.

Hidden visuals appear fully hidden in a viewing dashboard. In the **Analyses** pane, hidden visuals are visible with the message "Hidden based on rule". With this display, you can see where all of a dashboard's visuals are located.

Note

You can't create conditional rules that hide visuals that are already hidden by default or that show visuals that already appear by default. If you change the default appearance of a visual, existing rules that contradict the new default appearance will be disabled.

Setting a conditional rule

When you set up a conditional rule, you create a conditional statement that will hide or show a visual when a specific condition is met. You can currently create conditional rules that hide or show a visual. If you want to create a conditional rule that makes a hidden visual appear, choose **Hide this visual by default** in the **Rules** menu of the **Properties** pane.

Note

Before you begin, make a parameter and a corresponding parameter control to base your new conditional rule on. Supported parameters are string parameters and number parameters. For more information about parameters and parameter controls, see [Parameters in Amazon Quick Suite](#).

To set a conditional rule

1. From the Quick Suite homepage, choose **Analyses**, and then choose the analysis you want to customize.
2. Choose the visual that you want to add a rule to.
3. On the menu in the upper-right hand side of the visual, choose **Properties**.
4. In the **Properties** pane that appears on the left, choose **Interactions**, and then choose **Rules**.
5. Choose **ADD RULE**.
6. In the first menu in the **Add rule** pane, choose the parameter you want.
7. In the second menu in the **Add rule** pane, choose which condition you want. For string parameters, supported conditions are **Equals**, **Starts with**, **Contains**, and **Does not equal**. For number parameters, supported conditions are **Equals**, **Starts with**, **Contains**, and **Does not equal**.
8. Enter the value you want the conditional rule to meet.

Note

Values are case-sensitive.

9. Choose **Add rule** to apply the new conditional rule to the visual. To cancel the rule, choose **Cancel**.

Conditional rules can also be edited and deleted.

To edit a conditional rule

1. On the menu in the upper-right hand side of the visual, choose **Properties**.
2. In the **Properties** pane that appears on the left, choose **Interactions**, and then choose **Rules**.
3. Choose the menu icon on the right-hand side of the rule you want to edit, and choose **Edit**.
4. Make the changes that you want and choose **Save**.

To delete a conditional rule

1. On the menu in the upper-right hand side of the visual, choose **Properties**.
2. In the **Properties** pane that appears on the left, choose **Interactions**, and then choose **Rules**.

3. Choose the menu icon on the right-hand side of the rule you want to edit and choose **Delete**.

Using conditional rules

Once you have set up a conditional rule that is connected to a parameter and a parameter control, you can use the parameter control to enable or disable the conditional rules you have set.

To enable a conditional rule

1. From the Quick Suite homepage, choose **Analyses**, and then choose the analysis you want to customize.
2. On the **Controls** bar at the top of your workspace, choose the dropdown icon.
3. Choose the parameter control associated with the conditional rule you created.
4. Choose the value associated with the conditional rule that you created from the parameter's menu. You can also enter the value that you want into the **Search value** box.

Note

Values are case-sensitive.

Selecting the correct value causes the visual to appear or disappear depending on the rule you set.

You can also bring a parameter control to the sheet your visual is on. This is useful when you want a parameter control to be next to the visual it is associated with or when you want to add a conditional rule to the control so it appears only when specific conditions are met.

To bring a parameter control to a sheet

1. From the Quick Suite homepage, choose **Analyses**, and then choose the analysis you want to customize.
2. On the **Controls** bar at the top of your workspace, choose the control that you want to move.
3. At the upper right-hand side of the control, open the **Menu options** menu.
4. Choose **Move to sheet**.

To move a parameter control back to the Controls bar

1. On your dashboard, select the parameter control you want to move.
2. On the upper right-hand side of the control, open the **Menu options** menu.
3. Choose **Move to top of sheet**.

Parameters in Amazon Quick Suite

Parameters are named variables that can transfer a value for use by an action or an object. By using parameters, you can create an easier way for a dashboard user to interact with dashboard features in a less technical way. Parameters can also connect one dashboard to another, allowing a dashboard user to drill down into data that's in a different analysis.

For example, a dashboard user can use a list to choose a value. That value sets a parameter that in turn sets a filter, calculation, or URL action to the chosen value. Then the visuals in the dashboard react to the user's choices.

To make the parameters accessible to the dashboard viewer, you add a parameter control. You can set up cascading controls, so that a selection in one control filters the options that display in another control. A control can appear as a list of options, a slider, or a text entry area. If you don't create a control, you can still pass a value to your parameter in the dashboard URL.

For a parameter to work, it needs to be connected to something in your analysis, regardless of whether it has a related control. You can reference parameters in the following:

- Calculated fields (except for multivalue parameters)
- Filters
- Dashboard and analysis URLs
- Actions
- Titles and descriptions throughout an analysis

Some ways that you can use parameters are the following:

- Using a calculation, you can transform data that is shown in an analysis.
- If you add a control with a filter to an analysis you are publishing, the dashboard users can filter the data without creating their own filters.
- Using controls and custom actions, you can let dashboard users set values for the URL actions.

Topics

- [Setting up parameters in Amazon Quick Suite](#)
- [Using a control with a parameter in Amazon Quick Suite](#)
- [Creating parameter defaults in Amazon Quick Suite](#)
- [Connecting to parameters in Amazon Quick Suite](#)

Setting up parameters in Amazon Quick Suite

Use the following procedure to create or edit a basic parameter.

To create or edit a basic parameter

1. Choose an analysis to work with, then decide which field you want to parameterize.
2. Choose the **Parameters** icon from the icon list at the top of the page.
3. Add a new parameter by choosing the plus sign (+ **Add**) near the top of the pane.

Edit an existing parameter by first choosing the v-shaped icon near the parameter name and then choosing **Edit parameter**.

4. For **Name**, enter an alphanumeric value for the parameter.
5. For **Data type**, choose **String**, **Number**, **Integer**, or **Datetime**, and then complete the following steps.

- If you choose **String**, **Number**, or **Integer**, do the following:

1. For **Values**, choose **Single value** or **Multiple values**.

Choose the single value option for parameters that can contain only one value. Choose multiple values for parameters that can contain one or more values. Multivalue parameters can't be datetime data types. They also don't support dynamic default values.

To switch an existing parameter between single and multiple values, delete and recreate the parameter.

2. (Optional) For **Static default value** or **Static multiple default values**, enter one or more values.

This type of static value is used during the first page load if a dynamic default value or URL parameter isn't provided.

3. (Optional) Choose **Show as blank by default**.

Select this option to show the default value for multivalue lists as blank. This option only applies to multivalue parameters.

- If you choose **Datetime**, do the following:

1. For **Time granularity**, choose **Day**, **Hour**, **Minute**, or **Second**.

2. For **Default date**, select either **Fixed date** or **Relative date**, and then do the following:

- If you select **Fixed date**, enter a date and time by using the date and time picker.
- If you select **Relative date**, choose a rolling date. You can choose **Today**, **Yesterday**, or you can specify the **Filter condition** (start of or end of), **Range** (this, previous, or next), and **Period** (year, quarter, month, week, or day).

6. (Optional) Choose **Set a dynamic default** to create a default that is user-specific.

A *dynamic default* is a per-user default value for the first page load of the dashboard. Use a dynamic default to create a personalized view for each user.

Calculated fields can't be used as dynamic defaults.

Dynamic defaults don't prevent a user from selecting a different value. If you want to secure the data, you can add row-level locking. For more information, see [Using row-level security with user-based rules to restrict access to a dataset](#).

This option only appears if you choose a single value parameter. Multivalue parameters can't have dynamic defaults.

Note

If you choose a multivalue parameter, the screen changes to remove the default options. Instead, you see a box with the text **Enter values you want to use for this control**. You can enter multiple values in this box, each on a single line. These values are used as the default selected values in the parameter control. The values here are unioned with what you choose to enter for the parameter control. For more information on parameter controls, see [Parameter Controls](#).

7. (Optional) Set a reserved value to determine the value of the **Select all** value. The *reserved value* of a parameter is the value that is assigned to a parameter when you choose **Select all** as its value. When you set up a specific reserved value for your parameter, that value is no longer

considered a valid parameter value in your dataset. The reserved value can't be used in any *parameter consumers*, such as filters, controls, and calculated fields, and custom actions. Also, it does not appear in the parameter control list. You can choose from **Recommended value**, **Null**, and **Custom value**. **Recommended value** is the default. If you choose **Recommended value**, the reserved value is set to the following values based on the value type:

- Strings: "ALL_VALUES"
- Numbers: "Long.MIN_VALUE"-9,223,372,036,854,775,808
- "Integers: Int.MIN_VALUE"-2147483648

To set a reserved value in your new parameter, choose the **Advanced settings** dropdown list in either the **Create a new parameter** page or the **Edit parameter** page and select the value that you want.

8. Choose **Create** or **Update** to complete creating or updating the parameter.

After you create a parameter, you can use it in a variety of ways. You can create a control (such as a button) so that you can choose a value for your parameter. For more information, see the following sections.

Using a control with a parameter in Amazon Quick Suite

In dashboards, parameter controls appear at the top of the data sheet, which contains a set of visuals. Providing a control allows users to choose a value to use in a predefined filter or URL action. Dashboard users can use controls to apply filtering across all visuals datasets on a dashboard, without having to create the filters themselves.

The following rules apply:

- To create or edit a control for a parameter, make sure that the parameter exists.
- Multiselect list controls are compatible with analysis URLs, dashboard URLs, custom actions, and custom filters. The filter must be either equal or not equal to the values provided. No other comparisons are supported.
- Lists show up to 1,000 values. If there are more than 1,000 distinct values, a search box appears so you can filter the list. When the filtered list contains less than 1,001 values, the contents of the list appear as line items.

- The **Style** option displays only the style types that are appropriate for the parameter's data type and single or multivalue setting. If the style that you want to use isn't in the list, recreate your parameter with the appropriate settings and try again.
- If your parameter links to a dataset field, it must be an actual field. Calculated fields aren't supported.
- The values display alphabetically in the control, unless there are more than 1,000 distinct values. Then the control displays a search box instead. Each time you search for the value you want to use, it initiates a new query. If the results contain more than 1,000 values, you can scroll through the values with pagination. Wildcard search is supported. To learn more about wildcard search, see [Using wildcard search](#).

Use the following procedure to create or edit a control for an existing parameter.

To create or edit a control for an existing parameter

1. Choose an existing parameter's context menu, the  icon near the parameter name, and choose **Add control**.
2. Enter a name to give the new control a label. This label appears at the top of the workspace, and later at the top of the sheet that a dashboard displays on.
3. Choose a style for the control from the following:

- **Text field**

A text field lets you type in their own value. A text field works with numbers and text (strings).

- **Text field - multiline**

A multiline text field lets you type in their own values. With this option, you can choose to separate values you enter into the parameter control by a line break, comma, pipe (|), or semicolon. A text field works with numbers and text (strings).

- **Dropdown**

A dropdown list control that you can use to select a single value. A list control works with numbers and text (strings).

- **Dropdown multiselect**

A list control that you can use to select multiple values. A list control works with numbers and text (strings).

- **List**

A list control that you can use to select a single value. A list control works with numbers and text (strings).

- **List - multiselect**

A list control that you can use to select multiple values. A list control works with numbers and text (strings).

- **Slider**

A slider lets you select a numeric value by sliding the control from one end of the bar to another. A slider works with numbers.

- **Date-picker**

Using a date-picker, you can choose a date from a calendar control. When you choose to add a date-picker control, you can customize how to format dates in the control. To do so, for **Date format**, enter the date format that you want using the tokens described in [Customizing date formats in Quick Suite](#).

4. (Optional) If you choose a dropdown control, the screen expands so you can choose the values to display. You can either specify a list of values, or use a field in a dataset. Choose one of the following:

- **Specific values**

To create a list of specific values, type in one per line, with no separating spaces or commas, as shown in the following screenshot.

In the control, the values display alphabetically, not in the order that you typed them.

- **Link to a data set field**

To link to a field, choose the dataset that contains your field, then choose the field from the list.

If you change the default values in the parameter, choose **Reset** on the control to show the new values.

The values that you choose here are unioned with the static default values in the parameter settings.

5. (Optional) Enable the option **Hide [ALL] option from the control if the parameter has a default configured**. Doing this shows only the data values and removes the option to select all items in the control. If you don't configure a static default on the parameter, this option doesn't work. You can add a default after adding a control by choosing the parameter, and selecting **Edit parameter**.
6. (Optional) You can limit the values displayed in the controls, so they only show values that are valid for what is selected in other controls. This is called a cascading control.

To create one, choose **Show relevant values only**. Choose one or more controls that can change what displays in this control.

When creating cascading controls, the following limitations apply.

- Cascading controls must be tied to dataset columns from the same dataset.
 - The child control must be a dropdown or list control.
 - For parameter controls, the child control must be linked to a dataset column.
 - For filter controls, the child control must be linked to a filter (instead of showing only specific values).
 - The parent control must be one of the following.
 - A string, integer, or numeric parameter control.
 - A string filter control (EXCLUDING Top-Bottom filters).
 - A non-aggregated numeric filter control.
 - A date filter control (EXCLUDING Top-Bottom filters).
7. When you finish choosing options for your control, choose **Add**.

The finished control appears at the top of the workspace. The context menu, shaped like a v, offers four options:

- **Reset** restores the user's selection to its default state.
- **Refresh list** applies only to drop-downs that are linked to a field in a dataset. Choosing **Refresh list** queries the data to check for changes. Data used in the control is cached.
- **Edit** reopens the control creation screen so that you can change your settings.

Once you have the **Edit control** pane open, you can click on different visuals and controls to view formatting data for the specific visual or control. For more information about formatting a visual, see [Formatting in Amazon Quick Suite](#).

- **Delete** removes the control. You can recreate it by choosing the parameter context menu.

In the workspace, you can also resize and rearrange your controls. The dashboard users see them as you do, except without being able to edit or delete them.

Creating parameter defaults in Amazon Quick Suite

Use this section to learn more about the types of parameter defaults that are available, and how to set up each of them.

Each field can have a parameter and a control associated with it. When someone views a dashboard or email report, any sheet control that has a static default value configured uses the static default. The default value can change how data is filtered, how custom actions behave, and what text displays in a dynamic sheet title. Email reports also support dynamic defaults.

The simplest default is a static (unchanging) default, which shows the same value to everyone. As the designer of the dashboard, you choose the default value. It can't be changed by the person using the dashboard. However, that person can choose any value from the controls. Setting a default doesn't change this. To restrict the values that a person can select, consider using row-level security. For more information, see [Using row-level security with user-based rules to restrict access to a dataset](#).

To create or edit a static default value that applies to everyone's dashboard view

1. Choose the context menu (v) by the parameter that you want to edit, or create a new parameter by following the steps in [Setting up parameters in Amazon Quick Suite](#).
2. Enter a value for **Static default value** to set a static default.

To display a different default depending on who is viewing the dashboard, you create a dynamic default parameter (DDP). Using dynamic defaults involves some preparation to map people to their assigned defaults. First, you need to create a database query or a data file that contains information about the people, the fields, and the default values to display. You add this to a dataset, then add the dataset to your analysis. Following, you can find procedures that you can use to gather information, create the dataset, and add the dynamic default to the parameter.

Use the following guidelines when creating a dataset for dynamic default values:

- We recommend that you use a single dataset to contain all dynamic default definitions for a logical grouping of users or groups. If you can, maintain them in a single table or file.
- We also recommend that the fields in your dataset have names that closely match the field names in the analysis. Not all dataset fields need to be part of the analysis, for example if you're using the same dataset for the defaults in multiple dashboards. The fields can be in any order.
- We don't recommend that you combine both user and group names in the same column or even in the same dataset. This kind of configuration is more work to maintain and troubleshoot.
- If you use a comma-delimited file to create your dataset, make sure to remove any space between values in the file. The following example shows the correct comma-separated value (CSV) format. Enclose text (strings) that include nonalphanumeric characters—like spaces, apostrophes, and so on—in single or double quotation marks. You can enclose fields that are dates or times in quotation marks, but it isn't required. You can enclose numeric fields in quotation marks, for example if the numbers contain special characters, as shown following.

```
"Value includes spaces","Field contains ' other characters",12345.6789,"20200808"
ValueWithoutSpaces,"1000,67","Value 3",2020-AUG-08
```

- After you create the dataset, make sure to double-check the data types that Quick Suite selects for the fields.

Before you begin, you need a list of the user or group names for the people who are going to have dynamic defaults. To generate a list of users or groups, you can use the Amazon CLI to get the information. To run CLI commands, make sure that you have the Amazon CLI installed and configured. For more information, see [Installing the Amazon CLI](#) in the *Amazon CLI User Guide*.

This is just one example of how to get a list of user or group names. Use whatever method works best for you.

To identify people for a dynamic default parameter (DDP)

- List either individual user names or group names:
 - To list individual user names, include a column that identifies the people for your DDP. This column should contain each person's system user name that they use to connect from your identity provider to Quick Suite. This user name is often the same as a person's email alias before the @ sign, but not always.

To get a list of users, use the [ListUsers](#) Quick Suite API operation or Amazon CLI command. The CLI command is shown in the following example. Specify the Amazon Web Services Region for your identity provider, for example `us-east-1`.

```
awsacct1="111111111111"
namespace="default"
region="us-east-1"

aws quicksight list-users --aws-account-id $awsacct1 --namespace $namespace --
region $region
```

The following example alters the previous command by adding a query that limits the results to active users.

```
awsacct1="111111111111"
namespace="default"
region="us-east-1"

aws quicksight list-users --aws-account-id $awsacct1 --namespace $namespace --
region $region --query 'UserList[?Active==`true`]'
```

The result set looks similar to the following sample. This example is an excerpt from JSON output (`--output json`). People who have federated user names have principal IDs that start with the word `federated`.

```
[
  {
    "Arn": "arn:aws-cn:quicksight:us-east-1:111111111111:user/default/
anacasilva",
    "UserName": "anacarolinasilva",
    "Email": "anacasilva@example.com",
    "Role": "ADMIN",
    "Active": true,
    "PrincipalId": "federated/iam/AIDAJ64EIEIOPX5CEIEIO"
  },
  {
    "Arn": "arn:aws-cn:quicksight:us-east-1:111111111111:user/default/Reader/
liujie-stargate",
    "UserName": "Reader/liujie-stargate",
    "Role": "READER",
  }
]
```

```

        "Active": true,
        "PrincipalId": "federated/iam/AR0AIJSEIEIOMXTZEIEIO:liujie-stargate"
    },
    {
        "Arn": "arn:aws-cn:quicksight:us-east-1:111111111111:user/default/
embedding/cxoportal",
        "UserName": "embedding/cxoportal",
        "Email": "saanvisarkar@example.com",
        "Role": "AUTHOR",
        "Active": true,
        "PrincipalId": "federated/iam/AR0AJTGEIEIOWB6BEIEIO:cxoportal"
    },
    {
        "Arn": "arn:aws-cn:quicksight:us-east-1:111111111111:user/default/
zhangwei@example.com",
        "UserName": "zhangwei@example.com",
        "Email": "zhangwei@example.com",
        "Role": "AUTHOR",
        "Active": true,
        "PrincipalId": "user/d-96123-example-id-1123"
    }
]

```

- To list group names, include a column that identifies the groups containing the user names for your DDP. This column should contain the system group names that are used to connect from your identity provider to Quick Suite. To identify groups that you can add to the dataset, use one or more of the following Quick Suite API operations or CLI commands:
 - [ListGroups](#) – Lists Quick Suite groups by Amazon Web Services account ID and namespace for the Amazon Web Services Region that contains your identity provider.
 - [ListGroupMemberships](#) – Lists the users in the specified Quick Suite group.
 - [ListUserGroups](#) – Lists the Quick Suite groups that a Quick Suite user is a member of.

Or you can ask your network administrator to query your identity provider to get this information.

The next two procedures provide instructions on how to finish creating a dataset for dynamic default values. The first procedure is for creating a dataset for a single-value DDP. The second one is for creating a dataset for a multivalue DDP.

To create a dataset for a single-value DDP

1. Create dataset columns with single-value parameters. The first column in the query or file should be for the people using the dashboard. This field can contain user names or group names. However, support for groups is only available in Quick Suite Enterprise edition.
2. For each field that displays a dynamic default for a single-value parameter, add a column to the dataset. The name of the column doesn't matter—you can use the same name as the field or parameter.

Single-value parameters only work as specified if the combination of user entity and dynamic default is unique for that parameter's field. If there are multiple values a default field for a user entity, the single-value control for that field displays the static default instead. If no static default is defined, the control doesn't display a default value. Be careful if you use group names, because some user names can be members of multiple groups. If those groups have different default values, then this type of user name functions as a duplicate entry.

The following example shows a table that appears to contain two single-value parameters. We make this assumption because no user name is paired with multiple default values. To make this table easier to understand, we add the word 'default' in front of the field names from the analysis. Thus, you can read the table by making the following statement, changing the values for each row: When viewed by anacarolinasilva, the controls display a default region NorthEast and a default segment SMB.

Viewed-by	Default-region	Default-segment
anacarolinasilva	NorthEast	SMB
liujie	SouthEast	SMB
saanvisarkar	NorthCentral	SMB
zhangwei	SouthCentral	SMB

3. Import this data into Quick Suite, and save it as a new dataset.
4. In your analysis, add the dataset that you created. The analysis needs to use at least one other dataset that matches the columns you defined for the defaults. For more information, see [Adding a dataset to an analysis](#).

To create a dataset for a multivalue DDP

1. Create dataset columns with multivalue parameters. The first column in the query or file should be for the people using the dashboard. This field can contain user names or group names. However, support for groups is only available in Quick Suite Enterprise edition.
2. For each field that displays a dynamic default for a multivalue parameter, add a column to the dataset. The name of the column doesn't matter—you can use the same name as the field or parameter.

Unlike single-value parameters, multivalue parameters allow multiple values in the field that's associated with the parameter.

The following example shows a table that appears to contain a single-value parameter and a multivalue parameter. We can make this assumption because each user name has a unique value in one column, and some user names have multiple values in the other column. To make this table easier to understand, we add the word 'default' in front of the field names from the analysis. Thus, you can read the table by making the following statement, changing the values for each row: When viewed-by is liujie, the controls display a default-region value of SouthEast, and a default-city value of Atlanta. And if we read ahead one row, we see that liujie also has Raleigh in default-city.

Viewed-by	Default-region	Default-city
anacarolinasilva	NorthEast	New York
liujie	SouthEast	Atlanta
liujie	SouthEast	Raleigh
saanvisarkar	NorthCentral	Chicago
zhangwei	SouthCentral	Dallas
zhangwei	SouthCentral	Kansas City

In this example, the parameter that we apply default-region to works correctly whether it's a single-value or multivalue parameter. If it's a single-value parameter, two entries work for one user because both entries are the same value, SouthEast. If it's a multivalue parameter,

it still works, except that only one value is selected by default. However, if we change the parameter that's using `default-city` as its default from a multivalue to a single-value parameter, we don't see these defaults selected. Instead, the parameter uses the static default, if there is one defined. For example, if the static default is set to Atlanta, liujie has Atlanta selected in that control, but not Raleigh.

In some cases, your static default value might also be used as a dynamic default. If so, make sure to test the control for a user name that doesn't use a default value that can be both.

If a user name belongs to multiple groups, the named user sees a set of default values that is a union of the two groups' default values.

3. Import this data into Quick Suite, and save it as a new dataset.
4. In your analysis, add the dataset that you created. The analysis needs to use at least one other dataset that matches the columns you defined for the defaults. For more information, see [Adding a dataset to an analysis](#).

Use the following procedure to add a dynamic default parameter to your analysis. Before you begin, make sure that you have a dataset that contains the dynamic defaults for each user name or group name. Also make sure that your analysis is using this dataset. For help with these requirements, see the procedures preceding.

To add a DDP to your analysis

1. In the Quick Suite console, choose the **Parameters** icon at the top of the page and choose an existing parameter. Choose **Edit parameter** from the parameter's menu. To add a new parameter, choose the plus (+) sign near **Parameters**.
2. Choose **Set a dynamic default**.
3. Configure the following options with your settings:
 - **Dataset with default values and user information** – Choose the dataset that you created and added to your analysis.
 - **User name column** – To create defaults that are based on user names, choose the column in the dataset that contains the user names.
 - **Group name column** – To create defaults that are based on group names, choose the column in the dataset that contains the group names.

- **Column for default value** – Choose the column that contains default values for this parameter.
4. Choose **Apply** to save your setting changes, and then choose **Update** to save the parameter changes. To exit without saving changes, choose **Cancel** instead.
 5. Add a filter for each field that contains dynamic defaults to make the defaults work. To learn more about using filters with parameters, see [Using filters with parameters in Amazon Quick Suite](#)

Amazon Quick Suite uses the static default value for anyone whose user name doesn't exist in the dataset, doesn't have a default assigned, or doesn't have a unique default. Each person can have only one set of defaults. If you don't want to use dynamic defaults, you can set a static default instead.

Connecting to parameters in Amazon Quick Suite

Use this section after you have a parameter set up, to connect it and make it work.

After you create a parameter, you can create consumers of the parameters. *Parameter consumers* are components that consume the value of a parameter, such as filters, controls, calculated fields, or custom actions.

You can navigate to each of these options in another way, as follows:

- To create a filter, choose the **Filter** icon at the top of the page. In short, you create a **Custom Filter** and enable **Use parameters**. The list shows only eligible parameters.
- To add a new control for the parameter, choose the **Parameters** icon at the top of the page. In short, choose your parameter, and then **Add control**.
- To use a parameter in a calculated field, either edit an existing calculated field, or add a new one by choosing **Add** at the top left. The parameter list appears below the field list.

Note

You can't use multivalue parameters with calculated fields.

- To create a URL action, choose the v-shaped menu on a visual, and then choose **URL Actions**.

For more information on each of these topics, see the following sections.

Topics

- [Using filters with parameters in Amazon Quick Suite](#)
- [Using calculated fields with parameters in Amazon Quick Suite](#)
- [Using custom actions with parameters in Amazon Quick Suite](#)
- [Using parameters in a URL](#)
- [Using parameters in titles and descriptions in Amazon Quick Suite](#)

Using filters with parameters in Amazon Quick Suite

Use this section to filter the data in an analysis or dashboard by a single-value parameter value. To use a multivalued parameter—one with a multiselect drop-down control—create a custom filter that is equal (or not equal) to the values.

Before using a filter with a parameter, you should already know how to work with filters.

1. Verify that your analysis has a parameter already created. Choose **Edit** from either the parameter or the control menu to find out what settings are in use.
2. Choose the **Filter** pane from the left of the screen. If there is already a filter for the field that you want to use, choose it to open its settings. Otherwise, create a filter for the field that you want to filter by parameter.
3. Choose **Use Parameters**.
4. Choose your parameters from the list or lists below **Use parameters**. For text (string) fields, first choose **Custom Filter**, and then enable **Use Parameters**.

For date fields, choose the **Start date** and **End date** parameters, as shown in the following screenshot.

For fields with other data types, choose **Select a parameter** and then choose your parameter from the list.

Note

Parameters that can hold multiple values must use equal or not equal as the comparison type.

5. Choose **Apply** to save your changes.

Test your new filter by choosing the control near the top of the analysis. In this example, we use a basic parameter that has no defaults, and a dynamic control that is linked to the **Region** field in the sample dataset named **Sales Pipeline**. The control queries the data, returning all values.

If you delete or recreate a parameter that you are using in a filter, you can update the filter with the new parameter. To do this, open the filter, choose the new parameter that you want to use, and then choose **Apply**.

If you rename a parameter, you don't need to update the filter or any other consumers.

Using calculated fields with parameters in Amazon Quick Suite

You can pass the value of a parameter to a calculated field in an analysis. When you create a calculation, you can choose existing parameters from the list of parameters under **Parameter list**. You can't create a calculated field that contains a multivalued parameter—those with a multiselect drop-down control.

For the formula, you can use any of the available functions. You can pass the viewer's selection from the parameter control, to the `ifElse` function. In return, you get a metric. The following shows an example.

```
ifelse(  
  
  ${KPIMetric} = 'Sales',sum({Weighted Revenue}),  
  
  ${KPIMetric} = 'Forecast',sum({Forecasted Monthly Revenue}),  
  
  ${KPIMetric} = '# Active', distinct_count(ActiveItem),  
  
  NULL  
  
)
```

The preceding example creates a metric (a decimal) that you can use in a field well. Then, when a user chooses a value from the parameter control, the visual updates to reflect their selection.

Using custom actions with parameters in Amazon Quick Suite

A *custom action* enables you to launch URLs or filter visuals by selecting a data point in a visual or choosing the action name from the context menu. When you use a URL action with a parameter, you can pass or send parameters dynamically to the URL. To make this work, you set up a parameter, and then use it in the URL when you create a custom action with an action type of **URL**.

action. The parameters on both the sending and the receiving end must match in name and data type. All parameters are compatible with URL actions.

For details on creating a URL action, see [Creating and editing custom actions in Amazon Quick Sight](#). If you just want to use a parameter in a link without creating a URL action, see [Using parameters in a URL](#).

Using parameters in a URL

You can use a parameter name and value in a URL in Amazon Quick Suite to set a default value for that parameter in a dashboard or analysis.

The following example shows the URL of a dashboard that sets a parameter for another dashboard.

```
https://us-east-2.quicksight.aws.amazon.com/sn/dashboards/abc123-abc1-abc2-abc3-abcdefef1234#p.myParameter=12345
```

In the previous example, the first part is the link to the target dashboard: `https://us-east-2.quicksight.aws.amazon.com/sn/dashboards/abc123-abc1-abc2-abc3-abcdefef1234`. The hash sign (#) follows the first part to introduce the *fragments*, which contain the values that you want to set.

The values in the fragments aren't received or logged by Amazon servers. This functionality keeps your data values more secure.

The fragment after # follows these rules:

- Parameters are prefixed with p.. The names are the parameter name, not the control name. You can view the parameter name by opening the analysis, and choosing **Parameter** on the left sidebar.
- The value is set using equals (=). The following rules apply:
 - Literal values don't use quotation marks.
 - Spaces inside values are automatically encoded by the browser, so you don't need to use escape characters when manually creating a URL.
 - To return all values, set the parameter equal to "[ALL]".
 - To assign the parameter's value to null, set it equal to %00. For example, p.population=%00.

- In custom actions, target parameter names begin with \$, for example: <<\$passThroughParameter>>
- In custom actions, parameter values display with angle brackets << >>, for example <<dashboardParameter1>>). The dashboard user sees the lookup value, not the variable.
- For a custom URL action, multivalue parameters only need one instance of the same parameter in the fragment, for example: p.city=<<\$city>>
- For a direct URL, multiple values for a single parameter have two instances of the same parameter in the fragment. For an example, see following.
- Ampersands (&) separate multiple parameters. For an example, see following.

The server converts the date to UTC and sends it to the backend as a string without a time zone. To use Universal Coordinated Time (UTC) dates, exclude the time zone. Following are some examples of date formats that work:

- 2017-05-29T00%3A00%3A00
- 2018-04-04 14:51 -08:00
- Wed Apr 04 2018 22:51 GMT+0000

```
https://us-east-2.quicksight.aws.amazon.com/sn/dashboards/abc123-  
abc1-abc2-abc3-abcdefef1234#p.shipdate=2018-09-30 08:01&p.city=New  
York&p.city=Seattle&p.teamMember=12&p.percentageRank=2.3
```

In the browser, this code becomes the following.

```
https://us-east-2.quicksight.aws.amazon.com/sn/dashboards/abc123-  
abc1-abc2-abc3-abcdefef1234#p.shipdate=2018-09-30%2008:01&p.city=New  
%20York&p.city=Seattle&p.teamMember=12&p.percentageRank=2.3
```

The previous example sets four parameters:

- shipDate is a date parameter: Sept 30, 2018.
- city is a multivalued string parameter: New York, and Seattle
- teamMember is an integer parameter: 12.

- `percentageRank` is a decimal parameter: `2.3`.

The following example shows how to set values for a parameter that accepts multiple values.

```
https://us-east-2.quicksight.aws.amazon.com/sn/dashboards/abc123-abc1-abc2-abc3-  
abcdefef1234#p.MultiParam=WA&p.MultiParam=OR&p.MultiParam=CA
```

To pass values from one dashboard (or analysis) to another dashboard based on the user's data point selection, use custom URL actions. If you choose, you can also generate these URLs manually, and use them to share a specific view of the data.

For information on creating custom actions, see [Using custom actions for filtering and navigating](#).

Using parameters in titles and descriptions in Amazon Quick Suite

When you create parameters in Amazon Quick Suite, you can use them in titles and descriptions throughout your charts and analyses to dynamically display parameter values.

You can use parameters in the following areas of your analysis:

- Chart titles and subtitles
- Axis titles
- Legend titles
- Parameter control titles
- Sheet titles and descriptions

The following image shows a chart title that uses a parameter.

Use the following procedures to learn how to add parameters to areas throughout your analysis. For more information about parameters and how to create them, see [Parameters](#).

Adding parameters to chart titles and subtitles

Use the following procedure to learn how to add parameters to chart titles and subtitles.

To add a parameter to a chart title or subtitle

1. Open the **Properties** pane for the visual that you want to format.

2. In the **Properties** pane, choose the **Title** tab.
3. Select **Show title** or **Show subtitle**. These options might already be selected.
4. Choose the three dots at the right of **Edit title** or **Edit subtitle**, and then choose a parameter from the list.

The parameter is added to the title in the **Properties** pane. In the chart, the parameter value is displayed in the title.

For more information about editing titles and subtitles in visuals, see [Titles and subtitles on visual types in Quick Suite](#).

Adding parameters to axis titles

Use the following procedure to learn how to add parameters to axis titles.

To add a parameter to an axis title

1. Open the **Properties** pane for the visual that you want to format.
2. In the **Properties** pane, choose the axis that you want to format.
3. Select **Show title**.
4. Choose the three dots at the right of the default axis title, and then choose a parameter from the list.

The parameter is added to the axis title in the **Properties** pane. In the chart, the parameter value is displayed in the axis title.

For more information about editing axis titles, see [Axes and grid lines](#).

Adding parameters to legend titles

Use the following procedure to learn how to add parameters to legend titles.

To add a parameter to a legend title

1. Open the **Properties** pane for the visual that you want to format.
2. In the **Properties** pane, choose **Legend**.
3. Select **Show legend title**.
4. Choose the three dots at the right of **Legend title**, and then choose a parameter from the list.

The parameter is added to the legend title in the **Properties** pane. In the chart, the parameter value is displayed in the legend title.

For more information about formatting legends, see [Legends on visual types in Quick Suite](#).

Adding parameters to control titles

Use the following procedure to learn how to add parameters to parameter control titles.

To add a parameter to a parameter control title

1. Select the parameter control that you want to edit, choose the three dots at the right of the parameter control title, and then choose **Edit**.
2. In the **Edit control** page that opens, select **Show title**.
3. Choose the three dots at the right of **Display name**, and then choose a parameter from the list.

The parameter is added to the parameter control title.

For more information about using parameter controls, see [Parameter controls](#).

Adding parameters to sheet titles and descriptions

Use the following procedure to learn how to add parameters to sheet titles and descriptions in your analysis.

To add a parameter to a sheet title or description

1. On the analysis page, choose **Sheets** in the application bar and then choose **Add title** or **Add description**.

A sheet title or description appears on the sheet.

2. For **Sheet title** or for **Description**, choose the three dots at right, and then choose a parameter from the list.

The parameter is added to the sheet title or description and the parameter value appears in the text when you close the text box.

For more information about adding sheet titles and descriptions, see [Adding a title and description to an analysis](#).

Using custom actions for filtering and navigating

To add interactive options for dashboard subscribers (Quick Suite readers), you create custom actions on one or more visuals in your analysis. Enhancing dashboards with custom actions helps people explore data by adding more context from within the dataset. It can make it easier to drill into the details and to find new insights in the same dashboard, a different dashboard, or a different application. You can add up to 10 custom actions to each visual in a dashboard.

Before you begin, it's helpful to do some planning. For example, identify fields that are good candidates for filtering, for opening a different sheet, for opening a URL, or for sending email. For each sheet, identify the widgets that display these fields. Then decide which widgets are going to contain actions. It's also a good idea to create a naming scheme so the names of the actions are consistent throughout the entire analysis. Consistent names make it easier for the person using your analysis to figure out what the action will do, plus they make it easier for you to maintain actions that you might be duplicating throughout the analysis.

Actions only exist on the dashboard widget where you create them and they work in the context of that widget's parent sheet and child fields that it displays. You can create actions only on specific types of widget: visuals and insights. You can't add them to other widgets, for example filter or list controls. Custom actions can only be activated from the widget where you create them.

To activate an action, the person using the analysis can left-click (select) or right-click (use the context menu) on a data point. A *data point* is an item in the dataset, for example a point on a line chart, a cell in a pivot table, a slice on a pie chart, and so on. If the person clicks a visual element, the *select* action is activated. This is the action that is currently a member of the **On select** category of the **Actions** in an analysis. If the person instead right-clicks a visual element, they can choose from a list of *menu* actions. Any action listed is currently a member of the **Menu option** category of the **Actions** in an analysis. The **On select** category can contain one and only one member action.

By default, the first action you create becomes the select action—the one activated by left-clicking. To remove an action from the **On select** category, change the action's **Activation** setting to **Menu option**. After you save that change, you can set a different action's **Activation** setting to **Select**.

You can choose from three **Action types** when you configure an action:

- **Filter action** – Filter data included in visual or in the entire sheet. By default, filters are available for all fields in the parent visual. Cascading filters are enabled by default. Filter actions work across multiple datasets by using automatically generated field mappings.

If the analysis uses more than one dataset, you can view the automatically generated field mappings for fields that exist in multiple datasets. To do this, choose **View field mapping** at the end of the action settings, while you're editing an action. If you are viewing a list of actions, choose **View field mapping** from the menu for each action. The field mappings appear in a new screen that shows the mapping between the initial dataset and all the other datasets in the visual. If no fields are automatically mapped, a message displays with a link to [Mapping and Joining Fields](#).

- **Navigation actions** – Enable navigation between different sheets in the same analysis.
- **URL actions** – Open a link to another web page. If you want to open a different dashboard, use a URL action. You can use a URL action to send data points and parameters to other URLs. You can include any available field or parameter.

If the URL uses the `mailto` scheme, running the action opens your default email editor.

Topics

- [Adding one-click interactive filters](#)
- [Creating and editing custom actions in Amazon Quick Sight](#)
- [Repairing custom actions](#)
- [Understanding field mapping for custom actions in Amazon Quick Sight](#)

Adding one-click interactive filters

One-click interactive filtering provides point-and-click filtering that cascades from the clickable visual to all the other visuals and insights on a sheet. Add this to your analysis to start with summaries and drill down into the metrics, all within the same dashboard sheet.

After you set this up, when you click a data point (for example, a point in a line chart), you instantly filter using all mapped fields on all the other visuals on that sheet. If you have multiple datasets, all target fields must be mapped for this to work. Also, you can only have one action that works by clicking a data point; all other actions work from the context menu.

Use the following procedure to create a one-click filter in an analysis.

To create a one-click filter on a visual or insight

1. In your analysis, choose a visual or insight that you want to add interactive filtering to.
2. Choose **Actions** from the Menu options dropdown in the upper right corner.
3. Choose **Filter same-sheet visuals**. Doing this immediately adds one-click filtering.
4. Repeat this process for each visual that you wish to make interactive.

Creating and editing custom actions in Amazon Quick Sight

You create one action for each task that you want to be able to add to a visual. The actions you create become part of the functionality of each visual or insight.

The following table defines when to use each type of action.

Action to perform	Type of action
Add or customize an interactive filter action, including one-click filters	Filter action
Open another sheet in the same dashboard	Navigation action
Open a sheet in a different dashboard in the same Amazon Web Services account	URL action
Open a URL (https, http)	URL action
Send an email (mailto)	URL action

You can set the following attributes and options for a custom action:

- **Action name** – This is a descriptive name that you choose for the action. By default, actions are named **Action 1**, **Action 2**, and so on. If your custom action is activated from a context menu, this name displays in the menu when you right-click on a data point.

To make the action name dynamic, you can parameterize it. Use the plus icon near the action name header to display a list of available variables. Variables are enclosed in angle brackets << >>. Parameters are prefixed with a \$, for example <<\$parameterName>>. Field names have no prefix, for example <<fieldName>>.

- **Activation** – Available options are **Select** or **Menu option**. To use an action, you can *select* the data point (left-click) or navigate to the *menu option* in the context menu (right-click). Navigation actions and URL actions listed in the middle of the context menu, just above **Color** options. Actions that are activated by menu are also available from the legend on a visual.
- **Action type** – The type of action that you want. Settings that are specific to an action type only display after you choose the action type.
- **Filter action** settings include the following:
 - **Filter scope** – The fields to filter on. To filter on all fields, choose **All fields**. Otherwise, choose **Select fields** and then turn off the items you don't want to target.

The default is **All fields**.

- **Target visuals** – The dashboard widgets to target. To apply the filter to all of them, choose **All visuals**. Otherwise, choose **Select visuals** and then turn off the items you don't want to target. When you apply a filter action to other visuals, the effect is called *cascading filters*.

The default is **All visuals**.

A cascading filter applies all the visuals that are set up in the **Target visuals** section of a specific filter action. Amazon Quick Sight initially evaluates your visuals and preconfigures the settings for you. But you can change the defaults if you wish to do so. You can set up multiple cascading filters on multiple visuals in the same sheet or analysis. When you are using the analysis or dashboard, you can use multiple cascading filters at the same time, although you activate each of these one at a time.

A filter action requires at least one target visual, because a filter action requires a source and a target. To filter only the current visual, create a regular filter instead by choosing **Filter** at left.

- **Navigation action** settings include the following:
 - **Target sheet** – The sheet to target.
 - **Parameters** – The parameters to send to the target sheet. Choose the plus icon to add an existing parameter.
- **URL action** settings include the following:
 - **URL** – The URL to open. URL actions can be deep links into another application. Valid URL schemes include https, http, and mailto.

- **+ (Values)** – (Optional) The parameters to send to the target URL. Parameter names start with a \$. The parameters on both the sending and the receiving end must match in name and data type.
- **Open in** – Where to open the URL. You can choose **New browser tab**, **Same browser tab**, or **New browser window**.

Some types of actions enable you to include values from parameters or fields that are available in the visual or insight. You can type these in manually or choose **+** to select from a list. For the custom action to work, every field and parameter it references must be actively in use in the parent widget.

Use the following procedure to create, view, or edit a custom action in an analysis.

To create, view, or edit a custom action

1. With your analysis open, choose **Actions** from the **Menu options** dropdown in the upper right corner.

The existing actions, if any, display by activation type. To turn an existing action on or off, use the box to the right of the action's name.

2. (Optional) To edit or view an existing action, choose the menu icon next to the name of the action.

To edit the action, choose **Edit**.

To delete it, choose **Delete**.

3. To create a new action, choose either one of the following:

- The add icon near the **Actions** heading
- The **Define a custom action** button

4. For **Action name**, define an action name. To make the action name dynamic, use the plus icon to add parameter or field values.
5. For **Activation**, choose how the action runs.
6. For **Action type**, choose the action type you want to use.
7. For a **Filter action**, do the following:
 - a. For **Filter scope**, choose the scope of the filter.

- b. For **Target visuals**, choose how far the filter cascades
8. For a **Navigation action**, do the following:
 - a. For **Target sheet**, choose the target sheet.
 - b. For **Parameters**, choose the plus icon near the **Parameters** heading, select a parameter, and then choose a parameter value. You can choose all values, enter custom values, or select specific fields.
9. For a **URL action**, do the following:
 - a. For **URL**, enter the hyperlink.
 - b. Choose the plus icon near the **URL** heading. Then, add variables from the list.
 - c. For **Open in**, choose how to open the URL.
10. After you are finished with the action, choose one of the following at the bottom of the **Actions** panel (you might need to scroll down):
 - **Save** – Save your selections, and create the custom action.
 - **Close** – Close this custom action and discard your changes.
 - **Delete** – Delete this action.

Repairing custom actions

For a custom action to work, every field and parameter it references must be active in the parent widget. If a field is missing from the source widget, or if a parameter is missing from the analysis, the action for that field or parameter becomes unavailable. Menu actions are no longer included in the context menu. Select actions no longer respond to attempts to interact. However, in all other ways, the widget continues to function. No error displays to your users. You can fix broken filter actions and URL actions by adding the missing fields back to the broken visual or insight.

The following procedure explains how to fix an action that broke because someone removed a field or parameter without updating the action. These steps provide basic guidance how to fix this issue. However, use your own judgment on how or if you should make changes to the analysis. If you're not sure, it's better to ask a Amazon Quick Suite administrator for assistance before you change anything. For example, there might be a way to restore a previous version of the analysis, which might be safer if you aren't sure what happened to it.

To remove a field from a broken action

1. From the start page, choose **Analyses**. Then choose the analysis to fix.
2. Choose the visual or insight where the action no longer works. Make sure that it's highlighted on the sheet.
3. Choose **Actions** from the Menu options dropdown in the upper right corner.
4. Locate the action you want to fix and choose **Edit**.
5. If the action type is **Filter action**, and you see an error that says *the field used by this action was removed*, check the settings for **Filter scope**. **Selected fields** can only display fields that are in the visual. To disable selected fields that are removed, choose one of the following:
 - Change the **Filter scope** setting to **All fields**. Doing this enables the widget to filter on every field.
 - If you want to use a list of **Selected fields**, verify the list of fields. If you need to include another field, you need to add it to the visual first.
6. If the action type is **Navigation action**, follow the guidance on the error message, which reflects the type of change that caused the error.
7. If the action type is **URL action**, check the **URL** setting for variables marked with double angle brackets (<<FIELD-OR-\$PARAMETER>). Open the list of available variables by choosing the plus icon. Remove any fields or parameters that aren't in the list. Be sure you also remove the matching *URL parameter* and its separator (? for the first URL parameter, or & for subsequent parameters). The following examples show (in **bold**) which part is removed if you were removing the field named Product from the visual.

```
https://www.example.com/examplefunction?q=<<Product>
```

```
https://www.example.com/examplefunction?q=<<Product>&uact=<<$CSN>
```

```
https://www.example.com/examplefunction?pass=yes&q=<<Product>+<<City>&oq=<<Product>+<<City>&uact=<<$CSN>
```

Make sure to test the new URL.

8. (Optional) To delete the action, scroll to the end and choose **Delete**.
9. When you are finished, confirm your changes to the action. Scroll to the bottom of the **Action** pane and choose **Save**.

If the error also exists in an associated dashboard, share and publish the dashboard again to propagate the fix.

Understanding field mapping for custom actions in Amazon Quick Sight

Automated field mapping is based on identical fields. Fields with the same name and data type map automatically across datasets. Their field names and data types must be an exact match. This works similar to a join, except that it is automatically generated based on names and data types for every matching field. If you are missing fields, you can create them by using calculated fields in the dataset that's missing a field. If you don't want to have some of the fields mapped to each other, you can rename or remove them from the dataset.

It's important to make sure that all target fields are mapped if they are enabled for use with a filter action (in the **Filter scope**). Doing this allows filtering to apply automatically. If some target fields aren't mapped, the automatic filtering doesn't work.

Mapping is generated only when you create or save a custom action. So after every change that affects the mapping, make sure to return to it and save it again. When you create an action, mapping is based on the fields as they exist at that point. When you save an action, any mapped fields that you renamed since you created the custom action stay mapped. However, if you alter the data type of a mapped field, the mapping is removed.

If your mapping is missing some fields, you can do one of the following to fix it:

- Only target the mapped fields, by removing the unmapped fields from the **Filter scope**.
- Remove the visual in question from the target visuals.
- Create calculated fields to supply the missing fields for the mapping, and then save your custom action.
- Edit the dataset and rename the fields or change their data types, and then save your custom action.
- Edit the dataset and rename the fields or change their data types, and then resave your custom action.

Note

The information that displays on the mapping screen shows the configuration from the most recent time you saved it. To refresh or update the view, save the action again.

If you add or edit datasets, they aren't automatically mapped or remapped. This causes the filtering to work incorrectly. For example, suppose that you add a new dataset, then create visuals for it. The new visuals won't respond to filter actions, because there is no field mapping to connect them. When you make changes, remember to save your custom actions again to redo the field mappings.

If you remove a parameterized field or any other targeted field from the source visual, the action that uses it breaks. The action for the missing field either doesn't work when you select a data point, or it's hidden from the context menu.

For information about preparing your dataset for automated field mapping, see [Mapping fields](#).

Working with paginated reports in Amazon Quick Sight

With Amazon Quick Sight Paginated Reports, you can create, schedule, and share highly formatted multipage PDF reports. You can also schedule data exports as CSV files using Quick Sight's existing web interface. This unifies historically separate systems for dashboards and reports.

Report creators can use Quick Sight's browser-based authoring experience to connect to a broad range of supported data sources and create highly formatted reports. They can specify the exact page size, length, and arrangement of images, charts, and tables with pixel-level precision. Authors can then use Quick Sight's scheduling mechanisms to set up and schedule highly personalized report delivery to end users, or archive reports for future use.

Paginated reports are designed to be printed or distributed. Paginated report content is formatted to fit paper sizes and it displays all the data in a table and pivot table, even if the data spans multiple pages. Paginated reports are also referred to as pixel perfect because they are formatted for exact paper sizes and you can control page layout exactly. Each paginated report can generate a PDF of up to 1,000 pages.

Paginated reports provide all available data that is present when the report is published as a PDF or CSV. For example, let's say you have a table with 10,000 rows. A paginated report presents the entire report across multiple pages for readers to view in its entirety. If you include this same table in an interactive dashboard report, the generated PDF includes a snapshot of the table that fills in a single page that can be scrolled through. These customized reports can be sent out in email

bursts that generate up to thousands of personalized PDF or CSV reports to individual users and groups.

 Note

Paginated reports are not available in the eu-central-2 Europe (Zurich) region.

Topics

- [Getting started](#)
- [Creating reports from an analysis in Amazon Quick Sight](#)
- [Formatting reports in Amazon Quick Sight](#)
- [Consuming paginated reports in Amazon Quick Sight](#)
- [Unsubscribe from paginated reporting in Quick Sight](#)

Getting started

To get started working with Amazon Quick Sight paginated reports, first get the paginated reporting add-on for your Amazon Quick Suite account. Pricing for the add-on applies to your entire Quick Suite account and isn't specific to a Region. After you subscribe to Quick Suite reporting, authors can begin creating, scheduling, and sending paginated reports.

Get the Quick Suite paginated reports add-on

Before you can work with paginated reports in Amazon Quick Sight, you must add the **Paginated Reports add-on** to your Quick Suite subscription.

To get the paginated reporting add-on in Amazon Quick Sight

1. On the Quick Suite start page, choose your user name at the upper right, and then choose **Manage Quick Suite**.
2. Choose **Your subscriptions** on the left.
3. On the **Manage subscriptions** that opens, choose **Get Paginated Reports add-on**.
4. Choose the subscription plan that you want. You can choose between a monthly plan and an annual plan.
5. On the next page, review the Paginated Reports add-on pricing information and then choose **Confirm subscription**.

After you get the paginated reports add-on, it may take several minutes for your subscription to take effect. When your subscription takes effect, you will be able to begin creating paginated reports in Amazon Quick Sight.

Creating reports from an analysis in Amazon Quick Sight

Paginated reports are created at the sheet level of an analysis in Amazon Quick Sight. When you create a new analysis or a new sheet in an existing analysis, you choose whether to make the new sheet an **Interactive dashboard** or a **Paginated report**. This way, you can have analyses for interactive dashboards only, analyses for paginated reports only, or you can have an analysis that includes both interactive dashboards and paginated reports.

There are three ways to create a paginated report. You can create a new report from a new sheet in an analysis, you can duplicate an interactive sheet in a dashboard, or you can duplicate a paginated report that already exists. Use the procedures below to create a paginated report.

Creating reports from an analysis in Amazon Quick Sight

To create a paginated report from a new analysis

1. On the Quick Suite homepage, choose **Analyses**, and then choose **New analysis**.
2. Choose the dataset that you want to include in your new analysis, and then choose **USE IN ANALYSIS** in the top right.
3. In the **New sheet** pop-up that appears, choose **Paginated report**.
4. (Optional) Choose the paper size that you want for your paginated report. You can choose from the following options:
 - US letter (8.5 x 11 in.)
 - US legal (8.5 x 14 in.)
 - A0 (841 x 1189 mm)
 - A1 (594 x 841 mm)
 - A2 (420 x 594 mm)
 - A3 (297 x 420 mm)
 - A4 (210 x 297 mm)
 - A5 (148 x 210 mm)
 - Japan B4 (257 x 364 mm)

- Japan B5 (182 x 257 mm)

The default paper size is US letter (8.5 x 11 in.)

5. (Optional) Choose between a portrait and landscape arrangement for the sheet. The default option is portrait.
6. Choose **ADD**.

If you want to create a new paginated report in an existing analysis, choose the plus sign (+) icon to the right of the sheet tabs in your analysis and follow steps 3-6 from the preceding procedure.

Creating reports from an existing dashboard in Amazon Quick Sight

You can also create a paginated report by duplicating an interactive sheet and converting the duplicate sheet into a paginated report.

To create a paginated report from an interactive sheet

1. From the sheet that you want to duplicate in an analysis, choose the dropdown next to the name of the sheet that you want to convert.
2. Choose **Duplicate to report**.

You can convert an interactive sheet to a paginated report, but you can't convert a paginated report to an interactive sheet.

Duplicate an existing report in Amazon Quick Sight

This section will go over how to copy a report.

To copy a paginated report

1. From the sheet that you want to duplicate in an analysis, choose the dropdown next to the name of the sheet that you want to convert.
2. Choose **Duplicate**.

Formatting reports in Amazon Quick Sight

Use this section to learn how to format a paginated report in Amazon Quick Sight.

Topics

- [Working with sections](#)
- [Changing paper size, margins, and orientation](#)
- [Adding and removing page breaks to a report](#)
- [Adding and deleting visuals to a report](#)
- [Adding a text box to a report](#)
- [Setting up prompts for paginated reports](#)

Working with sections

A *section* is a container for different visuals that grow vertically to contain contents. Each section is rendered to completion, one after the other, accommodating configured page breaks and section settings. Headers and footers are special types of section that have a predefined size, location, and repetition throughout each page of a report.

Each section in a paginated report can be formatted independently from other sections in the report. Visuals can be dragged and dropped anywhere you want, similar to a free-form layout in an interactive sheet. Visuals can also be overlapped, resized, or brought forward or to the back of the section. Additionally, you can change the margins within a section to make the grouping of visuals stand out from the rest of the report.

Every report in Quick Sight needs at least one section. You can add multiple sections to group different sets of visuals together, or to control the rendering order for different groupings of visuals.

Each paginated report sheet supports up to 30 sections, including headers and footers.

Use the topics listed below to learn more about sections.

Topics

- [Adding, moving, and deleting sections](#)
- [Headers and footers](#)
- [Section padding](#)
- [Create repeating sections](#)

Adding, moving, and deleting sections

Add a new section

To add a new section to a paginated report, use the following procedure.

To add a new section to a paginated report

1. From the Quick Suite homepage, choose **Analyses** and then choose the analysis that contains the report that you want to add a section to.
2. Choose the sheet that contains the paginated report that you want to add a section to.
3. Choose the **ADD (+)** icon in the top left corner, and choose **Add section**.

You can also add a section by choosing the plus (+) icon at the bottom of an existing section and choosing **Add section**.

When you choose **Add section**, a new section is added to the bottom of the report.

You can't create a section inside of another section. If you select an existing section and then choose **Add section**, a new section will appear at the bottom of the report.

When you have multiple sections in a paginated report, they can be arranged in any order that you want.

Move a section

To move a section in a report

1. Choose the section that you want to move, and then choose the three-dot icon in the right corner to open the on-section menu.
2. Choose where you want to move your section. You can choose from the following options:
 - **Move section to top**
 - **Move section up**
 - **Move section down**
 - **Move section to bottom**

In some cases, you aren't able to select some of the preceding options. For example, if your section is already at the bottom of the report, you can't select **Move down** or **Move section to bottom**.

Sections are named according to their ascending order in the report. When you move a section up or down in a report, every section affected by the move is renamed according to the new ascending order.

When you delete a section from a paginated report, the names of the remaining sections may change depending on where the deleted section was located. For example, say you decide to delete Section 1. When you delete the section, the previous Section 2 will move up the report and become the new Section 1.

Delete a section

To delete a section from a report

1. Navigate to the section that you want to delete and choose the three-dot icon in the upper right corner to open the on-section menu.
2. Choose **Delete**.

Headers and footers

Headers and *Footers* are optional special sections located at the top and bottom of a paginated report. Headers and footers are commonly used to display basic information like the date the report was created or the page number. You can interact with headers and footers the same way you interact with a regular section in a report.

By default, every report in Amazon Quick Sight has a header and a footer. To remove the header or footer from your report, use the following procedure.

To remove a header or footer from a paginated report

1. In your paginated report, navigate to the header or footer that you want to delete and open the **On-section**.
2. Choose **Delete**.

When you delete a header or footer from your report, you are deleting the header or footer from every page of the report. You can't have a header or footer on some pages but not others.

If you have removed the header or footer from your report but want them to be visible again, use the following procedure.

To add a header or footer to a paginated report

1. Navigate to the paginated report that you want to add a header or footer to and choose **Insert** from the top menu.
2. Choose **Add header** or **Add footer**.

Section padding

You can use section padding to change the margins of an individual section in a paginated report. By default, all sections in a report use the page margins that are configured and applied to the entire report. You can also add section padding to a header or footer. With section padding, you can make a section stand out from other sections by creating another set of margins. Apply the new set of margins to the section on top of the page margins that the rest of the report uses.

To change the section padding of a section

1. Navigate to the section that you want to add section padding to and open the **Edit section**.
2. In the **Padding** section of the **Edit section**, enter the amount of padding you want in inches. You can customize the padding of every side of the section (top, bottom, left, and right).

You can't use section padding to decrease the margins of the section. For example, if the margins of the entire paginated report are 1 inch, you can only add to that value with section padding.

Create repeating sections

Use repeating sections to create duplicates of specific sections of a report to show one or more-dimension values. The data in the repeating section is sliced to match the dimensions of the section. Repeating sections can be replicated at scale to decrease the amount of time it takes to build reports.

Use the following procedures to create and configure a repeating section in a report.

To define a repeating section

1. Navigate to the section that you want to add a repeating behavior to and choose the **Edit repeating section** (triple panel).
2. In the **Edit section** pane that opens, choose **ADD DIMENSION**, and then choose the dimension that you want to add.

3. To add additional dimensions, repeat Step 2. You can add up to 3 dimensions in each repeating section configuration.

Considerations for repeating sections

The following limits apply to repeating sections.

- Insight visuals aren't supported for repeating sections.
- Repeating sections dimensions that are only from the last dataset that was selected for use in the analysis.

After you create a repeating section, you can define sorting and limits to the repeating section configuration. You can also use text boxes to add system parameters to repeating sections.

Define sorting in a repeating section

To define sorting in a repeating section

1. Navigate to the section that you want to add a repeating behavior to and choose the **Edit repeating section** (triple panel).
2. In the **Edit section** that opens, choose the ellipsis (three dots) next to the dimension that you want to change.
3. Navigate to the **Repeating** tab and choose the ellipsis (three dots) next to the dimension that you want to sort, and then choose **Edit**.
4. For **Sort by**, use the dropdown to choose the dimension that you want to sort by.
5. For **Aggregation** dropdown, choose the aggregation that you want to specify.
6. For **Sort order**, choose **Ascending** or **Descending**.

Define limits in a repeating section

You can set limits to show only a certain number of distinct dimension values for each dimension of a repeating section. You can choose to show between 1 and 1000 distinct values. The default limit is 50.

To define limits in a repeating section

1. Navigate to the section that you want to add a repeating behavior to and choose the **Edit repeating section** (triple panel).
2. In the **Edit section** that opens, choose the ellipsis (three dots) next to the dimension that you want to change.
3. For **Limit to**, enter the number of values that you want to limit the sorting to. You can enter a number between 1 and 1000.

Considerations for limits

The following limitations apply to limits in repeating sections.

- An *instance* is defined as a distinct value of a dimension or a unique combination of values of multiple dimensions.
- If the number of unique instances for a dimension in a repeating section exceeds 1000, the PDF report is NOT generated. If this occurs, try one of the following options.
 - Define a limit for your dimension.
 - Create a sheet level filter to restrict the dimension values.
 - Use row level security (RLS) to restrict the dimension values.
 - Apply dataset filters.

Add system parameters to repeating sections

You can use text boxes to add system parameters to your paginated report's repeating section. This makes it possible to access dimensions that have been used to configure repeating sections. Repeating sections and dimensions need to be configured before you can access the dimensions in a text box. The system parameters can only be used within a repeating section.

To add system parameters to a repeating section from a text box

1. Choose the text box visual that you want, and then choose the **System parameters** icon in the far right of the text box toolbar.
2. From the dropdown that appears, choose the parameter that you want.

Add page breaks to repeating sections

Similar to section page breaks, you can add page breaks to repeating sections.

To add a page break to a repeating section

1. Navigate to the section that contains the repeating behavior that you want to change and choose the **Edit repeating section** (triple panel) icon.
2. In the **Repeating** tab of the **Edit section** pane that appears, check the box titled **Page break after each instance**.

An instance is defined as a distinct value of a dimension or a unique combination of values of multidimensions. If you clear the **Page break after each instance** checkbox, the page break is removed.

Changing paper size, margins, and orientation

After you create a paginated report in Amazon Quick Sight, you can change the report format, orientation, and margins from the **Analysis settings** menu whenever you want.

To change the paper size of a paginated report

1. From the Quick Suite homepage, choose **Analyses**, and then choose the analysis that contains the paginated report that you want to change.
2. Choose **Sheets** in the file menu and select **Layout Settings**.
3. Open the **Paper size** dropdown menu and choose the paper size that you want. Choose from the following options:
 - US letter (8.5 x 11 in.)
 - US legal (8.5 x 14 in.)
 - A0 (841 x 1189 mm)
 - A1 (594 x 841 mm)
 - A2 (420 x 594 mm)
 - A3 (297 x 420 mm)
 - A4 (210 x 297 mm)
 - A5 (148 x 210 mm)
 - Japan B4 (257 x 364 mm)

- Japan B5 (182 x 257 mm)

4. Choose **Apply**.

To change the orientation of a report

1. From the Quick Suite homepage, choose **Analyses**, and then choose the analysis that contains the paginated report that you want to change.
2. Choose the **Settings** icon on the left.
3. Choose the orientation for your report, and then choose **Apply**.

To change the margins of a report

1. From the Quick Suite start page, choose **Analyses**, and then choose the analysis that contains the paginated report that you want to change.
2. Choose **Edit < Analysis Settings**.
3. Enter the margin values that you want your report to have, and then choose **Apply**.

Margin values are applied to every page of a paginated report. You can't set custom settings for specific pages in a report, but you can set custom margins for sections using section padding. For more information on section padding, see [Section padding](#). Margin values are expressed in inches. The default margins for all reports are 0.5 inches.

Adding and removing page breaks to a report

You can add page breaks between sections of a paginated report to organize the way data is rendered when the report is published by page. For example, let's say you have a report that contains two sections that are each 2.5 pages long. By default, Section 2 begins on the third page of the report directly following the end of Section 1. If you add a page break to the end of the Section 1, Section 2 begins on a new page, even if the last page of Section 1 only uses half of a page. This is useful when you don't want different sections to share pages, but you don't know how many pages each section will need.

To add or delete a page break

1. Select your section and choose the **Edit section** icon in the top left corner.
2. In the **Edit section** pane that opens on the left, select the **Page break after** check box.

3. Choose **Apply**.

When you check the **Page break after** box, a page break will appear at the end of the section. If you remove the check from the **Page break after** box, the page break is removed from the end of the section. Also, the proceeding section renders directly under the last page of the section, even if it causes the two sections to share a page.

You can also add or remove a page break from a report by choosing the plus (+) icon at the bottom of an existing section and choosing **Add page break** or **Remove page break**.

Adding and deleting visuals to a report

To add visuals to a section in a paginated report

1. In your paginated report, select the section that you want to add a visual to.
2. Choose the **ADD (+)** icon in **Visuals** pane.
3. Choose the visual type that you want to use in your report.

After you add a visual to a report, you can interact with it the same way you would if the visual was part of an interactive dashboard. You can drag and drop visuals anywhere you want, similar to a free-form layout in a Quick Sight interactive dashboard sheet. You can also overlap visuals, resize them, or bring them forward or to the back of the section. For more information on formatting visuals in Amazon Quick Sight, see [Formatting in Amazon Quick Suite](#).

To delete a visual

1. In the section that you want to delete a visual from, select the visual that you want to delete.
2. Choose the three-dot icon in the upper right corner of the visual to open the on-visual menu.
3. Choose **Delete**.

When you delete a visual from a section of a paginated report, you are only deleting that specific visual from the report. Any duplicate visuals that are located in different sections of the report will remain in the report.

Adding a text box to a report

You can add text boxes to your paginated reports to add context to your reports. Text box visuals can also be used boxes to add hyperlinks to external websites. To customize the font, font style,

text color, text spacing, text alignment, and text size, use the text box toolbar that appears when you select the visual.

To add a text box to a report

1. In your paginated report, select the section that you want to add a text box to.
2. Choose the **Text box** icon in the task bar.
3. The new text box appears in the section of the report that you selected.

To edit a text box, select the text box and begin typing what you want. A toolbar appears that you can use to make changes to the formatting and style of the text.

To delete a text box

1. In the section that you want to delete a text box from, select the text box that you want to delete.
2. Choose the three-dot icon in the upper right corner of the visual to open the on-text box menu.
3. Choose **Delete**.

Text box system parameters

Use text boxes to add system parameters to your paginated report's headers and footers. Text box system parameters appear on the far right side of the text box toolbar. You can add the following parameters to a header or footer of your report:

- Page numbers: The current page number of the report.
- Report print date: The date the report was generated.

To add a page number parameter to your text box, choose the number (#) icon on the far right side of the text box toolbar. To add a `PrintDate` parameter to your text box, choose the calendar icon on the far right side of the text box toolbar.

For more advanced parameter options, add an insight to your paginated report.

Setting up prompts for paginated reports

Amazon Quick Suite authors can create prompts on pixel-perfect reports to allow dashboard users to filter data in on-demand and scheduled reports. *Prompts* behave the same way a filter or control behaves in an interactive sheet.

To define a prompt in a paginated report

1. On a paginated sheet, define a filter control or a parameter control. For more information on filter controls to sheets, see [Adding filter controls to analysis sheets](#). For more information on parameter controls, see [Using a control with a parameter in Amazon Quick Suite](#).
2. In the new filter or parameter, choose the prompt values that you want. The new prompts are immediately reflected on the sheet.
3. To export the report with the new prompts, choose **File**, and then choose **Export to PDF**.

Prompts can't be moved to the sheet itself. Instead, they are displayed on the top panel.

After a prompt is created for a pixel-perfect report and is published as a dashboard, Quick Suite authors can use the new prompt to configure and schedule reports that are sent to Quick Suite dashboard viewers. Dashboard viewers can also use these prompts to create their own scheduled reports. For more information about reader generated reports, see [Creating a reader generated report in Amazon Quick Sight](#).

Consuming paginated reports in Amazon Quick Sight

When an Amazon Quick Suite author publishes and then sends a scheduled paginated report, Quick Suite will generate and save a snapshot of the report that is sent out. Whenever you go to view the paginated report's dashboard, you will see the generated snapshot of the most recently sent report. If you try to view your report's dashboard but you haven't sent an email report yet, you are prompted to schedule your first report to see the dashboard snapshot. For more information on scheduling an email report, see [Scheduling and sending Quick Sight reports by email](#).

If a Quick Suite author has set up a prompted report for a Quick Sight paginated report, Quick Suite readers can use the prompt to schedule their own on-demand reports for themselves. For more information about reader-generated reports, see [Creating a reader generated report in Amazon Quick Sight](#). For more information about prompts for paginated reports, see [Setting up prompts for paginated reports](#).

Users can't interact with a published paginated report the same way they can interact with a published interactive sheet. Unlike interactive sheets, paginated reports generate static snapshots of data that is presented in groups of visuals or text boxes. These static snapshots are generated at the time that the report is sent, so that the audience can see the latest version of the data in the report. Paginated reports are especially useful for generating invoices or weekly business reviews. Users can then compare the current paginated reports with reports that were generated in the past to better track their business data.

Viewing a report's snapshot history

Every time you send out a scheduled paginated report, Amazon Quick Suite saves a copy of the generated snapshot that is sent out for your reference. You can view these snapshots at any time in the Quick Suite console.

To view a report's snapshot history

1. From the Quick Suite homepage, choose **Dashboards**, and then choose the dashboard whose snapshot history you want to see.
2. Choose the **Scheduling** icon in the top right toolbar, and then choose **Recent snapshots**.
3. In the **Recent snapshots** pane that appears on the right, choose the snapshot to view, and then choose the download button next to the file that you want to download.

Unsubscribe from paginated reporting in Quick Sight

You can unsubscribe from Quick Sight paginated reporting at any time. Once you unsubscribe from paginated reporting, you will lose the ability to create and schedule paginated reports in Quick Sight. You are still able to access your existing paginated reports, but you won't be able to make changes or schedule new reports.

To unsubscribe from paginated reporting in Amazon Quick Sight

1. From any page in Quick Suite, choose your user name in the upper right, and choose **Manage Quick Suite**.
2. Choose **Your subscriptions** on the left.
3. On the **Your subscriptions** page, scroll down to the **Quick Suite Paginated Reports add-on** section and choose **Manage**.
4. Scroll down to your chosen subscription plan and choose **Cancel subscription**.

Working with items on sheets in Amazon Quick Sight analyses

Use this section to learn how to work with visuals and other items as you author sheets in Quick Sight

Topics

- [Adding visuals to Quick Sight analyses](#)
- [Using Topics on sheets in Amazon Quick Sight](#)
- [Visual types in Amazon Quick Sight](#)
- [Formatting in Amazon Quick Suite](#)
- [Customizing data presentation](#)

Adding visuals to Quick Sight analyses

A *visual* is a graphical representation of your data. You can create a wide variety of visuals in an analysis, using different datasets and visual types.

After you have created a visual, you can modify it in a range of ways to customize it to your needs. Possible customizations include changing what fields map to visual elements, changing the visual type, sorting visual data, or applying a filter.

Quick Sight supports up to 50 datasets in a single analysis, and up to 50 visuals in a single sheet, and a limit of 20 sheets per analysis.

You can create a visual in several ways. You can select the fields that you want and use AutoGraph to let Amazon Quick Sight determine the most appropriate visual type. Or you can choose a specific visual type and choose fields to populate it. If you aren't sure what questions your data can answer for you, you can choose **Suggested** on the tool bar and choose a visual that Amazon Quick Sight suggests. Suggested visuals are ones that we think are of interest, based on a preliminary examination of your data. For more information about AutoGraph, see [Using AutoGraph](#).

You can add more visuals to the workspace by choosing **Add**, then **Add visual**. Visuals created after June 21, 2018, are smaller in size, fitting two on each row. You can resize the visuals and drag them to rearrange them.

To create a useful visual, it helps to know what question you are trying to answer as specifically as possible. It also helps to use the smallest dataset that can answer that question. Doing so helps you create simpler visuals that are easier to analyze.

Fields as dimensions and measures

In the **Visuals** pane, dimension fields have blue icons and measure fields have orange icons. *Dimensions* are text or date fields that can be items, like products. Or they can be attributes that are related to measures and can be used to partition them, like sales date for sales figures. *Measures* are numeric values that you use for measurement, comparison, and aggregation. You typically use a combination of dimension and measure fields to produce a visual, for example sales totals (a measure) by sales date (a dimension). For more information about the types of fields expected by the different visual types, see the specific visual type topics in the [Visual types in Amazon Quick Sight](#) section. For more information about changing a field's measure or dimension setting, see [Setting fields as a dimensions or measures](#).

Field limitations

You can only use one date field per visual. This limitation applies to all visual types.

You can't use the same field for more than one dimension field well or drop target on a visual. For more information about how expected field type is indicated by field wells and drop targets, see [Using visual field controls](#).

Searching for fields

If you have a long field list in the **Fields list** pane, you can search to locate a specific field. To do so, choose the search icon at the top of the **Data** pane and then enter a search term into the search box. Any field whose name contains the search term is shown. Search is case-insensitive and wildcards aren't supported. Choose the cancel icon (X) to the right of the search box to return to viewing all fields.

Adding a visual

Use the following procedure to create a new visual.

To create a new visual

1. Open the [Quick Suite console](#).
2. On the Quick Suite homepage, choose the analysis that you want to add a visual to.
3. On the analysis page, choose the dataset that you want to use from the dataset list at the top of the **Data** pane. For more information, see [Adding a dataset to an analysis](#).
4. Open the **Visualize** pane, choose **Add**, and then choose **Add visual**.

A new, blank visual is created and receives focus.

5. Use one of the following options:

- Choose the fields to use from the **Data** pane at left. If the fields aren't visible, choose **Visualize** to display it. Amazon Quick Sight creates the visual, using the visual type it determines is most compatible with the data you selected.
- Choose the dropdown arrow next to the **ADD** button to choose a visual type. After the visual is created, choose the fields that you want to populate it.
 1. Choose the icon of a visual type from the **Visual types** pane.

The field wells display the fields that are visualized.

2. From the **Data** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the color of the target field well. If you choose to use a dimension field to populate a **Value** field well, the **Count** aggregate function is automatically applied to it to create a numeric value.

Amazon Quick Sight creates the visual using the visual type you selected.

- Create a visual using a suggestion.

On the tool bar, choose **Suggested**, then choose a suggested visual.

Importing Amazon Quick Sight visuals to an analysis

Quick Sight authors can import Quick Sight visuals from one analysis or dashboard to a new analysis that has access privileges. When you import a visual from a Quick Sight analysis or dashboard to another Quick Sight analysis, the following dependencies are imported along with the visual.

- Datasets associated with the visual
- All parameters that are configured to the visual
- Calculated fields that are configured to the visual
- Filter definitions
- Visual properties
- Conditional formatting rules

Use the following sections to learn more about importing Quick Sight visuals.

Topics

- [Considerations](#)
- [Import a visual](#)

Considerations

Before you import a visual, review the following limitations.

- The Quick Sight author that wants to import a visual must have ownership privileges to the analysis that they want to import the visual to
- Filter controls can't be imported
- Importing visuals from multiple sheets at a time is not supported
- Some user configurations including filter configurations that are maintained against bookmarks and alerts are not supported

Import a visual

Use the following procedure to import a visual from a source dashboard or analysis to a different analysis.

1. Open the [Quick Suite console](#).
2. Open the analysis that you want to import a visual to.
3. Choose **File**, and then choose **Import**. Alternatively, you can choose the **Import** icon in the **ADD** toolbar.
4. The **Asset explorer** modal opens. A list of all eligible source analyses and dashboards that you can access are displayed. Choose the artifact that you want to import a visual from, and then choose **LOAD**. Alternatively, enter the name of the source artifact that contains the visual that you want to import in the **Find source to insert** search bar. Choose the artifact that you want, and then choose **LOAD**.
5. In the **Select visuals to import** page that opens, choose the sheet that contains the visuals that you want to import, and then choose the visuals that you want to import. You can only import visuals from one sheet at a time. When you have chosen all visuals that you want to import, choose **IMPORT**.

After a successful import job, the imported visuals are added to the destination analysis. The imported visuals retain the original properties that were configured to them in the source dashboard or analysis. Imported visuals inherit the theme-level properties from the theme that is applied to the destination analysis.

Duplicating Quick Sight visuals

You can duplicate a visual to make a new copy of it on the same sheet or on a different sheet.

To duplicate a visual, on the **v**-shaped on-visual menu, choose **Duplicate visual to**, then choose the sheet where you want the visual to appear. The display automatically shows you the duplicated visual.

Duplicated visuals keep all the same filters and settings as the source visual. However, if you duplicate a visual onto a different sheet, all of its copied filters apply to the duplicate only. All copied filters are scoped down to apply only to that visual. If you want the filters to apply to more visuals on the new sheet, edit the filter and change the setting.

Parameters and controls apply to all sheets. To make parameter controls work with a visual that you duplicate to a different sheet, add filters on the target sheet and connect them to the parameter. To do this, choose **Custom filter** as the filter type.

Renaming Amazon Quick Sight visuals

Use the following procedure to rename a visual.

To rename a visual

1. On the analysis page, choose the visual that you want to rename.
2. Select the visual name at the top left of the visual and enter a new name.
3. Press **Enter** or click outside of the visual name field to save the new name.

Viewing visual data in Amazon Quick Sight

Amazon Quick Sight offers a variety of ways to see the details of the data being displayed in a visual. The axes or rows and columns of the visual (depending on the visual type) have labels. Hovering over any graphical element in a visual displays the data associated with that element. Some visual types use visual cues to emphasize the element that you are hovering over and make it easier to differentiate. For example, the visual type might change the color of the element or highlight it.

Use the following sections to learn more about viewing data in visuals.

Topics

- [Viewing visual details](#)
- [Scrolling through visual data](#)
- [Focusing on visual elements](#)
- [Excluding visual elements](#)
- [Searching for specific values in your data in Quick Sight](#)

Viewing visual details

When viewing a visual, you can hover your cursor over any graphical element to get details on that element. For example, when you hover over a single bar on a bar chart, information about that specific bar displays in a tooltip.

Hovering your cursor over a single data point on a scatter plot also displays information about that specific data point.

You can customize the information that appears when you hover your cursor over data in a chart. For more information, see [Tooltips](#).

Scrolling through visual data

For bar charts, line charts, and pivot tables, the content of the visual can be larger than the size that you want the visual to be.

In these cases, scrub bars appear so you can either reduce the data that is displayed or scrub through it. This process is similar to the way that you can scrub through a video.

To reduce the length of the scrub bar, hover over one end of it until the cursor changes shape. Then drag the widget to make the scrub bar larger or smaller. To scroll through the data, click and hold the scrub bar and slide it toward the end that you want to see.

Focusing on visual elements

When viewing visuals, you can choose data that you want to focus on or exclude. To perform this choice, choose an element such as a bar or bubble, or a row or column header.

Focusing on or excluding data causes Quick Sight to create a filter and show only the data that you selected.

To remove the filter, choose **Filters** at left and then disable or delete the filter. You can also use **Undo** to remove a filter.

If your visual has a legend that shows categories (dimensions), you can click on the values in the legend to see a menu of available actions. For example, suppose that your bar chart has a field in the **Color** or **Group/Color** field well. The bar chart menu displays the actions that you can choose by clicking or right-clicking on a bar, such as the following:

- Focusing on, or excluding, visual elements
- Changing colors of visual elements
- Drilling down into a hierarchy
- Custom actions activated from the menu, including filtering or URL actions

Excluding visual elements

When viewing visuals, you can choose an element on the visual, and then choose to focus on the element. Elements to focus on can include, for example, a bar or bubble, or a row or column header in the case of a pivot table. The exception is that you can't exclude elements that are mapped to date fields. You can exclude multiple elements on a single chart.

Excluding the element creates a filter that removes only that element from the visual.

To see the excluded element again, you can either choose **Undo** on the application bar, or you can disable or delete the filter.

For more information about filters, see [Filtering data in Amazon Quick Sight](#).

Searching for specific values in your data in Quick Sight

When filtering your visual data, previewing anomalies, or using list or dropdown controls in a dashboard, you can quickly search for values that interest you.

You can search for specific values or all values that contain a specific search query. For example, searching for *al* in a list of U.S. states returns **Alabama**, **Alaska**, and **California**.

You can also use wildcard search to search for all values that match a specific character pattern. For example, you can search for all U.S. states that end with the letters *ia* and narrow the results down to California, Georgia, Pennsylvania, Virginia, and West Virginia.

To search for values in a filter or control, enter a search query in the search bar.

Using wildcard search

The following wildcard characters can be used to find values in Quick Sight filters, list and dropdown controls, and anomaly previews.

- ***** - Use an asterisk symbol to search for values that match zero to many characters in a specific position.
- **?** - Use a question mark to match a single character in a specific position.
- **** - Use a backslash to escape the *****, **?**, or **** wildcard characters and search for them in your query. For example, you can search for phrases that end with a question mark.

Following are examples of how supported wildcard characters can be used in a Quick Sight search query.

- **a1** - This query searches for all values with **a1** and returns Alabama, Alaska, and California.
- **a1*** - This query searches for all values that begin with **a1** and end with zero to multiple characters. It returns Alabama, and Alaska in a list of U.S. states.
- ***ia** - This query searches for all values that begin with zero to multiple characters and end with letters **ia**. It returns California, Georgia, Pennsylvania, Virginia, and West Virginia.
- ***a1*** - This query searches for all values with zero to multiple characters before and after the letters **a1**. It returns Alabama, Alaska, and California.
- **a?a?a?a** - This query searches for all values with a single character in the exact positions between the **a** letters. It returns Alabama.
- **a?a*a** - This query searches for all values with a single character between the first two **a** letters and multiple characters between the second two **a** letters. It returns Alabama and Alaska.
- **How*\?** - This query searches for values that begin with **How**, followed by zero to multiple characters, and end with a question mark. The backslash (****) in this query informs Quick Sight to search for question marks in each value, rather than use the question mark symbol as a wildcard character. This query returns the questions, How are you? and, How is this possible?
- ****** - This query searches for values that begin with an asterisk and are followed by zero to multiple characters. The backslash (****) in this query informs Quick Sight to search for an actual asterisk in the values, rather than use the asterisk symbol as a wildcard character. This query returns values such as ***all**, ***above**, and ***below**.

- `\\*` - This query searches for values with a backslash, followed by zero to multiple characters. The first backslash (`\`) in this query informs Quick Sight to search for the second backslash (`\`) in each value, rather than use the backslash symbol as a wildcard character. This query returns results such as `\Home`.
- `???` - This query searches for values that contain three characters. It returns values such as `ant`, `bug`, and `car`.

Exporting data from visuals

Note

Export files can directly return information from the dataset import. This makes the files vulnerable to CSV injection if the imported data contains formulas or commands. For this reason, export files can prompt security warnings. To avoid malicious activity, turn off links and macros when reading exported files.

Using the Amazon Quick Suite console, you can export data from any type of chart or graph. The export contains only the data in the fields that are currently visible in the selected visualization. Any data that is filtered out is excluded from the export file. You can export data into the following formats:

- A text file containing comma-separated values (CSV), available for all visual types.
- A Microsoft Excel workbook file (.xlsx), available for pivot tables and table charts only.

The following rules apply:

- Exported files are downloaded to the default download directory configured in the browser that you're currently using.
- The downloaded file is named for the visualization that you exported it from. To make the file name unique, it has a sequential timestamp (a Unix epoch data type).
- Default limit for export to CSV format: 500 MB or 1M rows whichever comes first
- Default limit for export to Excel format:
 - from Pivot Table visual 400K cells or 50K rows
 - from Table visual 800K cells or 100K rows

Note

With a subscription to Paginated Reporting, you are able to [schedule the export of visuals in CSV and Excel formats](#) and export up to 3M rows (CSV) and 16M cells (Excel).

- You can't export data from an insight, because insights consume the data, but don't contain the data.
- Quick Sight doesn't support exporting data from more than a single visualization at a time. To export data from additional visuals in the same analysis or dashboard, repeat this process for each visual. To export all the data from a dashboard or analysis, you need to connect to the original data source using valid credentials and a tool that you can use to extract data.

Use the following procedure to export data from a visualization in Amazon Quick Sight. Before you begin, open the analysis or dashboard that contains the data that you want to export.

To export data from a visualization

1. Choose the visualization that you want to export. Make sure that it is selected and highlighted.
2. At top right on the visual, open the menu and choose one of the following:
 - To export to CSV, choose **Export to CSV**.
 - To export to XSLX, choose **Export to Excel**. This option is available only for pivot tables and table charts.
3. Depending on your browser settings, one of the following happens:
 - The file automatically goes to your default **Download** location.
 - A dialog box appears so you can choose a file name and location.
 - A dialog box appears so you can choose to open the file with the default software or to save to.

Refreshing visuals in Quick Sight

When you work in an Quick Sight analysis or dashboard, visuals refresh and reload when you change something that affects them, such as updating a parameter or filter control. If you switch to a new sheet after a parameter or filter changes, only the visuals affected by the change refresh

on the new sheet. Otherwise, visuals update every 30 minutes when you switch sheets. This is the default behavior for all analyses and dashboards.

If you want to refresh all visuals when you switch sheets, regardless of a change, you can do so for each analysis that you create.

To refresh all visuals each time that you switch sheets in an analysis

1. In Amazon Quick Suite, open the analysis.
2. In the analysis, choose **Edit > Analysis Settings**.
3. In the **Analysis Settings** pane that opens, for **Refresh Options**, toggle on **Reload visuals each time I switch sheets**.
4. Choose **Apply**.

Deleting Amazon Quick Sight visuals

Use the following procedure to delete a visual.

To delete a visual

1. On the analysis page, choose the visual that you want to delete.
2. Choose the on-visual menu at the upper-right corner of the visual, and then choose **Delete**.

Using Topics on sheets in Amazon Quick Sight

Amazon Quick Sight provides a guided workflow for creating topics. You can step out of the guided workflow and come back to it later, without disrupting your work.

By enabling one or more Quick Sight topics in your analysis workspace, you activate the ML-powered automated data prep , which speeds Natural Language (NL) topic creation. Automated data prep automatically selects high value fields, based on how they are used and on common Q&A needs. It automatically chooses user-friendly field names and synonyms, based on terms from existing analyses and on common dictionaries. It also automatically formats data, so it's immediately useful when presented.

Automated data prep binds the topic to your analysis and prepares an index for searching in natural language. A blue dot denotes this binding. Dashboard users find that the new Amazon Quick Sight topic is automatically selected, making it easier for them to query the dataset.

The following rules apply to working with topics:

- You must be an owner of the underlying dataset before you can create a topic using that dataset or an analysis that uses that dataset.
- You must be an owner of a topic before you can link the existing topic to an analysis.

To enable a topic

1. Open the analysis that you want to use with automated data prep .
2. On the top navigation bar, choose the topic icon.
3. Choose one of the following:
 - To activate a new topic, select **Create new topic** and enter a topic title and optional description.
 - To activate an existing topic, select **Update existing topic** and choose the topic from the list.
4. Choose **ENABLE TOPIC** to confirm your choice.
5. When the topic is finished processing, you can use what it learned from the analysis to ask questions in natural language.

Now, when users navigate to the dashboard, the linked topic is automatically selected in the search bar.

After a topic is linked to an analysis, further updates to the analysis are not automatically synced to the topic. Authors need to manage updating topics manually from the **Topics** page.

When you enable a topic for an analysis or dashboard, you are starting a process where automated data prep learns from how you analyze your data. Ask it questions, and provide feedback and further information by following the screen prompts. The more you interact with the topic, the better prepared it becomes to answer your questions.

To learn more, see <https://docs.amazonaws.cn/quicksight/latest/user/quicksight-q-starting-from-sheets.html>.

Visual types in Amazon Quick Sight

Amazon Quick Sight offers a range of visual types that you can use to display your data. Use the topics in this section to learn more about the capabilities of each visual type.

Topics

- [Measures and dimensions](#)
- [Display limits](#)
- [Hiding or displaying the other category](#)
- [Customizing the number of data points to display](#)
- [Using AutoGraph](#)
- [Using bar charts](#)
- [Using box plots](#)
- [Using combo charts](#)
- [Using custom visual content](#)
- [Using donut charts](#)
- [Using funnel charts](#)
- [Using gauge charts](#)
- [Using heat maps](#)
- [Using Highcharts](#)
- [Using histograms](#)
- [Using image components](#)
- [Using KPIs](#)
- [Using layer maps](#)
- [Using line charts](#)
- [Creating maps and geospatial charts](#)
- [Using small multiples](#)
- [Using pie charts](#)
- [Using pivot tables](#)
- [Using radar charts](#)
- [Using Sankey diagrams](#)
- [Using scatter plots](#)
- [Using tables as visuals](#)
- [Using text boxes](#)
- [Using tree maps](#)

- [Using waterfall charts](#)
- [Using word clouds](#)

Measures and dimensions

We use the term *measure* to refer to numeric values that you use for measurement, comparison, and aggregation in visuals. A measure can be either a numeric field, like product cost, or a numeric aggregate on a field of any data type, like count of transaction IDs.

We use the term *dimension* or *category* to refer to text or date fields that can be items, like products, or attributes that are related to measures and can be used to partition them. Examples are sales date for sales figures or product manufacturer for customer satisfaction numbers. Amazon Quick Sight automatically identifies a field as a measure or a dimension based on its data type.

Numeric fields can act as dimensions, for example ZIP codes and most ID numbers. It's helpful to give such fields a string data type during data preparation. This way, Amazon Quick Sight understands that they are to be treated as dimensions and are not useful for performing mathematical calculations.

You can change whether a field is displayed as a dimension or measure on an analysis-by-analysis basis instead. For more information, see [Fields as dimensions and measures](#).

Display limits

All visual types limit the number of data points they display, so that the visual elements (like lines, bars, or bubbles) are still easy to view and analyze. The visual selects the first n number of rows for display up to the limit for that visual type. The selection is either according to sort order, if one has been applied, or in default order otherwise.

The number of data points supported varies by visual type. To learn more about display limits for a particular visual type, see the topic for that type.

The visual title identifies the number of data points displayed if you have reached the display limit for that visual type. If you have a large dataset and want to avoid running into the visual display limit, use one or more filters to reduce the amount of data displayed. For more information about using filters with visuals, see [Filtering data in Amazon Quick Sight](#).

For dashboards and analyses, Amazon Quick Sight supports the following:

- 50 datasets per dashboard
- 20 sheets per dashboard
- 30 visualization objects per sheet

You can also choose to limit how many data points you want to display in your visual, before they are added to the **other** category. This category contains the aggregated data for all the data beyond the cutoff limit for the visual type you are using—either the one you impose, or the one based on display limits. You can use the on-visual menu to choose whether to display the **other** category. The **other** category doesn't show on scatter plots, heat maps, maps, tables (tabular reports), or key performance indicators (KPIs). It also doesn't show on line charts when the x-axis is a date. Drilling down into the **other** category is not supported.

Hiding or displaying the other category

Use the following procedure to hide or display the "other" category.

To hide or display the "other" category

1. On the analysis page, choose the visual that you want to modify.
2. Choose the on-visual menu at the upper-right corner of the visual, and then choose **Hide "other" category** or **Show "other" category**, as appropriate.

Customizing the number of data points to display

You can choose the number of data points to display on the main axis of some visuals. After this number is displayed in the chart, any additional data points are included in the "other" category. For example, if you choose to include 10 data points out of 200, 10 display in the chart and 190 become part of the "other" category.

To find this setting, choose the v-shaped on-visual menu, then choose **Format visual**. You can use the following table to determine which field well contains the data point setting and what number of data points the visual type displays by default.

Visual type	Where to find the data point setting	Default number of data points
Bar chart, horizontal	Y-axis – Number of data points displayed	10,000

Visual type	Where to find the data point setting	Default number of data points
Bar chart, vertical	X-axis – Number of data points displayed	10,000
Combo chart	X-axis – Number of data points displayed	2,500
Heat map	Rows – Number of rows displayed Columns – Number of columns displayed	100
Line chart	X-axis – Number of data points displayed	10,000
Pie chart	Group/Color – Number of slices displayed	20
Tree map	Group by – Number of squares displayed	100

Using AutoGraph

AutoGraph isn't a visual type itself, but instead lets you tell Amazon Quick Suite to choose the visual type for you. When you create a visual by choosing AutoGraph and then selecting fields, Amazon Quick Suite uses the most appropriate visual type for the number and data types of the fields you select.

Creating a visual using AutoGraph

Use the following procedure to create a visual using AutoGraph.

To create a visual using AutoGraph

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the AutoGraph icon.
4. On the **Fields list** pane, choose the fields that you want to use.

Using bar charts

Amazon Quick Suite supports the following types of bar charts, with either horizontal or vertical orientation:

- **Single-measure** – A *single-measure bar chart* shows values for a single measure for a dimension.
- **Multi-measure** – A *multi-measure bar chart* shows values for multiple measure for a dimension.
- **Clustered** – A *clustered bar chart* shows values for a single measure for a dimension, grouped by another dimension.
- **Stacked** – A *stacked bar chart* is similar to a clustered bar chart in that it displays a measure for two dimensions. However, instead of clustering bars for each child dimension by the parent dimension, it displays one bar per parent dimension. It uses color blocks within the bars to show the relative values of each item in the child dimension. The color blocks reflect the value of each item in the child dimension relative to the total for the measure. A stacked bar chart uses a scale based on the maximum value for the selected measure.
- **Stacked 100 percent** – A *stacked 100 percent bar chart* is similar to a stacked bar chart. However, in a stacked 100 percent bar chart, the color blocks reflect the percentage of each item in the child dimension, out of 100 percent.

Bar charts show up to 10,000 data points on the axis for visuals that don't use group or color. For visuals that do use group or color, they show up to 50 data points on the axis and up to 50 data points for group or color. For more information about how Amazon Quick Suite handles data that falls outside display limits, see [Display limits](#).

Creating single-measure bar charts

Use the following procedure to create a single-measure bar chart.

To create a single-measure bar chart

1. On the analysis page, choose **Visualize** on the toolbar at left.
2. On the application bar at upper left, choose **Add**, and then choose **Add visual**.
3. On the **Visual types** pane, choose the **Horizontal bar chart** or **Vertical bar chart** icon.
4. From the **Fields list** pane, drag a dimension to the **X-axis** or **Y-axis** field well.
5. From the **Fields list** pane, drag a measure to the **Value** field well.

Creating multi-measure bar charts

Use the following procedure to create a multi-measure bar chart.

To create a multi-measure bar chart

1. On the analysis page, choose **Visualize** on the toolbar at left.
2. On the application bar at upper-left, choose **Add**, and then choose **Add visual**.
3. On the **Visual types** pane, choose the **Horizontal bar chart** or **Vertical bar chart** icon.
4. From the **Fields list** pane, drag a dimension to the **X-axis** or **Y-axis** field well.
5. From the **Fields list** pane, drag two or more measures to the **Value** field well.

Creating clustered bar charts

Use the following procedure to create a clustered bar chart.

To create a clustered bar chart

1. On the analysis page, choose **Visualize** on the toolbar at left.
2. On the application bar at upper left, choose **Add**, and then choose **Add visual**.
3. On the **Visual types** pane, choose the **Horizontal bar chart** or **Vertical bar chart** icon.
4. From the **Fields list** pane, drag a dimension to the **X-axis** or **Y-axis** field well.
5. From the **Fields list** pane, drag a measure to the **Value** field well.
6. From the **Fields list** pane, drag a dimension to the **Group/Color** field well.

Creating stacked bar charts

Use the following procedure to create a stacked bar chart.

To create a stacked bar chart

1. On the analysis page, choose **Visualize** on the toolbar at left.
2. On the application bar at upper-left, choose **Add**, and then choose **Add visual**.
3. On the **Visual types** pane, choose the **Horizontal stacked bar chart** or **Vertical stacked bar chart** icon.
4. From the **Fields list** pane, drag a dimension to the **X-axis** or **Y-axis** field well.
5. From the **Fields list** pane, drag a dimension to the **Group/Color** field well.
6. From the **Fields list** pane, drag a measure to the **Value** field well.
7. (Optional) Add data labels and show totals:

- a. On the menu in the upper-right corner of the visual, choose the **Format visual** icon.
- b. In the **Visual** pane, choose **Data labels**.
- c. Toggle the switch to display data labels.

Labels for each measure value appear in the chart and the option to show totals appears in the pane.

- d. Check **Show totals**.

Totals appear for each bar in the chart.

Creating stacked 100 percent bar charts

Use the following procedure to create a stacked 100 percent bar chart.

To create a stacked 100 percent bar chart

- 1. On the analysis page, choose **Visualize** on the toolbar at left.
- 2. On the application bar at upper-left, choose **Add**, and then choose **Add visual**.
- 3. On the **Visual types** pane, choose the **Horizontal stacked 100% bar chart** or **Vertical stacked 100% bar chart** icon.
- 4. From the **Fields list** pane, drag a dimension to the **X-axis** or **Y-axis** field well.
- 5. From the **Fields list** pane, drag two or more measures to the **Value** field well.

Bar chart features

To understand the features supported by bar charts, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes, with exceptions	Multi-measure and clustered bar charts display a legend, while single-measure horizontal bar charts don't.	Legends on visual types in Quick Suite
Changing the title display	Yes		Titles and subtitles on visual types in Quick Suite

Feature	Supported?	Comments	For more information
Changing the axis range	Yes		Range and scale on visual types in Quick Suite
Showing or hiding axis lines, grid lines, axis labels, and axis sort icons	Yes		Axes and grid lines on visual types in Quick Suite
Changing the visual colors	Yes		Colors in visual types in Quick Suite
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude any bar on the chart, except when you are using a date field as the dimension for the axis. In that case, you can only focus on a bar, not exclude it.	Focusing on visual elements Excluding visual elements
Sorting	Yes	You can sort on the fields you choose for the axis and the values.	Sorting visual data in Amazon Quick Suite
Performing field aggregation	Yes	You must apply aggregation to the field or fields you choose for the value, and can't apply aggregation to the fields you choose for the axis or group/ color.	Changing field aggregation
Adding drill-downs	Yes	You can add drill-down levels to the axis and Group/Color field wells.	Adding drill-downs to visual data in Quick Sight

Feature	Supported?	Comments	For more information
Showing data labels	Yes		Data labels on visual types in Quick Suite
Showing stacked bar chart totals	Yes	Showing totals in a stacked bar chart is only available when you choose to show data labels.	Stacked bar charts

Using box plots

Box plots, also known as box and whisker plots, display data pooled from multiple sources into one visual, helping you make data-driven decisions. Use a box plot to visualize how data is distributed across an axis or over time, for example flights delayed over a 7-day time period. Typically, a box plot details information in quarters:

- **Minimum** – The lowest data point excluding outliers.
- **Maximum** – The highest data point excluding outliers.
- **Median** – The middle value of the dataset.
- **First Quartile** – The middle value between the smallest number and the median of the dataset. The first quartile doesn't include the minimum or the median.
- **Third Quartile** – The middle value between the largest number and the median of the dataset. The third quartile doesn't include the maximum or the median.

Outliers are extreme data points that aren't included in the calculation of a box plot's key values. Because outliers are calculated separately, their data points don't appear immediately after a box plot is created. Box plots display up to 10,000 data points. If a dataset contains more than 10,000 data points, a warning appears at the upper-right corner of the visual.

Box plots support up to five metrics and one group-by, but don't render if duplicate metrics are supplied.

Box plots support some calculated fields, but not all. Any calculated field that uses a window function, for example `avgOver1`, results in a SQL error.

Box plot visuals aren't compatible with MySQL 5.3 and earlier.

To create a basic box plot visual

1. Sign in to Amazon Quick Suite at <https://quicksight.aws.amazon.com/>.
2. Open Quick Suite and choose **Analyses** on the navigation pane at left.
3. Choose one of the following:
 - To create a new analysis, choose **New analysis** at upper right. For more information, see [Starting an analysis in Quick Sight](#).
 - To use an existing analysis, choose the analyses that you want to edit.
4. Choose **Add, Add visual**.
5. At lower left, choose the box plot icon from **Visual types**.
6. On the **Fields list** pane, choose the fields that you want to use for the appropriate field wells. Box plots require at least one unique measure field.
7. (Optional) Add drill-down layers by dragging one or more additional fields to the **Group/Color** field well. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

To understand the features supported by box plots, see [Analytics formatting per type in Quick Suite](#). For customization options, see [Formatting in Amazon Quick Suite](#).

Using combo charts

Using a combo chart, you can create one visualization that shows two different types of data, for example trends and categories. Combo charts are also known as line and column (bar) charts, because they combine a line chart with a bar chart. Bar charts are useful for comparing categories. Both bar charts and line charts are useful for displaying changes over time, although bar charts should show a greater difference between changes.

Amazon Quick Suite supports the following types of combo charts:

- **Clustered bar combo charts** – display sets of single-color bars where each set represents a parent dimension and each bar represents a child dimension. Use this chart to make it easy to determine values for each bar.
- **Stacked bar combo charts** – display multi-color bars where each bar represents a parent dimension and each color represents a child dimension. Use this chart to make it easy to see relationships between child dimensions within a parent dimension. This chart shows the total value for the parent dimension and how each child adds to the total value. To determine the

value for each child dimension, the chart reader must compare the size of the color section to the data labels for that axis.

Both types of combo chart require only one dimension on the **X axis**, but are usually more effective when also displaying at least one measure under **Lines**.

Use a combo chart only if you want to show a relationship between the bars and the lines. A good rule of thumb is that if you need to explain how the two chart types relate, you should probably use two separate charts instead.

Because each chart works differently, it can be helpful to understand the following points before you begin:

- The data points in each series render on different scales. Combo charts use a scale based on the maximum value for the selected measure.
- The distance between the numbers on the axis won't match between the lines and bars, even if you select the same scale for each chart type.
- For clarity, try to use different units for the measure in each data series.

The combo chart is like using two different types of visualization at the same time. Make sure that the data in the bars (or columns) directly relates to the data in the line or lines. This relationship is not technically enforced by the tool, so it's essential that you determine this relationship yourself. Without some relation between the lines and bars, the visual loses meaning.

You can use the combo chart visual type to create a single-measure or single-line chart. A single-measure combo chart shows one measure for one dimension.

To create a multi-measure chart, you can choose to add multiple lines, or multiple bars. A multi-measure bar chart shows two or more measures for one dimension. You can group the bars in clusters, or stack them.

For the bars, use a dimension for the axis and a measure for the value. The dimension is typically a text field that is related to the measure in some way and can be used to segment it to see more detailed information. Each bar in the chart represents a measure value for an item in the dimension you chose.

Bars and lines show up to 2,500 data points on the axis for visuals that don't use group or color. For visuals that do use group or color, bars show up to 50 data points on the axis and up to 50 data

points for group or color, while lines show 200 data points on the axis and up to 25 data points for group or color. For more information about how Amazon Quick Suite handles data that falls outside display limits, see [Display limits](#).

Combo chart features

To understand the features supported by combo charts, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes, with exceptions	Multi-measure combo charts display a legend, and single-measure combo charts don't.	Legends on visual types in Quick Suite
Changing the title display	Yes		Titles and subtitles on visual types in Quick Suite
Changing the axis range	Yes	You can set the range for the axis.	Range and scale on visual types in Quick Suite
Showing or hiding axis lines, grid lines, axis labels, and axis sort icons	Yes		Axes and grid lines on visual types in Quick Suite
Changing the visual colors	Yes		Colors in visual types in Quick Suite
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude any bar on the chart, except when you are using a date field as the dimension for the axis. In that case, you can only focus on a bar, not exclude it.	Focusing on visual elements Excluding visual elements

Feature	Supported?	Comments	For more information
Sorting	Yes	You can sort on the fields you choose for the axis and the values.	Sorting visual data in Amazon Quick Suite
Performing field aggregation	Yes	You must apply aggregation to the field or fields you choose for the value. You can't apply aggregation to the fields you choose for the axis or group/color.	Changing field aggregation
Adding drill-downs	Yes	You can add drill-down levels to the axis and Group/Color field wells.	Adding drill-downs to visual data in Quick Sight
Synchronizing y-axis	Yes	Synchronize the y-axes for both bars and lines into a single axis.	Range and scale on visual types in Quick Suite

Creating a combo chart

Use the following procedure to create a combo chart.

To create a combo chart

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose one of the combo chart icons.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value. You can create combo charts as follows:
 - Choose a dimension for the **X axis**.
 - To create a single-measure combo chart, choose one measure for either **Bars** or **Lines**.

- To create a multi-measure combo chart, choose two or more measures for the **Bars** or **Lines** field well.
 - Optionally, add a dimension to the **Group/Color** field well. If you have a field in **Group/Color**, you can't have more than one field under **Bars**.
5. (Optional) Add drill-down layers by dragging one or more additional fields to the **X axis** or **Group/Color** field wells. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

Using custom visual content

You can embed webpages and online videos, forms, and images in your Quick Suite dashboards using the custom visual content chart type.

For example, you can embed the image of your company logo in your dashboards. You can also embed an online video from your organization's latest conference, or embed an online form asking readers of the dashboard if the dashboard is helpful.

After you create custom visual content, you can use navigation actions to navigate within them. You can also use parameters to control what appears in them.

The following limitations apply to custom visual content:

- Only https URL schemes are supported.
- Custom visual content isn't supported in email reports.
- Images and websites that use hotlink protection won't load in custom visuals.

To embed a webpage, video, online form, or image in your dashboard, choose the custom visual content icon in the **Visual types** pane.

For more information about adding visuals to a dashboard, see [Adding a visual](#).

Use the following procedures to learn how to embed custom visuals in your dashboards.

Best practices for using custom visual content

When embedding web content using the custom visual content chart type, we recommend the following:

- Choose web content from sources that support viewing or opening the content in an IFrame. If the source of the web content doesn't support being viewed or opened in an IFrame, the content doesn't appear in Quick Suite, even if the URL is accurate.
- When possible, use embeddable URLs, especially for videos, online forms, spreadsheets, and documents. Embeddable URLs create a better experience for readers of your dashboard and make interacting with the content easier. You can usually find the embeddable URL for content when you choose to share the content from the source website.
- To embed internal URLs or URLs that you own, you might need to set them to be opened in an IFrame.
- When viewing custom visual content in an analysis or dashboard, make sure you enable all cookies. If third-party cookies are blocked in your browser, images that are part of the website that is embedded within the custom content visual do not render.

Note

Chrome has announced plans to deprecate all third-party cookies by the end of 2024. This means that websites that are embedded within Quick Suite custom content visuals will no longer show any contents that rely on third-party cookies in Chrome. For more information about Chrome's plans to deprecate third party cookies, see [Chrome is deprecating third-party cookies](#).

Embedding images in a dashboard

You can embed an online image in a dashboard using the image URL. Use the following procedure to embed an image using the custom visual content chart type.

Embedded images don't appear in a browser that has third-party cookies blocked. To see embedded images in a dashboard, enable third-party cookies in your browser settings.

To embed an image in a dashboard

1. In the **Visual types** pane, choose the custom visual content icon.
2. In the visual, choose **Customize visual**.
3. In the **Properties** pane that opens, under **Custom content**, enter the image URL for the image that you want to embed.
4. Choose **Apply**.

The image appears as a webpage in the visual.

5. Choose **Show as image**.

If the URL is an image, the image appears in the visual.

If the URL is not an image, such as a URL to a slide show, gallery, or webpage, the following message appears: This URL doesn't appear to be an image. Update the URL to an image. To do so, open the image that you want to embed in a separate browser tab, or choose an embeddable URL for the image (usually found when you choose to share the image).

6. (Optional) For **Image sizing options**, choose one of the following options:

- **Fit to width** – This option fits the image to the width of the visual.
- **Fit to height** – This option fits the image to the height of the visual.
- **Scale to visual** – This option scales the image to the width and height of the visual. This option might contort the image.
- **Do not scale** – This option keeps the image at its original scale and doesn't fit the image to the dimensions of the visual. With this option, the image is centered in the visual and the parts of the image that are within the width and height of the visual appear. Some parts of the image might not appear if the visual is smaller than the image. If the visual is larger than the image, however, the image is centered in the visual and is surrounded by white space.

Embedding online forms in a dashboard

You can embed an online form in a dashboard using the embeddable URL. Use the following procedure to embed an online form using the custom visual content chart type.

To embed an online form in a dashboard

1. In the **Visual types** pane, choose the custom visual content icon.
2. In the visual, choose **Customize visual**.
3. In the **Properties** pane that opens, under **Custom content**, enter the form URL for the online form that you want to embed.

If possible, use an embeddable URL for the form. Using an embeddable URL creates a better experience for readers of your dashboard who might want to interact with the form. You can

often find the embeddable URL when you choose to share the form on the site where you create it.

4. Choose **Apply**.

The form appears in the visual.

Embedding webpages in a dashboard

You can embed webpage in a dashboard using the URL. Use the following procedure to embed webpage using the custom visual content chart type.

To embed a webpage in a dashboard

1. In the **Visual types** pane, choose the custom visual content icon.
2. In the visual, choose **Customize visual**.
3. In the **Properties** pane that opens, under **Custom content**, enter the URL for the webpage that you want to embed.
4. Choose **Apply**.

The webpage appears in the visual.

Embedding online videos in a dashboard

You can embed an online video in a dashboard using the embeddable video URL. Use the following procedure to embed an online video using the custom visual content chart type.

To embed an online video in a dashboard

1. In the **Visual types** pane, choose the custom visual content icon.
2. In the visual, choose **Customize visual**.
3. In the **Properties** pane that opens, under **Custom content**, enter the embeddable URL for the video that you want to embed.

To find the embeddable URL for a video, share the video and copy the embed URL from IFrame code. The following is an example of an embed URL for a YouTube video: `https://www.youtube.com/embed/uniqueid`. For a Vimeo video, the following is an example of an embed URL: `https://player.vimeo.com/video/uniqueid`.

4. Choose **Apply**.

The video appears in the visual.

Using donut charts

Use donut charts to compare values for items in a dimension. The best use for this type of chart is to show a percentage of a total amount.

Each wedge in a donut chart represents one value in a dimension. The size of the wedge represents the proportion of the value for the selected measure that the item represents compared to the whole for the dimension. Donut charts are best when precision isn't important and there are few items in the dimension.

To learn how to use donut charts in Amazon Quick Suite, you can watch this video:

[Using Donut Charts](#)

To create a donut chart, use one dimension in the **Group/Color** field well. With only one field, the chart displays the division of values by row count. To display the division of dimension values by a metric value, you can add a metric field to the **Value** field well.

Donut charts show up to 20 data points for group or color. For more information about how Amazon Quick Suite handles data that falls outside display limits, see [Display limits](#).

Donut chart features

To understand the features supported by donut charts, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes		Legends on visual types in Quick Suite
Changing the title display	Yes		Titles and subtitles on visual types in Quick Suite
Changing the axis range	Not applicable		Range and scale on visual types in Quick Suite
Changing the visual colors	Yes		Colors in visual types in Quick Suite

Feature	Supported?	Comments	For more information
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude a wedge in a donut chart, except when you are using a date field as a dimension . In that case, you can only focus on a wedge, not exclude it.	Focusing on visual elements Excluding visual elements
Sorting	Yes	You can sort on the field that you choose for the value or the group or color.	Sorting visual data in Amazon Quick Suite
Performing field aggregation	Yes	You must apply aggregation to the field that you choose for the value, and can't apply aggregation to the field that you choose for group or color.	Changing field aggregation
Adding drill-downs	Yes	You can add drill-down levels to the Group/Color field well.	Adding drill-downs to visual data in Quick Sight
Choosing size	Yes	You can choose how thick the donut chart is: small, medium, and large.	Formatting in Amazon Quick Suite
Showing totals	Yes	You can choose to display or hide the aggregate of the Value field. By default, this displays the total count of the Group/Color field, or the total sum of the Value field.	Formatting in Amazon Quick Suite

Creating a donut chart

Use the following procedure to create a donut chart.

To create a donut chart

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the donut chart icon.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value.

To create a donut chart, drag a dimension to the **Group/Color** field well. Optionally, drag a measure to the **Value** field well.

5. (Optional) Add drill-down layers by dragging one or more additional fields to the **Group/Color** field well. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

Using funnel charts

Use a funnel chart to visualize data that moves across multiple stages in a linear process. In a funnel chart, each stage of a process is represented in blocks of different shapes and colors. The first stage, known as the *head*, is the largest block and is followed by the smaller stages, known as the *neck*, in a funnel shape. The size of the block representing each stage in a funnel chart is a percentage of the total, and is proportionate to its value. The bigger the size of the block, the bigger its value.

Funnel charts are often useful in business contexts because you can view trends or potential problem areas in each stage, such as bottlenecks. For example, they can help you visualize the amount of the potential revenue in each stage of a sale, from first contact to final sale and on through maintenance.

To create a basic funnel chart visual

1. Open Amazon Quick Suite and choose **Analyses** on the navigation pane at left.
2. Choose one of the following:

- To create a new analysis, choose **New analysis** at upper right. For more information, see [Starting an analysis in Quick Sight](#).
 - To use an existing analysis, choose the analysis that you want to edit.
3. Choose **Add (+), Add Visual**.
 4. At lower left, choose the funnel chart icon from **Visual types**.
 5. On the **Fields list** pane, choose the fields that you want to use for the appropriate field wells. Funnel charts require one dimension in **Group**.
 6. (Optional) Add drill-down layers by dragging one or more additional fields to the **Group/Color** field well. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

To understand the features supported by funnel charts, see [Analytics formatting per type in Quick Suite](#). For customization options, see [Formatting in Amazon Quick Suite](#).

Using gauge charts

Use gauge charts to compare values for items in a measure. You can compare them to another measure or to a custom amount.

A gauge chart is similar to a nondigital gauge, for example a gas gauge in an automobile. It displays how much there is of the thing you are measuring. In a gauge chart, this measurement can exist alone or in relation to another measurement. Each color section in a gauge chart represents one value. In the following example, we are comparing actual sales to the sales goal, and the gauge shows that we must sell an additional 33.27% to meet the goal.

To learn how to use gauge charts in Amazon Quick Suite, you can watch this video:

[Using Gauge Charts](#)

To create a gauge chart, you need to use at least one measure. Put the measure in the **Value** field well. If you want to compare two measures, put the additional measure in the **Target value** field well. If you want to compare a single measure to a target value that isn't in your dataset, you can use a calculated field that contains a fixed value.

You can choose a variety of formatting options for the gauge chart, including the following settings in **Format visual**.

- **Value displayed** – Hide value, display actual value, or display a comparison of two values

- **Comparison method** – Compare values as a percent, the actual difference between values, or difference as a percent
- **Axis style** –
 - **Show axis label** – Show or hide the axis label
 - **Range** – The numeric minimum and maximum range to display in the gauge chart
 - **Reserve padding (%)** – Added to the top of the range (target, actual value, or max)
- **Arc style** – Degrees the arc displays (180° to 360°)
- **Thickness** – Thickness of the arc (small, medial, or large)

Gauge chart features

To understand the features supported by gauge charts, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes		Legends on visual types in Quick Suite
Changing the title display	Yes		Titles and subtitles on visual types in Quick Suite
Formatting gauge	Yes	You can customize the value displayed, the comparison method, the axis style, the arc style, and the thickness of the gauge.	
Changing the axis range	No		
Changing the visual colors	Yes	The foreground color the filled area; it represents the Value . The background color the unfilled area; it represents the Target value if one is selected.	Colors in visual types in Quick Suite

Feature	Supported?	Comments	For more information
Focusing on or excluding elements	No		
Sorting	No		Sorting visual data in Amazon Quick Suite
Performing field aggregation	Yes		Changing field aggregation
Adding drill-downs	No		

Creating a gauge chart

Use the following procedure to create a gauge chart.

To create a gauge chart

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the gauge chart icon.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells.
To create a gauge chart, drag a measure to the **Value** field well. To add a comparison value, drag a different measure to the **Target value** field well.

Using heat maps

Use heat maps to show a measure for the intersection of two dimensions, with color-coding to easily differentiate where values fall in the range. Heat maps can also be used to show the count of values for the intersection of the two dimensions.

Each rectangle on a heat map represents the value for the specified measure for the intersection of the selected dimensions. Rectangle color represents where the value falls in the range for the measure, with darker colors indicating higher values and lighter colors indicating lower ones.

Heat maps and pivot tables display data in a similar tabular fashion. Use a heat map if you want to identify trends and outliers, because the use of color makes these easier to spot. Use a pivot table if you want to further analyze data on the visual, for example by changing column sort order or applying aggregate functions across rows or columns.

To create a heat map, choose at least two fields of any data type. Amazon Quick Suite populates the rectangle values with the count of the x-axis value for the intersecting y-axis value. Typically, you choose a measure and two dimensions.

Heat maps show up to 50 data points for rows and up to 50 data points for columns. For more information about how Amazon Quick Suite handles data that falls outside display limits, see [Display limits](#).

Heat map features

To understand the features supported by heat maps, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes		Legends on visual types in Quick Suite
Changing the title display	Yes		Titles and subtitles on visual types in Quick Suite
Changing the axis range	Not applicable		Range and scale on visual types in Quick Suite
Changing the visual colors	No		Colors in visual types in Quick Suite
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude a rectangle in a heat map, except when you are using a date field as the rows dimension. In that case, you can only focus on a rectangle, not exclude it.	Focusing on visual elements Excluding visual elements

Feature	Supported?	Comments	For more information
Sorting	Yes	You can sort by the fields you choose for the columns and the values.	Sorting visual data in Amazon Quick Suite
Performing field aggregation	Yes	You must apply aggregation to the fields you choose for the value, and can't apply aggregation to the fields you choose for the rows or columns.	Changing field aggregation
Adding drill-downs	Yes	You can add drill-down levels to the Rows and Columns field wells.	Adding drill-downs to visual data in Quick Sight
Conditional formatting	No		Conditional formatting on visual types in Quick Suite

Creating a heat map

Use the following procedure to create a heat map.

To create a heat map

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the heat map icon.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value.

To create a heat map, drag a dimension to the **Rows** field well, a dimension to the **Columns** field well, and a measure to the **Values** field well.

5. (Optional) Add drill-down layers by dragging one or more additional fields to the **Rows** or **Columns** field wells. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

Using Highcharts

Use Highcharts visuals to create custom chart types and visuals that use the [Highcharts Core library](#). Highcharts visuals provide Quick Suite authors direct access to the [Highcharts API](#).

To configure a Highcharts visual, Quick Suite authors need to add a Highcharts JSON schema to the visual in Quick Suite. Authors can use Quick Suite expressions to reference Quick Suite fields, and formatting options in the JSON schema that they use to generate the Highcharts visual. The JSON **Chart code** editor provides contextual assistance for autocomplete and real time validation to ensure that the input JSON schemas are configured properly. To maintain security, the Highcharts visual editor does not accept CSS, JavaScript, or HTML code input.

For more information about Highcharts visuals in Amazon Quick Suite, see the [Highcharts Visual QuickStart Guide](#) in [DemoCentral](#).

The following image shows a lipstick chart that is configured in the **Chart code** JSON editor of a Highcharts visual in Quick Suite.

For more examples of visuals that you can create with the Highcharts visual in Quick Suite, see [Highcharts demos](#).

Considerations

Before you start creating Highcharts visuals in Amazon Quick Suite, review the following limitations that apply to Highcharts visuals.

- The following JSON values are not supported in the Highcharts **Chart code** JSON editor:
 - Functions
 - Dates
 - Undefined values
- Links to GeoJSON files or other images are not supported for Highcharts visuals.
- Field colors are not available for Highcharts visuals. Default theme colors are applied to all Highcharts visuals.

Creating a Highcharts visual

Use the following procedure to create a Highcharts visual in Amazon Quick Suite.

1. Open the [Quick Suite console](#).
2. Open the Quick Suite analysis that you want to add a Highcharts visual to.
3. On the application bar, choose **Add**, and then choose **Add visual**.
4. On the **Visual types** pane, choose the Highcharts visual icon. An empty visual appears on the analysis sheet and the **Properties** pane opens on the left.
5. In the **Properties** pane, expand the **Display settings** section and perform the following actions:
 - a. For **Edit title**, choose the paintbrush icon, enter the title that you want the visual to have, and then choose **SAVE**. Alternatively, choose the eyeball icon to hide the title.
 - b. (Optional) For **Edit subtitle**, choose the paintbrush icon, enter the subtitle that you want the visual to have, and then choose **SAVE**. Alternatively, choose the eyeball icon to hide the subtitle.
 - c. (Optional) For **Alt text**, add the alt text that you want the visual to have.
6. Expand the **Data point limit** section. For **Number of data points to show**, enter the number of data points that you want the visual to show. Highcharts visuals can show up to 10,000 data points.
7. Expand the **Chart code** section.
8. Enter a JSON schema into the **Chart code** JSON editor. The editor provides contextual assistance and real time validation to ensure that your input JSON is configured properly. Any errors that Quick Suite identifies can be viewed in the **Errors** dropdown. The example below shows a JSON schema that creates a lipstick chart that shows current year sales by industry.

```
{
  "xAxis": {
    "categories": ["getColumn", 0]
  },
  "yAxis": {
    "min": 0,
    "title": {
      "text": "Amount ($)"
    }
  }
},
```

```

"tooltip": {
  "headerFormat": "<span style='font-size:10px'>{point.key}</span><table>",
  "pointFormat": "<tr><td style='color:{series.color};padding:0'>{series.name}:</td><td style='padding:0'><b>${point.y:,.0f}</b></td></tr>",
  "footerFormat": "</table>",
  "shared": true,
  "useHTML": true
},
"plotOptions": {
  "column": {
    "borderWidth": 0,
    "grouping": false,
    "shadow": false
  }
},
"series": [
  {
    "type": "column",
    "name": "Current Year Sales",
    "color": "rgba(124,181,236,1)",
    "data": ["getColumn", 1],
    "pointPadding": 0.3,
    "pointPlacement": 0.0
  }
]
}

```

9. Choose **APPLY CODE**. Quick Suite converts the JSON schema into a visual that appears in the analysis. To make changes to the rendered visual, update the appropriate properties in the JSON schema and choose **APPLY CODE**.
10. (Optional) Open the **Reference** dropdown to access links to helpful Highcharts reference material.

When you are happy with the rendered visual, close the properties pane. For more information about Quick Sight specific expressions that can be used to configure a Highcharts visual, see [Amazon Quick Suite JSON expression language for Highcharts visuals](#).

Interactive Highchart features

Highchart visualizations in Amazon Quick Sight support custom actions, highlighting, and custom field color consistencies, allowing you to create interactive and visually cohesive charts that integrate seamlessly with other Quick Sight visuals.

Custom actions

With custom actions, you can define specific behaviors for any data point in your Highchart visualizations. This feature seamlessly integrates with Quick Sight's existing action framework, enabling you to create interactive charts that respond to user clicks. The system currently supports single data point selection, giving you precise control over user interactions. Custom actions can be implemented across various chart types, including line charts, bar charts, and stacked bar charts, among others.

To implement custom actions, you'll need to modify your Highcharts JSON configuration. Add an event block to your series configuration, specifying the click event and the corresponding action. For example:

```
{
  "series": [{
    "type": "line",
    "data": ["getColumn", 1],
    "name": "value",
    "events": {
      "click": [
        "triggerClick", { "rowIndex": "point.index" }
      ]
    }
  ]
}
```

This configuration enables click events on your chart's data points, allowing Quick Sight to handle custom actions based on the selected data.

Cross-visual highlighting

Cross-visual highlighting enhances the interactivity of your dashboards by creating visual connections between different charts. When a user selects elements in one chart, related elements in other visuals are automatically highlighted, while unrelated elements are dimmed. This feature helps users quickly identify relationships and patterns across multiple visualizations, improving data comprehension and analysis.

To enable cross-visual highlighting and maintain field color consistency, use the `quicksight` clause in your Highcharts JSON configuration. This clause acts as a bridge between Highcharts rendering and Quick Suite's visual interaction system. Here's an example of how to set it up:

```
{
```

```

"quicksight": {
  "pointRender": ["updatePointAttributes", {
    "opacity": ["case",
      ["dataMarkMatch", ["getColumnName", 0], "series.name"],
      1, // Full opacity for matching elements
      0.1 // Dim non-matching elements
    ],
    "color": ["getColumnColorOverrides", ["getColumnName", 0], "series.name"]
  }]
}
}

```

This configuration uses Quick Sight's JSON expression language to dynamically modify visual properties like opacity and color based on user interactions and predefined color schemes.

For more complex scenarios, you can set up highlighting based on multiple conditions. This allows for more nuanced interactivity in your visualizations. The following example highlights elements based on either the quarter or day of the week:

```

{
  "quicksight": {
    "pointRender": ["updatePointAttributes", {
      "opacity": ["case",
        ["||",
          ["dataMarkMatch", "quarter", "series.name"],
          ["dataMarkMatch", "day_of_week", "point.name"]
        ],
        1, // Full opacity for matching elements
        0.1 // Dim non-matching elements
      ],
    }]
  }
}

```

Field-level color consistency

Maintaining visual coherence across your dashboard is crucial for effective data interpretation. The field-level color consistency feature ensures that colors assigned to specific dimensions persist across all visuals in your dashboard. This consistency helps users quickly recognize and track particular data categories across different chart types and views, enhancing the overall user experience and data comprehension.

Amazon Quick Suite JSON expression language for Highcharts visuals

Highcharts visuals accept most [valid JSON values](#), standard arithmetic operators, string operators, and conditional operators. The following JSON values are not supported for Highcharts visuals:

- Functions
- Dates
- Undefined values

Quick Suite authors can use JSON expression language create JSON schemas for a highcharts visual. JSON expression language is used to bind JSON to APIs or datasets to allow dynamic population and modification of JSON structures. Developers can also use JSON expression language to inflate and transform JSON data with concise and intuitive expressions.

With JSON expression language, expressions are represented as arrays, where the first element specifies the operation and subsequent elements are the arguments. For example, `["unique", [1, 2, 2]]` applies the unique operation to the array `[1, 2, 2]`, resulting in `[1, 2]`. This array-based syntax allows for flexible expressions, that allow complex transformations on JSON data.

JSON expression language supports *nested expressions*. Nested expressions are expressions that contain other expressions as arguments. For example `["split", ["toUpper", "hello world"], " "]` first converts the string `hello world` into an uppercase, then splits it into array of words, resulting in `["HELLO", "WORLD"]`.

Use the following sections to learn more about JSON expression language for Highcharts visuals in Amazon Quick Suite.

Topics

- [Arithmetics](#)
- [Array operations](#)
- [Amazon Quick Suite expressions](#)

Arithmetics

The following table shows arithmetic expressions that can be used with JSON expression language.

Operation	Expression	Input	Output
Addition	["+", operand1, operand2]	{ sum: ["+", 2, 4] }	{ sum: 6 }
Subtraction	["-", operand1, operand2]	{ difference: ["-", 10, 3] }	{ difference: 7 }
Multiplication	["*", operand1, operand2]	{ product: ["*", 5, 6] }	{ product: 30 }
Division	["/", operand1, operand2]	{ quotient: ["/", 20, 4] }	{ quotient: 5 }
Modulo	["%", operand1, operand2]	{ remainder: ["%", 15, 4] }	{ remainder: 3 }
Exponentiation	["**", base, exponent]	{ power: ["**", 2, 3] }	{ power: 8 }
Absolute Value	["abs", operand]	{ absolute: ["abs", -5] }	{ absolute: 5 }
Square Root	["sqrt", operand]	{ sqroot: ["sqrt", 16] }	{ sqroot: 4 }

Operation	Expression	Input	Output
Logarithm (base 10)	<code>["log10", operand]</code>	<code>{ log: ["log10", 100] }</code>	<code>{ log: 2 }</code>
Natural Logarithm	<code>["ln", operand]</code>	<code>{ ln: ["ln", Math.E] }</code>	<code>{ ln: 1 }</code>
Round	<code>["round", operand]</code>	<code>{ rounded: ["round", 3.7] }</code>	<code>{ rounded: 4 }</code>
Floor	<code>["floor", operand]</code>	<code>{ floor: ["floor", 3.7] }</code>	<code>{ floor: 3 }</code>
Ceiling	<code>["ceil", operand]</code>	<code>{ ceiling: ["ceil", 3.2] }</code>	<code>{ ceiling: 4 }</code>
Sine	<code>["sin", operand]</code>	<code>{ sine: ["sin", 0] }</code>	<code>{ sine: 0 }</code>
Cosine	<code>["cos", operand]</code>	<code>{ cosine: ["cos", 0] }</code>	<code>{ cosine: 1 }</code>
Tangent	<code>["tan", operand]</code>	<code>{ tangent: ["tan", Math.PI] }</code>	<code>{ tangent: 0 }</code>

Array operations

JSON expression language allows generic array manipulation for the following functions:

- **map** – Applies a mapping function to each element of an array and returns a new array with the transformed values.

For example, `["map", [1, 2, 3], ["*", ["item"], 2]]` maps each element of the array `[1, 2, 3]` by multiplying it by 2.

- **filter** – Filters an array based on a given condition and returns a new array containing only the elements that satisfy the condition

For example, `["filter", [1, 2, 3, 4, 5], ["==", ["%", ["item"], 2], 0]]` filters the array `[1, 2, 3, 4, 5]` to include only the even numbers.

- **reduce** – Reduces an array to a single value by applying a reducer function to each element and accumulating the result.

For example, `["reduce", [1, 2, 3, 4, 5], ["+", ["acc"], ["item"]], 0]` reduces the array `[1, 2, 3, 4, 5]` to the sum of its elements.

- **get** – Retrieves a value from an object or an array by specifying a key or index.

For example, `["get", ["item"], "name"]` retrieves the value of the "name" property from the current item.

- **unique** – Given an array returns only unique items inside this array.

For example, `["unique", [1, 2, 2]]` returns `[1, 2]`.

Amazon Quick Suite expressions

Amazon Quick Suite offers additional expressions to enhance the functionality of Highcharts visuals. Use the following sections to learn more about common Quick Suite expressions for highcharts visuals. For more information about JSON expression language in Amazon Quick Suite, see the [Highcharts Visual QuickStart Guide](#) in [DemoCentral](#).

Topics

- [getColumn](#)
- [formatValue](#)

getColumn

Use the `getColumn` expressions to return values from specified column indices. For example, the following table shows a list of products alongside their category, and price.

Product name	Category	Price
Product A	Technology	100

Product name	Category	Price
Product B	Retail	50
Product C	Retail	75

The following `getColumn` query generates an array that shows all product names alongside their price.

```
{
  product name: ["getColumn", 0],
  price: ["getColumn", 2]
}
```

The following JSON is returned:

```
{
  product name: ["Product A", "Product B", "Product C"],
  price: [100, 50, 75]
}
```

You can also pass multiple columns at once to generate an array of arrays, shown in the following example.

Input

```
{
  values: ["getColumn", 0, 2]
}
```

Output

```
{
  values: [["Product A", 100], ["Product B", 50], ["Product C", 75]]
}
```

Similar to `getColumn`, the following expressions can be used to return column values from field wells or themes:

- `getColumnFromGroupBy` returns columns from the group by field. The second argument is the index of the column to return. For example, `["getColumnFromGroupBy", 0]` returns values of the first field as an array. You can pass multiple indices to get an array of arrays where each element corresponds to the field in the group by field well.
- `getColumnFromValue` returns columns from the value field well. You can pass multiple indices to get an array of arrays where each element corresponds to the field in the values field well.
- `getColorTheme` returns the current color palette of a Quick Suite theme, shown in the following example.

```
{
  "color": ["getColorTheme"]
}
```

```
{
  "color": ["getPaletteColor", "secondaryBackground"]
}
```

Example

`getColumn` can access any column from the table:

- `["getColumn", 0]` - returns array `[1, 2, 3, 4, 5, ...]`
- `["getColumn", 1]` - returns array `[1, 1, 1, 1, 1, ...]`
- `["getColumn", 2]` - returns array `[1674, 7425, 4371, ...]`

`getColumnFromGroupBy` works similarly, but its index is limited to the columns in the group by field well:

- `["getColumnFromGroupBy", 0]` - returns array `[1, 2, 3, 4, 5, ...]`
- `["getColumnFromGroupBy", 1]` - returns array `[1, 1, 1, 1, 1, ...]`
- `["getColumnFromGroupBy", 2]` - does not work, since there are only two columns in the group by field well

`getColumnFromValue` works similarly, but its index is limited to the columns in the value field well:

- `["getColumnFromValue", 0]` - returns array `[1, 2, 3, 4, 5, ...]`
- `["getColumnFromValue", 1]` - does not work, since there is only one column in the value field well
- `["getColumnFromValue", 2]` - does not work, since there is only one column in the value field well

formatValue

Use the `formatValue` expression to apply Quick Suite formatting to your values. For example, the following expression formats the x-axis label with the format value that is specified in the first field of Quick Suite field wells.

```
"xAxis": {
  "categories": ["getColumn", 0],
  "labels": {
    "formatter": ["formatValue", "value", 0]
  }
}
```

Using histograms

Use a histogram chart in Amazon Quick Suite to display the distribution of continuous numerical values in your data. Amazon Quick Suite uses un-normalized histograms, which use an absolute count of the data points or events in each bin.

To create a histogram, you use one measure. A new histogram initially displays ten *bins* (also called *buckets*) across the X-axis. These appear as bars on the chart. You can customize the bins to suit your dataset. The Y-axis displays the absolute count of the values in each bin.

Make sure that you adjust the format settings so that you have a clearly identifiable shape. If your data contains outliers, this becomes clear if you spot one or more values off to the side of the X-axis. For information about how Amazon Quick Suite handles data that falls outside display limits, see [Display limits](#).

Histogram features

To understand the features supported by histograms, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	No		Legends on visual types in Quick Suite
Changing the title display	Yes		Titles and subtitles on visual types in Quick Suite
Changing the axis range	No	However, you can change the bin count or the bin interval width (range of distribution).	
Showing or hiding axis lines, grid lines, axis labels, and axis sort icons	Yes		Axes and grid lines on visual types in Quick Suite
Changing the visual colors	Yes		Colors in visual types in Quick Suite
Focusing on or excluding elements	No		
Sorting	No		
Performing field aggregation	No	Histograms use only the count aggregation.	
Adding drill-downs	No		

Creating a histogram

Use the following procedure to create a histogram.

To create a histogram

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the histogram icon.
4. On the **Fields list** pane, choose the field that you want to use in the **Value** field well. A **Count** aggregate is automatically applied to the value.

The resulting histogram shows the following:

- The X-axis displays 10 bins by default, representing the intervals in the measure that you choose. You can customize the bins in the next step.
 - The Y-axis displays the absolute count of individual values in each bin.
5. (Optional) Choose **Format** on the visual control to change the histogram format. You can format the bins either by count or width, not both together. The count setting changes how many bins display. The width setting changes how wide or long of an interval each bin contains.

Formatting a histogram

Use the following procedure to format a histogram.

To format a histogram

1. Choose the histogram chart that you want to work with. It should be the highlighted selection. The visual controls display on the top right of the histogram.
2. Choose the cog icon on the visual control menu to view the **Format visual** options.
3. On the **Properties** pane, set the following options to control the display of the histogram:
 - **Histogram** settings. Chose *one* of the following settings:
 - Bin count (option 1): The number of bins that display on the X-axis.
 - Bin width (option 1): The width (or length) of each interval. This setting controls the number of items or events to include in each bin. For example, if your data is in minutes, you can set this to 10 to show 10-minute intervals.
 - With the following settings, you can explore the best way to format the histogram for your dataset. For example, in some cases, you might have a tall peak in one bin, while most

of the other bins look sparse. This isn't a useful view. You can use the following settings individually or together:

- Change the **Number of data points displayed** in the **X-axis** settings.

Amazon Quick Suite displays up to 100 bins (buckets) by default. If you want to display more (up to 1,000), change the X-axis setting for **Number of data points displayed**.

- Enable **Logarithmic scale** in the **Y-axis** settings.

Sometimes your data doesn't fit the shape that you want and this can provide misleading results. For example, if the shape is skewed so far to the right that you can't read it properly, you can apply a log scale to it. Doing this doesn't normalize your data; however, it does reduce the skew.

- Display **Data labels**.

You can enable the display of data labels to see the absolute counts in the chart. Even if you don't want to display these in most cases, you can enable them while you're developing an analysis. The labels can help you decide on formatting and filtering options because they reveal counts in bins that are too small to stand out.

To see all the data labels, even if they overlap, enable **Allow labels to overlap**.

4. (Optional) Change other visual settings. For more information, see [Formatting in Amazon Quick Suite](#).

Understanding histograms

Although histograms look similar to bar charts, they are very different. In fact, the only similarity is their appearance because they use bars. On a histogram, each bar is called a *bin* or a *bucket*.

Each bin contains a range of values called an *interval*. When you pause on one of the bins, details about the interval appear in a tooltip that shows two numbers enclosed in glyphs. The type of enclosing glyphs indicates if the numbers inside them are part of the interval that's inside the selected bin, as follows:

- A square bracket next to a number means that the number is included.
- A parenthesis next to a number means that the number is excluded.

For example, let's say that the first bar in a histogram displays the following notation.

$[1, 10)$

The square bracket means that the number 1 is included in the first interval. The parenthesis means that the number 10 is excluded.

In the same histogram, a second bar displays the following notation.

 $[10, 20)$

In this case, 10 is included in the second interval, and 20 is excluded. The number 10 can't exist in both intervals, so the notation shows us which one includes it.

Note

The pattern used for marking intervals in a histogram comes from standard mathematical notation. The following examples show the possible patterns, using a set of numbers that includes 10, 20, and every number in between.

- $[10, 20]$ – This set is closed. It has hard boundaries on both ends.
- $[10, 21)$ – This set is half open. It has a hard boundary on the left and a soft boundary on the right.
- $(9, 20]$ – This set is half open. It has a soft boundary on the left and a hard boundary on the right.
- $(9, 21)$ – This set is open. It has soft boundaries on both ends.

Because the histogram uses quantitative data (numbers) rather than qualitative data, there's a logical order to the distribution of the data. This is called a *shape*. The shape is often described the qualities the shape possesses, based on the count in each bin. Bins that contain a higher number of values form a *peak*. Bins that contain a lower number of values form a *tail* on the edge of a chart, and a *valley* between peaks. Most histograms fall into one of the following shapes:

- Asymmetrical or *skewed* distributions have values that cluster near the left or the right—the low or high end of the X-axis. The direction of skewness is defined by where the longer tail of the data is, not by where the peak is. It's defined this way because this direction also describes the location of the mean (average). In skewed distributions, the mean and the median are two different numbers. The different types of skewed distribution are as follows:

- *Negatively skewed* or *left skewed* – A chart that has the mean to the left of the peak. It has a longer tail to the left and a peak to the right, sometimes followed by a shorter tail.
- *Positively skewed* or *right skewed* – A chart that has the mean to the right of the peak. It has a longer tail to the right and a peak to the left, sometimes preceded by a shorter tail.
- Symmetrical or *normal* distributions have a shape that's mirrored on each side of a center point (for example, a bell curve). In a normal distribution, the mean and the median are the same value. The different types of normal distribution are as follows:
 - Normal distribution, or *unimodal* – A chart that has one central peak representing the most common value. This is commonly called a bell curve, or a Gaussian distribution.
 - Bimodal – A chart that has two peaks representing the most common values.
 - Multimodal – A chart that has three or more peaks representing the most common values.
 - Uniform – A chart that has no peaks or valleys, with a relatively equal distribution of data.

The following table shows how a histogram differs from a bar chart.

Histogram	Bar chart
A histogram displays the distribution of values in one field.	A bar chart compares the values in one field, grouped by dimension.
A histogram sorts values into bins that represent a range of values, for example 1–10, 10–20, and so on.	A bar chart plots values that are grouped into categories.
The sum of all bins equals exactly 100% of the values in the filtered data.	A bar chart isn't required to display all of the available data. You can change display settings at the visual level. For example, a bar chart might show only the top 10 categories of data.
Rearranging bars detracts from the meaning of the chart as a whole.	Bars can be in any order without changing the meaning of the chart as a whole.
There are no spaces between the bars, to represent the fact this is continuous data.	There are spaces between the bars, to represent the fact that this is categorical data.

Histogram	Bar chart
If a line is included in a histogram, it represent s the general shape of the data.	If a line is included in a bar chart, it's called a combo chart, and the line represents a different measure than the bars.

Using image components

Use image components to upload static images from your desktop to a Quick Suite analysis. Each Quick Suite analysis sheet supports up to 10 image components. Image components are not included in the 50 visual per sheet limit. The file size of an image component can't exceed 1MB.

The following file formats are supported for image components:

- .bmp
- .jpg/.jpeg
- .png
- .tiff
- .webp

Use the following procedure to add an image component to a Quick Suite analysis:

To add an image component to a Quick Suite analysis

1. Open the [Quick Suite console](#).
2. Open the Quick Suite analysis that you want to add an image to.
3. Choose the **Add image** button in the toolbar at the top of the analysis.
4. The file selection of your desktop opens. Choose the file that you want to upload, and then choose **Open**. The file size of the image component can't exceed 1MB.
5. The image is uploaded to Quick Suite and appears in the analysis.
6. (Optional) To add alt text or update the image scaling options, choose the **Properties** icon at the top right of the image to open the **Properties** pane.
7. (Optional) To add a [custom tooltip](#) to the image, open the **Properties** pane, choose **Interactions**, and then choose **Add action**. Filter actions are not supported for image

components. You can also use the **Interactions** section to add custom navigation and URL actions to the image component.

8. (Optional) To duplicate or replace the image, choose the **More options** ellipsis (three dots) icon at the top right of the image, and then choose the action that you want to perform.

Using KPIs

Use a key performance indicator (KPI) to visualize a comparison between a key value and its target value.

A KPI displays a value comparison, the two values being compared, and a visual that provides context to the data that's displayed. You can choose from a set of predesigned layouts to suit your business needs. The following image shows an example of a KPI visual that uses a sparkline.

1. Choose **Add (+)** drop down in the **Visuals pane**.
2. Choose the KPI icon from Visual types menu.

KPI features

To understand the features supported by the KPI visual type in Amazon Quick Suite, use the following table.

Feature	Supported?	Comments	For more information
Changing the title display	Yes		Titles and subtitles on visual types in Quick Suite
Removing the title	Yes	You can choose not to display a title.	
Changing comparison method	Yes	By default, Amazon Quick Suite automatically chooses a method. The settings are auto, difference, percent, and difference as percent.	

Feature	Supported?	Comments	For more information
Changing the primary value displayed	Yes	You can choose comparison (default) or actual.	
Displaying or removing the progress bar	Yes	You can format the visual to either display (default) or not display a progress bar.	

For more information on KPI formatting options, see [KPI options](#).

Creating a KPI

Use the following procedure to create a KPI.

To create a KPI

1. Create a new analysis for your dataset.
2. In the **Visual types** pane, choose the KPI icon.
3. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. You must use measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value.

To create a KPI, drag a measure to the **Value** field well. To compare that value to a target value, drag a different measure to the **Target value** field well.

4. (Optional) Choose formatting options by selecting the on-visual menu at the upper-right corner of the visual, then choosing **Format visual**.

Changing a KPI's layout

Use the following procedure to change the layout for a KPI.

To change the layout of a KPI

1. Navigate to the KPI visual that you want to change and choose **KPI layouts**.
2. In the **KPI Layouts** pane, choose the KPI layout that you want to use.

Using layer maps

Use layer maps to visualize data with custom geographic boundaries, such as congressional districts, sales territories, or user-defined regions. With layer maps, Quick Suite authors upload GeoJSON files to Amazon Quick Suite that shape layers over a base map and join with Quick Suite data to visualize associated metrics and dimensions. Shape layers can be styled by color, border, and opacity. Quick Suite authors can also add interactivity to layer maps through tooltips and custom actions.

Note

Amazon Quick Suite layer map visuals only support polygon shapes. Line and point geometries are not supported.

The following image shows a layer map visual in Amazon Quick Suite.

Creating a shape layer with layered maps

Use the procedure below to create a shape layer with layer map visuals in Amazon Quick Suite.

1. Open the [Quick Suite console](#).
2. Open the Quick Suite analysis that you want to add a layer map to.
3. Choose **Add** on the application bar, and then choose **Add visual**.
4. On the **Visual types** pane, choose one of the layer map icon.
5. An empty map visual appears in the analysis and prompts you to continue configuring the layer. Choose **CONFIGURE LAYER** to continue configuring the layer map.
6. The **Layer properties** pane opens to the right. Navigate to the **Shape file** section, and then choose **UPLOAD SHAPE FILE**.
7. Choose the GeoJSON file that you want to visualize. The file must be in .geojson format and must not exceed 100 MB.
8. Navigate to the **Data** section.
9. For **Shape file key field**, choose the field that you want the shape to visualize.
10. (Optional) For **Dataset key field**, choose the dataset field that you want the shape to visualize. To assign color to the shapes, add a color field. If the color field is a measure, the shape uses gradient coloring. If the color field is a dimension, the shape uses categorical coloring. If a

color field is not assigned to the shape, use the fill color option in the **Styling** section of the **Layer properties** pane to set a common color for all shapes.

11. (Optional) To change the layer name, navigate to the **Layer options** section and enter a name in the **Layer name** input.
12. (Optional) To change the fill or border colors, navigate to the **Styling** section and choose the color switch next to the object that you want to change. To adjust the opacity of the color, enter a percentage amount in the input located next to the eye icon. If you do not assign a color field to the **Dataset key field**, the fill color can be used to set a common color for all shapes.

Using line charts

Use line charts to compare changes in measure values over period of time, for the following scenarios:

- One measure over a period of time, for example gross sales by month.
- Multiple measures over a period of time, for example gross sales and net sales by month.
- One measure for a dimension over a period of time, for example number of flight delays per day by airline.

Line charts show the individual values of a set of measures or dimensions against the range displayed by the Y axis. Area line charts differ from regular line charts in that each value is represented by a colored area of the chart instead of just a line, to make it easier to evaluate item values relative to each other.

The following screenshot shows a line chart.

The following screenshot shows an area line chart. In this version of a line chart, the area between the line and the x-axis is filled with color.

The following screenshot shows a stacked area line chart. In this version of a line chart, the area between the line and the **X axis** is filled with color. Also, the individual lines are layered to more clearly show the relationships between them. The values on the y-axis show the scale of the differences between data points.

Because a stacked area line chart works differently than other line charts, simplify it if you can. Then the audience won't try to interpret the numbers. Instead, they can focus on the relationships of each set of values to the whole. One way to simplify is to remove the numbers down the left side of the screen by reducing the step size for the axis. To do this, choose the **Options** icon from the on-visual menu. In **Format Options** under **Y-axis**, enter **2** as the **Step size**. The following screenshot shows the result.

Each line on the chart represents a measure value over a period of time. You can interactively view the values on the chart, as shown in the following screenshot. Hover over any line (1 in the screenshot) to see a pop-up legend that shows the values for each line on the **X axis**. If you hover over a data point (2), you can see the **Value** for that specific point on the **X axis**.

Use line charts to compare changes in values for one or more measures or dimensions over a period of time.

In regular line charts, each value is represented by a line, and in area line charts each value is represented by a colored area of the chart.

Use stacked area line charts to compare changes in values for one or more groups of measures or dimensions over a period of time. Stacked area line charts show the total value for each group on the x-axis. They use color segments to show the values of each measure or dimension in the group.

Line charts show up to 10,000 data points on the x-axis when no color field is selected. When color is populated, line charts show up to 400 data points on the x-axis and up to 25 data points for color. For more information about data that falls outside the display limit for this visual type, see [Display limits](#).

Line chart features

To understand the features supported by line charts, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes		Legends on visual types in Quick Suite
Changing the title display	Yes		Titles and subtitles on visual types in Quick Suite

Feature	Supported?	Comments	For more information
Changing the axis range	Yes	You can set the range for the Y axis.	Range and scale on visual types in Quick Suite
Showing or hiding axis lines, grid lines, axis labels, and axis sort icons	Yes		Axes and grid lines on visual types in Quick Suite
Adding a second Y-axis	Yes		Creating a dual-axis line chart
Changing the visual colors	Yes		Colors in visual types in Quick Suite
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude any line on the chart, except in the following cases: - You create a multi-dimension line chart and use a date field as the dimension for the line color. - You create a measure or multi-measure line chart and use a date field as the dimension for the X axis. In these cases, you can only focus on a line, not exclude it.	Focusing on visual elements Excluding visual elements

Feature	Supported?	Comments	For more information
Sorting	Yes, with exceptions	You can sort data for numeric measures in the X axis and Value field wells. Other data is automatically sorted in ascending order.	Sorting visual data in Amazon Quick Suite
Performing field aggregation	Yes	You must apply aggregation to the field that you choose for the value, and can't apply aggregation to the fields you choose for the X axis and color.	Changing field aggregation
Adding drill-downs	Yes	You can add drill-down levels to the X axis and Color field wells.	Adding drill-downs to visual data in Quick Sight

Creating a line chart

Use the following procedure to create a line chart.

To create a line chart

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose one of the line chart icons.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value.
 - To create a single-measure line chart, drag a dimension to the **X axis** field well and one measure to the **Value** field well.
 - To create a multi-measure line chart, drag a dimension to the **X axis** field well and two or more measures to the **Value** field well. Leave the **Color** field well empty.

- To create a multi-dimension line chart, drag a dimension to the **X axis** field well, one measure to the **Value** field well, and one dimension to the **Color** field well.
5. (Optional) Add drill-down layers by dragging one or more additional fields to the **X axis** or **Color** field wells. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

Creating a dual-axis line chart

If you have two or more metrics that you want to display in the same line chart, you can create a dual-axis line chart.

A *dual-axis chart* is a chart with two Y-axes (one axis at the left of the chart, and one axis at the right of the chart). For example, let's say you create a line chart. It shows the number of visitors who signed up for a mailing list and for a free service over a period of time. If the scale between those two measures varies widely over time, your chart might look something like the following line chart. Because the scale between measures varies so greatly, the measure with the smaller scale appears nearly flat at zero.

If you want to show these measures in the same chart, you can create a dual-axis line chart. The following is an example of the same line chart with two Y-axes.

To create a dual-axis line chart

1. In your analysis, create a line chart. For more information about creating line charts, see [Creating a line chart](#).
2. In the **Value field well**, choose a field drop-down menu, choose **Show on: Left Y-axis**, and then choose **Right Y-axis**.

Or you can create a dual-axis line chart using the **Properties** pane:

- a. On the menu in the upper-right corner of the line chart, choose the **Format visual** icon.
- b. In the **Properties** pane that opens, choose **Data series**.
- c. In the **Data series** section, choose the **Show on right axis** icon for the value that you want to place on a separate axis. Use the search bar to quickly find a value if you need to.

The icon updates to indicate that the value is being shown on the right axis. The chart updates with two axes.

The **Properties** pane updates with the following options:

- To synchronize the Y-axes for both lines back into a single axis, choose **Single Y-axis** at the top of the **Properties** pane.
- To format the axis at the left of the chart, choose **Left Y-axis**.
- To format the axis at the right of the chart, choose **Right Y-axis**.

For more information about formatting axis lines, see [Axes and grid lines](#). For more information about adjusting the range and scale of an axis, see [Range and scale](#).

Creating maps and geospatial charts

You can create two types of maps in Quick Suite: point maps and filled maps. *Point maps* show the difference between data values for each location by size. *Filled maps* show the difference between data values for each location by varying shades of color.

Important

Geospatial charts in Quick Suite currently aren't supported in some Amazon Web Services Regions, including in China.

For help with geospatial issues, see [Geospatial troubleshooting](#).

Before you get started creating maps, do the following:

- Make sure that your dataset contains location data. *Location data* is data that corresponds to latitudinal and longitudinal values. Location data can include a column for latitude and a column for longitude in your dataset. It can also include a column with city names. Quick Suite can chart latitude and longitude coordinates. It also recognizes geographic components such as country, state or region, county or district, city, and ZIP code or postal code.
- Make sure that your location data fields are marked as geospatial data types.
- Consider creating geographic hierarchies.

For more information about working with geospatial data, including changing field data types and creating geospatial hierarchies, see [Adding geospatial data](#).

To learn more about creating maps in Quick Suite, see the following.

Topics

- [Creating point maps](#)
- [Creating filled maps](#)
- [Interacting with maps](#)

Creating point maps

You can create point maps in Quick Suite to show the difference between data values for each location by size. Each point on this type of map corresponds to a geographic location in your data, such as a country, state or province, or city. The size of the points on the map represents the magnitude of the field in the **Size** field well, in relation to other values in the same field. The color of the points represents the values in the **Color** field well. The field values in the **Color** field well display in the legend, if you choose a field for color.

Use the following procedure to create a point map in Quick Suite.

To create point maps in Quick Suite, make sure that you have the following:

- One geospatial field (such as country, state or region, county or district, city, or ZIP code or postal code). Or you can use one latitude field and one longitude field.
- One numeric field (measure) for size.
- (Optional) A categorical field (dimension) for color.

For information on formatting geospatial maps, see [Map and geospatial chart formatting options](#).

Creating point maps

To create a point map

1. Add a new visual to your analysis. For more information about starting analyses, see [Starting an analysis in Quick Sight](#). For more information about adding visuals to analyses, see [Adding a visual](#).
2. For **Visual type**, choose the **Points on map** icon. It looks like a globe with a point on it.

3. Drag a geographic field from the **Fields list** pane to the **Geospatial** field well, for example **Country**. You can also choose a latitude or longitude field.

A point map appears with a point for each location in your data.

If the field is part of a geographic hierarchy, the hierarchy displays in the field well.

4. Drag a measure from the **Fields list** pane to the **Size** field well.

The points on the map update to show the magnitude of values for each location.

5. (Optional) Drag a dimension from the **Fields list** pane to the **Color** field well.

Each point updates to show a point for each categorical value in the dimension.

Creating filled maps

You can create filled maps in Quick Suite to show the difference between data values for each location by varying shades of color.

Use the following procedure to create a filled map in Quick Suite.

To create filled maps in Quick Suite, make sure that you have the following:

- One geospatial field (such as country, state or region, county or district, or ZIP code or postal code).
- (Optional) A numeric field (measure) for color.

Creating filled maps

To create a filled map

1. Add a new visual to your analysis. For more information about starting analyses, see [Starting an analysis in Quick Sight](#). For more information about adding visuals to analyses, see [Adding a visual](#).
2. For **Visual type**, choose the **Filled map** icon.
3. Drag a geographic field from the **Fields list** pane to the **Location** field well, for example **Country**.

A filled map appears with each location in your data filled in by the number of times they appear in your dataset (the count).

If the field is part of a geographic hierarchy, the hierarchy displays in the field well.

4. (Optional) Drag a measure from the **Fields list** pane to the **Color** field well, for example Sales.

Each location updates to show the sum of sales.

Interacting with maps

When you view a map visual in an Quick Suite analysis or published dashboard, you can interact with it to explore your data. You can pan, zoom in and out, and autozoom to all the data.

By default, map visuals are always zoomed based on the underlying data. When you pan around in the map or zoom to a different level, the zoom to data icon appears above the zoom in and out icons at bottom right of the map. Using this option, you can quickly zoom back to the underlying data.

To pan in a map visual

- Click anywhere on the map visual and drag your cursor in the direction that you want to pan the map.

To zoom in or out in a map visual

- On the map visual, choose the plus or minus icons at bottom right. Or you can double-click the map to zoom in, and shift-double-click to zoom out.

To zoom back to all the data

- On the map visual, choose the zoom to data icon. This icon appears when you pan or zoom in on a map.

Using small multiples

Use this feature when you need to set multiple comparative visuals in a row. When you activate the *small multiples* feature, Amazon Quick Suite creates a container or shelf of small visuals, presented side-by-side. Each copy of the visual contains a one view of the data. Using small multiples is a way to get a holistic view of your business, in an efficient and interactive way.

Small multiples aren't listed in the palette visualization icons. Instead, the option to create small multiples appears as a field well, in the visuals that support it.

To add small visuals to your analysis

1. On a line, bar, or pie charts, add a field to the **Small multiples** field well.
2. To see your small multiples, you need to enlarge the container that holds them, so you can see all of them at once.
3. To format the set of small multiples, choose Format visual (the pencil icon) from the menu on the visual. You can adjust the following settings:

- **Layout**

- **Visible rows**
- **Visible columns**
- **Number of panels**
- **Panel title options (toggle)**
 - Font size and color
 - Font weight
 - Text alignment
- **Panel order options (toggle)**

Line thickness, style, and color

- **Panel gutter (toggle)**

Spacing

- **Panel background (toggle)**

Background color

Using pie charts

Use pie charts to compare values for items in a dimension. The best use for this type of chart is to show a percentage of a total amount.

Each wedge in a pie chart represents one item in the dimension. Wedge size represents the proportion of the value for the selected measure that the item represents compared to the whole

for the dimension. Pie charts are best when precision isn't important and there are few items in the dimension.

To create a donut chart, use one dimension in the **Group/Color** field well. With only one field, the chart displays the division of values by row count. To display the division of dimension values by a metric value, you can add a metric field to the **Value** field well.

Pie charts show up to 20 data points for group or color. For more information about how Amazon Quick Suite handles data that falls outside display limits, see [Display limits](#).

Pie chart features

To understand the features supported by pie charts, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes		Legends on visual types in Quick Suite
Changing the title display	Yes		Titles and subtitles on visual types in Quick Suite
Changing the axis range	Not applicable		Range and scale on visual types in Quick Suite
Showing or hiding axis labels.	Yes		Axes and grid lines on visual types in Quick Suite
Changing the visual colors	Yes		Colors in visual types in Quick Suite
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude a wedge in a pie chart, except when you are using a date field as a dimension. In that case, you can only focus on a wedge, not exclude it.	Focusing on visual elements Excluding visual elements

Feature	Supported?	Comments	For more information
Sorting	Yes	You can sort on the field that you choose for the value or the group or color.	Sorting visual data in Amazon Quick Suite
Performing field aggregation	Yes	You must apply aggregation to the field that you choose for the value, and can't apply aggregation to the field that you choose for group or color.	Changing field aggregation
Adding drill-downs	Yes	You can add drill-down levels to the Group/Color field well.	Adding drill-downs to visual data in Quick Sight

Creating a pie chart

Use the following procedure to create a pie chart.

To create a pie chart

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the pie chart icon.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value.

To create a pie chart, drag a dimension to the **Group/Color** field well. Optionally, drag a measure to the **Value** field well.

5. (Optional) Add drill-down layers by dragging one or more additional fields to the **Group/Color** field well. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

Using pivot tables

Use pivot tables to show measure values for the intersection of two dimensions.

Heat maps and pivot tables display data in a similar tabular fashion. Use a heat map if you want to identify trends and outliers, because the use of color makes these easier to spot. Use a pivot table if you want to analyze data on the visual.

To create a pivot table, choose at least one field of any data type, and choose the pivot table icon. Amazon Quick Suite creates the table and populates the cell values with the count of the column value for the intersecting row value. Typically, you choose a measure and two dimensions measurable by that measure.

Pivot tables support scroll down and right. You can add up to 20 fields as rows and 20 fields as columns. Up to 500,000 records are supported.

Using a pivot table, you can do the following:

- Specify multiple measures to populate the cell values of the table, so that you can see a range of data
- Cluster pivot table columns and rows to show values for subcategories grouped by related dimension
- Sort values in pivot table rows or columns
- Apply statistical functions
- Add totals and subtotals to rows and columns
- Use infinite scroll
- Transpose fields used by rows and columns
- Create custom total aggregations

To easily transpose the fields used by the rows and columns of the pivot table, choose the orientation icon

() near the top right of the visual. To see options for showing and hiding totals and subtotals, formatting the visual, or exporting data to a CSV file, choose the Menu items icon at top right.

As with all visual types, you can add and remove fields. You can also change the field associated with a visual element, change field aggregation, and change date field granularity. In addition, you

can focus on or exclude rows or columns. For more information about how to make these changes to a pivot table, see [Changing fields used by a visual in Amazon Quick Suite](#).

For information on formatting pivot tables, see [Formatting in Amazon Quick Suite](#).

For information on custom total aggregations for pivot tables, see [Custom total values](#).

Topics

- [Pivot table features](#)
- [Creating a pivot table](#)
- [Orienting pivot table values](#)
- [Expanding and collapsing pivot table clusters](#)
- [Showing and hiding pivot table columns in Quick Suite](#)
- [Sorting pivot tables in Quick Suite](#)
- [Using table calculations in pivot tables](#)
- [Pivot table limitations](#)
- [Pivot table best practices](#)

Pivot table features

Pivot tables don't display a legend.

To understand the features supported by pivot tables, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	No		Legends on visual types in Quick Suite
Changing the title display	Yes		Titles and subtitles on visual types in Quick Suite
Changing the axis range	Not applicable		Range and scale on visual types in Quick Suite
Changing the visual colors	No		Colors in visual types in Quick Suite

Feature	Supported?	Comments	For more information
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude any column or row, except when you are using a date field as one of the dimensions. In that case, you can only focus on the column or row that uses the date dimension, not exclude it.	Focusing on visual elements Excluding visual elements
Sorting	Yes	You can sort fields in the Rows or Columns field wells alphabetically or by a metric in ascending or descending order.	Sorting visual data in Amazon Quick Suite Sorting pivot tables in Quick Suite
Performing field aggregation	Yes	You must apply aggregation to the field or fields you choose for the value. You can't apply aggregation to the fields that you choose for the rows or columns. If you choose to create a multi-measure pivot table, you can apply different types of aggregation to the different measures. For example, you can show the sum of the sales amount and the maximum discount amount.	Changing field aggregation
Adding drill-downs	No		Adding drill-downs to visual data in Quick Sight

Feature	Supported?	Comments	For more information
Showing and hiding totals and subtotals	Yes	You can show or hide totals and subtotals for rows and columns. Metrics automatically roll up to show subtotals when you collapse a row or column. If you use a table calculation, use aggregates to display roll-ups.	
Exporting or copying data	Yes	You can export all of the data to a CSV file. You can select and copy the content of the cells.	Exporting data from visuals
Conditional formatting	Yes	You can add conditional formatting for values, subtotals and totals.	Conditional formatting on visual types in Quick Suite

Creating a pivot table

Use the following procedure to create a pivot table.

To create a pivot table

1. On the analysis page, choose the **Visualize** icon on the tool bar.
2. On the **Visuals** pane, choose **+ Add**, and then choose the pivot table icon.
3. From the **Fields list** pane, choose the fields that you want to include. Amazon Quick Suite automatically places these into the field wells.

To change the placement of a field, drag it to the appropriate field wells. Typically, you use dimension or measure fields as indicated by the target field well. If you choose to use a

dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value.

- To create a single-measure pivot table, drag a dimension to the **Rows** field well, a dimension to the **Columns** field well, and a measure to the **Values** field well.
- To create a multi-measure pivot table, drag a dimension to the **Rows** field well, a dimension to the **Columns** field well, and two or more measures to the **Values** field well.
- To create a clustered pivot table, drag one or more dimensions to the **Rows** field well, one or more dimensions to the **Columns** field well, and a measure to the **Values** field well.

You can also select multiple fields for all of the pivot table field wells if you want to. Doing this combines the multi-measure and clustered pivot table approaches.

Note

To view roll-ups for calculated fields, make sure that you are using aggregates. For example, a calculated field with `field-1 / field-2` doesn't display a summary when rolled up. However, `sum(field-1) / sum(field-2)` does display a roll-up summary.

Choosing a layout

When you create a pivot table in Amazon Quick Suite, you can further customize the way your data is presented with Tabular and Hierarchy layout options. For pivot tables that use a tabular layout, each row field is displayed in its own column. For pivot tables that use a hierarchy layout, all row fields are displayed in a single column. Indentation is used to differentiate row headers of different fields. To change the layout of a pivot table visual, open the **Format visual** menu of the pivot table that you want to change and choose the layout option that you want from the **Pivot options** section.

Depending on the layout that you choose for your pivot table visual, different formatting options are available. For more information about formatting differences between tabular and hierarchy pivot tables, see [Table and pivot table formatting options in Quick Suite](#).

Orienting pivot table values

You can choose to display a pivot table in a columnar or row-based format. Columnar is the default. When you change to a row-based format, a column with the value name is added to the right of the row header column.

To change a pivot table format

1. On the analysis page, choose the pivot table visual that you want to edit.
2. Expand the **Field wells** pane by choosing the field wells at the top of the visual.
3. On the **Values** field well, choose one of the following options:
 - Choose **Column** for a columnar format.
 - Choose **Row** for a row format.

Note

If you use only one metric, you can eliminate the repeated header by formatting the visual and styling it with the **Hide single metric** option.

Expanding and collapsing pivot table clusters

If you are using grouped columns or rows in a pivot table, you can expand or collapse a group to show or hide its data in the visual.

To expand or collapse a pivot table group

1. On the analysis page, choose the pivot table visual that you want to edit.
2. Choose one of the following:
 - To collapse a group, choose the collapse icon near the name of the field.
 - To expand a group, choose the expand icon near the name of the field. The collapse icon shows a minus sign. The expand icon shows a plus sign.

In the following screenshot, Customer Region and the Enterprise segment are expanded, and SMB and Startup are collapsed. When a group is collapsed, its data is summarized in the row or column.

Showing and hiding pivot table columns in Quick Suite

By default, all columns, rows, and their field values appear when you create a pivot table. You can hide columns and rows that you don't want to appear in the pivot table without changing the pivot table values. When you have more than one measure in the pivot table, you can also hide values.

At any time, you can choose to show any hidden fields in the pivot table. When you publish the visual as part of a dashboard, anyone who subscribes to the dashboard can export the pivot table to a comma-separated value (CSV) or Microsoft Excel file. They can choose to export only the visible fields, or all fields. For more information, see [Exporting data from a dashboard to a CSV](#).

To hide a column or row in a pivot table

1. In your analysis, select the pivot table visual that you want to work with.
2. Choose the three-dot menu in the **Rows**, **Columns** or **Values** field wells, and then choose **Hide**.

To show all hidden fields in a pivot table

1. In your analysis, select the pivot table visual that you want to work with.
2. Choose any field in the **Fields well** and choose **Show all hidden fields**.

Sorting pivot tables in Quick Suite

In Amazon Quick Suite, you can sort values in a pivot table by fields in the **Rows** and **Columns** field wells or quickly by column headers in the pivot table. In pivot tables, you can sort rows and columns independently of each other in alphabetical order, or by a measure.

Note

You can't run Total, Difference, and Percent Difference table calculations when a pivot table is being sorted by a measure. For more information about using table calculations in pivot tables, see [Using table calculations in pivot tables](#).

Understanding sorting in pivot tables

When you have multiple panes in a pivot table, sorting is applied to each pane independently. For example, the Segment column in the pivot table on the left is being sorted in ascending order by Cost. Given that there are multiple panes, the sort starts over for each pane and the rows within each pane (for Segment) are ordered by lowest to highest cost. The table on the right has the same sort applied, but the sort is being applied across the entire table, as shown following.

When you apply multiple sorts to a pivot table, sorting is applied from the outside dimension to the inside dimension. Consider the following example image of a pivot table. The Customer Region column is sorted by Cost in descending order (as shown in orange). The Channel column is sorted by Revenue Goal in ascending order (as shown in blue).

Sorting pivot tables using row or column headers

Use the following procedure to sort a pivot table using Row or Column headers.

To sort values in a tabular pivot table using table headers

1. In a tabular pivot table chart, choose the header that you want to sort.
2. For **Sort by**, choose a field to sort by and a sort order.

You can sort dimension fields alphabetically a–z or z–a, or you can sort them by a measure in ascending or descending order.

Sorting pivot tables using value headers

Use the following procedure to sort a pivot table using value headers.

To sort a pivot table using value headers

1. In a pivot table chart, choose the value header that you want to sort.
2. Choose **Ascending** or **Descending**.

Sorting by value headers in a pivot table also works on subtotals.

Sorting tabular pivot tables using the field wells

Use the following procedure to sort values in a tabular pivot table using the field wells.

To sort values in a tabular pivot table using the field wells

1. On the analysis page, choose the tabular pivot table that you want to sort.
2. Expand the **Field wells**.
3. In the **Rows** or **Columns** field well, choose the field that you want to sort, and then choose how you want to sort the field for **Sort by**.

You can sort dimension fields in the **Rows** or **Columns** field wells alphabetically from a–z or z–a, or you can sort them by a measure in ascending or descending order. You also have the option to collapse all or expand all rows or columns for the field you choose in the field well. You can also remove the field, or to replace it with another field.

- To sort a dimension field alphabetically, hover your cursor over the field in the **Rows** or **Columns** field well, and then choose the a–z or z–a sort icon.
- To sort a dimension field by a measure, hover your cursor over the field in the **Rows** or **Columns** field well. Then choose a measure from the list, and then choose the ascending or descending sort icon.

Or, if you want more control over how the sort is applied to the pivot table, customize the sort options.

To create a sort using the sort options

1. On the analysis page, choose the pivot table that you want to sort.
2. Expand **Field wells**.
3. Choose the field that you want to sort in the **Rows** or **Columns** field well, and then choose **Sort options**.
4. In the **Sort options** pane that opens at left, specify the following options:
 - a. For **Sort by**, choose a field from the drop-down list.
 - b. For **Aggregation**, choose an aggregation from the list.

- c. For **Sort order**, select **Ascending** or **Descending**.
- d. Choose **Apply**.

Sorting hierarchy pivot tables using the field wells

For tabular pivot tables, each field in the **Rows** field well has a separate title cell. For hierarchy pivot tables, all row fields are displayed in a single column. To sort, collapse, and expand these row fields, select the **Rows** label to open the **Combined row fields** menu and choose the option that you want. Each field in a hierarchy pivot table can be individually sorted from the **Combined row fields** menu.

More advanced formatting options such as **Hide** and **Remove** are available from the field well menus.

Using table calculations in pivot tables

You can use table calculations to apply statistical functions to pivot table cells that contain measures (numeric values). Use the following sections to understand which functions you can use in calculations, and how to apply or remove them.

The data type of the cell value automatically changes to work for your calculation. For example, say that you apply the **Rank** function to a currency data type. The values display as integers rather than currency, because rank isn't measured as currency. Similarly, if you apply the **Percent difference** function instead, the cell values display as percentages.

Topics

- [Adding and deleting pivot table calculations](#)
- [Functions for pivot table calculations](#)
- [Ways to apply pivot table calculations](#)

Adding and deleting pivot table calculations

Use the following procedures to add, modify, and remove table calculation on a pivot table.

Topics

- [Adding a pivot table calculation](#)
- [Changing how a calculation is applied](#)

- [Removing a calculation](#)

Adding a pivot table calculation

Use the following procedure to add a table calculation to a pivot table.

To add a table calculation to a pivot table

1. Expand the **Field wells** pane by choosing the field wells near the bottom of the visual.
2. Choose the field in the **Values** well that you want to apply a table calculation to, choose **Add table calculation**, and then choose the function to apply.

Note

You can't run Total, Difference, and Percent Difference table calculations when a pivot table is being sorted by a measure. To use these table calculations, remove the sort from the pivot table.

Changing how a calculation is applied

Use the following procedure to change the way a table calculation is applied to a pivot table.

To change the way a table calculation is applied to a pivot table

1. Expand the **Field wells** pane by choosing field wells at the top of the visual.
2. Choose the field in the **Values** well that has the table calculation that you want to change, choose **Calculate as**, and then choose the way that you want the calculation applied.

Removing a calculation

Use the following procedure to remove a table calculation from a pivot table.

To remove a table calculation from a pivot table

1. Expand the **Field wells** pane by choosing the field wells near the bottom of the visual.
2. Choose the field in the **Values** well that you want to remove the table calculation from, and then choose **Remove**.

Functions for pivot table calculations

You can use the following functions in pivot table calculations.

Topics

- [Running total](#)
- [Difference](#)
- [Percentage difference](#)
- [Percent of total](#)
- [Rank](#)
- [Percentile](#)

You can apply functions listed to the following data:

Running total

The **Running total** function calculates the sum of a given cell value and the values of all cells prior to it. This sum is calculated as $Cell1=Cell1$, $Cell2=Cell1+Cell2$, $Cell3=Cell1+Cell2+Cell3$, and so on.

Applying the **Running total** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Difference

The **Difference** function calculates the difference between a cell value and value of the cell prior to it. This difference is calculated as $Cell1=Cell1-null$, $Cell2=Cell2-Cell1$, $Cell3=Cell3-Cell2$, and so on. Because $Cell1-null = null$, the Cell1 value is always empty.

Applying the **Difference** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Percentage difference

The **Percentage Difference** function calculates the percent difference between a cell value and the value of the cell prior to it, divided by the value of the cell prior to it. This value is calculated as $\text{Cell1}=(\text{Cell1}-\text{null})/\text{null}$, $\text{Cell2}=(\text{Cell2}-\text{Cell1})/\text{Cell1}$, $\text{Cell3}=(\text{Cell3}-\text{Cell2})/\text{Cell2}$, and so on. Because $(\text{Cell1}-\text{null})/\text{null} = \text{null}$, the Cell1 value is always empty.

Applying the **Percentage Difference** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Percent of total

The **Percent of Total** function calculates the percentage the given cell represents of the sum of all of the cells included in the calculation. This percentage is calculated as $\text{Cell1}=\text{Cell1}/(\text{sum of all cells})$, $\text{Cell2}=\text{Cell2}/(\text{sum of all cells})$, and so on.

Applying the **Percent of Total** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Rank

The **Rank** function calculates the rank of the cell value compared to the values of the other cells included in the calculation. Rank always shows the highest value equal to 1 and lowest value equal to the count of cells included in the calculation. If there are two or more cells with equal values, they receive the same rank but are considered to take up their own spots in the ranking. Thus, the next highest value is pushed down in rank by the number of cells at the rank above it, minus one. For example, if you rank the values 5,3,3,4,3,2, their ranks are 1,3,3,2,3,6.

For example, suppose that you have the following data.

Applying the **Rank** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Percentile

The **Percentile** function calculates the percent of the values of the cells included in the calculation that are at or below the value for the given cell.

This percent is calculated as follows.

```
percentile rank(x) = 100 * B / N
```

Where:

B = number of scores below x

N = number of scores

Applying the **Percentile** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Ways to apply pivot table calculations

You can apply table calculations in the ways described following. Table calculations are applied to only one field at a time. Thus, if you have a pivot table with multiple values, calculations are only applied to the cells representing the field that you applied the calculation to.

Topics

- [Table across](#)
- [Table down](#)
- [Table across down](#)
- [Table down across](#)
- [Group across](#)
- [Group down](#)
- [Group across down](#)
- [Group down across](#)

Table across

Using **Table across** applies the calculation across the rows of the pivot table, regardless of any grouping. This application is the default. For example, take the following pivot table.

Applying the **Running total** function using **Table across** gives you the following results, with row totals in the last column.

Table down

Using **Table down** applies the calculation down the columns of the pivot table, regardless of any grouping.

Applying the **Running total** function using **Table down** gives you the following results, with column totals in the last row.

Table across down

Using **Table across down** applies the calculation across the rows of the pivot table, and then takes the results and reapplies the calculation down the columns of the pivot table.

Applying the **Running total** function using **Table across down** gives you the following results. In this case, totals are summed both down and across, with the grand total in the lower-right cell.

In this case, suppose that you apply the **Rank** function using **Table across down**. Doing so means that the initial ranks are determined across the table rows and then those ranks are in turn ranked down the columns. This approach gives you the following results.

Table down across

Using **Table down across** applies the calculation down the columns of the pivot table. It then takes the results and reapplies the calculation across the rows of the pivot table.

You can apply the **Running total** function using **Table down across** to get the following results. In this case, totals are summed both down and across, with the grand total in the lower-right cell.

You can apply the **Rank** function using **Table down across** to get the following results. In this case, the initial ranks are determined down the table columns. Then those ranks are in turn ranked across the rows.

Group across

Using **Group across** applies the calculation across the rows of the pivot table within group boundaries, as determined by the second level of grouping applied to the columns. For example,

if you group by field-2 and then by field-1, grouping is applied at the field-2 level. If you group by field-3, field-2, and field-1, grouping is again applied at the field-2 level. When there is no grouping, **Group across** returns the same results as **Table across**.

For example, take the following pivot table where columns are grouped by Service Line and then by Consumption Channel.

You can apply the **Running total** function using **Group across** to get the following results. In this case, the function is applied across the rows, bounded by the columns for each service category group. The Mobile columns display the total for both Consumption Channel values for the given Service Line, for the Customer Region and Date (year) represented by the given row. For example, the highlighted cell represents the total for the APAC region for 2012, for all Consumption Channel values in the Service Line named Billing.

Group down

Using **Group down** applies the calculation down the columns of the pivot table within group boundaries, as determined by the second level of grouping applied to the rows. For example, if you group by field-2 and then by field-1, grouping is applied at the field-2 level. If you group by field-3, field-2, and field-1, grouping is again applied at the field-2 level. When there is no grouping, **Group down** returns the same results as **Table down**.

For example, take the following pivot table where rows are grouped by Customer Region and then by Date (year).

You can apply the **Running total** function using **Group down** to get the following results. In this case, the function is applied down the columns, bounded by the rows for each Customer Region group. The 2014 rows display the total for all years for the given Customer Region, for the Service Line and Consumption Channel represented by the given column. For example, the highlighted cell represents the total the APAC region, for the Billing service for the Mobile channel, for all the Date values (years) that display in the report.

Group across down

Using **Group across down** applies the calculation across the rows within group boundaries, as determined by the second level of grouping applied to the columns. Then the function takes the

results and reapplies the calculation down the columns of the pivot table. It does so within group boundaries as determined by the second level of grouping applied to the rows.

For example, if you group a row or column by field-2 and then by field-1, grouping is applied at the field-2 level. If you group by field-3, field-2, and field-1, grouping is again applied at the field-2 level. When there is no grouping, **Group across down** returns the same results as **Table across down**.

For example, take the following pivot table where columns are grouped by Service Line and then by Consumption Channel. Rows are grouped by Customer Region and then by Date (year).

You can apply the **Running total** function using **Group across down** to get the following results. In this case, totals are summed both down and across within the group boundaries. Here, these boundaries are Service Line for the columns and Customer Region for the rows. The grand total appears in the lower-right cell for the group.

You can apply the **Rank** function using **Group across down** to get the following results. In this case, the function is first applied across the rows bounded by each Service Line group. The function is then applied again to the results of that first calculation, this time applied down the columns bounded by each Customer Region group.

Group down across

Using **Group down across** applies a calculation down the columns within group boundaries, as determined by the second level of grouping applied to the rows. Then Amazon Quick Suite takes the results and reapplies the calculation across the rows of the pivot table. Again, it reapplies the calculation within group boundaries as determined by the second level of grouping applied to the columns.

For example, if you group a row or column by field-2 and then by field-1, grouping is applied at the field-2 level. If you group by field-3, field-2, and field-1, grouping is again applied at the field-2 level. When there is no grouping, **Group down across** returns the same results as **Table down across**.

For example, take the following pivot table. Columns are grouped by Service Line and then by Consumption Channel. Rows are grouped by Customer Region and then by Date (year).

You can apply the **Running total** function using **Group down across** to get the following results. In this case, totals are summed both down and across within the group boundaries. In this case, these are `Service Category` for the columns and `Customer Region` for the rows. The grand total is in the lower-right cell for the group.

You can apply the **Rank** function using **Group down across** to get the following results. In this case, the function is first applied down the columns bounded by each `Customer Region` group. The function is then applied again to the results of that first calculation, this time applied across the rows bounded by each `Service Line` group.

Pivot table limitations

The following limitations apply to pivot tables:

- You can create pivot tables with up to 500,000 records.
- You can add any combination of row and column field values that add up to 40. For example, if you have 10 row field values, then you can add up to 30 column field values.
- You can create pivot table calculations only on nonaggregated values. For example, if you create a calculated field that is a sum of a measure, you can't also add a pivot table calculation to it.
- If you are sorting by a custom metric, you can't add a table calculation until you remove the custom metric sort.
- If you are using a table calculation and then add a custom metric, you can't sort by the custom metric.
- Totals and subtotals are blank for table calculations on metrics aggregated by distinct count.

Pivot table best practices

It's best to deploy a minimal set of rows, columns, metrics, and table calculations, rather than offering all possible combinations in one pivot table. If you include too many, you risk overwhelming the viewer and you can also run into the computational limitations of the underlying database.

To reduce the level of complexity and reduce the potential for errors, you can take the following actions:

- Apply filters to reduce the data included in for the visual.
- Use fewer fields in the **Row** and **Column** field wells.
- Use as few fields as possible in the **Values** field well.
- Create additional pivot tables so that each displays fewer metrics.

In some cases, there's a business need to examine many metrics in relation to each other. In these cases, it can be better to use multiple visuals on the same dashboard, each showing a single metric. You can reduce the size of the visuals on the dashboard, and colocate them to form a grouping. If a decision the viewer makes based on one visual creates the need for a different view, you can deploy custom URL actions to launch another dashboard according to the choices made by the user.

It's best to think of visuals as building blocks. Rather than using one visual for multiple purposes, use each visual to facilitate one aspect of a larger business decision. The viewer should have enough data to make a well-informed decision, without being overwhelmed by the inclusion of all possibilities.

Using radar charts

You can use radar charts, which are also known as spider charts, to visualize multivariate data in Amazon Quick Suite. In a radar chart, one or more groups of values are plotted over multiple common variables. Each variable has its own axis, and each axis is arranged radially around a central point. The data points from a single observation are plotted on each axis and connected to each other to form a polygon. Multiple observations can be plotted in a single radar chart to display multiple polygons, which makes it easier to spot outlying values for multiple observations quickly.

In Quick Suite, you can organize a radar chart along its category, value, or color axes by dragging and dropping fields to the **Category**, **Value**, and **Color** field wells. How you choose to distribute fields among the field wells determines the axis that the data is plotted on.

The following image shows an example of a radar chart.

Radar chart features

To view the features supported by radar charts, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes		Legends on visual types in Quick Suite
Changing the title display	Yes		Titles and subtitles on visual types in Quick Suite
Changing the axis range	Yes		Range and scale on visual types in Quick Suite
Changing the visual colors	Yes		Colors in visual types in Quick Suite
Focusing on or excluding elements	Yes		Focusing on visual elements Excluding visual elements
Sorting	Limited	You can only sort data fields that are in the Category and Color field wells.	Sorting visual data in Amazon Quick Suite
Performing field aggregation	Yes		Changing field aggregation
Adding drill-downs	Not supported		Adding drill-downs to visual data in Quick Sight
Choosing size	Yes		Formatting in Amazon Quick Suite
Showing totals	Not supported		Formatting in Amazon Quick Suite

Creating a radar chart

Use the following procedure to create a radar chart.

To create a radar chart

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the radar chart icon.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. In most cases, you want to use dimension or measure fields as indicated by the target field well.

To create a radar chart, drag fields to the **Category**, **Value**, and **Group/Color** field wells. The axis that a radar chart is organized around is determined by the way that you organize fields into their respective field wells:

- In a radar chart that uses a **value axis**, dimension values are shown as lines and axes represent value fields. To create a radar chart that uses a value axis, add one category field to the **Color** field well and one or more values to the **Value** field well.
- In a radar chart that uses a **dimension axis**, group dimension values are shown as axes and value fields are shown as lines. All axes share a range and scale. To create a radar chart that uses a dimension axis, add one dimension to the **Group** field well and one or more values to the **Value** field well.
- In a radar chart that uses a **dimension-color axis**, group dimension values are shown as axes and color dimension values are shown as lines. All axes share a range and scale. To create a radar chart that uses a dimension-color axis, add one dimension to the **Category** field well, one value to the **Value** field well, and one dimension to the **Color** field well.

Using Sankey diagrams

Use Sankey diagrams to show flows from one category to another, or paths from one stage to the next.

For example, a Sankey diagram can show the number of people migrating from one country to another. A Sankey diagram can also show the path a web visitor takes from one page to the next on a company website, with possible stops along the way.

Data for Sankey diagrams

To create Sankey diagrams in Quick Suite, your dataset should contain a measure and two dimensions (one dimension containing source categories and another containing destination categories).

The following table is a simple example of data for a Sankey diagram.

Dimension (Source)	Dimension (Destination)	Measure (Weight)
A	W	500
A	X	23
A	Y	147

The following Sankey diagram is created when the dimensions and measure are added to the field well, with the A node on the left linking to the W, Y, and X nodes on the right. The width of each link between nodes is determined by the value in the Measure (Weight) column. The nodes are automatically ordered.

To create multilevel Sankey diagrams in Amazon Quick Suite, your dataset should still contain a measure and two dimensions (one for source and one for destination), but in this case your data values differ.

The following table is a simple example of data for a multilevel Sankey diagram with two stages.

Dimension (Source)	Dimension (Destination)	Measure (Weight)
A	W	500
A	X	23
A	Y	147
W	Z	300
X	Z	5

Dimension (Source)	Dimension (Destination)	Measure (Weight)
Y	Z	50

The following Sankey diagram is created when the dimensions and measure are added to the field well. Here, the A node on the left links to the W, Y, and X nodes in the middle, and the W, Y, and X nodes then link to the Z node on the right. The width of each link between nodes is determined by the value in the Measure (Weight) column.

Working with cyclical data

Sometimes, the data that you use for a Sankey diagram contains cycles. For example, suppose that you're visualizing user traffic flows between pages on a website. You might discover that users who come to page A move to page E, and then come back to page A. An entire flow might look something like A-E-A-B-A-E-A.

When your data contains cycles, the nodes in each cycle are repeated in Quick Suite. For example, if your data contains the flow A-E-A-B-A-E-A, the following Sankey diagram is created.

Preparing data for Sankey diagrams

If your dataset doesn't contain Source or Destination columns, prepare your data to include them. You can prepare data when creating a new dataset, or when editing an existing dataset. For more information about creating a new dataset and preparing it, see [Creating datasets](#). For more information about opening an existing dataset for data preparation, see [Editing datasets](#).

The following procedure uses an example table (illustrated in following) to demonstrate how to prepare your data for Sankey diagrams in Quick Suite. The table includes three columns: Customer ID, Time, and Action.

Customer ID	Time	Action
1	9:05 am	Step 1
1	9:06 am	Step 2

Customer ID	Time	Action
1	9:08 am	Step 3
2	11:44 am	Step 1
2	11:47 am	Step 2
2	11:48 am	Step 3

To create a Sankey diagram in Quick Suite using this data, first add Source and Destination columns to the table. Use the following procedure to learn how.

To add Source and Destination columns to your table

1. Add a Step Number column to the table to number or rank each row.

There are multiple ways to compute the Step Number column. If your data source is compatible with SQL and your database supports ROW_NUMBER or RANK functions, you can use custom SQL in Quick Suite to order the rows in the Step Number column. For more information about using custom SQL in Quick Suite, see [Using SQL to customize data](#).

Customer ID	Time	Action	Step Number
1	9:05 am	Step 1	1
1	9:06 am	Step 2	2
1	9:08 am	Step 3	3
2	11:44 am	Step 1	1

Customer ID	Time	Action	Step Number
2	11:47 am	Step 2	2
2	11:48 am	Step 3	3

2. Add a Next Row Number column to the table with values equal to Step Number plus one.

For example, in the first data row of the table, the value for Step Number is 1. To compute the value for Next Step Number for that row, add 1 to that value.

$$1 + 1 = 2$$

The value for Step Number in the second data row of the table is 2; therefore, the value for Next Step Number is 3.

$$2 + 1 = 3$$

Customer ID	Time	Action	Step Number	Next Step Number
1	9:05 am	Step 1	1	2
1	9:06 am	Step 2	2	3
1	9:08 am	Step 3	3	4
2	11:44 am	Step 1	1	2
2	11:47 am	Step 2	2	3
2	11:48 am	Step 3	3	4

3. Join the table with itself:

- a. For **Join type**, choose **Inner**.
- b. For **Join clauses**, do the following:
 - i. Choose **Customer ID = Customer ID**
 - ii. Choose **Next Step Number = Step Number**

For more information about joining data in Quick Suite, see [Joining data](#).

Joining the two tables creates two columns for Customer ID, Time, Action, Step Number and Next Step Number. The columns from the table at the left of the join are Source columns. The columns from the table at the right of the join are Destination columns.

4. (Optional) Rename columns to indicate sources and destinations.

The following is an example:

1. Rename the **Action** column on the left to **Source**.
2. Rename the **Action [copy]** column on the right to **Destination**.
3. Rename the **Time** column on the left to **Start Time**.
4. Rename the **Time [copy]** column on the right to **End Time**.

Your data is now ready to visualize.

Creating Sankey diagrams

Use the following procedure to create a Sankey diagram.

To create a Sankey diagram

1. On the analysis screen, choose **Visualize** on the left toolbar.
2. On the application bar, choose **Add**, and then choose **Add visual**.
3. On the **Visual types** pane, choose the Sankey diagram icon.
4. On the menu in the upper-right corner of the visual, choose the **Properties** icon.
5. In the **Properties** pane, choose either the **Source** or **Destination** section.

Customizing the number of nodes

Use the following procedure to customize the number of nodes that appear in a Sankey diagram. Quick Suite supports up to 100 Source/Destination nodes.

To customize the number of nodes that appear in a Sankey diagram

1. On the analysis page, choose the Sankey diagram visual that you want to format.
2. On the menu in the upper-right corner of the visual, select the **Format Visual** icon.
3. In the **Properties** pane that opens, choose either the **Source** or **Destination** tab.
4. For **Number of nodes displayed**, enter a number.

The nodes in the diagram update to the number that you specified. The top nodes are automatically shown. All other nodes are placed in an **Other** category.

 Note

Specifying the number of Source nodes controls how many Source nodes can appear overall in the diagram. Specifying the number of Destination nodes controls how many Destination nodes can appear per Source node. This means that if there is more than one Source node in your diagram, the overall number of Destination nodes will be higher than the number specified.
Quick Suite supports up to 100 Source/Destination nodes.

For example, the following Sankey diagram has a limit of three source nodes (out of five), so the top three are shown in the diagram. The other two source nodes are placed in the Other category.

To remove the **Other** category from the diagram, select it in the view and choose **Hide “other” categories**.

Sankey diagram features

To understand the features supported by Sankey diagrams, use the following table.

Feature	Supported?	For more information	
Changing the legend display	No		
Changing the title display	Yes	Titles and subtitles on visual types in Quick Suite	
Changing the axis range	No		
Changing the visual colors	No		
Focusing on or excluding elements	Yes	Focusing on visual elements Excluding visual elements	
Sorting	No		
Performing field aggregation	Yes	Changing field aggregation	
Adding drill-downs	No		
Conditional formatting	No		

Using scatter plots

Use scatter plots to visualize two or three measures across two dimensions.

Each bubble on the scatter plot represents one or two dimension values. The X and Y axes represent two different measures that apply to the dimension. A bubble appears on the chart at the point where the values for the two measures for an item in the dimension intersect. Optionally, you can also use bubble size to represent an additional measure.

Scatter plots show up to 2500 datapoints in aggregated and unaggregated scenarios regardless of whether a color or label dimension is used in the visual. Due to the order of limit operations, there may be cases where fewer datapoints for a dataset are shown. For more information about how Amazon Quick Suite handles data that falls outside display limits, see [Display limits](#).

Scatter plot features

To understand the features supported by scatter plots, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes, with exceptions	Scatter plots display a legend if you have the Group/Color field well populated.	Legends on visual types in Quick Suite
Changing the title display	Yes		Titles and subtitles on visual types in Quick Suite
Changing the axis range	Yes	You can set the range for both the X and Y axes.	Range and scale on visual types in Quick Suite
Showing or hiding axis lines, grid lines, axis labels, and axis sort icons	Yes		Axes and grid lines on visual types in Quick Suite
Changing the visual colors	Yes		Colors in visual types in Quick Suite
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude a bubble in a scatter plot, except when you are using a date field as a dimension. In that case, you can only focus on a bubble, not exclude it.	Focusing on visual elements Excluding visual elements

Feature	Supported?	Comments	For more information
Sorting	No		Sorting visual data in Amazon Quick Suite
Performing field aggregation	Yes	You must apply aggregation to the fields you choose for the X axis, Y axis, and size, and can't apply aggregation to the field that you choose for the group or color.	Changing field aggregation
Displaying unaggregated fields	Yes	On the field context menu, choose None to display unaggregated X and Y axis values. If your scatter plot shows unaggregated fields, you can't apply aggregations to the field that is in the color or label field well. Mixed aggregation is not supported for scatter plots.	
Adding drill-downs	Yes	You can add drill-down levels to the Group/Color field well.	Adding drill-downs to visual data in Quick Sight

Creating a scatter plot

Use the following procedure to create a scatter plot.

To create a scatter plot

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the scatter plot icon.

4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value.

To create a scatter plot, drag a measure to the **X axis** field well, a measure to the **Y axis** field well, and a dimension to the **Color** or **Label** field well. To represent another measure with bubble size, drag that measure to the **Size** field well.

5. (Optional) Add drill-down layers by dragging one or more additional fields to the **Color** field well. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

Scatter plot use cases

You can choose to plot unaggregated values even if you are using a field on Color by using the aggregate option **none** on the field menu, which also contains aggregation options like **sum**, **min**, and **max**. If one value is set to be aggregated, the other value will be automatically set as aggregated. The same applies to unaggregated scenarios. Mixed aggregation scenarios are not supported, meaning that one value cannot be set as aggregated while the other is unaggregated. Note that the unaggregated scenario, which is the **none** option, is supported only for numerical values, while categorical values, such as dates or dimensions, will display only aggregate values, such as **count** and **count distinct**.

Using the **none** option, you can choose to set both X and Y values to either aggregated or unaggregated from the **X axis** and **Y axis** field menus. This will define whether or not values will be aggregated by dimensions in the **Color** and **Label** field wells. To get started, add the required fields and choose the appropriate aggregation based on your use case, as shown in the following sections.

Unaggregated use cases

- Unaggregated X and Y values with Color
- Unaggregated X and Y values with Label
- Unaggregated X and Y values with Color and Label

Aggregated use cases

- Aggregated X and Y values with Color
- Aggregated X and Y values with Label
- Aggregated X and Y values with Color and Label

Using tables as visuals

Use a table visual to see a customized table view of your data. To create a table visual, choose at least one field of any data type. You can add as many columns as you need, up to 200. You can also add calculated columns.

Table visuals don't display a legend. You can hide or display the title on a table. You can also hide or display totals, and choose to show totals at the top or the bottom of the table. For more information, see [Analytics formatting per type in Quick Suite](#).

To create a table visual

1. Open Amazon Quick Suite and choose **Analyses** on the navigation pane at left.
2. Choose one of the following:
 - To create a new analysis, choose **New analysis** at upper right. For more information, see [Starting an analysis in Quick Sight](#).
 - To use an existing analysis, choose the analysis that you want to edit.
3. Choose **Insert** from the file menu and then **Add Visual**.
4. At lower left, choose the table icon from **Visual types**.
5. On the **Fields** list pane, choose the fields that you want to use. If you want to add a calculated field, choose **Insert** on the file menu and then **Add Calculated Field**.

To create a nonaggregated view of the data, add fields only to the **Value** field well. Doing this shows data without any aggregations.

To create an aggregated view of the data, choose the fields that you want to aggregate by, and then add them to the **Group by** field well.

To show or hide columns on a table

1. On your visual, choose the field that you want to hide, then choose **Hide column**.
2. To display hidden columns, choose any column, then choose **Show all hidden columns**.

To transpose columns to rows and rows to columns

- Choose the transpose icon () near the top right of the visual. It has two arrows at a 90 degree angle.

To vertically align columns

1. On your visual, choose the **Format visual** icon () near the top right of the visual.
2. In the **Properties** pane, choose **Table options**, and choose your table's vertical alignment.

To wrap the text for headers

1. On your visual, choose the **Format visual** icon () near the top right of the visual.
2. In the **Properties** pane, choose **Table options**, and select **Wrap header text**.

To rearrange columns in a table chart

1. Open the analysis with the visual that you want to sort. Visuals pane will be open by default.
2. Do one of the following:
 - Drag and drop one or more fields in **Field wells** to rearrange their order.
 - Select a field directly in the table and choose the left or right arrow on **Move column**.

Using field styling

You can render URLs in a table as links by using the **Field styling** pane of the format visual menu. You can add up to 500 rows of links for each page in a table. Only https and mailto hyperlinks are supported.

To add links to your tables

1. From the Quick Suite homepage, choose **Analyses**, and then choose the analysis that you want to customize.
2. Choose the table that you want to change.
3. On the menu at the upper right of the table, choose **Format visual**.
4. For **Format visual**, choose **Field styling**.
5. On the **Field styling** pane, choose the field that you want to style from the menu.
6. In the **Url options** section of the **Field styling** menu, choose **Make URLs hyperlinks**.

After you add links to your table, you can choose where you want the links to open when they're selected in the **Open in** section of the **Field style** pane. You can choose to have links open in a new tab, a new window, or in the same tab.

You can also choose how you want to style the link in the **Style as** section of **Field style** pane. Your links can appear as hyperlinks, icons, or plain text, or you can set a custom link.

To adjust the font size of a link icon or URL, change the **Font size** in the **Cells** section of the **Table options** pane of the **Format visual** menu.

You can set any URLs in your table that point to images to render in the table as images. Doing this can be useful when you want to include an image of a product as a part of a table.

To show URLs as images

1. From the Quick Suite home page, choose **Analyses**, and then choose the analysis that you want to customize.
2. Choose the table that you want to change.
3. On the menu at the upper-right of the table, choose **Format visual**.
4. In the **Format visual** menu, choose **Field styling**.
5. In the **Field styling** pane, choose the field that you want to style from the menu.

6. In the **Url options** section of the **Field styling** menu, choose **Show URLs as images**.

After rendering images in a table, you can choose how to size the images in the **Image sizing** section of the **Field style** pane. You can fit images to their cell's height or width, or you can choose not to scale the image. Images fit to a cell's height by default.

Freeze columns to table visuals

You can freeze columns to your table visuals to lock specific columns in place on screen. This allows essential information to remain visible while readers scroll across the table. You can freeze columns one at a time, or you can freeze groups of columns in one action. All pinned columns are fixed to the far left side of the table and stay visible on screen at all times. This allows Quick Suite readers to have a constant reference point for key data or information as they interact with other parts of the table.

To freeze columns to a table

1. On the table that you want to freeze a column to, choose the column that you want to pin.
2. Choose one of the following options.
 - To freeze a single column, choose **Freeze column**.
 - To freeze all columns up to the column that you choose, choose **Freeze up to this column**.

If your table has multiple pinned columns, you can reorder the columns in the order that you want. To adjust the order of the pinned columns on a table, choose the header of the column that you want to move, and then choose **Move** in the direction that you want.

To unfreeze columns from a table

1. On the table that you want to change, choose the pinned column that you want to unpin.
2. Choose one of the following options.
 - To unfreeze a single column, choose **Unfreeze column**.
 - To unfreeze all frozen columns, choose **Unfreeze all columns**.

Custom total values

Quick Suite authors can define the total and subtotal aggregations for their table or pivot table visuals from the field wells. For tables, the custom total menu is only available if totals are toggled on for the visual.

To change the aggregation of a total or subtotal

1. Navigate to the analysis that you want to change, and choose the table or pivot table visual whose total you want to define.
2. Choose the field that you want to change from the field wells.
3. Choose **Total**, and then choose the aggregation that you want. The following options are available.
 - **Default** – The total calculation uses the same aggregation as the metric field.
 - **Sum** – Calculates the sum of the data in the visual.
 - **Average** – Calculates the average of the data in the visual.
 - **Min** – Calculates the minimum value of the data in the visual.
 - **Max** – Calculates the maximum value of the data in the visual.
 - **None (HIDE)** – Totals are not calculated. When you choose this option, the total and subtotal cells in the visual are left blank. If the outer dimension is sorted with the metric field that calculates the total or subtotal, the dimension is sorted alphabetically. When you change the value from **None (HIDE)** to another value, the outer dimension is sorted by the subtotals that are calculated with the specified aggregation type.

The following limitations apply to custom totals.

- Conditional formatting is not supported for custom totals.
- Total aggregations aren't supported for string columns. Total aggregations include **Min**, **Max**, **Sum**, and **Average**.
- Date columns are incompatible with **Average** and **Sum** total aggregation functions.

Sorting tables

In Amazon Quick Suite, you can sort values in a table by fields in the columns headers of the table or with the **Sort visual** tool. You can sort up to 10 columns in a single table. Quick Suite can also

use an off-visual sort You can sort columns in an **Ascending** or a **Descending** order. The following image shows the **Sort visual** icon and pop over.

Single column sort options

Quick Suite Authors can access single column sort options from the field wells, the column headers, or from the **Sort visual** menu. Use the procedure below to use set up a single column sort on a table in Quick Suite.

1. Open the [Quick Suite console](#).
2. Open the analysis that you want to work in and navigate to the table that you want to sort.
3. Choose the header of the column that you want to sort.
4. For **Sort by**, choose the arrow icon, and then choose the field that you want to sort by.

You can also set up a single column sort in the **Sort visual** menu. To access the sort visual menu, choose the **Sort visual** icon in the on-visual menu. In the **Sort visual** menu, choose the field that you want to sort by, and then choose if you want the sort in an ascending or descending order. By default, new sorts are sorted in an ascending order. When you are finished, choose **APPLY**.

Tables that use single column sorting are sorted one column at a time. When a user chooses a new column to sort by, the previous sort order is overridden.

To make changes to a single column sort, open the **Sort visual** menu annd use the dropdown menus to choose a new field or sort order. When you are finished with your changes, choose **APPLY**.

To reset a table to its original state, open the **Sort visual** menu and choose **RESET**.

Multi column sort options

Quick Suite authors can access multi column sort options from the **Sort visual** menu. Use the procedure below to set up a multi column sort for a table.

1. Open the [Quick Suite console](#).
2. Open the analysis that you want to work in and navigate to the table that you want to sort.
3. Choose the **Sort visual** icon to open the **Sort visual** menu.
 - a. Alternatively, choose a header that you want to sort.

- b. For **Sort by**, choose the arrow icon, and then choose **Multiple fields**.
4. In the **Sort visual** menu that opens, choose a field from the **Sort by** dropdown, and then choose whether you want the field sorted in an ascending or descending order.
5. To add another sort, choose **ADD SORT**, and repeat the workflow from Step 4. You can add up to 10 sorts to each table.
6. When you are finished, choose **APPLY**.

Columns are sorted in the order that they are added to the **Sort visual** menu. To change the order that columns are sorted by, open the **Sort visual** menu and use the **Sort by** dropdowns to reorder the sorts. When you are finished, choose **APPLY** to apply the new sort order to the table.

To reset a table to its original state, open the **Sort visual** menu and choose **RESET**.

Off visual sort options

Quick Suite authors can configure an off-visual sort to sort the values in a table by a field and aggregation that is a part of the dataset that the table uses but not in one of the table's field wells. One off-field sort can be configured to a single table at a time.

Use the procedure below to configure an off-visual sort.

To add an off-visual sort to a table

1. Open the [Quick Suite console](#).
2. Open the analysis that you want to work in and navigate to the table that you want to sort.
3. Choose the header of any column in the table.
4. For **Sort by**, choose the arrow icon, and then choose **Off-visual field**.
5. In the **Off-visual field** pane that appears, open the **Sort by** dropdown menu and choose the field that you want to sort.
6. For **Aggregation** open the dropdown menu and choose the aggregation that you want to use.
7. For **Sort order**, choose if you want the sort to be in an ascending or descending order.
8. When you are finished, choose **Apply**.

After a off-visual sort is applied to a table, the sort is shown in the **Sort visual** menu. The sort order of a table that contains an off-visual sort depends on the sort configuration of the table

when the off-visual sort is added. If an off-visual sort is added to a table that already has a single or multi column sort configured, the off-visual sort overrides all other sorts. If the off-visual sort is applied before single or multi column sorts, you can add and reorder more sorts to the table.

Using text boxes

Add text to add context to sheets in an analysis by using a text box. Text can hold directions, descriptions, or even hyperlinks to external websites. The toolbar on the text box offers font settings so you can customize the font type, style, color, size, spacing, size in pixels, text highlights, and alignment. The text box itself has no format settings.

To add text to a new text box, simply select it and begin typing.

Using tree maps

To visualize one or two measures for a dimension, use tree maps.

Each rectangle on the tree map represents one item in the dimension. Rectangle size represents the proportion of the value for the selected measure that the item represents compared to the whole for the dimension. You can optionally use rectangle color to represent another measure for the item. Rectangle color represents where the value for the item falls in the range for the measure, with darker colors indicating higher values and lighter colors indicating lower ones.

Tree maps show up to 100 data points for the **Group by** field. For more information about how Amazon Quick Suite handles data that falls outside display limits, see [Display limits](#).

Tree map features

To understand the features supported by tree maps, use the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	Yes		Legends on visual types in Quick Suite
Changing the title display	Yes		Titles and subtitles on visual types in Quick Suite
Changing the axis range	Not applicable		Range and scale on visual types in Quick Suite

Feature	Supported?	Comments	For more information
Changing the visual colors	No		Colors in visual types in Quick Suite
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude a rectangle from a tree map, except when you are using a date field as the dimension. In that case, you can only focus on a rectangle, not exclude it.	Focusing on visual elements Excluding visual elements
Sorting	No	Default sorting is in descending order by the measure in the Size column.	Sorting visual data in Amazon Quick Suite
Performing field aggregation	Yes	You must apply aggregation to the fields you choose for size and color, and can't apply aggregation to the field that you choose to group by.	Changing field aggregation
Adding drill-downs	Yes	You can add drill-down levels to the Group by field well.	Adding drill-downs to visual data in Quick Sight

Creating a tree map

Use the following procedure to create a tree map.

To create a tree map

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the tree map icon.

4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is automatically applied to it to create a numeric value.

To create a tree map, drag a measure to the **Size** field well and a dimension to the **Group by** field well. Optionally, drag another measure to the **Color** field well.

5. (Optional) Add drill-down layers by dragging one or more additional fields to the **Group by** field well. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

Using waterfall charts

Use a waterfall chart to visualize a sequential summation as values are added or subtracted. In a waterfall chart, the initial value goes through a (positive or negative) change, with each change represented as a bar. The final total is represented by the last bar. Waterfall charts are also known as *bridges* because the connectors between the bars bridge the bars together, showing that they visually belong to the same story.

Waterfall charts are most commonly used to present financial data, because you can show change within one time period or from one time period to another. This way, you can visualize the different factors that have an impact your project cost. For example, you can use a waterfall chart to show gross sales to net income within the same month, or the difference in net income from last year to this year, and the factors that were responsible for this change.

You can also use waterfall charts to present statistical data, for example how many new employees you hired and how many employees left your company within a year.

The following screenshot shows a waterfall chart.

To create a basic waterfall chart visual

1. Open Amazon Quick Suite and choose **Analyses** on the navigation pane at left.
2. Choose one of the following:
 - To create a new analysis, choose **New analysis** at upper right. For more information, see [Starting an analysis in Quick Sight](#).
 - To use an existing analysis, choose the analysis that you want to edit.

3. Choose **Add (+), Add Visual**.
4. At lower left, choose the waterfall chart icon from **Visual types**.
5. On the **Fields list** pane, choose the fields that you want to use for the appropriate field wells. Waterfall charts require one category or measure in **Value**.
6. (Optional) Add drill-down layers by dragging one or more additional fields to the **Group/Color** field well. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

To understand the features supported by waterfall charts, see [Analytics formatting per type in Quick Suite](#). For customization options, see [Formatting in Amazon Quick Suite](#).

Using word clouds

As an engaging way to display how often a word is used in relation to other words in a dataset, use word clouds. The best use for this type of visual is to show word or phrase frequency. It can also make a fun addition to show trending items or actions. You can use a fixed dataset for creative purposes. For example, you might make one of team goals, motivational phrases, various translations of a specific word, or anything else that you want to draw attention to.

Each word in a word cloud represents one or more values in a dimension. The size of the word represents the frequency of a value's occurrence in a selected dimension, in proportion to the occurrences of other values in the same dimension. Word clouds are best when precision isn't important and there aren't a large number of distinct values.

The following screenshot shows an example of a word cloud.

To create a word cloud, use one dimension in the **Group by** field well. Optionally, you can add a metric to the **Size** field well.

Word clouds usually look better with 20–100 words or phrases, but the format settings offer a wide range of flexibility. If you choose too many words, they can become too small to be legible, depending on the size of your display. By default, word clouds display 100 distinct words. To show more, change the format setting for **Number of words**.

Word clouds are limited to 500 unique values for **Group by**. To avoid displaying the word **Other**, format the visual to hide the **Other** category. For more information about how Amazon Quick Suite handles data that falls outside display limits, see [Display limits](#).

Word cloud features

To understand the features supported by word clouds, see the following table.

Feature	Supported?	Comments	For more information
Changing the legend display	No		Legends on visual types in Quick Suite
Changing the title display	Yes		Titles and subtitles on visual types in Quick Suite
Changing the axis range	Not applicable		Range and scale on visual types in Quick Suite
Changing the visual colors	Yes	To change the color, choose a word and then choose a color.	Colors in visual types in Quick Suite
Focusing on or excluding elements	Yes		Focusing on visual elements Excluding visual elements
Sorting	Yes		Sorting visual data in Amazon Quick Suite
Performing field aggregation	Yes	You can't apply aggregation to the field that you choose for Group by . You must apply an aggregation to the field that you choose for Size .	Changing field aggregation
Adding drill-downs	Yes	You can add drill-down levels to the Group by field well.	Adding drill-downs to visual data in Quick Sight
Using format options	Yes	You can choose to allow vertical words, emphasize	Formatting in Amazon Quick Suite

Feature	Supported?	Comments	For more information
		scale, use a fluid layout, use lowercase, and set the amount of padding between words. You can set the maximum string length for the word cloud (default is 40). You can also choose the number of words for the Group by field (default is 100; maximum is 500).	
Showing totals	No		Formatting in Amazon Quick Suite

Creating a word cloud

Use the following procedure to create a word cloud.

To create a word cloud

1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the word cloud icon.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. Typically, you want to use dimension or measure fields as indicated by the target field well. If you choose to use a dimension field as a measure, the **Count** aggregate function is applied by default.

To create a word cloud, add a dimension to the **Group by** field well. Optionally, add a measure to the **Size** field well.

5. (Optional) Add drill-down layers by dragging one or more additional fields to the **Group by** field well. For more information about adding drill-downs, see [Adding drill-downs to visual data in Quick Sight](#).

Formatting in Amazon Quick Suite

You choose from a variety of options to format and style your data visualizations. To format a visual, select the visual that you want to format and choose the **Format visual** icon on the upper-right corner of the visual. Once you have the format visual pane open, you can click on different visuals and controls to view formatting data for the specific visual or control. For more information about formatting a visual control, see [Using a control with a parameter in Amazon Quick Suite](#).

Use the following sections to format and style your content:

Note

Any format changes applied from the field wells are applied only to the selected visual.

Topics

- [Analytics formatting per type in Quick Suite](#)
- [Table and pivot table formatting options in Quick Suite](#)
- [Adding data bars to tables in Quick Suite](#)
- [Map and geospatial chart formatting options in Quick Suite](#)
- [Axes and grid lines on visual types in Quick Suite](#)
- [Colors in visual types in Quick Suite](#)
- [Working with field level coloring in Amazon Quick Suite](#)
- [Conditional formatting on visual types in Quick Suite](#)
- [KPI options](#)
- [Labels on visual types in Quick Suite](#)
- [Formatting visual numeric data based on language settings in Quick Suite](#)
- [Legends on visual types in Quick Suite](#)
- [Line and marker styling on line charts in Quick Suite](#)
- [Missing data on visual types in Quick Suite](#)
- [Reference lines on visuals types in Quick Suite](#)
- [Formatting radar charts in Quick Suite](#)
- [Range and scale on visual types in Quick Suite](#)
- [Small multiples axis options](#)

- [Titles and subtitles on visual types in Quick Suite](#)
- [Tooltips on visual types in Quick Suite](#)

Analytics formatting per type in Quick Suite

Use the following list to see what type of formatting works in a visualization during analysis:

- Bar charts (both horizontal and vertical) support the following formatting:
 - Customize, display, or hide title, field labels, and data labels
 - Customize, display, or hide legend (exception: simple charts without clustering or multiple measures don't show a legend)
 - Specify axis range and steps on the x-axis for horizontal bar charts, and on the y-axis for vertical bar charts
 - Choose how many data points to display on the x-axis for vertical bar charts, and on the y-axis for horizontal bar charts
 - Show or hide axis lines, axis labels, axis sort icons, and chart grid lines
 - Customize, display, or remove reference lines
 - Show or hide the "other" category

Horizontal bar charts support sorting on the y-axis and **Value**. Vertical bar charts support sorting on the x-axis and **Value**.

Stacked bar charts support showing totals.

- Box plots support the following formatting:
 - Customize, display, or hide title
 - Customize, display, or hide legend
 - Specify axis range and label tick on the x-axis and axis range and step on the y-axis
 - Show or hide axis lines, axis labels, axis sort icons, and chart grid lines
 - Choose how many data points to display on the y-axis.
 - Show or hide the "other" category
 - Add reference lines

Box plots support sorting on **Group by**.

- **Combo charts** support the following formatting:

- Customize, display, or hide title, field labels, and data labels
- Customize, display, or hide legend (exception: simple charts without clustering, stacking, or multiple measures don't show a legend)
- Specify axis range on bars and lines
- Synchronize the Y axes for both bars and lines into a single axis.
- Choose how many data points to display on the x-axis
- Show or hide axis lines, axis labels, axis sort icons, and chart grid lines
- Customize, display, or remove reference lines
- Show or hide the "other" category

Combo charts support sorting on the x-axis, **Bars**, and **Lines**.

- Donut charts support the following formatting:
 - Customize, display, or hide title, data labels, and legend
 - Customize, display, or hide the labels for group or color and value fields
 - Choose how many slices to display from **Group/Color**
 - Show or hide the "other" category

Donut charts support sorting on **Group/Color** and **Value**.

- Filled maps support the following formatting:
 - Customize, display, or hide title.
 - Customize, display, or hide the legend

Filled maps support sorting on **Location** and **Color**.

- Funnel charts support the following formatting:
 - Customize, display, or hide title, and data labels
 - Customize, display, or hide the labels for group or color and value fields
 - Choose how many stages to display in the **Group by** field
 - Show or hide the "other" category

Funnel charts support sorting on **Group by** and **Value**.

- Gauge charts support the following formatting:
 - Customize, display, or hide title. Display or hide axis labels.
 - Customize how to display the value or values: hidden, actual value, comparison

- Choose the comparison method (available when you use two measures)
- Choose the axis range and padding to display in the gauge chart
- Choose the arc style (degrees from 180 to 360) and arc thickness

Gauge charts don't support sorting.

- Geospatial charts (maps) support the following formatting:
 - Customize, display, or hide title and legend
 - Choose the base map image.
 - Choose to display map points with or without clustering.

Geospatial charts don't support sorting.

- Heat maps support the following formatting:
 - Customize, display, or hide title, legend, and labels
 - Choose how many rows and columns to display
 - Choose colors or gradients.
 - Show or hide the "other" category

Heat maps support sorting on **Values** and **Columns**.

- Histogram charts support the following formatting:
 - Customize, display, or hide title, field labels, and data labels
 - Specify axis range, scale, and steps on the y-axis
 - Choose how many data points to display on the x-axis
 - Show or hide axis lines, axis labels, axis sort icons, and chart grid lines

Histogram charts don't support sorting.

- Key performance indicators (KPIs) support the following formatting:
 - Customize, display, or hide title
 - Display or hide trend arrows and progress bar
 - Customize comparison method as auto, difference, percent (%), or difference as percent (%)
 - Customize primary value displayed to be comparison or actual
 - Conditional formatting

- Line charts support the following formatting:
 - Customize, display, or hide title, field labels, and data labels
 - Customize, display, or hide legend (exception: simple charts don't show a legend)
 - Specify axis range and steps (on y-axis)
 - Choose how many data points to display on the x-axis
 - Show or hide axis lines, axis labels, axis sort icons, and chart grid lines
 - Customize, display, or remove reference lines
 - Customize the styling of lines and the markers for data points on a line
 - Show or hide the "other" category, except when the x-axis is a date

Line charts support sorting on the x-axis and **Value** for numeric purposes only.

- Pie charts support the following formatting:
 - Customize, display, or hide title, data labels, and legend
 - Customize, display, or hide the labels for group or color and value fields
 - Show metrics as values, percentages, or both
 - Choose how many slices to display from the **Group/Color** field
 - Show or hide the "other" category

Pie charts support sorting on **Value** and **Group/Color**.

- Pivot tables support the following formatting:
 - Customize, display, or hide title
 - Customize, display, or hide the labels for column, row, and value fields
 - Customize the font sizes for table headers and cells/body
 - Display or hide totals and subtotals on rows or columns
 - Custom labels for totals or subtotals
 - Choose additional styling options: fit table to view, hide +/- buttons, hide column field names, hide duplicate label when using single metric
 - Conditional formatting

Pivot tables support sorting on **Column** and **Row**. For more information on sorting pivot table data, see [Sorting pivot tables in Quick Suite](#).

- Scatter plots support the following formatting:
 - **Customize, display, or hide title, legend, field labels, and data labels**

- Customize, display, or remove reference lines
- Specify axis range (on x-axis and y-axis)
- Show or hide axis lines, axis labels, axis sort icons, and chart grid lines

Scatter plots don't support sorting.

- Tables support the following formatting:
 - Customize, display, or hide title, legend, and columns
 - Customize, display, or hide the column names for group-by and value fields
 - Customize the font sizes for table headers and cells/body
 - Display or hide totals at the top or bottom of the table
 - Provide a custom label for totals
 - Add conditional formatting

Tables support sorting on **Group by** and **Value**.

- Tree maps support the following formatting:
 - Customize, display, or hide title and legend
 - Customize, display, or hide the labels for group-by, size, and color fields
 - Choose colors or gradients.
 - Choose how many squares to display from the **Group by** field
 - Show or hide the "other" category

Line charts support sorting on **Size**, **Group by** and **Color**.

- Waterfall charts support the following formatting:
 - Customize, display, or hide title or subtitle
 - Customize the total label
 - Specify x-axis label size and orientation and y-axis label range and orientation.
 - Show or hide axis lines, axis labels, axis sort icons, and chart grid lines
 - Show or hide the "other" category
 - Customize the legend size and position.
 - Customize and display or hide data labels.

Waterfall charts support sorting on **Category** and **Value**.

- Word clouds support the following formatting:

- Customize, display, or hide title
- Customize the word color, and the number of words to display from the **Group by** field
- Show or hide the "other" category
- Choose additional styling options: allow vertical words, emphasize scale, or work with fluid layout, lowercase, padding level, or maximum string length

Word clouds support sorting on **Group by**.

Table and pivot table formatting options in Quick Suite

You can customize tables and pivot tables in Quick Suite to meet your business needs. You can customize table headers, cells, and totals by specifying the color, size, wrap, and alignment of text in each. You can also specify the height of rows in a table, add borders and grid lines, and add custom background colors. In addition, you can customize how to display totals and subtotals.

If you have applied conditional formatting to a table or pivot table, it takes precedence over any other styling you configure.

When you export table or pivot table visuals to Microsoft Excel, the formatting customizations that you applied to the visual aren't reflected in the downloaded Excel file.

To format a table or pivot table

- In your analysis, choose the table or pivot table that you want to customize, and then choose the **Format visual** icon.

The **Properties** pane opens at left.

Following, you can find descriptions for options for customizing each area of your table or pivot table in the **Properties** pane.

Topics

- [Headers](#)
- [Cell formatting](#)
- [Totals and subtotals](#)
- [Row and column size in tables and pivot tables in Quick Suite](#)

- [Customize pivot table data](#)

Headers

Expand all headers

You can choose to expand all headers in a pivot table to show all child and grandchild rows of a header.

To expand all headers of a pivot table

1. On the visual that you want to change, select any header to open the **On-visual** menu.
2. Choose **Expand all below**.

Header height

You can customize table header height.

To customize the height of headers in a table

1. In the **Properties** pane, choose **Headers**.
2. For **Row height**, enter a number in pixels. You can enter a whole number from 8 through 500.

To customize the height of headers in a pivot table

1. In the **Properties** pane, choose **Headers**.
2. In the **Columns** section, for **Row height**, enter a number in pixels. You can enter a whole number from 8 through 500.

Header text

You can customize table header text.

To customize header text in a table

1. In the **Properties** pane, choose **Headers**.
2. Navigate to the **TEXT** section and do one or more of the following:

- To change the color of the header text, choose the color swatch underneath **Text styling**, and then choose the color that you want the table text to be.
- To change the font or font size of the header text, open the **Font** or **Font size** dropdown and choose the font or font size that you want.
- To bold, italicize, or underline the header text, choose the appropriate icon from the style bar.
- To wrap text in headers that are too long to fit, select **Wrap text**. Wrapping text in a header doesn't automatically increase the height of the header. Follow the previous procedure for increasing header height.
- To change the horizontal alignment of text in the header, choose a horizontal alignment icon. You can choose left alignment, center alignment, right alignment, or automatic alignment.
- To change the vertical alignment of text in the header, choose a vertical alignment icon. You can choose top alignment, middle alignment, or bottom alignment.

To customize header text in a pivot table

1. In the **Properties** pane, choose **Headers**.

The Headers section expands to show options for customizing column and row headers.

2. In the **Headers** section, do one or more of the following:
 - To apply row styling to field names of the rows or columns, choose **Style rows label** or **Style columns label** depending on the label that you want to customize.
 - To customize the header font, navigate to the **TEXT** subsection of the **Rows** or **Columns** section and do one or more of the following:
 - To change the color of the header text, choose the color swatch underneath **Text styling**, and then choose the color that you want the table text to be.
 - To change the font or font size of the header text, open the **Font** or **Font size** dropdown and choose the font or font size that you want.
 - To bold, italicize, or underline the header text, choose the appropriate icon from the style bar.
 - To change the horizontal alignment of text in the header, choose an alignment icon. You can choose left alignment, center alignment, right alignment, or automatic alignment. You

can choose a horizontal alignment for column headers in the **Columns** section, and for row headers in the **Rows** section.

- To change the vertical alignment of text in the header, choose an alignment icon. You can choose top alignment, middle alignment, or bottom alignment. You can choose a vertical alignment for column headers in the **Columns** section, and row headers in the **Rows** section.
- To hide the rows label or column field names, choose the eye icon next to **Rows label** or **Column field names**.

Header background color

You can customize table headers' background color.

To customize the background color of table headers

1. In the **Properties** pane, choose **Headers**.
2. For **Background**, choose the background color icon, and then choose a color. You can choose one of the provided colors, reset the header text color to the default color, or create a custom color.

To customize the background color of pivot table headers

1. In the **Properties** pane, choose **Headers**.

The **Headers** section expands to show options for customizing column and row headers.

2. In the **Columns** section, choose the background color icon, and then choose a color.
3. In the **Rows** section, choose the background color icon, and then choose a color.

Header borders

You can customize header borders' color.

To customize header borders in a table

1. In the **Properties** pane, choose **Headers**.
2. For **Borders**, do one or more of the following:
 - To customize the type of border that you want, choose a border type icon. You can choose no borders, horizontal borders only, vertical borders only, or all borders.

- To customize the border thickness, choose a border thickness.
- To customize the border color, choose the border color icon, and then choose a color. You can choose one of the provided colors, reset the border color to the default color, or create a custom color.

To customize header borders in a pivot table

1. In the **Properties** pane, choose **Headers**.

The **Headers** section expands to show options for customizing column and row headers.

2. In the **Columns** and **Rows** sections, for **Borders**, do one or more of the following:
 - To customize the type of border that you want, choose a border type icon. You can choose no borders, horizontal borders only, vertical borders only, or all borders.
 - To customize the border thickness, choose a border thickness.
 - To customize the border color, choose the border color icon, and then choose a color. You can choose one of the provided colors, reset the border color to the default color, or create a custom color.

Header styling options for hierarchy pivot tables

You can hide or rename the **Rows** label of a hierarchy pivot table.

To make changes to the Rows label of a hierarchy pivot table

1. Select the hierarchy pivot table that you want to change and open the **Format visual** menu.
2. In the **Headers** section, you can perform the following tasks
 - Choose **Hide rows label** to hide the **Rows** label from your pivot table.
 - For **Rows label**, enter the label that you want displayed on the pivot table.

Cell formatting

Row height

You can customize table row height.

To customize the height of rows in a table or pivot table

1. In the **Properties** pane, choose **Cells**.

The **Cells** section expands to show options for customizing cells.

2. For **Row height**, enter a number in pixels. You can enter a whole number from 8 through 500.

Cell text

You can customize the formatting for cell text within a table.

To format the cell text in a table or pivot table

1. In the **Properties** pane, choose **Cells**.

The **Cells** section expands to show options for customizing cells.

2. For **Text**, do one or more of the following:
 - To change the color of the cell text, choose the color swatch underneath **Text styling**, and then choose the color that you want the table text to be.
 - To change the font or font size of the cell text, open the **Font** or **Font size** dropdown and choose the font or font size that you want.
 - To bold, italicize, or underline the cell text, choose the appropriate icon from the style bar.
 - To wrap text in headers that are too long to fit, select **Wrap text**. Wrapping text in cells doesn't automatically increase the row height. Follow the previous procedure for increasing row height.
 - To change the horizontal alignment of text in cells, choose a horizontal alignment icon. You can choose left alignment, center alignment, right alignment, or automatic alignment. Horizontal alignment can only be configured for the **Rows** fields of a hierarchy pivot table.
 - To change the vertical alignment of text in cells, choose a vertical alignment icon. You can choose top alignment, middle alignment, bottom alignment, or automatic. For tabular pivot tables, the value for **Automatic** is vertical. For hierarchy pivot tables, the value for **Automatic** is middle.

Cell background color

You can customize table cells' background color.

To customize the background color of cells in a table or pivot table

1. In the **Properties** pane, choose **Cells**.

The **Cells** section expands to show options for customizing cells.

2. For **Background**, do one or more of the following:
 - To alternate background colors between rows, select **Alternate row colors**. Clearing this option means that all cells have the same background color.
 - If you choose to alternate background colors between rows, choose a color for **Odd rows** and a color for **Even rows** by choosing the background color icon for each and selecting a color. You can choose one of the provided colors, reset the background color to the default color, or create a custom color.
 - If you choose not to alternate background colors between rows, choose the background color icon and select a color for all cells. You can choose one of the provided colors, reset the background color to the default color, or create a custom color.

Cell borders

You can customize table cells' borders.

To customize the borders for cells in a table or pivot table

1. In the **Properties** pane, choose **Cells**.

The **Cells** section expands to show options for customizing cells.

2. For **Borders**, do one or more of the following:
 - To customize the type of border that you want, choose a border type icon. You can choose no borders, horizontal borders only, vertical borders only, or all borders.
 - To customize the border thickness, choose a border thickness.
 - To customize the border color, choose the border color icon, and then choose a color. You can choose one of the provided colors, reset the border color to the default color, or create a custom color.

Totals and subtotals

On tables and pivot tables, you can configure the display of totals or subtotals. Tables can display totals at the top or the bottom of the visual. Pivot tables can display totals and subtotals on rows and columns.

Add totals and subtotals to tables and pivot tables in Quick Suite

You can add total columns to your table and pivot table visuals. You can also add subtotal columns to your pivot table visuals.

To display or hide totals and subtotals for a pivot table

1. To display totals, open the **Properties** pane and choose **Total**.
 - To show totals for rows, toggle the **ROWS** switch on. Totals are displayed on the bottom row of the visual. Choose **Pin totals** to keep the totals visible as you scroll through the table.
 - To show totals for columns, toggle the **COLUMNS** switch on. Totals are displayed on the last column of the visual.
2. To display totals, open the **Properties** pane and choose **Subtotal**.
 - To show subtotals for rows, toggle the **ROWS** switch on. Totals are displayed on the bottom row of the visual.
 - To show subtotals for columns, toggle the **COLUMNS** switch on.
 - For **Level**, choose one of the following:
 - Choose **Last** to only show the subtotal of the last field in the chart's hierarchy. This is the default option.
 - Choose **All** to show subtotals for every field.
 - Choose **Custom** to customize which fields show subtotals.

After you add row totals to your table or pivot table visual, you can also choose to position the totals at the top or bottom of the visual. You can also change the position of column totals in pivot tables.

To position row or column totals in a table or pivot table

1. In the **Properties** pane, choose **Total**.
2. (Optional) For **Rows**, choose **Show totals**.

3. (Optional) For **Columns**, choose **Show totals**.
4. (Optional) In the **Rows** menu, open the **Position** dropdown and choose the position that you want the totals to be displayed. Choose **Top** to position totals at the top of the table, or **Bottom** to position totals at the bottom of the table.
5. (Optional) In the **Columns** menu, open the **Position** dropdown and choose the position that you want the totals to be displayed. Choose **Left** to position totals at the left of the table, or **Right** to position totals at the right of the table.

You can't change the position of the subtotals of a pivot table visual. If your pivot table uses a hierarchy layout, the subtotal rows are positioned at the top of the table. Tabular pivot table subtotals are displayed at the bottom of the table.

Customize labels for totals and subtotals

You can rename the totals in table and pivot table visuals to provide better context for account readers. By default, the totals and subtotals appear without a label.

To rename totals in a table or pivot table visual

1. In the **Properties** pane, choose **Total** or **Subtotal**.
2. For **Label**, enter a word or short phrase that you want displayed for the total.

In pivot tables, you can also add labels to column totals and subtotals. To do so, enter a word or short phrase for **Label** in the **Columns** section.

3. (Optional) For tabular pivot tables, you can also add group names to subtotals. To add a group name to row subtotals, choose the **Plus (+)** icon next to the **Label** field to add the group name parameter that you want. You can also enter a word or short phrase to this field.

You can also make changes to the text size and font color of the total and subtotal labels of your table and pivot table visuals.

To format totals and subtotals text

1. In the **Properties** pane, choose **Total** or **Subtotal**.
2. For **Text**, do one or more of the following.
 - To change the color of the total or subtotal text, choose the color swatch underneath **Text styling**, and then choose the color that you want the table text to be.

- To change the font or font size of the total or subtotal text, open the **Font** or **Font size** dropdown and choose the font or font size that you want.
- To bold, italicize, or underline the total or subtotal text, choose the appropriate icon from the style bar.

In pivot tables, you can also add format text for column totals and subtotals. To do so, repeat the above steps in the **Columns** section.

Totals and subtotals background color

To customize the background color for totals and subtotals

1. In the **Properties** pane, choose **Total** or **Subtotal**.
2. For **Background**, choose the background color icon, and then choose a color. You can choose one of the provided colors, reset the background color to the default color, or create a custom color.

In pivot tables, you can also add background colors for column totals and subtotals. To do so, choose a the background color icon for **Background** in the **Columns** section.

Totals and subtotals borders

To customize the borders for totals and subtotals

1. In the **Properties** pane, choose **Total** or **Subtotal**.
2. For **Borders**, do one or more of the following:
 - To customize the type of border that you want, choose a border type icon. You can choose no borders, horizontal borders only, vertical borders only, or all borders.
 - To customize the border thickness, choose a border thickness.
 - To customize the border color, choose the border color icon, and then choose a color. You can choose one of the provided colors, reset the border color to the default color, or create a custom color.

In pivot tables, you can also add borders for column totals and subtotals. To do so, repeat the above steps in the **Columns** section.

Applying totals and subtotals styling to cells

In pivot tables, you can apply any text, background color, and border styling you apply to totals to cells in that same column or row. Row subtotals appear differently depending on the layout that your pivot table uses. For tabular pivot tables, explicit subtotal headers appear on the visual. For hierarchy pivot tables, explicit subtotal headers do not appear. Instead, authors apply subtotal styling to individual fields from the **Format visual** menu. Collapsed headers cannot be styled as subtotals.

To apply totals and subtotals styling to cells

1. In the **Properties** pane, choose **Total** or **Subtotal**.
2. For **Apply styling to**, choose the visual that you want to apply subtotal styling to. You can choose from the following options.
 - **None**– Removes styling options from all cells.
 - **Headers only**– Applies styling options to all headers in the pivot table.
 - **Cells only**– Applies styling options to all cells that aren't headers in the pivot table.
 - **Headers and cells**– Applies styling options to all cells in the pivot table.

Row and column size in tables and pivot tables in Quick Suite

Authors and readers can resize rows and columns in a table or pivot table visual. They can adjust both row height and column width. Authors can also set the default column width for columns in a pivot table visual.

To resize a row in a table or pivot table

- In the table or pivot table visual, hover your cursor over the line that you want to resize until you see the horizontal cursor appear. When it appears, select the line and drag it to a new height.

You can adjust the row height by selecting the horizontal lines on cells and row headers.

To resize a column width in a table or pivot table

- In the table or pivot table visual, hover your cursor over the line that you want to resize until you see the vertical cursor appear. When it appears, select the line and drag it to a new width.

You can adjust the column width by selecting the vertical lines on cells, column headers, and row headers.

To set the default column width for columns in a pivot table

1. Select the pivot table that you want to change and open the **Format visual** menu.
2. In the **Pivot options** section, navigate to the **Value column width (pixels)** field and enter the default value that you want in pixels.

Customize pivot table data

You can customize how Quick Suite readers view pivot tables so that they are easier to read and understand at a glance. You can choose to hide a pivot table's plus and minus icons, hide columns that only have a single-metric value, and hide collapsed columns from view. These options can help Quick Suite authors remove clutter from their pivot tables and provide an easier reader experience for Quick Suite users. This is not the same as choosing a pivot table layout. For more information on pivot table layout options, see [Choosing a layout](#).

These options can also be accessed from the **Combined row fields menu** of a pivot table. The layout that you choose for your pivot table determines how this menu is accessed. For more information on accessing the **Combined row fields** menu, see .

To make changes to a pivot table's layout

1. In the **Format visual** pane, choose **Pivot options**.
2. In the **Pivot options** menu, select the following options to customize the view:
 - **Hide +/- buttons** – Hide the plus and minus icons from your pivot table by default. Readers can still choose to show the plus and minus icons and expand or collapse columns and rows.
 - **Hide single metric** – Hide columns that only have a single metric value.
 - **Hide collapsed columns** – Automatically hide all collapsed columns in a pivot table. This option is only available for tabular pivot tables.

Adding data bars to tables in Quick Suite

You can use data bars to add visual context to your table visuals in Amazon Quick Suite. By injecting color into your tables, data bars can make it easier to visualize and compare data in a range of fields. *Data bars* are bars of different colors or shades that you add to the cells of a table. The bars are measured relative to the range of all cells in a single column, which is similar to a bar chart. You can use data bars to highlight a fluctuating trend, such as profit per quarter during the year.

You can only apply data bars to fields that are added to the **Values** field well of the visual. You can't apply data bars to items that are added to group bys.

You can create up to 200 different data bar configurations for a single table.

To add data bars to a table

1. On the analysis page, choose the visual that you want to format.
2. On the menu in the upper-right corner of the visual, select the **Format visual** icon. The **Format visual** pane opens.
3. In the **Properties** pane, open the **Visuals** dropdown list and choose **ADD DATA BARS**.
4. In the **Data bars** popup that appears, choose the value field that you want represented by the data bars. You can only choose from fields that are added to the **Values** field well of the visual.
5. (Optional) Choose the icon labeled **Positive color** to select the color that you want to represent positive value data bars. The default color is green.
6. (Optional) Choose the icon labeled **Negative color** to select the color that you want to represent negative value data bars. The default color is red.

When you create data bars, they are named for the field values that they are representing. For example, if you add data bars to represent the profit of a product over time, the data bar configuration is labeled "Profit". In the **Visuals** pane of the **Properties** menu, data bars are listed in the order that they are created.

To remove data bars from a visual

1. On the menu in the upper-right corner of the visual, select the **Format visual** icon. The **Properties** pane opens.

2. In the **Properties** pane, open the **Visuals** dropdown list and choose the data bar that you want to remove.
3. Choose **REMOVE DATA BARS**.

Map and geospatial chart formatting options in Quick Suite

In Amazon Quick Suite, you can choose from multiple formatting options for your maps and geospatial charts. You can view formatting options by opening the **Properties** pane from the on-visual menu located at the top right of the currently selected geospatial map.

Quick Suite authors and readers can also toggle the different formatting options of a geospatial map visual from the on visual menu.

Topics

- [Base maps on geospatial maps in Quick Suite](#)
- [Geospatial heatmaps in Amazon Quick Suite](#)
- [Marker clustering on geospatial point maps in Quick Suite](#)

Base maps on geospatial maps in Quick Suite

When you create a map visual in Quick Suite, you can change the base of the map. A *base map* is the style of map that appears beneath your data on a map. An example is a satellite view versus a street view.

In Quick Suite, there are four options for base maps: light gray canvas, dark gray canvas, streets, and imagery. The following list contains an example of each base map option:

Important

Only the light gray canvas is supported in the Asia Pacific (Mumbai) Amazon Region (ap-south-1).

- Light gray canvas
- Dark gray canvas

- Streets
- Imagery

Changing base maps

Use the following procedure to change a base map.

To change a base map

1. Create a point or filled map in an analysis. For more information, see [Creating maps and geospatial charts](#).
2. On the map visual, choose the **Format visual** icon.
3. In the **Properties** pane that opens, choose the **Base map** section and then choose the base map that you want.

Geospatial heatmaps in Amazon Quick Suite

Use geospatial heatmaps to reveal patterns of marker concentration in your geospatial visuals. Heat maps display concentrations of data points using a colored overlay that highlights the intensity or concentration of the visual's markers.

To turn a geospatial map into a heat map

1. Open your analysis and choose the geospatial map that you want to format. When you select a visual, it displays with a highlight around it.
2. To open the formatting pane, select the **Format visual** icon from the on-visual menu.
3. On the formatting pane at left, choose **Points**.
4. Choose **Heatmap**.
5. (Optional) For **Heatmap gradient**, choose a color that you want for the **High density** and **Low density** values.

Marker clustering on geospatial point maps in Quick Suite

Use marker clustering to improve readability of collocated points on a map. Geospatial locations on point maps are represented using markers. Usually, there is one marker per data point. However, if there are too many markers close together, the map becomes difficult to read. To make it easier to interpret the map, you can enable marker clustering to represent groupings of locations on the map. As the reader zooms in on the map, the clustered markers leave the area marker to display separately.

To add cluster points to a map

1. Open your analysis, and choose the geospatial map that you want to format. When you select a visual, it displays with a highlight around it.
2. To open the formatting pane, select the **Format visual** icon from the on-visual menu.
3. On the formatting pane at left, choose **Points**.
4. Choose one of the following options:
 - **Basic** – use the default display setting for map points.
 - **Cluster points** – cluster map points together when there are many in one area.

Axes and grid lines on visual types in Quick Suite

When you create a chart in Quick Suite, axis lines, axis labels, axis sort icons, and grid lines are added to the chart automatically. You can format your visuals to show or hide these if you want, and also customize the axis label size and orientation.

You can format axis lines, grid lines, and axis labels and axis sort icons for the following chart types:

- Bar charts
- Box plot charts
- Combo charts
- Histograms
- Line charts
- Scatter plots

- Waterfall charts

To format axis lines, axis labels, and grid lines in a chart

1. On the analysis page, choose the visual that you want to format.
2. On the menu in the upper-right corner of the visual, select the format visual icon.

The **Properties** pane opens at left.

To show or hide axis lines

1. In the **Properties** pane, choose the axis that you want to format.
2. Choose **Show axis line**. Clear the check box to hide the axis line for the chosen axis. Select the check box to show it.

To customize axis titles

1. In the **Properties** pane, choose the axis that you want to format.
2. Choose **Show title**. Clear the check box to hide the axis title and drop-down caret icon for the chosen axis. Select the check box to show them.
3. To change the title from the default field name, enter a title in the text box.

Note

In addition to the chart types listed previously in this topic, you can also customize the axis titles in pie charts, donut charts, funnel charts, heat maps, and tree maps.

To modify axis font settings

1. In the **Properties** pane, choose the axis that you want to format.
2. Adjust the following properties:
 - **Font family**
 - **Text size**
 - **Style** (Bold, Italic, Underline)

- **Color**

Note

- **Underline** is supported for axis titles, but not for axis labels
- Different chart types use different terminology:
 - **Bar/Line charts** - **X-axis** and **Y-axis**
 - **Pie charts** - **Values**
 - **Heat maps** - **Rows** and **Columns**

To show or hide the sort icon

1. In the **Properties** pane, choose the axis that you want to format.
2. Choose **Show sort**. Clear the check box to hide the sort icon for the chosen axis. Select the check box to show it.

When you choose to remove the sort icon, the sort icon is removed from the axis. Any sorts that were applied to the visual before removing the icon are not removed from the visual.

Note

In addition to the chart types listed previously in this topic, you can also show or hide the sort icon in pie charts, donut charts, funnel charts, heat maps, and tree maps.

To show or hide the data zoom

1. In the **Properties** pane, choose **X-axis**.
2. Choose **Show data zoom**. Clear the check box to hide the data zoom. Select the check box to show it.

The data zoom bar appears automatically on charts with an X-axis that contain more than one data point. Adjust the bar from the left and right to zoom to specific data points in the chart.

Note

If you zoom in or out using the data zoom bar, and then choose to hide the data zoom bar, the zoom position isn't maintained. The visual zooms completely out to include all data points. Showing the data zoom again returns the visual to its previous state.

To show or hide axis labels

1. In the **Properties** pane, choose the axis that you want to format.
2. Choose **Show labels**. Clear the check box to hide the axis labels for the chosen axis. Select the check box to show it.

To change the label size

1. In the **Properties** pane, choose the axis that you want to format.
2. For **Label size**, choose a size.

To change the label orientation

1. In the **Properties** pane, choose the axis that you want to format.
2. For **Label orientation**, choose an orientation.

To show or hide grid lines

1. In the **Properties** pane, choose the axis that you want to format.
2. Choose **Show grid lines**. Clear the check box to hide grid lines for the chosen axis. Select the check box to show it.

Colors in visual types in Quick Suite

You can change the color of one, some, or all elements on the following types of charts:

- Bar charts
- Donut charts

- Gauge charts
- Heat maps
- Line charts
- Scatter plots
- Tree maps

To change colors on bar charts, donut charts, gauge charts, line charts, and scatter plots, see [Changing colors on charts](#).

To change colors on heat maps and tree maps, see [Changing colors on heat maps and tree maps](#).

Changing colors on charts

You can change the chart color used by all elements on the chart, and also change the color of individual elements. When you set the color for an individual element, it overrides the chart color.

For example, suppose that you set the chart color to green.

All of the bars turn green. Even though you choose the first bar, the chart color applies to all the bars. Then you set the color for the **SMB** bar to blue.

Looking at the result, you decide that you need more contrast between the green and blue bars, so you change the chart color to orange. If you are changing the chart color, it doesn't matter which bar you choose to open the context menu from.

The **SMB** bar remains blue. This is because it was directly configured. The remaining bars turn orange.

When you change the color of an element that is grouped, the color for that element is changed in all of the groups. An example is a bar in a clustered bar chart. In the following example, Customer Segment is moved out of the **Y-axis** and into the **Group/Color** field well. Customer Region is added as the **Y-axis**. The chart color stays orange, and SMB stays blue for all Customer Regions.

If your visual has a legend that shows categories (dimensions), you can click on the values in the legend to see a menu of available actions. For example, suppose that your bar chart has a field in

the **Color** or **Group/Color** field well. The bar chart menu displays the actions that you can choose by clicking or right-clicking on a bar, such as the following:

- Focusing on, or excluding, visual elements
- Changing colors of visual elements
- Drilling down into a hierarchy
- Custom actions activated from the menu, including filtering or URL actions

Following is an example of using the legend to change the color for a dimension.

Setting new colors for a visual

Use the following procedure to change the colors for a visual.

To change the colors for a visual

1. On the analysis page, choose the visual that you want to modify.
2. To change the chart color, choose any element on the visual, and then choose **Chart Color**.

To select elements, do the following:

- On a bar chart, choose any bar.
 - On a line chart, choose the end of a line.
 - On a scatter plot, choose an element. The field must be in the **Group/Color** section of **Field wells**.
3. Choose the color that you want to use. You can choose a color from the existing palette, or you can choose a custom color. To use a custom color, enter the hexadecimal code for that color.

All elements on the visual are changed to use this color, except for any that have previously had their color individually set. In that case, the element color overrides the chart color.

4. To change the color for a single element on the visual, choose that element, choose **Color <field name>**, and then choose the color that you want to use. You can choose a color from the existing palette, or you can choose a custom color. To use a custom color, enter the hexadecimal code for that color.

Repeat this step until you have set the color on all elements that you want to modify. To change the color back to the color it was originally, choose **Reset to default**.

Setting visual colors back to defaults

Use the following procedure to return to using the default colors on a visual.

To return to default colors on a visual

1. On the analysis page, choose the visual that you want to modify.
2. Choose **Chart Color**, choose any element on the visual, and then choose **Reset to Default**. Doing this changes the chart color back to the default color for that visual type.

All elements on the visual are changed to the default color for the visual type, except for any that have previously had their color individually set. In that case, the element color setting overrides the chart color setting.

3. To change the color for a single element back to the default, choose that element, choose **Color <field name>**, and then choose **Reset to Default**.

The default color for individual elements is the chart color if you have specified one, or the default color for the visual type otherwise.

Changing colors on heat maps and tree maps

To change the colors that display on a heat map or a tree map

1. Choose the heat map or tree map that you want to edit.
2. Choose **Expand** for the settings menu, and choose the cog icon to open the **Properties** panel.
3. For **Color**, choose the settings that you want to use:
4. For **Gradient color** or **Discrete color**, choose the color square next to the color bar, and then choose the color that you want to use. Repeat for each color square. The bar holds two colors by default.
5. Select the **Enable 3 colors** check box if you want to add a third color. A new square appears in the middle of the color bar.

You can enter a number that defines the midpoint between the two main gradient colors. If you add a value, the middle color represents the number you entered. If you leave this blank, the middle color acts like the other colors in the gradient.

6. Select the **Enable steps** check box if you want to limit the chart to the colors that you chose. Doing this changes the label on the color bar from **Gradient color** to **Discrete color**.

7. For **Color for Null Value**, choose a color to depict NULL values. This option is only available on heat maps.

Working with field level coloring in Amazon Quick Suite

With field level coloring, you can assign specific colors to specific field values across all visuals in a Quick Suite analysis or dashboard. Colors are assigned on a per-field basis to simplify the process of setting colors and ensure consistency across all visuals that use the same field. For example, let's say you're a shipping company that wants to create a set of visuals that track shipping rates in different regions. With field level coloring, you can assign each region a different color to represent the field across all visuals in an analysis or dashboard. This way, account readers quickly learn what field colors they're looking for and have an easier time finding the information that they need.

Quick Suite authors can configure up to 50 field based colors per field. Colors that are defined at the visual level take precedence over field based colors. This means that if the author sets a color for a value on the visual, that color will override the field based colors configuration for that individual visual.

To apply field level coloring to a legacy account

1. In the **Fields** pane of the analysis, choose the ellipsis (three dots) next to the field that you want to assign a color to, and then choose **Edit field colors**.
2. In the **Edit field colors** pane that appears, choose the value that you want to assign a color to and choose the color that you want. You can apply colors to every value that appears in the **Field values** pane.
3. When you are finished assigning colors to the fields that you want, choose **Apply**.

If you want to reset the color value of a field, open the **Edit field colors** pane and choose the refresh icon next to the field that you want to reset. You can reset all color values in an analysis by choosing **RESET COLORS**.

You can view a list of unused colors that can be configured to new fields by choosing **Show unused colors** in the **Edit field colors** pane. When you reset a field's color, the discarded color is added to the **Unused colors** list and can be assigned to a new field.

Conditional formatting on visual types in Quick Suite

In some visual types, you can add conditional formatting to highlight some of your data. The conditional formatting options currently supported include changing text or background color and using symbolic icons. You can use icons from the provided set, or you can use Unicode icons instead.

Conditional formatting is available on the following visuals:

- Gauge charts
- Key performance indicators (KPIs)
- Pivot tables
- Tables

For tables and pivot tables, you can set multiple conditions for fields or supported aggregations, along with format options to apply to a target cell. For KPIs and gauge charts, you can format the primary value based on conditions that are applied to any dimension in the dataset. For gauge charts, you can also format the foreground color of the arc based on conditions.

To use conditional formatting on a visual

1. On the analysis page, choose the visual that you want to format.
2. On the visual, open the context menu on the down icon at the upper-right. Then choose **Conditional formatting**.

Options for formatting display on the left. Choose one of the following:

- **For pivot tables** – Begin by choosing a measure that you want to use. You can set conditional formatting on one or more fields. The selection is limited to the measures that are in the **Values** field well.
 - **For tables** – Begin by choosing a field that you want to use. You can set conditional formatting on one or more fields. You can also choose to apply formatting to the entire row. Formatting the entire row adds an option to **Apply on top**, which applies the row formatting in addition to formatting added by other conditions.
 - **For KPIs** – Apply formatting to the primary value or the progress bar or both.
3. For the remaining steps in this procedure, choose the features that you want to use. Not all options are available for all visuals.

4. (Optional) Choose **Add background color** to set a background color. If a background color is already added, choose **Background**.
 - **Fill type** – The background color can be **Solid** or **Gradient**. If you choose to use a gradient, additional color options display, enabling you to choose a minimum and maximum value for the gradient scale. The minimum value defaults to the lowest value, and the maximum value defaults to the highest value.
 - **Format field based on** – The field to use when applying the format.
 - **Aggregation** – The aggregation to use (displays only the available aggregations).
 - **Condition** – The comparison operator to use, for example "greater than".
 - **Value** – The value to use.
 - **Color** – The color to use.
 - **Additional options:** In pivot tables, you can set what you want to format by choosing options from the context menu (...): **Values**, **Subtotals**, and **Totals**.
5. (Optional) Choose **Add text color** to set a text color. If a text color is already added, choose **Text**.
 - **Format field based on** – The field or item to use when applying the format.
 - **Aggregation** – The aggregation to use (displays only the available aggregations). This option applies to tables and pivot tables.
 - **Condition** – The comparison operator to use, for example "greater than".
 - **Value** – The value to use.
 - **Color** – The color to use.
 - **Additional options:** In tables and pivot tables, you can set what you want to format by choosing options from the context menu (...): **Values**, **Subtotals**, and **Totals**.
6. (Optional) Choose **Add icons** to set an icon or icon set. If an icon is already added, choose **Icon**.
 - **Format field based on** – The field or item to use when applying the format.
 - **Aggregation** – The aggregation to use (displays only the available aggregations). This option applies to tables and pivot tables.
 - **Icon set** – The icon set to apply to field in **Format field based on**. This option applies to tables and pivot tables.
 - **Reverse colors** – Reverses the colors of the icons for tables and pivot tables.
 - **Custom conditions** – Provides more icon options for tables and pivot tables.

- **Condition** – The comparison operator to use.
 - **Value** – The value to use.
 - **Icon** – The icon to use. To choose an icon set, use the **Icon** symbol to choose the icons to use. Choose from the provided icon sets. In some cases, you can add your own. To use your own icon, choose **Use custom Unicode icon**. Paste in the Unicode glyph that you want to use as an icon. Choose **Apply** to save or choose **Cancel** to exit icon setup.
 - **Color** – The color to use.
 - **Show icon only** – Replaces the value with the icon for tables and pivot tables.
 - **Additional options:**
 - In tables and pivot tables, you can set what you want to format by choosing options from the context menu (...): **Values**, **Subtotals**, and **Totals**.
 - In pivot tables, enabling **Custom conditions** activates preset conditional formatting that you can keep, add to, or overwrite with your own settings.
7. (Optional) Choose **Add foreground color** to set the foreground color of a KPI progress bar. If a foreground color is already added, choose **Foreground**.
- **Format field based on** – The field to use when applying the format.
 - **Condition** – The comparison operator to use.
 - **Value** – The value to use.
 - **Color** – The color to use.
8. When you are finished configuring conditional formatting, choose one or more of the following:
- To save your work, choose **Apply**.
 - To cancel selections and return to the previous panel, choose **Cancel**.
 - To close the settings panel, choose **Close**.
 - To reset all settings on this panel, choose **Clear**.

KPI options

You can customize KPIs in Amazon Quick Suite to meet your business needs. You can add contextual sparklines or progress bars, assign primary and secondary values, and add conditional formatting to your KPIs.

To format a KPI in Quick Suite, navigate to the KPI that you want to change and choose the **Format visual** icon to open the **Format visual**.

Use the following procedures to perform formatting tasks for KPIs.

Add a visual to a KPI

You can choose to add an area sparkline, a sparkline, or a progress bar to any KPI in Quick Suite. Adding visuals to KPIs provides visual context to readers who are viewing KPI data. Use the following procedure to add a visual to a KPI.

To add a visual to a KPI

1. Navigate to the KPI that you want to change and open the format visual menu.
2. In the **Properties** menu, choose the **Visual** box to display a visual on your KPI chart.
3. (Optional) Open the **Visual** dropdown and choose the type of visual that you want to display on your KPI. You can choose to display an area sparkline, a sparkline, or a progress bar. To display a sparkline, make sure that your KPI has a value in the **Trend** field well. **Area sparkline** is the default value.
4. (Optional) To change the color of the sparkline, choose the color icon to the left of the **Visual** dropdown and choose the color that you want. Color formatting isn't supported for the progress bar.
5. (Optional) Choose **Add tooltip** to add a tooltip to the KPI visual.

Customizing primary and secondary values

Use the **Format visual** menu to customize the font, color, and to choose which primary value is displayed. You can also choose to display a secondary value.

To customize the primary and secondary values of a KPI

1. Navigate to the KPI that you want to change, open the **Format visual** menu, and navigate to the **KPI** section.
2. For **Primary value**, use the **Font** dropdown to choose the font size that you want. The default value is **Auto**.
3. (Optional) To change the color of the primary value's font, choose the color icon next to the **Font** dropdown, and then choose the color that you want.

4. For **Primary value displayed**, you can choose to display the actual value or the comparison value of the primary value.
5. To add a secondary value, choose **Secondary value**.
 - a. (Optional) Use the **Font** dropdown to choose the font size that you want. The default value is **Extra large**.
 - b. (Optional) To change the color of the secondary value's font, choose the color icon next to the **Font** dropdown, and then choose the color that you want.

Conditional formatting options for KPIs

Conditional formatting for KPIs is automatically set for comparison values. By default, positive values are represented in green and negative values are represented in red. You can customize the color values of these color values from the **Properties** pane.

To change the color of positive and negative values

1. In the **Properties** pane, open the **Conditional formatting** section and choose the comparison value that you want to change.
2. To change the color of the positive value, navigate to **Condition #1**, choose the **Color** icon, and then choose the color tht you want.
3. To change the color of the negative value, navigate to **Condition #2**, choose the **Color** icon, and then choose the color tht you want.
4. When you are finished making the changes that you want, choose **Apply**.

You can also add text colors and icons for the **Actual value** in thee **Conditional formatting** menu. To add a text color or icon to the actual value, choose **Add text color** or **Add icon** to set the new values.

Labels on visual types in Quick Suite

Use the following procedure to customize, display, or hide the labels for a visual.

To customize, display, or hide the labels for a visual

1. On the analysis page, choose the visual that you want to format. You can change the labels by choosing the label directly on the visual, and choosing **Rename**. To revert to the default name, delete your entry.

2. To see more options, choose the on-visual menu from the down icon at the upper-right corner of the visual, and then choose the **Format visual** icon.

For pivot tables, you can relabel row names, column names, and value names. Additionally, under **Styling**, you can choose to hide columns labels or metric labels (for single metrics only).

You can add the same value to the same visual multiple times. You can do so to show the same value with different aggregations or table calculations applied. By default, the fields all display the same label. You can edit the names by using the **Properties** panel, which you open by choosing the **V**-shaped icon at top right.

3. On the **Properties** pane, enable or disable **Show title**. This option removes the axis title.
4. Close the **Properties** pane by choosing the **X** icon in the upper-right corner of the pane.

Data labels on visual types in Quick Suite

To customize data labels on a visual, you can use the **Properties** pane to show data labels, and then use the settings to configure them. Data label customization is supported on bar, line, combo, area, scatterplot, donut, boxplot, waterfall, heatmap, treemap, histogram, funnel, sankey, gauge, radar, and pie charts.

You can customize the following options:

- Position, which determines where the label appears in relation to the data point (for bar, combo, and line charts):
 - For vertical bar charts, you can customize to set position:
 - Above bars
 - Inside of bars
 - Bottom of bars
 - Top of bars
 - For horizontal bar charts, you can customize to set position:
 - Right of bars
 - Inside of bars
 - For line charts, you can customize to set position:
 - Above lines
 - Left or right of points on lines

- Below lines
- For scatter charts, you can customize to set position:
 - Above points
 - Left or right of points
 - Below points
- Font size and color (for bar, combo, line, scatter, and pie charts)
- Label pattern, which determines how data is labeled (for bar, combo, line, and scatter charts):
 - For bar, combo, and scatter charts, you can label:
 - All
 - By group or color
 - For line charts, the following label options are available:
 - All
 - By group or color
 - Line ends
 - Minimum or maximum value only
 - Minimum and maximum values
 - For pie charts, the following label options are available:
 - Show category
 - Show metric
 - Choose to show the metric label as value, percent, or both
- Group selection (for bars and lines, when the label pattern is "by group/color")
- Allow labels to overlap (for bars and lines), for use with fewer data points
- For vertical bar, combo, and line charts, labels that are too long are angled by default. You can configure the degree of angle under the **X-axis** settings.

Note

If you add more than one measure to an axis, the data label displays the formatting for the first measure only.

To configure data labels

1. On the analysis page, choose the visual that you want to format.
2. Choose the on-visual menu from the down icon at the upper-right corner of the visual, and then choose the **Format visual** icon.
3. On the **Properties** pane, choose **Data Labels**.
4. Enable **Show data labels** to show and customize labels. Disable this option to hide data labels.
5. Choose the settings that you want to use. The settings offered are slightly different for each chart type. To see all available options, see the list before this procedure.

You can immediately view the effect of each change on the visual.

6. To modify the data label font settings, adjust the following properties:
 - **Font family**
 - **Text size**
 - **Style** (Bold, Italic)
 - **Color**
7. Close the **Properties** pane by choosing the X icon in the upper-right corner of the pane.

Formatting visual numeric data based on language settings in Quick Suite

In Amazon Quick Suite, you can choose how your numeric data values appear in visuals so that they align with the regional language that you have chosen.

As a Quick Suite author, you can choose the language formatting that best fits your audience. Amazon Quick Suite configures numeric data languages at the analysis level based on the language that you have chosen to view Quick Suite in. You can change the format of numbers, currencies, and dates. You can change your Quick Suite language settings in the **Language** dropdown list of the Quick Suite **User** menu in the top-right corner. You can change the language formatting for a field across every visual in a sheet, or you can change the language formatting at the individual visual level.

To change the numeric language formatting of all visuals in an analysis

1. On the **Visuals** pane of the analysis that you want to change, choose the more actions (three dots) icon next to the field that you want to change. From the menu that appears, open the **Format** dropdown list, and then choose **More formatting options**.

2. In the **Format data** pane that appears on the left, choose **Apply language format**.

You can reset the default language format of the field by reopening the **Format data** menu and choosing **Reset to defaults**. The default language format is American English.

To change the numeric language formatting of a single visual in an analysis

1. On the analysis page, choose the visual that you want to modify.
2. Navigate to the **Format data** pane using one of the following options:
 - On the visual that contains the data that you want to change, select the field that you want to change, open the **Format** dropdown list, and then choose **More formatting options**.
 - In the **Field wells** section of the analysis, open the dropdown next to the field that you want to change. Open the **Format** menu, and choose **More formatting options**.
3. In the **Format data** pane that appears, choose **Apply language format**.

You can reset the default language format of the visual by reopening the **Format data** menu and choosing **Reset to defaults**. The default language format is American English.

Legends on visual types in Quick Suite

The *visual legend* helps you identify what a visual element represents by mapping its value to a color. By default, the visual legend displays to the right of the visual. You can choose to hide or display the visual legend, and format the legend title and position. You can also customize the font settings for the legend title and items.

To display or hide a visual legend

1. Sign in to Quick Suite at <https://quicksight.aws.amazon.com/>.
2. On the analysis page, choose the visual that you want to format.
3. Choose the visual that you want to format, and then choose the **Properties** icon to open the Properties pane.
4. Toggle the **Legend** on to display the visual's legend. When shown, the legend displays the values in alphabetical order. To hide the legend, toggle the **Legend** switch off.

To customize a visual legend

1. Open the Properties pane and expand the **Legend** section.
2. Use the **Position** dropdown to customize the position of the legend in the visual.
3. For **Legend title**, enter a custom name for the legend and perform all or some of the following actions:
 - a. (Optional) To change the color of the legend title, choose the color swatch underneath the legend title, and then choose the color that you want the legend title to be.
 - b. (Optional) To change the font or font size of the legend title, open the **Font** or **Font size** dropdown and choose the font or font size that you want.
 - c. (Optional) To bold, italicize, or underline the legend title, choose the appropriate icon from the style bar.
4. For **Legend item**, perform all or some of the following actions:
 - a. (Optional) To change the color of the legend item font, choose the color swatch, and then choose the color that you want the legend title to be.
 - b. (Optional) To change the font or font size of the legend item, open the **Font** or **Font size** dropdown and choose the font or font size that you want.
 - c. (Optional) To bold, italicize, or underline the legend item font, choose the appropriate icon from the style bar.
5. Choose the **X** icon at upper right to close the **Properties** pane.

Line and marker styling on line charts in Quick Suite

In Quick Suite line charts, you have multiple options to emphasize what you want readers to focus on: color, line style, and markers. You can use these options together or separately to help readers understand your line charts more quickly under different circumstances. For example, if some of your readers won't see color differences—perhaps because of color blindness or because of monochrome printing—you can use line patterns to distinguish one or more lines in a chart.

In other cases, you could use step lines to call attention to abrupt changes or intervals between changes in data. For example, let's say you build a chart showing the changing price of postage stamps in the US, and you want to emphasize the amount of increase in price over time. You can use a step line, which remains flat between data points until the next price change occurs. The data story about abrupt increases in price is more clear to the reader with a step line. If you wanted

to show a story of gradual change over time, you'd be more likely to style the line with a smooth slope instead.

To customize the styling for a visualization

1. Open your analysis, and choose the chart that you want to format.
2. On the top right of the visual you want to format, select **Format visual**, which is represented by a pencil icon.
3. At left, choose **Data series**.
4. Choose one of the following options:
 - **Base style** – to edit the styling of all lines and markers on the chart
 - **Select series to style** – to edit the styling of the field that you choose from the list

Different options display depending on how many compatible fields are in the visual.

5. Toggle **Line** to turn line styling on or off.

You can customize the following line options:

- The weight or thickness of the line.
 - The style of the line: solid, dashed, or dotted.
 - The color of the line.
 - The type of line that it is: Linear, Smooth, or Stepped.
6. Toggle **Marker** to turn marker styling on or off.

You can customize the following marker options:

- The weight or thickness of the marker.
 - The style of the marker: circle, triangle, square, diamond, and so on.
 - The color of the marker.
7. For **Axis**, choose whether to display the axis on the left or the right.
 8. Your changes are saved automatically.
 9. (Optional) To undo customizations, choose one or more of the following options:
 - To undo one change, click the undo arrow at top left. Repeat as needed. There is also a redo arrow.

- To reset the base style for a data series, select **Base style** and then click **Reset to default**.
- To remove all styling from a data series, listed in **Styled series**, select a field and then click **Remove styling**.

Missing data on visual types in Quick Suite

You can customize how missing data points are visualized in your line charts and area charts. You can choose to have your missing data points appear in the following formats:

- *Broken line*: A disjointed line that breaks when a data point is missing. This is the default missing data format.
- *Continuous line*: Displays a continuous line by skipping over the missing data point and connecting the line to the next available data point in the series. To show a continuous line, the **Show date gaps** box on the **X axis** pane should be unchecked.
- *Show as zero*: Sets the value of the missing data point to zero.

To customize a visual's missing data settings

1. On the analysis page, choose the visual that you want to format.
2. Choose the **Format visual** icon in the upper right corner of the visual to access the **Format visual** menu.
3. Open the **Y axis** pane of the format visual menu and navigate to the **Missing data** section.
4. Select the missing data format that you want.

Reference lines on visuals types in Quick Suite

Reference lines are visual markings in a visual, similar to ruler lines. You typically use a reference line for a value that needs to be displayed with the data. You use the reference line to communicate thresholds or limits in values. The reference line isn't part of the data that's used to build a chart. Instead, it's based on a value that you enter or a field that you identify in the dataset used by a chart.

Quick Suite supports reference lines in the following:

- Bar charts

- Line charts
- Combo charts

You can create, change, and delete reference lines while designing an analysis. You can customize the line pattern, the label font, and the colors for each of those separately. You can show numeric values as numbers, currency, or percent. You can also customize a value's numerical format in the same way that you can customize a field in the field well.

There are two types of reference lines:

- A *constant line* displays at a position that's based on a value that you specify in the format settings. This value doesn't need to relate to any field. You can customize the formatting of the line.
- A *calculated line* displays at a position that's based on a value that is the result of a function. During configuration, you specify which measure (metric) you want to use and which aggregation to apply. These are the same aggregations you can apply to in the field wells. Then, you need to provide an aggregation to apply to the field calculation for the reference line, for example average, minimum, maximum, or percentile. The field needs to be in the dataset used by the chart, although it doesn't need to be displayed in the chart's field wells.

Calculated reference lines aren't supported in 100% stacked charts.

To add or edit a reference line (console)

1. Open the [Quick Suite console](#).
2. Open the analysis that you want to change.
3. Choose the visual that you want to change and open the **Properties** menu.
4. In the **Properties** pane that opens, open the **Reference lines** dropdown, and then choose **ADD NEW LINE**.
5. The **New reference line** menu opens. Use this menu to configure your new reference line. The list below describes all reference line properties that can be configured.
 - **Data**
 - **Type** – The type of reference line that you want to use. Choose one of the following options:
 - To create a constant line based on a single value that you enter, choose **Constant line**.

- To create a calculated line based on a field, choose **Calculated line**.
- **Value** – (For constant lines only) The value that you want to use. This becomes the location of the line on the visual. It appears immediately, so you can experiment with the setting.
- **Column** – (For calculated lines only) The column that you want to use for the reference line.
- **Aggregated as (column)** – (For calculated lines only) The aggregation that you want to apply to the selected column.
- **Calculate** – (For calculated lines only) The calculation that you want to apply to the aggregation.
- **Percentile value** – (Only if you set **Calculate** to **Percentile**) Enter a number from 1 through 100.
- **Chart type** – (For combo charts) Choose **Bars** or **Lines**.
- **Line style**
 - **Pattern** – The pattern used for the line. Valid options include **Dashed**, **Dotted**, and **Solid**.
 - **Color** – The color used for the line.
- **Label**
 - **Type** – The type of label to display. Valid options include **Value only**, **Custom text**, **Custom text and value**, **No label**. If you choose an option that includes custom text, enter the label text that you want to appear on the line.
 - **Enter custom text (text box)** – (Only if you set **Type** to **Custom text and value**) Choose where to show the value in relation to the label. Valid options are **Left** or **Right**.
 - **Position** – The position of the label in relation to the line. Valid options include a combination of the following: left, middle, right, above, and below.
 - **Value format** – The format to use for the value. Choose one of the following:
 - **Same as value** – Uses the formatting that's already selected for this field in the visualization.
 - **Show as** – Choose from the available options, for example number, currency, or percent.
 - **Format** – Choose from the available formatting options.
 - **Font size** – The font size to use for the label text.
 - **Color** – The color to use for the label text.

To list existing reference lines

1. Choose the visual that you want to change and open the **Properties** pane.
2. In the **Properties** pane, open the **Reference lines** dropdown, and then choose the ellipsis (three dots) next to the line that you want to change.
3. Choose **Edit**.
4. The **New reference line** menu opens. Use this menu to make changes to your reference line. When you are finished, choose **Done**.

To disable a reference line

1. Choose the visual that you want to change and open the **Properties** pane.
2. In the **Properties** pane, open the **Reference lines** dropdown, and then choose the ellipsis (three dots) next to the line that you want to change.
3. Choose **Disable**.

To delete a reference line

1. Choose the visual that you want to change and open the **Properties** pane.
2. In the **Properties** pane, open the **Reference lines** dropdown, and then choose the ellipsis (three dots) next to the line that you want to change.
3. Choose **Delete**.

Formatting radar charts in Quick Suite

You can customize radar charts in Amazon Quick Suite to arrange your data the way that you want. You can customize the series style, start angle, fill area, and grid shape of a radar chart.

To set the series style of a radar chart

1. Choose the radar chart visual that you want to change, and choose the **Format visual** icon on the top right corner of the visual.
2. In the **Properties** pane on the left, open the **Radar chart** dropdown list.
3. Under **Series style**, choose the style that you want. You can choose between the following styles:

- **LINE.** When selected, the polygons that are created by the data are outlined.
- **AREA.** When selected, the polygons that are created by the data are filled in.

The default selected value is **LINE**.

To choose the start angle of a radar chart

1. Choose the radar chart visual that you want to change, and choose the **Format visual** icon on the top right corner of the visual.
2. In the **Properties** pane on the left, open the **Radar chart** dropdown list.
3. Under **Start angle**, enter the start angle value that you want. The default value is 90 degrees.

To set the fill area of a radar chart

1. Choose the radar chart visual that you want to change, and choose the **Format visual** icon on the top right corner of the visual.
2. In the **Properties** pane on the left, open the **Axis** dropdown list.
3. Select the **Fill grid lines** check box.
4. (Optional) Select colors for the even and odd numbered grid lines.
 - Choose the **Even color** icon that appears, and then choose the color that you want the even numbered grid lines to be. The default color for this value is white.
 - Choose the **Odd color** icon that appears, and then choose the color that you want the odd numbered grid lines to be. The default color for this value is white.

To choose the grid shape of a radar chart

1. Choose the radar chart visual that you want to change, and choose the **Format visual** icon on the top right corner of the visual.
2. In the **Properties** pane on the left, open the **Radar chart** dropdown list.
3. Under **Grid shape**, choose the shape that you want the radar chart grid to be. You can choose between a **POLYGON** and a **CIRCLE**.

Range and scale on visual types in Quick Suite

To change the scale of the values shown on the visual, you can use the **Properties** pane to set the range for one or both axes of the visual. This option is available for the value axes on bar charts, combo charts, line charts, and scatter plots.

By default, the axis range starts at 0 and ends with the highest value for the measure being displayed. For the group-by axis, you can use the data zoom tool on the visual to dynamically adjust the scale.

To set the axis range for a visual

1. On the analysis page, choose the visual that you want to format.
2. Choose the control menu at the upper-right corner of the visual, and then choose the cog icon.
3. On the **Properties** pane, choose **X-Axis** or **Y-Axis**, depending on what type of visual you are customizing. This is the **X-Axis** section for horizontal bar charts, the **Y-Axis** section for vertical bar charts and line charts, and both axes are available for scatter plots. On combo charts, use **Bars** and **Lines** instead.
4. Enter a new name in the box to rename the axis. To revert to the default name, delete your entry.
5. Set the range for the axis by choosing one of the following options:
 - Choose **Auto (starting at 0)** to have the range start at 0 and end around the highest value for the measure being displayed.
 - Choose **Auto (based on data range)** to have the range start at the lowest value for the measure being displayed and end around the highest value for the measure being displayed.
 - Choose **Custom** to have the range start and end at values that you specify.

If you choose **Custom**, enter the start and end values in the fields in that section. Typically, you use integers for the range values. For stacked 100 percent bar charts, use a decimal value to indicate the percentage that you want. For example, if you want the range to be 0–30 percent instead of 0–100 percent, enter 0 for the start value and .3 for the end value.

6. For **Scale**, the default is linear scale. To show logarithmic scale, also called log scale, enable the logarithmic option. Quick Suite chooses the axis labels to display based on the range of values in that axis.

- On a linear scale, the axis labels are evenly spaced to show the arithmetical difference between them. The labels display the numbers in sets like {1000, 2000, 3000...} or {0, 50 million, 100 million...}, but not {10 thousand, 1 million, 1 billion...}.

Use a *linear scale* for the following cases:

- All the numbers that display on the chart are in the same order of magnitude.
- You want the axis labels to be evenly spaced.
- The axis values have a similar number of digits, for example 100, 200, 300, and so on.
- The rate of change between numbers is relatively slow and steady—in other words, your trend line never approaches becoming vertical.

Examples:

- Profits in different regions of the same country
- Costs incurred for manufacture of an item
- On a *logarithmic scale*, the axis values are spaced to show the orders of magnitude as a way of comparing them. The log scale is often used to display very large ranges of values or percentages, or to show exponential growth.

Use logarithmic scale for the following cases:

- The numbers that display on the chart aren't in the same order of magnitude.
- You want the axis labels to be flexibly spaced to reflect the wide range of values in that axis. This might mean that the axis values have a different number of digits, for example 10, 100, 1000, and so on. It might also mean that the axis labels are unevenly spaced.
- The rate of change between numbers is growing exponentially or is too large to display in a meaningful way.
- The customer of your chart understands how to interpret data on a log scale.
- The chart displays values that growing faster and faster. Moving given distance on the scale means the number has been multiplied by another number.

Examples:

- High yield stock prices over a long range of time
- Growth of pandemic infection rates

7. To customize the number of values to show on the axis labels, enter in an integer between 1 and 50.

8. For combo charts, choose **Single Y Axis** to synchronize the Y-axes for both bars and lines into a single axis.
9. Close the **Properties** pane by choosing the **X** icon in the upper-right corner of the pane.

Small multiples axis options

You can configure the x and y axes for each individual panel of a small multiples visual. You can group your data along an independent x-axis or an independent y-axis. You can also position the x and y axes inside or outside the chart to improve the readability of your data.

For small multiples visuals that use an independent x-axis, only the values that are relevant to each panel are shown on the axis. For example, say you have a small multiples visual that uses one panel to represent each region of the United States. With an independent x-axis, each panel only shows states in the region that the panel represents and hides states that are outside of the panel's region.

For small multiples visuals that use an independent y-axis, each panel uses its own y-axis scale that is determined by the range of the data it contains. By default, data labels appear on the inside of the panel.

To configure independent axes for small multiples visuals

1. Select the small multiples visual that you want to change and open the **Format visual** menu.
2. In the **Properties** pane that appears, open the **Multiples options** menu.
3. For **X-axis**, choose **Independent** from the dropdown.

Or, for **Y-axis**, choose **Independent** from the dropdown.

You can revert your changes by choosing **Shared** from the **X-axis** or **Y-axis** dropdown menus.

You can also configure the label positions of the x and y axes of all panels in a small multiples visual. You can choose to display axis labels inside or outside the panel.

To configure the axis label position for small multiples visuals

1. Select the small multiples visual that you want to change and open the **Format visual** menu.
2. In the **Properties** pane that appears, open the **Multiples options** menu.
3. For **X-axis labels**, choose **Inside** or **Outside** from the dropdown.

Or, for **Y-axis labels**, choose **Inside** or **Outside** from the dropdown.

Titles and subtitles on visual types in Quick Suite

In Quick Suite, you can format visual titles and subtitles to meet your business needs. Quick Suite offers rich text formatting for titles and subtitles, and the ability to add hyperlinks and parameters in titles. You can edit titles in the Properties pane, or by double-clicking on a title or subtitle in the visual.

Use the following procedure to customize the way the title and subtitle of a visual is displayed. The visual title is shown by default. After subtitles are created, they're also shown by default.

1. Sign in to Quick Suite at <https://quicksight.aws.amazon.com/>.
2. Open the analysis that you want to update.
3. On the analysis page, choose the visual that you want to format.
4. At the visual's right, choose the **Properties** icon.
5. In the **Properties** pane that opens, choose the **Display settings** tab.
6. To edit the title or subtitle of a visual, choose the paintbrush icon next to **Edit title** or **Edit subtitle**. Alternatively, you can choose the eyeball icon next to **Edit title** or **Edit subtitle** to hide the title or subtitle, shown in the following image.
7. In the **Edit title** or **Edit subtitle** popup that opens, you can use the following options to make the updates that you want:
 - To enter a custom title or subtitle, enter your title or subtitle text in the editor. Titles can be up to 120 characters long, including spaces. Subtitles can be up to 500 characters long.
 - To change the font type, choose a font type from the list at left.
 - To change the font size, choose a size from the list at right.
 - To change the font weight and emphasis, or to underline or strikethrough text, choose the bold, emphasis, underline, or strikethrough icons.
 - To change the font color, choose the color (Abc) icon, and then pick a color. You can also enter a hexadecimal number or RGB values.
 - To add an unordered list, choose the unordered list icon.
 - To change the text alignment, choose the left, center, or right alignment icons.

- To add a parameter to a title or subtitle, choose an existing parameter from the list under **Parameters** at right. For more information about how to create parameters, see [Setting up parameters in Amazon Quick Suite](#).
- To add a hyperlink, highlight the text that you want to link, choose the hyperlink icon, and then choose from the following options:
 - For **Enter link**, enter the URL that you want to link to.

Choose the + icon at right to add an existing parameter, function, or computation to the URL.

- To edit the display text, enter text for **Display text**.
- To open the hyperlink in the same browser tab as Quick Suite, select **Same tab**.
- To open the hyperlink in a new browser tab, select **New tab**.
- To delete the hyperlink, choose the delete icon at bottom left.

When finished configuring the hyperlink, choose **Save**.

8. When you are finished, choose **Save**.
9. For **Alt text**, enter the alt text that you want for the visual.
10. When you are finished, close the properties pane.

Tooltips on visual types in Quick Suite

When you hover your cursor over any graphical element in an Quick Suite visual, a tooltip appears with information about that specific element. For example, when you hover your cursor over dates in a line chart, a tooltip appears with information about those dates. By default, the fields in the Fields well determine what information displays in tooltips. Tooltips can display up to 10 fields.

You can provide your viewers with additional information about data in your visual, customizing what viewers can see. You can even prevent tooltips from appearing when viewers hover a cursor over an element. To do this, you can customize the tooltips for that visual.

Customizing tooltips in a visual

Use the following procedure to customize tooltips in a visual.

To customize tooltips in a visual

1. On the analysis page, choose the visual that you want to format.

2. On the menu in the upper-right corner of the visual, choose the **Format visual** icon.
3. In the **Properties** pane that opens, choose **Tooltip**.
4. For **Type**, choose **Detailed tooltip**. A new set of options appear.

To show or hide titles in a tooltip

- Choose **Use primary value as title**.

Clearing the option hides titles in the tooltip. Selecting the option shows the primary field value as the title in the tooltip.

To show or hide aggregations for fields in the tooltip

- Choose **Show aggregations**.

Clearing the option hides the aggregation for fields in the tooltip. Selecting the option shows the aggregation for fields in the tooltip.

To add a field to the tooltip

1. Choose **Add field**.
2. In the **Add field to tooltip** page that opens, choose **Select field** and then select a field from the list.

You can add up to 10 fields to tooltips.

3. (Optional) For **Label**, enter a label for the field. This option creates a custom label for the field in the tooltip.
4. (Optional) Depending on whether you add a dimension or a measure, choose how you want the aggregation to display in the tooltip. If you don't select an option, Quick Suite uses the default aggregation.

If you add a measure to the tooltip, you can select how you want the field to be aggregated. To do so, choose **Select aggregation**, and then select an aggregation from the list. For more information about the types of aggregations in Quick Suite, see [Changing field aggregation](#).

5. Choose **Save**.

A new field is added to the list of fields in your tooltip.

To remove a field from the tooltip

- Under the **Fields** list, select the field menu for the field that you want to remove (the three dots) and choose **Hide**.

To rearrange the order of the fields in the tooltip

- Under the **Fields** list, select the field menu for a field (the three dots) and choose either **Move up** or **Move down**.

To customize the label for a field in the tooltip

- Select the field menu for the field that you want to customize (the three dots) and choose **Edit**.
- In the **Edit tooltip field** page that opens, for **Label**, enter the label that you want to appear in the tooltip.
- Choose **Save**.

Hiding tooltips in a visual

If you don't want tooltips to appear when you hover your cursor over data in a visual, you can hide them.

To hide tooltips in a visual

- On the analysis page, choose the visual that you want to format.
- On the menu in the upper-right corner of the visual, choose the **Format visual** icon.
- In the **Properties** pane that opens, choose **Tooltip**.
- Choose **Show tooltip**.

Clearing the option hides tooltips for the visual. Selecting the option shows them.

Customizing data presentation

To gain further insight into your data when creating visuals (charts) in an Quick Suite analysis, you can sort and filter data in a visual. You can also change the granularity of date fields, data type, role, and format of fields in a visual.

Topics

- [Changing fields used by a visual in Amazon Quick Suite](#)
- [Sorting visual data in Amazon Quick Suite](#)

Changing fields used by a visual in Amazon Quick Suite

You can add or modify fields for a visual by using the **Fields list** pane, the field wells, or the on-visual editors or drop targets on the visual.

The field wells, on-visual editors, and drop targets available for a specific visual depends on the visual type selected. For details, see the appropriate visual type topic in the [Visual types in Amazon Quick Sight](#) section.

Important

You can also change the data type and format of numeric fields by using field wells and on-visual editors. If you change a field in this way, it changes for the selected visual only. For more information about changing numeric field data types and formats, see [Changing fields used by a visual in Amazon Quick Suite](#).

Use the following topics to learn more about adding, removing, and modifying fields on a visual.

Topics

- [Using visual field controls](#)
- [Adding or removing a field](#)
- [Changing the field associated with a visual element](#)
- [Changing field aggregation](#)
- [Changing date field granularity](#)
- [Customizing a field format](#)

Using visual field controls

You can edit the fields used by a visual with user interface (UI) controls.

You can use these controls as follows:

- Create a visual and assign fields to different elements on it by selecting fields in the **Fields list** pane, or dragging fields to field wells or drop targets.
- Change the field associated with a visual element by dragging a field to a drop target or field well, or selecting a different field in a field well or on-visual editor.
- Change field aggregation or date granularity by using the field wells or the on-visual editors.

The field wells, on-visual editors, and drop targets available on a specific visual depends on the visual type selected.

Dragging fields to drop targets or field wells

When you drag a field to either a drop target or field well, Amazon Quick Suite provides you with information about whether the target element expects a measure or a dimension. Amazon Quick Suite also provides you with information about whether that element is available for field assignment.

For example, when you drag a measure to the value drop target on a new single-measure line chart, you see the drop target color-coded green. That green color coding indicates that the drop target expects a measure. The drag label indicates that the target is available to add a field.

When you drag a dimension to the x-axis or color drop target on a new line chart, you see a label color-coded blue. That blue color coding indicates that the drop target expects a dimension. The drag label indicates that the target is available to add a field.

You can also drag a measure or dimension to a drop target on a line chart where the element is already associated with a field. In this case, the drag label indicates that you are replacing the field currently associated with the drop target.

Adding or removing a field

You can add a field to a visual by choosing it on the **Fields list** pane. You can also drag it to a drop target on the visual or to a field well. There is a 1:1 correspondence of drop targets to field wells for each visual type, so you can use either method.

On some charts, the **Axis title** field is hidden when there are two or more fields in the **Value** field on any side of the chart. This effect can happen with the following charts:

- Bar charts
- Line charts

- Box plots
- Combo charts
- Waterfall charts

To remove a field from a visual, clear selection from it in the **Fields list** pane. Or choose an on-visual editor or field well that uses that field, and then choose **Remove** from the context (right-click) menu.

Adding a field by selecting it in the fields list pane

You can also let Amazon Quick Suite map the field to the most appropriate visual element. To do so, choose the field in the **Fields list** pane. Amazon Quick Suite adds the field to the visual by populating the first empty field well that corresponds with that field type (either measure or dimension). If all of the visual elements are already populated, Amazon Quick Suite determines the most appropriate field well and replaces the field in it with the field you selected.

Adding a field by using a drop target

To add a field to a visual by using a drop target, first choose a field in the **Fields list** pane. Then drag the field to your chosen drop target on the visual, making sure the drop indicator shows that the field is being added.

Adding a field by using a field well

To add a field to a visual by using a field well, choose a field in the **Fields list** pane. Then drag the field to the target field well, making sure that the drop indicator shows that the field is being added.

1. Drag a field item into a **Field well**.
2. Drag the field that you want to add from the **Fields list** pane to the appropriate field well.

Note

You can add the same value to the same visual multiple times. You can do so to show the same value with different aggregations or table calculations applied. By default, the fields all display the same label. You can edit the names by using the **Properties** panel, which you open by choosing the **V-shaped icon** at top right.

Changing the field associated with a visual element

You can change the field assigned to an element in a visual by using the field wells, drop targets, or the on-visual editors on the visual. For pivot tables, use field wells or drop targets because this visual type doesn't provide on-visual editors.

Change a field mapping by using an on-visual editor

Use the following procedure to modify the mapping of a field to a visual element.

To modify the mapping of a field by using an on-visual editor

1. On the visual, choose the on-visual editor for the visual element for which you want to change the field.
2. On the on-visual editor menu, choose the field that you want to associate with that visual element.

Changing a field mapping by using a drop target

To modify the mapping of a field to a visual element by using a drop target, choose a field in the **Fields list** pane. Then drag the field to a drop target on the visual, making sure that the drop indicator shows that the field is being replaced.

Changing a field mapping by using a field well

Use the following procedure to modify the mapping of a field to a visual element.

To modify the mapping of a field by using a field well

1. Drag a field item into a **Field well**.
2. Choose the field well that represents the element that you want to remap, and then choose a new field from the menu that appears.

Changing field aggregation

You can apply functions to fields to display aggregate information, like the sum of the sales for a given product. You can apply an aggregate function by using the options in either an on-visual editor or a field well. The following aggregate functions are available in Amazon Quick Suite:

- **Average** – Calculates the average value for the selected field.

- **Count** – Provides a count of the number of records containing the selected measure for a given dimension. An example is a count of Order ID by State.
- **Distinct Count** – Provides a count of how many different values are in the selected measure, for the selected dimension or dimensions. An example is a count of Product by Region. A simple count can show how many products are sold for each region. A distinct count can show how many different products are sold for each region. You might have sold 2,000 items, but only two different types of items.
- **Max** – Calculates the maximum value for the selected field.
- **Min** – Calculates the minimum value for the selected field.
- **Median** – Calculates the median value of the specified measure, grouped by the chosen dimension or dimensions.
- **Sum** – Totals all of the values for the selected field.
- **Standard Deviation** – Calculates the standard deviation of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a sample or on a biased population.
- **Variance** – Calculates the variance of the set of numbers in the specified measure, grouped by the chosen dimension or dimensions, based on a sample or on a biased population.
- **Percentile** – Computes the *n*th percentile of the specified measure, grouped by the chosen dimension or dimensions.

All aggregate functions can be applied to numeric fields. *Count* is automatically applied to a dimension if you choose to use it in a field well that expects a measure. If you have used a dimension in that way, you can also change the aggregate function applied to it. You can't apply aggregate functions to fields in dimension field wells.

The visual elements that support aggregated fields varies by visual type.

Changing or adding aggregation on a field by using an on-visual editor

Use the following procedure to change or add aggregation on a field.

To change or add aggregation on a field

1. On the visual, choose the on-visual editor for the field that you want to apply aggregation to.
2. On the on-visual editor menu, choose **Aggregate**, then choose the aggregate function that you want to apply.

Changing or adding aggregation to a field by using a field well

Use the following procedure to add aggregation to a field for a pivot table visual.

To add aggregation to a field for a pivot table visual

1. Drag a field item into a **Field well**.
2. Choose the field well containing the field that you want to apply an aggregate function to.
3. On the field well menu, choose **Aggregate**, then choose the aggregate function that you want to apply.

Changing date field granularity

You can change the granularity for a date field on a visual to determine the intervals for which item values are shown. You can set the date field granularity to one of the following values:

- Year
- Quarter
- Month
- Week
- Day (this is the default)
- Hour
- Minute
- Second

Hour and minute are available only if the field contains time data.

Changing date field granularity by using an on-visual editor

Use the following procedure to change date field granularity by using an on-visual editor.

To change date field granularity with an on-visual editor

1. On the visual, choose the field well for the date field whose granularity you want to change.
2. On the field well menu, choose **Aggregate**, then choose the time interval that you want to apply, as shown following:

Changing date field granularity by using a field well

Use the following procedure to change date field granularity by using a field well.

To change date field granularity with a field well

1. Drag a field item into a **Field well**.
2. Choose the field well containing the date field, and then choose **Aggregate**. Choose the date granularity that you want to use.

Customizing a field format

Use the following procedure to customize the appearance of fields in an analysis.

To customize the appearance of fields in an analysis

1. In an analysis, choose a field to format, either by choosing it in the field well or in the **Fields list** of the **Visualize** pane.
2. Choose **Show as** to change how the field shows in the analysis, and choose from the options on the context menu. The list of available options varies based on the field's data type. If you choose a non-numeric field from the fields list, you can change the *count format*, which is the formatting used when the field is counted.
3. Choose **Format** to change the format of the field, and choose from the options on the context menu. If you don't see an option that you want to use, choose **More formatting options** from the context menu.

The **Format Data** pane opens, presenting options for the type of numeric or date field you chose.

The options for **Show as** from the context menu now appear in the drop-down list at the top of the **Format Data** pane. The rest of the options are specific to the data type and how you choose to show the field.

For date and time data, the default format pattern is YYYY-MM-DDTHH:mm:ssZZ, for example 2016-09-22T17:00:00-07:00.

For numbers, you can choose from the following units to display after the number:

- No unit suffix. This is the default.

- Thousands (K)
- Millions (M)
- Billions (B)
- Trillions (T)
- A custom unit prefix or suffix

For currency, you can choose from the following symbols:

- Dollars (\$)
- Euros (€)
- Pounds (£)
- Yen (¥)

Changing a field format

You can change the format of a field within the context of an analysis. The formatting options available for fields vary based on the field's data type.

Use menu options in the **Field list** pane or the visual field wells to make simple format changes, or use the **Format data** pane to make more extensive formatting changes.

Topics

- [Format a currency field](#)
- [Format a date field](#)
- [Format a number field](#)
- [Format a percent field](#)
- [Format a text field](#)
- [Return a field's format to default settings](#)

Format a currency field

When you format a currency field, you can either choose the currency symbol from a list of common options, or open the **Format data** pane and manually format the field. Manually

formatting the field allows you to choose which symbol to use, which separators to use, the number of decimal places to show, which units to use, and how to display negative numbers.

Changing a field format changes it for all visuals in the analysis, but does not change it in the underlying dataset.

If you want to choose the symbol for a currency field from a list of common options, you can access such a list in several ways. You can access it from the **Field list** pane, an on-visual editor, or a visual field well.

To select a currency field's symbol by choosing a list option

1. Choose one of the following options:
 - In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
 - On any visual that contains an on-visual editor associated with the currency field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.
2. Choose **Format**, and then choose the currency field that you want:
 - Display in dollars (\$).
 - Display in pounds (£).
 - Display in euros (€).
 - Display in yen or yuan (¥).

To manually change a currency field's format

1. Choose one of the following options:
 - In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
 - On any visual that contains an on-visual editor associated with the number field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.
2. Choose **Format**, and then choose **More Formatting Options**.

The **Format data** pane opens.

3. Expand the **Symbol** section and choose from the following options:
 - Display in dollars (\$). This is the default.
 - Display in pounds (£).
 - Display in euros (€).
 - Display in yen or yuan (¥).
4. Expand the **Separators** section and choose from the following options:
 - Under **Decimal**, choose a dot or a comma for the decimal separator. A dot is the default. If you choose a comma instead, use a dot or a space as the thousands separator.
 - Under **Thousands**, select or clear **Enabled** to indicate whether you want to use a thousands separator. **Enabled** is selected by default.
 - If you are using a thousands separator, choose whether to use a comma, dot, or space for the separator. A comma is the default. If you choose a dot instead, use a comma as the decimal separator.
5. Expand the **Decimal Places** section and choose the number of decimal places to use. The default is 2. Field values are rounded to the decimal places specified. For example, if you specify two decimal places, the value 6.728 is rounded to 6.73.
6. Expand the **Units** section and choose from the following options:
 - Choose the unit to use. Choosing a unit adds the appropriate suffix to the number value. For example, if you choose **Thousands**, a field value of 1234 displays as 1.234K.

The unit options are as follows:

- No unit suffix. This is the default.
 - Thousands (K)
 - Millions (M)
 - Billions (B)
 - Trillions (T)
 - If you want to use a custom prefix or suffix, specify it in the **Prefix** or **Suffix** box. Using a custom suffix is a good way to specify a currency suffix outside of those already offered by Amazon Quick Suite. You can specify both. You can also specify a custom prefix in addition to the suffix added by selecting a unit.
7. Expand the **Negatives** section and choose whether to display a negative value by using a minus sign or by enclosing it in parentheses. Using a minus sign is the default.

8. Expand the **Null values** section and choose whether to display null values as `null` or as a custom value. Using `null` is the default.

 Note

When using a table or pivot table, null values only display for fields that are placed in the **Rows**, **Columns**, or **Group by** field wells. Null values for fields in the **Values** field well appear empty in the table or pivot table.

Format a date field

When you format a date field, you can choose a list of common formatting options. Or you can open the **Format data** pane to choose from a list of common formats, or specify custom formatting for the date and time values.

Changing a field format changes it for all visuals in the analysis that use that dataset, but does not change it in the dataset itself.

If you want to format a date field by choosing from a list of common options, you can access such a list in several ways. You can access it from the **Field list** pane, a visual on-visual editor, or a visual field well.

To change a date field's format by choosing a list option

1. Choose one of the following options:
 - In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
 - On any visual that contains an on-visual editor associated with the number field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.
2. Choose **Format**, and then choose the format that you want. The following quick formatting options are offered for date fields:
 - Show the month, day, year, and time.
 - Show the month, day, and year.
 - Show the month and year.
 - Show the year.

To manually change a date field's format

1. Choose one of the following options:
 - In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
 - On any visual that contains an on-visual editor associated with the number field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.
2. Choose **Format**, and then choose **More Formatting Options**.

The **Format data** pane opens.

3. Expand the **Date** section. Choose an existing date format, or choose **Custom** and specify a format pattern in the **Custom** section lower down in the **Format data** pane. If you choose **Custom** for the **Date** section, you must also choose **Custom** for the following **Time** section. The pattern that you specify in the **Custom** section must include any date and time formatting that you want.

The default selection is **Custom**, with a default format pattern of MMM D, YYYY h:mma, for example Sep 20, 2022 5:30pm.

4. Expand the **Time** section. Choose an existing time format, or choose **Custom** and specify a format pattern in the **Custom** section lower down in the **Format data** pane. If you choose **Custom** for the **Time** section, you must also choose **Custom** for the preceding **Date** section. The pattern that you specify in the **Custom** section must include any date and time formatting that you want.

The default selection is **Custom**, with a default format pattern of MMM D, YYYY h:mma, for example Sep 20, 2022 5:30pm.

5. If you chose **Custom** in the **Date** and **Time** sections, expand the **Custom** section and specify the format pattern that you want, using the format pattern syntax specified in [Moment.js Display Format](#) in the Moment.js JavaScript documentation.

Note

The time zone related display token Z from the Moment.js library is supported in Quick Suite, but the z token is not.

If you chose something other than **Custom** in the **Date** and **Time** sections, **Custom** is populated with the format pattern that reflects your selections. For example, if you chose Jun 21, 2016 in the **Date** section and 17:00:00pm in the **Time** section, the **Custom** section shows the format pattern MMM D, YYYY H:mm:ssa.

6. (Optional) Expand the **Custom** section and use **Preview** to verify your specified format.
7. Expand the **Null values** section and choose whether to display null values as null or as a custom value. Using null is the default.

Customizing date formats in Quick Suite

In Quick Suite, you can customize how dates are formatted in your filter and parameter controls. For example, you can specify to format the date in a control as 20-09-2021, or, if you'd rather, as 09-20-2021. You can also specify to shorten the month in your dates (such as September) to three letters (Sep), among other customizations.

Following is a list of tokens you can use to create custom date formats. You can use them in combination with one another to control how dates appear in your controls.

List of supported tokens for formatting dates

Use the following tokens to customize the format of dates in Quick Suite.

Example	Description	Token
0–6	Numeric representation of a particular day of the week. 0 is Sunday and 6 is Saturday.	d
Mo–Su	A 2-character textual representation of a particular day of the week.	dd
Mon–Sun	A 3-character textual representation of a particular day of the week.	ddd

Example	Description	Token
Monday–Sunday	A textual representation of a particular day of the week.	dddd
99 or 21	A 2-digit representation of a year.	YY
1999 or 2021	A full, 4-digit numeric representation of a year.	YYYY
1–12	Number of a month, without leading zeros.	M
1st, 2nd, to 12th	Number of a month without leading zeros and with an ordinal suffix.	Mo
01–12	Number of a month with leading zeros.	MM
Jan–Dec	A 3-digit textual representation of a month.	MMM
January–December	A full textual representation of a month.	MMMM
1–4	A numeric representation of a quarter.	Q
1st–4th	A numeric representation of a quarter with an ordinal suffix.	Qo

Example	Description	Token
1–31	Day of the month without leading zeros.	D
1st, 2nd, to 31st	Day of the month without leading zeros and an ordinal suffix.	Do
01–31	A 2-digit day of the month with leading zeros.	DD
1–365	Day of the year without leading zeros.	DDD
001–365	Day of the year with leading zeros.	DDDD
1–53	Week of the year without leading zeros.	w
1st–53rd	The week of the year without leading zeros and with an ordinal suffix.	wo
01–53rd	Week of the year with leading zeros.	ww
1–23	Hours, in 24-hour format, without leading zeros.	H

Example	Description	Token
01–23	Hours, in 24-hour format, with leading zeros.	HH
1–12	Hours, in 12-hour format, without leading zeros.	h
01–12	Hours, in 12-hour format, with leading zeros.	hh
0–59	Minutes without leading zeros.	m
00–59	Minutes with leading zeros.	mm
0–59	Seconds without leading zeros.	s
00–59	Seconds with leading zeros.	ss
am or pm	am/pm	a
AM or PM	AM/PM	A
1632184215	Unix timestamp.	X
1632184215000	Millisecond Unix timestamp.	x
Z	Zero UTC offset.	Z

The following date types are not supported.

- Time zones offset with a colon. For example, +07:00.
- Time zones offset without a colon. For example, +0730.

Preset date formats

To quickly customize dates and times to appear as one of the following example formats, you can use the following Quick Suite preset tokens.

Example	Token
8:30 PM	LT
8:30:25 PM	LTS
August 2 1985	LL
Aug 2 1985	l1
August 2 1985 08:30 PM	LLL
Aug 2 1985 08:30 PM	l11
Thursday, August 2 1985 08:30 PM	LLLL
Thu, Aug 2 1985 08:30 PM	l111

Common date formats

Following are three common date examples and their associated token formats for your quick reference.

Example	Token Format
Sep 20, 2021	MMM DD, YYYY
20-09-21 5pm	DD-MM-YY ha
Monday, September 20, 2021 17:30:15	dddd, MMMM DD, YYYY HH:mm:ss

Adding words to dates

To include words in your date formats, such as the word "of" in *20th of Sep, 2021*, enter backslashes (\) before each character in the word. For example, for the 20th of Sep, 2021 date example, use the following token format: Do \o\f MMM, YYYY.

Example: Customizing the date format in a filter control

Use the following procedure to learn how to use date token formats to customize dates for a filter control.

To learn to customize dates for a filter control with data tokens

1. In a Quick Suite analysis, choose the filter control that you want to customize.
2. On the filter control, choose the **Edit control** icon.
3. On the **Edit control** page that opens, for **Date format**, enter the custom date format that you want. Use the tokens listed previously in this topic.

For example, let's say that you want to customize your dates using the following format: *Sep 3rd, 2020 at 5pm*. To do so, you can enter the following token format:

MMM Do, YYYY \a\t ha

A preview of the date format appears below the input field as you enter each token.

4. Choose **Apply**.

The dates in the control update to the format you specified.

Format a number field

When you format a number field, you can choose the decimal place and thousand separator format from a list of common options. Or you can open the **Format Data** pane and manually format the field. Manually formatting the field enables you to choose which separators to use and the number of decimal places to show. It also enables you to choose which units to use, and how to display negative numbers.

Changing a field format changes it for all visuals in the analysis, but does not change it in the underlying dataset.

If you want to format a number field by choosing from a list of common options, you can access such a list from the **Field list** pane, an on-visual editor, or a visual field well.

To change a number field's format by choosing a list option:

- Choose one of the following options:
 - In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
 - On any visual that contains an on-visual editor associated with the number field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.
- Choose **Format**, and then choose the format that you want. The following quick formatting options are offered for number fields:
 - Use commas to separate groups of thousands and use a decimal point to show the fractional part of the number, for example 1,234.56.
 - Use a decimal point to show the fractional part of the number, for example 1234.56.
 - Show the number as an integer and use commas to separate groups of thousands, for example 1,234.
 - Show the number as an integer, for example 1234.

To manually change a number field's format:

1. Choose one of the following options:
 - In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.

- On any visual that contains an on-visual editor associated with the number field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.
2. Choose **Format**, and then choose **More Formatting Options**.

The **Format data** pane opens.

3. Expand the **Separators** section and choose from the following options:
 - Under **Decimal**, choose a dot or a comma for the decimal separator. A dot is the default. If you choose a comma instead, use a dot or a space as the thousands separator.
 - Under **Thousands**, select or clear **Enabled** to indicate whether you want to use a thousands separator. **Enabled** is selected by default.
 - If you are using a thousands separator, choose whether to use a comma, dot, or space for the separator. A comma is the default. If you choose a dot instead, use a comma as the decimal separator.
4. Expand the **Decimal Places** section and choose from the following options:
 - Choose **Auto** to have Amazon Quick Suite automatically determine the appropriate number of decimal places, or choose **Custom** to specify a number of decimal places. **Auto** is the default.
 - If you chose **Custom**, enter the number of decimal places to use. Field values are rounded to the decimal places specified. For example, if you specify two decimal places, the value 6.728 is rounded to 6.73.
5. Expand the **Units** section and choose from the following options:
 - Choose the unit to use. Choosing a unit adds the appropriate suffix to the number value. For example, if you choose **Thousands**, a field value of 1234 displays as 1.234K.

The unit options are as follows:

- No unit suffix. This is the default.
- Thousands (K)
- Millions (M)
- Billions (B)
- Trillions (T)
- If you want to use a custom prefix or suffix, specify it in the **Prefix** or **Suffix** box. You can specify both. You can also specify a custom prefix in addition to the suffix added by selecting a unit.

6. Expand the **Negatives** section and choose whether to display a negative value by using a minus sign or by enclosing it in parentheses. Using a minus sign is the default.
7. Expand the **Null values** section and choose whether to display null values as `null` or as a custom value. Using `null` is the default.

 Note

When using a table or pivot table, null values only display for fields that are placed in the **Rows**, **Columns**, or **Group by** field wells. Null values for fields in the **Values** field well appear empty in the table or pivot table.

Format a percent field

When you format a percent field, you can choose the number of decimal places from a list of common options. Or you can open the **Format data** pane and manually format the field. Manually formatting the field enables you to choose which separators to use. It also enables you to choose the number of decimal places to show and how to display negative numbers.

Changing a field format changes it for all visuals in the analysis, but does not change it in the underlying dataset.

If you want to choose the number of decimal places for a percent field from a list of common options, you can access such a list in several ways. You can access it from the **Field list** pane, an on-visual editor, or a visual field well.

To change a percent field's number of decimal places by choosing a list option

1. Choose one of the following options:
 - In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
 - On any visual that contains an on-visual editor associated with the percent field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.
2. Choose **Format**, and then choose the number of decimal places that you want. The following quick formats are offered for percent fields:
 - Display the value with two decimal places.

- Display the value with one decimal place.
- Display the value with no decimal places.

To manually change a percent field's format

1. Choose one of the following options:

- In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
- On any visual that contains an on-visual editor associated with the number field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.

2. Choose **Format**, and then choose **More Formatting Options**.

The **Format data** pane opens.

3. Expand the **Separators** section and choose from the following options:

- Under **Decimal**, choose a dot or a comma for the decimal separator. A dot is the default. If you choose a comma instead, use a dot or a space as the thousands separator.
- Under **Thousands**, select or clear **Enabled** to indicate whether you want to use a thousands separator. **Enabled** is selected by default.
- If you are using a thousands separator, choose whether to use a comma, dot, or space for the separator. A comma is the default. If you choose a dot instead, use a comma as the decimal separator.

4. Expand the **Decimal Places** section and choose from the following options:

- Choose **Auto** to have Amazon Quick Suite automatically determine the appropriate number of decimal places, or choose **Custom** to specify a number of decimal places. **Auto** is the default.
- If you chose **Custom**, enter the number of decimal places to use. Field values are rounded to the decimal places specified. For example, if you specify two decimal places, the value 6.728 is rounded to 6.73.

5. Expand the **Negatives** section and choose whether to display a negative value by using a minus sign or by enclosing it in parentheses. Using a minus sign is the default.

6. Expand the **Null values** section and choose whether to display null values as null or as a custom value. Using null is the default.

 Note

When using a table or pivot table, null values only display for fields that are placed in the **Rows**, **Columns**, or **Group by** field wells. Null values for fields in the **Values** field well appear empty in the table or pivot table.

Format a text field

When you format a text field, you can choose how to display null values using the **Field list** pane, an on-visual editor, or a visual field well.

To choose how to display a text field's null values

1. Choose one of the following options:
 - In the **Field list** pane, choose the selector icon to the right of the number field that you want to format.
 - On any visual that contains an on-visual editor associated with the percent field that you want to format, choose that on-visual editor. Expand the **Field wells** pane, and then choose the field well associated with the number field that you want to change.
2. Choose **Format**, and then choose **More Formatting Options**.

The **Format data** pane opens.

3. Expand the **Null values** section and choose whether to display null values as `null` or as a custom value. Using `null` is the default.

Return a field's format to default settings

Use the following procedure to return a field's format to the default settings.

To return a field's format to the default settings

1. In the **Field list** pane, choose the selector icon to the right of the field that you want to reset.
2. Choose **Format**, and then choose **More Formatting options**.

The **Format data** pane opens.

3. Choose **Reset to defaults**.

Sorting visual data in Amazon Quick Suite

You can sort data using multiple methods for most visual types. You can choose the sort order of on-visual data by using the quick sort option or field wells. You can also use field wells to sort data by an off-visual metric. The visual element you can sort by depends on the visual type and whether sorting is supported for that visual. For more information on which visual types support sorting, see [Analytics formatting per type in Quick Suite](#).

Pivot tables behave differently than tables when sorting values. For more information about sorting pivot tables, see [Sorting pivot tables in Quick Suite](#).

For SPICE datasets, you can sort text strings of sizes up to the following limitations:

- Up to two million (2,000,000) unique values
- Up to 16 columns

When you exceed the limitations, the visual displays a notification at the upper right.

You can sort any visual type that supports sorting. If a visual type supports sorting, you can sort by using either the quick sort option or a field well.

To quickly sort dimensions and measures

- Do one of the following:
 - Choose the sort icon that appears near the field name on either axis. In direct queries, this icon appears for any data type. For SPICE, this icon is available only for datetime, numeric, and decimal data types.
 - Choose the field name and then choose the sort option from the menu. If the label doesn't display on the axis, check the visual format to see if the axis is set to display labels. The display labels are automatically hidden on smaller visuals. You might need to make the visual large enough to display labels.

To sort by using an off-visual metric

1. Open the analysis with the visual that you want to sort. Visuals pane will be open by default.
2. Choose a field well that supports sorting, then choose **Sort by, Sort options**.
3. On the **Sort options** pane, sort by specific fields, choose an aggregation, or sort ascending or descending, or do a combination of these.

4. Choose **Apply** to save your changes. Or choose **Clear** to start over or **Cancel** to go back.

To sort by using a field well

1. Open the analysis with the visual that you want to sort. Visuals pane will be open by default.
2. Choose a field well that supports sorting.
3. On the field well menu, choose **Sort**, and then choose the ascending or descending sort order icon.

Using themes in Amazon Quick Sight

In Amazon Quick Sight, a *theme* is a collection of settings that you can apply to multiple analyses and dashboards. Amazon Quick Sight includes some themes, and you can add your own by using the theme editor. You can share themes with permissions levels set to user or owner. Anyone who has access to the theme can apply it to analyses and dashboards, or use **Save as** to make their own copy of it. Theme owners can also edit the theme and share it with others.

An analysis can have only one theme applied. If you apply a theme to an analysis (by using the **Apply** button), it instantly changes it for everyone—both analysis and dashboard viewers. To explore and save color options without applying them, avoid editing and saving the applied theme.

All colors come in pairs of background and foreground colors. The foreground colors are meant to specifically appear above their matching background color, so choose something that contrasts well.

The following table defines the different settings.

Group	Setting	What the setting changes
Main	Primary background	The background color used for visuals and other high emphasis UI.
Main	Primary foreground	The color of text and other foreground elements that appear over the primary background regions such

Group	Setting	What the setting changes
		as grid lines, borders, table banding, icons, and so on.
Main	Secondary background	The background color used for the sheet background and sheet controls.
Main	Secondary foreground	The foreground color used for any sheet title, sheet control text, or UI that appears over the secondary background.
Main	Accent	This setting is used as an interactive hint for the following: • Buttons • Borders around the selected visual • Loading indicators • Narration customizations • Links • Filter panes for embedded dashboards
Main	Accent foreground	The foreground color applies to any text or other elements that appear over the accent color.

Group	Setting	What the setting changes
Main	Font	The font to use for all of the text. You can choose from a variety of fonts supported by Amazon Quick Sight.
Data	Data colors	These are the data colors that charts rotate through when assigning colors to groups. You can add or remove colors to this list, or choose a color to change it.
Data	Min max gradient	The default minimum and maximum gradient colors to use when a gradient is used as a scale, for example in heat maps.
Data	Empty fill color	This is the color used with your data colors to indicate a lack of data. For example, this color appears in the empty portion of the progress bars that are shown in key performance indication (KPI) and gauge charts, or for empty heat map cells.
Layout	Border	This setting toggles the border around the visuals that aren't currently selected. The selected visual's border still displays the accent color.

Group	Setting	What the setting changes
Layout	Margin	This setting toggles the space between the sheet boundaries and the visuals.
Layout	Gutter	This setting shows or hides the space between visuals in the grid.
Other	Success Success foreground	These colors are used for success messages, for example the check mark for a successful download.
Other	Warning Warning foreground	These colors are used for warning and informational messages.
Other	Danger Danger foreground	These colors are used for error messages.
Other	Dimension Dimension foreground	These colors are used for the names of fields that are identified as dimensions. This option also sets the color for dimensions in the filter panel of embedded dashboards.
Other	Measure Measure foreground	These colors are used for the names of fields that are identified as measures. These colors also apply to measures in the filter panel of embedded dashboards.

To take a short tour of the theme viewer and editor

1. Open the [Quick Suite console](#).
2. Open an analysis, or create a new one. You must have an analysis open to work with themes. However, the view you see with the theme applied is only a preview.

Themes are separate from analyses. No changes are made to your analysis, even when you save a theme.

3. Choose **Edit** from the application bar, and then choose **Themes**. The themes panel opens.
4. The list of themes shows the following:
 - **Applied theme** shows the theme that is currently applied to this analysis and its dashboards.
 - **My themes** shows themes that you created and themes that are shared with you.
 - **Starter themes** shows themes created by Amazon Quick Sight.
5. Each theme has context menu that you can access from the ... icon.

The actions that are available to you on each theme depend on your level of access.

- **Theme owners** – If you created the theme, or someone shared it with you and made you an owner, you can do the following:
 - **Edit** – Change the settings for the theme, and save them.
 - **Save** – Save changes you made to the theme. If you edit the applied theme save your changes, the new theme settings apply to all the analyses and dashboards that use it. An informational message displays before you overwrite an applied theme.
 - **Share** – Share the theme and assign user or owner permissions to other people.
 - **Delete** – Delete a theme. You can't undo this action. An informational message displays before you confirm deletion.
- **Theme users** – If someone shared the theme with you, or if it's an Amazon Quick Sight theme, you can do the following:
 - **Apply** – Apply the theme to the current analysis. This option also applies the theme to dashboards created from the analysis. An informational message displays before you overwrite an applied theme.
 - **Save as** – Save the current theme to another name, so you can edit it.
- **Analysis authors** – If you have access to the analysis, but not the theme, you can do the following:

- You can see the analysis with the theme applied.
 - You can see the theme in the **Theme** panel.
 - You can use **Save as** to create your own copy of the theme.
 - **Dashboard viewers** – If you have access to the dashboard, but not the theme, you can do the following:
 - You can see the dashboard with the theme applied.
 - You can't see the theme or its settings. Dashboard users can't see the **Theme** panel.
6. To explore a theme's settings, choose the icons on the left to see settings for colors.

The following procedure walks you through creating a theme. You can start on the analysis, or a copy of the analysis, that you want to use to preview the colors. Or you can start a new analysis. After you save the theme, you can apply it to the current analysis or to other analyses. If you share it, other people can use it too.

To use the theme editor

1. Open the [Quick Suite console](#).
2. Open an analysis, or create a new one. Choose **Edit** from the application bar, and then choose **Themes**. The **Themes** panel opens.

You must have an analysis open to work with themes. However, the view you see with the theme applied is only a preview. Themes are separate from analyses. No changes are made to your analysis, even when you save a theme.

3. Choose **Main**. The color picker used in each of these settings is the standard one used throughout Amazon Quick Sight.

Set colors for **Primary background** and **Primary foreground** to use in visuals and other high impact UI.

Set colors for **Secondary background** and **Secondary foreground** to use in sheets and sheet controls.

Set colors for **Accent** and **Accent foreground** to use in interactive hints including buttons, borders around selected visuals, loading indicators, narration customizations, links, and the filter pane in embedded dashboards.

4. Choose **Data**.

Set the **Colors** to use as data colors. Charts rotate through these when assigning colors. You can add or delete colors, or change the order they're in by dragging and dropping. To change an existing color, select it to open the color editor.

Set colors for **Min max gradient** to use when a gradient is used as a scale, for example in heat maps.

Set the color for **Empty fill** to use when showing a lack of data, for example the unfilled part of a progress bar.

5. Choose **Layout**.

Enable or disable the **Border** check box to show or hide the border around the visuals that aren't currently selected.

Enable or disable the **Margin** check box to show or hide the space between the sheet boundaries and the visuals.

Enable or disable the **Gutter** check box to show or hide the space between visuals in the grid.

6. Choose **Other**.

Set the color for **Success** to use in success messages, for example when you successfully download a .csv file. The success foreground color isn't currently used.

Set the color for **Warning** to use in warning and informational messages. The warning foreground color isn't currently used.

Set the color for **Danger** to use in error messages. The danger foreground color isn't currently used.

Set the color for **Dimension** to use for the names of fields that are identified as dimensions. This option also sets the color for dimensions in the filter panel of embedded dashboards.

Set the color for **Measure** to use for the names of fields that are identified as measures. This option also sets the color for measures in the filter panel of embedded dashboards.

7. To save the theme, choose **Main** and give the new theme a name, and then choose **Save** at the upper-right of the browser.

Saving a theme doesn't apply it to the analysis, even though you can see a preview of the colors that uses the current analysis.

8. To share the theme, save or close the theme you are viewing. Find the theme in your theme collection. Choose **Share** from the context menu (...).
9. To apply the theme, save or close the theme you are viewing. Find the theme in your theme collection. Choose **Apply** from the context menu (...).

Accessing Amazon Quick Sight using keyboard shortcuts

You can use the following keyboard shortcuts to navigate a Amazon Quick Sight dashboard or analysis:

- Use the **TAB** key to navigate among menu options or visuals.
- Use the **Shift+TAB** keys to move backward to the previous selection.
- Use the **Enter** key to select a visual or menu option.
- Use the **ESC** key to clear the selection from a visual or menu item.

Using shortcuts within a visual

You can use the **TAB**, **Shift+TAB**, and **Enter** keys to navigate and select different fields within a selected visual. For example, say that you want to use a link that's a part of your visuals title. To do this, select the visual that you want, then use the **TAB** key until just the link is selected. Then, use the **Enter** key to click on the link.

You can also use these keyboard shortcuts to navigate and enter the on-visual menu on the upper-right corner of a visual. To do this, select the visual that you want and use the **TAB** key to get to the field that you want to select. If you miss the field that you want, use the **Shift+TAB** keys to go back a field.

Sharing and subscribing to data in Amazon Quick Sight with dashboards and reports

A *dashboard* is a read-only snapshot of an analysis that you can share with other Amazon Quick Sight users for reporting purposes. A dashboard preserves the configuration of the analysis at the time you publish it, including such things as filtering, parameters, controls, and sort order. The

data used for the analysis isn't captured as part of the dashboard. When you view the dashboard, it reflects the current data in the data sets used by the analysis.

When you share a dashboard, you specify which users have access to it. Users who are dashboard viewers can view and filter the dashboard data. Any selections to filters, controls, or sorting that users apply while viewing the dashboard exist only while the user is viewing the dashboard, and aren't saved after it's closed. Users who are dashboard owners can edit and share the dashboard, and optionally can edit and share the analysis. If you want them to also edit and share the data set, you can set that up in the analysis.

A shared dashboard can also be embedded in a website or app, if you are using Enterprise edition. For more information about embedded dashboards, see [Embedded analytics for Amazon Quick Sight](#).

Use the following sections to learn how to publish and share dashboards, subscribe to threshold alerts, and send and subscribe to dashboard email reports.

Topics

- [Publishing dashboards](#)
- [Sharing Amazon Quick Sight dashboards](#)
- [Using Quick action connectors in dashboard visuals](#)
- [Sharing your view of a Amazon Quick Sight dashboard](#)
- [Scheduling and sending Quick Sight reports by email](#)
- [Subscribing to email reports in Amazon Quick Sight](#)
- [Working with threshold alerts in Amazon Quick Sight](#)
- [Printing a dashboard or analysis](#)
- [Exporting Amazon Quick Sight analyses or dashboards as PDFs](#)
- [Error codes for failed PDF export jobs](#)
- [Organizing assets into folders for Amazon Quick Sight](#)

Publishing dashboards

When you publish an analysis, that analysis becomes a dashboard that can be shared and interacted with by users of your Amazon Quick Suite account or, in some cases, with anonymous users that aren't on your account. You can choose to publish one sheet of an analysis, all sheets in the analysis, or any other combination of sheets that you want. When you publish an interactive

sheet, that sheet becomes an interactive dashboard that users can interact with. When you publish a paginated report sheet, the sheet becomes a paginated report that generates and saves a snapshot of the report's data when you schedule a report in Amazon Quick Sight. You can publish a dashboard that contains any combination of interactive sheets and paginated reports from the same analysis.

For more information about scheduling a report, see [Scheduling and sending Quick Sight reports by email](#).

For more information about viewing a report's snapshots, see [Consuming paginated reports in Amazon Quick Sight](#).

Use the following procedure to publish and optionally share a dashboard. You can also use this procedure to rename a published dashboard. A renamed dashboard retains its security and emailed report settings.

1. Open the analysis that you want to use. Choose **Publish**.
2. Do one of the following:
 - To create a new dashboard, choose **New dashboard**, and then type a dashboard name.
 - To replace an existing dashboard, do one of the following. Replacing a dashboard updates it without altering security or emailed report settings.
 - To update it with your changes, choose **Replace an existing dashboard** and then choose a dashboard from the list.
 - To rename it, choose **Replace an existing dashboard**, choose a dashboard from the list, and then select the pencil icon. Enter a new name to rename the existing dashboard and click the checkmark or press enter to confirm. When you publish a dashboard after renaming, it also saves any changes you made to the analysis. Changes to the analysis or dashboard are not persisted until you **Publish**. An initial version of a dashboard must be published in order to rename it.
3. (Optional) Choose the sheets that you want to publish in the **SHEETS** dropdown. When you select sheets to add to the new dashboard, the dropdown shows how many sheets are selected for publishing. The default option is **ALL SHEETS SELECTED**.

If you are replacing an existing dashboard, the sheets that are already published to the existing dashboard are pre-selected in the dropdown, unless you are publishing from an analysis you have not previously published from. You can make changes to this by selecting or de-selecting sheets from the dropdown list.

4. (Optional) Add comments on the changes you have made in the notes section, which is available to view under [Version History](#).
5. (Optional) To allow dashboard readers to share data stories, choose **Allow sharing data stories**. For more information about data stories, see [???](#).
6. (Optional) Open **More Settings**. These options are only available if at least one sheet in the new dashboard is an interactive sheet.

 Note

This is a scrollable window. Scroll down in the **Publish a dashboard** window to view all available options.

There are some options that you can turn off to simplify the experience for this dashboard, as follows:

- For **Dashboard options**:
 - Leave **Expand on-sheet controls by default** cleared to show a simplified view. This is disabled by default. To show the controls by default, turn on this option.
 - Clear **Enable advanced filtering on the left pane** to remove the ability for dashboard viewers to filter the data themselves. If they create their own filters, the filters exist only while the user is viewing the dashboard. Filters can't be saved or reused.
 - Clear **Enable on-hover tooltip** to turn off tooltips.
- For **Visual options**:
 - Clear **Enable visual menu**, to turn off the on-visual menu entirely.
 - Clear **Enable download options** if your dashboard viewers don't need to be able to download data from the visuals in the dashboard. The CSV file includes only what is currently visible in the visual at the time they download it. The viewer downloads data by using the on-visual menu on each individual visual.
 - Clear **Enable maximize visual option** to turn off the ability to enlarge visuals to fill the screen.
- For **Data point options**:
 - Clear **Enable drill up/down** if your dashboard doesn't offer drillable field hierarchies.
 - Clear **Enable on-click tooltip** to turn off tooltips that appear when the reader chooses (clicks on) a data point.

- Clear **Enable sort options** to turn off sorting controls.

7. Choose **Publish dashboard**.

If you renamed the existing dashboard, the top of the screen refreshes to show the new name.

8. (Optional) Do one of the following:

- To publish a dashboard without sharing, choose **x** at the upper right of the **Share dashboard with users** screen when it appears. You can always share the dashboard later by choosing **File>Share** from the application bar.
- To share the dashboard, follow the procedure in [Sharing Amazon Quick Sight dashboards](#).

After you complete these steps, you complete creating and sharing the dashboard. Subscribers of the dashboard receive email that contains a link to the dashboard. Groups don't receive invitation emails.

Copying an Amazon Quick Sight dashboard

If you have co-owner access or **Save as** privileges on an existing dashboard, you can copy it. To do this, create a new analysis from the dashboard and then create a new dashboard from the analysis that you copied.

After you save the original dashboard as a new analysis, you can collaborate on it by sharing the new analysis with other users. For example, you can use this workflow to preserve a production version of the dashboard, while also developing or testing a new version of it.

To copy a dashboard

1. Sign in to Quick Suite and choose **Dashboards** from the homepage.
2. Open the dashboard that you want to duplicate.
3. At upper right, choose **Save As**, and then enter a name for the new analysis. When you save an existing dashboard using **Save As**, it creates an analysis based on the dashboard.

Note

If you can't see **Save as**, check with your administrator that you have the right permissions.

4. (Optional) Make changes to the new analysis.
5. (Optional) Share the analysis with other users so you can collaborate on changes. All users who have access can make changes to the new analysis.

To share the analysis with other users, choose **Share** from the top right corner of the page, and then choose **Share analysis**.

6. (Optional) Create a new dashboard with your changes to the new analysis by choosing **Share**, and then choosing **Publish Dashboard**.

For more information, see the following:

- [Sharing Amazon Quick Sight dashboards](#)
- [Sharing Quick Sight analyses](#)

Deleting an Amazon Quick Sight dashboard

When you delete an Amazon Quick Sight dashboard, the dashboard is permanently removed from your account and all folders that the dashboard was a part of. You will no longer be able to access the deleted dashboard. You can only delete dashboards that you own or co-own. Use the following procedure to delete a dashboard.

To delete a dashboard

1. On the **Dashboards** tab of the Amazon Quick Suite homepage, choose the details icon (vertical dots :) on the dashboard that you want to delete.
2. Choose **Delete**. Then choose **Delete** again to confirm that you want to delete the dashboard. Deleting a dashboard permanently deletes the dashboard from your account, and the dashboard will disappear from all folders that it belonged to. You can still access and create other dashboards from the analysis that the deleted dashboard was published from.

Publishing a previous version of an Amazon Quick Sight dashboard

Each time you make updates to an analysis and publish it, a new version of the Amazon Quick Sight dashboard is created. To revert back to a previous version of a dashboard, you can search for it under the dashboard's **Version History** and publish the former version you are interested in. Each dashboard can store up to 1000 versions that are never deleted. Use the following procedure to publish a previous version of a dashboard.

To publish a previous version of a dashboard

1. On the **Dashboards** tab of the Amazon Quick Suite homepage, choose the dashboard that you want to manage.
2. Choose **Version History** on the toolbar on the right. The version of the dashboard that is currently published, as well as previous available versions, will appear in a list. Any comments that were added in the notes section will appear with the respective version.
3. Select the version of the dashboard you are interested in. You can see when this version was published and which user published it.
4. To revert to this version, select **Publish**. Click **Confirm** to publish the version.

Sharing Amazon Quick Sight dashboards

By default, dashboards in Amazon Quick Sight aren't shared with anyone and are only accessible to the owner. However, after you publish a dashboard, you can share it with other users or groups in your Amazon Quick Suite account. You can also choose to share the dashboard with everyone in your Quick Suite account and make the dashboard visible on the Quick Suite homepage for all users in your account. Additionally, you can copy a link to the dashboard to share with others who have access to it.

Important

Users who have access to the dashboard can also see the data used in the associated analysis.

After you share a dashboard, you can review the other users or groups that have access to it and control the type of access they have. You can revoke access to the dashboard for any user. You can also remove yourself from it.

You can also embed interactive dashboards and visuals in websites and apps by copying the dashboard or visual embed code and pasting it in your application. For more information, see [Embedding Amazon Quick Sight visuals and dashboards for registered users with a 1-click embed code](#).

Granting access to a dashboard

You can share dashboards and visuals with specific users or groups in your account or with everyone in your Amazon Quick Suite account. Or you share them with anyone on the internet. You can share dashboards and visuals by using the Quick Suite console or the Quick Sight API. Access to a shared visual depends on the sharing settings that are configured for the dashboard that the visual belongs to. To share and embed visuals to your website or application, adjust the sharing settings of the dashboard that it belongs to. For more information, see the following:

- [Granting individual Amazon Quick Sight users and groups access to a dashboard in Amazon Quick Sight](#)
- [Granting everyone in your Amazon Quick Sight account access to a dashboard](#)
- [Granting anyone on the internet access to an Amazon Quick Sight dashboard](#)
- [Granting everyone in your Amazon Quick Suite account access to a dashboard with the Quick Sight API](#)
- [Granting anyone on the internet access to an Amazon Quick Sight dashboard using the Quick Sight API](#)

Granting individual Amazon Quick Sight users and groups access to a dashboard in Amazon Quick Sight

Use the following procedure to grant access to a dashboard.

To grant users or groups access to a dashboard

1. Open the published dashboard and choose **Share** at upper right. Then choose **Share dashboard**.
2. In the **Share dashboard** page that opens, do the following:
 - a. For **Invite users and groups to dashboard** at left, enter a user email or group name in the search box.

Any users or groups that match your query appear in a list below the search box. Only active users and groups appear in the list.
 - b. For the user or group that you want to grant access to the dashboard, choose **Add**. Then choose the level of permissions that you want them to have.

You can select **Viewer** or **Co-owner**, depending on the user's Quick Suite role. The available permissions for each role are as follows:

- **Readers** – Quick Suite readers can only be granted **Viewer** access to dashboards. They can view, export, and print the dashboard, but they can't save the dashboard as an analysis. They can view, filter, and sort the dashboard data. They can also use any controls or custom actions that are on the dashboard. Any changes that they make to the dashboard exist only while they are viewing it, and aren't saved after they close the dashboard.
- **Authors** – Quick Suite authors can be granted **Viewer** or **Co-owner** access to dashboards.
 - Authors with Viewer access can view, export, and print the dashboard. They can view, filter, and sort the dashboard data. They can also use any controls or custom actions that are on the dashboard. Any changes that they make to the dashboard exist only while they are viewing it, and aren't saved after they close the dashboard.

However, they can save the dashboard as an analysis, unless the dashboard owner specifies otherwise. This privilege grants them read-only access to the datasets so that they can create new analyses from them. The owner has the option to provide them with the same permissions to the analysis. If the owner wants them also to edit and share the datasets, the owner can set that up inside the analysis.

- Authors with Co-owner access can view, export, and print the dashboard. They can also edit, share, and delete it. They can also save the dashboard as an analysis, unless the dashboard owner specifies otherwise. This privilege grants them read-only access to the datasets so that they can create new analyses from them. The owner has the option to provide them with the same permissions to the analysis. If the owner wants them to also edit and share the datasets, the owner can set that up inside the analysis.
- **Groups** – Quick Suite groups can only be granted **Viewer** access to dashboards. They can view, export, and print the dashboard, but they can't save the dashboard as an analysis.

After you add a user or group to the dashboard, you can see information about them in the **Manage permissions** section, under **Users & Groups**. You can see their user name, email, permission level, and "save as" privileges.

To allow a user or group to save the dashboard as an analysis, turn on **Allow "save as"** in the **Save as Analysis** column.

- c. To add more users to the dashboard, enter another user email or group name in the search box and repeat steps A and B.

Granting everyone in your Amazon Quick Sight account access to a dashboard

Alternatively, you can share your Amazon Quick Sight dashboard with everyone in your account. When you do this, everyone in your account can access the dashboard, even if they weren't granted access individually and assigned permissions. They can access the dashboard if they have a link to it (shared by you) or if it's embedded.

Sharing the dashboard with everyone in your account doesn't affect email reports. For example, suppose that you choose to share the dashboard with everyone in your account. Suppose also that you choose **Send email report to all users with access to dashboard** when setting up an email report for the same dashboard. In this case, the email report is sent only to people who have access to the dashboard. They receive access either through someone explicitly sharing it with them, through groups, or through shared folders.

To grant everyone in your account access to a dashboard

1. Open the published dashboard and choose **Share** at upper right. Then choose **Share dashboard**.
2. In the **Share dashboard** page that opens, for **Enable access for** at bottom left, toggle on **Everyone in this account**. Accounts that sign in with an Active Directory can't access the **Everyone in this account** switch. Accounts that use Active Directory can enable this setting with an `UpdateDashboardPermissions` API call. For more information on `UpdateDashboardPermissions`, see [UpdateDashboardPermissions](#) in the *Amazon Quick Sight API Reference*.
3. (Optional) Toggle on **Discoverable in Quick Sight**.

When you share a dashboard with everyone in the account, owners can also choose to make the dashboard discoverable in Quick Sight. A dashboard that's discoverable appears in everyone's list of dashboards on the **Dashboards** page. When this option is turned on, everyone in the account can see and search for the dashboard. When this option is turned off, they can only access the dashboard if they have a link or if it's embedded. The dashboard doesn't appear on the **Dashboards** page, and users can't search for it.

Granting anyone on the internet access to an Amazon Quick Sight dashboard

Applies to: Enterprise Edition

You can also share your Amazon Quick Sight dashboard with anyone on the internet from the **Share** menu in the Amazon Quick Suite console. When you do this, anyone on the internet will be able to access the dashboard, even if they aren't a registered user on your Quick Suite account, when you share the dashboard link or embed the dashboard.

Use the following sections to grant anyone on the internet access to dashboard when you share it.

Topics

- [Before you start](#)
- [Granting anyone on the internet access to a dashboard](#)
- [Updating a publicly shared dashboard](#)
- [Turning off public sharing settings](#)

Before you start

Before you can share a dashboard with anyone on the internet, make sure to do the following:

1. Turn on session capacity pricing on your account. If you have not turned on session capacity pricing on your account, you won't be able to update your account's public sharing settings.
2. Assign public sharing permissions to an administrative user in the IAM console. You can add these permissions with a new policy or you can add the new permissions to an existing user.

The following sample policy provides permissions for use with UpdatePublicSharingSettings.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "quicksight:UpdatePublicSharingSettings",
      "Resource": "*",
```



```
        "Effect": "Allow"
      }
    ]
  }
}
```

Accounts that don't want users with administrator access to use this feature can add an IAM policy that denies public sharing permissions. The following sample policy denies permissions for use with `UpdatePublicSharingSettings`.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "quicksight:UpdatePublicSharingSettings",
      "Resource": "*",
      "Effect": "Deny"
    }
  ]
}
```

For more information on using IAM with Quick Sight, see [Using Quick Suite with IAM](#).

You can also use the "Deny" policy as a Service Control Policy (SCP) if you don't want any of the accounts in your organization to have the public sharing feature. For more information, see [Service control policies \(SCPs\)](#) in the *Amazon Organizations User Guide*.

3. Turn on public sharing on your Amazon Quick Suite account.
 1. From the Amazon Quick Suite start page, choose your user icon at the upper right of your browser window, and then choose **Manage Quick Suite**.
 2. In the page that opens, scroll down to the **Permissions** section.
 3. Choose **Public access to dashboards** at left.
 4. On the page that opens, choose **Anyone on the internet**.

When you turn on this setting, a pop up will appear asking you to confirm your choice. Once you've confirmed your choice, you can grant the public access to specific dashboards and share those dashboards with them with a link or by embedding the dashboard in a public application, wiki, or portal.

Granting anyone on the internet access to a dashboard

To grant anyone on the internet access to a dashboard

1. In Quick Suite, open the published dashboard that you want to share. You must be the owner or a co-owner of the dashboard.
2. In the published dashboard, choose the **Share** icon at upper-right, and then choose **Share dashboard**.
3. In the **Share dashboard** page that opens, choose **Anyone on the internet (public)** in the **Enable access for** section at bottom-left.

This setting allows you to share the dashboard with anyone on the internet with the share link or when embedded. Turning on this switch also automatically turns on the **Everyone in this account** option, which means that the dashboard will be shared with anyone in your Quick Suite account. If you do not want this, turn off this option.

4. In the **Allow public access** pop-up that appears, enter confirm in the box to confirm your choice, and then choose **Confirm**.

After you confirm your dashboard's access settings, an orange **PUBLIC** tag appears at upper right of your dashboard in the Amazon Quick Suite console. Additionally, an eye icon appears on the dashboard on the Quick Sight Dashboards page, both in tile and list view.

Note that when public access is turned on, the dashboard can only be accessed using the link or when embedded using the embed code. For more information about sharing a link to the dashboard, see [Sharing a link a shared dashboard](#). For more information about embedding dashboards for anyone on the internet, see [Embedding Amazon Quick Sight visuals and dashboards for anonymous users with a 1-click embed code](#).

Updating a publicly shared dashboard

Use the following procedure to update a shared dashboard that can be accessed by anyone on the internet.

To update a public dashboard:

1. From the Amazon Quick Suite start page, choose the analysis that is tied to the dashboard that you want to update and make your desired changes. You must be the owner or a co-owner of the analysis.

2. In the analysis, choose **Publish**.
3. In the pop-up that appears, choose **Replace an existing dashboard** and select the public dashboard that you want to update.
4. To confirm your choice, enter `confirm` and then choose **Publish dashboard**.

Once you choose **Publish dashboard**, your public dashboard is updated to reflect the new changes.

Turning off public sharing settings

You can turn off public sharing settings for dashboards at anytime. You can turn off public sharing for an individual dashboard, or for all dashboards in your account. Visual sharing settings are determined at the dashboard level. If you turn off public sharing settings to a dashboard that holds a visual that you are embedding, users won't be able to access the visual.

The following table describes the different scenarios for when a dashboard is publicly available.

Account-level public setting	Dashboard-level public setting	Public access	Visual indicators
Off	Off	Off	None
On	Off	Off	None
On	On	Yes	An orange badge appears on the dashboard and an eye icon appears on the dashboard in the Dashboards page.
Off	On	No	A grey badge appears on the dashboard and an eye icon wi

Account-level public setting	Dashboard-level public setting	Public access	Visual indicators
			th a slash appears on the dashboard in the Dashboards page. It can take up to two minutes for a dashboard's public access to be revoked.

To turn off public sharing for a single dashboard

1. In Amazon Quick Suite, open the published dashboard that you no longer want to share. You must be the owner or a co-owner of the dashboard.
2. In the published dashboard, choose the **Share** icon at upper-right, and then choose **Share dashboard**.
3. In the **Share dashboard** page that opens, toggle off the **Anyone on the internet (public)** switch in the **Enable access for** section at bottom-left.

This action will remove public access to the dashboard. It will now only be accessible to users that it has been shared with.

To turn off public sharing settings for all dashboards in a Quick Suite user account

1. From the Amazon Quick Suite start page, choose your user icon at upper right of your browser window, and then choose **Manage Quick Suite**.
2. In the page that opens, scroll down to the **Permissions** section.
3. Choose **Public access to dashboards** at left.
4. On the page that opens, toggle off the **Anyone on the internet** switch.

When you disable public sharing settings from the **Public sharing** menu, a pop-up will appear asking you to confirm your choice. Select **I have read and acknowledge this change** and then choose **Confirm** to confirm your choice.

This action will remove public access to every dashboard on your account. Dashboards that were visible to anyone on the internet will now only be accessible to users that each dashboard has been shared with. Individual dashboards that have their public settings turned on will have a gray badge and the eye icon that appears on the **Dashboards** page will have a strike through it to indicate that the account level public settings are disabled and that the dashboard can't be viewed. It can take up to two minutes for a dashboard's public access to be revoked.

If your session capacity pricing subscription has expired, public sharing settings will be automatically removed across your account. Renew your subscription to restore access to public sharing settings.

Granting everyone in your Amazon Quick Suite account access to a dashboard with the Quick Sight API

Intended audience: Amazon Quick Suite developers

Alternatively, you can grant everyone in your account access to the dashboard with the Quick Sight API using the `UpdateDashboardPermissions` operation.

The following example API request illustrates how to do so using an Amazon CLI command. It grants link permissions on the dashboard in your account, and allows the following operations: `DescribeDashboard`, `QueryDashboard` and `ListDashboard`.

```
aws quicksight update-dashboard-permissions \  
--aws-account-id account-id \  
--region aws-directory-region \  
--dashboard-id dashboard-id \  
--grant-link-permissions  
Principal="arn:aws-cn:quicksight:aws-directory-region:account-id:namespace/default",  
Actions="quicksight:DescribeDashboard, quicksight:QueryDashboard,  
quicksight:ListDashboardVersions"
```

The response for the preceding request looks similar to the following.

```
{  
  "Status": 200,  
  "DashboardArn": "arn:aws-cn:quicksight:AWSDIRECTORYREGION:ACCOUNTID:dashboard/  
DASHBOARDID",  
}
```

```

"DashboardId": "DASHBOARDID",
"LinkSharingConfiguration": {
  "Permissions": [
    {
      "Actions": [
        "quicksight:DescribeDashboard",
        "quicksight:ListDashboardVersions",
        "quicksight:QueryDashboard"
      ],
      "Principal": "arn:aws-cn:quicksight:AWSDIRECTORYREGION:ACCOUNTID:namespace/default"
    }
  ],
  "Permissions": [
    // other dashboard permissions here
  ],
  "RequestId": "REQUESTID"
}

```

You can also prevent all users in your account from accessing the dashboard using the same API operation. The following example request illustrates how by using a CLI command.

```

aws quicksight update-dashboard-permissions \
--aws-account-id account-id \
--region aws-directory-region \
--dashboard-id dashboard-id \
--revoke-link-permissions
Principal="arn:aws-cn:quicksight:aws-directory-region:account-id:namespace/default",
Actions="quicksight:DescribeDashboard, quicksight:QueryDashboard,
quicksight:ListDashboardVersions"

```

For more information, see [UpdateDashboardPermissions](#) in the *Amazon Quick Suite API Reference*.

When all users in a Quick Suite user account are granted access to the dashboard, the following snippet is added to Amazon CloudTrail log as part of the eventName `UpdateDashboardAccess`, and the eventCategory `Management`.

```

"linkPermissionPolicies":
[
{
  "principal": "arn:aws-cn:quicksight:AWSDIRECTORYREGION:ACCOUNTID:

```

```

        namespace/default",
    "actions":
    [
        "quicksight:DescribeDashboard",
        "quicksight:ListDashboardVersions",
        "quicksight:QueryDashboard"
    ]
}
]
```

Granting anyone on the internet access to an Amazon Quick Sight dashboard using the Quick Sight API

Alternatively, you can grant anyone on the internet access to the dashboard with the Amazon Quick Sight API using the UpdateDashboardPermissions operation.

Before you begin, make sure to grant everyone in your account access to the dashboard. For more information, see [Granting everyone in your Amazon Quick Suite account access to a dashboard with the Quick Sight API](#).

The following example API request illustrates how to grant anyone on the internet access to a dashboard using an Amazon CLI command. It grants link permissions on the dashboard in your account, and allows the following operations: DescribeDashboard, QueryDashboard and ListDashboardVersions.

```

aws quicksight update-dashboard-permissions
--aws-account-id account-id
--region aws-directory-region
--dashboard-id dashboard-id
--grant-link-permissions
Principal="arn:aws-cn:quicksight::publicAnonymousUser/*",
Actions="quicksight:DescribeDashboard, quicksight:QueryDashboard,
quicksight:ListDashboardVersions"
```

The response for the preceding request looks similar to the following.

```

{
    "Status": 200,
    "DashboardArn": "arn:aws-cn:quicksight:AWSDIRECTORYREGION:ACCOUNTID:dashboard/
DASHBOARDID",
    "DashboardId": "DASHBOARDID",
    "LinkSharingConfiguration": {
```

```

    "Permissions": [
      {
        "Actions": [
          "quicksight:DescribeDashboard",
          "quicksight:ListDashboardVersions",
          "quicksight:QueryDashboard"
        ],
        "Principal": "arn:aws-
cn:quicksight:AWSDIRECTORYREGION:ACCOUNTID:namespace/default"
      },
      {
        "Principal": "arn:aws-cn:quicksight::publicAnonymousUser/*",
        "Actions": [
          "quicksight:DescribeDashboard",
          "quicksight:ListDashboardVersions",
          "quicksight:QueryDashboard"
        ]
      }
    ],
    "Permissions": [
      // other dashboard permissions here
    ],
    "RequestId": "REQUESTID"
  }

```

You can also prevent anyone on the internet from accessing the dashboard using the same API operation. The following example request illustrates how by using a CLI command.

```

aws quicksight update-dashboard-permissions \
--aws-account-id account-id \
--region aws-directory-region \
--dashboard-id dashboard-id \
--revoke-link-permissions
Principal="arn:aws-cn:quicksight::publicAnonymousUser/*",
Actions="quicksight:DescribeDashboard, quicksight:QueryDashboard,
quicksight:ListDashboardVersions"

```

For more information, see [UpdateDashboardPermissions](#) in the *Amazon Quick Suite API Reference*.

When anyone on the internet is granted access to the dashboard, the following snippet is added to Amazon CloudTrail log as part of the eventName UpdateDashboardAccess, and the eventCategory Management.


```
"linkPermissionPolicies":
[
{
  "principal": "arn:aws-cn:quicksight::publicAnonymousUser/*",
  "actions":
  [
    "quicksight:DescribeDashboard",
    "quicksight:ListDashboardVersions",
    "quicksight:QueryDashboard"
  ]
}
]
```

Sharing a link a shared dashboard

After you grant users access to a dashboard, you can copy a link to it and send it to them. Anyone with access to the dashboard can access the link and see the dashboard.

To send users a link to the dashboard

1. Open the published dashboard and choose **Share** at upper right. Then choose **Share dashboard**.
2. In the **Share dashboard** page that opens, choose **Copy link** at upper left.

The link to the dashboard is copied to your clipboard. It's similar to the following,

```
https://quicksight.aws.amazon.com/sn/accounts/accountid/
dashboards/dashboardid?directory_alias=account_directory_alias
```

Users and groups (or all users on your Quick Suite account) who have access to this dashboard can access it by using the link. If they are accessing Quick Suite for the first time, they will be asked to sign in with their email address or Quick Suite user name and password for the account. After they sign in, they will have access to the dashboard.

View who has access to a shared dashboard

Use the following procedure to see which users or groups have access to the dashboard.

1. Open the published dashboard and choose **Share** at upper right. Then choose **Share dashboard**.

2. In the **Share dashboard** page that opens, under **Manage permissions**, review the users and groups, and their roles and settings.

You can search to locate a specific user or group by entering their name or any part of their name in the search box at upper right. Searching is case-sensitive, and wildcards aren't supported. Delete the search term to return the view to all users.

Revoke access to a shared dashboard

Use the following procedure to revoke user access to a dashboard.

To revoke user access to a dashboard

1. Open the dashboard and choose **Share** at top right. Then choose **Share dashboard**.
2. In the **Share dashboard** page that opens, under **Manage permissions**, locate the user that you want to remove and choose the delete icon at far right.

Using Quick action connectors in dashboard visuals

Prerequisites

Before you begin, make sure to [create at least one action connector](#).

The connector must meet these requirements:

- Uses the **User Auth** authentication method
- Uses one of the following integrations:
 - Atlassian Jira Cloud
 - Microsoft Outlook
 - Microsoft Teams
 - Salesforce
 - ServiceNow
 - Slack

Enable Quick actions on a dashboard to use action connectors

To enable Quick actions on a dashboard to use action connectors

1. If a dashboard exists, go to the source analysis of the dashboard. Otherwise, [create a new analysis](#).
2. Choose **Publish**.
3. Choose between **New Dashboard** or **Replace existing dashboard**.
4. Choose the **Enable Quick actions** checkbox under **Dashboard options**.
5. Choose **Publish dashboard**.

Use action connectors on a visual

To use action connectors on a visual

1. Open a dashboard with the **Enable Quick actions** publishing option turned on.
2. Hover over a visual.
3. Choose the lightening bolt icon.
4. A menu appears with a list of all supported action connectors and actions.
5. Choose the desired action from the list.
6. If you have not used the connector before, or if your previous login credentials have expired, an authentication modal will appear. Log in with appropriate credentials for your organization.
7. An **Action** form appears in the right pane.
8. Enter all the information you need to include with the action.
9. Some fields allow the inclusion of autofill values. Choose **Autofill** to open the menu. Choose the values you need and they will be added to your entered text.
 - **Today's date**: Injects today's date
 - **Visual name**: Injects visual name
 - **All**: Injects both of the above
10. Some actions support the ability to include an attachment. You can optionally attach an image of the visual with these actions by selecting the **Visual image** checkbox.
11. Select the action button at the bottom of the form to invoke the action.

Security and customizations

Custom Permissions/Capability Customization

- **Actions** capability: You cannot see or use actions if your user or role is restricted the permission to use the Actions capability

To learn more about custom permissions, see [Creating a custom permissions profile in Amazon Quick Suite](#).

Row Level Security (RLS) / Column Level Security (CLS)

- You cannot see or use actions on visuals that are based on datasets that use RLS or CLS.

To learn more about RLS, see [Using row-level security in Amazon Quick Suite](#).

To learn more about CLS, see [Using column-level security to restrict access to a dataset](#).

Dashboard publishing options

- Enable Quick actions
 - You cannot see or use actions on any visuals of a dashboard that was published with the **Enable Quick actions** publishing option disabled.

To learn more about dashboard publishing options, see [Publishing dashboards](#).

Limitations

Visual Image attachment support

The following visual types do not support image attachments:

- High charts (when HTML is used)
- ML Insights (when HTML is used)
- Textbox and insights (when HTML is used)
- Custom content

 Note

For these visuals, the **Visual image** checkbox will not appear on the UI.

Sharing your view of a Amazon Quick Sight dashboard

While interacting with a published dashboard, you can choose to share a unique link to the dashboard with only your changes. For example, if you filter the data in the dashboard, you can share what you see with others who have permissions to see the dashboard. That way, they can see what you see, without your having to create a new dashboard.

When others access your view of the dashboard by using the link that you sent them, they see the dashboard exactly as it was when the link was created. They see any parameters, filters, or controls that you changed.

To share your view of a dashboard

1. Open the published dashboard, and make any changes that you want.
2. Choose **Share** at upper right, and then choose **Share this view**.
3. On the **Share using a link** page that opens, choose **Copy link**.
4. Paste the link in an email or IM message to share it with others.

Only people with permissions to see the dashboard in Quick Sight can access the link.

Scheduling and sending Quick Sight reports by email

 Important

Amazon Quick Sight in the China (Beijing) (cn-north-1) uses an internal email service (Amazon SES) in the China (Ningxia) (cn-northwest-1) to send emails to Quick Sight users. Customer data that's included in scheduled reports, alerts, and other features are passed by email from China (Beijing) to ZHY before it reaches Quick Sight users.

The following Quick Sight features send emails through Amazon SES in the China (Ningxia) region. If you don't want to pass data across these regions, avoid these features. There is no centralized way to disable these features.

- Anomaly detection alerts

- Emails that notify users that a failure has occurred
- Emails that use threshold alerts
- Scheduled reports

In Enterprise edition, you can send a dashboard in report form either once or on a schedule (daily, weekly, monthly, or yearly). You can email the reports to users or groups who share your Amazon Quick Suite subscription. To receive email reports, the users or group members must meet the following conditions:

- They are part of your Quick Suite subscription.
- You already shared the dashboard with them.
- Amazon Quick Sight can't send scheduled emails to more than 5,000 members.

Amazon Quick Sight generates a custom email snapshot for each user or group based on their data permissions, which are defined in the dashboard. Row Level Security (RLS), Column Level Security (CLS) and Dynamic Default Parameters for email reports works for both scheduled and ad hoc (one-time) emails.

Quick Suite authors can run scheduled reports with the **Report now** button in the Quick Suite console or with the [StartDashboardSnapshotJobSchedule](#) API.

Subscribers who are readers see an option for **Reports** on the dashboard when an email report is available for that dashboard. They can use the **Schedules** menu to subscribe to or unsubscribe from the emails. For more information, see [Subscribing to email reports in Amazon Quick Sight](#).

You can create up to five schedules for each dashboard.

Quick Sight dashboard viewers can also schedule their own reports for themselves from a Quick Sight dashboard. For more information about reader generated reports, see [Creating a reader generated report in Amazon Quick Sight](#).

Use the following topics to learn more about email report settings and report billing.

Topics

- [Configuring email report settings for a Quick Sight dashboard](#)
- [How billing works for email reports](#)

Configuring email report settings for a Quick Sight dashboard

Applies to: Enterprise Edition

In Amazon Quick Suite Enterprise edition, you can email a report from any sheet in a dashboard. You can send reports from interactive dashboards and paginated report sheets. Schedules include settings for when to send them, the contents to include, and who receives the email. You can view a sample report and a list of the datasets used in the report. To set up or change the schedule sent from a dashboard, make sure that you're an owner or co-owner of the dashboard.

If you have access to the dashboard, you can change your subscription options by opening your view of the dashboard. For more information on how this works, see [Subscribing to email reports in Amazon Quick Sight](#).

Scheduling options that are available for an email report include the following:

- **Once (Does not repeat)** – Sends the report only once at the date and time that you choose.
- **Daily** – Repeats daily at the time that you choose.
- **Weekly** – Repeats each week on the same day or days at the time that you choose. You can also use this option to send reports in weekly intervals, such as every other week or every three weeks.
- **Monthly** – Repeats each month on the same day of the month at the time that you choose. You can also use this option to send reports on specific days of the month, such as the second Wednesday or the last Friday of each month.
- **Yearly** – Repeats each year on the same day of the month or months selected at the time that you choose. You can also use this option to send reports on specific days or sets of days in selected months. For example, you can configure a report to be sent on the first Monday of January, March, and September, or on July 14th, or on the second day of February, April, and June each year.
- **Custom** – Configure your own scheduled report that best fits your business needs.

You can customize the title of the report, the optional email subject, and the body text.

Although you can configure the report so that everyone who has access receives a copy, this is not usually the best plan. We recommend limiting automated emails, especially those sent to groups.

You can start with a small number of subscribers by choosing specific people from the access list. Verify your company's policy before subscribing anyone to a subscription.

You can directly add people to a report subscription in these ways:

- (Recommended) Choose recipients from the provided access list to specify and maintain a list of people who you want to email reports to. You can use the search box to find people by email or group name.
- To send reports to all of the dashboard's subscribers, choose **Send email report to all users with access to dashboard** when prompted.

Anyone else who wants to get the emails can open the dashboard and set their own subscription options to either opt in or opt out.

Important

When you share the dashboard with new Quick Suite user names or groups, they automatically start receiving the email reports. If you don't want this to happen, you need to edit the report settings each time you add people to the dashboard.

For existing email schedules, you can pause the schedule in Amazon Quick Sight while you make changes. In the **Schedules** pane, you can pause or resume a scheduled report with the toggle that appears under each report. Pausing a report does not delete the report's schedule from Quick Sight.

If your report includes custom visuals, be aware that you can't include images from a private network in an email report, even if you can access the images. If you want to include an image, use a publicly available one.

Before you begin, make sure that you are using Amazon Quick Suite Enterprise edition and that you have shared the dashboard with intended recipients.

To create or change an email report

1. Open Quick Suite and choose **Dashboards** on the navigation pane at left.
2. Open a dashboard to configure its email report.
3. At top right, choose **Schedules**, and then choose **Schedules**.

4. Choose **ADD SCHEDULE**.
5. In the **New schedule** pane that appears, enter the schedule name. Optionally, add a description for the new schedule.
6. In the **Content** tab, toggle the **PDF**, **CSV**, or **Excel** switches to choose the report format. CSV and Excel format are currently supported for paginated reports.
7. In the **Sheet** dropdown on the **Content** tab, choose the sheet that you want to schedule a report for.

If you choose **CSV** or **Excel**, choose the table or pivot table visuals from any sheet of the dashboard that you want to include in the report. You can select up to 5 visuals for each schedule.

If you choose **Excel**, one Excel workbook is generated as a final output.

8. In the **Dates** tab, choose the frequency for the report in the **Repeat** dropdown. If you're not sure, choose **Send once (Does not repeat)**.
9. For **Start date**, choose the start date and runtime that you want to send the first report on.
10. For **Timezone**, choose the time zone from the dropdown.
11. In the **Email** tab, for **E-mail subject line**, enter a custom subject line, or leave it blank to use the report title.
12. Enter the email addresses of the Quick Suite group name of the users or groups that you want to receive the report. You can also select the **Send to all users with access** box to send the report to everyone that has access to the dashboard in your account.
13. For **Email header**, enter the header that you want the email report to show.
14. (Optional) For **E-mail body text**, leave it blank or enter a custom message to display at the beginning of the email.
15. (Optional) For PDF attachments, you can choose **Include sheet in email body** to show the first page of the PDF snapshot in the email's body.
16. Choose the method of attachment that you want the report to use. The following options are available.
 - **File attachment**– uploads an attachment of the snapshot to the email. The email size can't exceed 10 MB. This limit includes all attachments.
 - **Download link**– adds a link to the email body that users can access to download the snapshot report. When a user chooses the download link, they are prompted to sign in before the report starts to download. The link expires one year after the report is sent.

17. (Optional, recommended) To send a sample of the report before you save changes, choose **Send test report**. This option displays beside the user name of the owner of the dashboard.
18. Do one of the following:
 - (Recommended) Choose **Save** to confirm your entries.
 - To immediately send a report, choose **Save and run now**. The report is sent immediately, even if your schedule's start date is in the future.

How billing works for email reports

Authors and admins can receive any number of email reports at no extra charge.

For readers (users in the reader role), it costs one session per report, up to the monthly maximum. After receiving an email report, the reader gets a session credit to access the dashboard at no additional cost during the same month. Reader session credits don't carry over to the next billing month.

For a reader, charges for email reports and interactive sessions both accrue up to the monthly maximum charge. For readers who hit the monthly max charge, there are no further charges, and they can receive as many additional email reports as they need.

Subscribing to email reports in Amazon Quick Sight

In Enterprise edition, Amazon Quick Suite authors can set up subscriptions to a dashboard in report form. For more information, see [Scheduling and sending Quick Sight reports by email](#). Quick Suite readers and authors can then subscribe to a dashboard and adjust their report settings. For more information about subscribing to dashboards as a reader, see [Subscribing to Amazon Quick Sight dashboard emails and alerts](#).

Use the following procedure to change your subscription and report settings for a specific dashboard.

1. First, open a dashboard that is shared with you, or a dashboard that you own or co-own.
2. Choose the **Reports** icon at top right.
3. The **Change report preferences** screen appears. This screen shows the current report schedule, in addition to the subscription and optimization options.

For **Subscription**, choose **Subscribe** to start receiving reports, or **Unsubscribe** to stop receiving reports.

Under **Optimize**, choose the device you prefer to view the report on.

- If you usually use a mobile device or you prefer to view reports in a portrait format, choose **Viewing on a mobile device**. When you receive the report, the visuals display in a single vertical column.
- If you usually use a desktop or you prefer to view reports in a landscape format, choose **Viewing on a desktop**. When you receive the report, the visuals display in the same layout shown in your dashboard on your desktop.

4. Choose **Update** to confirm your choices, or choose **Cancel** to discard your changes.

Working with threshold alerts in Amazon Quick Sight

Applies to: Enterprise Edition

To stay informed about important changes in your data, you can create threshold alerts using KPI, Gauge, Table, and Pivot table visuals in an Amazon Quick Sight dashboard. With these alerts, you can set thresholds for your data and be notified by email when your data crosses them. You can also view and manage your alerts at anytime in a Quick Sight supported web browser.

For example, let's say that you're a customer success manager for a large organization and you want to know when the number of tickets in a support queue exceeds a certain number. Let's say too that you have a dashboard with a KPI, Gauge, Table or Pivot table visual that tracks the number of tickets in this queue. In this case, you can create an alert and be notified by email when the number exceeds the threshold you specified. That way, you can take action as soon as you're notified.

You can create multiple alerts for a single visual. If the visual is updated or deleted by the author after you create an alert, your alert settings don't change. When you create an alert, the alert takes on any filters applied to the visual at that time. If you or the author changes the filter, your existing alert doesn't change. However, if you create a new alert, your new alert takes on the new filter settings.

For example, let's say you have a dashboard with a filter control that you can use to switch the data for each visual in the dashboard from one US city to another. You have a KPI visual on the dashboard that shows average flight delays, and you're interested in delays for flights leaving

from Seattle, Washington, in the US. You change the filter control to Seattle and set an alert on the visual. This alert tracks flight delays from Seattle. Tomorrow, let's say that you want to also track flight delays from Portland, Oregon, so you change the filter control to Portland and create another alert. This new alert tracks flight delays from Portland. You now have two alerts, one on Seattle and one on Portland, working independently.

Threshold alerts are not available in the eu-central-2 Europe (Zurich) region.

For more information on KPI, Gauge, Table, or Pivot table visuals, see [Visual types in Amazon Quick Sight](#).

Note

You can't create alerts for visuals in an embedded dashboard or from the Quick Suite mobile app.

For table visuals, threshold alerts can't be created for values that are located in the Group by field well. Alerts can only be created for values that are located in the Value field well. KPI visuals that don't use a date-time field as a trend don't support alerts. An example is a KPI that shows the difference in flights between carriers X and Y instead of a KPI that shows the difference in flights between dates A and B.

Use the sections below to create and configure threshold alerts for KPI, Gauge, Table, and Pivot table visuals in Quick Sight.

Topics

- [Alert Permissions](#)
- [Creating Alerts](#)
- [Managing Threshold Alerts](#)
- [Investigating Alert Failures](#)
- [Alert Scheduling](#)
- [Using Quick action connectors in threshold alerts](#)

Alert Permissions

If you're an administrator, you can control who in your organization can set threshold alerts in Quick Sight by creating a custom permissions policy. To set custom permissions in Quick Suite,

choose your user name at the upper-right corner of any Quick Suite page, choose **Manage Quick Suite**, and then choose **Custom permissions**.

Creating Alerts

Use the following procedure to create threshold alerts for KPI or Gauge visuals in a dashboard.

To create an alert

1. Open Quick Suite and navigate to the dashboard that you want.

For more information about viewing dashboards as a dashboard subscriber in Quick Suite, see [Interacting with Amazon Quick Sight dashboards](#).

2. In the dashboard, select the visual that you want to create an alert for, open the menu at the upper-right, and choose **Create alert**.
3. On the menu at upper-right on the visual, choose the **Create alert** icon.

Alternatively, you can choose the alert icon in the blue toolbar at upper right. Then, in the **Create alert** page that opens, select the KPI, Gauge, Table or Pivot table visual that you want to create an alert for, and then choose **Next**.

You can also create alerts on table or pivot table visuals by selecting a cell and choosing **Create alert**. You can only create alerts for single cells. Alerts can't be created for entire columns or for values that use a custom aggregation. For more information about custom aggregations, see [Aggregate functions](#).

4. On the **Create alert** page that opens at right, do the following:

- a. For **Name**, enter a name for the alert.

By default, the visual name is used for the alert name. You can change it if you want.

- b. For **Value to track**, choose a value that you want to set the threshold for. The information presented will vary based on the visual type you're creating an alert for.

The values that are available for this option are based on the values the dashboard author sets in the visual. For example, let's say you have a KPI visual that shows a percent difference between two dates. Given that, you see two alert value options: percent difference and actual.

If there is only one value in the visual, you can't change this option. It is the current value and it is displayed here so that you can use it as a reference while you choose a threshold. For example, if you're setting an alert on average cost, this value will show you what the current average cost is (say, \$5). With this reference value you can make more informed decisions while setting your threshold.

- c. For **Condition**, choose a condition for the threshold.

You can choose the following conditions.

- **Is above** – Sets a rule that the alert triggers if the alert value goes above the threshold you set.
- **Is below** – Sets a rule that the alert triggers if the alert value goes below the threshold that you set.
- **Is equal to** – Sets a rule that the alert triggers if the alert value is equal to the threshold you set.

- d. For **Threshold**, enter a value to prompt the alert.
- e. For **Notification preference**, choose how often you want to be notified about a breach to the threshold you set.

You can choose from the following options.

- **As frequently as possible** - This option alerts you whenever the threshold is breached. If you choose this option, you might get alerts multiple times a day.
 - **Daily at most** - This option alerts you once per day when the threshold is breached.
 - **Weekly at most** - This option alerts you once per week when the threshold is breached.
- f. (Optional) Choose **Email me when there is no data** - When you select this option, you're notified when there's no data to check your alert rule against.
- g. Choose **Save**.

A message at upper-right appears indicating that the alert has been saved. If your data crosses the threshold you set, you get a notification by email at the address that's associated with your Quick Suite account.

Managing Threshold Alerts

You can edit your existing alerts, turn them on or off, or view the history of times when the alert was triggered. Use the following procedures to do so.

To edit an existing alert

1. Open Quick Suite, choose **Dashboards**, and then navigate to the dashboard that you want to edit an alert for.
2. On the Dashboards page, choose **Alerts** at upper-right.
3. On the **Manage alerts** page that opens, find the alert that you want to edit, and then choose **Edit** beneath the alert name.

You can edit the alert name, condition, and threshold.

4. Choose **Save**.

To view the history of when an alert was triggered

1. Open Quick Suite, choose **Dashboards**, and then navigate to the dashboard that you want to view alert history for.
2. On the Dashboards page, choose **Alerts** at upper-right.
3. On the **Manage alerts** page that opens, find the alert that you want to view the history for, and then choose **History** beneath the alert name.

To turn on or turn off an existing alert

1. Open Quick Suite, choose **Dashboards**, and navigate to the dashboard that you want to turn on or turn off an alert for.
2. On the Dashboards page, choose **Alerts** at upper-right.
3. On the **Manage alerts** page that opens, find the alert that you want to turn on or off, and then select or clear the toggle by the alert name.

The alert is turned on when the toggle is blue, and turned off when the toggle is gray.

To delete an existing alert

1. Open Quick Suite, choose **Dashboards**, and navigate to the dashboard that you want to delete an alert from.
2. On the Dashboards page, choose **Alerts** at upper-right.
3. On the **Manage alerts** page that opens, find the alert that you want to turn on or off, choose the three-dot menu next to the alert, and then choose **Delete** from the dropdown.

Investigating Alert Failures

When an alert fails, Quick Suite sends you an email notification about the failure. Alerts can fail for many reasons, including the following:

- The dataset the alert is using was deleted.
- The owner of the alert lost permissions to the dataset or to certain rows or columns in the dataset.
- The owner of the alert lost access to the dashboard.
- There is no data for the data tracked by the alert.

When a failure occurs, Quick Suite sends you a notification and disables the alert if the reason for the failure isn't likely to be fixed. For example, if the alert fails due to the loss of access to a dashboard, or if the dashboard was deleted. Otherwise, Quick Suite attempts to check your data for threshold breaches again. After four failures, Quick Suite turns off the alert and notifies you that the alert is turned off. If the alert can be checked again, Quick Suite sends you a notification.

To investigate why an alert failed, check that you still have access to the dashboard. Also check that you have permissions to the correct dataset and to the correct rows and columns in the dataset. If you have lost access or permissions, contact the dashboard owner. If you have the necessary access and permissions, you might need to edit your alert to avoid future alert failures.

Alert Scheduling

When you create an alert, Quick Suite checks your data for any breaches against the thresholds you set based on when your dataset is scheduled to refresh. The information presented in the alert varies based on the visual type that you're creating an alert for. For SPICE datasets, alert rules are checked after a successful refresh of your SPICE dataset. For direct query datasets, alert rules are

checked at a random time between 6:00 PM and 8:00 AM in the Amazon Web Services Region that holds the dataset by default.

If you're a dataset owner, you can set an alert evaluation schedule in the dataset settings. See the following procedure to learn how.

To set an alert evaluation schedule for a dataset

1. In Quick Suite, choose **Data** in the navigation bar at left.
2. Choose the dataset that you want to schedule alert evaluations for.
3. Choose **Set alert schedule**.
4. In the **Set alert schedule** page that opens, do the following.
 - For **Time zone**, choose a time zone.
 - For **Repeats**, choose how often you want the data to be evaluated.
 - For **Starts**, enter the time that you want the alert evaluation to start.

Using Quick action connectors in threshold alerts

Prerequisites

Before you begin, make sure to [create at least one action connector](#).

The connector must meet these requirements:

- Uses the **Service Auth** authentication method
- Uses one of the following integrations:
 - Atlassian Jira Cloud
 - Microsoft Outlook
 - Salesforce
 - ServiceNow

Enable Quick actions on a dashboard to use action connectors

To enable Quick actions on a dashboard to use action connectors

1. If a dashboard exists, go to the source analysis of the dashboard. Otherwise, [create a new analysis](#).

2. Choose **Publish**.
3. Choose between **New Dashboard** or **Replace existing dashboard**.
4. Choose the **Enable Quick actions** checkbox under **Dashboard options**.
5. Choose **Publish dashboard**.

Use action connectors in a threshold alert

To use action connectors in a threshold alert

1. Open a dashboard with the **Enable Quick actions** publishing option turned on.
2. Hover over a visual that supports threshold alerts. The types of visuals that support alerts can be found [here](#).
3. Choose the bell icon.
4. The **Create alert** pane opens at right.
5. Choose **Add Action**.
6. A menu appears with a list of all supported action connectors and actions.
7. Choose the desired action from the list.
8. An **Action** form appears in the right pane.
9. Enter all the information you need to include with the action.
10. Some fields allow the inclusion of autofill values. Choose **Autofill** to open the menu. Choose the values you need and they will be added to your entered text.
 - **Value**: Injects the current value that was used by the alert to evaluate the alert condition
 - **Alert name**: Injects the alert name
 - **Condition**: Injects the alert condition
 - **Threshold**: Injects the threshold value
 - **All**: Injects all of the above
11. Some actions support the ability to include an attachment. You can optionally attach a PDF of the current dashboard sheet with these actions by selecting the **Include this sheet as PDF** checkbox.
12. Choose **Add action** to add the action to the alert.
13. Back at the **Create alert** pane, the configured action is added to the alert at the bottom.

14. Configure all the other desired fields of the alert and choose **Save**.
15. When your configured threshold is breached, this action should get invoked. To learn more about when a threshold alert is evaluated, see [Alert Scheduling](#).

Security and customizations

Custom Permissions/Capability Customization

- **Actions** capability: You cannot see or use actions if your user or role is restricted the permission to use the Actions capability
- **Export To PDF** capability:
 - New actions on alerts: You will not see the option to attach a PDF of the sheet while adding a new action to an alert if your user or role is restricted the permission to use the **Export To PDF** capability.
 - Existing actions on alerts: If you have existing alerts with actions containing PDF attachments then those actions will be sent out without the PDF attachments when your user or role gets restricted from using the **Export To PDF** capability.

To learn more about custom permissions, see [Creating a custom permissions profile in Amazon Quick Suite](#).

Row Level Security (RLS) / Column Level Security (CLS)

- New actions on alerts: If your dashboard contains a dataset with RLS or CLS, then
 - You cannot add actions to new alerts that track the dataset with RLS or CLS
 - You can add actions to new alerts that track a different dataset without RLS or CLS, but you cannot include PDF attachments in these actions
- Existing actions on alerts: If you add RLS or CLS to a dataset after creating alerts with actions, then
 - Existing actions on alerts tracking that dataset will stop working completely
 - Existing actions on alerts tracking a different dataset on the same dashboard will be sent out without any PDF attachments

To learn more about RLS, see [Using row-level security in Amazon Quick Suite](#).

To learn more about CLS, see [Using column-level security to restrict access to a dataset](#).

Dashboard publishing options

- Enable PDF generation for interactive sheets
 - New actions on alerts: You will not see the option to attach a PDF of the sheet while adding a new action on alert if your dashboard has the **Enable PDF generation for interactive sheets** publishing option disabled.
 - Existing actions on alerts: If you have existing alerts with actions containing PDF attachments, then those actions will be sent out without the PDF attachments when the **Enable PDF generation for interactive sheets** publishing option gets disabled on your dashboard.
- Enable Quick actions
 - New actions on alerts: You will not see the option to add an action to the alert if your dashboard has the **Enable Quick actions** publishing option turned off.
 - Existing actions on alerts: Your existing actions on alerts will stop working completely when the **Enable Quick actions** publishing option gets disabled on your dashboard.

To learn more about dashboard publishing options, see [Publishing dashboards](#).

Printing a dashboard or analysis

You can print a dashboard or an analysis in Amazon Quick Sight.

Use the following procedure to print.

1. Open the dashboard or the analysis that you want to print.
2. Choose the **Print** icon at top right.
3. On the **Prepare for printing** screen, choose the paper size and orientation that you want to use.
4. Choose **Go to Preview**.
5. Do one of the following:
 - To proceed to printing, choose **Print** to open your operating system's print dialog.
 - To make changes to the paper size or orientation, choose **Configure**.
6. To exit the preview screen, choose **Exit preview**.

Exporting Amazon Quick Sight analyses or dashboards as PDFs

You can export content from a dashboard into a Portable Document Format file (PDF). Similar to a print-out, this format provides a snapshot of the current sheet as it appears on-screen at the time of download.

To export a dashboard sheet as a PDF

1. Open Quick Suite and choose **Dashboards** on the navigation pane at left.
2. Open the dashboard that you want to export.
3. At upper right, choose **Export, Download as PDF**. The download is prepared in the background.

When the file is ready to download, a message appears saying **Your PDF is ready..**

4. Choose **Download now** to download the file. Choose **Close** to close without downloading.

If you close this dialog box without downloading the file and want to recreate the file, repeat the previous step. Also, the downloadable file is available only temporarily for five minutes. If you wait too long to download it, the file expires. If this happens, Quick Sight instead displays an error message saying that the request has expired.

5. Repeat the previous steps for each sheet that you want to export.

You can also attach PDFs to dashboard email reports. For more information, see [Scheduling and sending Quick Sight reports by email](#).

Error codes for failed PDF export jobs

When you generate PDF reports in Amazon Quick Sight, you may encounter instances where your request to generate a PDF report fails. There are many reasons why a failure might occur. Quick Sight provides error codes that can help you understand why the error occurred and provide guidance to troubleshoot the issue. The following table lists the error codes that Quick Sight returns when a PDF export job fails.

Error code	Guidance
INVALID_DATAPREP_SYNTAX	Check the syntax for your calculated fields, and try again.

Error code	Guidance
POST_AGGREGATED_METRIC_AS_DIMENSION	Aggregated metrics/operands can't be used as visual's grouping dimensions. Choose a valid visual's grouping dimensions, and try again.
SPICE_TABLE_NOT_FOUND	The dataset has been deleted or is unavailable. Import a valid dataset, and try again.
FIELD_NOT_FOUND	A field is no longer available. Update or replace the missing fields in this dataset, and try again.
FIELD_ACCESS_DENIED	You don't have access to some fields in this dataset. Request access, and try again.
PERMISSIONS_DATASET_INVALID_COLUMN_VALUE	An invalid row level permission column value was found. Check your parent dataset rules, and try again.
COLUMN_NOT_FOUND	Replace the missing columns in your filters or parameters, and try again.
INVALID_COLUMN_TYPE	Some fields' data types have been changed and can not be automatically updated. Adjust these fields in your dataset, and try again.
PERMISSIONS_DATASET_USER_DENIED	You don't have access to this dataset. Request access to this dataset, and try again.
DATA_SOURCE_TIMEOUT	Your query has timed out. Reduce the amount of data, or import the data into SPICE, and try again.
MAX_PAGE_EXCEEDED_ERROR	Your file is ready but content is not complete. PDFs have a 1,000 page limit. Choose a shorter PDF, and try again.

Error code	Guidance
INSUFFICIENT_BODY_HEIGHT_ERROR	Adjust the header and footer to be less than the page height, and try again.
FIRST_PAGE_HEIGHT_TOO_SMALL_ERROR	Adjust sections to make room for your tables, and try again.
INTERNAL_ERROR	We can't create your PDF right now. Wait a few minutes, and try again.

Organizing assets into folders for Amazon Quick Sight

Applies to: Enterprise Edition

In Quick Suite Enterprise edition, your team members can create personal and shared folders to add hierarchical structure to Quick Sight asset management. Using folders, people can more easily organize, navigate through, and discover dashboards, analyses, datasets, data sources, and topics. Within a folder, you can still use your usual tools to search for assets or to add assets to your favorites list.

You can use the following types of folders with Quick Sight:

- Personal folders to organize work for yourself.

Personal folders are visible only to the person who owns them. You can't transfer ownership of personal folders to anyone else.

- Shared folders:
 - **Shared folders** organize work and simplify sharing among multiple people. To create and manage shared folders, you need to be a Quick Sight administrator.
 - **Shared restricted folders** are a type of shared folder in Quick Sight that ensure that assets remain in the shared folder. Assets that are created from assets that exist within a shared restricted folder must also stay in the restricted folder. Assets that are located in restricted folders can't be moved or shared outside of the restricted folder. For example, if you create a dataset that uses a data source that's located in a shared restricted folder, the new dataset can't be moved outside of that shared restricted folder.

Assets that are located in a restricted folder can be moved within the restricted folder tree into one or more subfolders. Subfolders of restricted folders behave like restricted folders, but dependent assets can exist in different subfolders under the same root restricted folder. The root restricted folder acts as a boundary that all assets in all subfolders can exist in as long as they remain within the root folder tree. For example, a dataset that is located in one subfolder can use a data source that is located either another subfolder in the same folder tree or in the root folder. Any supported asset type can be created in a root folder or in any of its subfolders. Users can have different roles in different subfolders. Subfolder permissions are inherited from the parent folders of that subfolder.

Restricted folders can only be created with the Quick Sight [CreateFolder](#) API operation.

- Users that are viewers on a folder and have the Author or Admin role in Quick Suite can view all asset types that are in the folder. Users that are viewers on a folder and have the Reader role in Quick Suite can only see dashboards and stories that are in the folder.

All shared folders are visible to people who have access to them.

Use the following topics to learn more about creating and configuring a folder or subfolder in Quick Sight.

Topics

- [Considerations for Quick Sight folders](#)
- [Overview of Quick Sight folders](#)
- [Permissions for Quick Sight shared folders](#)
- [Create and manage membership permissions for Quick Sight shared folders](#)
- [Creating Quick Sight scaled folders with the Quick Sight APIs](#)

Considerations for Quick Sight folders

Before you get started creating and modifying folders in Amazon Quick Sight, review the following limitations that apply to Quick Sight folders.

- You can't share folders in your Amazon account with people in other Amazon Web Services accounts.
- For people who have Quick Suite reader permissions, the following limitations apply:

- Readers can't own a personal or shared folder.
- Readers can't create or manage folders or folder content.
- Readers can't have the *contributor* access level.
- In shared folders, readers can only see dashboard assets.

In addition, these limitations apply specifically to shared folders:

- The name of a shared folder (at the top level of the tree) must be unique in your Amazon account.
- In a single folder, multiple assets can't have the same name. For example, in your top-level folder, you can't create two subfolders with the same name. In the same folder, you can't add two assets with the same name, even if they have different asset IDs. The path to each asset behaves like an Amazon S3 key name. It must be unique in your Amazon account.
- Restricted shared folders can only be created with the Quick Sight CLI.

See [Overview of Quick Sight folders](#) to learn more about the different types of folder available in Amazon Quick Sight.

Overview of Quick Sight folders

In Quick Sight, you can create personal and shared folders. You can also favorite your personal or shared folders for quick access by choosing the favorite (

) icon next to it.

You can do the following with personal folders:

- Create subfolders.
- Add assets to your folder, including analyses, dashboards, datasets, and data sources. To add assets to a personal folder, you must already have access to the assets. Multiple assets can have the same name.

Shared folders (unrestricted)

Quick Suite administrators can perform the following tasks with shared folders.

- Create or delete a shared folder and subfolders inside of it. You can move either of these around within the top-level folder.
- Add or remove owners, contributors, and viewers. When you make a person an *owner* of the folder, you give them ownership of every asset in the folder. For more information, see [Permissions for Quick Sight shared folders](#).

The following table summarizes the actions that a Quick Suite user can take when working with unrestricted shared folders based on their role.

Action	Owner	Contributor	Viewer
Share an asset in a folder with users that don't have access to the folder	Yes	No	No
Modify folder permissions	Yes	No	No
Create assets in the folder	Yes	Yes	No
Modify assets in the folder	Yes	Yes	No
Delete assets in the folder	Yes	Yes	No
Add an existing asset to a folder	Yes	Yes	No
Remove an asset from a shared folder	Yes	No	No
View assets in the folder	Yes	Yes	Yes

Action	Owner	Contributor	Viewer
Create downstream assets outside of the shared folder that use assets that are located in the shared folder	Yes	Yes	Yes*
Create downstream assets in the folder that use assets that are located outside of the folder	Yes	Yes	No
Create subfolders	Yes	Yes	No
Delete subfolders	Yes	No	No
Manage subfolder permissions	Yes	No	No
Add existing assets to subfolders	Yes	No	No
Create new assets in subfolders	Yes	Yes	No

Action	Owner	Contributor	Viewer
Delete assets in subfolders	Yes	Yes	No

*The user must be assigned an admin or author role to create assets.

Restricted shared folders

Restricted shared folders provide an additional security boundary that restricts the sharing of data outside of the folder. Administrators with the appropriate IAM permissions can perform the following tasks with restricted shared folders.

- Restricted folders can be created using the `CreateFolder` API operation. For more information about the `CreateFolder` API operation, see [CreateFolder](#).
- The contributor role is assigned to users that can create and edit assets within the restricted folders. Contributors can't manage the permissions of the folder or of the assets that are in the restricted folder.
- Administrators can assign folder contributor and viewer permissions to users with the `UpdateFolderPermissions` API operation. For more information about the `UpdateFolderPermissions` API operation, see [UpdateFolderPermissions](#).

The following table summarizes the actions that a Quick Sight user can take when working with restricted shared folders based on their role.

Action	Contributor	Viewer
Share an asset in a folder with users that don't have access to the folder	No	No
Modify folder permissions	No	No
Create assets in the folder	Yes	No

Action	Contributor	Viewer
Modify assets in the folder	Yes	No
Delete assets from the folder	Yes	No
Add an existing asset to a folder	No	No
Remove an asset from a shared folder	No	No
View assets in the folder	Yes	Yes
Create downstream assets outside of the shared folder that use assets that are located in the shared folder	No	No
Create downstream assets in the folder that use assets that are located outside of the folder	No	No
Create subfolders	Yes	No
Delete subfolders	No	No
Manage subfolder permissions	No	No
Add existing assets to subfolders	No	No

Action	Contributor	Viewer
Create new assets in subfolders	Yes	No
Delete assets from subfolders	Yes	No

The owner role is not supported for restricted shared folders.

After you choose which folder type best fits your use case, see [Permissions for Quick Sight shared folders](#) and [Create and manage membership permissions for Quick Sight shared folders](#) to create folders and set up folder permissions.

Permissions for Quick Sight shared folders

Shared folders have three permission levels. To set folder-level permissions for a user or group, see [Create and manage membership permissions for Quick Sight shared folders](#).

- **Owners** - The folder *owner* owns everything (folders, analyses, dashboards, datasets, data sources, topics) inside of the folder. They can create, edit, and delete the assets in the folder, modify permissions on the folder and its assets, and delete the folder entirely. The owner role is not supported for restricted shared folders.
- **Contributors** - A *contributor* can create, edit, and delete assets in a folder just like an owner. They can't delete the folder or modify permissions on the folder or on assets where they have contributor access that they inherited from the folder.
- **Viewers** - A *viewer* can only view the assets (folders, dashboards, datasets, data sources, topics) in the folder. A viewer can't edit or share those assets.

The following rules also apply to security for shared folders:

- Quick Suite readers' sharing status for a folder gets shared with the folder. However, a reader gets only read access to folders, and only dashboard access to visuals.
- Amazon security is enforced on every object within a folder. The folder applies the same type of security to the assets of whoever the folder is shared with according to their access level (admin, author, or reader).

- The *top-level folder* is the root folder of any subfolders. When a subfolder is shared at any level, the person whom the folder was shared with sees the root folder in the top-level folders view.
- The folder permission is the permission on the current folder, combined with permissions of all the folders leading to the root folder.
- A *shared asset* inherits its permission from the folder. A shared asset is created when an asset that belongs to the folder owner is added to a shared folder.
- If you own an unrestricted shared folder, you can transfer ownership of the folder to another Quick Suite admin.
- The owner role is not supported for restricted folders. The contributor role is assigned to authors that create and edit assets within the restricted folders. Folder contributors can't manage the permissions of the restricted folder or its assets.
- The correct IAM permissions are required to update the permissions of a restricted shared folder with the `UpdateFolderPermissions` API.

To create and manage permissions of a shared folder, see [Create and manage membership permissions for Quick Sight shared folders](#).

Create and manage membership permissions for Quick Sight shared folders

Shared folders (unrestricted)

To create a shared folder and to share the folder with one or more groups in the Quick Suite console, you must be an Amazon QuickSight administrator. You can also create a shared folder with the `CreateFolder` API operation. Use the following procedure to share or modify the membership permissions of a shared folder.

1. From the left navigation, choose **Folders** then **Shared Folders**. Find the folder that you want to share or manage permissions for.
2. To open the actions menu for that folder's row, choose the ellipsis (three dots).
3. Choose **Share**.
4. In the **Share folder** modal, add the groups and users with whom you want to share the contents of the folder.
5. For each user and group that you add, choose a permission level from the **Permissions** menu in that row.
6. To update the permission type for an existing user, choose **Manage folder access**.

7. When you're done setting user and group permissions for the folder, choose **Share**. Users are not notified that they now have access to the folder.

Restricted shared folders

Restricted shared folders can only be created with the CreateFolder API operation. The following example creates a restricted shared folder.

```
aws quicksight create-folder \  
--aws-account-id AWSACCOUNTID \  
--region us-east-1 \  
--folder-id example-folder-name \  
--folder-type RESTRICTED \  
--name "Example Folder" \
```

After you create a restricted shared folder, assign folder contributor and viewer permissions with a UpdateFolderPermissions API call. The following example updates the permissions of a restricted shared folder to grant contributor permissions to a user.

```
aws quicksight update-folder-permissions \  
--aws-account-id AWSACCOUNTID \  
--region us-east-1 \  
--folder-id example-folder-name \  
--grant-permissions Principal=arn:aws:quicksight::us-east-1::AWSACCOUNTID:user/default/:username,Actions=quicksight:CreateFolder  
,quicksight:DescribeFolder, \  
quicksight:CreateFolderMembership,quicksight>DeleteFolderMembership,qu  
icksight:DescribeFolderPermissions \
```

The permissions that you pass to the user depend on the type of folder role that you want to grant them. Use the following lists to determine which permissions are needed for the user that you want to grant folder access to.

Folder owner

- quicksight:CreateFolder
- quicksight:DescribeFolder
- quicksight:UpdateFolder
- quicksight>DeleteFolder

- quicksight:CreateFolderMembership
- quicksight>DeleteFolderMembership
- quicksight:DescribeFolderPermissions
- quicksight:UpdateFolderPermissions

Folder contributor

- quicksight:CreateFolder
- quicksight:DescribeFolder
- quicksight:CreateFolderMembership
- quicksight>DeleteFolderMembership
- quicksight:DescribeFolderPermissions

Folder viewer

- quicksight:DescribeFolder

After you create a shared folder, you can begin using the folder in Quick Sight.

You can also use the Quick Sight APIs to create special scaled folders that can be shared with up to 3000 namespaces. To learn more about creating a scaled folder, see [Creating Quick Sight scaled folders with the Quick Sight APIs](#).

Creating Quick Sight scaled folders with the Quick Sight APIs

You can use the Amazon Quick Sight APIs to create special scaled folders that can be shared with up to 3000 namespaces. Each namespace that is added to a folder can contain up to 100 principals. A *principal* is a user or a group of users. After you create a scaled folder and add the principals that you want, any QuickSight asset can be added to the folder. It can then be shared to every principal in the namespaces that the folder principals are assigned to. This streamlines the process to share Quick Sight assets with thousands of users.

Scaled folders can only be created with the Quick Sight APIs. When you create a scaled folder, you can share the folder with up to 100 principals that are in the same namespace. You can add principals that belong to a different namespace with an UpdateFolderPermissions API call.

After the folder is created, you can add and remove assets from the folder with the Quick Sight APIs or the Quick Suite console.

Each Amazon Quick Sight account holds up to 100 scaled folders. You can add up to 100 assets to a scaled folder. If you want to share a scaled folder with more than 3000 namespaces, contact [Amazon support](#).

Examples

The following examples show how to create a scaled folder with the Quick Sight APIs.

Prerequisites

Before you begin, verify that you have an Amazon Identity and Access Management role that grants the API user access to call the Quick Sight API operations. The following example shows an IAM policy that you can add to an existing IAM role to create, delete, or modify a scaled folder. With the sample policy, users can add dashboards, analyses, and datasets to a scaled folder.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:CreateFolder",
        "quicksight:CreateFolderMembership",
        "quicksight>DeleteFolderMembership",
        "quicksight>DeleteFolder",
        "quicksight:DescribeFolderPermissions",
        "quicksight:DescribeFolderResolvedPermissions",
        "quicksight:UpdateFolderPermissions",
        "quicksight:UpdateDashboardPermissions",
        "quicksight:UpdateAnalysisPermissions",
        "quicksight:UpdateDataSetPermissions"
      ],
      "Resource": "*"
    }
  ]
}
```

The following example creates a scaled folder.

```
aws quicksight create-folder \  
--aws-account-id "AWSACCOUNTID" \  
--region "us-east-1" \  
--name "eastcoast-users" \  
--sharing-model "NAMESPACE" \  
--folder-id "eastcoast-users"
```

After you create a scaled folder, share the folder with a principal in your account. You can only grant or revoke permissions to users and groups that are within the same namespace in each API call. The following example shares a scaled folder with a user in the same account that the folder exists in.

```
aws quicksight update-folder-permissions \  
--aws-account-id "AWSACCOUNTID" \  
--region "us-east-1" \  
--folder-id "eastcoast-users" \  
--grant-permissions \  
  '[  
    {"Actions":  
      ["quicksight:DescribeFolder",  
       "quicksight:UpdateFolder",  
       "quicksight:DeleteFolder",  
       "quicksight:DescribeFolderPermissions",  
       "quicksight:UpdateFolderPermissions",  
       "quicksight:CreateFolderMembership",  
       "quicksight:DeleteFolderMembership",  
       "quicksight:CreateFolder"  
      ],  
      "Principal": "arn:aws-cn:quicksight:us-east-1:AWSACCOUNTID:user/default/my-user"  
    }  
  ]'
```

After you share the folder with a new principal, validate the new folder permissions with a `describe-folder-permissions` API call.

```
aws quicksight describe-folder-permissions \  
--aws-account-id "AWSACCOUNTID" \  
--region "us-east-1" \  
--folder-id "eastcoast-users" \  

```

```
--namespace "default"
```

After you validate the new folder permissions, create a subfolder within the scaled folder. The subfolder inherits the permissions of the scaled folder that it's created in.

```
aws quicksight create-folder \  
--aws-account-id "AWSACCOUNTID" \  
--region "us-east-1" \  
--name "new-york-users" \  
--sharing-model "NAMESPACE" \  
--folder-id "new-york-users" \  
--parent-folder-arn "arn:aws-cn:quicksight:us-east-1:AWSACCOUNTID:folder/eastcoast-  
users"
```

The following example validates the inherited permissions of the new subfolder.

```
aws quicksight describe-folder-resolved-permissions \  
--aws-account-id "AWSACCOUNTID" \  
--region "us-east-1" \  
--folder-id "new-york-users" \  
--namespace "default"
```

After you validate the permissions of the subfolder, add the Quick Sight asset that you want to share to the folder. After you add the asset to the subfolder, the asset is shared with every principal that the subfolder is shared with. The following example adds a dashboard to a subfolder.

```
aws quicksight create-folder-membership \  
--aws-account-id "AWSACCOUNTID" \  
--folder-id "new-york-users" \  
--member-id "my-dashboard" \  
--member-type "DASHBOARD" \  
--region "us-east-1"
```

Exploring interactive dashboards in Amazon Quick Sight

Intended audience: Amazon Quick Suite Dashboard subscribers or viewership

In Amazon Quick Sight, a *data dashboard* is a collection of charts, graphs, and insights. It's like a newspaper that's all about the data that you're interested in, except it has digital pages. Instead of reading it, you interact with it.

Dashboards come in a wide variety of designs, depending on what you do and the analytics that you need to do it well. Using Quick Sight, you can interact with your data on a webpage or your mobile device. If you also subscribe by mail, you can see a static preview of it.

The story told by your data reflects the expertise of the analysts and data scientists who built the dashboards. They refine the data, add calculations, find angles on the story, and decide how to present it. The publisher designs the dashboard and fills it with interactive data visualizations and controls that adjust your view. Publishers can customize the level of interactivity that you have, including filter and search options. You can interact with the active items on the screen to filter, sort, drill down, or jump to another tool.

When you view a dashboard, you're seeing the most recently received data. As you interact with the items on the screen, any changes you make change your view of the dashboard, and no one else's. Thus, your device's privacy is assured, although the publisher can tell what you looked at. After you close the dashboard, your explorations aren't preserved and neither is the data. As always, while you're an Quick Suite reader, your monthly subscription is provided by the publishers of the dashboards at no additional cost to you.

If you're also a dashboard publisher—we call them authors, because they write reports—you can also save a copy of the dashboard for further analysis. If you find a new feature of the data that you want to publish, work with the original authors to update it. That way, everyone can see the same version of the story. However, you can also use your copy to learn how their design works or to inspire your work on something entirely new. Then, when you're finished, you can publish your analysis as a new dashboard.

To learn to set up dashboards, see [Sharing and subscribing to data in Amazon Quick Sight with dashboards and reports](#).

Topics

- [Interacting with Amazon Quick Sight dashboards](#)
- [Interacting with paginated reports in Amazon Quick Sight](#)
- [Subscribing to Amazon Quick Sight dashboard emails and alerts](#)
- [Creating a reader generated report in Amazon Quick Sight](#)

- [Bookmarking views of a Amazon Quick Sight dashboard](#)

Interacting with Amazon Quick Sight dashboards

To access a dashboard that you've been invited to share, follow the instructions in the invitation email. You can also access a dashboard if it's embedded into an application or website that you already have access to.

To fit the dashboard to your screen, open the **View** menu at upper right and select **Fit to window**.

Depending on how the dashboard is configured, you can find all or some of the following elements:

- The menu bar – This displays the name of the dashboard. Also, the menu bar shows what you can do with the dashboard, including **Undo**, **Redo**, and **Reset**, on the left. As you interact with the dashboard, you can use these as tools to help you explore, knowing that you can change your view without losing anything. On the right, you can find options to **Print** the dashboard, work with **Data**, choose a different Amazon **Region**, and open your **User Profile**. The user profile menu has options so you can choose the language that Amazon Quick Sight displays. It also has links to the Quick Suite **Community** and the online documentation (**Help**).
- The dashboard sheets – If your dashboard has multiple sheets, these display as tabs across the top of the dashboard.
- The **Filter** menu – This option displays to the left of the dashboard, if the dashboard publisher allows filtering.
- The **Controls** palette – If your dashboard includes controls, you can use them to choose the options (parameters) that you want to apply to your dashboard. Sometimes a control value is selected for you, and sometimes it's set to **ALL**.
- The dashboard title – If your dashboard has a title, it is usually a larger heading. It might have some status information or instructions below it.
- The dashboard widgets – The items on the screen can include charts, graphs, insights, narratives, or images. To see them all, you might need to scroll vertically or horizontally.

Using filters on Amazon Quick Sight dashboard data

You can use filters to refine the data displayed in a visual. Filters are applied to the data before any aggregate functions. If you have multiple filters, all top-level filters apply together using AND. If the filters are grouped inside a top-level filter, the filters in the group apply using OR.

Amazon Quick Sight applies all of the enabled filters to the field. For example, suppose that there is one filter of `state = WA` and another filter of `sales >= 500`. In this case, the dataset contains only records that meet both of those criteria. If you disable one of these, only one filter applies. Take care that multiple filters applied to the same field aren't mutually exclusive.

Viewing filters

To see the existing filters, choose **Filter** on the element settings menu, then choose to view filters. The filters display in the **Applied filters** panel in order of creation, with the oldest filter on top.

Understanding filter icons in an Amazon Quick Sight dashboard

Filters in the **Applied filters** panel display icons to indicate how they are scoped and whether they are enabled.

A filter that isn't enabled is grayed out, and you can't select its check box.

One of several scope icons displays to the right of the filter name to indicate the scope set on that filter. The scope icon resembled four boxes in a square. If all boxes are filled, the filter applies to all visuals on the analysis sheet. If only one box is filled, the filter applies to the selected visual only. If some boxes are filled, the filter applies to some of the visuals on the sheet, including the one currently selected.

The scope icons match the ones that display on the filter menu when you are choosing the scope for the filter.

Viewing filter details in an Amazon Quick Sight dashboard

To see filter details, choose **Filter** at left. The filter view retains your last selection. So when you open **Filter**, you see either the **Applied filters** or the **Edit filter** view.

In the **Applied filters** view, you can choose any filter to view its details. The filters in this list can change depending on the scope of the filter, and which visual you currently have selected.

You can close the **Edit filter** view by choosing the selector on the right. Doing this resets the **Filter** view.

Filtering data during your session in Amazon Quick Sight

While your dashboard session is active, you can filter data in three ways:

1. If your dashboard has controls at the top of the screen, you can use them to filter data by choosing from a preset list of values.
2. You can use the filter icon on each widget's settings menu.
3. You can create your own filters by using the filter panel on the left side of the page. The filter icon looks like the following.

To create a filter, choose the **Filter** icon at left.

The first step is to choose which dashboard element you want to filter.

Click on the item you choose, so that a highlight appears around the selected item. Also, if any filters are already there, they display in a list. If there aren't any filters, you can add one by using the plus sign (+) near **Filters**.

Filtering options vary depending on the data type of the field you want to filter, and on the options you choose inside the filter. The following screenshot shows some of the options available for a time-range date filter.

For each filter, you can choose whether to apply it to one, some, or all dashboard elements. You can also enable or disable filters by using the check box next to the name of the filter. To delete a filter, edit it and scroll to the bottom to see the options. Remember that your filters aren't saved from one session to the next.

For more detailed information on creating filters, see [Filtering data in Amazon Quick Sight](#).

Using the elements on the Amazon Quick Sight dashboard

Each widget has a settings menu that appears when you select that widget. This menu provides options to zoom in or out, filter the data, export the data, and more. The options vary depending on what type of widget the element is.

When you choose a data point, several actions are available. You can click or tap on a data point, for example on a bar in a bar chart, on a point where the line bends on a line chart, and so on. The available options vary based on what type of item it is.

These actions are as follows:

- Focus on or exclude.

You can focus on or exclude specific data in a field, for example regions, metrics, or dates.

- Drill up or drill down.

If your dashboard contains data on which you can drill down or up, you can drill up to a higher level or drill down to explore deeper details.

- Custom URL actions.

If your dashboard contains custom actions, you can activate them by choosing a data point or by right-clicking it. For example, you might be able to email someone directly from the dashboard. Or you might open another sheet, website, or application, and send it the value you chose from this one.

- Change chart colors or specific field colors.

You can change all the chart colors to a specific color. Or you can choose a specific field value to change its color of the element it's part of.

Sorting dashboard data in Amazon Quick Sight

You can sort data in three ways:

1. You can hover over the label for the field you want to sort by, and choose the sort icon.
2. You can choose the filter icon at the upper right of one of the dashboard elements.
3. You can click or tap on the field and choose **Sort** from the context menu.

Sorting for pivot tables is different; you specify the sort order by using the column sort icon on the pivot table.

Exporting and printing interactive Amazon Quick Sight dashboard reports

You can export or print a PDF version of an interactive dashboard. You can also export some visuals in a dashboard to a CSV. Exporting an entire dashboard to a CSV is not currently supported for interactive dashboards.

Exporting data from a dashboard to a PDF

To export an interactive dashboard report as a PDF

1. From the dashboard report that you want to export, choose the **Export** icon at the top right.

2. Choose **Generate PDF**.
3. When you choose **Generate PDF**, Quick Sight will begin preparing the dashboard report for download. Choose **View downloads** in the blue pop-up to open the **Downloads** pane on the right.
4. There are two ways to download your report:
 - Choose **DOWNLOAD NOW** in the green pop-up.
 - Choose the **Export** icon at the top right, and then choose **View downloads** to view and download every report that is ready to download.

To print an interactive dashboard report

1. From the report that you want to print, choose the **Export** icon at the top right, and then choose **Print**.
2. In the **Prepare for printing** pop-up that appears, choose the paper size and orientation that you want. You can optionally choose to include the background color by selecting **Print background color**.
3. Choose **GO TO PREVIEW**.
4. In the preview window that appears, choose **PRINT**.

Exporting data from a dashboard to a CSV

Note

Export files can directly return information from the dataset import. This makes the files vulnerable to CSV injection if the imported data contains formulas or commands. For this reason, export files can prompt security warnings. To avoid malicious activity, turn off links and macros when reading exported files.

To export data from an analysis or dashboard to a comma-separated value (CSV) file, use the settings menu at the upper right of a widget. Exports only include data that currently displays in the item that you choose.

In tables and pivot tables, you can export data to a comma-separated value (CSV) file or Microsoft Excel file. You can choose to export only visible fields or all fields.

To export only visible fields to a CSV or Excel file, choose the menu at upper-right of the visual. Choose either **Export to CSV** or **Export to Excel**, and then choose **Export visible fields to CSV** or **Export visible fields to Excel**.

To export all fields to a CSV or Excel file, choose the menu at upper-right of the visual. Choose either **Export to CSV** or **Export to Excel**, and then choose **Export all fields to CSV** or **Export all fields to Excel**.

Generate an executive summary of an Amazon Quick Sight dashboard

Dashboard readers can generate executive summaries that provide a summary of all insights that Quick Sight has generated for the dashboard. Executive summaries make it easier for readers to find key insights and information about a dashboard at a glance.

When readers are viewing a dashboard that uses executive summaries, the **Executive summary** option is available in the **Build** dropdown list that is located in the top right of the Dashboard's page. Use the procedure below to generate an executive summary. If a dashboard doesn't use executive summaries, the **Executive summary** option does not appear in the **Build** dropdown list.

To generate an executive summary

1. In the dashboard that you want to work in, choose **Build**, and then choose **Executive summary**.
2. Choose **Summarize**. The executive summary is generated and the appears on the left.

Executive summaries use the data of the current dashboard sheet and visual settings. If the dashboard or visual settings are updated, a warning appears at the top of an executive summary. To refresh the executive summary of an updated dashboard, generate a new executive summary.

After an executive summary is generated, Amazon Quick Suite readers can copy the summary to their clipboard in order to share with others, or include in a Quick Sight story. For more information about Quick Sight stories, see [???](#).

Interacting with paginated reports in Amazon Quick Sight

To access a paginated report that you've been invited to share, follow the instructions in the invitation email. You can also access a paginated report if it's embedded into an application or website that you already have access to.

To fit the paginated report to your screen, open the **View** menu at upper right and select **Fit to window**. You can also zoom in and out using the plus (+) and minus (-) icons on the top left corner of the report.

Exporting and printing Amazon Quick Sight reports

Paginated reports are designed to be viewed from a specific point of time. These reports, or snapshots, can be printed or downloaded as a PDF or CSV.

To export a paginated report as a PDF

1. From the paginated report that you want to export, choose the **Export** icon at the top right.
2. Choose **Generate PDF**.
3. When you choose **Generate PDF**, Quick Sight will begin preparing the paginated report for download. When the report is ready, a green pop up will appear that says **Your PDF is ready**.
4. There are two ways to download your report:
 - Choose **DOWNLOAD NOW** in the green pop-up.
 - Choose the **Export** icon at the top right, and then choose **View downloads** to view and download every report that is ready to download.

To export a paginated report as a CSV

1. From the report that you want to export, choose the **Scheduling** icon at the top right, and then choose **Recent snapshots**.
2. In the **Recent snapshots** menu that appears on the right, snapshots are sorted from most recently generated to the oldest. Snapshots are stored for up to 1 year. Find the report that you want to download and choose the download icon to the right of the report.
3. In the report pop-up that appears, choose the download icon next to the version of the report that you want to download. You can choose to download the report as a CSV, or you can download the report as a PDF.

To print a paginated report

1. From the report that you want to print, choose the **Export** icon at the top right, and then choose **Print**.

2. When you choose **Print**, your browser's printer pop-up appears. From here, you can print the PDF the same way you would print anything else on your browser.

Subscribing to Amazon Quick Sight dashboard emails and alerts

Using Amazon Quick Sight, you can subscribe to updates for certain events, such as dashboard updates and anomaly alerts.

Topics

- [Sign up for dashboard emails](#)
- [Sign up for anomaly alerts](#)

Sign up for dashboard emails

You can sign up to get a dashboard in report form, and receive it in an email. You can also configure your report settings.

To change subscription and report settings for a dashboard

1. Open a dashboard that is shared with you.
2. Choose the **Schedules** icon at upper right, and then choose **Schedules** in the dropdown.
3. The **Schedules** pane appears on the right. This pane shows all of the different scheduled reports that you are or can be subscribed to. Navigate to the report that you want and toggle the switch to subscribe or unsubscribe from the report.

Sign up for anomaly alerts

On a dashboard that has a narrative insight that's configured for anomaly detection, you can sign up to get alerts for anomalies and contribution analysis. You receive anomaly alerts when anomalies are updated. The alerts email displays the total number of anomalies, and provides detail on the top five, according to your personal alert configuration. You receive key driver contribution analysis when it's updated, provided that contribution analysis is configured to run with anomaly detection.

To set up anomaly alerts

1. Open a dashboard that is shared with you.

2. You can configure alerts from one of two screens. Choose one of the following, then go to the next step:
 - In the dashboard, locate the anomaly widget that you're interested in. Select it so that it has a highlighted box around it.
 - If you're in the dashboard and have the **Explore Anomalies** page open, you can configure the alert without returning to the dashboard view.
3. At upper right, choose **Configure alert**. The **Alert** configuration screen appears.
4. For **Severity**, choose the lowest level of significance that you want to see.

For **Direction**, choose to get alerts about anomalies that are **Higher than expected** or **Lower than expected**. You can also choose **[ALL]** to receive alerts about all anomalies.

5. Choose **OK** to confirm your choices.
6. To stop receiving to an anomaly alert, locate the anomaly widget in the dashboard and use the bell icon to unsubscribe. You can also use the **To manage this alert** link at the bottom of an alert email.

Creating a reader generated report in Amazon Quick Sight

If a Amazon Quick Suite author has set up a prompted report for a Quick Sight paginated report, Quick Sight dashboard viewers can use the prompt to schedule their own reports for themselves. For more information about prompts for paginated reports, see [Setting up prompts for paginated reports](#).

Use the following sections to learn how to create and modify a reader generated report.

Topics

- [Creating a reader generated report](#)
- [Loading a saved view of a Quick Sight reader generated report](#)
- [Updating the view of a scheduled reader generated report](#)
- [Updating a reader generated report schedule](#)

Creating a reader generated report

Use the following procedure to create a reader generated report.

To create a reader generated report

1. Open the [Quick Suite console](#).
2. Open the dashboard that you want to create a report for.
3. Choose the **Scheduling** at the top of the dashboard page.
4. The scheduling pane opens. To add a new report schedule, choose **Add**. If you do not see the **Add** button, the dashboard does not contain a paginated sheet, or your Quick Suite account does not have the Paginated reports add on. For more information about the paginated reports add on, see [Getting started](#).
5. For **Schedule name**, enter a name for the new schedule. The schedule name can be up to 100 characters long.
6. For **Description**, choose the view option that you want the report to use. You can choose from the following views:
 - **Custom view** – The current view of the dashboard.
 - **Original view** – The author published view of the dashboard.
7. For **Content**, choose the paginated report sheet that you want to generate a PDF report for.
8. For **Dates**, choose the frequency at which you want to receive the report. Scheduling options that are available for an email report include the following:
 - **Once (Does not repeat)** – Sends the report only once at the date and time that you choose.
 - **Daily** – Repeats daily at the time that you choose.
 - **Weekly** – Repeats each week on the same day or days at the time that you choose. You can also use this option to send reports in weekly intervals, such as every other week or every three weeks.
 - **Monthly** – Repeats each month on the same day of the month at the time that you choose. You can also use this option to send reports on specific days of the month, such as the second Wednesday or the last Friday of each month.
 - **Yearly** – Repeats each year on the same day of the month or months selected at the time that you choose. You can also use this option to send reports on specific days or sets of days in selected months. For example, you can configure a report to be sent on the first Monday of January, March, and September, or on July 14th, or on the second day of February, April, and June each year.
 - **Custom** – Configure your own scheduled report that best fits your business needs.

The scheduled report is sent within 1 hour from the specified time. Delays may occur during peak hours.

9. In the **Email** tab, for **E-mail subject line**, enter a custom subject line, or leave it blank to use the report title.
10. Enter the email addresses of the Quick Suite group name of the users or groups that you want to receive the report.
11. For **Email header**, enter the header that you want the email report to show.
12. (Optional) For **E-mail body text**, leave it blank or enter a custom message to display at the beginning of the email.
13. (Optional, recommended) To send a sample of the report before you save changes, choose **Send test report**.
14. Do one of the following:
 - (Recommended) Choose **Save** to confirm your entries.
 - To immediately send a report, choose **Save and run now**. The report is sent immediately, even if your schedule's start date is in the future.

After you save a report schedule, the schedule appears in the **Schedules** pane. Reader generated reports are only available to the user that created them and can't be shared.

Loading a saved view of a Quick Sight reader generated report

Amazon Quick Suite readers can use the **Schedules** pane to load a saved view of any scheduled paginated report they have created or received. Use the following procedure to load a saved review of a scheduled report.

To load a saved view of a scheduled report

1. Open the [Quick Suite console](#).
2. Open the dashboard that contains the report that you want to change.
3. Choose the **Scheduling** at the top of the dashboard page.
4. The scheduling pane opens. Locate the schedule that you want to change and choose the ellipsis (three dots) icon next to the report to open the schedule menu, and then choose **Details**.

5. Choose **Load saved view**. The saved view of the dashboard that was used for the selected schedule is rendered. All filter values that were active when the dashboard snapshot was taken are applied to the dashboard. When a saved view of a dashboard is loaded, the reader's current view of the dashboard is lost.

Updating the view of a scheduled reader generated report

After a Amazon Quick Suite reader has created a report in Quick Sight, they can use the **Schedules** pane to update the dashboard view that is used in the scheduled report. Use the following procedure to update the dashboard view of a scheduled report.

To change the dashboard view of a scheduled report

1. Open the [Quick Suite console](#).
2. Open the dashboard that contains the report that you want to change.
3. Choose the **Scheduling** at the top of the dashboard page.
4. The scheduling pane opens. Locate the schedule that you want to change and choose the ellipsis (three dots) icon next to the report to open the schedule menu, and then choose **Details**.
5. Choose **Load saved view**. The saved view of the dashboard that was used for the selected schedule is rendered. All filter values that were active when the dashboard snapshot was taken are applied to the dashboard. When a saved view of a dashboard is loaded, the reader's current view of the dashboard is lost.
6. Update the dashboard filters that you want to change.
7. Choose the **Scheduling** at the top of the dashboard page.
8. The scheduling pane opens. Locate the schedule that you want to change and choose the ellipsis (three dots) icon next to the report to open the schedule menu, and then choose **Edit**.
9. Navigate to the **Dashboard view** section, and then choose **Custom view**. The new filter values that you updated are applied to the dashboard report.
10. Choose **Save** to update the schedule.

Updating a reader generated report schedule

After they create a reader generated report, Amazon Quick Suite readers can use the **Schedules** pane to make a report schedule active or inactive. Use the following procedure to update active status of a reader generated report schedule.

1. Open the [Quick Suite console](#).
2. Open the dashboard that contains the report that you want to change.
3. Choose the **Scheduling** at the top of the dashboard page to open the **Schedules** pane.
4. Choose **Schedules**.
5. Navigate to the **My schedules** section and find the schedule that you want to update.
6. Use the toggle to set the report schedule to **Active** or **Inactive**.
7. When you are finished making changes to the report schedule, close the **Schedules** pane.

Bookmarking views of a Amazon Quick Sight dashboard

When you load a dashboard as an Amazon Quick Suite reader or author, you can create bookmarks to capture specific views of your interests. For example, you can create a bookmark for a dashboard with a specific filter setting that differs from the original published dashboard. By doing this, you can quickly return to the data that's relevant to you.

After you create a bookmark, you can set it as the default view of the dashboard that you see when you open the dashboard in a new session. This doesn't affect anyone else's view of the dashboard.

You can create up to 200 bookmarks for a dashboard and share them by a URL link with other subscribers of that dashboard.

Dashboard bookmarks are available on the Quick Suite console.

Dashboard bookmarks for paginated reports are currently not supported. For more information on paginated reports, see [Working with paginated reports in Amazon Quick Sight](#).

Use the following topics to learn how to use bookmarks.

Topics

- [Creating bookmarks in Amazon Quick Sight](#)
- [Updating bookmarks in Amazon Quick Sight](#)
- [Renaming bookmarks in Amazon Quick Sight](#)

- [Making a bookmark the default view in Amazon Quick Sight](#)
- [Sharing bookmarks in Amazon Quick Sight](#)
- [Deleting bookmarks in Amazon Quick Sight](#)

Creating bookmarks in Amazon Quick Sight

Use the following procedure to create a bookmark for a dashboard.

To create a bookmark for a dashboard

1. Open the published dashboard that you want to view and make changes to the filters or parameters, or select the sheet that you want. For example, you can filter to the Region that interests you, or you can select a specific date range using a sheet control on the dashboard.
2. Choose the bookmark icon at upper right, and then choose **Add bookmark**.
3. In the **Add a bookmark** pane that opens, enter a name for the bookmark, and then choose **Save**.

The bookmark is saved, and the dashboard name updates with the bookmark name (at top left).

You can return to the original dashboard view that the author published at any time by selecting **Original dashboard** in the **Bookmarks** pane at right.

Updating bookmarks in Amazon Quick Sight

At any time, you can change a bookmark dashboard view and update the bookmark to always reflect those changes.

To update a bookmark

1. Open the published dashboard and make needed changes to the filters or parameters, or select a sheet.
2. Choose the bookmark icon at upper right.
3. In the **Bookmarks** pane that opens, choose the context menu (the three vertical dots) for the bookmark that you want to update, and then choose **Update**.

A message appears, confirming the update.

Renaming bookmarks in Amazon Quick Sight

Use the following procedure to rename a bookmark.

To rename a bookmark

1. In a published dashboard, choose the bookmark icon at upper right to open the **Bookmarks** pane.
2. In the **Bookmarks** pane, choose the context menu (the three vertical dots) for the bookmark that you want to rename, and then choose **rename**.
3. In the **Rename bookmark** pane, enter a name for the bookmark, and then choose **Save**.

Making a bookmark the default view in Amazon Quick Sight

By default, when you update a dashboard, Quick Sight remembers those changes and keeps them after you close the dashboard. This way, you can pick up where you left off when you open the dashboard again. You can set a bookmark as the default view of a dashboard instead. If you do, anytime that you open the dashboard, the bookmark view is presented to you, regardless of the changes you made during your last session.

To set a bookmark as your default view of the dashboard

1. In a published dashboard, choose the bookmark icon at upper right to open the **Bookmarks** pane.
2. In the **Bookmarks** pane, choose the context menu (the three dots) for the bookmark that you want to set as your default view, and then choose **Set as default**.

Sharing bookmarks in Amazon Quick Sight

After you create a bookmark, you can share a URL link for the view with others who have permission to view the dashboard. They can then save that view as their own bookmark.

To share a bookmark with another dashboard subscriber

1. In a published dashboard, choose the bookmark icon at upper right to open the **Bookmarks** pane.
2. In the **Bookmarks** pane, choose the bookmark that you want to share so that the dashboard updates to that view.

3. Choose the share icon at upper right, and then choose **Share this view**.

You can copy the URL link that Quick Sight provides and paste it in an email or IM message to share it with others. The recipient of the URL link can then save the view as their own bookmark. For more information about sharing views of a dashboard, see [Sharing your view of a Amazon Quick Sight dashboard](#).

Deleting bookmarks in Amazon Quick Sight

Use the following procedure to delete a bookmark.

To delete a bookmark

1. In a published dashboard, choose the bookmark icon at upper right to open the **Bookmarks** pane.
2. In the **Bookmarks** pane, choose the context menu (the three vertical dots) for the bookmark that you want to delete, and then choose **Delete**.
3. In the **Delete Bookmark** pane that opens, choose **Yes, Delete Bookmark**.

Gaining insights with machine learning (ML) in Amazon Quick Sight

Amazon Quick Sight uses machine learning to help you uncover hidden insights and trends in your data, identify key drivers, and forecast business metrics. You can also consume these insights in natural language narratives embedded in dashboards.

Using machine learning (ML) and natural language capabilities, Amazon Quick Sight Enterprise Edition takes you beyond descriptive and diagnostic analysis, and launches you into forecasting and decision-making. You can understand your data at a glance, share your findings, and discover the best decisions to achieve your goals. You can do this without developing teams and technology to create the necessary machine learning models and algorithms.

You likely have already built visualizations that answer questions about what happened, when, where, and provide drill down for investigation and identification of patterns. With ML insights, you can avoid spending hours manually analyzing and investigating. You can select from a list of customized context-sensitive narratives, called *autonarratives*, and add them to your analysis. In addition to choosing autonarratives, you can choose to view forecasts, anomalies, and factors

contributing to these. You can also add autonarratives that explain the key takeaways in plain language, providing a single data-driven truth for your company.

As time passes and data flows through the system, Amazon Quick Sight continually learns so it can deliver ever more pertinent insights. Instead of deciding what the data means, you can decide what to do with the information it provides.

With a shared foundation based on machine learning, all of your analysts and stakeholders can see trends, anomalies, forecasts, and custom narratives built on millions of metrics. They can see root causes, consider forecasts, evaluate risks, and make well-informed, justifiable decisions.

You can create a dashboard like this with no manual analysis, no custom development skills, and no understanding of machine learning modeling or algorithms. All this capability is built into Amazon Quick Sight Enterprise Edition.

 Note

Machine learning capabilities are used as needed throughout the product. Features that actively use machine learning are labeled as such.

With ML Insights, Amazon Quick Sight provides three major features:

- **ML-powered anomaly detection** – Amazon Quick Sight uses Amazon's proven machine learning technology to continuously analyze all your data to detect anomalies (outliers). You can identify the top drivers that contribute to any significant change in your business metrics, such as higher-than-expected sales or a dip in your website traffic. Amazon Quick Sight uses the Random Cut Forest algorithm on millions of metrics and billions of data points. Doing this enables you to get deep insights that are often buried in the aggregates, inaccessible through manual analysis.
- **ML-powered forecasting** – Amazon Quick Sight enables nontechnical users to confidently forecast their key business metrics. The built-in ML Random Cut Forest algorithm automatically handles complex real-world scenarios such as detecting seasonality and trends, excluding outliers, and imputing missing values. You can interact with the data with point-and-click simplicity.
- **Autonarratives** – By using automatic narratives in Amazon Quick Sight, you can build rich dashboards with embedded narratives to tell the story of your data in plain language. Doing this can save hours of sifting through charts and tables to extract the key insights for reporting. It also creates a shared understanding of the data within your organization so you make decisions

faster. You can use the suggested autonarrative, or you can customize the computations and language to meet your unique requirements. Amazon Quick Sight is like providing a personal data analyst to all of your users.

Topics

- [Understanding the ML algorithm used by Amazon Quick Sight](#)
- [Dataset requirements for using ML insights with Amazon Quick Sight](#)
- [Working with insights in Amazon Quick Sight](#)
- [Creating autonarratives with Amazon Quick Sight](#)
- [Detecting outliers with ML-powered anomaly detection](#)
- [Forecasting and creating what-if scenarios with Amazon Quick Sight](#)

Understanding the ML algorithm used by Amazon Quick Sight

You don't need any technical experience in machine learning to use the ML-powered features in Amazon Quick Sight. This section dives into the technical aspects of the algorithm, for those who want the details about how it works. This information isn't required reading to use the features.

Amazon Quick Sight uses a built-in version of the Random Cut Forest (RCF) algorithm. The following sections explain what that means and how it is used in Amazon Quick Sight.

First, let's look at some of the terminology involved:

- **Anomaly** – Something that is characterized by its difference from the majority of the other things in the same sample. Also known as an outlier, an exception, a deviation, and so on.
- **Data point** – A discrete unit—or simply put, a row—in a dataset. However, a row can have multiple data points if you use a measure over different dimensions.
- **Decision Tree** – A way of visualizing the decision process of the algorithm that evaluates patterns in the data.
- **Forecast** – A prediction of future behavior based on current and past behavior.
- **Model** – A mathematical representation of the algorithm or what the algorithm learns.
- **Seasonality** – The repeating patterns of behavior that occur cyclically in time series data.

- Time series – An ordered set of date or time data in one field or column.

Topics

- [What's the difference between anomaly detection and forecasting?](#)
- [What is RCF?](#)
- [How RCF is applied to detect anomalies](#)
- [How RCF is applied to generate forecasts](#)
- [References for machine learning and RCF](#)

What's the difference between anomaly detection and forecasting?

Anomaly detection identifies outliers and their contributing drivers to answer the question "What happened that doesn't usually happen?" Forecasting answers the question "If everything continues to happen as expected, what happens in the future?" The math that allows forecasting also enables us to ask "If a few things change, what happens then?"

Both anomaly detection and forecasting begin by examining the current known data points. Amazon Quick Sight anomaly detection begins with what is known so it can establish what is outside the known set, and identify those data points as anomalous (outliers). Amazon Quick Sight forecasting excludes the anomalous data points, and sticks with the known pattern. Forecasting focuses on the established pattern of data distribution. In contrast, anomaly detection focuses on the data points that deviate from what is expected. Each method approaches decision-making from a different direction.

What is RCF?

A *random cut forest* (RCF) is a special type of *random forest* (RF) algorithm, a widely used and successful technique in machine learning. It takes a set of random data points, cuts them down to the same number of points, and then builds a collection of models. In contrast, a model corresponds to a decision tree—thus the name forest. Because RFs can't be easily updated in an incremental manner, RCFs were invented with variables in tree construction that were designed to allow incremental updates.

As an unsupervised algorithm, RCF uses cluster analysis to detect spikes in time series data, breaks in periodicity or seasonality, and data point exceptions. Random cut forests can work as a synopsis or sketch of a dynamic data stream (or a time-indexed sequence of numbers). The answers to our

questions about the stream come out of that synopsis. The following characteristics address the stream and how we make connections to anomaly detection and forecasting:

- A *streaming algorithm* is an online algorithm with a small memory footprint. An online algorithm makes its decision about the input point indexed by time t before it sees the $(t+1)$ -st point. The small memory allows nimble algorithms that can produce answers with low latency and allow a user to interact with the data.
- Respecting the ordering imposed by time, as in an *online* algorithm, is necessary in anomaly detection and forecasting. If we already know what will happen the day after tomorrow, then predicting what happens tomorrow isn't a forecast—it's just interpolating an unknown missing value. Similarly, a new product introduced today can be an anomaly, but it doesn't necessarily remain an anomaly at the end of the next quarter.

How RCF is applied to detect anomalies

A human can easily distinguish a data point that stands out from the rest of the data. RCF does the same thing by building a "forest" of decision trees, and then monitoring how new data points change the forest.

An *anomaly* is a data point that draws your attention away from normal points—think of an image of a red flower in a field of yellow flowers. This "displacement of attention" is encoded in the (expected) position of a tree (that is, a model in RCF) that would be occupied by the input point. The idea is to create a forest where each decision tree grows out of a partition of the data sampled for training the algorithm. In more technical terms, each tree builds a specific type of binary space partitioning tree on the samples. As Amazon Quick Sight samples the data, RCF assigns each data point an anomaly score. It gives higher scores to data points that look anomalous. The score is, in approximation, inversely proportional to the resulting depth of the point in the tree. The random cut forest assigns an anomaly score by computing the average score from each constituent tree and scaling the result with respect to the sample size.

The votes or scores of the different models are aggregated because each of the models by itself is a weak predictor. Amazon Quick Sight identifies a data point as anomalous when its score is significantly different from the recent points. What qualifies as an anomaly depends on the application.

The paper [Random Cut Forest Based Anomaly Detection On Streams](#) provides multiple examples of this state-of-the-art online anomaly detection (time-series anomaly detection). RCFs are used on contiguous segments or "shingles" of data, where the data in the immediate segment acts as

a context for the most recent one. Previous versions of RCF-based anomaly-detection algorithms score an entire shingle. The algorithm in Amazon Quick Sight also provides an approximate location of the anomaly in the current extended context. This approximate location can be useful in the scenario where there is delay in detecting the anomaly. Delays occur because any algorithm needs to characterize "previously seen deviations" to "anomalous deviations," which can unfold over some time.

How RCF is applied to generate forecasts

To forecast the next value in a stationary time sequence, the RCF algorithm answers the question "What would be the most likely completion, after we have a candidate value?" It uses a single tree in RCF to perform a search for the best candidate. The candidates across different trees are aggregated, because each tree by itself a weak predictor. The aggregation also allows the generation of quantile errors. This process is repeated t times to predict the t -th value in the future.

The algorithm in Amazon Quick Sight is called *BIFOCAL*. It uses two RCFs to create a CALibrated BI-Forest architecture. The first RCF is used to filter out anomalies and provide a weak forecast, which is corrected by the second. Overall, this approach provides significantly more robust forecasts in comparison to other widely available algorithms such as ETS.

The number of parameters in the Amazon Quick Sight forecasting algorithm is significantly fewer than for other widely available algorithms. This allows it to be useful out of the box, without human adjustment for a larger number of time series data points. As more data accumulates in a particular time series, the forecasts in Amazon Quick Sight can adjust to data drifts and changes of pattern. For time series that show trends, trend detection is performed first to make the series stationary. The forecast of that stationary sequence is projected back with the trend.

Because the algorithm relies on an efficient online algorithm (RCF), it can support interactive "what-if" queries. In these, some of the forecasts can be altered and treated as hypotheticals to provide conditional forecasts. This is the origin of the ability to explore "what-if" scenarios during analysis.

References for machine learning and RCF

To learn more about machine learning and this algorithm, we suggest the following resources:

- The article [Robust Random Cut Forest \(RRCF\): A No Math Explanation](#) provides a lucid explanation without the mathematical equations.

- The book [The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition \(Springer Series in Statistics\)](#) provides a thorough foundation on machine learning.
- [Random Cut Forest Based Anomaly Detection On Streams](#), a scholarly paper that dives deep into the technicalities of both anomaly detection and forecasting, with examples.

A different approach to RCF appears in other Amazon services. If you want to explore how RCF is used in other services, see the following:

- *Amazon Managed Service for Apache Flink SQL Reference*: [RANDOM_CUT_FOREST](#) and [RANDOM_CUT_FOREST_WITH_EXPLANATION](#)
- *Amazon SageMaker Developer Guide*: [Random Cut Forest \(RCF\) Algorithm](#). This approach is also explained in [The Random Cut Forest Algorithm](#), a chapter in [Machine Learning for Business](#) (October 2018).

Dataset requirements for using ML insights with Amazon Quick Sight

To begin using the machine learning capabilities of Amazon Quick Sight, you need to connect to or import your data. You can use an existing Amazon Quick Sight dataset or create a new one. You can directly query your SQL-compatible source, or ingest the data into SPICE.

The data must have the following properties:

- At least one metric (for example, sales, orders, shipped units, sign ups, and so on).
- At least one category dimension (for example, product category, channel, segment, industry, and so on). Categories with NULL values are ignored.
- Anomaly detection requires a minimum of 15 data points for training. For example, if the grain of your data is daily, you need at least 15 days of data. If the grain is monthly, you need at least 15 months of data.
- Forecasting work best with more data. Make sure that your dataset has enough historical data for optimal results. For example, if the grain of your data is daily, you need at least 38 days of data. If the grain is monthly, you need at least 43 months of data. Following are the requirements for each time grain:
 - Years: 32 data points
 - Quarters: 35 data points
 - Months: 43 data points

- Weeks: 35 data points
- Days: 38 data points
- Hours: 39 data points
- Minutes: 46 data points
- Seconds: 46 data points
- If you want to analyze anomalies or forecasts, you also need at least one date dimension.

If you don't have a dataset to get started, you can download this sample dataset: [ML Insights Sample Dataset VI](#). After you have a dataset ready, create a new analysis from the dataset.

Working with insights in Amazon Quick Sight

In Amazon Quick Sight, you can add ready-to-use analytical computations to your analysis as widgets. You can work with insights in two ways:

- **Suggested insights**

Amazon Quick Sight creates a list of suggested insights based on its interpretation of the data you put into your visuals. The list changes based on context. In other words, you can see different suggestions depending on what fields you add to your visual and what type of visual you choose. For example, if you have a time-series visualization, your insights might include period-over-period changes, anomalies, and forecasts. As you add more visualizations to your analysis, you generate more suggested insights.

- **Custom insights**

Custom insights enable you to create your own computation, using your own words to give context to the fields that appear in the widget. When you create a custom insight, you add it to the analysis, and then choose what type of calculation that you want to use. Then, you can add text and formatting to make it look how you want. You can also add more fields, calculations, and parameters.

You can add any combination of suggested and custom insights to your analysis, to create the decision-making environment that best serves your purposes.

Topics

- [Adding suggested insights](#)

- [Adding custom insights to your analysis](#)

Adding suggested insights

Use the following procedure to add suggested insights to your analysis.

Before you begin, make sure that your dataset meets the criteria outlined in [Dataset requirements for using ML insights with Amazon Quick Sight](#).

1. Begin with an analysis that has a few fields added to a visual.
2. On the left, choose **Insights**. The **Insights** panel opens and displays a list of ready-to-use suggested insights.

Each visual also displays a small box on its top border to display how many insights are available for that visual. You can choose this box to open the **Insights** panel, and it opens to whatever view you most recently had open.

Scroll down to preview more insights.

The insights that appear are controlled by the data type of the fields you choose to include in your visual. This list is generated each time you change your visual. If you make changes, check **Insights** to see what is new. To get a specific insight, see [Adding custom insights to your analysis](#).

3. (Optional) Open the context menu with more options for one of the insights. To do this, choose the ellipses on the top right of the insight (...).

The options are different for each type of insight. The options that you can interact with include the following:

- **Change the time series aggregation** – To year, quarter, month, week, day, hour, or minute.
 - **Analyze contributions to metrics** – Choose contributors and a time frame to analyze.
 - **Show all anomalies** – Browse anomalies in this time frame.
 - **Edit forecast** – Choose forecast length, prediction interval, and seasonality.
 - **Focus on** or **Exclude** – Zoom in or zoom out on your dimensional data.
 - **Show details** – View more information about a recent anomaly (outlier).
 - Provide feedback on the usefulness of the insight in your analysis.
4. Add a suggested insight to your analysis by choosing the plus sign (+) near the insight title.

5. (Optional) After you add an insight to your analysis, customize the narrative that you want it to display. To do this, choose the **v**-shaped on-visual menu, then choose **Customize narrative**. For more information, see [Creating autonarratives with Amazon Quick Sight](#).

If your insight is for anomalies (outliers), you can also change the settings for the anomaly detection job. To do this, choose **Configure anomaly**. For more information, see [Setting up ML-powered anomaly detection for outlier analysis](#).

6. (Optional) To remove the insight from your analysis, choose the **v**-shaped on-visual menu at the top right of the visual. Then choose **Delete**.

Adding custom insights to your analysis

If you don't want to use any of the suggested insights, you can create your own custom insight. Use the following procedure to create a custom computational insight.

1. Start with an existing analysis. On the top menu bar, choose **Add+**. Then choose **Add Insight**.

A container for the new insight is added to the analysis.

2. Do one of the following:

- Choose the computation that you want to use from the list. As you choose each item, an example of that insight's output displays. When you find the one that you want to use, choose **Select**.
- Exit this screen and customize the insight manually. An unconfigured insight has a **Customize insight** button. Choose the button to open the **Configure narrative** screen. For more information on using the expression editor, see [Creating autonarratives with Amazon Quick Sight](#).

Because you are initiating the creation of the insight, it's not based on an existing visual. When the insight is added to the analysis, it displays a note showing what kind of data it needs to complete your request. For example, it might ask for **1 dimension in Time**. In this case, you add a dimension to the **Time** field well.

3. After you have the correct data, follow any remaining screen prompts to finish creating the custom insight.
4. (Optional) To remove the insight from your analysis, choose the **v**-shaped on-visual menu at the top right of the visual. Then choose **Delete**.

Creating autonarratives with Amazon Quick Sight

An *autonarrative* is a natural-language summary widget that displays descriptive text instead of charts. You can embed these widgets throughout your analysis to highlight key insights and callouts. You don't have to sift through the visual, drilling down, comparing values, and rechecking ideas to extract a conclusion. You also don't have to try to understand what the data means, or discuss different interpretations with your colleagues. Instead, you can extrapolate the conclusion from the data, and display it in the analysis, stated plainly. A single interpretation can be shared by everyone.

Amazon Quick Sight automatically interprets the charts and tables in your dashboard and provides a number of suggested insights in natural language. The suggested insights that you can choose from are ready-made and come with words, calculations, and functions. But you can change them if you want to. You can also design your own. As the author of the dashboard, you have complete flexibility to customize the computations and language for your needs. You can use narratives to effectively tell the story of your data in plain language.

Note

Narratives are separate from machine learning. They only use ML if you add forecast or anomaly (outlier) computations to them.

Topics

- [Insights that include autonarratives](#)
- [Use the narrative expression editor](#)
- [The expression editor workspace](#)
- [Adding URLs](#)
- [Working with autonarrative computations](#)

Insights that include autonarratives

When you are adding an insight, also known as an autonarrative, to your analysis, you can choose from the following templates. In the following list, they are defined by example. Each definition includes a list of the minimum required fields for the autonarrative to work. If you are using only the suggested insights on the **Insights** tab, choose the appropriate fields to get an insight to show up in the suggested insights list.

For more information on customizing autonarratives, see [Working with autonarrative computations](#).

- **Bottom ranked** – For example, the bottom three states by sales revenue. Requires that you have at least one dimension in the **Categories** field well.
- **Bottom movers** – For example, the bottom three products sold, by sales revenue. Requires that you have at least one dimension in the **Time** field well and at least one dimension in the **Categories** field well.
- **Forecast** (*ML-powered insight*) – For example, "Total sales are forecasted to be \$58,613 for Jan 2016." Requires that you have at least one dimension in the **Time** field well.
- **Growth rate** – For example, "The 3-month compounded growth rate for sales is 22.23%." Requires that you have at least one dimension in the **Time** field well.
- **Maximum** – For example, "Highest month is Nov 2014 with sales of \$112,326." Requires that you have at least one dimension in the **Time** field well.
- **Metric comparison** – For example, "Total sales for Dec 2014 is \$90,474, 10% higher than target of \$81,426." Requires that you have at least one dimension in the **Time** field well and at least two measures in the **Values** field well.
- **Minimum** – For example, "Lowest month is Feb 2011 with sales of \$4,810." Requires that you have at least one dimension in the **Time** field well.
- **Anomaly detection** (*ML-powered insight*) – For example, top three outliers and their contributing drivers for total sales on January 3, 2019. Requires that you have at least one dimension in the **Time** field well, at least one measure in the **Values** field well, and at least one dimension in the **Categories** field well.
- **Period over period** – For example, "Total sales for Nov 2014 increased by 44.39% (\$34,532) from \$77,793 to \$112,326." Requires that you have at least one dimension in the **Time** field well.
- **Period to date** – For example, "Year-to-date sales for Nov 30, 2014 increased by 25.87% (\$132,236) from \$511,236 to \$643,472." Requires that you have at least one dimension in the **Time** field well.
- **Top ranked** – For example, top three states by sales revenue. Requires that you have at least one dimension in the **Categories** field well.
- **Top movers** – For example, top products by sales revenue for November 2014. Requires that you have at least one dimension in the **Time** field well and at least one dimension in the **Categories** field well.

- **Total aggregation** – For example, "Total revenue is \$2,297,200." Requires that you have at least one dimension in the **Time** field well and at least one measure in the **Values** field well.
- **Unique values** – For example, "There are 793 unique values in Customer_IDs." Requires that you have at least one dimension in the **Categories** field well.

Use the narrative expression editor

The following walkthrough shows an example of how to customize a narrative. For this example, we use a period over period computation type.

1. Begin with an existing analysis. Add a **period over period** insight to it. The easiest way to do this is to choose the + icon, then **Add insight**, then choose a type of insight from the list. To learn what type of computational insights you can add as autonarratives, see [Insights that include autonarratives](#).

After you choose a type of insight, choose **Select** to create the widget. To create an empty narrative, close this screen without choosing a template. To follow this example, choose **Period over period**.

If you had a visual selected when you added the insight, the field wells have preconfigured fields for the date, metric, and category. These come from the visualization that you chose when you created the insight. You can customize the fields as needed.

You can only customize a narrative for a new or existing insight (text-based) widget. You can't add one to an existing visual (chart based), because it's a different type of widget.

2. Edit the narrative in the expressions editor by choosing the on-visual menu, then choosing **Customize narrative**.

In this context, **Computations** are predefined calculations (period-over-period, period-to-date, growth rate, max, min, top movers, and so on) that you can reference in your template to describe your data. Currently, Amazon Quick Sight supports 13 different types of computations that you can add to your insight. In this example, **PeriodOverPeriod** is added by default because we chose the **Period Over Period** template from the suggested insights panel.

3. Choose **Add computation** at bottom right to add a new computation, and then choose one from the list. For this walkthrough, choose **Growth rate**, and then choose **Next**.
4. Configure the computation by choosing the number of periods that you want to compute over. The default is four, and that works for our example. Optionally, you can change the

name of the computation at the top of the screen. However, for our purposes, leave the name unchanged.

 Note

The computation names that you create are unique within the insight. You can reference multiple computations of the same type in your narrative template. For example, suppose that you have two metrics, sales revenue and units sold. You can create growth rate computations for each metric if they have different names. However, anomaly computations aren't compatible with any other computation type in the same widget. Anomaly detection must exist in an insight by itself. To use other computations in the same analysis, put them into insights separate from anomalies.

To proceed, choose **Add**.

5. Expand **Computations** on the right. The computations that are part of the narrative display in the list. In this case, it's **PeriodOverPeriod** and **GrowthRate**.
6. In the workspace, add the following text after the final period: **Compounded growth rate for the last**, then add a space.
7. Next, to add the computation leave your cursor after the space after the word **last**. On the right, under **GrowthRate**, choose the expression named **timePeriods** (click only once to add it).

Doing this inserts the expression **GrowthRate.timePeriods**, which is the number of periods you set in the configuration for **GrowthRate**.

8. Complete the sentence with **days is** (a space before and afterwards), and add the expression **GrowthRate.compoundedGrowthRate.formattedValue**, followed by a period (.). Choose the expression from the list, rather than typing it in. However, you can edit the contents of the expression after you add it.

 Note

The **formattedValue** expression returns a string that is formatted based on the formatting applied for the metric on the field. To perform metric math, use **value** instead, which returns the raw value as an integer or decimal.

9. Add a conditional statement and formatting. Place your cursor at the end of the template, after the `formattedValue` expression. Add a space if necessary. On the **Edit narrative** menu bar, choose **Insert code**, and then choose **Inline IF** from the list. An expression block opens.
10. With the expression block open, choose **GrowthRate**, **compoundedGrowthRate**, **value** from the expression list. Enter `>0` at the end of the expression. Choose **Save**. Don't move your cursor yet.

A prompt appears for the conditional content; enter **better than expected!** Then select the text you just entered, and use the formatting toolbar at the top to turn it green and bold.
11. Add another expression block for the case when the growth rate wasn't that great by repeating the previous step. But this time, make it `<0` and enter the text **worse than expected**. Make it red instead of green.
12. Choose **Save**. The customized narrative that we just created should look similar to the following.

The expression editor provides you with a sophisticated tool to customize your narratives. You can also reference the parameters you create for your analysis or dashboard, and use a set of built-in functions for further customization.

Tip

To create an empty narrative, add an insight using the **+** icon and then **Add insights**. But instead of choosing a template, simply close the screen.

The best way to get started with customizing narratives is to use the existing templates to learn the syntax.

The expression editor workspace

Use the expression editor to customize a narrative to best fit your business needs. The information below provides an overview of the expression editor workspace and lists all menu options that can be configured for your narrative. For a walkthrough that shows you how to create a custom narrative, see [Use the narrative expression editor](#).

On the right side of the screen, there's a list of items that you can add to the narrative:

- **Computations** – Use this to choose from the computations that are available in this insight. You can expand this list.
- **Parameters** – Use this to choose from the parameters that exist in your analysis. You can expand this list.
- **Functions** – Use this to choose from functions that you can add to a narrative. You can expand this list.
- **Add computation** – Use this button to create another computation. New computations appear in the **Computations** list, ready to add to the insight.

At the bottom of the narrative expression editor, there's a preview of the narrative that updates as you work. This area also shows an alert if you introduce an error into the narrative or if the narrative is empty. To see a preview of ML-powered insights like anomaly detection or forecasting, run your insight calculation at least once before customizing the narrative.

Editing tools are located across the top of the screen. They offer the following options:

- **Insert code** – You can insert the following code blocks from this menu:
 - **Expressions** – Add a free-form expression.
 - **Inline IF** – Add an IF statement that displays inline with the existing block of text.
 - **Inline FOR** – Add a FOR statement that displays inline with the existing block of text.
 - **Block IF** – Add an IF statement that displays in a separate block of text.
 - **Block FOR** – Add a FOR statement that displays in a separate block of text.

The IF and FOR statements enable you to create content that is conditionally formatted. For example, you might add a **block IF** statement, then configure it to compare an integer to a value from a calculation. To do this, you use the following steps, also demonstrated in [Use the narrative expression editor](#):

1. Open the calculations menu at right, and choose one of the blue highlighted items from one of the calculations. Doing this adds the item to the narrative.
2. Click once on the item to open it.
3. Enter the comparison that you want to make. The expression looks something like this:
`PeriodOverPeriod.currentMetricValue.value>0.`
4. Save this expression in the pop-up editor, which prompts you for **Conditional content**.
5. Enter what you want to display in the insight, and format it as you want it to appear. Or if you prefer, you can add an image or a URL—or add a URL to an image.

- **Paragraph** – This menu offers options for changes to the font size:
 - **H1 Large header**
 - H2 Header
 - H3 Small header
 - ¶1 Large paragraph
 - ¶2 Paragraph
 - ¶3 Small paragraph
- **Font** – Use this menu tray to choose options for text formatting. These include bold, italic, underline, strikethrough, foreground color of the text (the letters themselves), and background color of the text. Choose the icon to turn on an option; choose it again to toggle the option off.
- **Formatting** – Use this menu tray to choose options for paragraph formatting, including bulleted list, left justify, center, and right justify. Choose the icon to turn on an option, choose it again to toggle the option off.
- **Image** – Use this icon add an image URL. The image displays in your insight, provided the link is accessible. You can resize images. To display an image based on a condition, put the image inside an IF block.
- **URL** – Use this icon to add a static or dynamic URL. You can also add URLs to images. For example, you can add traffic light indicator images to an insight for an executive dashboard, with links to a new sheet for red, amber, and green conditions.

Adding URLs

Using the **URL** button on the editing menu of the narrative expression editor, you can add static and dynamic URLs (hyperlinks) into a narrative. You can also use the following keyboard shortcuts: ⌘+⬆+L or Ctrl+⬆+L.

A static URL is a link that doesn't change; it always opens the same URL. A dynamic URL is a link that changes based on the expressions or parameters that you provide when you set it up. It's built with dynamically evaluated expressions or parameters.

Following are of examples of when you might add a static link in your narrative:

- **In an IF statement, you might use the URL in the conditional content.** If you do and a metric fails to meet an expected value, your link might send the user to a wiki with a list of best practices to improve the metric.

- **You might use a static URL to create a link to another sheet in the same dashboard, by using the following steps:**

1. Go to the sheet that you want to make the link to.
2. Copy that sheet's URL.
3. Return to the narrative editor and create a link using the URL that you just copied.

Following are examples of when you might add a dynamic link in your narrative:

- **To search a website with a query, by using the following steps.**

1. Create a URL with the following link.

```
https://google.com?q=<<formatDate(now(), 'yyyy-MM-dd')>>
```

This link sends a query to Google with search text that is the evaluated value of the following.

```
formatDate(now(), 'yyyy-MM-dd')
```

If the value of `now()` is 02/02/2020, then the link on your narrative contains `https://google.com?q=2020-02-02`.

- **To create a link that updates a parameter.** To do this, create or edit a link and set the URL to the current dashboard or analysis URL. Then add the expression that sets the parameter value to at the end, for example `#p.myParameter=12345`.

Suppose that the following is the dashboard link that you start with.

```
https://us-east-1.quicksight.aws.amazon.com/sn/analyses/00000000-1111-2222-3333-44444444
```

If you add a parameter value assignment to it, it looks like the following.

```
https://us-east-1.quicksight.aws.amazon.com/sn/analyses/00000000-1111-2222-3333-44444444#p.myParameter=12345
```

For more information on parameters in URLs, see [Using parameters in a URL](#).

Working with autonarrative computations

Use this section to help you understand what functions are available to you when you are customizing an autonarrative. You only need to customize a narrative if you want to change or build on the default computation.

After you create an autonarrative, the expression editor opens. You can also activate the expression editor by choosing the on-visual menu, and then **Customize Narrative**. To add a computation while using the expression editor, choose **+ Add computation**.

You can use the following code expression to build your autonarrative. These are available from the list that's labeled **Insert code**. Code statements can display inline (in a sentence) or as a block (in a list).

- Expression – Create your own code expression.
- IF – An IF statement that includes an expression after evaluating a condition.
- FOR – A FOR statement that loops through values.

You can use the following computations to build your autonarrative. You can use the expression editor without editing any syntax, but you can also customize it if you want to. To interact with the syntax, open the computational widget in the autonarrative expression editor.

Topics

- [ML-powered anomaly detection for outliers](#)
- [Bottom movers computation](#)
- [Bottom ranked computation](#)
- [ML-powered forecasting](#)
- [Growth rate computation](#)
- [Maximum computation](#)
- [Metric comparison computation](#)
- [Minimum computation](#)
- [Period over period computation](#)
- [Period to date computation](#)
- [Top movers computation](#)
- [Top ranked computation](#)

- [Total aggregation computation](#)
- [Unique values computation](#)

ML-powered anomaly detection for outliers

The ML-powered anomaly detection computation searches your data for outliers. For example, you can detect the top three outliers for total sales on January 3, 2019. If you enable contribution analysis, you can also detect the key drivers for each outlier.

To use this function, you need at least one dimension in the **Time** field well, at least one measure in the **Values** field well, and at least one dimension in the **Categories** field well. The configuration screen provides an option to analyze the contribution of other fields as key drivers, even if those fields aren't in the field wells.

For more information, see [Detecting outliers with ML-powered anomaly detection](#).

Note

You can't add ML-powered anomaly detection to another computation, and you can't add another computation to an anomaly detection.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name that you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. You can use items displayed in **bold monospace font** following in the narrative.

- `timeField` – From the **Time** field well.
 - `name` – The formatted display name of the field.
 - `timeGranularity` – The time field granularity (**DAY**, **YEAR**, and so on).
- `categoryFields` – From the **Categories** field well.
 - `name` – The formatted display name of the field.

- **metricField** – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- **itemsCount** – The number of items included in this computation.
- **items** – Anomalous items.
 - **timeValue** – The values in the date dimension.
 - **value** – The date/time field at the point of the anomaly (outlier).
 - **formattedValue** – The formatted value in the date/time field at the point of the anomaly.
 - **categoryName** – The actual name of the category (cat1, cat2, and so on).
 - **direction** – The direction on the x-axis or y-axis that's identified as anomalous: HIGH or LOW. HIGH means "higher than expected." LOW means "lower than expected."

When iterating on items, `AnomalyDetection.items[index].direction` can contain either HIGH or LOW. For example, `AnomalyDetection.items[index].direction='HIGH'` or `AnomalyDetection.items[index].direction=LOW`.

`AnomalyDetection.direction` can have an empty string for ALL. An example is `AnomalyDetection.direction=' '`.

- **actualValue** – The metric's actual value at the point of the anomaly or outlier.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- **expectedValue** – The metric's expected value at the point of the anomaly (outlier).
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.

Bottom movers computation

The bottom movers computation counts the requested number of categories by date that rank in the bottom of the autonarrative's dataset. For example, you can create a computation to find the bottom three products sold, by sales revenue.

To use this function, at least one dimension in the **Time** field well and at least one dimension in the **Categories** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Date

The date dimension that you want to rank.

Category

The category dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Number of movers

The number of ranked results that you want to display.

Order by

The order that you want to use, percent difference or absolute difference.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

Note

These are the same output parameters as the ones that are returned by the top movers computation.

- **timeField** – From the **Time** field well.
 - **name** – The formatted display name of the field.
 - **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- **categoryField** – From the **Categories** field well.
 - **name** – The formatted display name of the field.
- **metricField** – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- **startTimeValue** – The value in the date dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.
- **endTimeValue** – The value in the date dimension.
 - **value** – The raw value.
 - **formattedValue** – The absolute value formatted by the datetime field.
- **itemsCount** – The number of items included in this computation.
- **items**: Bottom moving items.
 - **categoryField** – The category field.
 - **value** – The value (contents) of the category field.
 - **formattedValue** – The formatted value (contents) of the category field. If the field is null, this displays 'NULL'. If the field is empty, it displays '(empty)'.
 - **currentMetricValue** – The current value for the metric field.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
 - **previousMetricValue** – The previous value for the metric field.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
 - **percentDifference** – The percent difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the percent difference.

- **formattedValue** – The formatted value of the percent difference (for example, -42%).
- **formattedAbsoluteValue** – The formatted absolute value of the percent difference (for example, 42%).
- **absoluteDifference** – The absolute difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the absolute difference.
 - **formattedValue** – The absolute difference formatted by the settings in the metric field's format preferences.
 - **formattedAbsoluteValue** – The absolute value of the difference formatted by the metric field.

Bottom ranked computation

The bottom ranked computation calculates the requested number of categories by value that rank in the bottom of the autonarrative's dataset. For example, you can create a computation to find the bottom three states by sales revenue.

To use this function, you need at least one dimension in the **Categories** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Category

The category dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Number of results

The number of ranked results that you want to display.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

Note

These are the same output parameters as the ones that are returned by the top ranked computation.

- **categoryField** – From the **Categories** field well.
 - **name** – The formatted display name of the field.
- **metricField** – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- **itemsCount** – The number of items included in this computation.
- **items**: Bottom ranked items.
 - **categoryField** – The category field.
 - **value** – The value (contents) of the category field.
 - **formattedValue** – The formatted value (contents) of the category field. If the field is null, this displays 'NULL'. If the field is empty, it displays '(empty)'.
 - **metricValue** – The metric field.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.

Example

The following screenshot shows the default configuration for the bottom-ranked computation.

ML-powered forecasting

The ML-powered forecast computation forecasts future metrics based on patterns of previous metrics by seasonality. For example, you can create a computation to forecast total revenue for the next six months.

To use this function, you need at least one dimension in the **Time** field well.

For more information about working with forecasts, see [Forecasting and creating what-if scenarios with Amazon Quick Sight](#).

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Date

The date dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Periods forward

The number of time periods in the future that you want to forecast. Ranges from 1 to 1,000.

Periods backward

The number of time periods in the past that you want to base your forecast on. Ranges from 0 to 1,000.

Seasonality

The number of seasons included in the calendar year. The default setting, **automatic** detects this for you. Ranges from 1 to 180.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

- `timeField` – From the **Time** field well.
 - **name** – The formatted display name of the field.
 - **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- `metricField` – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- `metricValue` – The value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- `timeValue` – The value in the date dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the date field.
- **relativePeriodsToForecast** – The relative number of periods between latest datetime record and last forecast record.

Growth rate computation

The growth rate computation compares values over time periods. For example, you can create a computation to find the three-month compounded growth rate for sales, expressed as a percentage.

To use this function, you need at least one dimension in the **Time** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Date

The date dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Number of periods

The number of time periods in the future that you want to use to compute the growth rate.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

- `timeField` – From the **Time** field well.
 - **name** – The formatted display name of the field.
 - **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- `metricField` – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- `previousMetricValue` – The previous value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.

- **previousTimeValue** – The previous value in the datetime dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.
- **compoundedGrowthRate** – The percent difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the percent difference.
 - **formattedValue** – The formatted value of the percent difference (for example, -42%).
 - **formattedAbsoluteValue** – The formatted absolute value of the percent difference (for example, 42%).
- **absoluteDifference** – The absolute difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the absolute difference.
 - **formattedValue** – The absolute difference formatted by the settings in the metric field's format preferences.
 - **formattedAbsoluteValue** – The absolute value of the difference formatted by the metric field.

Maximum computation

The maximum computation finds the maximum dimension by value. For example, you can create a computation to find the month with the highest revenue.

To use this function, you need at least one dimension in the **Time** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Date

The date dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

Note

These are the same output parameters as the ones that are returned by the minimum computation.

- `timeField` – From the **Time** field well.
 - **name** – The formatted display name of the field.
 - **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- `metricField` – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- `metricValue` – The value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- `timeValue` – The value in the datetime dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.

Metric comparison computation

The metric comparison computation compares values in different measures. For example, you can create a computation to compare two values, such as actual sales compared to sales goals.

To use this function, you need at least one dimension in the **Time** field well and at least two measures in the **Values** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Date

The date dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Target value

The field that you want to compare to the value.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

- `timeField` – From the **Time** field well.
 - **name** – The formatted display name of the field.
 - **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- `fromMetricField` – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- `fromMetricValue` – The value in the metric dimension.

- **value** – The raw value.
- **formattedValue** – The value formatted by the metric field.
- **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- toMetricField – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- toMetricValue – The current value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- timeValue – The value in the datetime dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.
- percentDifference – The percent difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the percent difference.
 - **formattedValue** – The formatted value of the percent difference (for example, -42%).
 - **formattedAbsoluteValue** – The formatted absolute value of the percent difference (for example, 42%).
- absoluteDifference – The absolute difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the absolute difference.
 - **formattedValue** – The absolute difference formatted by the settings in the metric field's format preferences.
 - **formattedAbsoluteValue** – The absolute value of the difference formatted by the metric field.

Minimum computation

The minimum computation finds the minimum dimension by value. For example, you can create a computation to find the month with the lowest revenue.

To use this function, you need at least one dimension in the **Time** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Date

The date dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

Note

These are the same output parameters as the ones that are returned by the maximum computation.

- `timeField` – From the **Time** field well.
 - **name** – The formatted display name of the field.
 - **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- `metricField` – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- `metricValue` – The value in the metric dimension.

- **value** – The raw value.
- **formattedValue** – The value formatted by the metric field.
- **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- **timeValue** – The value in the datetime dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.

Period over period computation

The period over period computation compares values from two different time periods. For example, you can create a computation to find out how much sales increased or decreased since the previous time period.

To use this function, you need at least one dimension in the **Time** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Date

The date dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

- **timeField** – From the **Time** field well.
 - **name** – The formatted display name of the field.
 - **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- **metricField** – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- **previousMetricValue** – The previous value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- **previousTimeValue** – The previous value in the datetime dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.
- **currentMetricValue** – The current value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- **currentTimeValue** – The current value in the datetime dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.
- **percentDifference** – The percent difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the percent difference.
 - **formattedValue** – The formatted value of the percent difference (for example, -42%).
 - **formattedAbsoluteValue** – The formatted absolute value of the percent difference (for example, 42%).
- **absoluteDifference** – The absolute difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the absolute difference.
 - **formattedValue** – The absolute difference formatted by the settings in the metric field's format preferences.

- **formattedAbsoluteValue** – The absolute value of the difference formatted by the metric field.

Example

To create a Period over period computation

1. In the analysis that you want to change, choose **Add insight**.
2. For **Computation type**, choose **Period over period**, and then choose **Select**.
3. In the new insight that you created, add the time dimension and value dimension fields that you want to compare. In the screenshot below, **Order Date** and **Sales (Sum)** are added to the insight. With these two fields selected, Quick Sight shows the year to date sales of the latest month and the percentage difference compared with the previous month.
4. (Optional) To further customize the insight, open the on-visual menu and choose **Customize narrative**. In the **Edit narrative** window that appears, drag and drop the fields that you need from the **Computations** list, and then choose **Save**.

Period to date computation

The period to date computation evaluates values for a specified period to date. For example, you can create a computation to find out how much you've earned in year-to-date sales.

To use this function, you need at least one dimension in the **Time** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Date

The date dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Time granularity

The date granularity that you want to use for the computation, for example year to date.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

- `timeField` – From the **Time** field well.
 - **name** – The formatted display name of the field.
 - **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- `metricField` – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- `previousMetricValue` – The previous value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- `previousTimeValue` – The previous value in the datetime dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.
- `currentMetricValue` – The current value in the metric dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
- `currentTimeValue` – The current value in the datetime dimension.
 - **value** – The raw value.

- **formattedValue** – The value formatted by the datetime field.
- **periodGranularity** – The period granularity for this computation (**MONTH**, **YEAR**, and so on).
- **percentDifference** – The percent difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the percent difference.
 - **formattedValue** – The formatted value of the percent difference (for example, -42%).
 - **formattedAbsoluteValue** – The formatted absolute value of the percent difference (for example, 42%).
- **absoluteDifference** – The absolute difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the absolute difference.
 - **formattedValue** – The absolute difference formatted by the settings in the metric field's format preferences.
 - **formattedAbsoluteValue** – The absolute value of the difference formatted by the metric field.

Example

To create a Period to date computation

1. In the analysis that you want to change, choose **Add insight**.
2. For **Computation type**, choose **Period to date**, and then choose **Select**.
3. In the new insight that you created, add the time dimension and value dimension fields that you want to compare. In the screenshot below, **Order Date** and **Sales (Sum)** are added to the insight. With these two fields selected, Quick Sight shows the year to date sales of the latest month and the percentage difference compared with the previous month.
4. (Optional) To further customize the insight, open the on-visual menu and choose **Customize narrative**. In the **Edit narrative** window that appears, drag and drop the fields that you need from the **Computations** list, and then choose **Save**.

Top movers computation

The top movers computation counts the requested number of categories by date that rank in the top of the autonarrative's dataset. For example, you can create a computation to find the top products by sales revenue for a time period.

To use this function, you need at least one dimension in the **Time** field well and at least one dimension in the **Categories** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Category

The category dimension you want to rank.

Value

The aggregated measure that the computation is based on.

Number of results

The number of top ranking items you want to find.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

Note

These are the same output parameters as the ones that are returned by the bottom movers computation.

- **timeField** – From the **Time** field well.
 - **name** – The formatted display name of the field.
 - **timeGranularity** – The time field granularity (**DAY**, **YEAR**, and so on).
- **categoryField** – From the **Categories** field well.
 - **name** – The formatted display name of the field.
- **metricField** – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- **startTimeValue** – The value in the date dimension.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the datetime field.
- **endTimeValue** – The value in the date dimension.
 - **value** – The raw value.
 - **formattedValue** – The absolute value formatted by the datetime field.
- **itemsCount** – The number of items included in this computation.
- **items**: Top moving items.
 - **categoryField** – The category field.
 - **value** – The value (contents) of the category field.
 - **formattedValue** – The formatted value (contents) of the category field. If the field is null, this displays 'NULL'. If the field is empty, it displays '(empty)'.
 - **currentMetricValue** – The current value for the metric field.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
 - **previousMetricValue** – The previous value for the metric field.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.
 - **percentDifference** – The percent difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the percent difference.

- **formattedValue** – The formatted value of the percent difference (for example, -42%).
- **formattedAbsoluteValue** – The formatted absolute value of the percent difference (for example, 42%).
- **absoluteDifference** – The absolute difference between the current and previous values of the metric field.
 - **value** – The raw value of the calculation of the absolute difference.
 - **formattedValue** – The absolute difference formatted by the settings in the metric field's format preferences.
 - **formattedAbsoluteValue** – The absolute value of the difference formatted by the metric field.

Top ranked computation

The top ranked computation finds the top ranking dimensions by value. For example, you can create a computation to find the top three states by sales revenue.

To use this function, you need at least one dimension in the **Categories** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Category

The category dimension that you want to rank.

Value

The aggregated measure that the computation is based on.

Number of results

The number of top ranking items that you want to find.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

 Note

These are the same output parameters as the ones that are returned by the bottom ranked computation.

- **categoryField** – From the **Categories** field well.
 - **name** – The formatted display name of the field.
- **metricField** – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- **itemsCount** – The number of items included in this computation.
- **items**: Top ranked items.
 - **categoryField** – The category field.
 - **value** – The value (contents) of the category field.
 - **formattedValue** – The formatted value (contents) of the category field. If the field is null, this displays 'NULL'. If the field is empty, it displays '(empty)'.
 - **metricValue** – The metric field.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.

Total aggregation computation

The total aggregation computation creates a grand total of the value. For example, you can create a computation to find the total revenue.

To use this function, you need at least one dimension in the **Time** field well and at least one measure in the **Values** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Value

The aggregated measure that the computation is based on.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

- categoryField – The category field.
 - **name** – The display name of the category field.
- metricField – From the **Values** field well.
 - **name** – The formatted display name of the field.
 - **aggregationFunction** – The aggregation used for the metric (**SUM**, **AVG**, and so on).
- totalAggregate – The total value of the metric aggregation.
 - **value** – The raw value.
 - **formattedValue** – The value formatted by the metric field.
 - **formattedAbsoluteValue** – The absolute value formatted by the metric field.

Unique values computation

The unique values computation counts the unique values in a category field. For example, you can create a computation to count the number of unique values in a dimension, such as how many customers you have

To use this function, you need at least one dimension in the **Categories** field well.

Parameters

name

A unique descriptive name that you assign or change. A name is assigned if you don't create your own. You can edit this later.

Category

The category dimension that you want to rank.

Computation outputs

Each function generates a set of output parameters. You can add these outputs to the autonarrative to customize what it displays. You can also add your own custom text.

To locate the output parameters, open the **Computations** tab on the right, and locate the computation that you want to use. The names of the computations come from the name you provide when you create the insight. Choose the output parameter by clicking on it only once. If you click twice, you add the same output twice. Items displayed in **bold** can be used in the narrative.

- `categoryField` – The category field.
 - **name** – The display name of the category field.
- **uniqueGroupValuesCount** – The number of unique values included in this computation.

Detecting outliers with ML-powered anomaly detection

Amazon Quick Sight uses proven Amazon technology to continuously run ML-powered anomaly detection across millions of metrics to discover hidden trends and outliers in your data. This tool allows you to get deep insights that are often buried in the aggregates and not scalable with manual analysis. With ML-powered anomaly detection, you can find outliers in your data without the need for manual analysis, custom development, or ML domain expertise.

Amazon Quick Sight notifies you in your visuals if it detects that you can analyze an anomaly or do some forecasting on your data.

Anomaly detection is not available in the eu-central-2 Europe (Zurich) region.

⚠ Important

ML-powered anomaly detection is a compute-intensive task. Before you start using it, you can get an idea of costs by analyzing the amount of data that you want to use. We offer a tiered pricing model that is based on the number of metrics you process per month.

Topics

- [Concepts for anomaly or outlier detection](#)
- [Setting up ML-powered anomaly detection for outlier analysis](#)
- [Exploring outliers and key drivers with ML-powered anomaly detection and contribution analysis](#)

Concepts for anomaly or outlier detection

Amazon Quick Sight uses the word *anomaly* to describe data points that fall outside an overall pattern of distribution. There are many other words for anomalies, which is a scientific term, including outliers, deviations, oddities, exceptions, irregularities, quirks, and many more. The term that you use might be based on the type of analysis you do, or the type of data you use, or even just the preference of your group. These outlying data points represent an entity—a person, place, thing, or time—which is exceptional in some way.

Humans easily recognize patterns and spot things that aren't like the others. Our senses provide this information for us. If the pattern is simple, and there is only a little data, you can easily make a graph to highlight the outliers in your data. Some simple examples include the following:

- A red balloon in a group of blue ones
- A racehorse that is far ahead of the others
- A kid who isn't paying attention during class
- A day when online orders are up, but shipping is down
- A person who got well, where others didn't

Some data points represent a significant event, and others represent a random occurrence. Analysis uncovers which data is worth investigating, based on what driving factors (key drivers) contributed to the event. Questions are essential to data analysis. Why did it happen? What's it related to? Did it happen only once or many times? What can you do to encourage or discourage more like it?

Understanding how and why a variation exists, and whether there is a pattern in the variations, requires more thought. Without the assistance of machine learning, each person might come to a different conclusion, because they have different experience and information. Therefore, each person might make a slightly different business decision. If there is a lot of data or variables to consider, it can require an overwhelming amount of analysis.

ML-powered anomaly detection identifies the causations and correlations to enable you to make data-driven decisions. You still have control over defining how you want the job to work on your data. You can specify your own parameters, and choose additional options, such as identifying key drivers in a contribution analysis. Or you can use the default settings. The following section walks you through the setup process, and provides explanations for the options available.

Setting up ML-powered anomaly detection for outlier analysis

Use procedures in the following sections to start detecting outliers, detecting anomalies, and identifying the key drivers that contribute to them.

Topics

- [Viewing anomaly and forecast notifications](#)
- [Adding an ML insight to detect outliers and key drivers](#)
- [Using contribution analysis for key drivers](#)

Viewing anomaly and forecast notifications

Amazon Quick Sight notifies you on a visual where it detects an anomaly, key drivers, or a forecasting opportunity. You can follow the prompts to set up anomaly detection or forecasting based on the data in that visual.

1. In an existing line chart, look for an insight notification in the menu on the visual widget.
2. Choose the lightbulb icon to display the notification.
3. If you want more information about the ML insight, you can follow the screen prompts to add an ML insight.

Adding an ML insight to detect outliers and key drivers

You can add an ML insight that detects *anomalies*, which are outliers that seem significant. To get started, you create for your insight a widget, also known as an *autonarrative*. As you configure your options, you can view a limited screenshot of your insight in the **Preview** pane at screen right.

In your insight widget, you can add up to five dimension fields that are not calculated fields. In the field wells, values for **Categories** represent the dimensional values that Amazon Quick Sight uses to split the metric. For example, let's say that you are analyzing revenue across all product categories and product SKUs. There are 10 product categories, each with 10 product SKUs. Amazon Quick Sight splits the metric by the 100 unique combinations and runs anomaly detection on each combination for the split.

The following procedure shows how to do this, and also how to add contribution analysis to detect the key drivers that are causing each anomaly. You can add contribution analysis later, as described in [Using contribution analysis for key drivers](#).

To set up outlier analysis, including key drivers

1. Open your analysis and in the toolbar, choose **Insights**, then **Add**. From the list, choose **Anomaly detection** and **Select**.
2. Follow the screen prompt on the new widget, which tells you to choose fields for the insight. Add at least one date, one measure, and one dimension.
3. Choose **Get started** on the widget. The configuration screen appears.
4. Under **Compute options**, choose values for the following options:
 - a. For **Combinations to be analysed**, choose one of the following options:

- i. **Hierarchical**

Choose this option if you want to analyze the fields hierarchically. For example, if you chose a date (T), a measure (N), and three dimension categories (C1, C2, and C3), Quick Sight analyses the fields hierarchically, as shown following.

T-N, T-C1-N, T-C1-C2-N, T-C1-C2-C3-N

- ii. **Exact**

Choose this option if you want to analyze only the exact combination of fields in the Category field well, as they are listed. For example, if you chose a date (T), a measure (N), and three dimension categories (C1, C2, and C3), Quick Sight analyses only the exact combination of category fields in the order they are listed, as shown following.

T-C1-C2-C3-N

iii. All

Choose this option if you want to analyze all field combinations in the Category field well. For example, if you chose a date (T), a measure (N), and three dimension categories (C1, C2, and C3), Quick Sight analyses all combinations of fields, as shown following.

T-N, T-C1-N, T-C1-C2-N, T-C1-C2-C3-N, T-C1-C3-N, T-C2-N, T-C2-C3-N, T-C3-N

If you chose a date and a measure only, Quick Sight analyses the fields by date and then by measure.

In the **Fields to be analyzed** section, you can see a list of fields from the field wells for reference.

- b. For **Name**, enter a descriptive alphanumeric name with no spaces, or choose the default value. This provides a name for the computation.

If you plan on editing the narrative that automatically displays on the widget, you can use the name to identify this widget's calculation. Customize the name if you plan to edit the autonarrative and if you have other similar calculations in your analysis.

5. In the **Display options** section, choose the following options to customize what is displayed in your insight widget. You can still explore all your results, no matter what you display.
 - a. **Maximum number of anomalies to show** – The number of outliers you want to display in the narrative widget.
 - b. **Severity** – The minimum level of severity for anomalies that you want to display in the insight widget.

A level of severity is a range of anomaly scores that is characterized by the lowest actual anomaly score included in the range. All anomalies that score higher are included in the range. If you set severity to **Low**, the insight displays all of the anomalies that rank between low and very high. If you set the severity to **Very high**, the insight displays only the anomalies that have the highest anomaly scores.

You can use the following options:

- **Very high**

- **High and above**
 - **Medium and above**
 - **Low and above**
- c. **Direction** – The direction on the x-axis or y-axis that you want to identify as anomalous. You can choose from the following:
- **Higher than expected** to identify higher values as anomalies.
 - **Lower than expected** to identify lower values as anomalies.
 - **[ALL]** to identify all anomalous values, high and low (default setting).
- d. **Delta** – Enter a custom value to use to identify anomalies. Any amount higher than the threshold value counts as an anomaly. The values here change how the insight works in your analysis. In this section, you can set the following:
- **Absolute value** – The actual value to use. For example, suppose this is 48. Amazon Quick Sight then identifies values as anomalous when the difference between a value and the expected value is greater than 48.
 - **Percentage** – The percentage threshold to use. For example, suppose this is 12.5%. Amazon Quick Sight then identifies values as anomalous when the difference between a value and the expected value is greater than 12.5%.
- e. **Sort by** – Choose a sort method for your results. Some methods are based on the anomaly score that Amazon Quick Sight generates. Amazon Quick Sight gives higher scores to data points that look anomalous. You can use any of the following options:
- **Weighted anomaly score** – The anomaly score multiplied by the log of the absolute value of the difference between the actual value and the expected value. This score is always a positive number.
 - **Anomaly score** – The actual anomaly score assigned to this data point.
 - **Weighted difference from expected value** – The anomaly score multiplied by the difference between the actual value and the expected value (default).
 - **Difference from expected value** – The actual difference between the actual value and the expected value (that is, actual–expected).
 - **Actual value** – The actual value with no formula applied.
6. In the **Schedule options** section, set the schedule for automatically running the insight recalculation. The schedule runs only for published dashboards. In the analysis, you can run it manually as needed. Scheduling includes the following settings:

- **Occurrence** – How often that you want the recalculation to run: every hour, every day, every week, or every month.
 - **Start schedule on** – The date and time to start running this schedule.
 - **Timezone** – The time zone that the schedule runs in. To view a list, delete the current entry.
7. In the **Top contributors** section, set Amazon Quick Sight to analyze the key drivers when an outlier (anomaly) is detected.

For example, Amazon Quick Sight can show the top customers that contributed to a spike in sales in the US for home improvement products. You can add up to four dimensions from your dataset. These include dimensions that you didn't add to the field wells of this insight widget.

For a list of dimensions available for contribution analysis, choose **Select fields**.

8. Choose **Save** to confirm your choices. Choose **Cancel** to exit without saving.
9. From the insight widget, choose **Run now** to run the anomaly detection and view your insight.

The amount of time that anomaly detection takes to complete varies depending on how many unique data points you are analyzing. The process can take a few minutes for a minimum number of points, or it can take many hours.

While it's running in the background, you can do other work in your analysis. Make sure to wait for it to complete before you change the configuration, edit the narrative, or open the **Explore anomalies** page for this insight.

The insight widget needs to run at least once before you can see results. If you think the status might be out of date, you can refresh the page. The insight can have the following states.

Appears on the Page	Status
Run now button	The job has not yet started.
Message about Analyzing for anomalies	The job is currently running.
Narrative about the detected anomalies (outliers)	The job has run successfully. The message says when this widget's calculation was last updated.

Appears on the Page	Status
Alert icon with an exclamation point (!)	This icon indicates there was an error during the last run. If the narrative also displays, you can still use Explore anomalies to use data from the previous successful run.

Using contribution analysis for key drivers

Amazon Quick Sight can identify the dimensions (categories) that contribute to outliers in measures (metrics) between two points in time. The key driver that contributes to an outlier helps you to answer the question: What happened to cause this anomaly?

If you are already using anomaly detection without contribution analysis, you can enable the existing ML insight to find key drivers. Use the following procedure to add contribution analysis and identify the key drivers behind outliers. Your insight for anomaly detection needs to include a time field and at least one aggregated metric (SUM, AVERAGE, or COUNT). You can include multiple categories (dimension fields) if you wish, but you can also run contribution analysis without specifying any category or dimension field.

You can also use this procedure to change or remove fields as key drivers in your anomaly detection.

To add contribution analysis to identify key drivers

1. Open your analysis and locate an existing ML insight for anomaly detection. Select the insight widget to highlight it.
2. Choose **Menu Options (...)** from the menu on the visual.
3. Choose **Configure anomaly** to edit the settings.
4. The **Contribution analysis (optional)** setting allows Amazon Quick Sight to analyze the key drivers when an outlier (anomaly) is detected. For example, Amazon Quick Sight can show you the top customers that contributed to a spike in sales in the US for home improvement products. You can add up to four dimensions from your dataset, including dimensions that you didn't add to the field wells of this insight widget.

To view a list of dimensions available for contribution analysis, choose **Select fields**.

If you want to change the fields you're using as key drivers, change the fields that are enabled in this list. If you disable all of them, Quick Sight won't perform any contribution analysis in this insight.

5. To save your changes, scroll to the bottom of the configuration options, and choose **Save**. To exit without saving, choose **Cancel**. To completely remove these settings, choose **Delete**.

Exploring outliers and key drivers with ML-powered anomaly detection and contribution analysis

You can interactively explore the anomalies (also known as outliers) in your analysis, along with the contributors (key drivers). The analysis is available for you to explore after the ML-powered anomaly detection runs. The changes you make in this screen aren't saved when you go back to the analysis.

To begin, choose **Explore anomalies** in the insight. The following screenshot shows the anomalies screen as it appears when you first open it. In this example, contributors analysis is set up and shows two key drivers.

The sections of the screen include the following, from top left to bottom right:

- **Contributors** displays key drivers. To see this section, you need to have contributors set up in your anomaly configuration.
- **Controls** contains settings for anomaly exploration.
- **Number of anomalies** displays outliers detected over time. You can hide or show this chart section.
- **Your field names** for category or dimension fields act as titles for charts that show anomalies for each category or dimension.

The following sections provide detailed information for each aspect of exploring anomalies.

Topics

- [Exploring contributors \(key drivers\)](#)
- [Setting controls for anomaly detection](#)
- [Showing and hiding anomalies by date](#)

- [Exploring anomalies per category or dimension](#)

Exploring contributors (key drivers)

If your anomaly insight is set up to detect key drivers, Quick Sight runs the contribution analysis to determine which categories (dimensions) are influencing the outliers. The **Contributors** section appears on the left.

Contributors contains the following sections:

- **Narrative** – At top left, a summary describes any changes in the metrics.
- **Top contributors configuration** – Choose **Configure** to change the contributors and the date range to use in this section.
- **Sort by** – Sets the sort applied to the results that appear below. You can choose from the following:
 - **Absolute difference**
 - **Contribution percentage** (default)
 - **Deviation from expected**
 - **Percentage difference**
- **Top contributor results** – Displays the results of the top contributor analysis for the point in time selected on the timeline at right.

Contribution analysis identifies up to four of the top contributing factors or key drivers of an anomaly. For example, Amazon Quick Sight can show you the top customers that contributed to a spike in sales in the US for health products. This panel appears only if you choose to include fields in contribution analysis when you configure the anomaly.

If you don't see this panel and you want to display it, you can turn it on. To do so, go to the analysis, choose anomaly configuration from the insight's menu, and choose up to four fields to analyze for contributions. If you make changes in the sheet controls that exclude the contributing drivers, the **Contributions** panel closes.

Setting controls for anomaly detection

You can find the settings for anomaly detection in the **Controls** section of the screen. You can open and close this section by clicking the word **Controls**.

The settings include the following:

- **Controls** – The current settings appear at the top of the workspace. You can expand this section by choosing the double arrow icon on the right side. The following settings are available for exploring outliers generated by ML-powered anomaly detection:
 - **Severity** – Sets how sensitive your detector is to detected anomalies (outliers). You should expect to see more anomalies with the threshold set to **Low and above**, and fewer anomalies when the threshold is set to **High and above**. This sensitivity is determined based on standard deviations of the anomaly score generated by the RCF algorithm. The default is **Medium and above**.
 - **Direction** – The direction on the x-axis or y-axis that you want to identify as anomalous. The default is [ALL]. You can choose the following:
 - Set to **Higher than expected** to identify higher values as anomalies.
 - Set to **Lower than expected** to identify lower values as anomalies.
 - Set to **[ALL]** to identify all anomalous values, both high and low.
 - **Minimum Delta - absolute value** – Enter a custom value to use to as the absolute threshold to identify anomalies. Any amount higher than this value counts as an anomaly.
 - **Minimum Delta - percentage** – Enter a custom value to use to as the percentage threshold to identify anomalies. Any amount higher than this value counts as an anomaly.
 - **Sort by** – Choose the method that you want to apply to sorting anomalies. These are listed in preferred order on the screen. View the following list for a description of each method.
 - **Weighted anomaly score** – The anomaly score multiplied by the log of the absolute value of the difference between the actual value and the expected value. This score is always a positive number.
 - **Anomaly score** – The actual anomaly score assigned to this data point.
 - **Weighted difference from expected value** – (Default) The anomaly score multiplied by the difference between the actual value and the expected value.
 - **Difference from expected value** – The actual difference between the actual value and the expected value (actual–expected).
 - **Actual value** – The actual value with no formula applied.
 - **Categories** – One or more settings can appear at the end of the other settings. There is one for each category field that you added to the category field well. You can use category settings to limit the data that displays in the screen.

Showing and hiding anomalies by date

The **Number of anomalies** chart shows outliers detected over time. If you don't see this chart, you can display it by choosing **SHOW ANOMALIES BY DATE**.

This chart shows anomalies (outliers) for the most recent data point in the time series. When expanded, it displays the following components:

- **Anomalies** – The middle of the screen displays the anomalies for the most recent data point in the time series. One or more graphs appear with a chart showing variations in a metric over time. To use this graph, select a point along the timeline. The currently selected point in time is highlighted in the graph, and includes a menu offering you the option to analyze contributions to the current metric. You can also drag the cursor over the timeline without choosing a specific point to display the metric value for that point in time.
- **Anomalies by date** – If you choose **SHOW ANOMALIES BY DATE**, another graph appears that shows how many significant anomalies there were for each time point. You can see details in this chart on each bar's context menu.
- **Timeline adjustment** – Each graph has a timeline adjustor tool below the dates, which you can use to compress, expand, or choose a period of time to view.

Exploring anomalies per category or dimension

The main section of the **Explore anomalies** screen is locked to the lower right of the screen. It remains here no matter how many other sections of the screen are open. If multiple anomalies exist, you can scroll out to highlight them. The chart displays anomalies in color ranges and shows where they occur over a period of time.

Each category or dimension has a separate chart that uses the field name as the chart title. Each chart contains the following components:

- **Configure alerts** – If you are exploring anomalies from a dashboard, select this button to subscribe to alerts and contribution analysis (if configured). You can set up the alerts for the level of severity (medium, high, and so on). You can get the top five alerts for **Higher than expected**, **Lower than expected**, or **ALL**. Dashboard readers can configure alerts for themselves. If you open the **Explore Anomalies** page doesn't display this button if you opened the page from an analysis.

Note

The ability to configure alerts is available only in published dashboards.

- **Status** – Under the **Anomalies** header, the status label displays information on the last run. For example, you might see "Anomalies for Revenue on November 17, 2018." This label tells you how many metrics were processed and how long ago. You can choose the link to learn more about the details, such as how many metrics were ignored.

Forecasting and creating what-if scenarios with Amazon Quick Sight

Using ML-powered forecasting, you can forecast your key business metrics with point-and-click simplicity. No machine learning expertise is required. The built-in ML algorithm in Amazon Quick Sight is designed to handle complex real-world scenarios. Amazon Quick Sight uses machine learning to help provide more reliable forecasts than available by traditional means.

For example, suppose that you are a business manager. Suppose that you want to forecast sales to see if you are going to meet your goal by the end of the year. Or, suppose that you expect a large deal to come through in two weeks and you want to know how it's going to affect your overall forecast.

You can forecast your business revenue with multiple levels of seasonality (for example, sales with both weekly and quarterly trends). Amazon Quick Sight automatically excludes anomalies in the data (for example, a spike in sales due to price drop or promotion) from influencing the forecast. You also don't have to clean and reprep the data with missing values because Amazon Quick Sight automatically handles that. In addition, with ML-powered forecasting, you can perform interactive what-if analyses to determine the growth trajectory you need to meet business goals.

Using forecasts and what-if scenarios

You can add a forecasting widget to your existing analysis, and publish it as a dashboard. To analyze what-if scenarios, use an analysis, not a dashboard. With ML-powered forecasting, Amazon Quick Sight enables you to forecast complex, real-world scenarios such as data with multiple seasonality. It automatically excludes outliers that it identifies and imputes missing values.

Use the following procedure to add a graphical forecast to your analysis, and explore what-if scenarios.

Although the following procedure is for graphical forecasting, you can also add a forecast as a narrative in an insight widget. To learn more, see [Creating autonarratives with Amazon Quick Sight](#).

ML-powered forecasting is not compatible with [small multiples](#). To ensure accurate display of data and forecasts, avoid using small multiples in your visualizations.

To add a graphical forecast to your analysis

1. Create a visual that uses a single date field and up to three metrics (measures).
2. On the menu in the upper-right corner of the visual, choose the **Menu options** icon (the three dots), and then choose **Add forecast**.

Quick Sight automatically analyzes the historical data using ML, and displays a graphical forecast for the next 14 periods. Forecast properties apply to all metrics in your visual. If you want individual forecasts for each metric, consider creating a separate visual for each metric and adding a forecast to each.

3. On the **Forecast properties** panel at left, customize one or more of the following settings:
 - **Forecast length** – Set **Periods forward** to forecast, or set **Periods backward** to look for patterns to base the forecast on.
 - **Prediction interval** – Set the estimated range for the forecast. Doing this changes how wide the band of possibility is around the predicted line.
 - **Seasonality** – Set the number of time periods involved in the predictable seasonal pattern of data. The range is 1–180, and the default setting is **Automatic**.
 - **Forecast boundaries** – Set a minimum and/or maximum forecast value to prevent forecast values from going above or below a specified value. For example, if your forecasting predicts the number of new hires the company will make in the next month to be in the negative numbers, you can set a forecast boundary minimum to zero. This stops the forecasted values from ever going below zero.

To save your changes, choose **Apply**.

If your forecast contains multiple metrics, you can isolate one of the forecasts by selecting anywhere inside the orange band. When you do this, the other forecasts disappear. Select the isolated forecast band again to have them reappear.

4. Analyze what-if scenarios by choosing a forecasted data point (in the orange band) on the chart, and then choosing **What-if analysis** from the context menu.

The **What-if analysis** panel opens at left. Set the following options:

- **Scenario** – Set a target for a date, or set a target for a time range.
- **Dates** – If you are setting a target for a specific date, enter that date here. If you are using a time range, set the start and end dates.
- **Target** – Set a target value for the metric.

Amazon Quick Sight adjusts the forecast to meet the target.

 Note

The **What-if analysis** option isn't available for multiple-metric forecasts. If you want to perform a what-if scenario on your forecast, your visual should contain only one metric.

5. Keep your changes by choosing **Apply**. To discard them, close the **What-if analysis** panel.

If you keep your changes, you see the new forecast adjusted for the target, alongside the original forecast without the what-if.

The what-if analysis is represented on the visual as a dot on the metric line. You can hover over the data points on the forecasting line to see the details.

Here are other things you can do:

- To interact with or remove a what-if analysis, choose the dot on the metric line.
- To create additional what-if scenarios, close the what-if analysis before choosing a new point on the line.

 Note

What-if analyses can exist inside an analysis only, not inside a dashboard.

Troubleshooting Amazon Quick Sight

Use this information to help you diagnose and fix common issues that you can encounter when using Amazon Quick Sight.

Note

Need more help? You can visit the Amazon Quick Sight [User Community](#) or the [Amazon forums](#). See also the [Quick Sight Resource Library](#).

Topics

- [Resolving Amazon Quick Sight issues and error messages](#)
- [Connectivity issues when using Amazon Athena with Amazon Quick Sight](#)
- [Data source connectivity issues for Amazon Quick Sight](#)
- [Login issues with Quick Sight](#)
- [Visual issues with Quick Sight](#)

Resolving Amazon Quick Sight issues and error messages

If you are having difficulties or receiving an error message, there's a few ways that you can go about resolving the issue. Following are some resources that can help:

- For errors during dataset ingestion (importing data), see [SPICE ingestion error codes](#).
- For technical user questions, visit the [User Community](#).
- For administrator questions, see the [Amazon forums](#).
- If you need more customized assistance, contact Amazon Support. To do this while you are signed in to your Amazon Web Services account, choose **Support** at upper right, and then choose **Support Center**.

Connectivity issues when using Amazon Athena with Amazon Quick Sight

Following, you can find information about troubleshooting issues that you might encounter when using Amazon Athena with Amazon Quick Sight.

Before you try troubleshooting anything else for Athena, make sure that you can connect to Athena. For information about troubleshooting Athena connection issues, see [I can't connect to Amazon Athena](#).

If you can connect but have other issues, it can be useful to run your query in the Athena console (<https://console.amazonaws.cn/athena/>) before adding your query to Amazon Quick Sight. For additional troubleshooting information, see [Troubleshooting](#) in the *Athena User Guide*.

Topics

- [Column not found when using Athena with Amazon Quick Sight](#)
- [Invalid data when using Athena with Amazon Quick Sight](#)
- [Query timeout when using Athena with Amazon Quick Sight](#)
- [Staging bucket no longer exists when using Athena with Amazon Quick Sight](#)
- [Table incompatible when using Amazon Glue with Athena in Amazon Quick Sight](#)
- [Table not found when using Athena with Amazon Quick Sight](#)
- [Workgroup and output errors when using Athena with Quick Sight](#)

Column not found when using Athena with Amazon Quick Sight

You can receive a "column not found" error if the columns in an analysis are missing from the Athena data source.

In Amazon Quick Sight, open your analysis. On the **Visualize** tab, choose **Choose dataset**, **Edit analysis data sets**.

On the **Data sets in this analysis** screen, choose **Edit** near your dataset to refresh the dataset. Amazon Quick Sight caches the schema for two minutes. So it can take two minutes before the latest changes display.

To investigate how the column was lost in the first place, you can go to the Athena console (<https://console.amazonaws.cn/athena/>) and check the query history to find queries that edited the table.

If this error happened when you were editing a custom SQL query in preview, verify that the name of the column in the query, and check for any other syntax errors. For example, check that the column name isn't enclosed in single quotation marks, which are reserved for strings.

If you still have the issue, verify that your tables, columns, and queries comply with Athena requirements. For more information, see [Names for Tables, Databases, and Columns](#) and [Troubleshooting](#) in the *Athena User Guide*.

Invalid data when using Athena with Amazon Quick Sight

An invalid data error can occur when you use any operator or function in a calculated field. To address this, verify that the data in the table is consistent with the format that you supplied to the function.

For example, suppose that you are using the function `parseDate(expression, ['format'], ['time_zone'])` as `parseDate(date_column, 'MM/dd/yyyy')`. In this case, all values in `date_column` must conform to 'MM/dd/yyyy' format ('05/12/2016'). Any value that isn't in this format ('2016/12/05') can cause an error.

Query timeout when using Athena with Amazon Quick Sight

If your query times out, you can try these options to resolve your problem.

If the failure was generated while working on an analysis, remember that the Amazon Quick Sight timeout for generating any visual is two minutes. If you're using a custom SQL query, you can simplify your query to optimize running time.

If you are in direct query mode (not using SPICE), you can try importing your data to SPICE. However, if your query exceeds the Athena 30-minute timeout, you might get another timeout while importing data into SPICE. For the most current information on Athena limits, see [Amazon Athena Limits](#) in the *Amazon Web Services General Reference*.

Staging bucket no longer exists when using Athena with Amazon Quick Sight

Use this section to help solve this error: **"The staging bucket for this query result no longer exists in the underlying data source."**

When you create a dataset using Athena, Amazon Quick Sight creates an Amazon S3 bucket. By default, this bucket has a name similar to "aws-athena-query-results-**<REGION>**-**<ACCOUNTID>**". If you remove this bucket, then your next Athena query might fail with an error saying the staging bucket no longer exists.

To fix this error, create a new bucket with the same name in the correct Amazon Web Services Region.

Table incompatible when using Amazon Glue with Athena in Amazon Quick Sight

If you are getting errors when using Amazon Glue tables in Athena with Amazon Quick Sight, it might be because you're missing some metadata. Follow these steps to find out if your tables don't have the `TableType` attribute that Amazon Quick Sight needs for the Athena connector to work. Usually, the metadata for these tables wasn't migrated to the Amazon Glue Data Catalog. For more information, see [Upgrading to the Amazon Glue Data Catalog Step-by-Step](#) in the *Amazon Glue Developer Guide*.

If you don't want to migrate to the Amazon Glue Data Catalog at this time, you have two options. You can recreate each Amazon Glue table through the Amazon Glue Management Console. Or you can use the Amazon CLI scripts listed in the following procedure to identify and update tables with missing `TableType` attributes.

If you prefer to use the CLI to do this, use the following procedure to help you design your scripts.

To use the CLI to design scripts

1. Use the CLI to learn which Amazon Glue tables have no `TableType` attributes.

```
aws glue get-tables --database-name <your_database_name>;
```

For example, you can run the following command in the CLI.

```
aws glue get-table --database-name "test_database" --name  
"table_missing_table_type"
```

Following is a sample of what the output looks like. You can see that the table `"table_missing_table_type"` doesn't have the `TableType` attribute declared.

```
{  
  "TableList": [  
    {  
      "Retention": 0,  
      "UpdateTime": 1522368588.0,  
      "PartitionKeys": [  
        {  
          "Name": "year",  
          "Type": "string"  
        },  
      ],  
    },  
  ],  
}
```

```
{
  "Name": "month",
  "Type": "string"
},
{
  "Name": "day",
  "Type": "string"
}
],
"LastAccessTime": 1513804142.0,
"Owner": "owner",
"Name": "table_missing_table_type",
"Parameters": {
  "delimiter": ",",
  "compressionType": "none",
  "skip.header.line.count": "1",
  "sizeKey": "75",
  "averageRecordSize": "7",
  "classification": "csv",
  "objectCount": "1",
  "typeOfData": "file",
  "CrawlerSchemaDeserializerVersion": "1.0",
  "CrawlerSchemaSerializerVersion": "1.0",
  "UPDATED_BY_CRAWLER": "crawl_date_table",
  "recordCount": "9",
  "columnsOrdered": "true"
},
"StorageDescriptor": {
  "OutputFormat": "org.apache.hadoop.hive ql.io.HiveIgnoreKeyTextOutputFormat",
  "SortColumns": [],
  "StoredAsSubDirectories": false,
  "Columns": [
    {
      "Name": "col1",
      "Type": "string"
    },
    {
      "Name": "col2",
      "Type": "bigint"
    }
  ]
},
"Location": "s3://myAthenatest/test_dataset/",
"NumberOfBuckets": -1,
"Parameters": {
```

```

    "delimiter": ",",
    "compressionType": "none",
    "skip.header.line.count": "1",
    "columnsOrdered": "true",
    "sizeKey": "75",
    "averageRecordSize": "7",
    "classification": "csv",
    "objectCount": "1",
    "typeOfData": "file",
    "CrawlerSchemaDeserializerVersion": "1.0",
    "CrawlerSchemaSerializerVersion": "1.0",
    "UPDATED_BY_CRAWLER": "crawl_date_table",
    "recordCount": "9"
  },
  "Compressed": false,
  "BucketColumns": [],
  "InputFormat": "org.apache.hadoop.mapred.TextInputFormat",
  "SerdeInfo": {
    "Parameters": {
      "field.delim": ",",
    },
    "SerializationLibrary": "org.apache.hadoop.hive.serde2.lazy.LazySimpleSerDe"
  }
}
}
]
}

```

2. Edit the table definition in your editor to add **"TableType": "EXTERNAL_TABLE"** to the table definition, as shown in the following example.

```

{
  "Table": {
    "Retention": 0,
    "TableType": "EXTERNAL_TABLE",
    "PartitionKeys": [
      {
        "Name": "year",
        "Type": "string"
      },
      {
        "Name": "month",
        "Type": "string"
      }
    ]
  }
}

```

```
    },
    {
      "Name": "day",
      "Type": "string"
    }
  ],
  "UpdateTime": 1522368588.0,
  "Name": "table_missing_table_type",
  "StorageDescriptor": {
    "BucketColumns": [],
    "SortColumns": [],
    "StoredAsSubDirectories": false,
    "OutputFormat": "org.apache.hadoop.hive ql.io.HiveIgnoreKeyTextOutputFormat",
    "SerdeInfo": {
      "SerializationLibrary": "org.apache.hadoop.hive.serde2.lazy.LazySimpleSerDe",
      "Parameters": {
        "field.delim": ",",
      }
    }
  },
  "Parameters": {
    "classification": "csv",
    "CrawlerSchemaSerializerVersion": "1.0",
    "UPDATED_BY_CRAWLER": "crawl_date_table",
    "columnsOrdered": "true",
    "averageRecordSize": "7",
    "objectCount": "1",
    "sizeKey": "75",
    "delimiter": ",",
    "compressionType": "none",
    "recordCount": "9",
    "CrawlerSchemaDeserializerVersion": "1.0",
    "typeOfData": "file",
    "skip.header.line.count": "1"
  },
  "Columns": [
    {
      "Name": "col1",
      "Type": "string"
    },
    {
      "Name": "col2",
      "Type": "bigint"
    }
  ],
]
```

```

    "Compressed": false,
    "InputFormat": "org.apache.hadoop.mapred.TextInputFormat",
    "NumberOfBuckets": -1,
    "Location": "s3://myAthenatest/test_date_part/"
  },
  "Owner": "owner",
  "Parameters": {
    "classification": "csv",
    "CrawlerSchemaSerializerVersion": "1.0",
    "UPDATED_BY_CRAWLER": "crawl_date_table",
    "columnsOrdered": "true",
    "averageRecordSize": "7",
    "objectCount": "1",
    "sizeKey": "75",
    "delimiter": ",",
    "compressionType": "none",
    "recordCount": "9",
    "CrawlerSchemaDeserializerVersion": "1.0",
    "typeOfData": "file",
    "skip.header.line.count": "1"
  },
  "LastAccessTime": 1513804142.0
}

```

3. You can adapt the following script to update the table input, so that it includes the `TableType` attribute.

```

aws glue update-table --database-name <your_datebase_name> --table-input
<updated_table_input>

```

The following shows an example.

```

aws glue update-table --database-name test_database --table-input '
{
  "Retention": 0,
  "TableType": "EXTERNAL_TABLE",
  "PartitionKeys": [
    {
      "Name": "year",
      "Type": "string"
    },
  ],
  {

```

```
    "Name": "month",
    "Type": "string"
  },
  {
    "Name": "day",
    "Type": "string"
  }
],
"Name": "table_missing_table_type",
"StorageDescriptor": {
  "BucketColumns": [],
  "SortColumns": [],
  "StoredAsSubDirectories": false,
  "OutputFormat": "org.apache.hadoop.hive.ql.io.HiveIgnoreKeyTextOutputFormat",
  "SerdeInfo": {
    "SerializationLibrary": "org.apache.hadoop.hive.serde2.lazy.LazySimpleSerDe",
    "Parameters": {
      "field.delim": ",",
    }
  },
},
"Parameters": {
  "classification": "csv",
  "CrawlerSchemaSerializerVersion": "1.0",
  "UPDATED_BY_CRAWLER": "crawl_date_table",
  "columnsOrdered": "true",
  "averageRecordSize": "7",
  "objectCount": "1",
  "sizeKey": "75",
  "delimiter": ",",
  "compressionType": "none",
  "recordCount": "9",
  "CrawlerSchemaDeserializerVersion": "1.0",
  "typeOfData": "file",
  "skip.header.line.count": "1"
},
"Columns": [
  {
    "Name": "col1",
    "Type": "string"
  },
  {
    "Name": "col2",
    "Type": "bigint"
  }
]
```

```
],
  "Compressed": false,
  "InputFormat": "org.apache.hadoop.mapred.TextInputFormat",
  "NumberOfBuckets": -1,
  "Location": "s3://myAthenatest/test_date_part/"
},
"Owner": "owner",
"Parameters": {
  "classification": "csv",
  "CrawlerSchemaSerializerVersion": "1.0",
  "UPDATED_BY_CRAWLER": "crawl_date_table",
  "columnsOrdered": "true",
  "averageRecordSize": "7",
  "objectCount": "1",
  "sizeKey": "75",
  "delimiter": ",",
  "compressionType": "none",
  "recordCount": "9",
  "CrawlerSchemaDeserializerVersion": "1.0",
  "typeOfData": "file",
  "skip.header.line.count": "1"
},
"LastAccessTime": 1513804142.0
}'
```

Table not found when using Athena with Amazon Quick Sight

You can receive a "table not found" error if the tables in an analysis are missing from the Athena data source.

In the Athena console (<https://console.amazonaws.cn/athena/>), check for your table under the corresponding schema. You can recreate the table in Athena and then create a new dataset in Amazon Quick Sight on that table. To investigate how the table was lost in the first place, you can use the Athena console to check the query history. Doing this helps you find the queries that dropped the table.

If this error happened when you were editing a custom SQL query in preview, verify that the name of the table in the query, and check for any other syntax errors. Amazon Quick Sight can't infer the schema from the query. The schema must be specified in the query.

For example, the following statement works.


```
select from my_schema.my_table
```

The following statement fails because it's missing the schema.

```
select from my_table
```

If you still have the issue, verify that your tables, columns, and queries comply with Athena requirements. For more information, see [Names for Tables, Databases, and Columns](#) and [Troubleshooting](#) in the *Athena User Guide*.

Workgroup and output errors when using Athena with Quick Sight

To verify that workgroups are set up properly, check the following settings:

- **The Athena workgroup that's associated with the data source must exist.**

To fix this, you can return to the Athena data source settings and choose a different workgroup. For more information, see [Setting Up Workgroups](#) in the *Athena User Guide*.

Another solution is to have the Amazon Web Services account administrator recreate the workgroup in the Athena console.

- **The Athena workgroup that's associated with the data source must be enabled.**

An Amazon Web Services account administrator needs to enable the workgroup in the Athena console. Open the Athena console by using this direct link: <https://console.amazonaws.cn/athena/>. Then choose the appropriate workgroup in the **Workgroup** panel and view its settings. Choose **Enable workgroup**.

- **Make sure that you have access to the Amazon S3 output location that's associated with the Athena workgroup.**

To grant Amazon Quick Sight permissions to access the S3 output location, the Amazon Quick Sight administrator can edit **Security & Permissions** in the **Manage QuickSight** screen.

- **The Athena workgroup must have an associated S3 output location.**

An Amazon Web Services account administrator needs to associate an S3 bucket with the workgroup in the Athena console. Open the Athena console by using this direct link: <https://console.amazonaws.cn/athena/>. Then choose the appropriate workgroup in the **Workgroup** panel and view its settings. Set **Query result location**.

Data source connectivity issues for Amazon Quick Sight

Use the following section to help you troubleshoot connections to data sources. Before you continue, verify that your database is currently available. Also, verify that you have the correct connection information and valid credentials.

Topics

- [I can't connect although my data source connection options look right \(SSL\)](#)
- [I can't connect to Amazon Athena](#)
- [I can't connect to Amazon S3](#)
- [I can't create or refresh a dataset from an existing Adobe Analytics data source](#)
- [I need to validate the connection to my data source, or change data source settings](#)
- [I can't connect to MySQL \(issues with SSL and authorization\)](#)
- [I can't connect to RDS](#)

I can't connect although my data source connection options look right (SSL)

Problems connecting can occur when Secure Sockets Layer (SSL) is incorrectly configured. The symptoms can include the following:

- You can connect to your database in other ways or from other locations but not in this case.
- You can connect to a similar database but not this one.

Before continuing, rule out the following circumstances:

- Permissions issues
- Availability issues
- An expired or invalid certificate
- A self-signed certificate
- Certificate chain in the wrong order
- Ports not enabled
- Firewall blocking an IP address
- Web Sockets are blocked
- A virtual private cloud (VPC) or security group not configured correctly.

To help find issues with SSL, you can use an online SSL checker, or a tool like OpenSSL.

The following steps walk through troubleshooting a connection where SSL is suspect. The administrator in this example has already installed OpenSSL.

Example

1. The user finds an issue connecting to the database. The user verifies that they can connect a different database in another Amazon Web Services Region. They check other versions of the same database and can connect easily.
2. The administrator reviews the issue and decides to verify that the certificates are working correctly. The administrator searches online for an article on using OpenSSL to troubleshoot or debug SSL connections.
3. Using OpenSSL, the administrator verifies the SSL configuration in the terminal.

```
echo quit
openssl s_client -connect <host>:port
```

The result shows that the certificate is not working.

```
...
...
...
CONNECTED(00000003)
012345678901234:error:140770FC:SSL routines:SSL23_GET_SERVER_HELLO:unknown
protocol:s23_clnt.c:782:
---
no peer certificate available
---
No client certificate CA names sent
---
SSL handshake has read 7 bytes and written 278 bytes
---
New, (NONE), Cipher is (NONE)
Secure Renegotiation IS NOT supported
SSL-Session:
    Protocol  : TLSv1.2
    Cipher    : 0000
```

```
Session-ID:  
Session-ID-ctx:  
Master-Key:  
Key-Arg    : None  
PSK identity: None  
PSK identity hint: None  
Start Time: 1497569068  
Timeout    : 300 (sec)  
Verify return code: 0 (ok)
```

4. The administrator corrects the problem by installing the SSL certificate on the user's database server.

For more detail on the solution in this example, see [Using SSL to Encrypt a Connection to a DB Instance](#) in the *Amazon RDS User Guide*.

I can't connect to Amazon Athena

Intended audience: Amazon Quick Suite administrators

Use this section to help troubleshoot connecting to Athena.

If you can't connect to Amazon Athena, you might get an insufficient permissions error when you run a query, showing that the permissions aren't configured. To verify that you can connect Amazon Quick Sight to Athena, check the following settings:

- Amazon resource permissions inside of Amazon Quick Sight
- Amazon Identity and Access Management (IAM) policies
- Amazon S3 location
- Query results location
- Amazon KMS key policy (for encrypted datasets only)

For details, see following. For information about troubleshooting other Athena issues, see [Connectivity issues when using Amazon Athena with Amazon Quick Sight](#).

Make sure that you authorized Amazon Quick Sight to use Athena

Intended audience: Amazon Quick Suite administrators

Use the following procedure to make sure that you successfully authorized Amazon Quick Sight to use Athena. Permissions to Amazon resources apply to all Amazon Quick Sight users.

To perform this action, you must be an Amazon Quick Sight administrator. To check if you have access, verify that you see the **Manage QuickSight** option when you open the menu from your profile at upper right.

To authorize Amazon Quick Sight to access Athena

1. Choose your profile name (upper right). Choose **Manage Quick Sight**, and then scroll down to the **Custom permissions** section.
2. Choose **Amazon resources** then choose **Add or remove**.
3. Find Athena in the list. Clear the box by Athena, then select it again to enable Athena.

Then choose **Connect both**.

4. Choose the buckets that you want to access from Amazon Quick Sight.

The settings for S3 buckets that you access here are the same ones that you access by choosing Amazon S3 from the list of Amazon Web Services services. Be careful that you don't inadvertently disable a bucket that someone else uses.

5. Choose **Finish** to confirm your selection. Or choose **Cancel** to exit without saving.
6. Choose **Update** to save your new settings for Amazon Quick Sight access to Amazon Web Services services. Or choose **Cancel** to exit without making any changes.
7. Make sure that you are using the correct Amazon Web Services Region when you are finished.

If you had to change your Amazon Web Services Region as part of the first step of this process, change it back to the Amazon Web Services Region that you were using before.

Make sure that your IAM policies grant the right permissions

Intended audience: System administrators

Your Amazon Identity and Access Management (IAM) policies must grant permissions to specific actions. Your IAM user or role must be able to read and write both the input and the output of the S3 buckets that Athena uses for your query.

If the dataset is encrypted, the IAM user needs to be a key user in the specified Amazon KMS key's policy.

To verify that your IAM policies have permission to use S3 buckets for your query

1. Open the IAM console at <https://console.amazonaws.cn/iam/>.
2. Locate the IAM user or role that you are using. Choose the user or role name to see the associated policies.
3. Verify that your policy has the correct permissions. Choose a policy that you want to verify, and then choose **Edit policy**. Use the visual editor, which opens by default. If you have the JSON editor open instead, choose the **Visual editor** tab.
4. Choose the S3 entry in the list to see its contents. The policy needs to grant permissions to list, read, and write. If S3 is not in the list, or it doesn't have the correct permissions, you can add them here.

For examples of IAM policies that work with Quick Sight, see [IAM policy examples for Quick Suite](#).

Make sure that the IAM user has read/write access to your S3 location

Intended audience: Amazon Quick Suite administrators

To access Athena data from Quick Sight, first make sure that Athena and its S3 location are authorized in **Manage QuickSight** screen. For more information, see [Make sure that you authorized Amazon Quick Sight to use Athena](#).

Next, verify the relevant IAM permissions. The IAM user for your Athena connection needs read/write access to the location where your results go in S3. Start by verifying that the IAM user has an attached policy that [allows access to Athena](#), such as AmazonAthenaFullAccess. Let Athena create the bucket using the name that it requires, and then add this bucket to the list of buckets that QuickSight can access. If you change the default location of the results bucket (aws-athena-query-results-*), be sure that the IAM user has permission to read and write to the new location.

Verify that you don't include the Amazon Web Services Region code in the S3 URL. For example, use `s3://awsexamplebucket/path` and not `s3://us-east-1.amazonaws.com/awsexamplebucket/path`. Using the wrong S3 URL causes an Access Denied error.

Also verify that the bucket policies and object access control lists (ACLs) [allow the IAM user to access the objects in the buckets](#). If the IAM user is in a different Amazon Web Services account, see [Cross-account Access](#) in the *Amazon Athena User Guide*.

If the dataset is encrypted, verify that the IAM user is a key user in the specified Amazon KMS key's policy. You can do this in the Amazon KMS console at <https://console.amazonaws.cn/kms>.

To set permissions to your Athena query results location

1. Open the Athena console at <https://console.amazonaws.cn/athena/>.
2. Verify that you have selected the workgroup you want to use:
 - Examine the **Workgroup** option at the top. It has the format **Workgroup: *group-name***. If the group name is the one that you want to use, skip to the next step.
 - To choose a different workgroup, choose **Workgroup** at the top. Choose the workgroup that you want to use, and choose **Switch workgroup**.
3. Choose **Settings** at upper right.

(Not common) If you get an error that your workgroup is not found, use these steps to fix it:

- a. Ignore the error message for now, and instead find **Workgroup: *group-name*** on the **Settings** page. Your workgroup's name is a hyperlink. Open it.
- b. On the **Workgroup: *<groupname>*** page, choose **Edit workgroup** at left. Now close the error message.
- c. Near **Query result location**, open the S3 location selector by choosing the **Select** button that has the file folder icon.
- d. Choose the small arrow at the end of the name of the S3 location for Athena. The name must begin with `aws-athena-query-results`.
- e. (Optional) Encrypt query results by selecting the **Encrypt results stored in S3** check box.
- f. Choose **Save** to confirm your choices.
- g. If the error doesn't reappear, return to **Settings**.

Occasionally, the error might appear again. If so, take the following steps:

1. Choose the workgroup and then choose **View details**.
2. (Optional) To preserve your settings, take notes or a screenshot of the workgroup configuration.
3. Choose **Create workgroup**.
4. Replace the workgroup with a new one. Configure the correct S3 location and encryption options. Note the S3 location because you need it later.
5. Choose **Save** to proceed.
6. When you no longer need the original workgroup, disable it. Make sure to carefully read the warning that appears, because it tells you what you lose if you choose to disable it.
4. If you didn't get this by troubleshooting in the previous step, choose **Settings** at upper right and get the S3 location value shown as **Query result location**.
5. If **Encrypt query results** is enabled, check whether it uses SSE-KMS or CSE-KMS. Note the key.
6. Open the S3 console at <https://console.amazonaws.cn/s3/>, open the correct bucket, and then choose the **Permissions** tab.
7. Check that your IAM user has access by viewing **Bucket Policy**.

If you manage access with ACLs, make sure that the access control lists (ACLs) are set up by viewing **Access Control List**.

8. If your dataset is encrypted (**Encrypt query results** is selected in the workgroup settings), make sure that the IAM user or role is added as a key user in that Amazon KMS key's policy. You can access Amazon KMS settings at <https://console.amazonaws.cn/kms>.

To grant access to the S3 bucket used by Athena

1. Open the Amazon S3 console at <https://console.amazonaws.cn/s3/>.
2. Choose the S3 bucket used by Athena in the **Query result location**.
3. On the **Permissions** tab, verify the permissions.

For more information, see the Amazon Support article [When I run an Athena query, I get an "Access Denied" error](#).

I can't connect to Amazon S3

To successfully connect to Amazon S3, make sure that you configure authentication and create a valid manifest file inside the bucket you are trying to access. Also, make sure that the file described by the manifest is available.

To verify authentication, make sure that you authorized Amazon Quick Sight to access the S3 account. It's not enough that you, the user, are authorized. Amazon Quick Sight must be authorized separately.

To authorize Amazon Quick Sight to access your Amazon S3 bucket

1. In the Amazon Web Services Region list at upper right, choose the US East (N. Virginia) Region. You use this Amazon Web Services Region temporarily while you edit your account permissions.
2. Inside Amazon Quick Sight, choose your profile name (upper right). Choose **Manage Quick Sight**, and then scroll down to the **Custom permissions** section.
3. Choose **Amazon resources** then choose **Add or remove**.
4. Locate Amazon S3 in the list. Choose one of the following actions to open the screen where you can choose S3 buckets:
 - If the check box is clear, select the check box next to Amazon S3.
 - If the check box is selected, choose **Details**, and then choose **Select S3 buckets**.
5. Choose the buckets that you want to access from Amazon Quick Sight. Then choose **Select**.
6. Choose **Update**.
7. If you changed your Amazon Web Services Region during the first step of this process, change it back to the Amazon Web Services Region that you want to use.

We strongly recommend that you make sure that your manifest file is valid. If Amazon Quick Sight can't parse your file, it gives you an error message. That might be something like "We can't parse the manifest file as valid JSON" or "We can't connect to the S3 bucket."

To verify your manifest file

1. Open your manifest file. You can do this directly from the Amazon S3 console at <https://console.amazonaws.cn/s3/>. Go to your manifest file and choose **Open**.

2. Make sure that the URI or URLs provided inside the manifest file indicate the file or files that you want connect to.
3. Make sure that your manifest file is formed correctly, if you use a link to the manifest file rather than uploading the file. The link shouldn't have any additional phrases after the word `.json`. You can get the correct link to an S3 file by viewing its **Link** value in the details on the S3 console.
4. Make sure that the content of the manifest file is valid by using a JSON validator, like the one at <https://jsonlint.com>.
5. Verify permissions on your bucket or file. In the <https://console.amazonaws.cn/s3/>, navigate to your Amazon S3 bucket, choose the **Permissions** tab, and add the appropriate permissions. Make sure that the permissions are at the right level, either on the bucket or on the file or files.
6. If you are using the `s3://` protocol, rather than `https://`, make sure that you reference your bucket directly. For example, use `s3://awsexamplebucket/myfile.csv` instead of `s3://s3-us-west-2.amazonaws.com/awsexamplebucket/myfile.csv`. Doubly specifying Amazon S3, by using `s3://` and also `s3-us-west-2.amazonaws.com`, causes an error.

For more information about manifest files and connecting to Amazon S3, see [Supported formats for Amazon S3 manifest files](#).

In addition, verify that your Amazon S3 dataset was created according to the steps in [Creating a dataset using Amazon S3 files](#).

If you use Athena to connect to Amazon S3, see [I can't connect to Amazon Athena](#).

I can't create or refresh a dataset from an existing Adobe Analytics data source

As of May 1, 2022, Quick Sight no longer supports legacy OAuth and version 1.3 and SOAP API operations in Adobe Analytics. If you experience failures while trying to create or refresh a dataset from an existing Adobe Analytics data source, you might have a stale access token.

To troubleshoot failures while creating or refreshing a dataset from an existing Adobe Analytics data source

1. Open Quick Sight and choose **Data** at left.
2. Choose **New** then **Dataset**.
3. On the **Create a dataset** page, choose the Adobe Analytics data source that you want to update from the list of existing data sources.

4. Choose **Edit data source**.
5. On the **Edit Adobe Analytics data source** page that opens, choose **Update data source** to reauthorize the Adobe Analytics connection.
6. Try recreating or refreshing the dataset again. The dataset creation or refresh should succeed.

I need to validate the connection to my data source, or change data source settings

In some cases, you might need to update your data source, or you got a connection error and need to check your settings. If so, take the following steps.

To validate your connection to the data source

1. From the Quick Sight homepage, choose **Data** at left.
2. Choose **New** then **Dataset**.
3. You will see a list of existing data sources.
4. Choose the data source that you want to test or change.
5. If the option is offered, choose **Edit/Preview data**.
6. Choose **Validate connection**.
7. Make any changes that you want to make, then choose **Update data source**.

I can't connect to MySQL (issues with SSL and authorization)

To check on some common connection issues in MySQL, use the following steps. This procedure helps you find out if you have enabled SSL and granted usage rights.

To find solutions for some common connection issues in MySQL

1. Check `/etc/my.cnf` to make sure SSL is enabled for MySQL.
2. In MySQL, run the following command.

```
show status like 'Ssl%';
```

If SSL is working, you see results like the following.

Variable_name	Value
Ssl_accept_renegotiates	0
Ssl_accepts	1
Ssl_callback_cache_hits	0
Ssl_cipher	
Ssl_cipher_list	
Ssl_client_connects	0
Ssl_connect_renegotiates	0
Ssl_ctx_verify_depth	18446744073709551615
Ssl_ctx_verify_mode	5
Ssl_default_timeout	0
Ssl_finished_accepts	0
Ssl_finished_connects	0
Ssl_session_cache_hits	0
Ssl_session_cache_misses	0
Ssl_session_cache_mode	SERVER
Ssl_session_cache_overflows	0
Ssl_session_cache_size	128
Ssl_session_cache_timeouts	0
Ssl_sessions_reused	0
Ssl_used_session_cache_entries	0
Ssl_verify_depth	0
Ssl_verify_mode	0
Ssl_version	

If SSL is disabled, you see results like the following.

Variable_name	Value
Ssl_accept_renegotiates	0
Ssl_accepts	0
Ssl_callback_cache_hits	0
Ssl_cipher	
Ssl_cipher_list	
Ssl_client_connects	0

Ssl_connect_renegotiates	0	
Ssl_ctx_verify_depth	0	
Ssl_ctx_verify_mode	0	
Ssl_default_timeout	0	
Ssl_finished_accepts	0	
Ssl_finished_connects	0	
Ssl_session_cache_hits	0	
Ssl_session_cache_misses	0	
Ssl_session_cache_mode	NONE	
Ssl_session_cache_overflows	0	
Ssl_session_cache_size	0	
Ssl_session_cache_timeouts	0	
Ssl_sessions_reused	0	
Ssl_used_session_cache_entries	0	
Ssl_verify_depth	0	
Ssl_verify_mode	0	
Ssl_version		
+-----+-----+		

3. Make sure that you have installed a supported SSL certificate on the database server.
4. Grant usage for the specific user to connect using SSL.

```
GRANT USAGE ON *.* TO 'encrypted_user'@'%' REQUIRE SSL;
```

For more detail on the solution in this example, see the following:

- [SSL Support for MySQL DB Instances](#) in the *Amazon RDS User Guide*.
- [Using SSL to Encrypt a Connection to a DB Instance](#) in the *Amazon RDS User Guide*.
- [MySQL documentation](#)

I can't connect to RDS

For details on troubleshooting connections to Amazon RDS, see [Creating a dataset from a database](#).

You can also refer to the Amazon RDS documentation on troubleshooting connections, [Cannot Connect to Amazon RDS DB Instance](#).

Login issues with Quick Sight

Use the following section to help you troubleshoot login and access issues with the Quick Sight console.

Topics

- [Insufficient permissions when using Athena with Amazon Quick Sight](#)
- [Amazon Quick Sight isn't working in my browser](#)
- [How do I delete my Amazon Quick Sight account?](#)
- [Individuals in my organization get an "External Login is Unauthorized" message when they try to access Quick Sight](#)
- [My email sign-in stopped working](#)

Insufficient permissions when using Athena with Amazon Quick Sight

If you receive an error message that says you have insufficient permissions, try the following steps to resolve your problem.

You need administrator permissions to troubleshoot this issue.

To resolve an insufficient permissions error

1. Make sure that Amazon Quick Sight can access the Amazon S3 buckets used by Athena:
 - a. To do this, choose your profile name (upper right). Choose **Manage Quick Sight**, and then scroll down to the **Custom permissions** section.
 - b. Choose **Amazon resources** then choose **Add or remove**.
 - c. Locate Athena in the list. Clear the check box by Athena, then select it again to enable Athena.

Choose **Connect both**.

- d. Choose the buckets that you want to access from Amazon Quick Sight.

The settings for S3 buckets that you access here are the same ones that you access by choosing Amazon S3 from the list of Amazon Web Services services. Be careful that you don't inadvertently disable a bucket that someone else uses.

- e. Choose **Select** to save your S3 buckets.

- f. Choose **Update** to save your new settings for Amazon Quick Sight access to Amazon Web Services services. Or choose **Cancel** to exit without making any changes.
2. If your data file is encrypted with an Amazon KMS key, grant permissions to the Amazon Quick Sight IAM role to decrypt the key. The easiest way to do this is to use the Amazon CLI.

You can run the [create-grant](#) command in Amazon CLI to do this.

```
aws kms create-grant --key-id <Amazon KMS key ARN> --grantee-principal <Your Amazon Quick Sight Role ARN> --operations Decrypt
```

The Amazon Resource Name (ARN) for the Amazon Quick Sight role has the format `arn:aws-cn:iam::<account id>:role/service-role/aws-quicksight-service-role-v<version number>` and can be accessed from the IAM console. To find your Amazon KMS key ARN, use the S3 console. Go to the bucket that contains your data file and choose the **Overview** tab. The key is located near **KMS key ID**.

For Amazon Athena, Amazon S3, and Athena Query Federation connections, Quick Sight uses the following IAM role by default:

```
arn:aws-cn:iam::AWS-ACCOUNT-ID:role/service-role/aws-quicksight-s3-consumers-role-v0
```

If the `aws-quicksight-s3-consumers-role-v0` is not present, then Quick Sight uses:

```
arn:aws-cn:iam::AWS-ACCOUNT-ID:role/service-role/aws-quicksight-service-role-v0
```

Amazon Quick Sight isn't working in my browser

If you can't view Amazon Quick Sight correctly in your Google Chrome browser, take the following steps to fix the problem.

To view Amazon Quick Sight in your Chrome browser

1. Open Chrome and go to `chrome://flags/#touch-events`.
2. If the option is set to **Automatic**, change it to **Disabled**.
3. Close and reopen Chrome.

How do I delete my Amazon Quick Sight account?

In some cases, you might need to delete your Amazon Quick Sight account even when you can't access Amazon Quick Sight to unsubscribe. If so, sign in to Amazon and use the following link to open [the unsubscribe screen](https://us-east-1.quicksight.aws.amazon.com/sn/console/unsubscribe): <https://us-east-1.quicksight.aws.amazon.com/sn/console/unsubscribe>. This approach works no matter what Amazon Web Services Regions that you use. It deletes all data, analyses, Amazon Quick Sight users, and Amazon Quick Sight administrators. If you have further difficulty, contact support.

Individuals in my organization get an "External Login is Unauthorized" message when they try to access Quick Sight

Intended audience: Amazon Quick Suite administrators

When an individual in your organization is federating into Quick Sight using **AssumeRoleWithWebIdentity**, Quick Sight maps a single role-based user to a single external login. In some cases, that individual might be authenticated through an external login (such as Amazon Cognito) that's different from the originally mapped user. If so, they can't access Quick Sight and get the following unexpected error message.

The external login used for federation is unauthorized for the Quick Sight user.

To learn how to troubleshoot this issue, see the following sections:

- [Why is this happening?](#)
- [How can I fix it?](#)

Why is this happening?

You are using a simplified Amazon Cognito flow

If you're using Amazon Cognito to federate into Quick Sight, the single sign-on (IAM Identity Center) setup might use the `CognitoIdentityCredentials` API operation to assume the Quick Sight role. This method maps all users in the Amazon Cognito identity pool to a single Quick Sight user and isn't supported by Quick Sight.

We recommend that you use the `AssumeRoleWithWebIdentity` API operation instead, which specifies the role session name.

You're using unauthenticated Amazon Cognito users

Amazon Cognito IAM Identity Center is set up for unauthenticated users in the Amazon Cognito identity pool. The Quick Sight role trust policy is set up like the following example.

This setup allows a temporary Amazon Cognito user to assume a role session mapped to a unique Quick Sight user. Because unauthenticated identities are temporary, they aren't supported by Quick Sight.

We recommend that you don't use this setup, which setup isn't supported by Quick Sight. For Quick Sight, make sure that the Amazon Cognito IAM Identity Center uses authenticated users.

You deleted and recreated an Amazon Cognito user with the same user name attributes

In this case, the associated Amazon Cognito user that's mapped to the Quick Sight user was deleted and recreated. The newly created Amazon Cognito user has a different underlying subject. Depending on how the role session name is mapped to the Quick Sight user, the session name might correspond to the same Quick Sight role-based user.

We recommend that you remap the Quick Sight user to the updated Amazon Cognito user subject by using the `UpdateUser` API operation. For more information, see the following [UpdateUser API example](#).

You're mapping multiple Amazon Cognito user pools in different Amazon Web Services accounts to one identity pool and with Quick Sight

Mapping multiple Amazon Cognito user pools in different Amazon Web Services accounts to one identity pool and Quick Sight isn't supported by Quick Sight.

How can I fix it?

You can use Quick Sight public API operations to update the external login information for your users. Use the following options to learn how.

Use `RegisterUser` to create users with external login information

If the external login provider is Amazon Cognito, use the following CLI code to create users.

```
aws quicksight register-user --aws-account-id account-id --namespace namespace --  
email user-email --user-role user-role --identity-type IAM  
--iam-arn arn:aws-cn:iam::account-id:role/cognito-associated-iam-role  
--session-name cognito-username --external-login-federation-provider-type COGNITO
```

```
--external-login-id cognito-identity-id --region identity-region
```

The external-login-id should be the identity ID for the Amazon Cognito user. The format is <identity-region>:<cognito-user-sub>, as shown in the following example.

```
aws quicksight register-user --aws-account-id 111222333 --namespace default --email  
cognito-user@amazon.com --user-role ADMIN --identity-type IAM  
--iam-arn arn:aws-cn:iam::111222333:role/CognitoQuickSightRole  
--session-name cognito-user --external-login-federation-provider-type COGNITO  
--external-login-id us-east-1:12345678-1234-1234-abc1-a1b1234567 --region us-east-1
```

If the external login provider is a custom OpenID Connect (OIDC) provider, use the following CLI code to create users.

```
aws quicksight register-user --aws-account-id account-id --namespace namespace  
--email user-email --user-role user-role --identity-type IAM  
--iam-arn arn:aws-cn:iam::account-id:role/identity-provider-associated-iam-role  
--session-name identity-username --external-login-federation-provider-type CUSTOM_OIDC  
--custom-federation-provider-url custom-identity-provider-url  
--external-login-id custom-provider-identity-id --region identity-region
```

The following is an example.

```
aws quicksight register-user --aws-account-id 111222333 --namespace default  
--email identity-user@amazon.com --user-role ADMIN --identity-type IAM  
--iam-arn arn:aws-cn:iam::111222333:role/CustomIdentityQuickSightRole  
--session-name identity-user --external-login-federation-provider-type CUSTOM_OIDC  
--custom-federation-provider-url idp.us-east-1.amazonaws.com/us-east-1_ABCDE  
--external-login-id 12345678-1234-1234-abc1-a1b1234567 --region us-east-1
```

To learn more about using RegisterUser in the CLI, see [RegisterUser](#) in the *Amazon Quick Suite API Reference*.

Use DescribeUser to check external login information for users

If a user is a role-based federated user from an external login provider, use the DescribeUser API operation to check the external login information for it, as shown in the following code.

```
aws quicksight describe-user --aws-account-id account-id --namespace namespace  
--user-name identity-provider-associated-iam-role/identity-username  
--region identity-region
```

The following is an example.

```
aws quicksight describe-user --aws-account-id 111222333 --namespace default --user-name IdentityQuickSightRole/user --region us-west-2
```

The result contains the external login information fields if there are any. Following is an example.

```
{
  "Status": 200,
  "User": {
    "Arn": "arn:aws-cn:quicksight:us-east-1:111222333:user-default-IdentityQuickSightRole-user",
    "UserName": "IdentityQuickSightRole-user",
    "Email": "user@amazon.com",
    "Role": "ADMIN",
    "IdentityType": "IAM",
    "Active": true,
    "PrincipalId": "federated-iam-AR0AAAAAAAAAAAAAAAAA:user",
    "ExternalLoginFederationProviderType": "COGNITO",
    "ExternalLoginFederationProviderUrl": "cognito-identity.amazonaws.com",
    "ExternalLoginId": "us-east-1:123abc-1234-123a-b123-12345678a"
  },
  "RequestId": "12345678-1234-1234-abc1-a1b1234567"
}
```

To learn more about using `DescribeUser` in the CLI, see [DescribeUser](#) in the *Amazon Quick Suite API Reference*.

Use `UpdateUser` to update external login information for users

In some cases, you might find that the external login information saved for the user from the `DescribeUser` result isn't correct or the external login information is missing. If so, you can use the `UpdateUser` API operation to update it. Use the following examples.

For Amazon Cognito users, use the following.

```
aws quicksight update-user --aws-account-id account-id --namespace namespace
--user-name cognito-associated-iam-role/cognito-username
--email user-email --role user-role
--external-login-federation-provider-type COGNITO
--external-login-id cognito-identity-id --region identity-region
```

The following is an example.

```
aws quicksight update-user --aws-account-id 111222333 --namespace default
--user-name CognitoQuickSightRole/cognito-user --email cognito-user@amazon.com
--role ADMIN --external-login-federation-provider-type COGNITO
--external-login-id us-east-1:12345678-1234-1234-abc1-a1b1234567 --region us-west-2
```

For custom OIDC provider users, use the following.

```
aws quicksight update-user --aws-account-id account-id --namespace namespace
--user-name identity-provider-associated-iam-role/identity-username
--email user-email --role user-role
--external-login-federation-provider-type CUSTOM_OIDC
--custom-federation-provider-url custom-identity-provider-url
--external-login-id custom-provider-identity-id --region identity-region
```

The following is an example.

```
aws quicksight update-user --aws-account-id 111222333 --namespace default
--user-name IdentityQuickSightRole/user --email user@amazon.com --role ADMIN
--external-login-federation-provider-type CUSTOM_OIDC
--custom-federation-provider-url idp.us-east-1.amazonaws.com/us-east-1_ABCDE
--external-login-id 123abc-1234-123a-b123-12345678a --region us-west-2
```

If you want to delete the external login information for the user, use NONE external login federation provider type. Use the following CLI command to delete external login information.

```
aws quicksight update-user --aws-account-id account-id --namespace namespace
--user-name identity-provider-associated-iam-role/identity-username
--email user-email --role user-role
--external-login-federation-provider-type NONE --region identity-region
```

The following is an example.

```
aws quicksight update-user --aws-account-id 111222333 --namespace default
--user-name CognitoQuickSightRole/cognito-user --email cognito-user@amazon.com --role
ADMIN --external-login-federation-provider-type NONE --region us-west-2
```

To learn more about using UpdateUser in the CLI, see the [UpdateUser](#) in the *Amazon Quick Suite API Reference*.

My email sign-in stopped working

Currently, emails are case-sensitive. If yours isn't working, ask your administrator to check it for a mix of upper and lowercase letters. Use your email as it was entered.

Visual issues with Quick Sight

Use the following section to help you troubleshoot problems with visuals and their formatting.

Topics

- [I can't see my visuals](#)
- [I get a feedback bar across my printed documents](#)
- [My map charts don't show locations](#)
- [My pivot table stops working](#)
- [My visual can't find missing columns](#)
- [My visual can't find the query table](#)
- [My visual doesn't update after I change a calculated field](#)
- [Values in a Microsoft Excel file with scientific notation don't format correctly in Quick Sight](#)

I can't see my visuals

Use the following section to help you troubleshoot missing visuals. Before you continue, check to make sure you can still access your data source. If you can't connect to your data source, see [Data source connectivity issues for Amazon Quick Sight](#).

- If you are having trouble adding a visual to an analysis, try the following:
 - Check your connectivity and confirm that you have access to all domains that Quick Sight uses for access. To see a list of all URLs Quick Sight uses, see [Domains accessed by Quick Sight](#).
 - Check that you aren't trying to add more objects than the quota allows. Amazon Quick Sight supports up to 30 datasets in a single analysis, up to 30 visuals in a single sheet, and a limit of 20 sheets per analysis.
 - Suppose that you are editing an analysis for a selected data source and the connection to the data source ends unexpectedly. The resulting error state can prevent further changes to the analysis. In this case, you can't add more visuals to the analysis. Check for this state.
- If your visuals don't load, try the following:

- If you are using a corporate network, seek out help from your network administrator and verify that the network's firewall settings permit traffic from *.aws.amazon.com, amazonaws.com, wss://*.aws.amazon.com, and cloudfront.net.
- Add exceptions to your ad blocker for *.aws.amazon.com, amazonaws.com, wss://*.aws.amazon.com, and cloudfront.net.
- If you are using a proxy server, verify that *.quicksight.aws.amazon.com and cloudfront.net are added to the list of approved domains (the allow list).

I get a feedback bar across my printed documents

The browser sometimes prints the document feedback bar across the page, blocking some printed content.

To avoid this problem, use the twirl-down icon on the bottom left of the screen (shown following) to minimize the feedback bar. Then print your document.

My map charts don't show locations

For automatic mapping, called geocoding, to work on map charts, make sure that your data is prepared following specific rules. For help with geospatial issues, see [Geospatial troubleshooting](#). For help with preparing data for geospatial charts, see [Adding geospatial data](#).

My pivot table stops working

If your pivot table exceeds the computational limitations of the underlying database, this is usually caused by the combination of items in the field wells. That is, it's caused by a combination of rows, columns, metrics, and table calculations. To reduce the level of complexity and the potential for errors, simplify your pivot table. For more information, see [Pivot table best practices](#).

My visual can't find missing columns

The visuals in my analysis aren't working as expected. The error message says "The column(s) used in this visual do not exist."

The most common cause of this error is that your data source schema changed. For example, it's possible a column name changed from a_column to b_column.

Depending on how your dataset accesses the data source, choose one of the following.

- If the dataset is based on custom SQL, do one or more of the following:

- Edit the dataset.
- Edit the SQL statement.

For example, if the table name changed from `a_column` to `b_column`, you can update the SQL statement to create an alias: `SELECT b_column as a_column`. By using the alias to maintain the same field name in the dataset, you avoid having to add the column to your visuals as a new entity.

When you're done, choose **Save & visualize**.

- If the dataset isn't based on custom SQL, do one or more of the following:
 - Edit the dataset.
 - For fields that now have different names, rename them in the dataset. You can use the field names from your original dataset. Then open your analysis and add the renamed fields to the affected visuals.

When you're done, choose **Save & visualize**.

My visual can't find the query table

In this case, the visuals in your analysis aren't working as expected. The error message says "Amazon Quick Sight can't find the query table."

The most common cause of this error is that your data source schema changed. For example, it's possible a table name changed from `x_table` to `y_table`.

Depending on how the dataset accesses the data source, choose one of the following.

- If the dataset is based on custom SQL, do one or more of the following:
 - Edit the dataset.
 - Edit the SQL statement.

For example, if the table name changed from `x_table` to `y_table`, you can update the `FROM` clause in the SQL statement to refer to the new table instead.

When you're done, choose **Save & visualize**, then choose each visual and readd the fields as needed.

- If the dataset isn't based on custom SQL, do the following:

1. Create a new dataset using the new table, `y_table` for example.
2. Open your analysis.
3. Replace the original dataset with the newly created dataset. If there are no column changes, all the visuals should work after you replace the dataset. For more information, see [Replacing datasets](#).

My visual doesn't update after I change a calculated field

When you update a calculated field that many other fields depend on, the consuming entities might not update as expected. For example, when you update a calculated field that's used by a field being visualized, the visual doesn't update as expected.

To resolve this issue, refresh your internet browser.

Values in a Microsoft Excel file with scientific notation don't format correctly in Quick Sight

When you connect to a Microsoft Excel file that has a number column that contains values with scientific notation, they might not format correctly in Quick Sight. For example, the value `1.59964E+11`, which is actually `159964032802`, formats as `159964000000` in Quick Sight. This can lead to an incorrect analysis.

To resolve this issue, format the column as Text in Microsoft Excel, and then upload the file to Quick Sight.

Developing with Amazon Quick Sight

We provide API operations for Amazon Quick Sight, and also software development kits (SDKs) for Amazon that enable you to access Amazon Quick Sight from your preferred programming language. Currently, you can manage users and groups. In Enterprise edition, you can also embed dashboards in your webpage or app.

To monitor the calls made to the Amazon Quick Sight API for your account, including calls made by the Amazon Web Services Management Console, command line tools, and other services, use Amazon CloudTrail. For more information, see the [Amazon CloudTrail User Guide](#).

Required knowledge

If you plan to access Amazon Quick Sight through an API, you should be familiar with the following:

- JSON
- Web services
- HTTP requests
- One or more programming languages, such as JavaScript, Java, Python, or C#.

We recommend visiting the Amazon [Getting Started Resource Center](#) for a tour of what Amazon SDKs and toolkits have to offer.

Although you can use a terminal and your favorite text editor, you might benefit from the more visual UI experience you get in an integrated development environment (IDE). We provide a list of IDEs in the *Amazon Getting Started Resource Center* in the [IDE and IDE Toolkits](#) section. This site provides Amazon toolkits that you can download for your preferred IDE. Some IDEs also offer tutorials to help you learn more about programming languages.

Available API operations for Amazon Quick Sight

Amazon provides libraries, sample code, tutorials, and other resources for software developers who prefer to build applications using language-specific API operations instead of submitting a request over HTTPS. These libraries provide basic functions that automatically take care of tasks such as cryptographically signing your requests, retrying requests, and handling error responses. These libraries help make it easier for you to get started.

For more information about downloading the Amazon SDKs, see [Amazon SDKs and Tools](#). The following links are a sample of the language-specific API documentation available.

Amazon Command Line Interface

- [Amazon CLI QuickSight Command Reference](#)
- [Amazon CLI User Guide](#)
- [Amazon CLI Command Reference](#)

Amazon SDK for .NET

- [Amazon.Quicksight](#)
- [Amazon.Quicksight.Model](#)

Amazon SDK for C++

- [Aws::QuickSight::QuickSightClient Class Reference](#)

Amazon SDK for Go

- [quicksight](#)

Amazon SDK for Java

- [com.amazonaws.services.quicksight](#)
- [com.amazonaws.services.quicksight.model](#)

Amazon SDK for JavaScript

- [AWS.QuickSight](#)

Amazon SDK for PHP

- [QuickSightClient](#)

Amazon SDK for Python (Boto3)

- [Amazon Quick Sight](#)

Amazon SDK for Ruby

- [Aws::QuickSight](#)

Terminology and concepts

This section provides a list of terms for development in Amazon Quick Sight.

Anonymous Amazon Quick Sight User: – A temporary Amazon Quick Sight user identity that virtually belongs to a namespace, and is usable only with embedding. You can use tag-based rules to implement row-level security for such users.

Caller identity: – The identity of the Amazon Identity and Access Management user making an API request. The identity of the caller is determined by Amazon Quick Sight using the signature attached to the request. Through the use of our provided SDK clients, no manual steps are necessary to generate the signature or attach it to the requests. However, you can do it manually if you want to.

Invoker identity: – In addition to the caller identity, but not as a replacement for it, you can assume a caller's identity through the IAM AssumeRole API when making calls to Amazon Quick Sight. Amazon approves callers through their invoker's identity. This is done to avoid having to explicitly add multiple accounts belonging to the same Amazon Quick Sight subscription.

Namespace: – a logical container that allows you to isolate user pools so that you can organize clients, subsidiaries, teams, and so on. For more information, see [Supporting multitenancy with isolated namespaces](#)

QuickSight ARN: – Amazon Resource Name (ARN). Amazon Quick Sight resources are identified using their name or ARN. For example, these are the ARNs for a group named MyGroup1, a user named User1, and a dashboard with the ID 1a1ac2b2-3fc3-4b44-5e5d-c6db6778df89:

```
arn:aws-cn:quicksight:us-east-1:111122223333:group/default/MyGroup1
arn:aws-cn:quicksight:us-east-1:111122223333:user/default/User1
arn:aws-cn:quicksight:us-west-2:111122223333:dashboard/1a1ac2b2-3fc3-4b44-5e5d-
c6db6778df89
```

The following examples show ARNs for a template named MyTemplate and a dashboard named MyDashboard.

1. Sample ARN for a template

```
arn:aws-cn:quicksight:us-east-1:111122223333:template/MyTemplate
```

2. Sample ARN for a template, referencing a specific version of the template

```
arn:aws-cn:quicksight:us-east-1:111122223333:template/MyTemplate/version/10
```

3. Sample ARN for a template alias

```
arn:aws-cn:quicksight:us-east-1:111122223333:template/MyTemplate/alias/STAGING
```

4. Sample ARN for a dashboard

```
arn:aws-cn:quicksight:us-east-1:111122223333:dashboard/MyDashboard
```

5. Sample ARN for a dashboard, referencing a specific version of the dashboard

```
arn:aws-cn:quicksight:us-east-1:111122223333:dashboard/MyDashboard/version/10
```

Depending on the scenario, you might need to provide an entity's name, ID, or ARN. You can retrieve the ARN if you have the name, using some of the Amazon Quick Sight API operations.

Amazon Quick Sight dashboard: – An entity which identifies Amazon Quick Sight reports, created from analyses or templates. Amazon Quick Sight dashboards are sharable. With the right permissions, scheduled email reports can be created from them. The `CreateDashboard` and `DescribeDashboard` API Operations act on the dashboard entity.

Amazon Quick Sight template: – An entity which encapsulates the metadata required to create an analysis or a dashboard. It abstracts the dataset associated with the analysis by replacing it with placeholders. Templates can be used to create dashboards by replacing dataset placeholders with datasets that follow the same schema that was used to create the source analysis and template.

Amazon Quick Sight user: – This is an Amazon Quick Sight user identity acted upon by your API call. This user isn't identical to the caller identity but might be the one that maps to the user within Amazon Quick Sight.

Using the developer portal

The Dev Portal helps you learn by example how to use the Amazon Quick Sight API in your web site or application. In this initial offering, the dev portal focuses on API operations for embedded analytics.

The dev portal provides easy-to-use code samples to get you started. You can choose from the following three different use cases:

- Displaying embedded dashboards to everyone (non-authenticated users)
- Personalizing dashboards for your users
- Embedding dashboard authoring

The portal itself displays dashboards by using embedding for everyone.

To get started with the dev portal

1. Open the Dev Portal and choose **Try it** on the use case you want to view.
2. To view code examples, choose **How to embed it** in the menu bar. Then choose each of the following from the navigation pane at left:
 - Configure permissions
 - Get embedding URL (code samples in Java, JavaScript, and Python)
 - Embed URL in your application
3. To download all of the code in a zip file, choose **Download all code**.
4. To customize the dashboard, choose **How to customize it**. This screen is interactive, so you can choose any item in the navigation pane to view the changes live.
5. You can also view and download the html code at bottom left.
6. To return to the start page, click on the Amazon Quick Sight icon, top left.

Developing applications with the Amazon Quick Sight API

You can manage most aspects of your deployment by using the Amazon SDKs to access an API that's tailored to the programming language or platform that you're using. For more information, see [Amazon SDKs](#).

For more information on the API operations, see [Amazon Quick Sight API Reference](#).

Before you can call the Amazon Quick Sight API operations, you need the `quicksight:operation-name` permission in a policy attached to your IAM identity. For example, to call `list-users`, you need the permission `quicksight:ListUsers`. The same pattern applies to all operations.

If you're not sure what the necessary permission is, you can attempt to make a call. The client then tells you what the missing permission is. You can use asterisk (*) in the Resource field of your

permission policy instead of specifying explicit resources. However, we recommended that you restrict each permission as much as possible. You can restrict user access by specifying or excluding resources in the policy, using their Amazon Quick Sight Amazon Resource Name (ARN) identifier.

For more information, see the following:

- [IAM policy examples for Amazon Quick Sight](#)
- [Actions, Resources, and Condition Keys](#)
- [IAM JSON Policy Elements](#)

To retrieve the ARN of a user or a group, use the `Describe` operation on the relevant resource. You can also add conditions in IAM to further restrict access to an API in some scenarios. For instance, when adding `User1` to `Group1`, the main resource is `Group1`, so you can allow or deny access to certain groups, but you can also add a condition by using the IAM Amazon Quick Sight key `quicksight:UserName` to allow or prevent certain users from being added to that group.

Following is an example policy. It means that the caller with this policy attached, is able to invoke the `CreateGroupMembership` operation on any group, provided that the user name they are adding to the group is not `user1`.

```
{
  "Effect": "Allow",
  "Action": "quicksight:CreateGroupMembership",
  "Resource": "arn:aws-cn:quicksight:us-east-1:aws-account-id:group/default/*",
  "Condition": {
    "StringNotEquals": {
      "quicksight:UserName": "user1"
    }
  }
}
```

Amazon CLI

The following procedure explains how to interact with Amazon Quick Sight API operations through the Amazon CLI. The following instructions have been tested in Bash but should be identical or similar in other command-line environments.

1. Install Amazon SDK in your environment. Instructions on how to do that are located here: [Amazon Command line Interface](#).

2. Set up your Amazon CLI identity and region using the following command and follow-up instructions. Use the credentials for an IAM identity or role that has the proper permissions.

```
aws configure
```

3. Look at the Amazon Quick Sight SDK help by issuing the following command:

```
aws quicksight help
```

4. To get detailed instructions on how to use an API, enter its name followed by help, like so:

```
aws quicksight list-users help
```

5. Now you can call an Amazon Quick Sight API operation. This example returns a list of Amazon Quick Sight users in your account.

```
aws quicksight list-users --aws-account-id aws-account-id --namespace default --  
region us-east-1
```

Java SDK

Use the following procedure to set up a Java app that interacts with Amazon Quick Sight.

1. To get started, create a Java project in your IDE.
2. Import the Amazon Quick Sight SDK into your new project, for example:
`AWSQuickSightJavaClient-1.11.x.jar`
3. Once your IDE indexes the Amazon Quick Sight SDK, you should be able to add an import line as follows:

```
import com.amazonaws.services.quicksight.AmazonQuickSight;
```

If your IDE doesn't recognize this as valid, verify that you imported the SDK.

4. Like other Amazon SDKs, Amazon Quick Sight SDK requires external dependencies to perform many of its functions. You need to download and import those into the same project. The following dependencies are required:

- `aws-java-sdk-1.11.402.jar` (Amazon Java SDK and credentials setup) — See [Set up the Amazon SDK for Java](#)

- commons-logging-1.2.jar — See https://commons.apache.org/proper/commons-logging/download_logging.cgi
 - jackson-annotations-2.9.6.jar, jackson-core-2.9.6.jar, and jackson-databind-2.9.6.jar — See <http://repo1.maven.org/maven2/com/fasterxml/jackson/core/>
 - httpclient-4.5.6.jar, httpcore-4.4.10.jar — See <https://hc.apache.org/downloads.cgi>
 - joda-time-2.1.jar — See <https://mvnrepository.com/artifact/joda-time/joda-time/2.1>
5. Now, you are ready to create an Amazon Quick Sight client. You can use a default public endpoint that the client can communicate with or you can reference the endpoint explicitly. There are multiple ways to provide your Amazon credentials. In the following example, we provide a direct, simple approach. The following client method is used to make all the API calls that follow:

```
private static AmazonQuickSight getClient() {
    final AWSCredentialsProvider credsProvider = new AWSCredentialsProvider() {
        @Override
        public AWSCredentials getCredentials() {
            // provide actual IAM access key and secret key here
            return new BasicAWSCredentials("access-key", "secret-key");
        }

        @Override
        public void refresh() {}
    };

    return AmazonQuickSightClientBuilder
        .standard()
        .withRegion(Regions.US_EAST_1.getName())
        .withCredentials(credsProvider)
        .build();
}
```

6. Now, we can use the above client to list all the users in our Amazon Quick Sight account.

Note

You have to provide the Amazon account ID that you used to subscribe to Amazon Quick Sight. This must match the Amazon account ID of the caller's identity. Cross-account calls aren't supported at this time. Furthermore, the required parameter namespace should always be set to *default*.

```
getClient().listUsers(new ListUsersRequest()
    .withAwsAccountId("relevant_AWS_account_ID")
    .withNamespace("default"))
    .getUserList().forEach(user -> {
        System.out.println(user.getArn());
    });
```

7. To see a list of all possible API operations and the request objects they use, you can **CTRL-click** on the client object in your IDE in order to view the Amazon Quick Sight interface. Alternatively, find it within the `com.amazonaws.services.quicksight` package in the Amazon Quick Sight JavaClient JAR file.

JavaScript (Node.js) SDK

Use the following procedure to interact with Amazon Quick Sight using Node.js.

1. Set up your node environment using the following commands:
 - `npm install aws-sdk`
 - `npm install aws4`
 - `npm install request`
 - `npm install url`
2. For information on configuring the Node.js with Amazon SDK and setting your credentials, see--> the [Amazon SDK for JavaScript Developer Guide for SDK v2](#).
3. Use the following code sample to test your setup. HTTPS is required. The sample displays a full listing of Amazon Quick Sight operations along with their URL request parameters, followed by a list of Amazon Quick Sight users in your account.

```
const Amazon = require('aws-sdk');
```

```
const https = require('https');

var quicksight = new Amazon.Service({
  apiConfig: require('./quicksight-2018-04-01.min.json'),
  region: 'us-east-1',
});

console.log(quicksight.config.apiConfig.operations);

quicksight.listUsers({
  // Enter your actual Amazon account ID
  'AwsAccountId': 'relevant_AWS_account_ID',
  'Namespace': 'default',
}, function(err, data) {
  console.log('---');
  console.log('Errors: ');
  console.log(err);
  console.log('---');
  console.log('Response: ');
  console.log(data);
});
```

Python3 SDK

Use the following procedure to create a custom built botocore package to interact with Amazon Quick Sight.

1. Create a credentials file in the Amazon directory for your environment. In a Linux/Mac-based environment, that file is called `~/.aws/credentials` and looks like this:

```
[default]
aws_access_key_id = Your_IAM_access_key
aws_secret_access_key = Your_IAM_secret_key
```

2. Unzip the folder `botocore-1.12.10`. Change directory into `botocore-1.12.10` and enter the Python3 interpreter environment.
3. Responses come back as a dictionary object. They each have a `ResponseMetadata` entry that contains request IDs and response status. Other entries are based on what type of operation you run.

4. The following example is a sample app that first creates, deletes, and lists groups. Then, it lists users in a Quicksight account:

```
import boto3.session
default_namespace = 'default'
account_id = 'relevant_AWS_Account'

session = boto3.session.get_session()
client = session.create_client("quicksight", region_name='us-east-1')

print('Creating three groups: ')
client.create_group(AwsAccountId = account_id, Namespace=default_namespace,
    GroupName='MyGroup1')
client.create_group(AwsAccountId = account_id, Namespace=default_namespace,
    GroupName='MyGroup2')
client.create_group(AwsAccountId = account_id, Namespace=default_namespace,
    GroupName='MyGroup3')

print('Retrieving the groups and listing them: ')
response = client.list_groups(AwsAccountId = account_id,
    Namespace=default_namespace)
for group in response['GroupList']:
    print(group)

print('Deleting our groups: ')
client.delete_group(AwsAccountId = account_id, Namespace=default_namespace,
    GroupName='MyGroup1')
client.delete_group(AwsAccountId = account_id, Namespace=default_namespace,
    GroupName='MyGroup2')
client.delete_group(AwsAccountId = account_id, Namespace=default_namespace,
    GroupName='MyGroup3')

response = client.list_users(AwsAccountId = account_id,
    Namespace=default_namespace)
for user in response['UserList']:
    print(user)
```

.NET/C# SDK

Use the following procedure to interact with Amazon Quick Sight using C#.NET. This example is constructed on Microsoft Visual for Mac; the instructions can vary slightly based on your IDE and platform. However, they should be similar.

1. Unzip the `nuget.zip` file into a folder called `nuget`.
2. Create a new **Console app** project in Visual Studio.
3. Under your solution, locate app **Dependencies**, then open the context (right-click menu) and choose **Add Packages**.
4. In the sources list, choose **Configure Sources**.
5. Choose **Add**, and name the source `QuickSightSDK`. Browse to the `nuget` folder and choose **Add Source**.
6. Choose **OK**. Then, with `QuickSightSDK` selected, select all three Amazon Quick Sight packages:
 - `AWSSDK.QuickSight`
 - `AWSSDK.Extensions.NETCore.Setup`
 - `AWSSDK.Extensions.CognitoAuthentication`
7. Click **Add Package**.
8. Copy and paste the following sample app into your console app editor.

```
using System;
using Amazon.QuickSight.Model;
using Amazon.QuickSight;

namespace DotNetQuickSightSDKTest
{
    class Program
    {
        private static readonly string AccessKey = "insert_your_access_key";
        private static readonly string SecretAccessKey =
            "insert_your_secret_key";
        private static readonly string AccountID = "AWS_account_ID";
        private static readonly string Namespace = "default"; // leave this as
        default

        static void Main(string[] args)
        {
            var client = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
                Amazon.RegionEndpoint.USEast1);
```

```
var listUsersRequest = new ListUsersRequest
{
    AwsAccountId = AccountID,
    Namespace = Namespace
};

client.ListUsersAsync(listUsersRequest).Result.UserList.ForEach(
    user => Console.WriteLine(user.Arn)
);

var listGroupsRequest = new ListGroupsRequest
{
    AwsAccountId = AccountID,
    Namespace = Namespace
};

client.ListGroupsAsync(listGroupsRequest).Result.GroupList.ForEach(
    group => Console.WriteLine(group.Arn)
);
    }
}
}
```

Amazon Quick Sight events integration

With Amazon EventBridge, you can respond automatically to events in Amazon Quick Sight such as new dashboard creation or updates. These events are delivered to EventBridge in near real time. As a developer, you can write simple rules to indicate which events are of interest, and what actions to take when an event matches a rule. By using events, you can complete use cases such as continuous backup and deployment.

Topics

- [Supported events](#)
- [Example event payload](#)
- [Creating rules to send Amazon Quick Sight events to Amazon CloudWatch](#)
- [Creating rules to send Amazon Quick Sight events to Amazon Lambda](#)

Supported events

Amazon Quick Sight currently supports the following events.

Asset type	Action	Event detail type	Event detail
Dashboard	Create	Amazon Quick Sight Dashboard Creation Successful	<pre>{ "dashboardId": "6fdbbc328-ebbd-457f-aa02-9780173afc83", "versionNumber": 1 }</pre>
Dashboard	Create	Amazon Quick Sight Dashboard Creation Failed	<pre>{ "dashboardId": "6fdbbc328-ebbd-457f-aa02-9780173afc83", "versionNumber": 1, "errors": [{ "Type": "PARAMETER_NOT_FOUND", "Message": "Missing property abc" }, { "Type": "DATA_SET_NOT_FOUND", "Message": "Cannot find dataset with id abc" }] }</pre>

Asset type	Action	Event detail type	Event detail
			<pre>] }</pre>
Dashboard	Create	Amazon Quick Sight Dashboard Permissions Updated	<pre>{ "dashboardId": "6fdbbc328-ebbd-457f-aa02-9780173afc83" }</pre>
Dashboard	Update	Amazon Quick Sight Dashboard Update Successful	<pre>{ "dashboardId": "6fdbbc328-ebbd-457f-aa02-9780173afc83", "versionNumber": 1 }</pre>

Asset type	Action	Event detail type	Event detail
Dashboard	Update	Amazon Quick Sight Dashboard Update Failed	<pre>{ "dashboardId": "6fdbbc328-ebbd-457f-aa02-9780173afc83", "versionNumber": 1, "errors": [{ "Type": "PARAMETER_NOT_FOUND", "Message": "Missing property abc" }, { "Type": "DATA_SET_NOT_FOUND", "Message": "Cannot find dataset with id abc" }] }</pre>
Dashboard	Update	Amazon Quick Sight Dashboard Permissions Updated	<pre>{"dashboardId": "6fdbbc328-ebbd-457f-aa02-9780173afc83"}</pre>

Asset type	Action	Event detail type	Event detail
Dashboard	Publish	Amazon Quick Sight Dashboard Published Version Updated	<pre>{ "dashboardId": "6fdbbc328-ebbd-457f-aa02-9780173afc83", "versionNumber": 2 }</pre>
Dashboard	Delete	Amazon Quick Sight Dashboard Deleted	<pre>{ "dashboardId": "6fdbbc328-ebbd-457f-aa02-9780173afc83" }</pre>
Analysis	Create	Amazon Quick Sight Analysis Creation Successful	<pre>{ "analysisId": "e5f37119-e24c-4874-901a-af9032b729b5" }</pre>

Asset type	Action	Event detail type	Event detail
Analysis	Create	Amazon Quick Sight Analysis Creation Failed	<pre>{ "analysisId": "e5f37119-e24c-4874-901a-af9032b729b5", "errors": [{ "Type": "PARAMETER_NOT_FOUND", "Message": "Missing property abc" }, { "Type": "DATA_SET_NOT_FOUND", "Message": "Cannot find dataset with id abc" }] }</pre>
Analysis	Create	Amazon Quick Sight Analysis Permissions Updated	<pre>{"analysisId": "e5f37119-e24c-4874-901a-af9032b729b5" }</pre>
Analysis	Delete	Amazon Quick Sight Analysis Deleted	<pre>{ "analysisId": "e5f37119-e24c-4874-901a-af9032b729b5" }</pre>

Asset type	Action	Event detail type	Event detail
Analysis	Update	Amazon Quick Sight Analysis update successful	<pre>{ "analysis Id": "e5f37119- e24c-4874-901a- af9032b729b5" }</pre>

Asset type	Action	Event detail type	Event detail
Analysis	Update	Amazon Quick Sight Analysis update failed	<pre>{ "analysis Id": "e5f37119- e24c-4874-901a- af9032b729b5", "errors": [{ "Type": "PARAMETE R_NOT_FOUND", "Message" : "Missing property abc" }, { "Type": "DATA_SET _NOT_FOUND", "Message": "Cannot find dataset with id abc" }] }</pre>

Asset type	Action	Event detail type	Event detail
Analysis	Update	Amazon Quick Sight Analysis Permissions Updated	<pre>{ "analysisId": "e5f37119- e24c-4874-901a- af9032b729b5" }</pre>
VPC connection	Create	Amazon Quick Sight VPC Connection Creation Successful	<pre>{ "vpcConne ctionId": "53d34238 -57e7-488d- b99a-a0037d275a4 e", "availabi lityStatu s": "CREATION _SUCCESSFUL" }</pre>
VPC connection	Create	Amazon Quick Sight VPC Connection Creation Failed	<pre>{ "vpcConne ctionId": "53d34238 -57e7-488d- b99a-a0037d275a4 e", "availabi lityStatu s": "CREATION _FAILED" }</pre>

Asset type	Action	Event detail type	Event detail
VPC connection	Update	Amazon Quick Sight VPC Connection Update Successful	<pre>{ "vpcConnectionId": "53d34238-57e7-488d-b99a-a0037d275a4e", "availabilityStatus": "UPDATE_SUCCESSFUL" }</pre>
VPC connection	Update	Amazon Quick Sight VPC Connection Update Failed	<pre>{ "vpcConnectionId": "53d34238-57e7-488d-b99a-a0037d275a4e", "availabilityStatus": "UPDATE_FAILED" }</pre>
VPC connection	Delete	Amazon Quick Sight VPC Connection Deletion Successful	<pre>{ "vpcConnectionId": "53d34238-57e7-488d-b99a-a0037d275a4e", "availabilityStatus": "DELETED" }</pre>

Asset type	Action	Event detail type	Event detail
VPC connection	Delete	Amazon Quick Sight VPC Connection Deletion Failed	<pre>{ "vpcConnectionId": "53d34238-57e7-488d-b99a-a0037d275a4e", "availabilityStatus": "DELETION_FAILED" }</pre>
Folder	Create	Amazon Quick Sight Folder Created	<pre>{ "folderId": "77e307e8-b41b-472a-90e8-fe3f471537be", "parentFolderArn": "arn:aws-cn:quicksight:us-east-1:123456789012:folder/098765432134" }</pre>
Folder	Create	Amazon Quick Sight Folder Permissions Updated	<pre>{"folderId": "77e307e8-b41b-472a-90e8-fe3f471537be" }</pre>

Asset type	Action	Event detail type	Event detail
Folder	Update	Amazon Quick Sight Folder Updated	<pre>{ "folderId": "77e307e8-b41b-472a-90e8-fe3f471537be" }</pre>
Folder	Update	Amazon Quick Sight Folder Permissions Updated	<pre>{"folderId": "77e307e8-b41b-472a-90e8-fe3f471537be" }</pre>
Folder	Delete	Amazon Quick Sight Folder Deleted	<pre>{ "folderId": "77e307e8-b41b-472a-90e8-fe3f471537be" }</pre>
Folder	Membership update	Amazon Quick Sight Folder Membership Updated	<pre>{ "folderId": "77e307e8-b41b-472a-90e8-fe3f471537be", "membersAdded": ["arn:aws-cn:quick-sight:us-east-1:123456789012:analysis/e5f37119-e24c-4874-901a-af9032b729b5"], "membersRemoved": [] }</pre>

Asset type	Action	Event detail type	Event detail
Dataset	Create	Amazon Quick Sight Dataset Created	<pre>{ "datasetId": "a6553a81-f97e-4ffa-a860-baea63196efa" }</pre>
Dataset	Create	Amazon Quick Sight Dataset Permissions Updated	<pre>{"datasetId": "a6553a81-f97e-4ffa-a860-baea63196efa" }</pre>
Dataset	Update	Amazon Quick Sight Dataset Updated	<pre>{ "datasetId": "a6553a81-f97e-4ffa-a860-baea63196efa" }</pre>
Dataset	Update	Amazon Quick Sight Dataset Permissions Updated	<pre>{"datasetId": "a6553a81-f97e-4ffa-a860-baea63196efa" }</pre>
Dataset	Delete	Amazon Quick Sight Dataset Deleted	<pre>{ "datasetId": "a6553a81-f97e-4ffa-a860-baea63196efa" }</pre>

Asset type	Action	Event detail type	Event detail
DataSource	Create	Amazon Quick Sight DataSource Creation Successful	<pre>{ "datasourceId": "230caa6e-dc87-406b-91fb-037f29c32824" }</pre>
DataSource	Create	Amazon Quick Sight DataSource Creation Failed	<pre>{ "datasourceId": "230caa6e-dc87-406b-91fb-037f29c32824", "error": { "message": "AMAZON_ELASTICSEARCH engine version 7.4 is lower than minimum supported version 7.7", "type": "ENGINE_VERSION_NOT_SUPPORTED" } }</pre>
DataSource	Create	Amazon Quick Sight DataSource Permissions Updated	<pre>{"datasourceId": "230caa6e-dc87-406b-91fb-037f29c32824" }</pre>

Asset type	Action	Event detail type	Event detail
DataSource	Update	Amazon Quick Sight DataSource Update Successful	<pre>{ "datasourceId": "230caa6e-dc87-406b-91fb-037f29c32824" }</pre>
DataSource	Update	Amazon Quick Sight DataSource Update Failed	<pre>{ "datasourceId": "230caa6e-dc87-406b-91fb-037f29c32824", "error": { "message": "AMAZON_ELASTICSEARCH engine version 7.4 is lower than minimum supported version 7.7", "type": "ENGINE_VERSION_NOT_SUPPORTED" } }</pre>
DataSource	Update	Amazon Quick Sight DataSource Permissions Updated	<pre>{"datasourceId": "230caa6e-dc87-406b-91fb-037f29c32824" }</pre>

Asset type	Action	Event detail type	Event detail
DataSource	Delete	Amazon Quick Sight DataSource Deleted	<pre>{ "datasourceId": "230caa6e-dc87-406b-91fb-037f29c32824" }</pre>
Theme	Create	Amazon Quick Sight Theme Creation Successful	<pre>{ "themeId": "6fdb328-ebbd-457f-aa02-9780173afc83", "versionNumber": 1 }</pre>
Theme	Create	Amazon Quick Sight Theme Creation Failed	<pre>{ "themeId": "6fdb328-ebbd-457f-aa02-9780173afc83", "versionNumber": 1 }</pre>
Theme	Create	Amazon Quick Sight Theme Permissions Updated	<pre>{"themeId": "6fdb328-ebbd-457f-aa02-9780173afc83" }</pre>

Asset type	Action	Event detail type	Event detail
Theme	Update	Amazon Quick Sight Theme Update Successful	<pre>{ "themeId" : "6fdb328 -ebbd-457f- aa02-9780173afc8 3", "versionN umber": 2 }</pre>
Theme	Update	Amazon Quick Sight Theme Update Failed	<pre>{ "themeId" : "6fdb328 -ebbd-457f- aa02-9780173afc8 3", "versionN umber": 2 }</pre>
Theme	Update	Amazon Quick Sight Theme Permissions Updated	<pre>{"themeId": "6fdb328- ebbd-457f- aa02-9780173afc8 3" }</pre>
Theme	Delete	Amazon Quick Sight Theme deleted	<pre>{ "themeId" : "6fdb328 -ebbd-457f- aa02-9780173afc8 3" }</pre>

Asset type	Action	Event detail type	Event detail
Theme	Alias Create	Amazon Quick Sight Theme Alias Created	<pre>{ "themeId" : "6fdb328 -ebbd-457f- aa02-9780173afc8 3", "aliasName": "MyThemeAlias" "versionN umber": 2 }</pre>
Theme	Alias Update	Amazon Quick Sight Alias Updated	<pre>{ "themeId" : "6fdb328 -ebbd-457f- aa02-9780173afc8 3", "aliasName": "MyThemeAlias" "versionN umber": 4 }</pre>
Theme	Alias Delete	Amazon Quick Sight Theme Alias Deleted	<pre>{ "themeId" : "6fdb328 -ebbd-457f- aa02-9780173afc8 3", "aliasName": "MyThemeAlias" "versionN umber": 2 }</pre>

Example event payload

All events follow the standard EventBridge [object structure](#). The detail field is a JSON object that contains more information about the event.

```
{
  "version": "0",
  "id": "3acb26c8-397c-4c89-a80a-ce672a864c55",
  "detail-type": "QuickSight Dashboard Creation Successful",
  "source": "aws.quicksight",
  "account": "123456789012",
  "time": "2023-10-30T22:06:31Z",
  "region": "us-east-1",
  "resources": ["arn:aws-cn:quicksight:us-east-1:123456789012:dashboard/6fdb328-ebbd-457f-aa02-9780173afc83"],
  "detail": {
    "dashboardId": "6fdb328-ebbd-457f-aa02-9780173afc83",
    "versionNumber": 1
  }
}
```

Creating rules to send Amazon Quick Sight events to Amazon CloudWatch

You can write simple rules to indicate which Amazon Quick Sight events interest you and which automated actions to take when an event matches a rule. For example, you can configure Amazon Quick Sight to send events to Amazon CloudWatch whenever a Amazon Quick Sight asset is placed in a folder. For more information, see the [Amazon EventBridge user guide](#).

1. Sign in to the Amazon Web Services Management Console and open the CloudWatch console at <https://console.amazonaws.cn/cloudwatch/>.
2. Under **Events** in the navigation pane, choose **Rules**.
3. Choose **Create rule**.
4. Enter a name and description for the rule. The rule name must be unique within this Region. For example, enter QuickSightAssetChangeRuleCloudWatch.
5. Choose **default** Event bus.
6. Choose **Rule with an event pattern**, and then choose **Next**.
7. For **Event source**, choose **Amazon events or EventBridge partner events**.
8. In the **Creation method** section, choose **Custom pattern (JSON editor)**.

9. In the **Event pattern** text box, enter the following snippet and choose **Next**.

```
{
  "source": ["aws.quicksight"]
}
```

Alternatively, you can create the rule that only subscribes to a subset of event types in Amazon Quick Sight. For example, the following rule will only triggered when an asset is added to or removed from a folder with id 77e307e8-b41b-472a-90e8-fe3f471537be.

```
{
  "source": ["aws.quicksight"],
  "detail-type": ["QuickSight Folder Membership Updated"],
  "detail": {
    "folderId": "77e307e8-b41b-472a-90e8-fe3f471537be"
  }
}
```

10. For **Targets**, choose **Amazon service > CloudWatch log group**.
11. Choose from an existing log group or create a new one by entering a new log group name.
12. Optionally, you can add another target for this rule.
13. In **Configure tags**, choose **Next**.
14. Choose **Create rule**.

For more information, see [Creating Amazon EventBridge rule that reacts To events](#) in the Amazon EventBridge user guide.

Creating rules to send Amazon Quick Sight events to Amazon Lambda

In this tutorial, you create an Amazon Lambda function that logs the asset events in the Amazon Quick Sight account. You then create a rule that runs the function whenever there is an asset change. This tutorial assumes that you already signed up for Amazon Quick Sight.

Step 1: Create a Lambda function

Create a Lambda function to log the state change events. You specify this function when you create your rule.

1. Sign in to the Amazon Web Services Management Console and open the Amazon Lambda console at <https://console.amazonaws.cn/lambda/>.
2. If you're new to Lambda, you see a welcome page. Choose **Get Started Now**. Otherwise, choose **Create function**.
3. Choose **Author from scratch**.
4. On the Create function page, enter a name and description for the Lambda function. For example, name the function `QuickSightAssetChangeFn`.
5. In **Runtime**, select **Node.js 18.x**.
6. For **Architecture**, choose **x86_64**.
7. For **Execution role**, choose either **Create a new role with basic Lambda permissions** or **Use an existing role** and choose the role you want.
8. Choose **Create function**.
9. On the **QuickSightAssetChange** page, choose **index.js**.
10. In the **index.js** pane, delete the existing code.
11. Enter the following code snippet.

```
console.log('Loading function');
exports.handler = async (event, context) => {
  console.log('Received QuickSight event:', JSON.stringify(event));
};
```

12. Choose **Deploy**.

Step 2: Create a rule

Create a rule to run your Lambda function whenever you create/update/delete a Amazon Quick Sight asset.

1. Sign in to the Amazon Web Services Management Console and open the Amazon EventBridge console at <https://console.amazonaws.cn/events/>.
2. In the navigation pane, choose **Rules**.
3. Choose **Create rule**.
4. Enter a name and description for the rule. For example, enter `QuickSightAssetChangeRule`.
5. Select **default** Event bus.

6. Choose **Rule with an event pattern**, and then choose **Next**.
7. For **Event source**, choose **Amazon events or EventBridge partner events**.
8. In the **Creation method** section, choose **Custom pattern (JSON editor)**.
9. In the **Event pattern** text box, enter the following snippet and choose **Next**.

```
{
  "source": ["aws.quicksight"]
}
```

Alternatively, you can create the rule that only subscribes to a subset of event types in Amazon Quick Sight. For example, the following rule will only triggered when an asset is added to or removed from a folder with id 77e307e8-b41b-472a-90e8-fe3f471537be.

```
{
  "source": ["aws.quicksight"],
  "detail-type": ["QuickSight Folder Membership Updated"],
  "detail": {
    "folderId": "77e307e8-b41b-472a-90e8-fe3f471537be"
  }
}
```

10. For **Target types**, choose **Amazon service** and **Lambda function**.
11. For **Function**, choose the Lambda function that you created. Then choose **Next**.
12. In **Configure tags**, choose **Next**.
13. Review the steps in your rule. Then choose **Create rule**.

Step 3: Test the rule

To test your rule, create an analysis. After waiting a minute, verify that your Lambda function was invoked.

1. Open the Amazon Quick Sight console at <https://quicksight.aws.amazon.com/>.
2. Create a new analysis.
3. In the navigation pane, choose **Rules**, choose the name of the rule that you created.
4. In **Rule details**, choose **Monitoring**.
5. You will be redirected to the Amazon CloudWatch console. If you are not redirected, choose **View the metrics in CloudWatch**.

6. In **All metrics**, choose the name of the rule that you created. The graph indicates that the rule was invoked.
7. In the navigation pane, choose **Log groups**.
8. Choose the name of the log group for your Lambda function. For example, `/aws/lambda/function-name`.
9. Choose the name of the log stream to view the data provided by the function for the instance that you launched. You should see a received event similar to the following:

```
{
  "version": "0",
  "id": "3acb26c8-397c-4c89-a80a-ce672a864c55",
  "detail-type": "QuickSight Analysis Creation Successful",
  "source": "aws.quicksight",
  "account": "123456789012",
  "time": "2023-10-30T22:06:31Z",
  "region": "us-east-1",
  "resources": ["arn:aws-cn:quicksight:us-east-1:123456789012:analysis/e5f37119-e24c-4874-901a-af9032b729b5"],
  "detail": {
    "analysisId": "e5f37119-e24c-4874-901a-af9032b729b5"
  }
}
```

For an example of Amazon Quick Sight event in JSON format, see [Overview of events for Amazon Quick Sight](#).

Embedded analytics for Amazon Quick Sight

Important

Amazon Quick Sight has new API operations for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` API operations to embed dashboards and the Amazon Quick Sight console, but they don't contain the latest embedding capabilities. For more information about embedding using the old API operations, see [Embedding analytics using the `GetDashboardEmbedURL` and `GetSessionEmbedURL` API operations](#).

Applies to: Enterprise Edition**Intended audience:** Amazon Quick Suite developers

With Amazon Quick Sight embedded analytics, you can seamlessly integrate data-driven experiences into your software applications. You can style the embedded components to match your brand. This capability brings the power of Amazon Quick Sight to your end users, where they can analyze and interact with data without ever leaving the application. Improving the user experience by reducing cognitive complexity gives users a better opportunity for deeper understanding and effectiveness.

Amazon Quick Sight supports embedding for these elements:

- Amazon Quick Sight console (full authoring experience for registered users)
- Amazon Quick Sight dashboards and visuals (for registered users, anonymous users, public end users)
- Amazon Quick Sight Q search bar (for registered users and anonymous users)

With an embedded Amazon Quick Sight console, you embed the full Amazon Quick Sight experience. Doing this makes it possible to use Amazon Quick Sight authoring tools as part of your application, rather than in the context of the Amazon Web Services Management Console or a standalone website. Users of an embedded Amazon Quick Sight console need to be registered as Amazon Quick Sight authors or admins in your Amazon Web Services account. They also need to be authenticated into the same Amazon Web Services account, using any of the Amazon Quick Sight-supported authentication methods.

With an embedded Amazon Quick Sight dashboard or visual, readers get the same functionality and interactivity as they do in a published dashboard or visual. To use an embedded dashboard or visual, readers (viewers) can include any of the following:

- Amazon Quick Sight users authenticated in your Amazon Web Services account by any method supported by Amazon Quick Sight.
- Unauthenticated visitors to a website or application – This option requires session packs with capacity pricing . For information about subscription types, see [Understanding Amazon Quick Sight subscriptions and roles](#).

- Multiple end users viewing a display on monitors or large screens by programmatic access.

If your app also resides in Amazon, the app doesn't need to reside on the same Amazon Web Services account as the Amazon Quick Sight subscription. However, the app needs to be able to assume the Amazon Identity and Access Management (IAM) role that you use for the API calls.

Before you can embed content, make sure that you're using Amazon Quick Sight Enterprise edition in the Amazon Web Services account where you plan to use embedding.

Amazon Quick Sight embedding is available in all supported Amazon Web Services Regions.

Topics

- [Embedding Amazon Quick Sight analytics into your applications](#)
- [Embedding custom Amazon Quick Sight assets into your application](#)
- [Embedding Amazon Quick Sight visuals and dashboards with a 1-click embed code](#)
- [Embedding with the Amazon Quick Sight APIs](#)

Embedding Amazon Quick Sight analytics into your applications

Applies to: Enterprise Edition

To embed analytics, you can run the Amazon Quick Sight embedding API to generate the embed code. Alternatively for dashboards, you can copy an embed code when you share the dashboard in Amazon Quick Sight. Each option is described below.

1-click embedding for registered users

When you share a dashboard with registered users in your account, you can copy an embed code for the dashboard and paste it into your internal application's HTML.

Using 1-click enterprise embedding is best for when you want to embed a Amazon Quick Sight dashboard in an internal application that users need to authenticate in to. When you copy the embed code, you get a static embed code that doesn't change.

For more information, see [Embedding Amazon Quick Sight visuals and dashboards for registered users with a 1-click embed code](#).

Embedding with the Amazon Quick Sight APIs

Embedding with the Amazon Quick Sight API is best for when you want to embed the Amazon Quick Sight experience in an internal application that users must authenticate in to, or an external application that anyone can access. When you use the embedding API operations to generate an embed code, you get a one-time code.

For more information, see [Embedding with the Amazon Quick Sight APIs](#).

Embedding custom Amazon Quick Sight assets into your application

You can use Amazon Quick Sight embedded analytics to embed custom Amazon Quick Sight assets into your application that are tailored to meet your business needs. For embedded dashboards and visuals, Amazon Quick Sight authors can add filters and drill downs that readers can access as they navigate the dashboard or visual. Amazon Quick Sight developers can also use the Amazon Quick Sight SDKs to build tighter integrations between their SaaS applications and their Amazon Quick Sight embedded assets to add datapoint callback actions to visuals on a dashboard at runtime.

For more information about the Amazon Quick Sight SDKs, see the `amazon-quicksight-embedding-sdk` on [GitHub](#) or [NPM](#).

Following, you can find descriptions of how to use the Amazon Quick Sight SDKs to customize your Amazon Quick Sight embedded analytics.

Topics

- [Adding embedded callback actions at runtime in Amazon Quick Sight](#)
- [Filtering data at runtime for Amazon Quick Sight embedded dashboards and visuals](#)
- [Customize the look and feel of Amazon Quick Sight embedded dashboards and visuals](#)
- [Using the Amazon Quick Sight Embedding SDK to enable shareable links to embedded dashboard views](#)

Adding embedded callback actions at runtime in Amazon Quick Sight

Use embedded datapoint callback actions to build tighter integrations between your software as a service (SaaS) application and your Amazon Quick Sight embedded dashboards and visuals. Developers can register datapoints to be called back with the Amazon Quick Sight embedding SDK. When you register a callback action for a visual, readers can select a datapoint on the visual to receive a callback that provides data specific to the selected data point. This information can be

used to flag key records, compile raw data specific to the datapoint, capture records, and compile data for backend processes.

Embedded callbacks aren't supported for custom visual content, text boxes, or insights.

Before you begin registering datapoints for callback, update the Embedding SDK to version 2.3.0. For more information about using the Amazon Quick Sight Embedding SDK, see the [amazon-quicksight-embedding-sdk](#) on GitHub.

A datapoint callback can be registered to one or more visuals at runtime through the Amazon Quick Sight SDK. You can also register a datapoint callback to any interaction supported by the [VisualCustomAction](#) API structure. This allows the datapoint callback to initiate when the user selects the datapoint on the visual or when the datapoint is selected from the datapoint context menu. The following example registers a datapoint callback that the reader initiates when they select a datapoint on the visual.

```
/const MY_GET_EMBED_URL_ENDPOINT =  
  "https://my.api.endpoint.domain/MyGetEmbedUrlApi"; // Sample URL  
  
// The dashboard id to embed  
const MY_DASHBOARD_ID = "my-dashboard"; // Sample ID  
  
// The container element in your page that will have the embedded dashboard  
const MY_DASHBOARD_CONTAINER = "#experience-container"; // Sample ID  
  
// SOME HELPERS  
  
const ActionTrigger = {  
  DATA_POINT_CLICK: "DATA_POINT_CLICK",  
  DATA_POINT_MENU: "DATA_POINT_MENU",  
};  
  
const ActionStatus = {  
  ENABLED: "ENABLED",  
  DISABLED: "DISABLED",  
};  
  
// This function makes a request to your endpoint to obtain an embed url for a given  
// dashboard id  
// The example implementation below assumes the endpoint takes dashboardId as request  
// data  
// and returns an object with EmbedUrl property
```

```

const myGetEmbedUrl = async (dashboardId) => {
  const apiOptions = {
    dashboardId,
  };
  const apiUrl = new URL(MY_GET_EMBED_URL_ENDPOINT);
  apiUrl.search = new URLSearchParams(apiOptions).toString();
  const apiResponse = await fetch(apiUrl.toString());
  const apiResponseData = await apiResponse.json();
  return apiResponseData.EmbedUrl;
};

// This function constructs a custom action object
const myConstructCustomActionModel = (
  customActionId,
  actionName,
  actionTrigger,
  actionStatus
) => {
  return {
    Name: actionName,
    CustomActionId: customActionId,
    Status: actionStatus,
    Trigger: actionTrigger,
    ActionOperations: [
      {
        CallbackOperation: {
          EmbeddingMessage: {},
        },
      },
    ],
  };
};

// This function adds a custom action on the first visual of first sheet of the
// embedded dashboard
const myAddVisualActionOnFirstVisualOfFirstSheet = async (
  embeddedDashboard
) => {
  // 1. List the sheets on the dashboard
  const { SheetId } = (await embeddedDashboard.getSheets())[0];
  // If you'd like to add action on the current sheet instead, you can use
  // getSelectedSheetId method
  // const SheetId = await embeddedDashboard.getSelectedSheetId();

```



```

// 2. List the visuals on the specified sheet
const { VisualId } = (await embeddedDashboard.getSheetVisuals(SheetId))[0];

// 3. Add the custom action to the visual
try {
  const customActionId = "custom_action_id"; // Sample ID
  const actionName = "Flag record"; // Sample name
  const actionTrigger = ActionTrigger.DATA_POINT_CLICK; // or
ActionTrigger.DATA_POINT_MENU
  const actionStatus = ActionStatus.ENABLED;
  const myCustomAction = myConstructCustomActionModel(
    customActionId,
    actionName,
    actionTrigger,
    actionStatus
  );
  const response = await embeddedDashboard.addVisualActions(
    SheetId,
    VisualId,
    [myCustomAction]
  );
  if (!response.success) {
    console.log("Adding visual action failed", response.errorCode);
  }
} catch (error) {
  console.log("Adding visual action failed", error);
}
};

const parseDatapoint = (visualId, datapoint) => {
  datapoint.Columns.forEach((Column, index) => {
    // FIELD | METRIC
    const columnType = Object.keys(Column)[0];

    // STRING | DATE | INTEGER | DECIMAL
    const valueType = Object.keys(Column[columnType])[0];
    const { Column: columnMetadata } = Column[columnType][valueType];

    const value = datapoint.RawValues[index][valueType];
    const formattedValue = datapoint.FormattedValues[index];

    console.log(
      `Column: ${columnMetadata.ColumnName} has a raw value of ${value}
        and formatted value of ${formattedValue.Value} for visual: ${visualId}`
    );
  });
};

```

```

    );
  });
};

// This function is used to start a custom workflow after the end user selects a
// datapoint
const myCustomDatapointCallbackWorkflow = (callbackData) => {
  const { VisualId, Datapoints } = callbackData;

  parseDatapoint(VisualId, Datapoints);
};

// EMBEDDING THE DASHBOARD

const main = async () => {
  // 1. Get embed url
  let url;
  try {
    url = await myGetEmbedUrl(MY_DASHBOARD_ID);
  } catch (error) {
    console.log("Obtaining an embed url failed");
  }

  if (!url) {
    return;
  }

  // 2. Create embedding context
  const embeddingContext = await createEmbeddingContext();

  // 3. Embed the dashboard
  const embeddedDashboard = await embeddingContext.embedDashboard(
    {
      url,
      container: MY_DASHBOARD_CONTAINER,
      width: "1200px",
      height: "300px",
      resizeHeightOnSizeChangedEvent: true,
    },
    {
      onMessage: async (messageEvent) => {
        const { eventName, message } = messageEvent;
        switch (eventName) {
          case "CONTENT_LOADED": {

```

```

        await myAddVisualActionOnFirstVisualOfFirstSheet(embeddedDashboard);
        break;
    }
    case "CALLBACK_OPERATION_INVOKED": {
        myCustomDatapointCallbackWorkflow(message);
        break;
    }
    },
    },
    );
};

main().catch(console.error);

```

You can also configure the preceding example to initiate datapoint callback when the user opens the context menu. To do this with the preceding example, set the value of `actionTrigger` to `ActionTrigger.DATA_POINT_MENU`.

After a datapoint callback is registered, it's applied to most datapoints on the specified visual or visuals. Callbacks don't apply to totals or subtotals on visuals. When a reader interacts with a datapoint, a `CALLBACK_OPERATION_INVOKED` message is emitted to the Amazon Quick Sight embedding SDK. This message is captured by the `onMessage` handler. The message contains the raw and display values for the full row of data associated with the datapoint that is selected. It also contains the column metadata for all columns in the visual that the datapoint is contained in. The following is an example of a `CALLBACK_OPERATION_INVOKED` message.

```

{
  CustomActionId: "custom_action_id",
  DashboardId: "dashboard_id",
  SheetId: "sheet_id",
  VisualId: "visual_id",
  DataPoints: [
    {
      RawValues: [
        {
          String: "Texas" // 1st raw value in row
        },
        {
          Integer: 1000 // 2nd raw value in row
        }
      ],
    },
  ],
}

```

```

    FormattedValues: [
      {Value: "Texas"}, // 1st formatted value in row
      {Value: "1,000"} // 2nd formatted value in row
    ],
    Columns: [
      { // 1st column metadata
        Dimension: {
          String: {
            Column: {
              ColumnName: "State",
              DatsetIdentifier: "...",
            }
          }
        }
      },
      { // 2nd column metadata
        Measure: {
          Integer: {
            Column: {
              ColumnName: "Cancelled",
              DatsetIdentifier: "...",
            },
            AggregationFunction: {
              SimpleNumericalAggregation: "SUM"
            }
          }
        }
      }
    ]
  }
}

```

Filtering data at runtime for Amazon Quick Sight embedded dashboards and visuals

You can use filter methods in the Amazon Quick Sight embedding SDK to leverage the power of Amazon Quick Sight filters within your software as a service (SaaS) application at runtime. Runtime filters allow business owners to integrate their application with their embedded Amazon Quick Sight dashboards and visuals. To accomplish this, create customized filter controls in your application and apply filter presets based on data from your application. Then, developers can personalize filter configurations for end users at runtime.

Developers can create, query, update, and remove Amazon Quick Sight filters on an embedded dashboard or visual from their application with the Amazon Quick Sight Embedding SDK. Create Amazon Quick Sight filter objects in your application with the [FilterGroup](#) data model and apply them to embedded dashboards and visuals using the filter methods. For more information about using the Amazon Quick Sight Embedding SDK, see the [amazon-quicksight-embedding-sdk](#) on GitHub.

Prerequisites

Before you can get started, make sure that you are using the Amazon Quick Sight Embedding SDK version 2.5.0 or higher.

Terminology and concepts

The following terminology can be useful when working with embedded runtime filtering.

- *Filter group* – A group of individual filters. Filters that are located within a `FilterGroup` are OR-ed with each other. Filters within a [FilterGroup](#) are applied to the same sheets or visuals.
- *Filter* – A single filter. The filter can be a category, numeric, or datetime filter type. For more information on filters, see [Filter](#).

Setting up

Before you begin, make sure that you have the following assets and information prepared.

- The sheet ID of the sheet that you want to scope the `FilterGroup` to. This can be obtained with the `getSheets` method in the Embedding SDK.
- The dataset and column identifier of the dataset that you want to filter. This can be obtained through the [DescribeDashboardDefinition](#) API operation.

Depending on the column type that you use, there might be restrictions on the types of filters that can be added to an embedded asset. For more information on filter restrictions, see [Filter](#).

- The visual ID of the visual that you want to scope the `FilterGroup` to, if applicable. This can be obtained by using the `getSheetVisuals` method in the Embedding SDK.

In addition to the `getSheetVisuals` method, the `FilterGroup` that you add can only be scoped to the currently selected sheet.

To use this feature, you must already have a dashboard or visual embedded into your application through the Amazon Quick Sight Embedding SDK. For more information about using the Amazon Quick Sight Embedding SDK, see the [amazon-quicksight-embedding-sdk](#) on GitHub.

SDK method interface

Dashboard embedding getter methods

The following table describes different dashboard embedding getter methods that developers can use.

Method	Description
<code>getFilterGroupsForSheet(sheetId: string)</code>	Returns all <code>FilterGroups</code> that are currently scoped to the sheet that is supplied in the parameter.
<code>getFilterGroupsForVisual(sheetId: string, visualId: string)</code>	Returns all <code>FilterGroups</code> that are scoped to the visual that is supplied in the parameter.

If the sheet that is supplied in the parameter is not the currently selected sheet of the embedded dashboard, the above methods return an error.

Visual embedding getter methods

The following table describes different visual embedding getter methods that developers can use.

Method	Description
<code>getFilterGroups()</code>	Returns all <code>FilterGroups</code> that are currently scoped to the embedded visual.

Setter methods

The following table describes different setter methods that developers can use for dashboard or visual embedding.

Method	Description
<code>addFilterGroups(filterGroups: FilterGroup[])</code>	Adds and applies the supplied FilterGroups to the embedded dashboard or visual. A <code>ResponseMessage</code> that indicates whether the addition was successful is returned.
<code>updateFilterGroups(filterGroups: FilterGroup[])</code>	Updates the <code>FilterGroups</code> on the embedded experience that contains the same <code>FilterGroupId</code> as the <code>FilterGroup</code> that is supplied in the parameter. A <code>ResponseMessage</code> that indicates whether the update was successful is returned.
<code>removeFilterGroups(filterGroupsOrIds: FilterGroup[] string[])</code>	Removes the supplied <code>FilterGroups</code> from the dashboard and returns a <code>ResponseMessage</code> that indicates whether the removal attempt is successful.

The `FilterGroup` that is supplied must be scoped to the embedded sheet or visual that is currently selected.

Customize the look and feel of Amazon Quick Sight embedded dashboards and visuals

You can use the Amazon Quick Sight Embedding SDK (version 2.5.0 and higher) to make changes to the theming of your embedded Amazon Quick Sight dashboards and visuals at runtime. Runtime theming makes it easier to integrate your Software as a service (SaaS) application with your Amazon Quick Sight embedded assets. Runtime theming allows you to synchronize the theme of your embedded content with the themes of the parent application that your Amazon Quick Sight assets are embedded into. You can also use runtime theming to add customization options for readers. Theming changes can be applied to embedded assets at initialization or throughout the lifetime of your embedded dashboard or visual.

For more information about themes, see [Using themes in Amazon Quick Sight](#). For more information about using the Amazon Quick Sight Embedding SDK, see the [amazon-quicksight-embedding-sdk](#) on GitHub.

Prerequisites

Before you can get started, make sure that you have the following prerequisites.

- You are using the Amazon Quick Sight Embedding SDK version 2.5.0 or higher.
- Permissions to access the theme that you want to work with. To grant permissions to a theme in Amazon Quick Sight, make a `UpdateThemePermissions` API call or use the **Share** icon next to the theme in the Amazon Quick Sight console's analysis editor.

Terminology and concepts

The following terminology can be useful when working with embedded runtime theming.

- *Theme* – A collection of settings that you can apply to multiple analyses and dashboards that change how the content is displayed.
- *ThemeConfiguration* – A configuration object that contains all of the display properties for a theme.
- *Theme Override* – A *ThemeConfiguration* object that is applied to the active theme to override some or all aspects of how content is displayed.
- *Theme ARN* – An Amazon Resource Name (ARN) that identifies a Amazon Quick Sight theme. Following is an example of custom theme ARN.

```
arn:aws-cn:quicksight:region:account-id:theme/theme-id
```

Amazon Quick Sight provided starter themes don't have a region in their theme ARN. Following is an example of a starter theme ARN.

```
arn:aws-cn:quicksight::aws:theme/CLASSIC
```

Setting up

Make sure that you have the following information ready to get started working with runtime theming.

- The theme ARNs of the themes that you want to use. You can choose an existing theme, or you can create a new one. To obtain a list all themes and theme ARNs in your Amazon Quick Sight account, make a call to the [ListThemes](#) API operation. For information on preset Amazon Quick Sight themes, see [Setting a default theme for Amazon Quick Suite analyses with the Amazon Quick Suite APIs](#).

- If you are using registered user embedding, make sure that the user has access to the themes that you want to use.

If you are using anonymous user embedding, pass a list of theme ARNs to the `AuthorizedResourceArns` parameter of the `GenerateEmbedUrlForAnonymousUser` API. Anonymous users are granted access to any theme that is listed in the `AuthorizedResourceArns` parameter.

SDK method interface

Setter methods

The following table describes different setter methods that developers can use for runtime theming.

Method	Description
<code>setTheme(themeArn: string)</code>	Replaces the active theme of a dashboard or visual with another theme. If applied, the theme override is removed. An error is returned if you don't have access to the theme or if the theme doesn't exist.
<code>setThemeOverride(themeOverride: ThemeConfiguration)</code>	Sets a dynamic <code>ThemeConfiguration</code> to override the current active theme. This replaces the previously set theme override. Any values that are not supplied in the new <code>ThemeConfiguration</code> are defaulted to the values in the currently active theme. An error is returned if the <code>ThemeConfiguration</code> that you supply is invalid.

Initializing embedded content with a theme

To initialize an embedded dashboard or visual with a non-default theme, define a `themeOptions` object on the `DashboardContentOptions` or `VisualContentOptions` parameters, and set the `themeArn` property within `themeOptions` to the desired theme ARN.

The following example initializes an embedded dashboard with the MIDNIGHT theme.

```
import { createEmbeddingContext } from 'amazon-quicksight-embedding-sdk';

const embeddingContext = await createEmbeddingContext();

const {
  embedDashboard,
} = embeddingContext;

const frameOptions = {
  url: '<YOUR_EMBED_URL>',
  container: '#experience-container',
};

const contentOptions = {
  themeOptions: {
    themeArn: "arn:aws-cn:quicksight::aws:theme/MIDNIGHT"
  }
};

// Embedding a dashboard experience
const embeddedDashboardExperience = await embedDashboard(frameOptions, contentOptions);
```

Initializing embedded content with a theme override

Developers can use theme overrides to define the theme of an embedded dashboard or visual at runtime. This allows the dashboard or visual to inherit a theme from a third-party application without the need to pre-configure a theme within Amazon Quick Sight. To initialize an embedded dashboard or visual with a theme override, set the `themeOverride` property within `themeOptions` in either the `DashboardContentOptions` or `VisualContentOptions` parameters. The following example overrides the font of a dashboard's theme from the default font to Amazon Ember.

```
import { createEmbeddingContext } from 'amazon-quicksight-embedding-sdk';

const embeddingContext = await createEmbeddingContext();
```

```
const {
  embedDashboard,
} = embeddingContext;

const frameOptions = {
  url: '<YOUR_EMBED_URL>',
  container: '#experience-container',
};
const contentOptions = {
  themeOptions: {
    "themeOverride":{"Typography":{"FontFamilies":[{"FontFamily":"Comic Neue"]}}}
  }
};

// Embedding a dashboard experience
const embeddedDashboardExperience = await embedDashboard(frameOptions, contentOptions);
```

Initializing embedded content with preloaded themes

Developers can configure a set of dashboard themes to be preloaded on initialization. This is most beneficial for quick toggling between different views, for example dark and light modes. An embedded dashboard or visual can be initialized with up to 5 preloaded themes. To use preloaded themes, set the `preloadThemes` property in either `DashboardContentOptions` or `VisualContentOptions` with an array of up to 5 theme ARNs. The following example preloads the Midnight and Rainier starter themes to a dashboard.

```
import { createEmbeddingContext } from 'amazon-quicksight-embedding-sdk';

const embeddingContext = await createEmbeddingContext();

const {
  embedDashboard,
} = embeddingContext;

const frameOptions = {
  url: '<YOUR_EMBED_URL>',
  container: '#experience-container',
};
const contentOptions = {
  themeOptions: {
    "preloadThemes": ["arn:aws:quicksight::aws:theme/RAINIER",
    "arn:aws:quicksight::aws:theme/MIDNIGHT"]
```

```
}  
};  
  
// Embedding a dashboard experience  
const embeddedDashboardExperience = await embedDashboard(frameOptions, contentOptions);
```

Using the Amazon Quick Sight Embedding SDK to enable shareable links to embedded dashboard views

Amazon Quick Sight developers can use the Amazon Quick Sight Embedding SDK (version 2.8.0 and higher) to allow readers of embedded dashboards to receive and distribute shareable links to their view of an embedded dashboard. Developers can use dashboard or console embedding to generate a shareable link to their application page with Amazon Quick Sight's reference encapsulated using the Amazon Quick Sight Embedding SDK. Amazon Quick Sight Readers can then send this shareable link to their peers. When their peer accesses the shared link, they are taken to the page on the application that contains the embedded Amazon Quick Sight dashboard. Developers can also generate and save shareable links of dashboard views that can be used as a bookmarks for anonymous readers of Amazon Quick Sight when using anonymous embedding.

Prerequisites

Before you get started, make sure that you are using the Amazon Quick Sight Embedding SDK version 2.8.0 or higher

Topics

- [Enabling the SharedView feature configuration for Amazon Quick Sight embedded analytics](#)
- [Creating a shared view with the Amazon Quick Sight createSharedView API](#)
- [Consuming a shared Amazon Quick Sight view](#)

Enabling the SharedView feature configuration for Amazon Quick Sight embedded analytics

When you create an the embedded instance with the Amazon Quick Sight API, set the value of SharedView in the FeatureConfigurations payload to true, as shown in the example below. SharedView overrides the StatePersistence configurations for registered users who access embedded dashboards. If a dashboard user has StatePersistence disabled and SharedView enabled, their state will persist.

```
const generateNewEmbedUrl = async () => {  
  const generateUrlPayload = {
```

```

    experienceConfiguration: {
      QuickSightConsole: {
        FeatureConfigurations: {
          "SharedView": {
            "Enabled": true
          },
        },
      },
    },
  },
}
const result: GenerateEmbedUrlResult = await
generateEmbedUrlForRegisteredUser(generateUrlPayload);
return result.url;
};

```

Creating a shared view with the Amazon Quick Sight createSharedView API

After you update the Embedding SDK to version 2.8.0 or higher, use the `createSharedView` API to create a new shared view. Record the `sharedViewId` and the `dashboardId` that the operation returns. The example below creates a new shared view.

```

const response = await embeddingFrame.createSharedView();
const sharedViewId = response.message.sharedViewId;
const dashboardId = response.message.dashboardId;

```

`createSharedView` can only be called when a user views a dashboard. For console-specific shared view creation, make sure that users are on the dashboard page before you enable the `createSharedView` action. You can do this with the `PAGE_NAVIGATION` event, shown in the example below.

```

const contentOptions = {
  onMessage: async (messageEvent, metadata) => {
    switch (messageEvent.eventName) {
      case 'CONTENT_LOADED': {
        console.log("Do something when the embedded experience is fully
loaded.");
        break;
      }
      case 'ERROR_OCCURRED': {
        console.log("Do something when the embedded experience fails
loading.");
        break;
      }
    }
  }
}

```

```

        case 'PAGE_NAVIGATION': {
            setPageType(messageEvent.message.pageType);
            if (messageEvent.message.pageType === 'DASHBOARD') {
                setShareEnabled(true);
            } else {
                setShareEnabled(false);
            }
            break;
        }
    }
}
};

```

Consuming a shared Amazon Quick Sight view

After you create a new shared view, use the Embedding SDK to make the shared view consumable for other users. The examples below set up a consumable shared view for an embedded dashboard in Amazon Quick Sight.

With an appended URL

Append the `sharedViewId` to the embed URL, under `/views/{viewId}`, and expose this URL to your users. Users can use this URL to will navigate to that shared view.

```

const response = await dashboardFrame.createSharedView();
const newEmbedUrl = await generateNewEmbedUrl();
const formattedUrl = new URL(newEmbedUrl);
formattedUrl.pathname = formattedUrl.pathname.concat('/views/' +
    response.message.sharedViewId);
const baseUrl = formattedUrl.href;
alert("Click to view this QuickSight shared view", baseUrl);

```

With the contentOptions SDK

Pass a `viewId` to the `contentOptions` to open the experience with the given `viewId`.

```

const contentOptions = {
    toolbarOptions: {
        ...
    },
    viewId: sharedViewId,
};

```

```
const embeddedDashboard = await embeddingContext.embedDashboard(
  {container: containerRef.current},
  contentOptions
);
```

With the InitialPath property

```
const shareView = async() => {
  const returnValue = await consoleFrame.createSharedView();
  const {dashboardId, sharedViewId} = returnValue.message;
  const newEmbedUrl = await generateNewEmbedUrl(`/dashboards/${dashboardId}/views/${sharedViewId}`);
  setShareUrl(newEmbedUrl);
};

const generateNewEmbedUrl = async (initialPath) => {
  const generateUrlPayload = {
    experienceConfiguration: {
      QuickSightConsole: {
        InitialPath: initialPath,
        FeatureConfigurations: {
          "SharedView": {
            "Enabled": true
          },
        },
      },
    },
  }
  const result: GenerateEmbedUrlResult = await
  generateEmbedUrlForRegisteredUser(generateUrlPayload);
  return result.url;
};
```

Embedding Amazon Quick Sight visuals and dashboards with a 1-click embed code

You can embed a visual or dashboard in your application using an embed code. You get this code when you share the dashboard or from the **Embed visual** menu in Amazon Quick Sight.

You can embed a visual or dashboard in your internal application for your registered users. Or you can turn on public sharing in the Amazon Quick Sight console. Doing this grants anyone on the

internet access to a shared visual or dashboard that is embedded in a public application, wiki, or portal.

Following, you can find descriptions about how to embed visuals and dashboards using the 1-click visual or dashboard embed code.

Topics

- [Embedding Amazon Quick Sight visuals and dashboards for registered users with a 1-click embed code](#)
- [Embedding Amazon Quick Sight visuals and dashboards for anonymous users with a 1-click embed code](#)

Embedding Amazon Quick Sight visuals and dashboards for registered users with a 1-click embed code

Applies to: Enterprise Edition

You can embed a visual or dashboard in your internal application for registered users of your Amazon Quick Sight account. You do so using the embed code that you get when you share the dashboard or from the **Embed visual** menu in Amazon Quick Sight. You don't have to run the Amazon Quick Sight embedding API to generate the embed code. You can copy the embed code from Amazon Quick Sight and paste it in your internal application's HTML code.

When users and groups (or all users on your Amazon Quick Sight account) who have access to the dashboard that you want to embed or that holds the visual that you want to embed access your internal application, they're prompted to sign in to the Amazon Quick Sight account with their credentials. After they are authenticated, they can access the visual or dashboard on their internal page. If you have single sign-on enabled, users aren't prompted to sign in again.

Following, you can find descriptions about how to embed a visual or dashboard for registered users using the visual or dashboard embed code.

Before you start

Before you get started, make sure of the following:

- Your internet browser settings contain one of the following to allow communication between the popup and the iframe:

- Native support for the Mozilla Broadcast Channel API. For more information, see [Broadcast Channel API](#) in the Mozilla documentation.
- IndexedDB support.
- LocalStorage support.
- Your internet browser's "block all cookies" settings is turned off.

Step 1: Grant access to the dashboard

For users to access your embedded dashboard, grant them access to view it. You can grant individual users and groups access to a dashboard, or you can grant everyone in your account access. Visual permissions are determined at the dashboard level. To grant access to embedded visuals, grant access to the dashboard that the visual belongs to. For more information, see [Granting access to a dashboard](#).

Step 2: Put the domain where you want to embed the visual or dashboard on your allow list

To embed visuals and dashboards in your internal application, make sure that the domain where you're embedding is allow-listed in your Amazon Quick Sight account. For more information, see [Allow listing static domains](#).

Step 3: Get the embed code

Use the following procedure to get the visual or dashboard embed code.

To get the dashboard embed code

1. Open the published dashboard in Amazon Quick Sight and choose **Share** at upper right. Then choose **Share dashboard**.
2. In the **Share dashboard** page that opens, choose **Copy embed code** at upper left.

The embed code is copied to your clipboard and is similar to the following. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
<iframe
  width="960"
  height="720"
  src="https://quicksightdomain/sn/embed/share/accounts/accountid/
dashboards/dashboardid?directory_alias=account_directory_alias">
```

```
</iframe>
```

To get the visual embed code

1. Open the published dashboard in Amazon Quick Sight and choose the visual that you want to embed. Then open the on-visual menu at the upper right of the visual and choose **Embed visual**.
2. In the **Embed visual** pane that opens, choose **Copy code**.

The embed code is copied to your clipboard and is similar to the following. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
<iframe
  width="600"
  height="400"
  src="https://quicksightdomain/sn/embed/share/accounts/111122223333/
dashboards/DASHBOARDID/sheets/SHEETID>/visuals/VISUALID">
</iframe>
```

Step 4: Paste the code into your internal application's HTML page

Use the following procedure to paste the embed code into your internal application's HTML page

To paste the code in your internal application's HTML page

- Open the HTML code for any page where you want to embed the dashboard and paste the embed code in.

The following example shows what this might look like for an embedded dashboard. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
<!DOCTYPE html>
<html>
<body>

<h2>Example.com - Employee Portal</h2>
<h3>Current shipment stats</h3>
```

```
<iframe
  width="960"
  height="720"
  src="https://quicksightdomain/sn/embed/share/accounts/accountid/
dashboards/dashboardid?directory_alias=account_directory_alias">
</iframe>

</body>
</html>
```

The following example shows what this might look like for an embedded visual. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
<!DOCTYPE html>
<html>
<body>

<h2>Example.com - Employee Portal</h2>
<h3>Current shipment stats</h3>
  <iframe
    width="600"
    height="400"
    src="https://quicksightdomain/sn/embed/share/accounts/111122223333/
dashboards/DASHBOARDID/sheets/SHEETID>visuals/VISUALID?
directory_alias=account_directory_alias">
  </iframe>

</body>
</html>
```

For example, let's say that you want to embed your visual or dashboard in an internal Google Sites page. You can open the page on Google Sites and paste the embed code in an embed widget.

If you want to embed your visual or dashboard in an internal Microsoft SharePoint site, you can create a new page and then paste the embed code in an Embed web part.

Embedding Amazon Quick Sight visuals and dashboards for anonymous users with a 1-click embed code

Applies to: Enterprise Edition

You can embed a visual or dashboard in public sites using the embed code that you get when you share the visual or dashboard in Amazon Quick Sight. You can also turn on public sharing by using the Amazon Quick Sight console and automatically grant access to a shared visual or dashboard to anyone on the internet.

Following, you can find how to turn on public sharing for a visual or dashboard and embed the visual or dashboard for anyone on the internet to see. In both cases, you do this by using the 1-click embed code.

Before you start

Before you get started, make sure of the following:

- Your internet browser settings contain one of the following to allow communication between the popup and the iframe that sharing uses:
 - Native support for the Mozilla Broadcast Channel API. For more information, see [Broadcast Channel API](#) in the Mozilla documentation.
 - IndexedDB support.
 - LocalStorage support.
- Your internet browser's "block all cookies" settings is turned off.

Step 1: Turn on public access for the dashboard

For anyone on the internet to access your embedded visual or dashboard, first turn on public access for the dashboard. Visual permissions are determined at the dashboard level. To grant access to embedded visuals, grant access to the dashboard that the visual belongs to. For more information, see [Granting anyone on the internet access to an Amazon Quick Sight dashboard](#).

Step 2: Put the domain where you want to embed the visual or dashboard on your allow list

To embed visuals and dashboards in a public application, wiki, or portal, make sure that the domain where you're embedding it is on the allow list for your Amazon Quick Sight account.

Step 3: Get the embed code

Use the following procedure to get the visual or dashboard embed code.

To get the dashboard embed code

1. Open the published dashboard in Amazon Quick Sight and choose **Share** at upper right. Then choose **Share dashboard**.
2. In the **Share dashboard** page that opens, choose **Copy embed code** at upper left.

The embed code is copied to your clipboard and is similar to the following. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
<iframe
  width="960"
  height="720"
  src="https://quicksightdomain/sn/
    embed/share/accounts/accountid/dashboards/dashboardid">
</iframe>
```

To get the visual embed code

1. Open the published dashboard in Amazon Quick Sight and choose the visual that you want to embed. Then open the on-visual menu in the top right corner of the visual and choose **Embed visual**.
2. In the **Embed visual** pane that opens, choose **Copy code**.

The embed code is copied to your clipboard and is similar to the following. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
<iframe
  width="600"
  height="400"
  src="https://quicksightdomain/sn/embed/share/accounts/111122223333/
dashboards/DASHBOARDID/sheets/SHEETID>visuals/VISUALID">
</iframe>
```

Step 4: Paste the embed code into an HTML page, wiki page, or portal

Use the following procedure to paste the embed code into an HTML page, wiki page, or portal.

To paste the embed code

- Open the HTML code for the location where you want to embed the visual or dashboard, and paste the embed code in.

The following example shows what this might look like for an embedded dashboard. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
<!DOCTYPE html>
<html>
<body>

<h2>Example.com - Employee Portal</h2>
<h3>Current shipment stats</h3>
  <iframe
    width="960"
    height="720"
    src="https://quicksightdomain/sn/
        embed/share/accounts/accountid/dashboards/dashboardid">
  </iframe>

</body>
</html>
```

The following example shows what this might look like for an embedded visual. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
<!DOCTYPE html>
<html>
<body>

<h2>Example.com - Employee Portal</h2>
<h3>Current shipment stats</h3>
  <iframe
    width="600"
    height="400"
```

```
src="https://quicksightdomain/sn/embed/share/accounts/111122223333/
dashboards/DASHBOARDID/sheets/SHEETID>/visuals/VISUALID">
</iframe>

</body>
</html>
```

If your public-facing applications are built on Google Sites, open the page on Google Sites and then paste the embed code using the embed widget.

Make sure that the following domains in Amazon Quick Sight are on your allow list when you embed in Google Sites:

- <https://googleusercontent.com> (turns on subdomains)
- <https://www.gstatic.com>
- <https://sites.google.com>

After you embed the visual or dashboard in your application, anyone who can access your application can access the embedded visual or dashboard. To update a dashboard that's shared with the public, see [Updating a publicly shared dashboard](#). To turn off public sharing, see [Turning off public sharing settings](#).

When you turn off public sharing, no one from the internet can access a dashboard or dashboards that you have embedded on a public application or shared with a link. The next time anyone tries to view such a dashboard from the internet, they receive a message that they don't have access to view it.

Embedding with the Amazon Quick Sight APIs

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

There are only a few steps involved in the actual process of embedding analytics using the Amazon Quick Sight APIs.

Before you begin, make sure to have the following items in place:

- Set up the required IAM permissions for the caller identity used by your application that will use the Amazon SDK to make API calls. For example, grant permission to allow the `quicksight:GenerateEmbedUrlForAnonymousUser` or `quicksight:GenerateEmbedUrlForRegisteredUser` action.
- To embed for registered users, share Amazon Quick Sight assets with them beforehand. For new authenticating users, know how to grant access to the assets. One way to do this is by adding all the assets to a Amazon Quick Sight folder. If you prefer to use the Amazon Quick Sight API, use the `DescribeDashboardPermissions` and `UpdateDashboardPermissions` API operations. For more information, see [DescribeDashboardPermissions](#) or [UpdateDashboardPermissions](#) in the *Amazon Quick Suite API Reference*. If you want to share the dashboard with all users in a namespace or group, you can share the dashboard with namespace or group.
- If you're embedding dashboards, make sure to have the ID of the dashboards you want to embed. The dashboard ID is the code in the URL of the dashboard. You can also get it from the dashboard URL.
- A Amazon Quick Sight administrator must explicitly enable domains where you plan to embed your Amazon Quick Sight analytics. You can do this by using the **Manage Amazon Quick Sight, Domains and Embedding** from the profile menu, or you can use the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` or `GenerateEmbedUrlForRegisteredUser` API call.

This option is only visible to Amazon Quick Sight administrators. You can also add subdomains as part of a domain. For more information, see [Allow listing domains at runtime with the Amazon Quick Suite API](#).

All domains in your static allow list (such as development, staging, and production) must be explicitly allowed, and they must use HTTPS. You can add up to 100 domains to the allow list. You can add domains at runtime with Amazon Quick Sight API operations.

After all the prerequisites are complete, embedding Amazon Quick Sight involves the following steps, which are explained in greater detail later:

1. For authentication, use your application server to authenticate the user. After authentication in your server, generate the embedded dashboard URL using the Amazon SDK that you need.
2. In your web portal or application, embed Amazon Quick Sight using the generated URL. To simplify this process, you can use the Amazon Quick Sight Embedding SDK, available on [NPMJS](#)

and [GitHub](#). This customized JavaScript SDK is designed to help you efficiently integrate Amazon Quick Sight into your application pages, set defaults, connect controls, get callbacks, and handle errors.

You can use Amazon CloudTrail auditing logs to get information about the number of embedded dashboards, users of an embedded experience, and access rates.

Topics

- [Embedding Amazon Quick Sight dashboards with the Amazon Quick Sight API](#)
- [Embedding Amazon Quick Sight visuals with the Amazon Quick Sight APIs](#)
- [Embedding the full functionality of the Amazon Quick Sight console for registered users](#)
- [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#)
- [Embedding the Amazon Quick Sight Q search bar \(Classic\)](#)
- [Embedding analytics using the GetDashboardEmbedURL and GetSessionEmbedURL API operations](#)

Embedding Amazon Quick Sight dashboards with the Amazon Quick Sight API

Use the following topics to learn about embedding dashboards with the Amazon Quick Sight API.

Topics

- [Embedding Amazon Quick Sight dashboards for registered users](#)
- [Embedding Amazon Quick Sight dashboards for anonymous \(unregistered\) users](#)
- [Enabling executive summaries in embedded dashboards](#)

Embedding Amazon Quick Sight dashboards for registered users

Important

Amazon Quick Sight has new API operations for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` API operations to embed dashboards and the Amazon Quick Sight console, but they don't contain the latest embedding capabilities. For more information about embedding using the

old API operations, see [Embedding analytics using the GetDashboardEmbedURL and GetSessionEmbedURL API operations](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

In the following sections, you can find detailed information about how to set up embedded Amazon Quick Sight dashboards for registered users of Amazon Quick Sight.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the dashboard URL](#)

Step 1: Set up permissions

In the following section, you can find out how to set up permissions for the backend application or web server. This task requires administrative access to IAM.

Each user who accesses a dashboard assumes a role that gives them Amazon Quick Sight access and permissions to the dashboard. To make this possible, create an IAM role in your Amazon Web Services account. Associate an IAM policy with the role to provide permissions to any user who assumes it. The IAM role needs to provide permissions to retrieve embedding URLs for a specific user pool. With the help of the wildcard character *, you can grant the permissions to generate a URL for all users in a specific namespace, or for a subset of users in specific namespaces. For this, you add `quicksight:GenerateEmbedUrlForRegisteredUser`.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForRegisteredUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants you as a developer the option to override the static domains that are configured in the **Manage Amazon Quick Sight** menu. Instead, you can list up to three domains or subdomains that can access the generated URL. This URL is then embedded in the website that you create. Only the domains that are listed in the parameter

can access the embedded visual. Without this condition, you can list any domain on the internet in the `AllowedDomains` parameter.

To limit the domains that developers can use with this parameter, add an `AllowedEmbeddingDomains` condition to your IAM policy. For more information about the `AllowedDomains` parameter, see [GenerateEmbedUrlForRegisteredUser](#) in the *Amazon Quick Sight API Reference*.

The following sample policy provides these permissions.

Additionally, if you are creating first-time users who will be Amazon Quick Sight readers, make sure to add the `quicksight:RegisterUser` permission in the policy.

The following sample policy provides permission to retrieve an embedding URL for first-time users who are to be Amazon Quick Sight readers.

Finally, your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf and provision the user in Amazon Quick Sight. The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

```
]
}
```

For more information regarding trust policies for OpenID Connect or SAML authentication, see the following sections of the *IAM User Guide*:

- [Creating a Role for Web Identity or OpenID Connect Federation \(Console\)](#)
- [Creating a Role for SAML 2.0 Federation \(Console\)](#)

Step 2: Generate the URL with the authentication code attached

In the following section, you can find out how to authenticate your user and get the embeddable dashboard URL on your application server. If you plan to embed dashboards for IAM or Amazon Quick Sight identity types, share the dashboard with the users.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then it adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

Performing these steps ensures that each viewer of the dashboard is uniquely provisioned in Amazon Quick Sight. It also enforces per-user settings, such as the row-level security and dynamic defaults for parameters.

The following examples perform the IAM authentication on the user's behalf. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import
com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserRequest;
import
com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserResult;
import
com.amazonaws.services.quicksight.model.RegisteredUserEmbeddingExperienceConfiguration;
```

```

import
com.amazonaws.services.quicksight.model.RegisteredUserDashboardEmbeddingConfiguration;

/**
 * Class to call QuickSight Amazon SDK to get url for dashboard embedding.
 */
public class GetQuicksightEmbedUrlRegisteredUserDashboardEmbedding {

    private final AmazonQuickSight quickSightClient;

    public GetQuicksightEmbedUrlRegisteredUserDashboardEmbedding() {
        this.quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AWSCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret key here
                    return new BasicAWSCredentials("access-key", "secret-key");
                }

                @Override
                public void refresh() {}
            })
            .build();
    }

    public String getQuicksightEmbedUrl(
        final String accountId, // Amazon Account ID
        final String dashboardId, // Dashboard ID to embed
        final List<String> allowedDomains, // Runtime allowed domain for
embedding
        final String userArn // Registered user arn to use for embedding. Refer
to Get Embed Url section in developer portal to find out how to get user arn for a
QuickSight user.
    ) throws Exception {
        final RegisteredUserEmbeddingExperienceConfiguration
experienceConfiguration = new RegisteredUserEmbeddingExperienceConfiguration()
            .withDashboard(new
RegisteredUserDashboardEmbeddingConfiguration().withInitialDashboardId(dashboardId));
        final GenerateEmbedUrlForRegisteredUserRequest
generateEmbedUrlForRegisteredUserRequest = new
GenerateEmbedUrlForRegisteredUserRequest();

```

```

        generateEmbedUrlForRegisteredUserRequest.setAwsAccountId(accountId);
        generateEmbedUrlForRegisteredUserRequest.setUserArn(userArn);
        generateEmbedUrlForRegisteredUserRequest.setAllowedDomains(allowedDomains);

generateEmbedUrlForRegisteredUserRequest.setExperienceConfiguration(experienceConfiguration);

        final GenerateEmbedUrlForRegisteredUserResult
generateEmbedUrlForRegisteredUserResult =
quickSightClient.generateEmbedUrlForRegisteredUser(generateEmbedUrlForRegisteredUserRequest);

        return generateEmbedUrlForRegisteredUserResult.getEmbedUrl();
    }
}

```

JavaScript

```

global.fetch = require('node-fetch');
const Amazon = require('aws-sdk');

function generateEmbedUrlForRegisteredUser(
    accountId,
    dashboardId,
    openIdToken, // Cognito-based token
    userArn, // registered user arn
    roleArn, // IAM user role to use for embedding
    sessionName, // Session name for the roleArn assume role
    allowedDomains, // Runtime allowed domain for embedding
    getEmbedUrlCallback, // GetEmbedUrl success callback method
    errorCallback // GetEmbedUrl error callback method
) {
    const stsClient = new AWS.STS();
    let stsParams = {
        RoleSessionName: sessionName,
        WebIdentityToken: openIdToken,
        RoleArn: roleArn
    }

    stsClient.assumeRoleWithWebIdentity(stsParams, function(err, data) {
        if (err) {
            console.log('Error assuming role');
            console.log(err, err.stack);
            errorCallback(err);
        } else {

```

```

const getDashboardParams = {
  "AwsAccountId": accountId,
  "ExperienceConfiguration": {
    "Dashboard": {
      "InitialDashboardId": dashboardId
    }
  },
  "UserArn": userArn,
  "AllowedDomains": allowedDomains,
  "SessionLifetimeInMinutes": 600
};

const quicksightClient = new AWS.QuickSight({
  region: process.env.AWS_REGION,
  credentials: {
    accessKeyId: data.Credentials.AccessKeyId,
    secretAccessKey: data.Credentials.SecretAccessKey,
    sessionToken: data.Credentials.SessionToken,
    expiration: data.Credentials.Expiration
  }
});

quicksightClient.generateEmbedUrlForRegisteredUser(getDashboardParams,
function(err, data) {
  if (err) {
    console.log(err, err.stack);
    errorCallback(err);
  } else {
    const result = {
      "statusCode": 200,
      "headers": {
        "Access-Control-Allow-Origin": "*", // Use your website
domain to secure access to GetEmbedUrl API
        "Access-Control-Allow-Headers": "Content-Type"
      },
      "body": JSON.stringify(data),
      "isBase64Encoded": false
    }
    getEmbedUrlCallback(result);
  }
});
}
});

```

```
}
```

Python3

```
import json
import boto3
from botocore.exceptions import ClientError

sts = boto3.client('sts')

# Function to generate embedded URL
# accountId: Amazon account ID
# dashboardId: Dashboard ID to embed
# userArn: arn of registered user
# allowedDomains: Runtime allowed domain for embedding
# roleArn: IAM user role to use for embedding
# sessionName: session name for the roleArn assume role
def getEmbeddingURL(accountId, dashboardId, userArn, allowedDomains, roleArn,
    sessionName):
    try:
        assumedRole = sts.assume_role(
            RoleArn = roleArn,
            RoleSessionName = sessionName,
        )
    except ClientError as e:
        return "Error assuming role: " + str(e)
    else:
        assumedRoleSession = boto3.Session(
            aws_access_key_id = assumedRole['Credentials']['AccessKeyId'],
            aws_secret_access_key = assumedRole['Credentials']['SecretAccessKey'],
            aws_session_token = assumedRole['Credentials']['SessionToken'],
        )
        try:
            quicksightClient = assumedRoleSession.client('quicksight', region_name='us-
west-2')
            response = quicksightClient.generate_embed_url_for_registered_user(
                AwsAccountId=accountId,
                ExperienceConfiguration = {
                    "Dashboard": {
                        "InitialDashboardId": dashboardId
                    }
                },
                UserArn = userArn,
```



```
        AllowedDomains = allowedDomains,
        SessionLifetimeInMinutes = 600
    )

    return {
        'statusCode': 200,
        'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-Headers": "Content-Type"},
        'body': json.dumps(response),
        'isBase64Encoded': bool('false')
    }
except ClientError as e:
    return "Error generating embedding url: " + str(e)
```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
const Amazon = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({
    apiConfig: require('./quicksight-2018-04-01.min.json'),
    region: 'us-east-1',
});

quicksightClient.generateEmbedUrlForRegisteredUser({
    'AwsAccountId': '111122223333',
    'ExperienceConfiguration': {
        'Dashboard': {
            'InitialDashboardId': '1c1fe111-e2d2-3b30-44ef-a0e111111cde'
        }
    },
    'UserArn': 'REGISTERED_USER_ARN',
    'AllowedDomains': allowedDomains,
    'SessionLifetimeInMinutes': 100
}, function(err, data) {
    console.log('Errors: ');
    console.log(err);
});
```

```
        console.log('Response: ');
        console.log(data);
    });
```

Example

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
    //readability and added ellipsis to indicate that it's incomplete.
    {
        Status: 200,
        EmbedUrl: 'https://quicksightdomain/embed/12345/dashboards/67890...'
        RequestId: '7bee030e-f191-45c4-97fe-d9faf0e03713'
    }
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;

namespace GenerateDashboardEmbedUrlForRegisteredUser
{
    class Program
    {
        static void Main(string[] args)
        {
            var quicksightClient = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
                SessionToken,
                Amazon.RegionEndpoint.USEast1);
            try
            {
                RegisteredUserDashboardEmbeddingConfiguration
registeredUserDashboardEmbeddingConfiguration
                    = new RegisteredUserDashboardEmbeddingConfiguration
```

```

        {
            InitialDashboardId = "dashboardId"
        };
        RegisteredUserEmbeddingExperienceConfiguration
registeredUserEmbeddingExperienceConfiguration
        = new RegisteredUserEmbeddingExperienceConfiguration
        {
            Dashboard = registeredUserDashboardEmbeddingConfiguration
        };

        Console.WriteLine(
            quicksightClient.GenerateEmbedUrlForRegisteredUserAsync(new
GenerateEmbedUrlForRegisteredUserRequest
            {
                AwsAccountId = "111122223333",
                ExperienceConfiguration =
registeredUserEmbeddingExperienceConfiguration,
                UserArn = "REGISTERED_USER_ARN",
                AllowedDomains = allowedDomains,
                SessionLifetimeInMinutes = 100
            }).Result.EmbedUrl
        );
    } catch (Exception ex) {
        Console.WriteLine(ex.Message);
    }
}
}
}
}
}

```

Amazon CLI

To assume the role, choose one of the following Amazon Security Token Service (Amazon STS) API operations:

- [AssumeRole](#) – Use this operation when you're using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you're using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you're using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForRegisteredUser`. If you are

taking a just-in-time approach to add users when they first open a dashboard, the role also needs permissions enabled for `quicksight:RegisterUser`.

```
aws sts assume-role \  
    --role-arn "arn:aws-cn:iam::111122223333:role/  
embedding_quicksight_dashboard_role" \  
    --role-session-name john.doe@example.com
```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. If you're using a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"  
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"  
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_dashboard_role/john.doe@example.com`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. *Throttling* is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations.

The role session ID also becomes the user name in Amazon Quick Sight. You can use this pattern to provision your users in Amazon Quick Sight ahead of time, or to provision them the first time they access the dashboard.

The following example shows the CLI command that you can use to provision a user. For more information about [RegisterUser](#), [DescribeUser](#), and other Amazon Quick Sight API operations, see the [Amazon Quick Sight API Reference](#).

```
aws quicksight register-user \  
    --aws-account-id 111122223333 \  
    --namespace default \  
    --identity-type IAM \  
    --iam-arn "arn:aws-cn:iam::111122223333:role/  
embedding_quicksight_dashboard_role" \  
    --user-role READER \  
    --user-name jhnd \  
    --session-name "john.doe@example.com" \  
    --email john.doe@example.com \  
    --region us-east-1 \  
    --custom-permissions-name TeamA1
```

If your user is authenticated through Microsoft AD, you don't need to use `RegisterUser` to set them up. Instead, they should be automatically subscribed the first time they access Amazon Quick Sight. For Microsoft AD users, you can use `DescribeUser` to get the user ARN.

The first time a user accesses Amazon Quick Sight, you can also add this user to the group that the dashboard is shared with. The following example shows the CLI command to add a user to a group.

```
aws quicksight create-group-membership \  
    --aws-account-id=111122223333 \  
    --namespace=default \  
    --group-name=financeusers \  
    --member-name="embedding_quicksight_dashboard_role/john.doe@example.com"
```

You now have a user of your app who is also a user of Amazon Quick Sight, and who has access to the dashboard.

Finally, to get a signed URL for the dashboard, call `generate-embed-url-for-registered-user` from the app server. This returns the embeddable dashboard URL. The following example shows how to generate the URL for an embedded dashboard using a server-side call for users authenticated through Amazon Managed Microsoft AD or single sign-on (IAM Identity Center).

```
aws quicksight generate-embed-url-for-registered-user \  
    --aws-account-id 111122223333 \  
    --session-lifetime-in-minutes 600 \
```

```
--user-arn arn:aws-cn:quicksight:us-east-1:111122223333:user/default/  
embedding_quicksight_visual_role/embeddingsession \  
--allowed-domains '["domain1", "domain2"]' \  
--experience-configuration  
Dashboard={InitialDashboardId=1a1ac2b2-3fc3-4b44-5e5d-c6db6778df89}
```

For more information about using this operation, see [GenerateEmbedUrlForRegisteredUser](#). You can use this and other API operations in your own code.

Step 3: Embed the dashboard URL

In the following section, you can find out how you can use the [Amazon Quick Sight Embedding SDK](#) (JavaScript) to embed the dashboard URL from step 3 in your website or application page. With the SDK, you can do the following:

- Place the dashboard on an HTML page.
- Pass parameters into the dashboard.
- Handle error states with messages that are customized to your application.

Call the `GenerateEmbedUrlForRegisteredUser` API operation to generate the URL that you can embed in your app. This URL is valid for 5 minutes, and the resulting session is valid for up to 10 hours. The API operation provides the URL with an `auth_code` that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-registered-user`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string  
for  
//readability and added ellipsis to indicate that it's incomplete.  
{  
  "Status": "200",  
  "EmbedUrl": "https://quicksightdomain/embed/12345/dashboards/67890..",  
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"  
}
```

Embed this dashboard in your webpage by using the [Amazon Quick Sight Embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as

your window size changes. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the dashboard and receive callbacks in terms of page load completion and errors.

The domain that is going to host embedded dashboards must be on the *allow list*, the list of approved domains for your QuickSight subscription. This requirement protects your data by keeping unapproved domains from hosting embedded dashboards. For more information about adding domains for embedded dashboards, see [Allow listing domains at runtime with the Amazon Quick Sight API](#).

The following example shows how to use the generated URL. This code is generated on your app server.

SDK 2.0

```
<!DOCTYPE html>
<html>

  <head>
    <title>Dashboard Embedding Example</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.0.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      const embedDashboard = async() => {
        const {
          createEmbeddingContext,
        } = QuickSightEmbedding;

        const embeddingContext = await createEmbeddingContext({
          onChange: (changeEvent, metadata) => {
            console.log('Context received a change', changeEvent,
metadata);
          },
        });

        const frameOptions = {
          url: '<YOUR_EMBED_URL>',
          container: '#experience-container',
          height: "700px",
          width: "1000px",
          onChange: (changeEvent, metadata) => {
            switch (changeEvent.eventName) {
              case 'FRAME_MOUNTED': {
```

```

        console.log("Do something when the experience frame is
mounted.");
        break;
    }
    case 'FRAME_LOADED': {
        console.log("Do something when the experience frame is
loaded.");
        break;
    }
},
};

const contentOptions = {
  parameters: [
    {
      Name: 'country',
      Values: [
        'United States'
      ],
    },
    {
      Name: 'states',
      Values: [
        'California',
        'Washington'
      ]
    }
  ],
  locale: "en-US",
  sheetOptions: {
    initialSheetId: '<YOUR_SHEETID>',
    singleSheet: false,
    emitSizeChangedEventOnSheetChange: false,
  },
  toolbarOptions: {
    export: false,
    undoRedo: false,
    reset: false
  },
  attributionOptions: {
    overlayContent: false,
  },
  onMessage: async (messageEvent, experienceMetadata) => {

```



```

        switch (messageEvent.eventName) {
            case 'CONTENT_LOADED': {
                console.log("All visuals are loaded. The title of the
document:", messageEvent.message.title);
                break;
            }
            case 'ERROR_OCCURRED': {
                console.log("Error occurred while rendering the
experience. Error code:", messageEvent.message.errorCode);
                break;
            }
            case 'PARAMETERS_CHANGED': {
                console.log("Parameters changed. Changed parameters:",
messageEvent.message.changedParameters);
                break;
            }
            case 'SELECTED_SHEET_CHANGED': {
                console.log("Selected sheet changed. Selected sheet:",
messageEvent.message.selectedSheet);
                break;
            }
            case 'SIZE_CHANGED': {
                console.log("Size changed. New dimensions:",
messageEvent.message);
                break;
            }
            case 'MODAL_OPENED': {
                window.scrollTo({
                    top: 0 // iframe top position
                });
                break;
            }
        }
    },
    );
    const embeddedDashboardExperience = await
embeddingContext.embedDashboard(frameOptions, contentOptions);

    const selectCountryElement = document.getElementById('country');
    selectCountryElement.addEventListener('change', (event) => {
        embeddedDashboardExperience.setParameters([
            {
                Name: 'country',
                Values: event.target.value
            }
        ])
    })
}

```

```

        }
    });
});
};
</script>
</head>

<body onload="embedDashboard()">
    <span>
        <label for="country">Country</label>
        <select id="country" name="country">
            <option value="United States">United States</option>
            <option value="Mexico">Mexico</option>
            <option value="Canada">Canada</option>
        </select>
    </span>
    <div id="experience-container"></div>
</body>

</html>

```

SDK 1.0

```

<!DOCTYPE html>
<html>

    <head>
        <title>Basic Embed</title>
        <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@1.0.15/dist/
quicksight-embedding-js-sdk.min.js"></script>
        <script type="text/javascript">
            var dashboard
            function onDashboardLoad(payload) {
                console.log("Do something when the dashboard is fully loaded.");
            }

            function onError(payload) {
                console.log("Do something when the dashboard fails loading");
            }

            function embedDashboard() {
                var containerDiv = document.getElementById("embeddingContainer");
                var options = {

```

```

        // replace this dummy url with the one generated via embedding API
        url: "https://us-east-1.quicksight.aws.amazon.com/sn/dashboards/
dashboardId?isauthcode=true&identityprovider=quicksight&code=authcode",
        container: containerDiv,
        parameters: {
            country: "United States"
        },
        scrolling: "no",
        height: "700px",
        width: "1000px",
        locale: "en-US",
        footerPaddingEnabled: true
    };
    dashboard = QuickSightEmbedding.embedDashboard(options);
    dashboard.on("error", onError);
    dashboard.on("load", onDashboardLoad);
}

function onCountryChange(obj) {
    dashboard.setParameters({country: obj.value});
}
</script>
</head>

<body onload="embedDashboard()">
    <span>
        <label for="country">Country</label>
        <select id="country" name="country" onchange="onCountryChange(this)">
            <option value="United States">United States</option>
            <option value="Mexico">Mexico</option>
            <option value="Canada">Canada</option>
        </select>
    </span>
    <div id="embeddingContainer"></div>
</body>

</html>

```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded dashboard on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight Embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.

- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Embedding Amazon Quick Sight dashboards for anonymous (unregistered) users

Important

Amazon Quick Sight has new API operations for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` API operations to embed dashboards and the Amazon Quick Sight console, but they don't contain the latest embedding capabilities. For more information about embedding using the old API operations, see [Embedding analytics using the `GetDashboardEmbedURL` and `GetSessionEmbedURL` API operations](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

In the following sections, you can find detailed information about how to set up embedded Amazon Quick Sight dashboards for anonymous (unregistered) users.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the dashboard URL](#)

Step 1: Set up permissions

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

In the following section, you can find out how to set up permissions for the backend application or web server. This task requires administrative access to IAM.

Each user who accesses a dashboard assumes a role that gives them Amazon Quick Sight access and permissions to the dashboard. To make this possible, create an IAM role in your Amazon Web Services account. Associate an IAM policy with the role to provide permissions to any user who assumes it.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants you as a developer the option to override the static domains that are configured in the **Manage Amazon Quick Sight** menu. Instead, you can list up to three domains or subdomains that can access a generated URL. This URL is then embedded in the website that you create. Only the domains that are listed in the parameter can access the embedded dashboard. Without this condition, you can list any domain on the internet in the `AllowedDomains` parameter.

To limit the domains that developers can use with this parameter, add an `AllowedEmbeddingDomains` condition to your IAM policy. For more information about the `AllowedDomains` parameter, see [GenerateEmbedUrlForAnonymousUser](#) in the *Amazon Quick Sight API Reference*.

The following sample policy provides these permissions for use with `GenerateEmbedUrlForAnonymousUser`. For this approach to work, you also need a session pack, or session capacity pricing, for your Amazon Web Services account. Otherwise, when a user tries to access the dashboard, the error `UnsupportedPricingPlanException` is returned.

Your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf to open the dashboard. The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

For more information regarding trust policies, see [Temporary security credentials in IAM](#) in the *IAM User Guide*.

Step 2: Generate the URL with the authentication code attached

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

In the following section, you can find how to authenticate on behalf of the anonymous visitor and get the embeddable dashboard URL on your application server.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then it adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

The following examples perform the IAM authentication on the user's behalf. It passes an identifier as the unique role session ID. This code runs on your app server.

Java

```
import java.util.List;
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import
com.amazonaws.services.quicksight.model.RegisteredUserDashboardEmbeddingConfiguration;
import
com.amazonaws.services.quicksight.model.AnonymousUserEmbeddingExperienceConfiguration;
import
com.amazonaws.services.quicksight.model.GenerateEmbedUrlForAnonymousUserRequest;
import
com.amazonaws.services.quicksight.model.GenerateEmbedUrlForAnonymousUserResult;
import com.amazonaws.services.quicksight.model.SessionTag;

/**
 * Class to call QuickSight Amazon SDK to generate embed url for anonymous user.
 */
public class GenerateEmbedUrlForAnonymousUserExample {

    private final AmazonQuickSight quickSightClient;

    public GenerateEmbedUrlForAnonymousUserExample() {
        quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AWSCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret key here
                    return new BasicAWSCredentials("access-key", "secret-key");
                }
            })
            .build();
    }
}
```

```

        }

        @Override
        public void refresh() {
        }
    }
)
    .build();
}

public String GenerateEmbedUrlForAnonymousUser(
    final String accountId, // YOUR Amazon ACCOUNT ID
    final String initialDashboardId, // DASHBOARD ID TO WHICH THE CONSTRUCTED
URL POINTS.
    final String namespace, // ANONYMOUS EMBEDDING REQUIRES SPECIFYING A VALID
NAMESPACE FOR WHICH YOU WANT THE EMBEDDING URL
    final List<String> authorizedResourceArns, // DASHBOARD ARN LIST TO EMBED
    final List<String> allowedDomains, // RUNTIME ALLOWED DOMAINS FOR EMBEDDING
    final List<SessionTag> sessionTags // SESSION TAGS USED FOR ROW-LEVEL
SECURITY
) throws Exception {
    AnonymousUserEmbeddingExperienceConfiguration experienceConfiguration = new
AnonymousUserEmbeddingExperienceConfiguration();
    AnonymousUserDashboardEmbeddingConfiguration dashboardConfiguration = new
AnonymousUserDashboardEmbeddingConfiguration();
    dashboardConfiguration.setInitialDashboardId(initialDashboardId);
    experienceConfiguration.setDashboard(dashboardConfiguration);

    GenerateEmbedUrlForAnonymousUserRequest
generateEmbedUrlForAnonymousUserRequest = new
GenerateEmbedUrlForAnonymousUserRequest()
    .withAwsAccountId(accountId)
    .withNamespace(namespace)
    .withAuthorizedResourceArns(authorizedResourceArns)
    .withExperienceConfiguration(experienceConfiguration)
    .withSessionTags(sessionTags)
    .withSessionLifetimeInMinutes(600L); // OPTIONAL: VALUE CAN BE
[15-600]. DEFAULT: 600
    .withAllowedDomains(allowedDomains);

    GenerateEmbedUrlForAnonymousUserResult dashboardEmbedUrl =
quickSightClient.generateEmbedUrlForAnonymousUser(generateEmbedUrlForAnonymousUserRequest);

    return dashboardEmbedUrl.getEmbedUrl();
}

```



```

    }

}

```

JavaScript

```

global.fetch = require('node-fetch');
const Amazon = require('aws-sdk');

function generateEmbedUrlForAnonymousUser(
  accountId, // YOUR Amazon ACCOUNT ID
  initialDashboardId, // DASHBOARD ID TO WHICH THE CONSTRUCTED URL POINTS
  quicksightNamespace, // VALID NAMESPACE WHERE YOU WANT TO DO NOAUTH EMBEDDING
  authorizedResourceArns, // DASHBOARD ARN LIST TO EMBED
  allowedDomains, // RUNTIME ALLOWED DOMAINS FOR EMBEDDING
  sessionTags, // SESSION TAGS USED FOR ROW-LEVEL SECURITY
  generateEmbedUrlForAnonymousUserCallback, // GENERATEEMBEDURLFORANONYMOUSUSER SUCCESS
    CALLBACK METHOD
  errorCallback // GENERATEEMBEDURLFORANONYMOUSUSER ERROR CALLBACK METHOD
) {
  const experienceConfiguration = {
    "DashboardVisual": {
      "InitialDashboardVisualId": {
        "DashboardId": "dashboard_id",
        "SheetId": "sheet_id",
        "VisualId": "visual_id"
      }
    }
  };

  const generateEmbedUrlForAnonymousUserParams = {
    "AwsAccountId": accountId,
    "Namespace": quicksightNamespace,
    "AuthorizedResourceArns": authorizedResourceArns,
    "AllowedDomains": allowedDomains,
    "ExperienceConfiguration": experienceConfiguration,
    "SessionTags": sessionTags,
    "SessionLifetimeInMinutes": 600
  };

  const quicksightClient = new AWS.QuickSight({
    region: process.env.AWS_REGION,
    credentials: {

```

```

        accessKeyId: AccessKeyId,
        secretAccessKey: SecretAccessKey,
        sessionToken: SessionToken,
        expiration: Expiration
    }
});

quicksightClient.generateEmbedUrlForAnonymousUser(generateEmbedUrlForAnonymousUserParams,
function(err, data) {
    if (err) {
        console.log(err, err.stack);
        errorCallback(err);
    } else {
        const result = {
            "statusCode": 200,
            "headers": {
                "Access-Control-Allow-Origin": "*", // USE YOUR WEBSITE DOMAIN TO
SECURE ACCESS TO THIS API
                "Access-Control-Allow-Headers": "Content-Type"
            },
            "body": JSON.stringify(data),
            "isBase64Encoded": false
        }
        generateEmbedUrlForAnonymousUserCallback(result);
    }
});
}

```

Python3

```

import json
import boto3
from botocore.exceptions import ClientError
import time

# Create QuickSight and STS clients
quicksightClient = boto3.client('quicksight', region_name='us-west-2')
sts = boto3.client('sts')

# Function to generate embedded URL for anonymous user
# accountId: YOUR AWS ACCOUNT ID
# quicksightNamespace: VALID NAMESPACE WHERE YOU WANT TO DO NOAUTH EMBEDDING
# authorizedResourceArns: DASHBOARD ARN LIST TO EMBED

```

```

# allowedDomains: RUNTIME ALLOWED DOMAINS FOR EMBEDDING
# dashboardId: DASHBOARD ID TO WHICH THE CONSTRUCTED URL POINTS
# sessionTags: SESSION TAGS USED FOR ROW-LEVEL SECURITY
def generateEmbedUrlForAnonymousUser(accountId, quicksightNamespace,
  authorizedResourceArns, allowedDomains, dashboardId, sessionTags):
try:
    response = quicksightClient.generate_embed_url_for_anonymous_user(
        AwsAccountId = accountId,
        Namespace = quicksightNamespace,
        AuthorizedResourceArns = authorizedResourceArns,
        AllowedDomains = allowedDomains,
        ExperienceConfiguration = {
            "Dashboard": {
                "InitialDashboardId": dashboardId
            }
        },
        SessionTags = sessionTags,
        SessionLifetimeInMinutes = 600
    )

    return {
        'statusCode': 200,
        'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-Headers":
"Content-Type"},
        'body': json.dumps(response),
        'isBase64Encoded': bool('false')
    }
except ClientError as e:
    print(e)
    return "Error generating embeddedURL: " + str(e)

```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```

const Amazon = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({

```

```
    apiConfig: require('./quicksight-2018-04-01.min.json'),
    region: 'us-east-1',
  });

quicksightClient.generateEmbedUrlForAnonymousUser({
  'AwsAccountId': '111122223333',
  'Namespace' : 'default',
  'AuthorizedResourceArns': authorizedResourceArns,
  'AllowedDomains': allowedDomains,
  'ExperienceConfiguration': experienceConfiguration,
  'SessionTags': sessionTags,
  'SessionLifetimeInMinutes': 600

}, function(err, data) {
  console.log('Errors: ');
  console.log(err);
  console.log('Response: ');
  console.log(data);
});
```

Example

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
  //readability and added ellipsis to indicate that it's incomplete.
  {
    Status: 200,
    EmbedUrl: 'https://quicksightdomain/embed/12345/dashboards/67890..',
    RequestId: '7bee030e-f191-45c4-97fe-d9faf0e03713'
  }
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;
```

```

var quicksightClient = new AmazonQuickSightClient(
    AccessKey,
    SecretAccessKey,
    sessionToken,
    Amazon.RegionEndpoint.USEast1);

try
{
    Console.WriteLine(
        quicksightClient.GenerateEmbedUrlForAnonymousUserAsync(new
GenerateEmbedUrlForAnonymousUserRequest
    {
        AwsAccountId = "111122223333",
        Namespace = default,
        AuthorizedResourceArns = authorizedResourceArns,
        AllowedDomains = allowedDomains,
        ExperienceConfiguration = experienceConfiguration,
        SessionTags = sessionTags,
        SessionLifetimeInMinutes = 600,
    }).Result.EmbedUrl
    );
} catch (Exception ex) {
    Console.WriteLine(ex.Message);
}

```

Amazon CLI

To assume the role, choose one of the following Amazon Security Token Service (Amazon STS) API operations:

- [AssumeRole](#) – Use this operation when you're using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you're using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you're using Security Assertion Markup Language (SAML) to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForAnonymousUser`.

```

aws sts assume-role \
    --role-arn "arn:aws-cn:iam::111122223333:role/QuickSightEmbeddingAnonymousPolicy" \

```

```
--role-session-name anonymous caller
```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

 Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. If you're using a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_dashboard_role/QuickSightEmbeddingAnonymousPolicy`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each visiting user. It also keeps each session separate and distinct. If you're using an array of web servers, for example for load balancing, and a session is reconnected to a different server, a new session begins.

To get a signed URL for the dashboard, call `generate-embed-url-for-anonymous-user` from the app server. This returns the embeddable dashboard URL. The following example shows how to generate the URL for an embedded dashboard using a server-side call for users who are making anonymous visits to your web portal or app.

```
aws quicksight generate-embed-url-for-anonymous-user \
--aws-account-id 111122223333 \
--namespace default-or-something-else \
```

```
--session-lifetime-in-minutes 15 \
--authorized-resource-arns '["dashboard-arn-1","dashboard-arn-2"]' \
--allowed-domains '["domain1","domain2"]' \
--session-tags '{"Key": tag-key-1, "Value": tag-value-1, {"Key": tag-key-1, "Value": tag-value-1}}' \
--experience-configuration
'DashboardVisual={InitialDashboardVisualId={DashboardId=dashboard_id, SheetId=sheet_id, VisualId=
```

For more information about using this operation, see [GenerateEmbedUrlForAnonymousUser](#). You can use this and other API operations in your own code.

Step 3: Embed the dashboard URL

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

In the following section, you can find out how you can use the [Amazon Quick Sight Embedding SDK](#) (JavaScript) to embed the dashboard URL from step 2 in your website or application page. With the SDK, you can do the following:

- Place the dashboard on an HTML page.
- Pass parameters into the dashboard.
- Handle error states with messages that are customized to your application.

Call the `GenerateEmbedUrlForAnonymousUser` API operation to generate the URL that you can embed in your app. This URL is valid for 5 minutes, and the resulting session is valid for 10 hours. The API operation provides the URL with an `auth_code` that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-anonymous-user`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.
{
    "Status": "200",
    "EmbedUrl": "https://quicksightdomain/embed/12345/dashboards/67890..",
```

```
    "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
  }
```

Embed this dashboard in your web page by using the [Amazon Quick Sight Embedding SDK](#) or by adding this URL into an iframe. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the dashboard and receive callbacks in terms of page load completion and errors.

The domain that is going to host embedded dashboards must be on the *allow list*, the list of approved domains for your QuickSight subscription. This requirement protects your data by keeping unapproved domains from hosting embedded dashboards. For more information about adding domains for embedded dashboards, see [Allow listing domains at runtime with the Amazon Quick Sight API](#).

The following example shows how to use the generated URL. This code resides on your app server.

SDK 2.0

```
<!DOCTYPE html>
<html>

<head>
  <title>Dashboard Embedding Example</title>
  <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.0.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
  <script type="text/javascript">
    const embedDashboard = async() => {
      const {
        createEmbeddingContext,
      } = QuickSightEmbedding;

      const embeddingContext = await createEmbeddingContext({
        onChange: (changeEvent, metadata) => {
          console.log('Context received a change', changeEvent, metadata);
        },
      });

      const frameOptions = {
        url: '<YOUR_EMBED_URL>',
```



```

        container: '#experience-container',
        height: "700px",
        width: "1000px",
        onChange: (changeEvent, metadata) => {
            switch (changeEvent.eventName) {
                case 'FRAME_MOUNTED': {
                    console.log("Do something when the experience frame is
mounted.");
                    break;
                }
                case 'FRAME_LOADED': {
                    console.log("Do something when the experience frame is
loaded.");
                    break;
                }
            }
        },
    };

    const contentOptions = {
        parameters: [
            {
                Name: 'country',
                Values: [
                    'United States'
                ],
            },
            {
                Name: 'states',
                Values: [
                    'California',
                    'Washington'
                ]
            }
        ],
        locale: "en-US",
        sheetOptions: {
            initialSheetId: '<YOUR_SHEETID>',
            singleSheet: false,
            emitSizeChangedEventOnSheetChange: false,
        },
        toolbarOptions: {
            export: false,
            undoRedo: false,

```

```

        reset: false
    },
    attributionOptions: {
        overlayContent: false,
    },
    onMessage: async (messageEvent, experienceMetadata) => {
        switch (messageEvent.eventName) {
            case 'CONTENT_LOADED': {
                console.log("All visuals are loaded. The title of the
document:", messageEvent.message.title);
                break;
            }
            case 'ERROR_OCCURRED': {
                console.log("Error occurred while rendering the experience.
Error code:", messageEvent.message.errorCode);
                break;
            }
            case 'PARAMETERS_CHANGED': {
                console.log("Parameters changed. Changed parameters:",
messageEvent.message.changedParameters);
                break;
            }
            case 'SELECTED_SHEET_CHANGED': {
                console.log("Selected sheet changed. Selected sheet:",
messageEvent.message.selectedSheet);
                break;
            }
            case 'SIZE_CHANGED': {
                console.log("Size changed. New dimensions:",
messageEvent.message);
                break;
            }
            case 'MODAL_OPENED': {
                window.scrollTo({
                    top: 0 // iframe top position
                });
                break;
            }
        }
    },
};

const embeddedDashboardExperience = await
embeddingContext.embedDashboard(frameOptions, contentOptions);

```

```

        const selectCountryElement = document.getElementById('country');
        selectCountryElement.addEventListener('change', (event) => {
            embeddedDashboardExperience.setParameters([
                {
                    Name: 'country',
                    Values: event.target.value
                }
            ]);
        });
    });
</script>
</head>

<body onload="embedDashboard()">
    <span>
        <label for="country">Country</label>
        <select id="country" name="country">
            <option value="United States">United States</option>
            <option value="Mexico">Mexico</option>
            <option value="Canada">Canada</option>
        </select>
    </span>
    <div id="experience-container"></div>
</body>

</html>

```

SDK 1.0

```

<!DOCTYPE html>
<html>

<head>
    <title>Basic Embed</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@1.0.15/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
        var dashboard
        function onDashboardLoad(payload) {
            console.log("Do something when the dashboard is fully loaded.");
        }

        function onError(payload) {

```

```

        console.log("Do something when the dashboard fails loading");
    }

    function embedDashboard() {
        var containerDiv = document.getElementById("embeddingContainer");
        var options = {
            // replace this dummy url with the one generated via embedding API
            url: "https://us-east-1.quicksight.aws.amazon.com/sn/dashboards/
dashboardId?isauthcode=true&identityprovider=quicksight&code=authcode",
            container: containerDiv,
            parameters: {
                country: "United States"
            },
            scrolling: "no",
            height: "700px",
            width: "1000px",
            locale: "en-US",
            footerPaddingEnabled: true
        };
        dashboard = QuickSightEmbedding.embedDashboard(options);
        dashboard.on("error", onError);
        dashboard.on("load", onDashboardLoad);
    }

    function onCountryChange(obj) {
        dashboard.setParameters({country: obj.value});
    }
</script>
</head>

<body onload="embedDashboard()">
    <span>
        <label for="country">Country</label>
        <select id="country" name="country" onchange="onCountryChange(this)">
            <option value="United States">United States</option>
            <option value="Mexico">Mexico</option>
            <option value="Canada">Canada</option>
        </select>
    </span>
    <div id="embeddingContainer"></div>
</body>

</html>

```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded dashboard on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight Embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest Amazon Quick Sight Embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Enabling executive summaries in embedded dashboards

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

You can enable executive summaries in your embedded dashboards. When enabled, registered users can generate executive summaries that provide a summary of all insights that Amazon Quick Sight has generated for the dashboard. Executive summaries make it easier for readers to find key insights and information about a dashboard. For more information about how users generate an executive summary of a dashboard, see [Generate an executive summary of an Amazon Quick Sight dashboard](#).

Note

Executive summaries are only available in embedded dashboards for registered users, and cannot be enabled in embedded dashboards for anonymous or unregistered users.

To enable executive summaries in embedded dashboards for registered users

- Follow the steps in [Embedding Amazon Quick Sight dashboards for registered users](#) to embed a dashboard with the following changes:

- a. When generating the URL in Step 2, set `Enabled: true` in the `ExecutiveSummary` parameter in the [GenerateEmbedUrlForRegisteredUser](#) or [GenerateEmbedUrlForRegisteredUserWithIdentity](#) as shown in the following example:

```
ExperienceConfiguration: {
  Dashboard: {
    InitialDashboardId: dashboard_id,
    FeatureConfigurations: {
      AmazonQInQuickSight: {
        ExecutiveSummary: {
          Enabled: true
        }
      }
    }
  }
}
```

- b. When embedding the dashboard URL with the Amazon Quick Sight Embedding SDK in Step 3, set `executiveSummary: true` in `contentOptions`, as shown in the following example:

```
const contentOptions = {
  toolbarOptions: {
    executiveSummary: true
  }
};
```

Embedding Amazon Quick Sight visuals with the Amazon Quick Sight APIs

You can embed individual visuals that are a part of a published dashboard in your application with the Amazon Quick Sight API.

Topics

- [Embedding Amazon Quick Sight visuals for registered users](#)
- [Embedding Amazon Quick Sight visuals for anonymous \(unregistered\) users](#)

Embedding Amazon Quick Sight visuals for registered users

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

In the following sections, you can find detailed information about how to set up embedded Amazon Quick Sight visuals for registered users of Amazon Quick Sight.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the visual URL](#)

Step 1: Set up permissions

In the following section, you can find out how to set up permissions for the backend application or web server. This task requires administrative access to IAM.

Each user who accesses a visual assumes a role that gives them Amazon Quick Sight access and permissions to the visual. To make this possible, create an IAM role in your Amazon Web Services account. Associate an IAM policy with the role to provide permissions to any user who assumes it. The IAM role needs to provide permissions to retrieve embedding URLs for a specific user pool. With the help of the wildcard character *, you can grant the permissions to generate a URL for all users in a specific namespace, or for a subset of users in specific namespaces. For this, you add `quicksight:GenerateEmbedUrlForRegisteredUser`.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants you as a developer the option to override the static domains that are configured in the **Manage Amazon Quick Sight** menu. Instead, you can list up to three domains or subdomains that can access a generated URL. This URL is then embedded in the website that you create. Only the domains that are listed in the parameter can access the embedded dashboard. Without this condition, you can list any domain on the internet in the `AllowedDomains` parameter.

To limit the domains that developers can use with this parameter, add an `AllowedEmbeddingDomains` condition to your IAM policy. For more information about the `AllowedDomains` parameter, see [GenerateEmbedUrlForRegisteredUser](#) in the *Amazon Quick Sight API Reference*.

The following sample policy provides these permissions.

Additionally, if you are creating first-time users who will be Amazon Quick Sight readers, make sure to add the `quicksight:RegisterUser` permission in the policy.

The following sample policy provides permission to retrieve an embedding URL for first-time users who are to be Amazon Quick Sight readers.

Finally, your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf and provision the user in Amazon Quick Sight. The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```


For more information regarding trust policies for OpenID Connect or SAML authentication, see the following sections of the *IAM User Guide*:

- [Creating a Role for Web Identity or OpenID Connect Federation \(Console\)](#)
- [Creating a Role for SAML 2.0 Federation \(Console\)](#)

Step 2: Generate the URL with the authentication code attached

In the following section, you can find out how to authenticate your Amazon Quick Sight user and get the embeddable visual URL on your application server. If you plan to embed visuals for IAM or Amazon Quick Sight identity types, share the visual with the Amazon Quick Sight users.

When a Amazon Quick Sight user accesses your app, the app assumes the IAM role on the Amazon Quick Sight user's behalf. Then it adds the user to Amazon Quick Sight, if that Amazon Quick Sight user doesn't already exist. Next, it passes an identifier as the unique role session ID.

Performing the described steps ensures that each viewer of the visual is uniquely provisioned in Amazon Quick Sight. It also enforces per-user settings, such as the row-level security and dynamic defaults for parameters.

The following examples perform the IAM authentication on the Amazon Quick Sight user's behalf. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import com.amazonaws.services.quicksight.model.DashboardVisualId;
import
    com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserRequest;
import com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserResult;
import
    com.amazonaws.services.quicksight.model.RegisteredUserDashboardVisualEmbeddingConfiguration;
import
    com.amazonaws.services.quicksight.model.RegisteredUserEmbeddingExperienceConfiguration;

import java.util.List;
```

```

/**
 * Class to call QuickSight Amazon SDK to get url for Visual embedding.
 */
public class GenerateEmbedUrlForRegisteredUserTest {

    private final AmazonQuickSight quickSightClient;

    public GenerateEmbedUrlForRegisteredUserTest() {
        this.quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AWSCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret key here
                    return new BasicAWSCredentials("access-key", "secret-key");
                }

                @Override
                public void refresh() {
                }
            })
            .build();
    }

    public String getEmbedUrl(
        final String accountId, // Amazon Account ID
        final String dashboardId, // Dashboard ID of the dashboard to embed
        final String sheetId, // Sheet ID of the sheet to embed
        final String visualId, // Visual ID of the visual to embed
        final List<String> allowedDomains, // Runtime allowed domains for embedding
        final String userArn // Registered user arn of the user that you want to
        provide embedded visual. Refer to Get Embed Url section in developer portal to find
        out how to get user arn for a QuickSight user.
    ) throws Exception {
        final DashboardVisualId dashboardVisual = new DashboardVisualId()
            .withDashboardId(dashboardId)
            .withSheetId(sheetId)
            .withVisualId(visualId);

        final RegisteredUserDashboardVisualEmbeddingConfiguration
registeredUserDashboardVisualEmbeddingConfiguration
            = new RegisteredUserDashboardVisualEmbeddingConfiguration()
                .withInitialDashboardVisualId(dashboardVisual);
    }
}

```

```

        final RegisteredUserEmbeddingExperienceConfiguration
registeredUserEmbeddingExperienceConfiguration
        = new RegisteredUserEmbeddingExperienceConfiguration()

.withDashboardVisual(registeredUserDashboardVisualEmbeddingConfiguration);
        final GenerateEmbedUrlForRegisteredUserRequest
generateEmbedUrlForRegisteredUserRequest
        = new GenerateEmbedUrlForRegisteredUserRequest()
            .withAwsAccountId(accountId)
            .withUserArn(userArn)

.withExperienceConfiguration(registeredUserEmbeddingExperienceConfiguration)
            .withAllowedDomains(allowedDomains);

        final GenerateEmbedUrlForRegisteredUserResult
generateEmbedUrlForRegisteredUserResult =
quickSightClient.generateEmbedUrlForRegisteredUser(generateEmbedUrlForRegisteredUserRequest);

        return generateEmbedUrlForRegisteredUserResult.getEmbedUrl();
    }
}

```

JavaScript

```

global.fetch = require('node-fetch');
const Amazon = require('aws-sdk');

function generateEmbedUrlForRegisteredUser(
    accountId, // Your Amazon account ID
    dashboardId, // Dashboard ID to which the constructed URL points
    sheetId, // Sheet ID to which the constructed URL points
    visualId, // Visual ID to which the constructed URL points
    openIdToken, // Cognito-based token
    userArn, // registered user arn
    roleArn, // IAM user role to use for embedding
    sessionName, // Session name for the roleArn assume role
    allowedDomains, // Runtime allowed domain for embedding
    getEmbedUrlCallback, // GetEmbedUrl success callback method
    errorCallback // GetEmbedUrl error callback method
) {
    const stsClient = new AWS.STS();
    let stsParams = {
        RoleSessionName: sessionName,

```

```

    WebIdentityToken: openIdToken,
    RoleArn: roleArn
  }

  stsClient.assumeRoleWithWebIdentity(stsParams, function(err, data) {
    if (err) {
      console.log('Error assuming role');
      console.log(err, err.stack);
      errorCallback(err);
    } else {
      const getDashboardParams = {
        "AwsAccountId": accountId,
        "ExperienceConfiguration": {
          "DashboardVisual": {
            "InitialDashboardVisualId": {
              "DashboardId": dashboardId,
              "SheetId": sheetId,
              "VisualId": visualId
            }
          }
        },
        "UserArn": userArn,
        "AllowedDomains": allowedDomains,
        "SessionLifetimeInMinutes": 600
      };

      const quicksightGetDashboard = new AWS.QuickSight({
        region: process.env.AWS_REGION,
        credentials: {
          accessKeyId: data.Credentials.AccessKeyId,
          secretAccessKey: data.Credentials.SecretAccessKey,
          sessionToken: data.Credentials.SessionToken,
          expiration: data.Credentials.Expiration
        }
      });

      quicksightGetDashboard.generateEmbedUrlForRegisteredUser(getDashboardParams,
        function(err, data) {
          if (err) {
            console.log(err, err.stack);
            errorCallback(err);
          } else {
            const result = {

```

```

        "statusCode": 200,
        "headers": {
            "Access-Control-Allow-Origin": "*", // Use your website
domain to secure access to GetEmbedUrl API
            "Access-Control-Allow-Headers": "Content-Type"
        },
        "body": JSON.stringify(data),
        "isBase64Encoded": false
    }
    getEmbedUrlCallback(result);
}
});
}
});
}
}

```

Python3

```

import json
import boto3
from botocore.exceptions import ClientError

sts = boto3.client('sts')

# Function to generate embedded URL
# accountId: Amazon account ID
# dashboardId: Dashboard ID to embed
# sheetId: SHEET ID to embed from the dashboard
# visualId: Id for the Visual you want to embedded from the dashboard sheet.
# userArn: arn of registered user
# allowedDomains: Runtime allowed domain for embedding
# roleArn: IAM user role to use for embedding
# sessionName: session name for the roleArn assume role
def getEmbeddingURL(accountId, dashboardId, sheetId, visualId, userArn, allowedDomains,
    roleArn, sessionName):
    try:
        assumedRole = sts.assume_role(
            RoleArn = roleArn,
            RoleSessionName = sessionName,
        )
    except ClientError as e:
        return "Error assuming role: " + str(e)
    else:

```

```

    assumedRoleSession = boto3.Session(
        aws_access_key_id = assumedRole['Credentials']['AccessKeyId'],
        aws_secret_access_key = assumedRole['Credentials']['SecretAccessKey'],
        aws_session_token = assumedRole['Credentials']['SessionToken'],
    )
    try:
        quicksightClient = assumedRoleSession.client('quicksight', region_name='us-
west-2')
        response = quicksightClient.generate_embed_url_for_registered_user(
            AwsAccountId=accountId,
            ExperienceConfiguration = {
                'DashboardVisual': {
                    'InitialDashboardVisualId': {
                        'DashboardId': dashboardId,
                        'SheetId': sheetId,
                        'VisualId': visualId
                    }
                },
            },
            UserArn = userArn,
            AllowedDomains = allowedDomains,
            SessionLifetimeInMinutes = 600
        )

        return {
            'statusCode': 200,
            'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-
Headers": "Content-Type"},
            'body': json.dumps(response),
            'isBase64Encoded': bool('false')
        }
    except ClientError as e:
        return "Error generating embedding url: " + str(e)

```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
const Amazon = require('aws-sdk');
```

```
const https = require('https');

var quicksightClient = new AWS.Service({
  apiConfig: require('./quicksight-2018-04-01.min.json'),
  region: 'us-east-1',
});

quicksightClient.generateEmbedUrlForRegisteredUser({
  'AwsAccountId': '111122223333',
  'ExperienceConfiguration': {
    'DashboardVisual': {
      'InitialDashboardVisualId': {
        'DashboardId': 'dashboard_id',
        'SheetId': 'sheet_id',
        'VisualId': 'visual_id'
      }
    }
  },
  'UserArn': 'REGISTERED_USER_ARN',
  'AllowedDomains': allowedDomains,
  'SessionLifetimeInMinutes': 100
}, function(err, data) {
  console.log('Errors: ');
  console.log(err);
  console.log('Response: ');
  console.log(data);
});
```

Example

```
//The URL returned is over 900 characters. For this example, we've shortened the string
//for
//readability and added ellipsis to indicate that it's incomplete.
{
  "Status": "200",
  "EmbedUrl": "https://quicksightdomain/embed/12345/dashboards/67890/
sheets/12345/visuals/67890...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;

namespace GenerateDashboardEmbedUrlForRegisteredUser
{
    class Program
    {
        static void Main(string[] args)
        {
            var quicksightClient = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
                SessionToken,
                Amazon.RegionEndpoint.USEast1);
            try
            {
                DashboardVisualId dashboardVisual = new DashboardVisualId
                {
                    DashboardId = "dashboard_id",
                    SheetId = "sheet_id",
                    VisualId = "visual_id"
                };

                RegisteredUserDashboardVisualEmbeddingConfiguration
registeredUserDashboardVisualEmbeddingConfiguration
                    = new RegisteredUserDashboardVisualEmbeddingConfiguration
                    {
                        InitialDashboardVisualId = dashboardVisual
                    };

                RegisteredUserEmbeddingExperienceConfiguration
registeredUserEmbeddingExperienceConfiguration
                    = new RegisteredUserEmbeddingExperienceConfiguration
```



```

        {
            DashboardVisual =
registeredUserDashboardVisualEmbeddingConfiguration
        };

        Console.WriteLine(
            quicksightClient.GenerateEmbedUrlForRegisteredUserAsync(new
GenerateEmbedUrlForRegisteredUserRequest
            {
                AwsAccountId = "111122223333",
                ExperienceConfiguration =
registeredUserEmbeddingExperienceConfiguration,
                UserArn = "REGISTERED_USER_ARN",
                AllowedDomains = allowedDomains,
                SessionLifetimeInMinutes = 100
            }).Result.EmbedUrl
        );
    } catch (Exception ex) {
        Console.WriteLine(ex.Message);
    }
}
}
}
}

```

Amazon CLI

To assume the role, choose one of the following Amazon Security Token Service (Amazon STS) API operations:

- [AssumeRole](#) – Use this operation when you're using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you're using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you're using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForRegisteredUser`. If you are taking a just-in-time approach to add users when they first open a dashboard, the role also needs permissions enabled for `quicksight:RegisterUser`.

```
aws sts assume-role \
    --role-arn "arn:aws-cn:iam::111122223333:role/embedding_quicksight_visual_role" \
```

```
--role-session-name john.doe@example.com
```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

 Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. If you're using a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_visual_role/john.doe@example.com`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. *Throttling* is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations.

The role session ID also becomes the user name in Amazon Quick Sight. You can use this pattern to provision your users in Amazon Quick Sight ahead of time, or to provision them the first time they access the dashboard.

The following example shows the CLI command that you can use to provision a user. For more information about [RegisterUser](#), [DescribeUser](#), and other Amazon Quick Sight API operations, see the [Amazon Quick Sight API Reference](#).

```
aws quicksight register-user \
```

```
--aws-account-id 111122223333 \
--namespace default \
--identity-type IAM \
--iam-arn "arn:aws-cn:iam::111122223333:role/embedding_quicksight_visual_role" \
--user-role READER \
--user-name jhnd \
--session-name "john.doe@example.com" \
--email john.doe@example.com \
--region us-east-1 \
--custom-permissions-name TeamA1
```

If the user is authenticated through Microsoft AD, you don't need to use `RegisterUser` to set them up. Instead, they should be automatically subscribed the first time they access Amazon Quick Sight. For Microsoft AD users, you can use `DescribeUser` to get the user ARN.

The first time a user accesses Amazon Quick Sight, you can also add this user to the group that the visual is shared with. The following example shows the CLI command to add a user to a group.

```
aws quicksight create-group-membership \
--aws-account-id=111122223333 \
--namespace=default \
--group-name=financeusers \
--member-name="embedding_quicksight_visual_role/john.doe@example.com"
```

You now have a user of your app who is also a user of Amazon Quick Sight, and who has access to the visual.

Finally, to get a signed URL for the visual, call `generate-embed-url-for-registered-user` from the app server. This returns the embeddable visual URL. The following example shows how to generate the URL for an embedded visual using a server-side call for users authenticated through Amazon Managed Microsoft AD or single sign-on (IAM Identity Center).

```
aws quicksight generate-embed-url-for-registered-user \
--aws-account-id 111122223333 \
--session-lifetime-in-minutes 600 \
--user-arn arn:aws-cn:quicksight:us-east-1:111122223333:user/default/
embedding_quicksight_visual_role/embeddingsession \
--allowed-domains '["domain1","domain2"]' \
--experience-configuration
'DashboardVisual={InitialDashboardVisualId={DashboardId=dashboard_id,SheetId=sheet_id,VisualId=visual_id}}
```

For more information about using this operation, see [GenerateEmbedUrlForRegisteredUser](#). You can use this and other API operations in your own code.

Step 3: Embed the visual URL

In the following section, you can find out how you can use the [Amazon Quick Sight Embedding SDK](#) (JavaScript) to embed the visual URL from step 3 in your website or application page. With the SDK, you can do the following:

- Place the visual on an HTML page.
- Pass parameters into the visual.
- Handle error states with messages that are customized to your application.

Call the `GenerateEmbedUrlForRegisteredUser` API operation to generate the URL that you can embed in your app. This URL is valid for 5 minutes, and the resulting session is valid for up to 10 hours. The API operation provides the URL with an `auth_code` that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-registered-user`. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
//for
//readability and added ellipsis to indicate that it's incomplete.
{
  "Status": "200",
  "EmbedUrl": "https://quicksightdomain/embed/12345/dashboards/67890/
sheets/12345/visuals/67890...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed this visual in your webpage by using the [Amazon Quick Sight Embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your

window size changes. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the visual and receive callbacks in terms of page load completion and errors.

The domain that is going to host embedded visuals and dashboards must be on the *allow list*, the list of approved domains for your QuickSight subscription. This requirement protects your data by keeping unapproved domains from hosting embedded visuals and dashboards. For more information about adding domains for embedded visuals and dashboards, see [Allow listing domains at runtime with the Amazon Quick Sight API](#).

The following example shows how to use the generated URL. This code is generated on your app server.

SDK 2.0

```
<!DOCTYPE html>
<html>

  <head>
    <title>Visual Embedding Example</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.0.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      const embedVisual = async() => {
        const {
          createEmbeddingContext,
        } = QuickSightEmbedding;

        const embeddingContext = await createEmbeddingContext({
          onChange: (changeEvent, metadata) => {
            console.log('Context received a change', changeEvent,
metadata);
          },
        });

        const frameOptions = {
          url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
          container: '#experience-container',
          height: "700px",
          width: "1000px",
          onChange: (changeEvent, metadata) => {
            switch (changeEvent.eventName) {
              case 'FRAME_MOUNTED': {
```

```

        console.log("Do something when the experience frame is
mounted.");
        break;
    }
    case 'FRAME_LOADED': {
        console.log("Do something when the experience frame is
loaded.");
        break;
    }
},
];

const contentOptions = {
    parameters: [
        {
            Name: 'country',
            Values: ['United States'],
        },
        {
            Name: 'states',
            Values: [
                'California',
                'Washington'
            ]
        }
    ],
    locale: "en-US",
    onMessage: async (messageEvent, experienceMetadata) => {
        switch (messageEvent.eventName) {
            case 'CONTENT_LOADED': {
                console.log("All visuals are loaded. The title of the
document:", messageEvent.message.title);
                break;
            }
            case 'ERROR_OCCURRED': {
                console.log("Error occurred while rendering the
experience. Error code:", messageEvent.message.errorCode);
                break;
            }
            case 'PARAMETERS_CHANGED': {
                console.log("Parameters changed. Changed parameters:",
messageEvent.message.changedParameters);
                break;
            }
        }
    }
};

```

```

        }
        case 'SIZE_CHANGED': {
            console.log("Size changed. New dimensions:",
messageEvent.message);
            break;
        }
    },
    },
    };
    const embeddedVisualExperience = await
embeddingContext.embedVisual(frameOptions, contentOptions);

    const selectCountryElement = document.getElementById('country');
    selectCountryElement.addEventListener('change', (event) => {
        embeddedVisualExperience.setParameters([
            {
                Name: 'country',
                Values: event.target.value
            }
        ]);
    });
};
</script>
</head>

<body onload="embedVisual()">
    <span>
        <label for="country">Country</label>
        <select id="country" name="country">
            <option value="United States">United States</option>
            <option value="Mexico">Mexico</option>
            <option value="Canada">Canada</option>
        </select>
    </span>
    <div id="experience-container"></div>
</body>

</html>

```

SDK 1.0

```

<!DOCTYPE html>
<html>

```

```
<head>
  <title>Visual Embedding Example</title>
  <!-- You can download the latest QuickSight embedding SDK version from https://
www.npmjs.com/package/amazon-quicksight-embedding-sdk -->
  <!-- Or you can do "npm install amazon-quicksight-embedding-sdk", if you use
npm for javascript dependencies -->
  <script src="./quicksight-embedding-js-sdk.min.js"></script>
  <script type="text/javascript">
    let embeddedVisualExperience;
    function onVisualLoad(payload) {
      console.log("Do something when the visual is fully loaded.");
    }

    function onError(payload) {
      console.log("Do something when the visual fails loading");
    }

    function embedVisual() {
      const containerDiv = document.getElementById("embeddingContainer");
      const options = {
        url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
        container: containerDiv,
        parameters: {
          country: "United States"
        },
        height: "700px",
        width: "1000px",
        locale: "en-US"
      };
      embeddedVisualExperience = QuickSightEmbedding.embedVisual(options);
      embeddedVisualExperience.on("error", onError);
      embeddedVisualExperience.on("load", onVisualLoad);
    }

    function onCountryChange(obj) {
      embeddedVisualExperience.setParameters({country: obj.value});
    }
  </script>
</head>

<body onload="embedVisual()">
  <span>
```



```
<label for="country">Country</label>
<select id="country" name="country" onchange="onCountryChange(this)">
  <option value="United States">United States</option>
  <option value="Mexico">Mexico</option>
  <option value="Canada">Canada</option>
</select>
</span>
<div id="embeddingContainer"></div>
</body>

</html>
```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded visual on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight Embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Embedding Amazon Quick Sight visuals for anonymous (unregistered) users

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

In the following sections, you can find detailed information about how to set up embedded Amazon Quick Sight visuals for anonymous (unregistered) users.

Topics

- [Step 1: Set up permissions](#)

- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the visual URL](#)

Step 1: Set up permissions

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

In the following section, you can find out how to set up permissions for the backend application or web server. This task requires administrative access to IAM.

Each user who accesses a visual assumes a role that gives them Amazon Quick Sight access and permissions to the visual. To make this possible, create an IAM role in your Amazon Web Services account. Associate an IAM policy with the role to provide permissions to any user who assumes it.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants you as a developer the option to override the static domains that are configured in the **Manage Amazon Quick Sight** menu. Instead, you can list up to three domains or subdomains that can access a generated URL. This URL is then embedded in the website that you create. Only the domains that are listed in the parameter can access the embedded dashboard. Without this condition, you can list any domain on the internet in the `AllowedDomains` parameter.

To limit the domains that developers can use with this parameter, add an `AllowedEmbeddingDomains` condition to your IAM policy. For more information about the `AllowedDomains` parameter, see [GenerateEmbedUrlForAnonymousUser](#) in the *Amazon Quick Sight API Reference*.

Your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf to open the visual. The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

For more information regarding trust policies, see [Temporary security credentials in IAM](#) in the *IAM User Guide*.

Step 2: Generate the URL with the authentication code attached

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

In the following section, you can find how to authenticate on behalf of the anonymous visitor and get the embeddable visual URL on your application server.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then it adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

The following examples perform the IAM authentication on the user's behalf. It passes an identifier as the unique role session ID. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import
    com.amazonaws.services.quicksight.model.AnonymousUserDashboardVisualEmbeddingConfiguration;
import
    com.amazonaws.services.quicksight.model.AnonymousUserEmbeddingExperienceConfiguration;
import com.amazonaws.services.quicksight.model.DashboardVisualId;
import com.amazonaws.services.quicksight.model.GenerateEmbedUrlForAnonymousUserRequest;
import com.amazonaws.services.quicksight.model.GenerateEmbedUrlForAnonymousUserResult;
import com.amazonaws.services.quicksight.model.SessionTag;

import java.util.List;

/**
 * Class to call QuickSight Amazon SDK to get url for Visual embedding.
 */
public class GenerateEmbedUrlForAnonymousUserTest {
    private final AmazonQuickSight quickSightClient;

    public GenerateEmbedUrlForAnonymousUserTest() {
        this.quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AWSCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret key here
                    return new BasicAWSCredentials("access-key", "secret-key");
                }
            })
    }
```

```

        @Override
        public void refresh() {
        }
    }
)
.build();
}

public String getEmbedUrl(
    final String accountId, // Amazon Account ID
    final String namespace, // Anonymous embedding required specifying a valid
namespace for which you want the embedding URL
    final List<String> authorizedResourceArns, // Dashboard arn list of
dashboard visuals to embed
    final String dashboardId, // Dashboard ID of the dashboard to embed
    final String sheetId, // Sheet ID of the sheet to embed
    final String visualId, // Visual ID of the visual to embed
    final List<String> allowedDomains, // Runtime allowed domains for embedding
    final List<SessionTag> sessionTags // Session tags used for row-level
security
) throws Exception {
    final DashboardVisualId dashboardVisual = new DashboardVisualId()
        .withDashboardId(dashboardId)
        .withSheetId(sheetId)
        .withVisualId(visualId);
    final AnonymousUserDashboardVisualEmbeddingConfiguration
anonymousUserDashboardVisualEmbeddingConfiguration
        = new AnonymousUserDashboardVisualEmbeddingConfiguration()
            .withInitialDashboardVisualId(dashboardVisual);
    final AnonymousUserEmbeddingExperienceConfiguration
anonymousUserEmbeddingExperienceConfiguration
        = new AnonymousUserEmbeddingExperienceConfiguration()

.withDashboardVisual(anonymousUserDashboardVisualEmbeddingConfiguration);
    final GenerateEmbedUrlForAnonymousUserRequest
generateEmbedUrlForAnonymousUserRequest
        = new GenerateEmbedUrlForAnonymousUserRequest()
            .withAwsAccountId(accountId)
            .withNamespace(namespace)
            // authorizedResourceArns should contain ARN of dashboard used below in
ExperienceConfiguration
            .withAuthorizedResourceArns(authorizedResourceArns)

.withExperienceConfiguration(anonymousUserEmbeddingExperienceConfiguration)

```

```

        .withAllowedDomains(allowedDomains)
        .withSessionTags(sessionTags)
        .withSessionLifetimeInMinutes(600L);

    final GenerateEmbedUrlForAnonymousUserResult
generateEmbedUrlForAnonymousUserResult
        =
quickSightClient.generateEmbedUrlForAnonymousUser(generateEmbedUrlForAnonymousUserRequest);

    return generateEmbedUrlForAnonymousUserResult.getEmbedUrl();
}
}

```

JavaScript

```

global.fetch = require('node-fetch');
const Amazon = require('aws-sdk');

function generateEmbedUrlForAnonymousUser(
    accountId, // Your Amazon account ID
    dashboardId, // Dashboard ID to which the constructed url points
    sheetId, // Sheet ID to which the constructed url points
    visualId, // Visual ID to which the constructed url points
    quicksightNamespace, // valid namespace where you want to do embedding
    authorizedResourceArns, // dashboard arn list of dashboard visuals to embed
    allowedDomains, // runtime allowed domains for embedding
    sessionTags, // session tags used for row-level security
    generateEmbedUrlForAnonymousUserCallback, // success callback method
    errorCallback // error callback method
) {
    const experienceConfiguration = {
        "DashboardVisual": {
            "InitialDashboardVisualId": {
                "DashboardId": dashboardId,
                "SheetId": sheetId,
                "VisualId": visualId
            }
        }
    };

    const generateEmbedUrlForAnonymousUserParams = {
        "AwsAccountId": accountId,
        "Namespace": quicksightNamespace,

```

```

    // authorizedResourceArns should contain ARN of dashboard used below in
    ExperienceConfiguration
    "AuthorizedResourceArns": authorizedResourceArns,
    "AllowedDomains": allowedDomains,
    "ExperienceConfiguration": experienceConfiguration,
    "SessionTags": sessionTags,
    "SessionLifetimeInMinutes": 600
  };

  const quicksightClient = new AWS.QuickSight({
    region: process.env.AWS_REGION,
    credentials: {
      accessKeyId: AccessKeyId,
      secretAccessKey: SecretAccessKey,
      sessionToken: SessionToken,
      expiration: Expiration
    }
  });

  quicksightClient.generateEmbedUrlForAnonymousUser(generateEmbedUrlForAnonymousUserParams,
  function(err, data) {
    if (err) {
      console.log(err, err.stack);
      errorCallback(err);
    } else {
      const result = {
        "statusCode": 200,
        "headers": {
          "Access-Control-Allow-Origin": "*", // USE YOUR WEBSITE DOMAIN TO
          SECURE ACCESS TO THIS API
          "Access-Control-Allow-Headers": "Content-Type"
        },
        "body": JSON.stringify(data),
        "isBase64Encoded": false
      }
      generateEmbedUrlForAnonymousUserCallback(result);
    }
  });
}

```

Python3

```
import json
import boto3
from botocore.exceptions import ClientError
import time

# Create QuickSight and STS clients
quicksightClient = boto3.client('quicksight', region_name='us-west-2')
sts = boto3.client('sts')

# Function to generate embedded URL for anonymous user
# accountId: YOUR AWS ACCOUNT ID
# quicksightNamespace: VALID NAMESPACE WHERE YOU WANT TO DO NOAUTH EMBEDDING
# authorizedResourceArns: DASHBOARD ARN LIST TO EMBED
# allowedDomains: RUNTIME ALLOWED DOMAINS FOR EMBEDDING
# experienceConfiguration: DASHBOARD ID, SHEET ID and VISUAL ID TO WHICH THE
#   CONSTRUCTED URL POINTS
# Example experienceConfig -> 'DashboardVisual': {
#   'InitialDashboardVisualId': {
#     'DashboardId': 'dashboardId',
#     'SheetId': 'sheetId',
#     'VisualId': 'visualId'
#   }
# },
# sessionTags: SESSION TAGS USED FOR ROW-LEVEL SECURITY
def generateEmbedUrlForAnonymousUser(accountId, quicksightNamespace,
    authorizedResourceArns, allowedDomains, experienceConfiguration, sessionTags):
    try:
        response = quicksightClient.generate_embed_url_for_anonymous_user(
            AwsAccountId = accountId,
            Namespace = quicksightNamespace,
            AuthorizedResourceArns = authorizedResourceArns,
            AllowedDomains = allowedDomains,
            ExperienceConfiguration = experienceConfiguration,
            SessionTags = sessionTags,
            SessionLifetimeInMinutes = 600
        )

        return {
            'statusCode': 200,
            'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-Headers": "Content-Type"},
            'body': json.dumps(response),
```



```

        'isBase64Encoded': bool('false')
    }
except ClientError as e:
    print(e)
    return "Error generating embeddedURL: " + str(e)

```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```

const Amazon = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({
    apiConfig: require('./quicksight-2018-04-01.min.json'),
    region: 'us-east-1',
});

quicksightClient.generateEmbedUrlForAnonymousUser({
    'AwsAccountId': '111122223333',
    'Namespace' : 'default',
    // authorizedResourceArns should contain ARN of dashboard used below in
    ExperienceConfiguration
    'AuthorizedResourceArns': authorizedResourceArns,
    'ExperienceConfiguration': {
        'DashboardVisual': {
            'InitialDashboardVisualId': {
                'DashboardId': 'dashboard_id',
                'SheetId': 'sheet_id',
                'VisualId': 'visual_id'
            }
        }
    },
    'AllowedDomains': allowedDomains,
    'SessionTags': sessionTags,
    'SessionLifetimeInMinutes': 600
}, function(err, data) {
    console.log('Errors: ');

```

```
console.log(err);
console.log('Response: ');
console.log(data);
});
```

Example

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.
{
    "Status": "200",
    "EmbedUrl": "https://quicksightdomain/embed/12345/dashboards/67890/
sheets/12345/visuals/67890...",
    "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;

namespace GenerateDashboardEmbedUrlForAnonymousUser
{
    class Program
    {
        static void Main(string[] args)
        {
            var quicksightClient = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
                SessionToken,
                Amazon.RegionEndpoint.USEast1);

            try
            {
                DashboardVisualId dashboardVisual = new DashboardVisualId
```

```

        {
            DashboardId = "dashboard_id",
            SheetId = "sheet_id",
            VisualId = "visual_id"
        };

        AnonymousUserDashboardVisualEmbeddingConfiguration
anonymousUserDashboardVisualEmbeddingConfiguration
            = new AnonymousUserDashboardVisualEmbeddingConfiguration
            {
                InitialDashboardVisualId = dashboardVisual

            };

        AnonymousUserEmbeddingExperienceConfiguration
anonymousUserEmbeddingExperienceConfiguration
            = new AnonymousUserEmbeddingExperienceConfiguration
            {
                DashboardVisual =
anonymousUserDashboardVisualEmbeddingConfiguration
            };

        Console.WriteLine(
            quicksightClient.GenerateEmbedUrlForAnonymousUserAsync(new
GenerateEmbedUrlForAnonymousUserRequest
            {
                AwsAccountId = "111222333444",
                Namespace = default,
                // authorizedResourceArns should contain ARN of dashboard used
below in ExperienceConfiguration
                AuthorizedResourceArns = { "dashboard_id" },
                ExperienceConfiguration =
anonymousUserEmbeddingExperienceConfiguration,
                SessionTags = sessionTags,
                SessionLifetimeInMinutes = 600,
            }).Result.EmbedUrl
        );
    } catch (Exception ex) {
        Console.WriteLine(ex.Message);
    }
}
}
}
}

```

Amazon CLI

To assume the role, choose one of the following Amazon Security Token Service (Amazon STS) API operations:

- [AssumeRole](#) – Use this operation when you're using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you're using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you're using Security Assertion Markup Language (SAML) to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForAnonymousUser`.

```
aws sts assume-role \  
  --role-arn "arn:aws-cn:iam::11112222333:role/QuickSightEmbeddingAnonymousPolicy" \  
  --role-session-name anonymous caller
```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. If you're using a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"  
  export AWS_SECRET_ACCESS_KEY = "secret_key_from_assume_role"  
  export AWS_SESSION_TOKEN    = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_visual_role/QuickSightEmbeddingAnonymousPolicy`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each visiting user. It also keeps each session separate and distinct. If you're using an array of web servers, for example for load balancing, and a session is reconnected to a different server, a new session begins.

To get a signed URL for the visual, call `generate-embed-url-for-anonymous-user` from the app server. This returns the embeddable visual URL. The following example shows how to generate the URL for an embedded visual using a server-side call for users who are making anonymous visits to your web portal or app.

```
aws quicksight generate-embed-url-for-anonymous-user \
  --aws-account-id 111122223333 \
  --namespace default-or-something-else \
  --session-lifetime-in-minutes 15 \
  --authorized-resource-arns '["dashboard-arn-1","dashboard-arn-2"]' \
  --allowed-domains '["domain1","domain2"]' \
  --session-tags '[{"Key": tag-key-1,"Value": tag-value-1},{"Key": tag-
key-1,"Value": tag-value-1}]' \
  --experience-configuration
'DashboardVisual={InitialDashboardVisualId={DashboardId=dashboard_id,SheetId=sheet_id,VisualId=
```

For more information about using this operation, see [GenerateEmbedUrlForAnonymousUser](#). You can use this and other API operations in your own code.

Step 3: Embed the visual URL

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

In the following section, you can find out how you can use the [Amazon Quick Sight Embedding SDK](#) (JavaScript) to embed the visual URL from step 2 in your website or application page. With the SDK, you can do the following:

- Place the visual on an HTML page.
- Pass parameters into the visual.
- Handle error states with messages that are customized to your application.

Call the `GenerateEmbedUrlForAnonymousUser` API operation to generate the URL that you can embed in your app. This URL is valid for 5 minutes, and the resulting session is valid for 10 hours. The API operation provides the URL with an authorization (auth) code that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-anonymous-user`. The *quicksightdomain* in this example is the URL that you use to access your Amazon Quick Sight account.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.
{
    "Status": "200",
    "EmbedUrl": "https://quicksightdomain/embed/12345/dashboards/67890/
sheets/12345/visuals/67890...",
    "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed this visual in your web page by using the Amazon Quick Sight [Embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the visual and receive callbacks in terms of visual load completion and errors.

The domain that is going to host embedded visual must be on the *allow list*, the list of approved domains for your QuickSight subscription. This requirement protects your data by keeping unapproved domains from hosting embedded visuals and dashboards. For more information about adding domains for embedded visuals and dashboards, see [Allow listing domains at runtime with the Amazon Quick Sight API](#).

The following example shows how to use the generated URL. This code resides on your app server.

SDK 2.0

```

<!DOCTYPE html>
<html>

  <head>
    <title>Visual Embedding Example</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.0.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      const embedVisual = async() => {
        const {
          createEmbeddingContext,
        } = QuickSightEmbedding;

        const embeddingContext = await createEmbeddingContext({
          onChange: (changeEvent, metadata) => {
            console.log('Context received a change', changeEvent,
metadata);
          },
        });

        const frameOptions = {
          url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
          container: '#experience-container',
          height: "700px",
          width: "1000px",
          onChange: (changeEvent, metadata) => {
            switch (changeEvent.eventName) {
              case 'FRAME_MOUNTED': {
                console.log("Do something when the experience frame is
mounted.");
                break;
              }
              case 'FRAME_LOADED': {
                console.log("Do something when the experience frame is
loaded.");
                break;
              }
            }
          },
        };
      };
    
  

```

```

const contentOptions = {
  parameters: [
    {
      Name: 'country',
      Values: ['United States'],
    },
    {
      Name: 'states',
      Values: [
        'California',
        'Washington'
      ]
    }
  ],
  locale: "en-US",
  onMessage: async (messageEvent, experienceMetadata) => {
    switch (messageEvent.eventName) {
      case 'CONTENT_LOADED': {
        console.log("All visuals are loaded. The title of the
document:", messageEvent.message.title);
        break;
      }
      case 'ERROR_OCCURRED': {
        console.log("Error occurred while rendering the
experience. Error code:", messageEvent.message.errorCode);
        break;
      }
      case 'PARAMETERS_CHANGED': {
        console.log("Parameters changed. Changed parameters:",
messageEvent.message.changedParameters);
        break;
      }
      case 'SIZE_CHANGED': {
        console.log("Size changed. New dimensions:",
messageEvent.message);
        break;
      }
    }
  },
};

const embeddedVisualExperience = await
embeddingContext.embedVisual(frameOptions, contentOptions);

const selectCountryElement = document.getElementById('country');

```



```

        selectCountryElement.addEventListener('change', (event) => {
            embeddedVisualExperience.setParameters([
                {
                    Name: 'country',
                    Values: event.target.value
                }
            ]);
        });
    });
</script>
</head>

<body onload="embedVisual()">
    <span>
        <label for="country">Country</label>
        <select id="country" name="country">
            <option value="United States">United States</option>
            <option value="Mexico">Mexico</option>
            <option value="Canada">Canada</option>
        </select>
    </span>
    <div id="experience-container"></div>
</body>

</html>

```

SDK 1.0

```

<!DOCTYPE html>
<html>

    <head>
        <title>Visual Embedding Example</title>
        <!-- You can download the latest QuickSight embedding SDK version from https://
www.npmjs.com/package/amazon-quicksight-embedding-sdk -->
        <!-- Or you can do "npm install amazon-quicksight-embedding-sdk", if you use
npm for javascript dependencies -->
        <script src="./quicksight-embedding-js-sdk.min.js"></script>
        <script type="text/javascript">
            let embeddedVisualExperience;
            function onVisualLoad(payload) {
                console.log("Do something when the visual is fully loaded.");
            }

```

```

function onError(payload) {
    console.log("Do something when the visual fails loading");
}

function embedVisual() {
    const containerDiv = document.getElementById("embeddingContainer");
    const options = {
        url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
        container: containerDiv,
        parameters: {
            country: "United States"
        },
        height: "700px",
        width: "1000px",
        locale: "en-US"
    };
    embeddedVisualExperience = QuickSightEmbedding.embedVisual(options);
    embeddedVisualExperience.on("error", onError);
    embeddedVisualExperience.on("load", onVisualLoad);
}

function onCountryChange(obj) {
    embeddedVisualExperience.setParameters({country: obj.value});
}
</script>
</head>

<body onload="embedVisual()">
    <span>
        <label for="country">Country</label>
        <select id="country" name="country" onchange="onCountryChange(this)">
            <option value="United States">United States</option>
            <option value="Mexico">Mexico</option>
            <option value="Canada">Canada</option>
        </select>
    </span>
    <div id="embeddingContainer"></div>
</body>

</html>

```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded visual on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight Embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest QuickSight embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Embedding the full functionality of the Amazon Quick Sight console for registered users

Important

Amazon Quick Sight has new API operations for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` API operations to embed dashboards and the Amazon Quick Sight console, but they don't contain the latest embedding capabilities. For more information about embedding using the old API operations, see [Embedding analytics using the `GetDashboardEmbedURL` and `GetSessionEmbedURL` API operations](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

With Enterprise edition, in addition to providing read-only dashboards you can also provide the Amazon Quick Sight console experience in a custom-branded authoring portal. Using this approach, you allow your users to create data sources, datasets, and analyses. In the same interface, they can create, publish, and view dashboards. If you want to restrict some of those permissions, you can also do that.

Users who access Amazon Quick Sight through an embedded console need to belong to the author or admin security cohort. Readers don't have enough access to use the Amazon Quick Sight console for authoring, regardless of whether it's embedded or part of the Amazon Web Services Management Console. However, authors and admins can still access embedded dashboards. If you want to restrict permissions to some of the authoring features, you can add a custom permissions profile to the user with the [UpdateUser](#) API operation. Use the [RegisterUser](#) API operation to add a new user with a custom permission profile attached. For more information, see the following sections:

- For information about creating custom roles by defining custom console permissions, see [Customizing Access to the Amazon Quick Sight Console](#).
- For information about using namespaces to isolate multitenancy users, groups, and Amazon Quick Sight assets, see [Amazon Quick Sight Namespaces](#).
- For information about adding your own branding to an embedded Amazon Quick Sight console, see [Using Themes in Amazon Quick Sight](#) and the [QuickSight Theme API Operations](#).

In the following sections, you can find detailed information about how to set up embedded Amazon Quick Sight dashboards for registered users.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the console session URL](#)
- [Enabling Generative BI features in embedded consoles for registered users](#)

Step 1: Set up permissions

In the following section, you can find out how to set up permissions for the backend application or web server. This task requires administrative access to IAM.

Each user who accesses a Amazon Quick Sight assumes a role that gives them Amazon Quick Sight access and permissions to the console session. To make this possible, create an IAM role in your Amazon account. Associate an IAM policy with the role to provide permissions to any user who assumes it. Add `quicksight:RegisterUser` permissions to ensure that the reader can access Amazon Quick Sight in a read-only fashion, and not have access to any other data or creation

capability. The IAM role also needs to provide permissions to retrieve console session URLs. For this, you add `quicksight:GenerateEmbedUrlForRegisteredUser`.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants you as a developer the option to override the static domains that are configured in the **Manage Amazon Quick Sight** menu. Instead, you can list up to three domains or subdomains that can access a generated URL. This URL is then embedded in the website that you create. Only the domains that are listed in the parameter can access the embedded dashboard. Without this condition, you can list any domain on the internet in the `AllowedDomains` parameter.

The following sample policy provides these permissions.

The following sample policy provides permission to retrieve a console session URL. You can use the policy without `quicksight:RegisterUser` if you are creating users before they access an embedded session.

Finally, your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf and provision the user in Amazon Quick Sight. The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      }
    }
  ]
}
```

```
        },  
        "Action": "sts:AssumeRole"  
    }  
]  
}
```

For more information regarding trust policies for OpenID Connect or SAML authentication, see the following sections of the *IAM User Guide*:

- [Creating a Role for Web Identity or OpenID Connect Federation \(Console\)](#)
- [Creating a Role for SAML 2.0 Federation \(Console\)](#)

Step 2: Generate the URL with the authentication code attached

In the following section, you can find out how to authenticate your user and get the embeddable console session URL on your application server.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then it adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

Performing the described steps ensures that each viewer of the console session is uniquely provisioned in Amazon Quick Sight. It also enforces per-user settings, such as the row-level security and dynamic defaults for parameters.

The following examples perform the IAM authentication on the user's behalf. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;  
import com.amazonaws.auth.BasicAWSCredentials;  
import com.amazonaws.auth.AWSCredentialsProvider;  
import com.amazonaws.regions.Regions;  
import com.amazonaws.services.quicksight.AmazonQuickSight;  
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;  
import  
    com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserRequest;  
import com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserResult;  
import  
    com.amazonaws.services.quicksight.model.RegisteredUserEmbeddingExperienceConfiguration;
```

```

import
    com.amazonaws.services.quicksight.model.RegisteredUserQuickSightConsoleEmbeddingConfiguration;

/**
 * Class to call QuickSight Amazon SDK to get url for QuickSight console embedding.
 */
public class GetQuicksightEmbedUrlRegisteredUserQSConsoleEmbedding {

    private final AmazonQuickSight quickSightClient;

    public GetQuicksightEmbedUrlRegisteredUserQSConsoleEmbedding() {
        this.quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AWSCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret key here
                    return new BasicAWSCredentials("access-key", "secret-key");
                }

                @Override
                public void refresh() {
                }
            })
            .build();
    }

    public String getQuicksightEmbedUrl(
        final String accountId,
        final String userArn, // Registered user arn to use for embedding. Refer
        // to Get Embed Url section in developer portal to find out how to get user arn for a
        // QuickSight user.
        final List<String> allowedDomains, // Runtime allowed domain for embedding
        final String initialPath
    ) throws Exception {
        final RegisteredUserEmbeddingExperienceConfiguration experienceConfiguration =
            new RegisteredUserEmbeddingExperienceConfiguration()
                .withQuickSightConsole(new
                    RegisteredUserQuickSightConsoleEmbeddingConfiguration().withInitialPath(initialPath));
        final GenerateEmbedUrlForRegisteredUserRequest
            generateEmbedUrlForRegisteredUserRequest = new
                GenerateEmbedUrlForRegisteredUserRequest();
    }

```

```

        generateEmbedUrlForRegisteredUserRequest.setAwsAccountId(accountId);
        generateEmbedUrlForRegisteredUserRequest.setUserArn(userArn);
        generateEmbedUrlForRegisteredUserRequest.setAllowedDomains(allowedDomains);

generateEmbedUrlForRegisteredUserRequest.setExperienceConfiguration(experienceConfiguration);

        final GenerateEmbedUrlForRegisteredUserResult
generateEmbedUrlForRegisteredUserResult =
quickSightClient.generateEmbedUrlForRegisteredUser(generateEmbedUrlForRegisteredUserRequest);

        return generateEmbedUrlForRegisteredUserResult.getEmbedUrl();
    }
}

```

JavaScript

```

global.fetch = require('node-fetch');
const Amazon = require('aws-sdk');

function generateEmbedUrlForRegisteredUser(
    accountId,
    dashboardId,
    openIdToken, // Cognito-based token
    userArn, // registered user arn
    roleArn, // IAM user role to use for embedding
    sessionName, // Session name for the roleArn assume role
    allowedDomains, // Runtime allowed domain for embedding
    getEmbedUrlCallback, // GetEmbedUrl success callback method
    errorCallback // GetEmbedUrl error callback method
) {
    const stsClient = new AWS.STS();
    let stsParams = {
        RoleSessionName: sessionName,
        WebIdentityToken: openIdToken,
        RoleArn: roleArn
    }

    stsClient.assumeRoleWithWebIdentity(stsParams, function(err, data) {
        if (err) {
            console.log('Error assuming role');
            console.log(err, err.stack);
            errorCallback(err);
        } else {

```



```

const getDashboardParams = {
  "AwsAccountId": accountId,
  "ExperienceConfiguration": {
    "QuickSightConsole": {
      "InitialPath": '/start'
    }
  },
  "UserArn": userArn,
  "AllowedDomains": allowedDomains,
  "SessionLifetimeInMinutes": 600
};

const quicksightGetDashboard = new AWS.QuickSight({
  region: process.env.AWS_REGION,
  credentials: {
    accessKeyId: data.Credentials.AccessKeyId,
    secretAccessKey: data.Credentials.SecretAccessKey,
    sessionToken: data.Credentials.SessionToken,
    expiration: data.Credentials.Expiration
  }
});

quicksightGetDashboard.generateEmbedUrlForRegisteredUser(getDashboardParams,
function(err, data) {
  if (err) {
    console.log(err, err.stack);
    errorCallback(err);
  } else {
    const result = {
      "statusCode": 200,
      "headers": {
        "Access-Control-Allow-Origin": "*", // Use your website
domain to secure access to GetEmbedUrl API
        "Access-Control-Allow-Headers": "Content-Type"
      },
      "body": JSON.stringify(data),
      "isBase64Encoded": false
    }
    getEmbedUrlCallback(result);
  }
});
}
});

```

```
}
```

Python3

```
import json
import boto3
from botocore.exceptions import ClientError

# Create QuickSight and STS clients
qs = boto3.client('quicksight', region_name='us-east-1')
sts = boto3.client('sts')

# Function to generate embedded URL
# accountId: Amazon account ID
# userArn: arn of registered user
# allowedDomains: Runtime allowed domain for embedding
# roleArn: IAM user role to use for embedding
# sessionName: session name for the roleArn assume role
def generateEmbeddingURL(accountId, userArn, allowedDomains, roleArn, sessionName):
    try:
        assumedRole = sts.assume_role(
            RoleArn = roleArn,
            RoleSessionName = sessionName,
        )
    except ClientError as e:
        return "Error assuming role: " + str(e)
    else:
        assumedRoleSession = boto3.Session(
            aws_access_key_id = assumedRole['Credentials']['AccessKeyId'],
            aws_secret_access_key = assumedRole['Credentials']['SecretAccessKey'],
            aws_session_token = assumedRole['Credentials']['SessionToken'],
        )
        try:
            quickSightClient = assumedRoleSession.client('quicksight', region_name='us-east-1')

            experienceConfiguration = {
                "QuickSightConsole": {
                    "InitialPath": "/start"
                }
            }
            response = quickSightClient.generate_embed_url_for_registered_user(
                AwsAccountId = accountId,
```

```

        ExperienceConfiguration = experienceConfiguration,
        UserArn = userArn,
        AllowedDomains = allowedDomains,
        SessionLifetimeInMinutes = 600
    )

    return {
        'statusCode': 200,
        'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-Headers": "Content-Type"},
        'body': json.dumps(response),
        'isBase64Encoded': bool('false')
    }
except ClientError as e:
    return "Error generating embedding url: " + str(e)

```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded console session. You can use this URL in your website or app to display the console session.

Example

```

const Amazon = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({
    apiConfig: require('./quicksight-2018-04-01.min.json'),
    region: 'us-east-1',
});

quicksightClient.generateEmbedUrlForRegisteredUser({
    'AwsAccountId': '111122223333',
    'ExperienceConfiguration': {
        'QuickSightConsole': {
            'InitialPath': '/start'
        }
    },
    'UserArn': 'REGISTERED_USER_ARN',
    'AllowedDomains': allowedDomains,
    'SessionLifetimeInMinutes': 100
}, function(err, data) {

```

```
console.log('Errors: ');
console.log(err);
console.log('Response: ');
console.log(data);
});
```

Example

```
// The URL returned is over 900 characters. For this example, we've shortened the
// string for
// readability and added ellipsis to indicate that it's incomplete.
{
    Status: 200,
    EmbedUrl: 'https://quicksightdomain/embed/12345/dashboards/67890..',
    RequestId: '7bee030e-f191-45c4-97fe-d9faf0e03713'
}
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded console session. You can use this URL in your website or app to display the console.

Example

```
using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;

namespace GenerateDashboardEmbedUrlForRegisteredUser
{
    class Program
    {
        static void Main(string[] args)
        {
            var quicksightClient = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
                SessionToken,
                Amazon.RegionEndpoint.USEast1);

            try
            {
```

```

        RegisteredUserQuickSightConsoleEmbeddingConfiguration
registeredUserQuickSightConsoleEmbeddingConfiguration
    = new RegisteredUserQuickSightConsoleEmbeddingConfiguration
    {
        InitialPath = "/start"
    };
    RegisteredUserEmbeddingExperienceConfiguration
registeredUserEmbeddingExperienceConfiguration
    = new RegisteredUserEmbeddingExperienceConfiguration
    {
        QuickSightConsole =
registeredUserQuickSightConsoleEmbeddingConfiguration
    };

    Console.WriteLine(
        quicksightClient.GenerateEmbedUrlForRegisteredUserAsync(new
GenerateEmbedUrlForRegisteredUserRequest
    {
        AwsAccountId = "111122223333",
        ExperienceConfiguration =
registeredUserEmbeddingExperienceConfiguration,
        UserArn = "REGISTERED_USER_ARN",
        AllowedDomains = allowedDomains,
        SessionLifetimeInMinutes = 100
    }).Result.EmbedUrl
    );
} catch (Exception ex) {
    Console.WriteLine(ex.Message);
}
}
}
}
}

```

Amazon CLI

To assume the role, choose one of the following Amazon Security Token Service (Amazon STS) API operations:

- [AssumeRole](#) – Use this operation when you're using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you're using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you're using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForRegisteredUser`. If you are taking a just-in-time approach to add users when they first open Amazon Quick Sight, the role also needs permissions enabled for `quicksight:RegisterUser`.

```
aws sts assume-role \  
  --role-arn "arn:aws-cn:iam::111122223333:role/embedding_quicksight_dashboard_role" \  
  --role-session-name john.doe@example.com
```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. If you're using a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"  
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"  
export AWS_SESSION_TOKEN     = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_console_session_role/john.doe@example.com`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. Throttling is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations.

The role session ID also becomes the user name in Amazon Quick Sight. You can use this pattern to provision your users in Amazon Quick Sight ahead of time, or to provision them the first time they access a console session.

The following example shows the CLI command that you can use to provision a user. For more information about [RegisterUser](#), [DescribeUser](#), and other Amazon Quick Sight API operations, see the [Amazon Quick Sight API Reference](#).

```
aws quicksight register-user \  
  --aws-account-id 111122223333 \  
  --namespace default \  
  --identity-type IAM \  
  --iam-arn "arn:aws-cn:iam::111122223333:role/embedding_quicksight_dashboard_role" \  
  --user-role READER \  
  --user-name jhnd \  
  --session-name "john.doe@example.com" \  
  --email john.doe@example.com \  
  --region us-east-1 \  
  --custom-permissions-name TeamA1
```

If the user is authenticated through Microsoft AD, you don't need to use RegisterUser to set them up. Instead, they should be automatically subscribed the first time they access Amazon Quick Sight. For Microsoft AD users, you can use DescribeUser to get the user ARN.

The first time a user accesses Amazon Quick Sight, you can also add this user to the appropriate group. The following example shows the CLI command to add a user to a group.

```
aws quicksight create-group-membership \  
  --aws-account-id=111122223333 \  
  --namespace=default \  
  --group-name=financeusers \  
  --member-name="embedding_quicksight_dashboard_role/john.doe@example.com"
```

You now have a user of your app who is also a user of Amazon Quick Sight, and who has access to the Amazon Quick Sight console session.

Finally, to get a signed URL for the console session, call generate-embed-url-for-registered-user from the app server. This returns the embeddable console session URL. The following example shows how to generate the URL for an embedded console session using a

server-side call for users authenticated through Amazon Managed Microsoft AD or single sign-on (IAM Identity Center).

```
aws quicksight generate-embed-url-for-registered-user \
  --aws-account-id 111122223333 \
  --entry-point the-url-for--the-console-session \
  --session-lifetime-in-minutes 600 \
  --user-arn arn:aws-cn:quicksight:us-east-1:111122223333:user/default/
embedding_quicksight_dashboard_role/embeddingsession
--allowed-domains '["domain1","domain2"]' \
  --experience-configuration QuickSightConsole={InitialPath="/start"}
```

For more information about using this operation, see [GenerateEmbedUrlForRegisteredUser](#). You can use this and other API operations in your own code.

Step 3: Embed the console session URL

In the following section, you can find out how you can use the [Amazon Quick Sight Embedding SDK](#) (JavaScript) to embed the console session URL from step 3 in your website or application page. With the SDK, you can do the following:

- Place the console session on an HTML page.
- Pass parameters into the console session.
- Handle error states with messages that are customized to your application.

Call the `GenerateEmbedUrlForRegisteredUser` API operation to generate the URL that you can embed in your app. This URL is valid for 5 minutes, and the resulting session is valid for up to 10 hours. The API operation provides the URL with an `auth_code` that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-registered-user`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.
{
  "Status": "200",
  "EmbedUrl": "https://quicksightdomain/embedding/12345/start...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```


Embed this console session in your webpage by using the Amazon Quick Sight [Embedding SDK](#) or by adding this URL into an iframe. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the console session and receive callbacks in terms of page load completion and errors.

The domain that is going to host embedded dashboards must be on the *allow list*, the list of approved domains for your QuickSight subscription. This requirement protects your data by keeping unapproved domains from hosting embedded dashboards. For more information about adding domains for an embedded console, see [Allow listing domains at runtime with the Amazon Quick Sight API](#).

The following example shows how to use the generated URL. This code is generated on your app server.

SDK 2.0

```
<!DOCTYPE html>
<html>

  <head>
    <title>Console Embedding Example</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.0.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      const embedSession = async() => {
        const {
          createEmbeddingContext,
        } = QuickSightEmbedding;

        const embeddingContext = await createEmbeddingContext({
          onChange: (changeEvent, metadata) => {
            console.log('Context received a change', changeEvent,
metadata);
          },
        });

        const frameOptions = {
          url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
```

```

        container: '#experience-container',
        height: "700px",
        width: "1000px",
        onChange: (changeEvent, metadata) => {
            switch (changeEvent.eventName) {
                case 'FRAME_MOUNTED': {
                    console.log("Do something when the experience frame is
mounted.");

                    break;
                }
                case 'FRAME_LOADED': {
                    console.log("Do something when the experience frame is
loaded.");

                    break;
                }
            }
        },
    };

    const contentOptions = {
        onMessage: async (messageEvent, experienceMetadata) => {
            switch (messageEvent.eventName) {
                case 'ERROR_OCCURRED': {
                    console.log("Do something when the embedded experience
fails loading.");

                    break;
                }
            }
        }
    };

    const embeddedConsoleExperience = await
embeddingContext.embedConsole(frameOptions, contentOptions);
    };
</script>
</head>

<body onload="embedSession()">
    <div id="experience-container"></div>
</body>

</html>

```

SDK 1.0

```
<!DOCTYPE html>
<html>

  <head>
    <title>QuickSight Console Embedding</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@1.0.15/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      var session

      function onError(payload) {
        console.log("Do something when the session fails loading");
      }

      function embedSession() {
        var containerDiv = document.getElementById("embeddingContainer");
        var options = {
          // replace this dummy url with the one generated via embedding API
          url: "https://us-east-1.quicksight.aws.amazon.com/sn/dashboards/
dashboardId?isauthcode=true&identityprovider=quicksight&code=authcode", // replace this
dummy url with the one generated via embedding API
          container: containerDiv,
          parameters: {
            country: "United States"
          },
          scrolling: "no",
          height: "700px",
          width: "1000px",
          locale: "en-US",
          footerPaddingEnabled: true,
          defaultEmbeddingVisualType: "TABLE", // this option only applies to
QuickSight console embedding and is not used for dashboard embedding
        };
        session = QuickSightEmbedding.embedSession(options);
        session.on("error", onError);
      }

      function onCountryChange(obj) {
        session.setParameters({country: obj.value});
      }
    </script>
  </head>
```

```
<body onload="embedSession()">
  <span>
    <label for="country">Country</label>
    <select id="country" name="country" onchange="onCountryChange(this)">
      <option value="United States">United States</option>
      <option value="Mexico">Mexico</option>
      <option value="Canada">Canada</option>
    </select>
  </span>
  <div id="embeddingContainer"></div>
</body>

</html>
```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded console session on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight Embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Enabling Generative BI features in embedded consoles for registered users

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

You can enable the following Generative BI features in your embedded console:

- **Executive summaries:** When enabled, registered Author Pro and Reader Pro users can generate executive summaries that provide a summary of all insights that Amazon Quick Sight has generated for the dashboard to easily discover key insights.
- **Authoring:** When enabled, Author Pro users can use Generative BI to build calculated fields and build and refine visuals.
- **Q&A:** When enabled, Author Pro and Reader Pro users can use the AI-powered Q&A to both suggest and answer questions related their data.
- **Data stories:** When enabled, Author Pro and Reader Pro users can provide details to quickly generate a first draft of their data story.

To enable Generative BI features in embedded consoles for registered users

- Follow the steps in [Embedding the full functionality of the Amazon Quick Sight console for registered users](#) to embed a console with the following changes:
 - a. When generating the URL in Step 2, set Enabled: true in the FeatureConfigurations parameter for each of the features you want to enable in the [GenerateEmbedUrlForRegisteredUser](#) or [GenerateEmbedUrlForRegisteredUserWithIdentity](#) APIs, as shown in the following example. If no configuration is provided, the features are disabled by default.

```
ExperienceConfiguration: {
  QuickSightConsole: {
    InitialPath: "initial_path",
    AmazonQInQuickSight: {
      FeatureConfigurations: {
        COMMENT: Enable executive summaries
        ExecutiveSummary: {
          Enabled: true
        },
        COMMENT: Enable Generative BI authoring
        GenerativeAuthoring: {
          Enabled: true
        },
        COMMENT: Enable Q&A
        DataQnA: {
          Enabled: true
        },
        COMMENT: Enable data stories
```

```
    DataStories: {  
      Enabled: true  
    }  
  }  
}
```

- b. When embedding the console URL with the Amazon Quick Sight Embedding SDK in Step 3, set the values in the following example as desired. If no configuration is provided, the features are disabled by default.

Note

There is no SDK option for enabling data stories. If data stories are enabled with the API as shown in the previous step, they will be available to registered users.

```
const contentOptions = {  
  toolbarOptions: {  
    executiveSummary: true, // Enable executive summaries  
    buildVisual: true, // Enable Generative BI authoring  
    dataQnA: true // Enable Q&A  
  }  
};
```

Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience

Intended audience: Amazon Quick Suite developers

In the following sections, you can find detailed information about how to set up an embedded Generative Q&A experience that uses enhanced NLQ capabilities powered by LLMs. The Generative Q&A experience is the recommended replacement for the embedded Q Search Bar and provides an updated BI experience for users.

Topics

- [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience for registered users](#)
- [Embedding the Amazon Q in Quick Suite Generative Q&A experience for anonymous \(unregistered\) users](#)

Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience for registered users

In the following sections, you can find detailed information about how to set up an embedded Generative Q&A experience for registered users of Amazon Quick Sight.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the Generative Q&A experience URL](#)
- [Optional embedded Generative Q&A experience functionalities](#)

Step 1: Set up permissions

In the following section, you can find how to set up permissions for your backend application or web server to embed the Generative Q&A experience. This task requires administrative access to Amazon Identity and Access Management (IAM).

Each user who accesses a Generative Q&A experience assumes a role that gives them Amazon Quick Sight access and permissions. To make this possible, create an IAM role in your Amazon Web Services account. Associate an IAM policy with the role to provide permissions to any user who assumes it. The IAM role needs to provide permissions to retrieve embedding URLs for a specific user pool.

With the help of the wildcard character *, you can grant the permissions to generate a URL for all users in a specific namespace. Or you can grant permissions to generate a URL for a subset of users in specific namespaces. For this, you add `quicksight:GenerateEmbedUrlForRegisteredUser`.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForRegisteredUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants developers the option to override

the static domains that are configured in the **Manage Amazon Quick Sight** menu and instead list up to three domains or subdomains that can access a generated URL. This URL is then embedded in a developer's website. Only the domains that are listed in the parameter can access the embedded Generative Q&A experience. Without this condition, developers can list any domain on the internet in the `AllowedDomains` parameter.

To limit the domains that developers can use with this parameter, add an `AllowedEmbeddingDomains` condition to your IAM policy. For more information about the `AllowedDomains` parameter, see [GenerateEmbedUrlForRegisteredUser](#) in the *Amazon Quick Sight API Reference*.

The following sample policy provides these permissions.

Also, if you're creating first-time users who will be Amazon Quick Sight readers, make sure to add the `quicksight:RegisterUser` permission in the policy.

The following sample policy provides permission to retrieve an embedding URL for first-time users who are to be Amazon Quick Sight readers.

Finally, your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf and provision the user in Amazon Quick Sight.

The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
```



```
        "Principal": {  
            "Service": "ec2.amazonaws.com"  
        },  
        "Action": "sts:AssumeRole"  
    }  
]  
}
```

For more information regarding trust policies for OpenID Connect or Security Assertion Markup Language (SAML) authentication, see the following sections of the *IAM User Guide*:

- [Creating a role for web identity or OpenID Connect federation \(console\)](#)
- [Creating a role for SAML 2.0 federation \(console\)](#)

Step 2: Generate the URL with the authentication code attached

In the following section, you can find how to authenticate your user and get the embeddable Q topic URL on your application server. If you plan to embed the Generative Q&A experience for IAM or Amazon Quick Sight identity types, share the Q topic with the users.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then the app adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

Performing the described steps ensures that each viewer of the Q topic is uniquely provisioned in Amazon Quick Sight. It also enforces per-user settings, such as the row-level security and dynamic defaults for parameters. Tag-based row-level security can be used for anonymous user embedding of the Q bar.

The following examples perform the IAM authentication on the user's behalf. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;  
import com.amazonaws.auth.BasicAWSCredentials;  
import com.amazonaws.auth.AWSCredentialsProvider;  
import com.amazonaws.regions.Regions;  
import com.amazonaws.services.quicksight.AmazonQuickSight;  
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
```

```

import
    com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserRequest;
import com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserResult;
import
    com.amazonaws.services.quicksight.model.RegisteredUserEmbeddingExperienceConfiguration;
import
    com.amazonaws.services.quicksight.model.RegisteredUserGenerativeQnAEmbeddingConfiguration;

/**
 * Class to call QuickSight AWS SDK to get url for embedding Generative Q&A experience.
 */
public class RegisteredUserGenerativeQnAEmbeddingSample {

    private final AmazonQuickSight quickSightClient;

    public RegisteredUserGenerativeQnAEmbeddingSample() {
        this.quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AmazonCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret key here
                    return new BasicAWSCredentials("access-key", "secret-
key");
                }

                @Override
                public void refresh() {
                }
            })
            .build();
    }

    public String getQuicksightEmbedUrl(
        final String accountId, // Amazon Account ID
        final String topicId, // Topic ID to embed
        final List<String> allowedDomains, // Runtime allowed domain for embedding
        final String userArn // Registered user arn to use for embedding. Refer to
        Get Embed Url section in developer portal to find how to get user arn for a QuickSight
        user.
    ) throws Exception {

```

```

        final RegisteredUserEmbeddingExperienceConfiguration experienceConfiguration =
new RegisteredUserEmbeddingExperienceConfiguration()
        .withGenerativeQnA(new
RegisteredUserGenerativeQnAEmbeddingConfiguration().withInitialTopicId(topicId));
        final GenerateEmbedUrlForRegisteredUserRequest
generateEmbedUrlForRegisteredUserRequest = new
GenerateEmbedUrlForRegisteredUserRequest();
        generateEmbedUrlForRegisteredUserRequest.setAwsAccountId(accountId);
        generateEmbedUrlForRegisteredUserRequest.setUserArn(userArn);
        generateEmbedUrlForRegisteredUserRequest.setAllowedDomains(allowedDomains);

generateEmbedUrlForRegisteredUserRequest.setExperienceConfiguration(experienceConfiguration);

        final GenerateEmbedUrlForRegisteredUserResult
generateEmbedUrlForRegisteredUserResult =
quickSightClient.generateEmbedUrlForRegisteredUser(generateEmbedUrlForRegisteredUserRequest);

        return generateEmbedUrlForRegisteredUserResult.getEmbedUrl();
    }
}

```

JavaScript

Note

Embed URL generation APIs cannot be called from browsers directly. Refer to the Node.JS example instead.

Python3

```

import json
import boto3
from botocore.exceptions import ClientError

sts = boto3.client('sts')

# Function to generate embedded URL
# accountId: Amazon account ID
# topicId: Topic ID to embed
# userArn: arn of registered user
# allowedDomains: Runtime allowed domain for embedding
# roleArn: IAM user role to use for embedding

```

```

# sessionName: session name for the roleArn assume role
def getEmbeddingURL(accountId, topicId, userArn, allowedDomains, roleArn, sessionName):
    try:
        assumedRole = sts.assume_role(
            RoleArn = roleArn,
            RoleSessionName = sessionName,
        )
    except ClientError as e:
        return "Error assuming role: " + str(e)
    else:
        assumedRoleSession = boto3.Session(
            aws_access_key_id = assumedRole['Credentials']['AccessKeyId'],
            aws_secret_access_key = assumedRole['Credentials']['SecretAccessKey'],
            aws_session_token = assumedRole['Credentials']['SessionToken'],
        )
        try:
            quicksightClient = assumedRoleSession.client('quicksight', region_name='us-
west-2')
            response = quicksightClient.generate_embed_url_for_registered_user(
                AwsAccountId=accountId,
                ExperienceConfiguration = {
                    'GenerativeQnA': {
                        'InitialTopicId': topicId
                    }
                },
                UserArn = userArn,
                AllowedDomains = allowedDomains,
                SessionLifetimeInMinutes = 600
            )

            return {
                'statusCode': 200,
                'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-
Headers": "Content-Type"},
                'body': json.dumps(response),
                'isBase64Encoded': bool('false')
            }
        except ClientError as e:
            return "Error generating embedding url: " + str(e)

```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
const Amazon = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({
    region: 'us-east-1'
});

quicksightClient.generateEmbedUrlForRegisteredUser({
    'AwsAccountId': '111122223333',
    'ExperienceConfiguration': {
        'GenerativeQnA': {
            'InitialTopicId': 'U4zJMVZ2n2stZflc80u3iKySEb3BEV6f'
        }
    },
    'UserArn': 'REGISTERED_USER_ARN',
    'AllowedDomains': allowedDomains,
    'SessionLifetimeInMinutes': 100
}, function(err, data) {
    console.log('Errors: ');
    console.log(err);
    console.log('Response: ');
    console.log(data);
});
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded Q search bar. You can use this URL in your website or app to display the Q search bar.

Example

```
using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;
```

```

namespace GenerateGenerativeQnAEmbedUrlForRegisteredUser
{
    class Program
    {
        static void Main(string[] args)
        {
            var quicksightClient = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
                SessionToken,
                Amazon.RegionEndpoint.USEast1);

            try
            {
                RegisteredUserGenerativeQnAEmbeddingConfiguration
registeredUserGenerativeQnAEmbeddingConfiguration
                = new RegisteredUserGenerativeQnAEmbeddingConfiguration
                {
                    InitialTopicId = "U4zJMVZ2n2stZflc80u3iKySEb3BEV6f"
                };

                RegisteredUserEmbeddingExperienceConfiguration
registeredUserEmbeddingExperienceConfiguration
                = new RegisteredUserEmbeddingExperienceConfiguration
                {
                    GenerativeQnA =
registeredUserGenerativeQnAEmbeddingConfiguration
                };

                Console.WriteLine(
                    quicksightClient.GenerateEmbedUrlForRegisteredUserAsync(new
GenerateEmbedUrlForRegisteredUserRequest
                    {
                        AwsAccountId = "111122223333",
                        ExperienceConfiguration =
registeredUserEmbeddingExperienceConfiguration,
                        UserArn = "REGISTERED_USER_ARN",
                        AllowedDomains = allowedDomains,
                        SessionLifetimeInMinutes = 100
                    }).Result.EmbedUrl
                );
            } catch (Exception ex) {
                Console.WriteLine(ex.Message);
            }
        }
    }
}

```

```
}  
}
```

Amazon CLI

To assume the role, choose one of the following Amazon Security Token Service (Amazon STS) API operations:

- [AssumeRole](#) – Use this operation when you are using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you are using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you are using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForRegisteredUser`. If you are taking a just-in-time approach to add users when they use a topic in the Q search bar, the role also needs permissions enabled for `quicksight:RegisterUser`.

```
aws sts assume-role \  
  --role-arn "arn:aws:iam::111122223333:role/  
embedding_quicksight_q_generative_qna_role" \  
  --role-session-name john.doe@example.com
```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. For a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_q_search_bar_role/john.doe@example.com`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. *Throttling* is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations.

The role session ID also becomes the user name in Amazon Quick Sight. You can use this pattern to provision your users in Amazon Quick Sight ahead of time, or to provision them the first time that they access the Generative Q&A experience.

The following example shows the CLI command that you can use to provision a user. For more information about [RegisterUser](#), [DescribeUser](#), and other Amazon Quick Sight API operations, see the [Amazon Quick Sight API reference](#).

```
aws quicksight register-user \
  --aws-account-id 111122223333 \
  --namespace default \
  --identity-type IAM \
  --iam-arn "arn:aws:iam::111122223333:role/
embedding_quicksight_q_generative_qna_role" \
  --user-role READER \
  --user-name jhnd \
  --session-name "john.doe@example.com" \
  --email john.doe@example.com \
  --region us-east-1 \
  --custom-permissions-name TeamA1
```

If the user is authenticated through Microsoft AD, you don't need to use `RegisterUser` to set them up. Instead, they should be automatically subscribed the first time that they access Amazon Quick Sight. For Microsoft AD users, you can use `DescribeUser` to get the user Amazon Resource Name (ARN).

The first time a user accesses Amazon Quick Sight, you can also add this user to the group that the dashboard is shared with. The following example shows the CLI command to add a user to a group.

```
aws quicksight create-group-membership \  
  --aws-account-id 111122223333 \  
  --namespace default \  
  --group-name financeusers \  
  --member-name "embedding_quicksight_q_generative_qna_role/john.doe@example.com"
```

You now have a user of your app who is also a user of Amazon Quick Sight, and who has access to the dashboard.

Finally, to get a signed URL for the dashboard, call `generate-embed-url-for-registered-user` from the app server. This returns the embeddable dashboard URL. The following example shows how to generate the URL for an embedded dashboard using a server-side call for users authenticated through Amazon Managed Microsoft AD or single sign-on (IAM Identity Center).

```
aws quicksight generate-embed-url-for-anonymous-user \  
  --aws-account-id 111122223333 \  
  --namespace default-or-something-else \  
  --authorized-resource-arns '["topic-arn-topicId1","topic-arn-topicId2"]' \  
  --allowed-domains '["domain1","domain2"]' \  
  --experience-configuration 'GenerativeQnA={InitialTopicId="topicId1"}' \  
  --session-tags '["Key": tag-key-1,"Value": tag-value-1,{"Key": tag-key-1,"Value": tag-value-1}]' \  
  --session-lifetime-in-minutes 15
```

For more information about using this operation, see [GenerateEmbedUrlForRegisteredUser](#). You can use this and other API operations in your own code.

Step 3: Embed the Generative Q&A experience URL

In the following section, you can find how to embed the Generative Q&A experience URL in your website or application page. You do this with the [Amazon Quick Sight embedding SDK](#) (JavaScript). With the SDK, you can do the following:

- Place the Generative Q&A experience on an HTML page.
- Customize the layout and appearance of the embedded experience to fit your application needs.
- Handle error states with messages that are customized to your application.

To generate the URL that you can embed in your app, call the `GenerateEmbedUrlForRegisteredUser` API operation. This URL is valid for 5 minutes, and the resulting session is valid for up to 10 hours. The API operation provides the URL with an `auth_code` value that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-registered-user`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.
{
  "Status": "200",
  "EmbedUrl": "https://quicksightdomain/embedding/12345/q/search...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed the Generative Q&A experience in your webpage by using the [Amazon Quick Sight embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes.

Make sure that the domain to host the embedded Generative Q&A experience is on the *allow list*, the list of approved domains for your Amazon Quick Sight subscription. This requirement protects your data by keeping unapproved domains from hosting embedded dashboards. For more information about adding domains for an embedded Generative Q&A experience, see [Managing domains](#).

You can use the Amazon Quick Sight Embedding SDK to customize the layout and appearance of the embedded Generative Q&A experience to fit your application. Use the `panelType` property to configure the landing state of the Generative Q&A experience when it renders in your application. Set the `panelType` property to `'FULL'` to render the full Generative Q&A experience panel. This panel resembles the experience that Amazon Quick Sight users have in the Amazon Quick Sight console. The frame height of the panel is not changed based on user interaction and respects the value that you set in the `frameOptions.height` property. The image below shows the Generative Q&A experience panel that renders when you set the `panelType` value to `'FULL'`.

Set the `panelType` property to `'SEARCH_BAR'` to render the Generative Q&A experience as a search bar. This search bar resembles the way that the Q Search Bar renders when it is embedded

into an application. The Generative Q&A search bar expands to a larger panel that displays topic selection options, the question suggestion list, the answer panel or the pinboard.

The default minimum height of the Generative Q&A search bar is rendered when the embedded asset loads. It is recommended that you set the `frameOptions.height` value to "38px" to optimize the search bar experience. Use the `focusedHeight` property to set the optimal size of the topic selection dropdown and the question suggestion list. Use the `expandedHeight` property to set the optimal size of the answer panel and pinboard. If you choose the 'SEARCH_BAR' option, it is recommended that you style the parent container with position; absolute to avoid unwanted content shifting in your application. The image below shows the Generative Q&A experience search bar that renders when you set the `panelType` value to 'SEARCH_BAR'.

After you configure the `panelType` property, use the Amazon Quick Sight embedding SDK to customize the following properties of the Generative Q&A experience.

- The title of the Generative Q&A panel (Applies only to the `panelType`: FULL option).
- The search bar's placeholder text.
- Whether topic selection is allowed.
- Whether topic names are shown or hidden.
- Whether the Amazon Q icon is shown or hidden (Applies only to the `panelType`: FULL option).
- Whether the pinboard is shown or hidden.
- Whether users can maximize the Generative Q&A panel to fullscreen.
- The theme of the Generative Q&A panel. A custom theme ARN can be passed in the SDK to change the appearance of the frame's content. Amazon Quick Sight starter themes are not supported for embedded Generative BI panels. To use a Amazon Quick Sight starter theme, save it as a custom theme in Amazon Quick Sight.

When you use the Amazon Quick Sight Embedding SDK, the Generative Q&A experience on your page is dynamically resized based on the state. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the Generative Q&A experience and receive callbacks in terms of page load completion, state changes, and errors.

The following example shows how to use the generated URL. This code is generated on your app server.

SDK 2.0

```

<!DOCTYPE html>
<html>
  <head>
    <title>Generative Q&A Embedding Example</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.7.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      const embedGenerativeQnA = async() => {
        const {createEmbeddingContext} = QuickSightEmbedding;

        const embeddingContext = await createEmbeddingContext({
          onChange: (changeEvent, metadata) => {
            console.log('Context received a change', changeEvent,
metadata);
          },
        });

        const frameOptions = {
          url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
          container: '#experience-container',
          height: "700px",
          width: "1000px",
          onChange: (changeEvent, metadata) => {
            switch (changeEvent.eventName) {
              case 'FRAME_MOUNTED': {
                console.log("Do something when the experience frame is
mounted.");
                break;
              }
              case 'FRAME_LOADED': {
                console.log("Do something when the experience frame is
loaded.");
                break;
              }
            }
          },
        };

        const contentOptions = {
          // Optional panel settings. Default behavior is equivalent to
{panelType: 'FULL'}

```

```

        panelOptions: {
            panelType: 'FULL',
            title: 'custom title', // Optional
            showQIcon: false, // Optional, Default: true
        },
        // Use SEARCH_BAR panel type for the landing state to be similar to
embedQSearchBar
        // with generative capability enabled topics
        /*
        panelOptions: {
            panelType: 'SEARCH_BAR',
            focusedHeight: '250px',
            expandedHeight: '500px',
        },
        */
        showTopicName: false, // Optional, Default: true
        showPinboard: false, // Optional, Default: true
        allowTopicSelection: false, // Optional, Default: true
        allowFullscreen: false, // Optional, Default: true
        searchPlaceholderText: "custom search placeholder", // Optional
        themeOptions: { // Optional
            themeArn: 'arn:aws:quicksight:<Region>:<AWS-Account-ID>:theme/
<Theme-ID>'
        }
        onMessage: async (messageEvent, experienceMetadata) => {
            switch (messageEvent.eventName) {
                case 'Q_SEARCH_OPENED': {
                    // called when pinboard is shown / visuals are rendered
                    console.log("Do something when SEARCH_BAR type panel is
expanded");

                    break;
                }
                case 'Q_SEARCH_FOCUSED': {
                    // called when question suggestions or topic selection
dropdown are shown

                    console.log("Do something when SEARCH_BAR type panel is
focused");

                    break;
                }
                case 'Q_SEARCH_CLOSED': {
                    // called when shrunk to initial bar height
                    console.log("Do something when SEARCH_BAR type panel is
collapsed");

                    break;

```

```

        }
        case 'Q_PANEL_ENTERED_FULLSCREEN': {
            console.log("Do something when panel enters full screen
mode");

            break;
        }
        case 'Q_PANEL_EXITED_FULLSCREEN': {
            console.log("Do something when panel exits full screen
mode");

            break;
        }
        case 'CONTENT_LOADED': {
            console.log("Do something after experience is loaded");
            break;
        }
        case 'ERROR_OCCURRED': {
            console.log("Do something when experience fails to
load");

            break;
        }
    }
}
};
const embeddedGenerativeQnExperience = await
embeddingContext.embedGenerativeQnA(frameOptions, contentOptions);
};
</script>
</head>

<body onload="embedGenerativeQnA()">
    <div id="experience-container"></div>
</body>

</html>

```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded Generative Q&A experience on your website with JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.

- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Optional embedded Generative Q&A experience functionalities

The following optional functionalities are available for the embedded Generative Q&A experience with the embedding SDK.

Invoke Generative Q&A search bar actions

- Set a question — This feature sends a question to the Generative Q&A experience and immediately queries the question.

```
embeddedGenerativeQnExperience.setQuestion('show me monthly revenue');
```

- Close the answer panel (applies to the Generative Q&A search bar option) — This feature closes the answer panel and returns the iframe to the original search bar state.

```
embeddedGenerativeQnExperience.close();
```

For more information, see the [Amazon Quick Sight embedding SDK](#).

Embedding the Amazon Q in Quick Suite Generative Q&A experience for anonymous (unregistered) users

Intended audience: Amazon Quick Suite developers

In the following sections, you can find detailed information about how to set up an embedded Generative Q&A experience for anonymous (unregistered) users.

Topics

- [Step 1: Set up permissions](#)

- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the Generative Q&A experience URL](#)
- [Optional embedded Generative Q&A experience functionalities](#)

Step 1: Set up permissions

In the following section, you can find how to set up permissions for your backend application or web server to embed the Generative Q&A experience. This task requires administrative access to Amazon Identity and Access Management (IAM).

Each user who accesses a Generative Q&A experience assumes a role that gives them Amazon Quick Sight access and permissions. To make this possible, create an IAM role in your Amazon Web Services account. Associate an IAM policy with the role to provide permissions to any user who assumes it. The IAM role needs to provide permissions to retrieve embedding URLs for a specific user pool.

With the help of the wildcard character *, you can grant the permissions to generate a URL for all users in a specific namespace. Or you can grant permissions to generate a URL for a subset of users in specific namespaces. For this, you add `quicksight:GenerateEmbedUrlForAnonymousUser`.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants developers the option to override the static domains that are configured in the **Manage Amazon Quick Sight** menu and instead list up to three domains or subdomains that can access a generated URL. This URL is then embedded in a developer's website. Only the domains that are listed in the parameter can access the embedded Q search bar. Without this condition, developers can list any domain on the internet in the `AllowedDomains` parameter.

To limit the domains that developers can use with this parameter, add an `AllowedEmbeddingDomains` condition to your IAM policy. For more information about the `AllowedDomains` parameter, see [GenerateEmbedUrlForAnonymousUser](#) in the *Amazon Quick Sight API Reference*.

Your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf to load the Generative Q&A experience. The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

For more information regarding trust policies, see [Temporary security credentials in IAM](#) in the *IAM User Guide*

Step 2: Generate the URL with the authentication code attached

In the following section, you can find how to authenticate your user and get the embeddable Q topic URL on your application server.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then the app adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

Java

```
import java.util.List;
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
```

```

import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import
    com.amazonaws.services.quicksight.model.AnonymousUserGenerativeQnAEmbeddingConfiguration;
import
    com.amazonaws.services.quicksight.model.AnonymousUserEmbeddingExperienceConfiguration;
import com.amazonaws.services.quicksight.model.GenerateEmbedUrlForAnonymousUserRequest;
import com.amazonaws.services.quicksight.model.GenerateEmbedUrlForAnonymousUserResult;
import com.amazonaws.services.quicksight.model.SessionTag;

/**
 * Class to call QuickSight Amazon SDK to generate embed url for anonymous user.
 */
public class GenerateEmbedUrlForAnonymousUserExample {

    private final AmazonQuickSight quickSightClient;

    public GenerateEmbedUrlForAnonymousUserExample() {
        quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AWSCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret key here
                    return new BasicAWSCredentials("access-key", "secret-key");
                }

                @Override
                public void refresh() {
                }
            })
            .build();
    }

    public String GenerateEmbedUrlForAnonymousUser(
        final String accountId, // YOUR Amazon ACCOUNT ID
        final String initialTopicId, // Q TOPIC ID TO WHICH THE CONSTRUCTED URL POINTS
        AND EXPERIENCE PREPOPULATES INITIALLY
        final String namespace, // ANONYMOUS EMBEDDING REQUIRES SPECIFYING A VALID
        NAMESPACE FOR WHICH YOU WANT THE EMBEDDING URL
    ) {

```

```

        final List<String> authorizedResourceArns, // Q TOPIC ARN LIST TO EMBED
        final List<String> allowedDomains, // RUNTIME ALLOWED DOMAINS FOR EMBEDDING
        final List<SessionTag> sessionTags // SESSION TAGS USED FOR ROW-LEVEL SECURITY
    ) throws Exception {
        AnonymousUserEmbeddingExperienceConfiguration experienceConfiguration = new
AnonymousUserEmbeddingExperienceConfiguration();
        AnonymousUserGenerativeQnAEmbeddingConfiguration generativeQnAConfiguration =
new AnonymousUserGenerativeQnAEmbeddingConfiguration();
        generativeQnAConfiguration.setInitialTopicId(initialTopicId);
        experienceConfiguration.setGenerativeQnA(generativeQnAConfiguration);

        GenerateEmbedUrlForAnonymousUserRequest generateEmbedUrlForAnonymousUserRequest
= new GenerateEmbedUrlForAnonymousUserRequest()
            .withAwsAccountId(accountId)
            .withNamespace(namespace)
            .withAuthorizedResourceArns(authorizedResourceArns)
            .withExperienceConfiguration(experienceConfiguration)
            .withSessionTags(sessionTags)
            .withSessionLifetimeInMinutes(600L); // OPTIONAL: VALUE CAN BE [15-600].
DEFAULT: 600
            .withAllowedDomains(allowedDomains);

        GenerateEmbedUrlForAnonymousUserResult result =
quickSightClient.generateEmbedUrlForAnonymousUser(generateEmbedUrlForAnonymousUserRequest);

        return result.getEmbedUrl();
    }
}

```

JavaScript

Note

Embed URL generation APIs cannot be called from browsers directly. Refer to the Node.JS example instead.

Python3

```

import json
import boto3
from botocore.exceptions import ClientError

```

```

import time

# Create QuickSight and STS clients
quicksightClient = boto3.client('quicksight', region_name='us-west-2')
sts = boto3.client('sts')

# Function to generate embedded URL for anonymous user
# accountId: YOUR Amazon ACCOUNT ID
# topicId: Topic ID to embed
# quicksightNamespace: VALID NAMESPACE WHERE YOU WANT TO DO NOAUTH EMBEDDING
# authorizedResourceArns: TOPIC ARN LIST TO EMBED
# allowedDomains: RUNTIME ALLOWED DOMAINS FOR EMBEDDING
# sessionTags: SESSION TAGS USED FOR ROW-LEVEL SECURITY
def generateEmbedUrlForAnonymousUser(accountId, quicksightNamespace,
    authorizedResourceArns, allowedDomains, sessionTags):
    try:
        response = quicksightClient.generate_embed_url_for_anonymous_user(
            AwsAccountId = accountId,
            Namespace = quicksightNamespace,
            AuthorizedResourceArns = authorizedResourceArns,
            AllowedDomains = allowedDomains,
            ExperienceConfiguration = {
                'GenerativeQnA': {
                    'InitialTopicId': topicId
                }
            },
            SessionTags = sessionTags,
            SessionLifetimeInMinutes = 600
        )

        return {
            'statusCode': 200,
            'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-Headers": "Content-Type"},
            'body': json.dumps(response),
            'isBase64Encoded': bool('false')
        }
    except ClientError as e:
        print(e)
        return "Error generating embeddedURL: " + str(e)

```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
const AWS = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({
    region: 'us-east-1',
});

quicksightClient.generateEmbedUrlForAnonymousUser({
    'AwsAccountId': '111122223333',
    'Namespace': 'DEFAULT',
    'AuthorizedResourceArns': '["topic-arn-topicId1","topic-arn-topicId2"]',
    'AllowedDomains': allowedDomains,
    'ExperienceConfiguration': {
        'GenerativeQnA': {
            'InitialTopicId': 'U4zJMVZ2n2stZflc80u3iKySEb3BEV6f'
        }
    },
    'SessionTags': '["Key": tag-key-1,"Value": tag-value-1,{"Key": tag-key-1,"Value": tag-value-1}]',
    'SessionLifetimeInMinutes': 15
}, function(err, data) {
    console.log('Errors: ');
    console.log(err);
    console.log('Response: ');
    console.log(data);
});
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded Q search bar. You can use this URL in your website or app to display the Q search bar.

Example

```
using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;

namespace GenerateGenerativeQnAEmbedUrlForAnonymousUser
{
    class Program
    {
        static void Main(string[] args)
        {
            var quicksightClient = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
                SessionToken,
                Amazon.RegionEndpoint.USEast1);
            try
            {
                AnonymousUserGenerativeQnAEmbeddingConfiguration
anonymousUserGenerativeQnAEmbeddingConfiguration
                = new AnonymousUserGenerativeQnAEmbeddingConfiguration
                {
                    InitialTopicId = "U4zJMVZ2n2stZflc80u3iKySEb3BEV6f"
                };
                AnonymousUserEmbeddingExperienceConfiguration
anonymousUserEmbeddingExperienceConfiguration
                = new AnonymousUserEmbeddingExperienceConfiguration
                {
                    GenerativeQnA =
anonymousUserGenerativeQnAEmbeddingConfiguration
                };

                Console.WriteLine(
                    quicksightClient.GenerateEmbedUrlForAnonymousUserAsync(new
GenerateEmbedUrlForAnonymousUserRequest
                {
                    AwsAccountId = "111122223333",
                    Namespace = "DEFAULT",
                    AuthorizedResourceArns ["topic-arn-topicId1","topic-arn-
topicId2"]',
                    AllowedDomains = allowedDomains,
                    ExperienceConfiguration =
anonymousUserEmbeddingExperienceConfiguration,
```

```

        SessionTags = '[{"Key": tag-key-1,"Value": tag-value-1,{"Key":
tag-key-1,"Value": tag-value-1}}]',
        SessionLifetimeInMinutes = 15,
    })).Result.EmbedUrl
    );
} catch (Exception ex) {
    Console.WriteLine(ex.Message);
}
}
}
}
}

```

Amazon CLI

To assume the role, choose one of the following Amazon Security Token Service (Amazon STS) API operations:

- [AssumeRole](#) – Use this operation when you are using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you are using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you are using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForAnonymousUser`.

```

aws sts assume-role \
    --role-arn "arn:aws:iam::111122223333:role/
embedding_quicksight_generative_qna_role" \
    --role-session-name anonymous caller

```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`

- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. For a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_q_search_bar_role/QuickSightEmbeddingAnonymousPolicy`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. *Throttling* is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations. In addition, it keeps each session separate and distinct. If you're using an array of web servers, for example for load balancing, and a session is reconnected to a different server, a new session begins.

To get a signed URL for the dashboard, call `generate-embed-url-for-anonymous-user` from the app server. This returns the embeddable dashboard URL. The following example shows how to generate the URL for an embedded dashboard using a server-side call for users who are making anonymous visits to your web portal or app.

```
aws quicksight generate-embed-url-for-anonymous-user \
--aws-account-id 111122223333 \
--namespace default-or-something-else \
--authorized-resource-arns '["topic-arn-topicId","topic-arn-topicId2"]' \
--allowed-domains '["domain1","domain2"]' \
--experience-configuration 'GenerativeQnA={InitialTopicId="topicId1"}' \
--session-tags '["Key": tag-key-1,"Value": tag-value-1,{"Key": tag-key-1,"Value": tag-value-1}]' \
--session-lifetime-in-minutes 15
```

For more information about using this operation, see [GenerateEmbedUrlForAnonymousUser](#). You can use this and other API operations in your own code.

Step 3: Embed the Generative Q&A experience URL

In the following section, you can find how to embed the Generative Q&A experience URL in your website or application page. You do this with the [Amazon Quick Sight embedding SDK](#) (JavaScript). With the SDK, you can do the following:

- Place the Generative Q&A experience on an HTML page.
- Customize the layout and appearance of the embedded experience to fit your application needs.
- Handle error states with messages that are customized to your application.

To generate the URL that you can embed in your app, call the `GenerateEmbedUrlForAnonymousUser` API operation. This URL is valid for 5 minutes, and the resulting session is valid for up to 10 hours. The API operation provides the URL with an `auth_code` value that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-anonymous-user`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.{
  "Status": "200",
  "EmbedUrl": "https://quicksightdomain/embedding/12345/q/search...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed the Generative Q&A experience in your webpage with the [Amazon Quick Sight embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes.

Make sure that the domain to host the Generative Q&A experience is on the *allow list*, the list of approved domains for your Amazon Quick Sight subscription. This requirement protects your data by keeping unapproved domains from hosting embedded Generative Q&A experiences. For more information about adding domains for an embedded Generative Q&A experience, see [Managing domains](#).

You can use the Amazon Quick Sight Embedding SDK to customize the layout and appearance of the embedded Generative Q&A experience to fit your application. Use the `panelType` property to

configure the landing state of the Generative Q&A experience when it renders in your application. Set the `panelType` property to `'FULL'` to render the full Generative Q&A experience panel. This panel resembles the experience that Amazon Quick Sight users have in the Amazon Quick Sight console. The frame height of the panel is not changed based on user interaction and respects the value that you set in the `frameOptions.height` property. The image below shows the Generative Q&A experience panel that renders when you set the `panelType` value to `'FULL'`.

Set the `panelType` property to `'SEARCH_BAR'` to render the Generative Q&A experience as a search bar. This search bar resembles the way that the Q Search Bar renders when it is embedded into an application. The Generative Q&A search bar expands to a larger panel that displays topic selection options, the question suggestion list, the answer panel or the pinboard.

The default minimum height of the Generative Q&A search bar is rendered when the embedded asset loads. It is recommended that you set the `frameOptions.height` value to `"38px"` to optimize the search bar experience. Use the `focusedHeight` property to set the optimal size of the topic selection dropdown and the question suggestion list. Use the `expandedHeight` property to set the optimal size of the answer panel and pinboard. If you choose the `'SEARCH_BAR'` option, it is recommended that you style the parent container with `position: absolute` to avoid unwanted content shifting in your application. The image below shows the Generative Q&A experience search bar that renders when you set the `panelType` value to `'SEARCH_BAR'`.

After you configure the `panelType` property, use the Amazon Quick Sight embedding SDK to customize the following properties of the Generative Q&A experience.

- The title of the Generative Q&A panel (Applies only to the `panelType: FULL` option).
- The search bar's placeholder text.
- Whether topic selection is allowed.
- Whether topic names are shown or hidden.
- Whether the Amazon Q icon is shown or hidden (Applies only to the `panelType: FULL` option).
- Whether the pinboard is shown or hidden.
- Whether users can maximize the Generative Q&A panel to fullscreen.
- The theme of the Generative Q&A panel. A custom theme ARN can be passed in the SDK to change the appearance of the frame's content. Amazon Quick Sight starter themes are not supported for embedded Generative BI panels. To use a Amazon Quick Sight starter theme, save it as a custom theme in Amazon Quick Sight.

When you use the Amazon Quick Sight Embedding SDK, the Generative Q&A experience on your page is dynamically resized based on the state. With the Amazon Quick Sight Embedding SDK, you can also control parameters within the Generative Q&A experience and receive callbacks in terms of page load completion, state changes, and errors.

The following example shows how to use the generated URL. This code is generated on your app server.

SDK 2.0

```
<!DOCTYPE html>
<html>
  <head>
    <title>Generative Q&A Embedding Example</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.7.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      const embedGenerativeQnA = async() => {
        const {createEmbeddingContext} = QuickSightEmbedding;

        const embeddingContext = await createEmbeddingContext({
          onChange: (changeEvent, metadata) => {
            console.log('Context received a change', changeEvent,
metadata);
          },
        });

        const frameOptions = {
          url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
          container: '#experience-container',
          height: "700px",
          width: "1000px",
          onChange: (changeEvent, metadata) => {
            switch (changeEvent.eventName) {
              case 'FRAME_MOUNTED': {
                console.log("Do something when the experience frame is
mounted.");
                break;
              }
              case 'FRAME_LOADED': {
                console.log("Do something when the experience frame is
loaded.");
            }
```

```

        break;
    }
    },
};

const contentOptions = {
    // Optional panel settings. Default behavior is equivalent to
    {panelType: 'FULL'}
    panelOptions: {
        panelType: 'FULL',
        title: 'custom title', // Optional
        showQIcon: false, // Optional, Default: true
    },
    // Use SEARCH_BAR panel type for the landing state to be similar to
    embedQSearchBar
    // with generative capability enabled topics
    /*
    panelOptions: {
        panelType: 'SEARCH_BAR',
        focusedHeight: '250px',
        expandedHeight: '500px',
    },
    */
    showTopicName: false, // Optional, Default: true
    showPinboard: false, // Optional, Default: true
    allowTopicSelection: false, // Optional, Default: true
    allowFullscreen: false, // Optional, Default: true
    searchPlaceholderText: "custom search placeholder", // Optional
    themeOptions: { // Optional
        themeArn: 'arn:aws:quicksight:<Region>:<Amazon-Account-
ID>:theme/<Theme-ID>'
    }
    onMessage: async (messageEvent, experienceMetadata) => {
        switch (messageEvent.eventName) {
            case 'Q_SEARCH_OPENED': {
                // called when pinboard is shown / visuals are rendered
                console.log("Do something when SEARCH_BAR type panel is
expanded");

                break;
            }
            case 'Q_SEARCH_FOCUSED': {
                // called when question suggestions or topic selection
                dropdown are shown

```

```

        console.log("Do something when SEARCH_BAR type panel is
focused");
        break;
    }
    case 'Q_SEARCH_CLOSED': {
        // called when shrunk to initial bar height
        console.log("Do something when SEARCH_BAR type panel is
collapsed");
        break;
    }
    case 'Q_PANEL_ENTERED_FULLSCREEN': {
        console.log("Do something when panel enters full screen
mode");
        break;
    }
    case 'Q_PANEL_EXITED_FULLSCREEN': {
        console.log("Do something when panel exits full screen
mode");
        break;
    }
    case 'CONTENT_LOADED': {
        console.log("Do something after experience is loaded");
        break;
    }
    case 'ERROR_OCCURRED': {
        console.log("Do something when experience fails to
load");
        break;
    }
}
}
};
const embeddedGenerativeQnExperience = await
embeddingContext.embedGenerativeQnA(frameOptions, contentOptions);
};
</script>
</head>

<body onload="embedGenerativeQnA()">
    <div id="experience-container"></div>
</body>

</html>

```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded Generative Q&A experience on your website with JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Optional embedded Generative Q&A experience functionalities

The following optional functionalities are available for the embedded Generative Q&A experience with the embedding SDK.

Invoke Generative Q&A search bar actions

- Set a question — This feature sends a question to the Generative Q&A experience and immediately queries the question.

```
embeddedGenerativeQnExperience.setQuestion('show me monthly revenue');
```

- Close the answer panel (applies to the Generative Q&A search bar option) — This feature closes the answer panel and returns the iframe to the original search bar state.

```
embeddedGenerativeQnExperience.close();
```

For more information, see the [Amazon Quick Sight embedding SDK](#).

Embedding the Amazon Quick Sight Q search bar (Classic)

Intended audience: Amazon Quick Suite developers

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

Use the following topics to learn about embedding the Amazon Quick Sight Q search bar with the Amazon Quick Sight APIs.

Topics

- [Embedding the Amazon Quick Sight Q search bar for registered users](#)
- [Embedding the Amazon Quick Sight Q search bar for anonymous \(unregistered\) users](#)

Embedding the Amazon Quick Sight Q search bar for registered users

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

In the following sections, you can find detailed information about how to set up an embedded Amazon Quick Sight Q search bar for registered users of Amazon Quick Sight.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the Q search bar URL](#)
- [Optional Amazon Quick Sight Q search bar embedding functionalities](#)

Step 1: Set up permissions

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

In the following section, you can find how to set up permissions for your backend application or web server to embed the Q search bar. This task requires administrative access to Amazon Identity and Access Management (IAM).

Each user who accesses a dashboard assumes a role that gives them Amazon Quick Sight access and permissions to the dashboard. To make this possible, create an IAM role in your Amazon Web Services account. Associate an IAM policy with the role to provide permissions to any user who assumes it. The IAM role needs to provide permissions to retrieve embedding URLs for a specific user pool.

With the help of the wildcard character *, you can grant the permissions to generate a URL for all users in a specific namespace. Or you can grant permissions to generate a URL for a subset of users in specific namespaces. For this, you add `quicksight:GenerateEmbedUrlForRegisteredUser`.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForRegisteredUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants developers the option to override the static domains that are configured in the **Manage Amazon Quick Sight** menu and instead list up to three domains or subdomains that can access a generated URL. This URL is then embedded in a developer's website. Only the domains that are listed in the parameter can access

the embedded Q search bar. Without this condition, developers can list any domain on the internet in the `AllowedDomains` parameter.

To limit the domains that developers can use with this parameter, add an `AllowedEmbeddingDomains` condition to your IAM policy. For more information about the `AllowedDomains` parameter, see [GenerateEmbedUrlForRegisteredUser](#) in the *Amazon Quick Sight API Reference*.

The following sample policy provides these permissions.

Also, if you're creating first-time users who will be Amazon Quick Sight readers, make sure to add the `quicksight:RegisterUser` permission in the policy.

The following sample policy provides permission to retrieve an embedding URL for first-time users who are to be Amazon Quick Sight readers.

Finally, your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf and provision the user in Amazon Quick Sight.

The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

```
}  
]  
}
```

For more information regarding trust policies for OpenID Connect or Security Assertion Markup Language (SAML) authentication, see the following sections of the *IAM User Guide*:

- [Creating a role for web identity or OpenID Connect federation \(console\)](#)
- [Creating a role for SAML 2.0 federation \(console\)](#)

Step 2: Generate the URL with the authentication code attached

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

In the following section, you can find how to authenticate your user and get the embeddable Q topic URL on your application server. If you plan to embed the Q bar for IAM or Amazon Quick Sight identity types, share the Q topic with the users.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then the app adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

Performing the described steps ensures that each viewer of the Q topic is uniquely provisioned in Amazon Quick Sight. It also enforces per-user settings, such as the row-level security and dynamic defaults for parameters.

The following examples perform the IAM authentication on the user's behalf. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;
```

```

import com.amazonaws.auth.BasicAWSCredentials;
    import com.amazonaws.auth.AWSCredentialsProvider;
    import com.amazonaws.regions.Regions;
    import com.amazonaws.services.quicksight.AmazonQuickSight;
    import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import
    com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserRequest;
import com.amazonaws.services.quicksight.model.GenerateEmbedUrlForRegisteredUserResult;
import
    com.amazonaws.services.quicksight.model.RegisteredUserEmbeddingExperienceConfiguration;
import
    com.amazonaws.services.quicksight.model.RegisteredUserQSearchBarEmbeddingConfiguration;

    /**
    * Class to call QuickSight Amazon SDK to get url for embedding the Q search bar.
    */
public class RegisteredUserQSearchBarEmbeddingConfiguration {

    private final AmazonQuickSight quickSightClient;

    public RegisteredUserQSearchBarEmbeddingConfiguration() {
        this.quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AWSCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret key here
                    return new BasicAWSCredentials("access-key", "secret-
key");
                }

                @Override
                public void refresh() {
                }
            })
            .build();
    }

    public String getQuicksightEmbedUrl(
        final String accountId, // Amazon Account ID
        final String topicId, // Topic ID to embed
        final List<String> allowedDomains, // Runtime allowed domain for embedding

```

```

        final String userArn // Registered user arn to use for embedding. Refer to
        Get Embed Url section in developer portal to find how to get user arn for a QuickSight
        user.

        ) throws Exception {
            final RegisteredUserEmbeddingExperienceConfiguration experienceConfiguration =
            new RegisteredUserEmbeddingExperienceConfiguration()
                .withQSearchBar(new
            RegisteredUserQSearchBarEmbeddingConfiguration().withInitialTopicId(topicId));
            final GenerateEmbedUrlForRegisteredUserRequest
            generateEmbedUrlForRegisteredUserRequest = new
            GenerateEmbedUrlForRegisteredUserRequest();
            generateEmbedUrlForRegisteredUserRequest.setAwsAccountId(accountId);
            generateEmbedUrlForRegisteredUserRequest.setUserArn(userArn);
            generateEmbedUrlForRegisteredUserRequest.setAllowedDomains(allowedDomains);

            generateEmbedUrlForRegisteredUserRequest.setExperienceConfiguration(QSearchBar);

            final GenerateEmbedUrlForRegisteredUserResult
            generateEmbedUrlForRegisteredUserResult =
            quickSightClient.generateEmbedUrlForRegisteredUser(generateEmbedUrlForRegisteredUserRequest);

            return generateEmbedUrlForRegisteredUserResult.getEmbedUrl();
        }
    }
}

```

JavaScript

```

global.fetch = require('node-fetch');
const Amazon = require('aws-sdk');

function generateEmbedUrlForRegisteredUser(
    accountId,
    topicId, // Topic ID to embed
    openIdToken, // Cognito-based token
    userArn, // registered user arn
    roleArn, // IAM user role to use for embedding
    sessionName, // Session name for the roleArn assume role
    allowedDomains, // Runtime allowed domain for embedding
    getEmbedUrlCallback, // GetEmbedUrl success callback method
    errorCallback // GetEmbedUrl error callback method
) {
    const stsClient = new AWS.STS();
    let stsParams = {

```

```

    RoleSessionName: sessionName,
    WebIdentityToken: openIdToken,
    RoleArn: roleArn
  }

  stsClient.assumeRoleWithWebIdentity(stsParams, function(err, data) {
    if (err) {
      console.log('Error assuming role');
      console.log(err, err.stack);
      errorCallback(err);
    } else {
      const getQSearchBarParams = {
        "AwsAccountId": accountId,
        "ExperienceConfiguration": {
          "QSearchBar": {
            "InitialTopicId": topicId
          }
        },
        "UserArn": userArn,
        "AllowedDomains": allowedDomains,
        "SessionLifetimeInMinutes": 600
      };

      const quicksightGetQSearchBar = new AWS.QuickSight({
        region: process.env.AWS_REGION,
        credentials: {
          accessKeyId: data.Credentials.AccessKeyId,
          secretAccessKey: data.Credentials.SecretAccessKey,
          sessionToken: data.Credentials.SessionToken,
          expiration: data.Credentials.Expiration
        }
      });

      quicksightGetQSearchBar.generateEmbedUrlForRegisteredUser(getQSearchBarParams,
        function(err, data) {
          if (err) {
            console.log(err, err.stack);
            errorCallback(err);
          } else {
            const result = {
              "statusCode": 200,
              "headers": {

```

```

        "Access-Control-Allow-Origin": "*", // Use your website
        domain to secure access to GetEmbedUrl API
        "Access-Control-Allow-Headers": "Content-Type"
    },
    "body": JSON.stringify(data),
    "isBase64Encoded": false
}
    getEmbedUrlCallback(result);
}
});
}
});
}
}
}

```

Python3

```

import json
import boto3
from botocore.exceptions import ClientError

sts = boto3.client('sts')

# Function to generate embedded URL
# accountId: Amazon account ID
# topicId: Topic ID to embed
# userArn: arn of registered user
# allowedDomains: Runtime allowed domain for embedding
# roleArn: IAM user role to use for embedding
# sessionName: session name for the roleArn assume role
def getEmbeddingURL(accountId, topicId, userArn, allowedDomains, roleArn, sessionName):
    try:
        assumedRole = sts.assume_role(
            RoleArn = roleArn,
            RoleSessionName = sessionName,
        )
    except ClientError as e:
        return "Error assuming role: " + str(e)
    else:
        assumedRoleSession = boto3.Session(
            aws_access_key_id = assumedRole['Credentials']['AccessKeyId'],
            aws_secret_access_key = assumedRole['Credentials']['SecretAccessKey'],
            aws_session_token = assumedRole['Credentials']['SessionToken'],
        )

```

```

    try:
        quicksightClient = assumedRoleSession.client('quicksight', region_name='us-
west-2')
        response = quicksightClient.generate_embed_url_for_registered_user(
            AwsAccountId=accountId,
            ExperienceConfiguration = {
                "QSearchBar": {
                    "InitialTopicId": topicId
                }
            },
            UserArn = userArn,
            AllowedDomains = allowedDomains,
            SessionLifetimeInMinutes = 600
        )

        return {
            'statusCode': 200,
            'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-
Headers": "Content-Type"},
            'body': json.dumps(response),
            'isBase64Encoded': bool('false')
        }
    except ClientError as e:
        return "Error generating embedding url: " + str(e)

```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```

const Amazon = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({
    apiConfig: require('./quicksight-2018-04-01.min.json'),
    region: 'us-east-1',
});

quicksightClient.generateEmbedUrlForRegisteredUser({
    'AwsAccountId': '111122223333',

```

```

    'ExperienceConfiguration': {
      'QSearchBar': {
        'InitialTopicId': 'U4zJMVZ2n2stZflc80u3iKySEb3BEV6f'
      }
    },
    'UserArn': 'REGISTERED_USER_ARN',
    'AllowedDomains': allowedDomains,
    'SessionLifetimeInMinutes': 100
  }, function(err, data) {
    console.log('Errors: ');
    console.log(err);
    console.log('Response: ');
    console.log(data);
  });

```

Example

```

//The URL returned is over 900 characters. For this example, we've shortened the string
//for
//readability and added ellipsis to indicate that it's incomplete.
{
  Status: 200,
  EmbedUrl: "https://quicksightdomain/embed/12345/dashboards/67890/sheets/12345/
visuals/67890...",
  RequestId: '7bee030e-f191-45c4-97fe-d9faf0e03713'
}

```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded Q search bar. You can use this URL in your website or app to display the Q search bar.

Example

```

using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;

namespace GenerateDashboardEmbedUrlForRegisteredUser
{
    class Program
    {

```



```

static void Main(string[] args)
{
    var quicksightClient = new AmazonQuickSightClient(
        AccessKey,
        SecretAccessKey,
        SessionToken,
        Amazon.RegionEndpoint.USEast1);

    try
    {
        RegisteredUserQSearchBarEmbeddingConfiguration
registeredUserQSearchBarEmbeddingConfiguration
        = new RegisteredUserQSearchBarEmbeddingConfiguration
        {
            InitialTopicId = "U4zJMVZ2n2stZflc80u3iKySEb3BEV6f"
        };

        RegisteredUserEmbeddingExperienceConfiguration
registeredUserEmbeddingExperienceConfiguration
        = new RegisteredUserEmbeddingExperienceConfiguration
        {
            QSearchBar = registeredUserQSearchBarEmbeddingConfiguration
        };

        Console.WriteLine(
            quicksightClient.GenerateEmbedUrlForRegisteredUserAsync(new
GenerateEmbedUrlForRegisteredUserRequest
            {
                AwsAccountId = "111122223333",
                ExperienceConfiguration =
registeredUserEmbeddingExperienceConfiguration,
                UserArn = "REGISTERED_USER_ARN",
                AllowedDomains = allowedDomains,
                SessionLifetimeInMinutes = 100
            }).Result.EmbedUrl
        );
    } catch (Exception ex) {
        Console.WriteLine(ex.Message);
    }
}
}
}

```

Amazon CLI

To assume the role, choose one of the following Amazon Security Token Service (Amazon STS) API operations:

- [AssumeRole](#) – Use this operation when you are using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you are using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you are using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForRegisteredUser`. If you are taking a just-in-time approach to add users when they use a topic in the Q search bar, the role also needs permissions enabled for `quicksight:RegisterUser`.

```
aws sts assume-role \  
    --role-arn "arn:aws-cn:iam::111122223333:role/  
embedding_quicksight_q_search_bar_role" \  
    --role-session-name john.doe@example.com
```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. For a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"
```

```
export AWS_SECRET_ACCESS_KEY = "secret_key_from_assume_role"
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_q_search_bar_role/john.doe@example.com`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. *Throttling* is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations.

The role session ID also becomes the user name in Amazon Quick Sight. You can use this pattern to provision your users in Amazon Quick Sight ahead of time, or to provision them the first time that they access the Q search bar.

The following example shows the CLI command that you can use to provision a user. For more information about [RegisterUser](#), [DescribeUser](#), and other Amazon Quick Sight API operations, see the [Amazon Quick Sight API reference](#).

```
aws quicksight register-user \
  --aws-account-id 111122223333 \
  --namespace default \
  --identity-type IAM \
  --iam-arn "arn:aws-cn:iam::111122223333:role/
embedding_quicksight_q_search_bar_role" \
  --user-role READER \
  --user-name jhnd \
  --session-name "john.doe@example.com" \
  --email john.doe@example.com \
  --region us-east-1 \
  --custom-permissions-name TeamA1
```

If the user is authenticated through Microsoft AD, you don't need to use `RegisterUser` to set them up. Instead, they should be automatically subscribed the first time that they access Amazon Quick Sight. For Microsoft AD users, you can use `DescribeUser` to get the user Amazon Resource Name (ARN).

The first time a user accesses Amazon Quick Sight, you can also add this user to the group that the dashboard is shared with. The following example shows the CLI command to add a user to a group.

```
aws quicksight create-group-membership \
  --aws-account-id=111122223333 \
```

```
--namespace=default \  
--group-name=financeusers \  
--member-name="embedding_quicksight_q_search_bar_role/john.doe@example.com"
```

You now have a user of your app who is also a user of Amazon Quick Sight, and who has access to the dashboard.

Finally, to get a signed URL for the dashboard, call `generate-embed-url-for-registered-user` from the app server. This returns the embeddable dashboard URL. The following example shows how to generate the URL for an embedded dashboard using a server-side call for users authenticated through Amazon Managed Microsoft AD or single sign-on (IAM Identity Center).

```
aws quicksight generate-embed-url-for-registered-user \  
--aws-account-id 111122223333 \  
--session-lifetime-in-minutes 600 \  
--user-arn arn:aws-cn:quicksight:us-east-1:111122223333:user/default/  
embedding_quicksight_q_search_bar_role/embeddingsession \  
--allowed-domains '["domain1", "domain2"]' \  
--experience-configuration QSearchBar={InitialTopicId=U4zJMVZ2n2stZflc80u3iKySEb3BEV6f}
```

For more information about using this operation, see [GenerateEmbedUrlForRegisteredUser](#). You can use this and other API operations in your own code.

Step 3: Embed the Q search bar URL

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

In the following section, you can find how to embed the Q search bar URL from step 3 in your website or application page. You do this with the [Amazon Quick Sight embedding SDK](#) (JavaScript). With the SDK, you can do the following:

- Place the Q search bar on an HTML page.
- Pass parameters into the Q search bar.

- Handle error states with messages that are customized to your application.

To generate the URL that you can embed in your app, call the `GenerateEmbedUrlForRegisteredUser` API operation. This URL is valid for 5 minutes, and the resulting session is valid for up to 10 hours. The API operation provides the URL with an `auth_code` value that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-registered-user`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.
{
  "Status": "200",
  "EmbedUrl": "https://quicksightdomain/embedding/12345/q/search...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed the Q search bar in your webpage by using the [Amazon Quick Sight embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes.

To do this, make sure that the domain to host the embedded Q search bar is on the *allow list*, the list of approved domains for your Amazon Quick Sight subscription. This requirement protects your data by keeping unapproved domains from hosting embedded dashboards. For more information about adding domains for an embedded Q search bar, see [Managing domains and embedding](#).

When you use the Amazon Quick Sight Embedding SDK, the Q search bar on your page is dynamically resized based on the state. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the Q search bar and receive callbacks in terms of page load completion and errors.

The following example shows how to use the generated URL. This code is generated on your app server.

SDK 2.0

```
<!DOCTYPE html>
```

```

<html>

  <head>
    <title>Q Search Bar Embedding Example</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.0.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      const embedQSearchBar = async() => {
        const {
          createEmbeddingContext,
        } = QuickSightEmbedding;

        const embeddingContext = await createEmbeddingContext({
          onChange: (changeEvent, metadata) => {
            console.log('Context received a change', changeEvent,
metadata);
          },
        });

        const frameOptions = {
          url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
          container: '#experience-container',
          height: "700px",
          width: "1000px",
          onChange: (changeEvent, metadata) => {
            switch (changeEvent.eventName) {
              case 'FRAME_MOUNTED': {
                console.log("Do something when the experience frame is
mounted.");
                break;
              }
              case 'FRAME_LOADED': {
                console.log("Do something when the experience frame is
loaded.");
                break;
              }
            }
          },
        };

        const contentOptions = {
          hideTopicName: false,
          theme: '<YOUR_THEME_ID>',

```

```

        allowTopicSelection: true,
        onMessage: async (messageEvent, experienceMetadata) => {
            switch (messageEvent.eventName) {
                case 'Q_SEARCH_OPENED': {
                    console.log("Do something when Q Search content
expanded");

                    break;
                }
                case 'Q_SEARCH_CLOSED': {
                    console.log("Do something when Q Search content
collapsed");

                    break;
                }
                case 'Q_SEARCH_SIZE_CHANGED': {
                    console.log("Do something when Q Search size changed");
                    break;
                }
                case 'CONTENT_LOADED': {
                    console.log("Do something when the Q Search is
loaded.");

                    break;
                }
                case 'ERROR_OCCURRED': {
                    console.log("Do something when the Q Search fails
loading.");

                    break;
                }
            }
        }
    };
    const embeddedDashboardExperience = await
embeddingContext.embedQSearchBar(frameOptions, contentOptions);
    };
</script>
</head>

<body onload="embedQSearchBar()">
    <div id="experience-container"></div>
</body>

</html>

```

SDK 1.0

```
<!DOCTYPE html>
<html>

  <head>
    <title>QuickSight Q Search Bar Embedding</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@1.18.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      var session

      function onError(payload) {
        console.log("Do something when the session fails loading");
      }

      function onOpen() {
        console.log("Do something when the Q search bar opens");
      }

      function onClose() {
        console.log("Do something when the Q search bar closes");
      }

      function embedQSearchBar() {
        var containerDiv = document.getElementById("embeddingContainer");
        var options = {
          url: "https://us-east-1.quicksight.aws.amazon.com/sn/dashboards/
dashboardId?isauthcode=true&identityprovider=quicksight&code=authcode", // replace this
dummy url with the one generated via embedding API
          container: containerDiv,
          width: "1000px",
          locale: "en-US",
          qSearchBarOptions: {
            expandCallback: onOpen,
            collapseCallback: onClose,
            iconDisabled: false,
            topicNameDisabled: false,
            themeId: 'bdb844d0-0fe9-4d9d-b520-0fe602d93639',
            allowTopicSelection: true
          }
        };
        session = QuickSightEmbedding.embedQSearchBar(options);
        session.on("error", onError);
      }
    </script>
  </head>
</html>
```



```
    }

    function onCountryChange(obj) {
        session.setParameters({country: obj.value});
    }
</script>
</head>

<body onload="embedQSearchBar()">
    <div id="embeddingContainer"></div>
</body>

</html>
```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded dashboard on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Optional Amazon Quick Sight Q search bar embedding functionalities

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

The following optional functionalities are available for the embedded Q search bar using the embedding SDK.

Invoke Q search bar actions

The following options are only supported for Q search bar embedding.

- Set a Q search bar question — This feature sends a question to the Q search bar and immediately queries the question. It also automatically opens the Q popover.

```
qBar.setQBarQuestion('show me monthly revenue');
```

- Close the Q popover — This feature closes the Q popover and returns the iframe to the original Q search bar size.

```
qBar.closeQPopover();
```

For more information, see the [Amazon Quick Sight embedding SDK](#).

Embedding the Amazon Quick Sight Q search bar for anonymous (unregistered) users

Intended audience: Amazon Quick Suite developers

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

In the following sections, you can find detailed information about how to set up an embedded Amazon Quick Sight Q search bar for anonymous (unregistered) users.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Generate the URL with the authentication code attached](#)
- [Step 3: Embed the Q search bar URL](#)

- [Optional Amazon Quick Sight Q search bar embedding functionalities](#)

Step 1: Set up permissions

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

In the following section, you can find how to set up permissions for your backend application or web server to embed the Q search bar. This task requires administrative access to Amazon Identity and Access Management (IAM).

Each user who accesses a Q search bar assumes a role that gives them Amazon Quick Sight access and permissions to the Q search bar. To make this possible, create an IAM role in your Amazon Web Services account. Associate an IAM policy with the role to provide permissions to any user who assumes it. The IAM role needs to provide permissions to retrieve embedding URLs for a specific user pool.

With the help of the wildcard character *, you can grant the permissions to generate a URL for all users in a specific namespace. Or you can grant permissions to generate a URL for a subset of users in specific namespaces. For this, you add `quicksight:GenerateEmbedUrlForAnonymousUser`.

You can create a condition in your IAM policy that limits the domains that developers can list in the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` API operation. The `AllowedDomains` parameter is an optional parameter. It grants developers the option to override the static domains that are configured in the **Manage Amazon Quick Sight** menu and instead list up to three domains or subdomains that can access a generated URL. This URL is then embedded in a developer's website. Only the domains that are listed in the parameter can access the embedded Q search bar. Without this condition, developers can list any domain on the internet in the `AllowedDomains` parameter.

To limit the domains that developers can use with this parameter, add an `AllowedEmbeddingDomains` condition to your IAM policy. For more information about the

AllowedDomains parameter, see [GenerateEmbedUrlForAnonymousUser](#) in the *Amazon Quick Sight API Reference*.

Your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf to open the Q search bar. The following example shows a sample trust policy.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowLambdaFunctionsToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "lambda.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    },
    {
      "Sid": "AllowEC2InstancesToAssumeThisRole",
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}
```

For more information regarding trust policies, see [Temporary security credentials in IAM](#) in the *IAM User Guide*

Step 2: Generate the URL with the authentication code attached

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

In the following section, you can find how to authenticate your user and get the embeddable Q topic URL on your application server.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then the app adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

For more information, see [AnonymousUserQSearchBarEmbeddingConfiguration](#).

Java

```
import java.util.List;
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import
com.amazonaws.services.quicksight.model.AnonymousUserQSearchBarEmbeddingConfiguration;
import
com.amazonaws.services.quicksight.model.AnonymousUserEmbeddingExperienceConfiguration;
import
com.amazonaws.services.quicksight.model.GenerateEmbedUrlForAnonymousUserRequest;
import
com.amazonaws.services.quicksight.model.GenerateEmbedUrlForAnonymousUserResult;
import com.amazonaws.services.quicksight.model.SessionTag;

/**
 * Class to call QuickSight Amazon SDK to generate embed url for anonymous user.
```

```

    */
    public class GenerateEmbedUrlForAnonymousUserExample {

        private final AmazonQuickSight quickSightClient;

        public GenerateEmbedUrlForAnonymousUserExample() {
            quickSightClient = AmazonQuickSightClientBuilder
                .standard()
                .withRegion(Regions.US_EAST_1.getName())
                .withCredentials(new AWSCredentialsProvider() {
                    @Override
                    public AWSCredentials getCredentials() {
                        // provide actual IAM access key and secret key here
                        return new BasicAWSCredentials("access-key", "secret-
key");
                    }

                    @Override
                    public void refresh() {
                    }
                })
                .build();
        }

        public String GenerateEmbedUrlForAnonymousUser(
            final String accountId, // YOUR Amazon ACCOUNT ID
            final String initialTopicId, // Q TOPIC ID TO WHICH THE CONSTRUCTED URL
POINTS AND SEARCHBAR PREPOPULATES INITIALLY
            final String namespace, // ANONYMOUS EMBEDDING REQUIRES SPECIFYING A
VALID NAMESPACE FOR WHICH YOU WANT THE EMBEDDING URL
            final List<String> authorizedResourceArns, // Q SEARCHBAR TOPIC ARN
LIST TO EMBED
            final List<String> allowedDomains, // RUNTIME ALLOWED DOMAINS FOR
EMBEDDING
            final List<SessionTag> sessionTags // SESSION TAGS USED FOR ROW-LEVEL
SECURITY
        ) throws Exception {
            AnonymousUserEmbeddingExperienceConfiguration experienceConfiguration =
new AnonymousUserEmbeddingExperienceConfiguration();
            AnonymousUserQSearchBarEmbeddingConfiguration qSearchBarConfiguration =
new AnonymousUserQSearchBarEmbeddingConfiguration();
            qSearchBarConfiguration.setInitialTopicId(initialTopicId);
            experienceConfiguration.setQSearchBar(qSearchBarConfiguration);

```

```

        GenerateEmbedUrlForAnonymousUserRequest
generateEmbedUrlForAnonymousUserRequest = new
GenerateEmbedUrlForAnonymousUserRequest()
        .withAwsAccountId(accountId)
        .withNamespace(namespace)
        .withAuthorizedResourceArns(authorizedResourceArns)
        .withExperienceConfiguration(experienceConfiguration)
        .withSessionTags(sessionTags)
        .withSessionLifetimeInMinutes(600L); // OPTIONAL: VALUE CAN BE
[15-600]. DEFAULT: 600
        .withAllowedDomains(allowedDomains);

        GenerateEmbedUrlForAnonymousUserResult qSearchBarEmbedUrl =
quickSightClient.generateEmbedUrlForAnonymousUser(generateEmbedUrlForAnonymousUserRequest);

        return qSearchBarEmbedUrl.getEmbedUrl();
    }
}

```

JavaScript

```

global.fetch = require('node-fetch');
const Amazon = require('aws-sdk');

function generateEmbedUrlForAnonymousUser(
    accountId, // YOUR Amazon ACCOUNT ID
    initialTopicId, // Q TOPIC ID TO WHICH THE CONSTRUCTED URL POINTS
    quicksightNamespace, // VALID NAMESPACE WHERE YOU WANT TO DO NOAUTH EMBEDDING
    authorizedResourceArns, // Q SEARCHBAR TOPIC ARN LIST TO EMBED
    allowedDomains, // RUNTIME ALLOWED DOMAINS FOR EMBEDDING
    sessionTags, // SESSION TAGS USED FOR ROW-LEVEL SECURITY
    generateEmbedUrlForAnonymousUserCallback, // SUCCESS CALLBACK METHOD
    errorCallback // ERROR CALLBACK METHOD
) {
    const experienceConfiguration = {
        "QSearchBar": {
            "InitialTopicId": initialTopicId // TOPIC ID CAN BE FOUND IN THE URL ON THE
TOPIC AUTHOR PAGE
        }
    };
};

```

```

const generateEmbedUrlForAnonymousUserParams = {
  "AwsAccountId": accountId,
  "Namespace": quicksightNamespace,
  "AuthorizedResourceArns": authorizedResourceArns,
  "AllowedDomains": allowedDomains,
  "ExperienceConfiguration": experienceConfiguration,
  "SessionTags": sessionTags,
  "SessionLifetimeInMinutes": 600
};

const quicksightClient = new AWS.QuickSight({
  region: process.env.AWS_REGION,
  credentials: {
    accessKeyId: AccessKeyId,
    secretAccessKey: SecretAccessKey,
    sessionToken: SessionToken,
    expiration: Expiration
  }
});

quicksightClient.generateEmbedUrlForAnonymousUser(generateEmbedUrlForAnonymousUserParams,
function(err, data) {
  if (err) {
    console.log(err, err.stack);
    errorCallback(err);
  } else {
    const result = {
      "statusCode": 200,
      "headers": {
        "Access-Control-Allow-Origin": "*", // USE YOUR WEBSITE DOMAIN TO
SECURE ACCESS TO THIS API
        "Access-Control-Allow-Headers": "Content-Type"
      },
      "body": JSON.stringify(data),
      "isBase64Encoded": false
    }
    generateEmbedUrlForAnonymousUserCallback(result);
  }
});
}

```


Python3

```
import json
import boto3
from botocore.exceptions import ClientError
import time

# Create QuickSight and STS clients
quicksightClient = boto3.client('quicksight', region_name='us-west-2')
sts = boto3.client('sts')

# Function to generate embedded URL for anonymous user
# accountId: YOUR AWS ACCOUNT ID
# quicksightNamespace: VALID NAMESPACE WHERE YOU WANT TO DO NOAUTH EMBEDDING
# authorizedResourceArns: TOPIC ARN LIST TO EMBED
# allowedDomains: RUNTIME ALLOWED DOMAINS FOR EMBEDDING
# experienceConfiguration: configuration which specifies the TOPIC ID to point URL to
# sessionTags: SESSION TAGS USED FOR ROW-LEVEL SECURITY
def generateEmbedUrlForAnonymousUser(accountId, quicksightNamespace,
    authorizedResourceArns, allowedDomains, experienceConfiguration, sessionTags):
    try:
        response = quicksightClient.generate_embed_url_for_anonymous_user(
            AwsAccountId = accountId,
            Namespace = quicksightNamespace,
            AuthorizedResourceArns = authorizedResourceArns,
            AllowedDomains = allowedDomains,
            ExperienceConfiguration = experienceConfiguration,
            SessionTags = sessionTags,
            SessionLifetimeInMinutes = 600
        )

        return {
            'statusCode': 200,
            'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-Headers": "Content-Type"},
            'body': json.dumps(response),
            'isBase64Encoded': bool('false')
        }
    except ClientError as e:
        print(e)
        return "Error generating embeddedURL: " + str(e)
```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to generate the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
const Amazon = require('aws-sdk');
const https = require('https');

var quicksightClient = new AWS.Service({
  apiConfig: require('./quicksight-2018-04-01.min.json'),
  region: 'us-east-1',
});

quicksightClient.generateEmbedUrlForAnonymousUser({
  'AwsAccountId': '111122223333',
  'Namespace': 'DEFAULT'
  'AuthorizedResourceArns': '["topic-arn-topicId1","topic-arn-topicId2"]',
  'AllowedDomains': allowedDomains,
  'ExperienceConfiguration': {
    'QSearchBar': {
      'InitialTopicId': 'U4zJMVZ2n2stZf1c80u3iKySEb3BEV6f'
    }
  },
  'SessionTags': '["Key": tag-key-1,"Value": tag-value-1,{"Key": tag-key-1,"Value": tag-value-1}]',
  'SessionLifetimeInMinutes': 15
}, function(err, data) {
  console.log('Errors: ');
  console.log(err);
  console.log('Response: ');
  console.log(data);
});
```

Example

```
//The URL returned is over 900 characters. For this example, we've shortened the string
//for
//readability and added ellipsis to indicate that it's incomplete.
{
  Status: 200,
```

```
    EmbedUrl : 'https://quicksightdomain/embed/12345/dashboards/67890/sheets/12345/visuals/67890...',
    RequestId: '7bee030e-f191-45c4-97fe-d9faf0e03713'
}
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to generate the URL for the embedded Q search bar. You can use this URL in your website or app to display the Q search bar.

Example

```
using System;
using Amazon.QuickSight;
using Amazon.QuickSight.Model;

namespace GenerateQSearchBarEmbedUrlForAnonymousUser
{
    class Program
    {
        static void Main(string[] args)
        {
            var quicksightClient = new AmazonQuickSightClient(
                AccessKey,
                SecretAccessKey,
                SessionToken,
                Amazon.RegionEndpoint.USEast1);
            try
            {
                AnonymousUserQSearchBarEmbeddingConfiguration
anonymousUserQSearchBarEmbeddingConfiguration
                = new AnonymousUserQSearchBarEmbeddingConfiguration
                {
                    InitialTopicId = "U4zJMVZ2n2stZflc80u3iKySEb3BEV6f"
                };
                AnonymousUserEmbeddingExperienceConfiguration
anonymousUserEmbeddingExperienceConfiguration
                = new AnonymousUserEmbeddingExperienceConfiguration
                {
                    QSearchBar = anonymousUserQSearchBarEmbeddingConfiguration
                };
            }
        }
    }
}
```

```

        Console.WriteLine(
            quicksightClient.GenerateEmbedUrlForAnonymousUserAsync(new
GenerateEmbedUrlForAnonymousUserRequest
            {
                AwsAccountId = "111122223333",
                Namespace = "DEFAULT",
                AuthorizedResourceArns '["topic-arn-topicId1","topic-arn-
topicId2"]',
                AllowedDomains = allowedDomains,
                ExperienceConfiguration =
anonymousUserEmbeddingExperienceConfiguration,
                SessionTags = '["Key": tag-key-1,"Value": tag-value-1,{"Key":
tag-key-1,"Value": tag-value-1}]',
                SessionLifetimeInMinutes = 15,
            }).Result.EmbedUrl
        );
    } catch (Exception ex) {
        Console.WriteLine(ex.Message);
    }
}
}
}
}

```

Amazon CLI

To assume the role, choose one of the following Amazon Security Token Service (Amazon STS) API operations:

- [AssumeRole](#) – Use this operation when you are using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you are using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you are using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GenerateEmbedUrlForAnonymousUser`.

```

aws sts assume-role \
    --role-arn "arn:aws-cn:iam::111122223333:role/
embedding_quicksight_q_search_bar_role" \
    --role-session-name anonymous caller

```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. For a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"  
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"  
export AWS_SESSION_TOKEN     = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_q_search_bar_role/QuickSightEmbeddingAnonymousPolicy`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. *Throttling* is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations. In addition, it keeps each session separate and distinct. If you're using an array of web servers, for example for load balancing, and a session is reconnected to a different server, a new session begins.

To get a signed URL for the dashboard, call `generate-embed-url-for-anonymous-user` from the app server. This returns the embeddable dashboard URL. The following example shows how to generate the URL for an embedded dashboard using a server-side call for users who are making anonymous visits to your web portal or app.

```
aws quicksight generate-embed-url-for-anonymous-user \  
--aws-account-id 111122223333 \  
--namespace default-or-something-else \  

```

```
--authorized-resource-arns '["topic-arn-topicId1","topic-arn-topicId2"]' \
--allowed-domains '["domain1","domain2"]' \
--experience-configuration 'QSearchBar={InitialTopicId="topicId1"}' \
--session-tags '{"Key": tag-key-1,"Value": tag-value-1,{"Key": tag-key-1,"Value": tag-value-1}}' \
--session-lifetime-in-minutes 15
```

For more information about using this operation, see [GenerateEmbedUrlForRegisteredUser](#). You can use this and other API operations in your own code.

Step 3: Embed the Q search bar URL

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

In the following section, you can find how to embed the Q search bar URL from step 3 in your website or application page. You do this with the [Amazon Quick Sight embedding SDK](#) (JavaScript). With the SDK, you can do the following:

- Place the Q search bar on an HTML page.
- Pass parameters into the Q search bar.
- Handle error states with messages that are customized to your application.

To generate the URL that you can embed in your app, call the `GenerateEmbedUrlForAnonymousUser` API operation. This URL is valid for 5 minutes, and the resulting session is valid for up to 10 hours. The API operation provides the URL with an `auth_code` value that enables a single-sign on session.

The following shows an example response from `generate-embed-url-for-anonymous-user`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.
```

```
{
  "Status": "200",
  "EmbedUrl": "https://quicksightdomain/embedding/12345/q/search...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed the Q search bar in your webpage by using the [Amazon Quick Sight embedding SDK](#) or by adding this URL into an iframe. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes.

To do this, make sure that the domain to host the embedded Q search bar is on the *allow list*, the list of approved domains for your Amazon Quick Sight subscription. This requirement protects your data by keeping unapproved domains from hosting embedded Q search bar. For more information about adding domains for an embedded Q search bar, see [Managing domains and embedding](#).

When you use the Amazon Quick Sight Embedding SDK, the Q search bar on your page is dynamically resized based on the state. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the Q search bar and receive callbacks in terms of page load completion and errors.

The following example shows how to use the generated URL. This code is generated on your app server.

SDK 2.0

```
<!DOCTYPE html>
<html>

  <head>
    <title>Q Search Bar Embedding Example</title>
    <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@2.0.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
      const embedQSearchBar = async() => {
        const {
          createEmbeddingContext,
        } = QuickSightEmbedding;

        const embeddingContext = await createEmbeddingContext({
          onChange: (changeEvent, metadata) => {
```

```

        console.log('Context received a change', changeEvent,
metadata);
    },
    });

    const frameOptions = {
        url: "<YOUR_EMBED_URL>", // replace this value with the url
generated via embedding API
        container: '#experience-container',
        height: "700px",
        width: "1000px",
        onChange: (changeEvent, metadata) => {
            switch (changeEvent.eventName) {
                case 'FRAME_MOUNTED': {
                    console.log("Do something when the experience frame is
mounted.");

                    break;
                }
                case 'FRAME_LOADED': {
                    console.log("Do something when the experience frame is
loaded.");

                    break;
                }
            }
        },
    };

    const contentOptions = {
        hideTopicName: false,
        theme: '<YOUR_THEME_ID>',
        allowTopicSelection: true,
        onMessage: async (messageEvent, experienceMetadata) => {
            switch (messageEvent.eventName) {
                case 'Q_SEARCH_OPENED': {
                    console.log("Do something when Q Search content
expanded");

                    break;
                }
                case 'Q_SEARCH_CLOSED': {
                    console.log("Do something when Q Search content
collapsed");

                    break;
                }
                case 'Q_SEARCH_SIZE_CHANGED': {

```



```

        console.log("Do something when Q Search size changed");
        break;
    }
    case 'CONTENT_LOADED': {
        console.log("Do something when the Q Search is
loaded.");
        break;
    }
    case 'ERROR_OCCURRED': {
        console.log("Do something when the Q Search fails
loading.");
        break;
    }
}
};
const embeddedDashboardExperience = await
embeddingContext.embedQSearchBar(frameOptions, contentOptions);
};
</script>
</head>

<body onload="embedQSearchBar()">
    <div id="experience-container"></div>
</body>

</html>

```

SDK 1.0

```

<!DOCTYPE html>
<html>

    <head>
        <title>QuickSight Q Search Bar Embedding</title>
        <script src="https://unpkg.com/amazon-quicksight-embedding-sdk@1.18.0/dist/
quicksight-embedding-js-sdk.min.js"></script>
        <script type="text/javascript">
            var session

            function onError(payload) {
                console.log("Do something when the session fails loading");
            }

```

```

function onOpen() {
    console.log("Do something when the Q search bar opens");
}

function onClose() {
    console.log("Do something when the Q search bar closes");
}

function embedQSearchBar() {
    var containerDiv = document.getElementById("embeddingContainer");
    var options = {
        url: "https://us-east-1.quicksight.aws.amazon.com/sn/dashboards/
dashboardId?isauthcode=true&identityprovider=quicksight&code=authcode", // replace this
        dummy url with the one generated via embedding API
        container: containerDiv,
        width: "1000px",
        locale: "en-US",
        qSearchBarOptions: {
            expandCallback: onOpen,
            collapseCallback: onClose,
            iconDisabled: false,
            topicNameDisabled: false,
            themeId: 'bdb844d0-0fe9-4d9d-b520-0fe602d93639',
            allowTopicSelection: true
        }
    };
    session = QuickSightEmbedding.embedQSearchBar(options);
    session.on("error", onError);
}

function onCountryChange(obj) {
    session.setParameters({country: obj.value});
}
</script>
</head>

<body onload="embedQSearchBar()">
    <div id="embeddingContainer"></div>
</body>

</html>

```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded Q search bar on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Optional Amazon Quick Sight Q search bar embedding functionalities

Note

The embedded Amazon Quick Sight Q search bar provides the classic Amazon Quick Sight Q&A experience. Amazon Quick Sight integrates with Amazon Q Business to launch a new Generative Q&A experience. Developers are recommended to use the new Generative Q&A experience. For more information on the embedded Generative Q&A experience, see [Embedding the Amazon Q in Amazon Quick Sight Generative Q&A experience](#).

The following optional functionalities are available for the embedded Q search bar using the embedding SDK.

Invoke Q search bar actions

The following options are only supported for Q search bar embedding.

- Set a Q search bar question — This feature sends a question to the Q search bar and immediately queries the question. It also automatically opens the Q popover.

```
qBar.setQBarQuestion('show me monthly revenue');
```

- Close the Q popover — This feature closes the Q popover and returns the iframe to the original Q search bar size.

```
qBar.closeQPopover();
```

For more information, see the [Amazon Quick Sight embedding SDK](#).

Embedding analytics using the `GetDashboardEmbedURL` and `GetSessionEmbedURL` API operations

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

The following API operations for embedding Amazon Quick Sight dashboards and the Amazon Quick Sight console have been replaced by the `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser` API operations. You can still use them to embed analytics in your application, but they are no longer maintained and do not contain the latest embedding features or functionality. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#)

- The [GetDashboardEmbedUrl](#) API operation embeds interactive dashboards.
- The [GetSessionEmbedUrl](#) API operation embeds the Amazon Quick Sight console.

Topics

- [Embedding dashboards for everyone using `GetDashboardEmbedURL` \(old API\)](#)
- [Embedding dashboards for registered users using `GetDashboardEmbedUrl` \(old API\)](#)
- [Embedding the Amazon Quick Sight console using `GetSessionEmbedUrl` \(old API\)](#)

Embedding dashboards for everyone using `GetDashboardEmbedURL` (old API)

Important

Amazon Quick Sight has new APIs for embedding analytics:
`GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`.

You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

In the following sections, you can find detailed information on how to set up embedded Amazon Quick Sight dashboards for everyone (nonauthenticated users) using `GetDashboardEmbedURL`.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Get the URL with the authentication code attached](#)
- [Step 3: Embed the dashboard URL](#)

Step 1: Set up permissions

Important

Amazon Quick Sight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

In the following section, you can find out how to set up permissions for the backend application or web server. This task requires administrative access to IAM.

Each user who accesses a dashboard assumes a role that gives them Amazon Quick Sight access and permissions to the dashboard. To make this possible, create an IAM role in your Amazon account. Associate an IAM policy with the role to provide permissions to any user who assumes it.

The following sample policy provides these permissions for use with `IdentityType=ANONYMOUS`. For this approach to work, you also need a session pack, or session capacity pricing, on your Amazon account. Otherwise, when a user tries to access the dashboard, the error `UnsupportedPricingPlanException` is returned.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GetDashboardEmbedUrl",
        "quicksight:GetAnonymousUserEmbedUrl"
      ],
      "Resource": "*"
    }
  ]
}
```

Your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf to open the dashboard. The following example shows a role called `QuickSightEmbeddingAnonymousPolicy`, which has the sample policy preceding as its resource.

For more information regarding trust policies, see [Temporary security credentials in IAM](#) in the *IAM User Guide*.

Step 2: Get the URL with the authentication code attached

Important

Amazon Quick Sight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

In the following section, you can find how to authenticate on behalf of the anonymous visitor and get the embeddable dashboard URL on your application server.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then it adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

The following examples perform the IAM authentication on the user's behalf. It passes an identifier as the unique role session ID. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import com.amazonaws.services.quicksight.model.GetDashboardEmbedUrlRequest;
import com.amazonaws.services.quicksight.model.GetDashboardEmbedUrlResult;
```

```

/**
 * Class to call QuickSight Amazon SDK to get url for dashboard embedding.
 */
public class GetQuicksightEmbedUrlNoAuth {

    private static String ANONYMOUS = "ANONYMOUS";

    private final AmazonQuickSight quickSightClient;

    public GetQuicksightEmbedUrlNoAuth() {
        this.quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AWSCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret
                    key here
                    return new BasicAWSCredentials("access-
key", "secret-key");
                }

                @Override
                public void refresh() {}
            })
            .build();
    }

    public String getQuicksightEmbedUrl(
        final String accountId, // YOUR Amazon ACCOUNT ID
        final String dashboardId, // YOUR DASHBOARD ID TO EMBED
        final String additionalDashboardIds, // ADDITIONAL DASHBOARD-1 ADDITIONAL
DASHBOARD-2
        final boolean resetDisabled, // OPTIONAL PARAMETER TO ENABLE DISABLE
RESET BUTTON IN EMBEDDED DASHBAORD
        final boolean undoRedoDisabled // OPTIONAL PARAMETER TO ENABLE DISABLE
UNDO REDO BUTTONS IN EMBEDDED DASHBAORD
    ) throws Exception {
        GetDashboardEmbedUrlRequest getDashboardEmbedUrlRequest = new
GetDashboardEmbedUrlRequest()
            .withDashboardId(dashboardId)
            .withAdditionalDashboardIds(additionalDashboardIds)
            .withAwsAccountId(accountId)

```



```

        .withNamespace("default") // Anonymous embedding requires specifying
a valid namespace for which you want the embedding url
        .withIdentityType(ANONYMOUS)
        .withResetDisabled(resetDisabled)
        .withUndoRedoDisabled(undoRedoDisabled);

    GetDashboardEmbedUrlResult dashboardEmbedUrl =
quickSightClient.getDashboardEmbedUrl(getDashboardEmbedUrlRequest);

    return dashboardEmbedUrl.getEmbedUrl();
}
}

```

JavaScript

```

global.fetch = require('node-fetch');
const Amazon = require('aws-sdk');

function getDashboardEmbedURL(
    accountId, // YOUR Amazon ACCOUNT ID
    dashboardId, // YOUR DASHBOARD ID TO EMBED
    additionalDashboardIds, // ADDITIONAL DASHBOARD-1 ADDITIONAL DASHBOARD-2
    quicksightNamespace, // VALID NAMESPACE WHERE YOU WANT TO DO NOAUTH EMBEDDING
    resetDisabled, // OPTIONAL PARAMETER TO ENABLE DISABLE RESET BUTTON IN EMBEDDED
DASHBAORD
    undoRedoDisabled, // OPTIONAL PARAMETER TO ENABLE DISABLE UNDO REDO BUTTONS IN
EMBEDDED DASHBAORD
    getEmbedUrlCallback, // GETEMBEDURL SUCCESS CALLBACK METHOD
    errorCallback // GETEMBEDURL ERROR CALLBACK METHOD
) {
    const getDashboardParams = {
        AwsAccountId: accountId,
        DashboardId: dashboardId,
        AdditionalDashboardIds: additionalDashboardIds,
        Namespace: quicksightNamespace,
        IdentityType: 'ANONYMOUS',
        ResetDisabled: resetDisabled,
        SessionLifetimeInMinutes: 600,
        UndoRedoDisabled: undoRedoDisabled
    };

    const quicksightGetDashboard = new AWS.QuickSight({
        region: process.env.AWS_REGION,

```

```

    });

    quicksightGetDashboard.getDashboardEmbedUrl(getDashboardParams, function(err,
data) {
    if (err) {
        console.log(err, err.stack);
        errorCallback(err);
    } else {
        const result = {
            "statusCode": 200,
            "headers": {
                "Access-Control-Allow-Origin": "*", // USE YOUR WEBSITE DOMAIN
                "Access-Control-Allow-Headers": "Content-Type"
            },
            "body": JSON.stringify(data),
            "isBase64Encoded": false
        }
        getEmbedUrlCallback(result);
    }
    });
}

```

Python3

```

import json
import boto3
from botocore.exceptions import ClientError
import time

# Create QuickSight and STS clients
qs = boto3.client('quicksight', region_name='us-east-1')
sts = boto3.client('sts')

# Function to generate embedded URL
# accountId: YOUR Amazon ACCOUNT ID
# dashboardId: YOUR DASHBOARD ID TO EMBED
# additionalDashboardIds: ADDITIONAL DASHBOARD-1 ADDITIONAL DASHBOARD-2 WITHOUT
# COMMAS
# quicksightNamespace: VALID NAMESPACE WHERE YOU WANT TO DO NOAUTH EMBEDDING
# resetDisabled: PARAMETER TO ENABLE DISABLE RESET BUTTON IN EMBEDDED DASHBAORD
# undoRedoDisabled: OPTIONAL PARAMETER TO ENABLE DISABLE UNDO REDO BUTTONS IN
# EMBEDDED DASHBAORD

```

```
def getDashboardURL(accountId, dashboardId, quicksightNamespace, resetDisabled,
undoRedoDisabled):
    try:
        response = qs.get_dashboard_embed_url(
            AwsAccountId = accountId,
            DashboardId = dashboardId,
            AdditionalDashboardIds = additionalDashboardIds,
            Namespace = quicksightNamespace,
            IdentityType = 'ANONYMOUS',
            SessionLifetimeInMinutes = 600,
            UndoRedoDisabled = undoRedoDisabled,
            ResetDisabled = resetDisabled
        )

        return {
            'statusCode': 200,
            'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-
Headers": "Content-Type"},
            'body': json.dumps(response),
            'isBase64Encoded': bool('false')
        }
    except ClientError as e:
        print(e)
        return "Error generating embeddedURL: " + str(e)
```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to get the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
const Amazon = require('aws-sdk');
const https = require('https');

var quicksight = new AWS.Service({
    apiConfig: require('./quicksight-2018-04-01.min.json'),
    region: 'us-east-1',
});

quicksight.getDashboardEmbedUrl({
    'AwsAccountId': '111122223333',
```

```

        'DashboardId': 'dashboard-id',
        'AdditionalDashboardIds': 'added-dashboard-id-1 added-dashboard-id-2
added-dashboard-id-3'
        'Namespace' : 'default',
        'IdentityType': 'ANONYMOUS',
        'SessionLifetimeInMinutes': 100,
        'UndoRedoDisabled': false,
        'ResetDisabled': true

    }, function(err, data) {
        console.log('Errors: ');
        console.log(err);
        console.log('Response: ');
        console.log(data);
    });

```

Example

```

//The URL returned is over 900 characters. For this example, we've shortened the
string for
    //readability and added ellipsis to indicate that it's incomplete.
        { Status: 200,
          EmbedUrl: 'https://dashboards.example.com/
embed/620bef10822743fab329fb3751187d2d...
          RequestId: '7bee030e-f191-45c4-97fe-d9faf0e03713' }

```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to get the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```

var client = new AmazonQuickSightClient(
    AccessKey,
    SecretAccessKey,
    sessionToken,
    Amazon.RegionEndpoint.USEast1);
try
{

```

```

        Console.WriteLine(
            client.GetDashboardEmbedUrlAsync(new GetDashboardEmbedUrlRequest
            {
                AwsAccountId = "111122223333",
                DashboardId = "dashboard-id",
                AdditionalDashboardIds = "added-dashboard-id-1 added-
dashboard-id-2 added-dashboard-id-3",
                Namespace = default,
                IdentityType = IdentityType.ANONYMOUS,
                SessionLifetimeInMinutes = 600,
                UndoRedoDisabled = false,
                ResetDisabled = true
            }).Result.EmbedUrl
        );
    } catch (Exception ex) {
        Console.WriteLine(ex.Message);
    }

```

Amazon CLI

To assume the role, choose one of the following Amazon Security Token Service (Amazon STS) API operations:

- [AssumeRole](#) – Use this operation when you are using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you are using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you are using Security Assertion Markup Language (SAML) to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GetDashboardEmbedURL`.

```

aws sts assume-role \
    --role-arn "arn:aws-
cn:iam::11112222333:role/QuickSightEmbeddingAnonymousPolicy" \
    --role-session-name anonymous caller

```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an ExpiredToken error when calling the AssumeRole operation, this is probably because the previous SESSION_TOKEN is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. If you are using a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"  
export AWS_SECRET_ACCESS_KEY = "secret_key_from_assume_role"  
export AWS_SESSION_TOKEN     = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_dashboard_role/QuickSightEmbeddingAnonymousPolicy`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each visiting user. It also keeps each session separate and distinct. If you're using an array of web servers, for example for load balancing, and a session is reconnected to a different server, a new session begins.

To get a signed URL for the dashboard, call `get-dashboard-embed-url` from the app server. This returns the embeddable dashboard URL. The following example shows how to get the URL for an embedded dashboard using a server-side call for users who are making anonymous visits to your web portal or app.

```
aws quicksight get-dashboard-embed-url \  
  --aws-account-id 111122223333 \  
  --dashboard-id dashboard-id \  
  --additional-dashboard-ids added-dashboard-id-1 added-dashboard-id-2 added-  
  dashboard-id-3 \  
  --namespace default-or-something-else \  
  --identity-type ANONYMOUS \  
  --session-lifetime-in-minutes 30 \  
  --output-type text
```

```
--undo-redo-disabled true \  
--reset-disabled true \  
--user-arn arn:aws-cn:quicksight:us-east-1:111122223333:user/  
default/QuickSightEmbeddingAnonymousPolicy/embeddingession
```

For more information on using this operation, see [GetDashboardEmbedUrl](#). You can use this and other API operations in your own code.

Step 3: Embed the dashboard URL

Important

Amazon Quick Sight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

In the following section, you can find out how you can use the [Amazon Quick Sight embedding SDK](#) (JavaScript) to embed the dashboard URL from step 2 in your website or application page. With the SDK, you can do the following:

- Place the dashboard on an HTML page.
- Pass parameters into the dashboard.
- Handle error states with messages that are customized to your application.

Call the `GetDashboardEmbedUrl` API operation to get the URL that you can embed in your app. This URL is valid for 5 minutes, and the resulting session is valid for 10 hours. The API operation provides the URL with an `auth_code` that enables a single-sign on session.

The following shows an example response from `get-dashboard-embed-url`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
for
//readability and added ellipsis to indicate that it's incomplete.
{
    "Status": "200",
    "EmbedUrl": "https://dashboards.example.com/
embed/620bef10822743fab329fb3751187d2d...",
    "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed this dashboard in your web page by using the Amazon Quick Sight [Embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the dashboard and receive callbacks in terms of page load completion and errors.

The following example shows how to use the generated URL. This code resides on your app server.

```
<!DOCTYPE html>
<html>

<head>
    <title>Basic Embed</title>
    <!-- You can download the latest QuickSight embedding SDK version from https://
www.npmjs.com/package/amazon-quicksight-embedding-sdk -->
    <!-- Or you can do "npm install amazon-quicksight-embedding-sdk", if you use npm
for javascript dependencies -->
    <script src="./quicksight-embedding-js-sdk.min.js"></script>
    <script type="text/javascript">
        var dashboard;

        function embedDashboard() {
            var containerDiv = document.getElementById("embeddingContainer");
```



```
var options = {
    // replace this dummy url with the one generated via embedding API
    url: "https://us-east-1.quicksight.aws.amazon.com/sn/dashboards/
dashboardId?isauthcode=true&identityprovider=quicksight&code=authcode",
    container: containerDiv,
    scrolling: "no",
    height: "700px",
    width: "1000px",
    footerPaddingEnabled: true
};
dashboard = QuickSightEmbedding.embedDashboard(options);
}
</script>
</head>

<body onload="embedDashboard()">
    <div id="embeddingContainer"></div>
</body>

</html>
```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded dashboard on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest QuickSight embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Embedding dashboards for registered users using GetDashboardEmbedUrl (old API)

Important

Amazon Quick Sight has new APIs for embedding analytics:

GenerateEmbedUrlForAnonymousUser and GenerateEmbedUrlForRegisteredUser.

You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

In the following sections, you can find detailed information on how to set up embedded Amazon Quick Sight dashboards for registered users using `GetDashboardEmbedUrl`.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Get the URL with the authentication code attached](#)
- [Step 3: Embed the dashboard URL](#)

Step 1: Set up permissions

Important

Amazon Quick Sight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

In the following section, you can find out how to set up permissions for the backend application or web server. This task requires administrative access to IAM.

Each user who accesses a dashboard assumes a role that gives them Amazon Quick Sight access and permissions to the dashboard. To make this possible, create an IAM role in your Amazon account. Associate an IAM policy with the role to provide permissions to any user who assumes it. The IAM role needs to provide permissions to retrieve dashboard URLs. For this, you add `quicksight:GetDashboardEmbedUrl`.

The following sample policy provides these permissions for use with `IdentityType=IAM`.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GetDashboardEmbedUrl"
      ],
      "Resource": "*"
    }
  ]
}
```

The following sample policy provides permission to retrieve a dashboard URL. You use the policy with `quicksight:RegisterUser` if you are creating first-time users who are to be Amazon Quick Sight readers.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "quicksight:RegisterUser",
      "Resource": "*",
      "Effect": "Allow"
    },
    {
      "Action": "quicksight:GetDashboardEmbedUrl",
      "Resource": "*",
      "Effect": "Allow"
    }
  ]
}
```

If you use QUICKSIGHT as your `identityType` and provide the user's Amazon Resource Name (ARN), you also need to allow the `quicksight:GetAuthCode` action in your policy. The following sample policy provides this permission.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GetDashboardEmbedUrl",
        "quicksight:GetAuthCode"
      ],
      "Resource": "*"
    }
  ]
}
```

Your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf and provision the user in Amazon Quick Sight. The following example shows a role called `embedding_quicksight_dashboard_role`, which has the sample policy preceding as its resource.

For more information regarding trust policies for OpenID Connect or SAML authentication, see the following sections of the *IAM User Guide*:

- [Creating a role for web identity or OpenID Connect federation \(console\)](#)
- [Creating a role for SAML 2.0 federation \(console\)](#)

Step 2: Get the URL with the authentication code attached

Important

Amazon Quick Sight has new APIs for embedding analytics:
`GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`.

You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

In the following section, you can find out how to authenticate your user and get the embeddable dashboard URL on your application server.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then it adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

Performing the described steps ensures that each viewer of the dashboard is uniquely provisioned in Amazon Quick Sight. It also enforces per-user settings, such as the row-level security and dynamic defaults for parameters.

The following examples perform the IAM authentication on the user's behalf. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.AWSStaticCredentialsProvider;
import com.amazonaws.auth.BasicSessionCredentials;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import com.amazonaws.services.quicksight.model.GetDashboardEmbedUrlRequest;
import com.amazonaws.services.quicksight.model.GetDashboardEmbedUrlResult;
import com.amazonaws.services.securitytoken.AWSSecurityTokenService;
import com.amazonaws.services.securitytoken.model.AssumeRoleRequest;
import com.amazonaws.services.securitytoken.model.AssumeRoleResult;

/**
 * Class to call QuickSight Amazon SDK to get url for dashboard embedding.
 */
public class GetQuicksightEmbedUrlIAMAuth {

    private static String IAM = "IAM";
```

```

private final AmazonQuickSight quickSightClient;

private final AWSSecurityTokenService awsSecurityTokenService;

public GetQuicksightEmbedUrlIAMAuth(final AWSSecurityTokenService
awsSecurityTokenService) {
    this.quickSightClient = AmazonQuickSightClientBuilder
        .standard()
        .withRegion(Regions.US_EAST_1.getName())
        .withCredentials(new AWSCredentialsProvider() {
            @Override
            public AWSCredentials getCredentials() {
                // provide actual IAM access key and secret
                return new BasicAWSCredentials("access-
key here", "secret-key");
            }

            @Override
            public void refresh() {}
        })
        .build();
    this.awsSecurityTokenService = awsSecurityTokenService;
}

public String getQuicksightEmbedUrl(
    final String accountId, // YOUR Amazon ACCOUNT ID
    final String dashboardId, // YOUR DASHBOARD ID TO EMBED
    final String openIdToken, // TOKEN TO ASSUME ROLE WITH ROLEARNS
    final String roleArn, // IAM USER ROLE TO USE FOR EMBEDDING
    final String sessionName, // SESSION NAME FOR THE ROLEARNS ASSUME ROLE
    final boolean resetDisabled, // OPTIONAL PARAMETER TO ENABLE DISABLE
    RESET BUTTON IN EMBEDDED DASHBAORD
    final boolean undoRedoDisabled // OPTIONAL PARAMETER TO ENABLE DISABLE
    UNDO REDO BUTTONS IN EMBEDDED DASHBAORD
) throws Exception {
    AssumeRoleRequest request = new AssumeRoleRequest()
        .withRoleArn(roleArn)
        .withRoleSessionName(sessionName)
        .withTokenCode(openIdToken)
        .withDurationSeconds(3600);

```

```

        AssumeRoleResult assumeRoleResult =
awsSecurityTokenService.assumeRole(request);

        AWSCredentials temporaryCredentials = new BasicSessionCredentials(
            assumeRoleResult.getCredentials().getAccessKeyId(),
            assumeRoleResult.getCredentials().getSecretAccessKey(),
            assumeRoleResult.getCredentials().getSessionToken());
        AWSStaticCredentialsProvider awsStaticCredentialsProvider = new
AWSStaticCredentialsProvider(temporaryCredentials);

        GetDashboardEmbedUrlRequest getDashboardEmbedUrlRequest = new
GetDashboardEmbedUrlRequest()
            .withDashboardId(dashboardId)
            .withAwsAccountId(accountId)
            .withIdentityType(IAM)
            .withResetDisabled(resetDisabled)
            .withUndoRedoDisabled(undoRedoDisabled)
            .withRequestCredentialsProvider(awsStaticCredentialsProvider);

        GetDashboardEmbedUrlResult dashboardEmbedUrl =
quickSightClient.getDashboardEmbedUrl(getDashboardEmbedUrlRequest);

        return dashboardEmbedUrl.getEmbedUrl();
    }
}

```

JavaScript

```

global.fetch = require('node-fetch');
const Amazon = require('aws-sdk');

function getDashboardEmbedURL(
    accountId, // YOUR Amazon ACCOUNT ID
    dashboardId, // YOUR DASHBOARD ID TO EMBED
    openIdToken, // TOKEN TO ASSUME ROLE WITH ROLEARN
    roleArn, // IAM USER ROLE TO USE FOR EMBEDDING
    sessionName, // SESSION NAME FOR THE ROLEARN ASSUME ROLE
    resetDisabled, // OPTIONAL PARAMETER TO ENABLE DISABLE RESET BUTTON IN EMBEDDED
    DASHBAORD
    undoRedoDisabled, // OPTIONAL PARAMETER TO ENABLE DISABLE UNDO REDO BUTTONS IN
    EMBEDDED DASHBAORD
    getEmbedUrlCallback, // GETEMBEDURL SUCCESS CALLBACK METHOD
    errorCallback // GETEMBEDURL ERROR CALLBACK METHOD
) {
    // Implementation details
}

```

```

    ) {
    const stsClient = new AWS.STS();
    let stsParams = {
        RoleSessionName: sessionName,
        WebIdentityToken: openIdToken,
        RoleArn: roleArn
    }

    stsClient.assumeRoleWithWebIdentity(stsParams, function(err, data) {
        if (err) {
            console.log('Error assuming role');
            console.log(err, err.stack);
            errorCallback(err);
        } else {
            const getDashboardParams = {
                AwsAccountId: accountId,
                DashboardId: dashboardId,
                IdentityType: 'IAM',
                ResetDisabled: resetDisabled,
                SessionLifetimeInMinutes: 600,
                UndoRedoDisabled: undoRedoDisabled
            };

            const quicksightGetDashboard = new AWS.QuickSight({
                region: process.env.AWS_REGION,
                credentials: {
                    accessKeyId: data.Credentials.AccessKeyId,
                    secretAccessKey: data.Credentials.SecretAccessKey,
                    sessionToken: data.Credentials.SessionToken,
                    expiration: data.Credentials.Expiration
                }
            });

            quicksightGetDashboard.getDashboardEmbedUrl(getDashboardParams,
function(err, data) {
    if (err) {
        console.log(err, err.stack);
        errorCallback(err);
    } else {
        const result = {
            "statusCode": 200,
            "headers": {
                "Access-Control-Allow-Origin": "*", // USE YOUR WEBSITE
DOMAIN TO SECURE ACCESS TO GETEMBEDURL API

```



```

        "Access-Control-Allow-Headers": "Content-Type"
    },
    "body": JSON.stringify(data),
    "isBase64Encoded": false
}
getEmbedUrlCallback(result);
}
});
}
});
}
}

```

Python3

```

import json
import boto3
from botocore.exceptions import ClientError

# Create QuickSight and STS clients
qs = boto3.client('quicksight', region_name='us-east-1')
sts = boto3.client('sts')

# Function to generate embedded URL
# accountId: YOUR Amazon ACCOUNT ID
# dashboardId: YOUR DASHBOARD ID TO EMBED
# openIdToken: TOKEN TO ASSUME ROLE WITH ROLEARN
# roleArn: IAM USER ROLE TO USE FOR EMBEDDING
# sessionName: SESSION NAME FOR THE ROLEARN ASSUME ROLE
# resetDisabled: PARAMETER TO ENABLE DISABLE RESET BUTTON IN EMBEDDED DASHBAORD
# undoRedoDisabled: PARAMETER TO ENABLE DISABLE UNDO REDO BUTTONS IN EMBEDDED
# DASHBAORD
def getDashboardURL(accountId, dashboardId, openIdToken, roleArn, sessionName,
    resetDisabled, undoRedoDisabled):
    try:
        assumedRole = sts.assume_role(
            RoleArn = roleArn,
            RoleSessionName = sessionName,
            WebIdentityToken = openIdToken
        )
    except ClientError as e:
        return "Error assuming role: " + str(e)
    else:
        assumedRoleSession = boto3.Session(

```

```

        aws_access_key_id = assumedRole['Credentials']['AccessKeyId'],
        aws_secret_access_key = assumedRole['Credentials']['SecretAccessKey'],
        aws_session_token = assumedRole['Credentials']['SessionToken'],
    )
    try:
        quickSight = assumedRoleSession.client('quicksight', region_name='us-
east-1')

        response = quickSight.get_dashboard_embed_url(
            AwsAccountId = accountId,
            DashboardId = dashboardId,
            IdentityType = 'IAM',
            SessionLifetimeInMinutes = 600,
            UndoRedoDisabled = undoRedoDisabled,
            ResetDisabled = resetDisabled
        )

        return {
            'statusCode': 200,
            'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-
Allow-Headers": "Content-Type"},
            'body': json.dumps(response),
            'isBase64Encoded': bool('false')
        }
    except ClientError as e:
        return "Error generating embeddedURL: " + str(e)

```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to get the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```

const Amazon = require('aws-sdk');
const https = require('https');

var quicksight = new AWS.Service({
    apiConfig: require('./quicksight-2018-04-01.min.json'),
    region: 'us-east-1',
});

```

```
quicksight.getDashboardEmbedUrl({
  'AwsAccountId': '111122223333',
  'DashboardId': '1c1fe111-e2d2-3b30-44ef-a0e111111cde',
  'IdentityType': 'IAM',
  'ResetDisabled': true,
  'SessionLifetimeInMinutes': 100,
  'UndoRedoDisabled': false,
  'StatePersistenceEnabled': true
}, function(err, data) {
  console.log('Errors: ');
  console.log(err);
  console.log('Response: ');
  console.log(data);
});
```

Example

```
//The URL returned is over 900 characters. For this example, we've shortened the
string for
//readability and added ellipsis to indicate that it's incomplete.
{ Status: 200,
  EmbedUrl: 'https://dashboards.example.com/
embed/620bef10822743fab329fb3751187d2d...
  RequestId: '7bee030e-f191-45c4-97fe-d9faf0e03713' }
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to get the URL for the embedded dashboard. You can use this URL in your website or app to display the dashboard.

Example

```
var client = new AmazonQuickSightClient(
    AccessKey,
    SecretAccessKey,
    sessionToken,
    Amazon.RegionEndpoint.USEast1);
try
{
```

```

        Console.WriteLine(
            client.GetDashboardEmbedUrlAsync(new GetDashboardEmbedUrlRequest
            {
                AwsAccountId = "111122223333",
                DashboardId = "1c1fe111-e2d2-3b30-44ef-a0e111111cde",
                IdentityType = EmbeddingIdentityType.IAM,
                ResetDisabled = true,
                SessionLifetimeInMinutes = 100,
                UndoRedoDisabled = false,
                StatePersistenceEnabled = true
            }).Result.EmbedUrl
        );
    } catch (Exception ex) {
        Console.WriteLine(ex.Message);
    }
}

```

Amazon CLI

To assume the role, choose one of the following Amazon Security Token Service (Amazon STS) API operations:

- [AssumeRole](#) – Use this operation when you are using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you are using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you are using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GetDashboardEmbedURL`. If you are taking a just-in-time approach to add users when they first open a dashboard, the role also needs permissions enabled for `quicksight:RegisterUser`.

```

aws sts assume-role \
    --role-arn "arn:aws-cn:iam::111122223333:role/
embedding_quicksight_dashboard_role" \
    --role-session-name john.doe@example.com

```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. If you are using a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"
export AWS_SESSION_TOKEN     = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_dashboard_role/john.doe@example.com`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. *Throttling* is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations.

The role session ID also becomes the user name in Amazon Quick Sight. You can use this pattern to provision your users in Amazon Quick Sight ahead of time, or to provision them the first time they access the dashboard.

The following example shows the CLI command that you can use to provision a user. For more information about [RegisterUser](#), [DescribeUser](#), and other Amazon Quick Sight API operations, see the [Amazon Quick Sight API reference](#).

```
aws quicksight register-user \
  --aws-account-id 111122223333 \
  --namespace default \
  --identity-type IAM \
  --iam-arn "arn:aws-cn:iam::111122223333:role/
embedding_quicksight_dashboard_role" \
```

```
--user-role READER \
--user-name jhnd \
--session-name "john.doe@example.com" \
--email john.doe@example.com \
--region us-east-1 \
--custom-permissions-name TeamA1
```

If the user is authenticated through Microsoft AD, you don't need to use `RegisterUser` to set them up. Instead, they should be automatically subscribed the first time they access Amazon Quick Sight. For Microsoft AD users, you can use `DescribeUser` to get the user ARN.

The first time a user accesses Amazon Quick Sight, you can also add this user to the group that the dashboard is shared with. The following example shows the CLI command to add a user to a group.

```
aws quicksight create-group-membership \
  --aws-account-id=111122223333 \
  --namespace=default \
  --group-name=financeusers \
  --member-name="embedding_quicksight_dashboard_role/john.doe@example.com"
```

You now have a user of your app who is also a user of Amazon Quick Sight, and who has access to the dashboard.

Finally, to get a signed URL for the dashboard, call `get-dashboard-embed-url` from the app server. This returns the embeddable dashboard URL. The following example shows how to get the URL for an embedded dashboard using a server-side call for users authenticated through Amazon Managed Microsoft AD or IAM Identity Center.

```
aws quicksight get-dashboard-embed-url \
  --aws-account-id 111122223333 \
  --dashboard-id 1a1ac2b2-3fc3-4b44-5e5d-c6db6778df89 \
  --identity-type IAM \
  --session-lifetime-in-minutes 30 \
  --undo-redo-disabled true \
  --reset-disabled true \
  --state-persistence-enabled true \
  --user-arn arn:aws-cn:quicksight:us-east-1:111122223333:user/default/
embedding_quicksight_dashboard_role/embeddingsession
```

For more information on using this operation, see [GetDashboardEmbedUrl](#). You can use this and other API operations in your own code.

Step 3: Embed the dashboard URL

Important

Amazon Quick Sight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

In the following section, you can find out how you can use the [Amazon Quick Sight embedding SDK](#) (JavaScript) to embed the dashboard URL from step 3 in your website or application page. With the SDK, you can do the following:

- Place the dashboard on an HTML page.
- Pass parameters into the dashboard.
- Handle error states with messages that are customized to your application.

Call the `GetDashboardEmbedUrl` API operation to get the URL that you can embed in your app. This URL is valid for 5 minutes, and the resulting session is valid for 10 hours. The API operation provides the URL with an `auth_code` that enables a single-sign on session.

The following shows an example response from `get-dashboard-embed-url`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
//for
//readability and added ellipsis to indicate that it's incomplete.
{
  "Status": "200",
  "EmbedUrl": "https://dashboards.example.com/
embed/620bef10822743fab329fb3751187d2d...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed this dashboard in your webpage by using the [Amazon Quick Sight embedding SDK](#) or by adding this URL into an iframe. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the dashboard and receive callbacks in terms of page load completion and errors.

The following example shows how to use the generated URL. This code is generated on your app server.

```
<!DOCTYPE html>
<html>

<head>
  <title>Basic Embed</title>

  <script src="./quicksight-embedding-js-sdk.min.js"></script>
  <script type="text/javascript">
    var dashboard;

    function embedDashboard() {
      var containerDiv = document.getElementById("embeddingContainer");
      var options = {
        // replace this dummy url with the one generated via embedding API
        url: "https://us-east-1.quicksight.aws.amazon.com/sn/dashboards/
dashboardId?isauthcode=true&identityprovider=quicksight&code=authcode",
        container: containerDiv,
        scrolling: "no",
        height: "700px",
        width: "1000px",
        footerPaddingEnabled: true
      };
      dashboard = QuickSightEmbedding.embedDashboard(options);
    }
  </script>
</head>

<body onload="embedDashboard()">
  <div id="embeddingContainer"></div>
</body>
```



```
</html>
```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded dashboard on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Embedding the Amazon Quick Sight console using GetSessionEmbedUrl (old API)

Important

Amazon Quick Sight has new APIs for embedding analytics:

`GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`.

You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

Applies to: Enterprise Edition

Intended audience: Amazon Quick Suite developers

In the following sections, you can find detailed information on how to provide the Amazon Quick Sight console experience in a custom-branded authoring portal for registered users using the `GetSessionEmbedUrl` API.

Topics

- [Step 1: Set up permissions](#)
- [Step 2: Get the URL with the authentication code attached](#)
- [Step 3: Embed the console session URL](#)

Step 1: Set up permissions

Important

Amazon Quick Sight has new APIs for embedding analytics:

`GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

In the following section, you can find out how to set up permissions for the backend application or web server. This task requires administrative access to IAM.

Each user who accesses a Amazon Quick Sight assumes a role that gives them Amazon Quick Sight access and permissions to the console session. To make this possible, create an IAM role in your Amazon account. Associate an IAM policy with the role to provide permissions to any user who assumes it. Add `quicksight:RegisterUser` permissions to ensure that the reader can access Amazon Quick Sight in a read-only fashion, and not have access to any other data or creation capability. The IAM role also needs to provide permissions to retrieve console session URLs. For this, you add `quicksight:GetSessionEmbedUrl`.

The following sample policy provides these permissions for use with `IdentityType=IAM`.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "quicksight:RegisterUser",
```

```
    "Resource": "*",
    "Effect": "Allow"
  },
  {
    "Action": "quicksight:GetSessionEmbedUrl",
    "Resource": "*",
    "Effect": "Allow"
  }
]
```

The following sample policy provides permission to retrieve a console session URL. You use the policy without `quicksight:RegisterUser` if you are creating users before they access an embedded session.

JSON

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GetSessionEmbedUrl"
      ],
      "Resource": "*"
    }
  ]
}
```

If you use `QUICKSIGHT` as your `identityType` and provide the user's Amazon Resource Name (ARN), you also need to allow the `quicksight:GetAuthCode` action in your policy. The following sample policy provides this permission.

JSON

```
{
  "Version": "2012-10-17",
```

```
"Statement": [  
  {  
    "Effect": "Allow",  
    "Action": [  
      "quicksight:GetSessionEmbedUrl",  
      "quicksight:GetAuthCode"  
    ],  
    "Resource": "*"  
  }  
]
```

Your application's IAM identity must have a trust policy associated with it to allow access to the role that you just created. This means that when a user accesses your application, your application can assume the role on the user's behalf and provision the user in Amazon Quick Sight. The following example shows a role called `embedding_quicksight_console_session_role`, which has the sample policy preceding as its resource.

For more information regarding trust policies for OpenID Connect or SAML authentication, see the following sections of the *IAM User Guide*:

- [Creating a role for web identity or OpenID Connect federation \(console\)](#)
- [Creating a role for SAML 2.0 federation \(console\)](#)

Step 2: Get the URL with the authentication code attached

Important

Amazon Quick Sight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

In the following section, you can find out how to authenticate your user and get the embeddable console session URL on your application server.

When a user accesses your app, the app assumes the IAM role on the user's behalf. Then it adds the user to Amazon Quick Sight, if that user doesn't already exist. Next, it passes an identifier as the unique role session ID.

Performing the described steps ensures that each viewer of the console session is uniquely provisioned in Amazon Quick Sight. It also enforces per-user settings, such as the row-level security and dynamic defaults for parameters.

The following examples perform the IAM authentication on the user's behalf. This code runs on your app server.

Java

```
import com.amazonaws.auth.AWSCredentials;
import com.amazonaws.auth.BasicAWSCredentials;
import com.amazonaws.auth.AWSCredentialsProvider;
import com.amazonaws.regions.Regions;
import com.amazonaws.services.quicksight.AmazonQuickSight;
import com.amazonaws.services.quicksight.AmazonQuickSightClientBuilder;
import com.amazonaws.services.quicksight.model.GetSessionEmbedUrlRequest;
import com.amazonaws.services.quicksight.model.GetSessionEmbedUrlResult;

/**
 * Class to call QuickSight Amazon SDK to get url for session embedding.
 */
public class GetSessionEmbedUrlQSAuth {

    private final AmazonQuickSight quickSightClient;

    public GetSessionEmbedUrlQSAuth() {
        this.quickSightClient = AmazonQuickSightClientBuilder
            .standard()
            .withRegion(Regions.US_EAST_1.getName())
            .withCredentials(new AWSCredentialsProvider() {
                @Override
                public AWSCredentials getCredentials() {
                    // provide actual IAM access key and secret
                    key here
                    return new BasicAWSCredentials("access-
key", "secret-key");
                }
            })
            .build();
    }

    @Override
```

```

        public void refresh() {}
    }

    )
    .build();
}

public String getQuicksightEmbedUrl(
    final String accountId, // YOUR Amazon ACCOUNT ID
    final String userArn // REGISTERED USER ARN TO USE FOR EMBEDDING.
    REFER TO GETEMBEDURL SECTION IN DEV PORTAL TO FIND OUT HOW TO GET USER ARN FOR A
    QUICKSIGHT USER
) throws Exception {
    GetSessionEmbedUrlRequest getSessionEmbedUrlRequest = new
    GetSessionEmbedUrlRequest()
        .withAwsAccountId(accountId)
        .withEntryPoint("/start")
        .withUserArn(userArn);

    GetSessionEmbedUrlResult sessionEmbedUrl =
    quickSightClient.getSessionEmbedUrl(getSessionEmbedUrlRequest);

    return sessionEmbedUrl.getEmbedUrl();
}
}

```

JavaScript

```

global.fetch = require('node-fetch');
const Amazon = require('aws-sdk');

function getSessionEmbedURL(
    accountId, // YOUR Amazon ACCOUNT ID
    userArn, // REGISTERED USER ARN TO USE FOR EMBEDDING. REFER TO GETEMBEDURL
    SECTION IN DEV PORTAL TO FIND OUT HOW TO GET USER ARN FOR A QUICKSIGHT USER
    getEmbedUrlCallback, // GETEMBEDURL SUCCESS CALLBACK METHOD
    errorCallback // GETEMBEDURL ERROR CALLBACK METHOD
) {
    const getSessionParams = {
        AwsAccountId: accountId,
        EntryPoint: "/start",
        UserArn: userArn,
        SessionLifetimeInMinutes: 600,
    };
};

```

```

const quicksightGetSession = new AWS.QuickSight({
  region: process.env.AWS_REGION,
});

quicksightGetSession.getSessionEmbedUrl(getSessionParams, function(err, data) {
  if (err) {
    console.log(err, err.stack);
    errorCallback(err);
  } else {
    const result = {
      "statusCode": 200,
      "headers": {
        "Access-Control-Allow-Origin": "*", // USE YOUR WEBSITE DOMAIN
        "Access-Control-Allow-Headers": "Content-Type"
      },
      "body": JSON.stringify(data),
      "isBase64Encoded": false
    }
    getEmbedUrlCallback(result);
  }
});
}

```

Python3

```

import json
import boto3
from botocore.exceptions import ClientError
import time

# Create QuickSight and STS clients
qs = boto3.client('quicksight', region_name='us-east-1')
sts = boto3.client('sts')

# Function to generate embedded URL
# accountId: YOUR Amazon ACCOUNT ID
# userArn: REGISTERED USER ARN TO USE FOR EMBEDDING. REFER TO GETEMBEDURL SECTION IN
# DEV PORTAL TO FIND OUT HOW TO GET USER ARN FOR A QUICKSIGHT USER
def getSessionEmbedURL(accountId, userArn):
    try:
        response = qs.get_session_embed_url(

```

```

        AwsAccountId = accountId,
        EntryPoint = "/start",
        UserArn = userArn,
        SessionLifetimeInMinutes = 600
    )

    return {
        'statusCode': 200,
        'headers': {"Access-Control-Allow-Origin": "*", "Access-Control-Allow-Headers": "Content-Type"},
        'body': json.dumps(response),
        'isBase64Encoded': bool('false')
    }
except ClientError as e:
    print(e)
    return "Error generating embeddedURL: " + str(e)

```

Node.js

The following example shows the JavaScript (Node.js) that you can use on the app server to get the URL for the embedded console session. You can use this URL in your website or app to display the console session.

Example

```

const Amazon = require('aws-sdk');
const https = require('https');

var quicksight = new AWS.Service({
    apiConfig: require('./quicksight-2018-04-01.min.json'),
    region: 'us-east-1',
});

quicksight.getSessionEmbedUrl({
    'AwsAccountId': '111122223333',
    'EntryPoint': 'https://url-for-console-page-to-open',
    'SessionLifetimeInMinutes': 600,
    'UserArn': 'USER_ARN'
}, function(err, data) {
    console.log('Errors: ');
    console.log(err);
    console.log('Response: ');

```



```
        console.log(data);
    });
```

Example

//The URL returned is over 900 characters. For this example, we've shortened the string for

```
    //readability and added ellipsis to indicate that it's incomplete.
        { Status: 200,
          EmbedUrl: 'https://dashboards.example.com/
embed/620bef10822743fab329fb3751187d2d...
          RequestId: '7bee030e-f191-45c4-97fe-d9faf0e03713' }
```

.NET/C#

The following example shows the .NET/C# code that you can use on the app server to get the URL for the embedded console session. You can use this URL in your website or app to display the console.

Example

```
var client = new AmazonQuickSightClient(
    AccessKey,
    SecretAccessKey,
    sessionToken,
    Amazon.RegionEndpoint.USEast1);
try
{
    Console.WriteLine(
        client.GetSessionEmbedUrlAsync(new GetSessionEmbedUrlRequest
        {
            'AwsAccountId': '111122223333',
            'EntryPoint': 'https://url-for-console-page-to-open',
            'SessionLifetimeInMinutes': 600,
            'UserArn': 'USER_ARN'
            AwsAccountId = 111122223333,
            EntryPoint = https://url-for-console-page-to-open,
            SessionLifetimeInMinutes = 600,
            UserArn = 'USER_ARN'
        }).Result.EmbedUrl
    );
```

```
    } catch (Exception ex) {  
        Console.WriteLine(ex.Message);  
    }
```

Amazon CLI

To assume the role, choose one of the following Amazon Security Token Service (Amazon STS) API operations:

- [AssumeRole](#) – Use this operation when you are using an IAM identity to assume the role.
- [AssumeRoleWithWebIdentity](#) – Use this operation when you are using a web identity provider to authenticate your user.
- [AssumeRoleWithSaml](#) – Use this operation when you are using SAML to authenticate your users.

The following example shows the CLI command to set the IAM role. The role needs to have permissions enabled for `quicksight:GetSessionEmbedUrl`. If you are taking a just-in-time approach to add users when they first open Amazon Quick Sight, the role also needs permissions enabled for `quicksight:RegisterUser`.

```
aws sts assume-role \  
    --role-arn "arn:aws-cn:iam::111122223333:role/  
embedding_quicksight_dashboard_role" \  
    --role-session-name john.doe@example.com
```

The `assume-role` operation returns three output parameters: the access key, the secret key, and the session token.

Note

If you get an `ExpiredToken` error when calling the `AssumeRole` operation, this is probably because the previous `SESSION_TOKEN` is still in the environment variables. Clear this by setting the following variables:

- `AWS_ACCESS_KEY_ID`
- `AWS_SECRET_ACCESS_KEY`
- `AWS_SESSION_TOKEN`

The following example shows how to set these three parameters in the CLI. If you are using a Microsoft Windows machine, use `set` instead of `export`.

```
export AWS_ACCESS_KEY_ID      = "access_key_from_assume_role"
export AWS_SECRET_ACCESS_KEY  = "secret_key_from_assume_role"
export AWS_SESSION_TOKEN      = "session_token_from_assume_role"
```

Running these commands sets the role session ID of the user visiting your website to `embedding_quicksight_console_session_role/john.doe@example.com`. The role session ID is made up of the role name from `role-arn` and the `role-session-name` value. Using the unique role session ID for each user ensures that appropriate permissions are set for each user. It also prevents any throttling of user access. Throttling is a security feature that prevents the same user from accessing Amazon Quick Sight from multiple locations.

The role session ID also becomes the user name in Amazon Quick Sight. You can use this pattern to provision your users in Amazon Quick Sight ahead of time, or to provision them the first time they access a console session.

The following example shows the CLI command that you can use to provision a user. For more information about [RegisterUser](#), [DescribeUser](#), and other Amazon Quick Sight API operations, see the [Amazon Quick Sight API reference](#).

```
aws quicksight register-user \
  --aws-account-id 111122223333 \
  --namespace default \
  --identity-type IAM \
  --iam-arn "arn:aws-cn:iam::111122223333:role/
embedding_quicksight_dashboard_role" \
  --user-role READER \
  --user-name jhnd \
  --session-name "john.doe@example.com" \
  --email john.doe@example.com \
  --region us-east-1 \
  --custom-permissions-name TeamA1
```

If the user is authenticated through Microsoft AD, you don't need to use `RegisterUser` to set them up. Instead, they should be automatically subscribed the first time they access Amazon Quick Sight. For Microsoft AD users, you can use `DescribeUser` to get the user ARN.

The first time a user accesses Amazon Quick Sight, you can also add this user to the appropriate group. The following example shows the CLI command to add a user to a group.

```
aws quicksight create-group-membership \  
  --aws-account-id=111122223333 \  
  --namespace=default \  
  --group-name=financeusers \  
  --member-name="embedding_quicksight_dashboard_role/john.doe@example.com"
```

You now have a user of your app who is also a user of Amazon Quick Sight, and who has access to the Amazon Quick Sight console session.

Finally, to get a signed URL for the console session, call `get-session-embed-url` from the app server. This returns the embeddable console session URL. The following example shows how to get the URL for an embedded console session using a server-side call for users authenticated through Amazon Managed Microsoft AD or Single Sign-on (IAM Identity Center).

```
aws quicksight get-dashboard-embed-url \  
  --aws-account-id 111122223333 \  
  --entry-point the-url-for--the-console-session \  
  --session-lifetime-in-minutes 600 \  
  --user-arn arn:aws-cn:quicksight:us-east-1:111122223333:user/  
default/embedding_quicksight_dashboard_role/embeddingession
```

For more information on using this operation, see [GetSessionEmbedUrl](#). You can use this and other API operations in your own code.

Step 3: Embed the console session URL

Important

Amazon Quick Sight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the Amazon Quick Sight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding Amazon Quick Sight analytics into your applications](#).

In the following section, you can find out how you can use the [Amazon Quick Sight embedding SDK](#) (JavaScript) to embed the console session URL from step 3 in your website or application page. With the SDK, you can do the following:

- Place the console session on an HTML page.
- Pass parameters into the console session.
- Handle error states with messages that are customized to your application.

Call the `GetSessionEmbedUrl` API operation to get the URL that you can embed in your app. This URL is valid for 5 minutes, and the resulting session is valid for 10 hours. The API operation provides the URL with an `auth_code` that enables a single-sign on session.

The following shows an example response from `get-dashboard-embed-url`.

```
//The URL returned is over 900 characters. For this example, we've shortened the string
//for readability and added ellipsis to indicate that it's incomplete.
{
  "Status": "200",
  "EmbedUrl": "https://dashboards.example.com/
embed/620bef10822743fab329fb3751187d2d...",
  "RequestId": "7bee030e-f191-45c4-97fe-d9faf0e03713"
}
```

Embed this console session in your webpage by using the Amazon Quick Sight [Embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), Amazon Quick Sight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, Amazon Quick Sight provides a responsive layout that is modified as your window size changes. By using the Amazon Quick Sight Embedding SDK, you can also control parameters within the console session and receive callbacks in terms of page load completion and errors.

The following example shows how to use the generated URL. This code is generated on your app server.

```
<!DOCTYPE html>
<html>

<head>
  <title>Basic Embed</title>

  <script src="./quicksight-embedding-js-sdk.min.js"></script>
  <script type="text/javascript">
    var dashboard;
```

```
function embedDashboard() {
    var containerDiv = document.getElementById("embeddingContainer");
    var options = {
        // replace this dummy url with the one generated via embedding API
        url: "https://us-east-1.quicksight.aws.amazon.com/sn/dashboards/
dashboardId?isauthcode=true&identityprovider=quicksight&code=authcode",
        container: containerDiv,
        scrolling: "no",
        height: "700px",
        width: "1000px",
        footerPaddingEnabled: true
    };
    dashboard = QuickSightEmbedding.embedDashboard(options);
}
</script>
</head>

<body onload="embedDashboard()">
    <div id="embeddingContainer"></div>
</body>

</html>
```

For this example to work, make sure to use the Amazon Quick Sight Embedding SDK to load the embedded console session on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon Quick Sight embedding SDK](#) from GitHub. This repository is maintained by a group of Amazon Quick Sight developers.
- Download the latest embedding SDK version from <https://www.npmjs.com/package/amazon-quicksight-embedding-sdk>.
- If you use npm for JavaScript dependencies, download and install it by running the following command.

```
npm install amazon-quicksight-embedding-sdk
```

Using Amazon Quick Research

Amazon Quick Research is a feature of Amazon Quick Suite that enables you to conduct comprehensive research by analyzing multiple data sources and generating detailed reports. Quick Research uses artificial intelligence to help you gather, analyze, and synthesize information from various sources including web search, uploaded files, and connected data spaces.

With Quick Research, you can define research objectives, select relevant data sources, and receive AI-generated research reports with proper citations and source tracing. This helps you make informed decisions based on comprehensive analysis of available information.

What you can do with Quick Research

Quick Research supports a wide range of research activities and use cases:

- Market research and competitive analysis
- Literature reviews and academic research
- Business intelligence and trend analysis
- Product research and feature comparison
- Policy and regulatory research
- Technical documentation and best practices research
- Customer insights and feedback analysis
- Industry reports and white paper synthesis

Quick Research helps you save time by automatically gathering information from multiple sources, analyzing the content, and presenting findings in a structured format with proper citations and source references.

Starting new research

Creating a new research project in Amazon Quick Research involves a structured workflow that guides you from initial objective setting through final report generation. This process ensures comprehensive research coverage while maintaining focus on your specific goals.

Enter research objective

The first step in creating a new research project is to define your research objective. A well-written research objective helps Quick Research understand what information you're looking for and guides the AI in selecting relevant sources and generating focused results.

To enter your research objective

1. In the main navigation bar, choose **Research**.
2. In the Quick Research interface, choose **New Research**.

New research

Research objective

I need to understand emerging trends in sustainable packaging for the food delivery industry. This research will help our product team design new packaging solutions. We're particularly interested in biodegradable materials that maintain food temperature for at least 30 minutes. The findings will be presented to senior management for budget allocation decisions.

7036 characters remaining

Research materials

☒ Web search ⓘ ☐

☒ File uploads [Upload](#)

Pro tip: Working with larger collections? Spaces helps you organize and optimize extensive research materials. [Dismiss](#)

☒ Quick Suite assets [Browse](#)

Pro tip: For optimal results, consider selecting the most relevant resources for your specific research question. [Dismiss](#)

No resources added, click Browse to add some resources.

[Share feedback](#) [Usage is subject to AWS Responsible AI Policy](#) [Create plan](#)

The quick and easy way to get quality research done fast.

You're in control – I'll follow your expertise and guidance throughout our collaboration.

Set up the Research Document

Share your research objective in a few sentences to describe your goal, frameworks, desired outcomes and any other essential information including your industry context and any specific requirements or constraints.

Example

Evaluate renewable energy investment opportunities for our corporate sustainability portfolio to support Q1 2025 capital allocation decisions. We're allocating \$50M across North American markets with a 7-year investment horizon. Focus on solar, wind, and battery storage sectors from 2023-2024 performance data. Key metrics: IRR projections, regulatory incentive impacts, market growth trends, and risk-adjusted returns. Cost sensitivity is critical – prioritize solutions under \$2M per MW capacity.

Primarily build research on our internal investment analysis from Document: "Renewable_Portfolio_Performance_2024.pdf" in Knowledge base: "Investment Research", which is added in Space: "Financial Analysis", and Chart: "Investment Performance Trends" in Dashboard: "Sustainability Portfolio Dashboard". Use web research to supplement with recent policy changes, emerging technology costs, and competitor strategies not covered in our internal data.

Next, select upload any helpful documents, and select Quick Suite assets as called out in the research objective to generate a plan.

3. In the **Research objective** field, enter a clear description of what you want to research.

Describe your research goal by stating what you want to achieve, for whom, and why. For best results, mention your intended audience, and any key metrics or focus areas that will help deliver exactly what you need.

Share your research objective in a few sentences to describe your goal frameworks, desired outcomes and any other essential information including your industry context and any specific

requirements or constraints. If you have preferences for specific websites to include or exclude from your research, mention those as well.

Example: I need to understand emerging trends in sustainable packaging for the food industry. This research will help our product team design new packaging solutions. We're particularly interested in biodegradable materials that maintain food temperature for at least 30 minutes. The findings will be presented to senior management for budget allocation decisions.

4. After entering your research objective, select what materials to use in **Research materials**, then choose **Create plan** to create your draft plan.

Tips for writing a good research objective

Follow these guidelines to create effective research objectives:

- Be specific about what you want to learn or understand
- Include relevant context about your industry, market, or domain
- Specify the time frame if relevant (e.g., "recent trends," "2024 data")
- Mention the intended use of the research (e.g., "for strategic planning," "to inform product decisions")
- Use clear, concise language and avoid ambiguous terms
- Include any specific aspects or angles you want to explore

Example of a well-written research objective: "Analyze the current state of artificial intelligence adoption in healthcare, focusing on diagnostic imaging applications, regulatory challenges, and market opportunities for 2024-2025."

Select research materials

After defining your research objective, you can select the data sources that Quick Research will use to gather information. You can choose from multiple types of research materials to ensure comprehensive coverage of your topic.

To select research materials

1. In the material selection interface, review the available data source options.
2. Select the types of materials you want to include:

- a. Toggle **Web search** to include online sources
- b. Choose **File upload** to add specific documents

For best results, include key documents that are most relevant to your research goals. Adding more materials may increase research time and lead to less focused insights.

If you're working with larger collections, consider using Spaces to organize and optimize extensive research materials for better results.

- c. Select **Quick Suite assets** to include data spaces, dashboards, and knowledge bases. These are collections of files, documents, and analytics that you've organized in Quick Suite for easy access and analysis. For more information about creating and managing data spaces, see [Organize and collaborate using Amazon Quick Suite spaces](#).
3. Choose **Create plan** to proceed to research plan review.


Web search

Enable web search to allow Quick Research to gather information from publicly available online sources. This includes academic papers, industry reports, news articles, and other relevant web content related to your research objective.

Amazon Quick Suite uses the internet to enhance your results. Web search queries will be processed securely in an Amazon region located in the US. For more information, see [Amazon Quick Suite User Guide](#).

Uploaded files

Upload specific documents, PDFs, spreadsheets, or other files that you want Quick Research to analyze as part of your research. This is useful when you have particular sources or documents that are directly relevant to your research objective.



To upload files

1. Choose **Upload** to open the **File Uploads** window.
2. Add files using one of the following methods:
 - a. Choose **Add Files** to navigate to your files in File Explorer
 - b. Drag and drop up to 20 files into the upload area

If you need to include more than 20 files in your research, consider organizing them in Spaces instead.

3. When you're finished selecting files, choose **Done**.

Accepted file formats include .pdf, .txt, .doc, .xls, and .csv files. The file size limit is 25MB per file.

Quick Suite assets

Connect Quick Research to your existing data spaces to include internal documents, reports, and knowledge bases in your research. This allows you to combine external web sources with your organization's proprietary information.

For optimal results, consider selecting the most relevant assets for your specific research question. Choose **Browse** to add assets from Quick Suite.

New research

Research objective

I need to understand emerging trends in sustainable packaging for the food delivery industry. This research will help our product team design new packaging solutions. We're particularly interested in biodegradable materials that maintain food temperature for at least 30 minutes. The findings will be presented to senior management for budget allocation decisions.

7036 characters remaining

Research materials

☒ Web search ⓘ ☐

☒ File uploads [Upload](#)

Pro tip: Working with larger collections? Spaces helps you organize and optimize extensive research materials. [Dismiss](#)

☒ Quick Suite assets [Browse](#)

Pro tip: For optimal results, consider selecting the most relevant resources for your specific research question. [Dismiss](#)

[Create plan](#)

Usage is subject to [AWS Responsible AI Policy](#)

[Share feedback](#)

Add Quick Suite assets

Provide relevant knowledge to your research by adding Quick Suite assets

Search resources here...

Name	Owner	Last mo...
☐ TPW Space	Others	28 minutes ago
☐ Hiring Manager Compa...	Others	15 hours ago
☐ ASP_2025	Others	21 hours ago
☐ AWS One Team 2 Space	Others	21 hours ago
☒ Account Executive Enha...	Others	21 hours ago
☐ APN Space	Others	21 hours ago
☐ OP1-OP2 Strategic Pla...	Others	21 hours ago
☐ LASER Agent	Others	21 hours ago
☐ CET MLDA Quality Team	Others	a day ago
☒ Amazon Quick Suite Us...	Others	2 days ago

Selected: 2 resources selected [Cancel](#) [Add](#)

To add Quick Suite assets

1. Choose **Browse** to open the **Add assets** window in the right pane.
2. Select from the available resource type tabs:
 - a. **Recent** - Recently accessed assets
 - b. **Space** - Available data spaces (maximum of 2 selections)
 - c. **Dashboard** - Dashboard resources (maximum of 2 selections)
 - d. **Topic** - Topic resources (maximum of 2 selections)

- e. **Knowledgebase** - Knowledge base resources
3. Review the available assets, which display the Name, Owner, and Last Modified date and time for each asset.
4. Select the assets you want to include in your research.
5. When you're finished specifying all assets, choose **Add**.

Review and finalize research plan

Before Quick Research begins gathering and analyzing information, you have the opportunity to review and refine your research plan. This step helps ensure that the research will meet your specific needs and expectations.

The screenshot shows the 'Review and finalize research plan' interface in Amazon Quick Suite. The top navigation bar includes 'Explore', 'Amazon Quick Suite', 'Research', and the current topic: 'Biodegradable Food Delivery Packaging: Thermal Performance Trends and Investment Opportunities 2025'. A 'Start researching' button is in the top right.

Research objective: I need to understand emerging trends in sustainable packaging for the food delivery industry. This research will help our product team design new packaging solutions. We're particularly interested i...

Research materials:

- ☒ Web search 🔍 🔧
- ☒ File uploads 📁
No files uploaded
- ☒ Quick Suite assets 📦 🔧
 - ☐ All data and apps 🔍
 - ☒ Select data and apps

Usage is subject to [AWS Responsible AI Policy](#). [Share feedback](#)

Topics to explore

Biodegradable Food Delivery Packaging: Thermal Performance Trends and Investment Opportunities 2025

I need to understand emerging trends in sustainable packaging for the food delivery industry. This research will help our product team design new packaging solutions. We're particularly interested in biodegradable materials that maintain food temperature for at least 30 minutes. The findings will be presented to senior management for budget allocation decisions.

Current Market Landscape of Sustainable Packaging in Food Delivery

This section will provide an overview of the existing sustainable packaging solutions in the food delivery industry as of 2025. It will analyze market size, key players, and the most widely adopted sustainable packaging technologies currently in use.

Biodegradable Materials Analysis for Food Packaging

This section will examine various biodegradable materials suitable for food delivery packaging, including their composition, degradation timeframes, and food safety certifications. It will evaluate materials like PLA, bagasse, mushroom packaging, seaweed-based solutions, and other plant-based alternatives specifically for food contact applications.

Temperature Retention Technologies in Sustainable Packaging

This section will focus on innovations in insulation properties of biodegradable materials that maintain food temperature for 30+ minutes. It will analyze thermal performance data of different sustainable materials and compare them with traditional packaging in various environmental conditions.

Regulatory Framework and Compliance Requirements

This section will outline current and upcoming regulations affecting sustainable

Research

You may consider following dimensions to further improve the research plan:

- Specify geographic markets of interest for targeted regulatory and consumer preference analysis.
- Identify budget constraints to focus cost analysis on feasible packaging solutions.
- Determine specific food types being delivered to address unique packaging requirements.

What can I do to make it better?

[Revise plan](#)

Revision history

To review and finalize your research plan

1. Review the AI-generated suggestions in the **Revise Plan** pane on the right for improving your research plan.

2. (Optional) To make revisions, enter your changes in the **How should I change the plan?** field. You can add comments and notes to provide additional context or specific instructions for the AI, such as clarifying particular aspects you want emphasized, specifying exclusions, or providing domain-specific guidance that will help generate more relevant results.
3. Choose **Revise plan** to update the plan with your revisions.
4. Choose **Start Research** to begin the analysis process.

Research suggestions

Quick Research analyzes your research objective and selected materials to provide intelligent suggestions for improving your research plan. These suggestions appear in the **Revise Plan** pane and might include additional data sources, refined search terms, or alternative approaches to your research question.

For example, if your research objective is about sustainable packaging for the food industry, suggestions might include specifying geographic regions for regulatory analysis, identifying specific food types requiring temperature maintenance, or providing budget constraints for cost comparison analysis.

Revision history

The revision history section tracks all changes made to your research plan during the review process. This allows you to see what modifications have been applied and provides a record of how your research plan has evolved.

View research report

Once you select to create the research, Quick Research will work through a series of research tasks. Research typically takes 20-40 minutes to complete. The window displays the percentage complete along with the specific tasks Quick Research is working on and which sources are being used for each research task.

Explore

Amazon Quick Suite

Research

Biodegradable Food Delivery Packaging: Thermal Performance Trends for 2025-2026

Biodegradable Food Delivery Packaging: Thermal Performance Trends for 2025...

Version 1

Cancel research

Research

I need to understand emerging trends in sustainable packaging for the food delivery industry. This research will help our product team design new packaging solutions. We're particularly interested i...

Materials

Topics

Research materials

Web search

File uploads

Quick Suite assets

No files uploaded

All data and apps

Select data and apps

Usage is subject to [AWS Responsible AI Policy](#)

Share feedback

Your research has begun and will take about 30-60 minutes to complete. Feel free to leave this page - you'll be notified when it's ready.

2%

Working through research tasks

Websites

patents.justia.comwww.thermal-engineering.orgipcpack.com

Investigating biodegradable thermal insulation sources by retrieving and analyzing patent, engineering, and commercial content related to packaging technologies....

Prepare comprehensive research documentation for seaweed-based polymers in advanced bioplastics packaging with thermal performance analysis for food delivery applications by 2025

Websites

www.mdpi.comzephhr.newscientist.com

Searching for detailed information on mushroom packaging thermal requirements and food delivery temperature standards to support research investigation....

Prepare a comprehensive research strategy for investigating sustainable packaging trends in the food delivery industry for 2025-2026, focusing on market adoption rates and emerging technologies.

Websites

www.mordorintelligence.comwww.biodeaderpack.comwww.restaurantware.com

Investigating sustainability initiatives for major food delivery companies by searching and analyzing online resources for packaging innovations and future trends...

Prepare comprehensive research on plant-based foam insulation for sustainable food packaging with thermal retention capabilities, focusing on recycled fiber insulators and air-pocket design innovations.

Websites

www.packagingstrategies.commse.washington.educmsgreen.com

Investigating detailed content retrieval from multiple URLs related to foam packaging and thermal protection technologies...

After Quick Research completes its analysis, you can view the comprehensive research report that synthesizes information from all selected sources. The report presents findings in a structured format with supporting evidence and source references.

View research report

1606

The screenshot displays the Amazon Quick Suite interface for a research report titled "Biodegradable Food Delivery Packaging: Thermal Performance Trends 2025". The interface is divided into several sections:

- Top Navigation:** Includes "Explore", "Amazon Quick Suite", "Research", and the report title. On the right, there are buttons for "Version 1", "Share", and "Summarize".
- Left Sidebar:**
 - Research Pane:** Contains a text box with the research objective: "I need to understand emerging trends in sustainable packaging for the food delivery industry. This research will help our product team design new packaging solutions. We're particularly interested i...".
 - Topics Tab:** Lists the following topics: "Biodegradable Food Delivery Packaging: Thermal Performance Trends 2025", "Executive Summary", "Current Market Landscape of Sustainable Packaging in Food Delivery", "Biodegradable Materials Innovation and Performance", "Temperature Retention Technologies in Sustainable Packaging", and "Usage is subject to AWS Responsible AI Policy".
- Main Report Area:**
 - Rate report:** Displayed as five stars.
 - Download Button:** A button to download the report.
 - Biodegradable Food Delivery Packaging: Thermal Performance Trends 2025**
 - Executive Summary:** "The sustainable packaging market in the food delivery industry is experiencing significant growth, valued at USD 65.73 billion in 2025 and projected to reach USD 88.46 billion by 2030, growing at a CAGR of 6.12% +2. This expansion is driven by increasing consumer demand, evolving regulatory frameworks, and technological innovations that are narrowing the performance and cost gaps between traditional and sustainable packaging options."
 - Key Materials and Performance:** "Our analysis of biodegradable materials reveals several options capable of maintaining food temperature for at least 30 minutes, with bagasse (sugarcane fiber) emerging as a particularly promising solution. Bagasse offers exceptional temperature resistance (-25°C to 220°C) +2, making it suitable for hot foods while being fully biodegradable and competitively priced at approximately \$0.08 per unit 5. Other viable materials include PHA (Polyhydroxyalkanoates), which offers good thermal stability, and molded fiber alternatives that balance performance with environmental benefits. Advanced temperature retention technologies are further enhancing the capabilities of sustainable packaging. Phase Change Materials (PCMs) can maintain hot food temperatures above 60°C for up to 90 minutes and cold food below 5°C for up to 180 minutes. Vacuum Insulation Panels (VIPs) provide
- Right Sidebar:**
 - Comments (1):** A section for user comments.
 - Research:** A section titled "Here's how you can make this research even better:" with a list of suggestions:
 - Highlight any text to add specific comments (like exploring a claim further or requesting more evidence)
 - Add overall comment in this panel (new angles to explore or additional context to consider)
 - When you're ready, create the next version. I'll update the report based on all your comments.
 - Comment history (1):** A section for the history of comments.
 - Add a comment:** A text input field for adding a new comment.

You can use the **Download** button at the top to download the report as a PDF or Word document. There is also a **Summarize** button in the upper right. For more information, see [the section called "Summarizing the research"](#). You can also select **Reading mode** to hide sidebars and focus on the report content.


When you're ready to share your research with others, use the **Share** button. For more information, see [the section called "Sharing the research"](#). If you want to refine the research further before sharing, you can make updates and iterate on the report. For more information, see [the section called "Updating research"](#).

In the left pane, under your research objective you can see a **Topics** tab which lists and allows you to navigate through the major sections of the report.

The **Research** pane on the right of the window accepts comments up to 400 characters. You can now add comments that will be queued for revision, as they require deeper investigation by Quick Research.

Viewing citations

Throughout the report, numbered citations are present. Clicking on a citation shows you a popup window with the source article and a hyperlink to the source page.



Understanding statements

You can also click the **Understand the statement** icon (three horizontal lines and a plus) to launch an explanation window that shows how a statement in the report was determined, including a summary of the evidence and a breakdown of the conclusion.

The screenshot displays the Amazon Quick Suite interface for a research report titled "Biodegradable Food Delivery Packaging: Thermal Performance Trends 2025". The report is rated 4.5 stars and has a "Download" button. The main content area shows the "Executive Summary" and "Key Materials and Performance" sections. A sidebar on the right, titled "Understand the statement", provides a detailed breakdown of the report's findings, including evidence 1 through 9 and a conclusion. The sidebar also includes a "Close" button.

Rate report: ☆☆☆☆

Download

Biodegradable Food Delivery Packaging: Thermal Performance Trends 2025

Executive Summary

The sustainable packaging market in the food delivery industry is experiencing significant growth, valued at USD 65.73 billion in 2025 and projected to reach USD 88.46 billion by 2030, growing at a CAGR of 6.12% [+2](#).

This expansion is driven by increasing consumer demand, evolving regulatory frameworks, and technological innovations that are narrowing the performance and cost gaps between traditional and sustainable packaging options.

Key Materials and Performance

Our analysis of biodegradable materials reveals several options capable of maintaining food temperature for at least 30 minutes, with bagasse (sugarcane fiber) emerging as a particularly promising solution. Bagasse offers exceptional temperature resistance (-25°C to 220°C) [+2](#), making it suitable for hot foods while being fully biodegradable and competitively priced at approximately \$0.08 per unit [5](#). Other viable materials include PHA

Understand the statement

This expansion is driven by increasing consumer demand, evolving regulatory frameworks, and technological innovations that are narrowing the performance and cost gaps between traditional and sustainable packaging options.

Evidence 1 establishes the significant growth in the sustainable food packaging market. Evidence 2 and 3 demonstrate how regulatory changes are driving this growth through concrete policies and timelines. Evidence 4 and 5 quantify strong consumer demand, particularly among younger demographics willing to pay premium prices. Evidence 6 identifies the current cost challenges, while Evidence 7 shows how technological innovations are creating new solutions that reduce plastic usage. Evidence 8 highlights performance challenges being addressed by these innovations. Evidence 9 connects these factors to show how the combination of technological advancement and scaling production is reducing both performance and cost differences compared to traditional packaging.

Here's a breakdown of the conclusion:

- "The Sustainable Foodservice Packaging Market size is worth USD 65.73 Billion in 2025, growing at an 6.12% CAGR and is forecast to hit USD 88.46 Billion by 2030."

This data demonstrates significant market growth in sustainable food packaging, with a substantial compound annual growth rate (CAGR) of 6.12%. This rapid expansion indicates strong market forces driving adoption of sustainable packaging solutions in the food service industry.

www.mordorintelligence.com

Close

Updating research

Research is an iterative process that often requires refinement and updates based on new findings or changing requirements. Quick Research provides flexible tools to modify your research objectives, incorporate additional sources, and track changes through version management.

Making updates

After reviewing your initial research report, you can make updates to refine and improve the results. Quick Research provides multiple ways to provide feedback and request modifications to better meet your needs.

Once you've added all your comments to the queue in the Research pane, the **Revise** button becomes available in the upper right. When you choose **Revise**, you see a message that "Review revision started". Quick Research goes back to work through its research tasks with your new comments in mind, analyzing the existing report's structure and content to determine how to apply your comments.

The screenshot displays the Amazon Quick Suite Research interface. The top navigation bar includes 'Explore', 'Amazon Quick Suite', 'Research', and the current report title 'Biodegradable Food Delivery Packaging: Thermal Performance Trends 2025'. The main content area is divided into three panels:

- Research Panel (Left):** Shows the report title 'Biodegradable Food Delivery Packaging: Thermal Performance Trends 2025' with a 'Version 2' badge. Below the title is a summary: 'I need to understand emerging trends in sustainable packaging for the food delivery industry. This research will help our product team design new packaging solutions...'. Under 'Research materials', there are options for 'Web search' (disabled), 'File uploads' (No files uploaded), and 'Quick Suite assets'. A 'Pro tip' box suggests selecting the most relevant resources. At the bottom are links for 'Share feedback' and 'Report an Issue'.
- Working through research tasks Panel (Center):** A progress bar at the top shows '6%' completion. The tasks listed are:
 - Initial Report Content Exploration and Structural Analysis:** The comprehensive report reveals a robust sustainable packaging market in food delivery, valued at \$65.73 billion in 2025, with innovative biodegradable materials like bagasse and PHA emerging as cost-competitive solutions that address environmental concerns while maintaining performance across thermal, regulatory, and consumer preference dimensions.
 - Report Section Analysis for Structural Understanding and Revision Planning:** The report contains 9 sections, providing a comprehensive structural overview for targeted revisions focusing on the United States and EU regions.
 - Assessing Report Length for Potential Revision Strategy:** Completed a computational analysis that yielded a numeric output of 9824, indicating progress in the current research phase.
 - Sustainable Packaging Trends Research for US and EU Food Delivery:** (Partially visible)
- Research Panel (Right):** Shows a 'Research' section with a note: 'I'll queue all the comments to change upon revision, as they'll require deeper investigation.' Below is a 'Comment history (1)' section showing a comment: 'Focus on the United States and EU regions.' At the bottom is an 'Add a comment' input field with a character count of '400 characters remaining'.

The version increments to Version 2 when the revision is complete.

In-line feedback

Provide specific feedback on individual sections or findings within your research report by highlighting the word or section you want to comment on. In-line feedback allows you to request clarification, additional detail, or corrections for specific parts of the report without affecting the entire document.

The screenshot displays the Amazon Quick Suite interface for a research report titled "Biodegradable Food Delivery Packaging: Thermal Performance Trends 2025". The interface is divided into three main sections:

- Left Sidebar (Research):** Contains a search bar, a list of topics (Materials, Topics), and a list of research items. The selected item is "Biodegradable Food Delivery Packaging: Thermal Performance Trends 2025". Below the list, there is a section for "Executive Summary" and "Current Market Landscape of Sustainable Packaging in Food Delivery".
- Main Content Area:** Displays the "Rate report" with a star rating of 5 stars. The title is "Biodegradable Food Delivery Packaging: Thermal Performance Trends 2025". The "Executive Summary" section states: "The sustainable packaging market in the food delivery industry is experiencing growth in both the United States and European Union, valued at USD 65.73 billion in 2025 and projected to reach USD 88.46 billion by 2030, growing at a CAGR of 6.12%." It also mentions that this expansion is driven by increasing consumer demand, evolving regulatory frameworks, and technological innovations. The summary further details that the analysis of biodegradable materials reveals several options capable of maintaining food temperature for at least 30 minutes across both US and EU markets. Bagasse (sugarcane fiber) is highlighted as particularly promising, offering exceptional temperature resistance (-25°C to 220°C) and being fully biodegradable and competitively priced at approximately \$0.08 per unit. Other viable materials include PHA (Polyhydroxyalkanoates) with good thermal stability and molded fiber alternatives that balance performance with environmental benefits.
- Right Panel (Comments):** Shows "Comments (1)" and a "Research" section. The research section provides instructions on how to make the research even better, including highlighting text to add specific comments, adding overall comments, and creating the next version. Below this, there is a "Comment history (1)" section showing a comment: "Focus on the United States and EU regions."

High level feedback

Submit broader feedback about the overall direction, focus, or approach of your research. High-level feedback is useful when you want to adjust the research scope, change the emphasis, or request additional analysis across the entire report.

Research versions

Quick Research maintains versions of your research reports as you make updates and refinements. This version control system allows you to track changes, compare different iterations, and revert to previous versions if needed.

Each time you submit feedback or request updates, Quick Research creates a new version of your report while preserving the previous versions. You can view the version history, compare changes between versions, and choose which version to use as your final report.

Version management helps you maintain a clear audit trail of your research process and ensures that you don't lose valuable insights from earlier iterations of your work.

Summarizing the research

After completing your research report, you can create focused summaries for different audiences. To access the summarizing feature, choose the **Summarize** button in the upper right of your research report.

Quick Research provides multiple summary formats to help you present your findings in different contexts and for various audiences. Choose the summary type that best fits your intended use and audience needs.

The screenshot displays the Amazon Quick Suite interface. At the top, the navigation bar includes 'Explore', 'Amazon Quick Suite', 'Research', and the current report title 'Biodegradable Food Delivery Packaging: Thermal Performance Trends 2025'. The main content area is divided into two panels. The left panel, titled 'Options', contains a 'Back to research' link, tabs for 'Options' and 'Previous', and a section for 'Create a custom summary' with a text input field and a '5000 characters remaining' indicator. Below this, there are two preset summary styles: 'Executive summary' (VP-oriented tone, 2-page max, no citations) and 'General share out' (Business-friendly tone, 6-page max, essential citations). A 'Begin' button is at the bottom of the left panel. The right panel, titled 'Biodegradable Food Delivery Packaging: Ther...', shows a 'Download' button and a 'Rate report' section with five stars. The main content of the right panel is an 'Executive Summary' titled 'Emerging Trends in Sustainable Packaging for Food Delivery: Executive Summary'. It includes a 'Purpose' section stating the summary analyzes emerging trends in sustainable packaging for the food delivery industry, focusing on biodegradable materials that maintain food temperature for at least 30 minutes. The 'Key Insights' section lists five points: 1. Market Growth and Material Innovation, 2. Temperature Retention Technology, 3. Cost-Benefit Analysis, 4. Consumer Demand, and 5. Regulatory Pressure.

In the **Options** tab, you can choose to create a custom summary by entering up to 5,000 characters about how you would like to summarize, or select one of two preset summary styles:

- *Executive summary* - VP-oriented tone, 2-page maximum, no citations
- *General share out* - Business-friendly tone, 6-page maximum, essential citations

Once you choose a summary style, you can choose **Begin** to start producing the summary. The **Previous** tab shows previous summary reports. You can download the summary as PDF or Word using the **Download** button at the top.

Custom summary

Create a tailored summary by entering up to 5,000 characters about how you would like to summarize your research. Custom summaries allow you to specify the key points, length, and focus areas you want emphasized for your particular use case.

Executive summary

Generate a concise executive summary with a VP-oriented tone, limited to 2 pages maximum with no citations. This format presents the most important findings and recommendations suitable for senior leadership and decision-makers.

General share out

Create a comprehensive summary document with a business-friendly tone, up to 6 pages maximum with essential citations included. This format balances depth with accessibility for team presentations and stakeholder briefings.

Sharing the research

After completing your research, you can download it in various formats or share it with others. Amazon Quick Research provides flexible options for both downloading and sharing your research findings.

Download the research

You can export your research in multiple formats to suit different needs and workflows.

Export to PDF

PDF format is ideal for sharing formatted documents that maintain their appearance across different devices and platforms.

To export your research to PDF:

- Navigate to your completed research report

- Choose *Download* from the options menu
- Select *PDF* as the export format
- The PDF file will be generated and downloaded to your device

Export to Word

Word format allows for further editing and collaboration using Microsoft Word or compatible applications.

To export your research to Word format:

- Navigate to your completed research report
- Choose *Download* from the options menu
- Select *Word* as the export format
- The Word document will be generated and downloaded to your device

Sharing the research with stakeholders

Quick Research provides several ways to share your research with colleagues and stakeholders, enabling collaboration and knowledge distribution.

To share your research:

- Open the research report you want to share
- Choose *Share* from the options menu
- Select your preferred sharing method from the available options
- Configure sharing permissions and access levels as needed
- Send the share link or invitation to your intended recipients

Shared research maintains all formatting, citations, and interactive elements, allowing recipients to view the complete research experience.

Admin controls

Administrators can configure various settings and controls for Quick Research to ensure it meets organizational policies and security requirements.

Turn on/off web search

Administrators can enable or disable web search functionality for their organization. When web search is disabled, Quick Research will only use internal data sources such as uploaded files and connected spaces, ensuring that research remains within the organization's controlled data environment.

This control is useful for organizations with strict data governance policies or when working with sensitive information that should not be supplemented with external web sources.

Status definitions

Quick Research projects progress through various status states during their lifecycle. Understanding these status definitions helps you track the progress of your research and know what actions you can take at each stage.

- *Draft* - Research project is being created and configured but has not yet been submitted for processing
- *In Progress* - Quick Research is actively gathering and analyzing information from your selected sources
- *Review* - Initial research is complete and available for your review and feedback
- *Updating* - Quick Research is processing your feedback and generating an updated version of the report
- *Complete* - Research is finalized and ready for download, sharing, or archiving
- *Error* - An issue occurred during processing that requires attention or retry

You can view the current status of all your research projects from the main Quick Research dashboard.

Using Amazon Quick Suite automations

Amazon Quick Suite helps business users make better decisions, faster and act on them, by unifying AI agents for research, business insights, and automation into a single experience. Quick Suite offers comprehensive automation capabilities, from streamlining everyday tasks, to automating complex business processes.

Types of automations

Quick Suite automation capabilities include:

- **Quick Flows** empowers anyone to create and share multi-step workflows for routine tasks using simple natural language, without requiring any technical skills. Ideal for streamlining everyday work like troubleshooting customer tickets or generating brand-compliant content.
- **Quick Automate** tackles sophisticated enterprise automation through its multi-agent architecture, handling complex business-critical processes.

Getting started

To begin using Quick Flows or Quick Automate, familiarize yourself with the core concepts and choose the appropriate automation type for your needs. Start with simple workflows to understand the capabilities, then explore building additional automations as your expertise grows.

Your organization's administrators control access to automation features and configure available integrations. Contact your administrators if you need access to specific capabilities or encounter permission-related issues.

Using Amazon Quick Flows

Amazon Quick Flows is a capability within Amazon Quick Suite that empowers business users to transform their everyday tasks into workflows for individual and team productivity. It enables any user to create, customize, and share purpose-built workflows utilizing their data, insights, and actions available within Amazon Quick Suite. These intelligent flows can be generated from conversations with chat agents or by describing requirements in natural language, without requiring any technical skills. They can also be published to an admin-managed library and shared with other Amazon Quick Suite users in the organization.

Creators can describe a pre-defined set of steps that accomplish tasks through a predictable, reusable flow. These flows can fetch information, take action, generate content, and handle process-specific requirements through creator-specified logic.

Building on top of Amazon Quick Suite - that brings together their data, analysis, and actions securely in one place, business users can describe the repetitive tasks they need automated. With Amazon Quick Flows, business users can streamline their routine tasks like generate responses for RFPs, review statements of work (SOW), or collate the latest industry trends and insights into a sales pitch.

Amazon Quick Flows offers agentic runtime support enabling end users to chat with their workflow and request updates (like summarizing an output or skipping a step). The building experience offers simple patterns to generate responses from various sources of data, and allows creators to specify instructions and "logic" within a flow. Creators can also add the ability to repeat a subset of steps, apply if-then-else logic, validate input before generating output and so on, all through simple, natural language instructions. Amazon Quick Flows offers expanded data options to generate AI-powered responses from Amazon Quick Sight to spaces to latest updates from web search. Amazon Quick Flows empower users to process multi-modal inputs (including images and video) and refine their outputs for 'faster' or more 'balanced and versatile' responses. With the UI agent feature in Amazon Quick Flows, users can navigate public websites (that don't require a login) and perform tasks (like scroll to find information, fill forms). All this while giving IT teams peace of mind with governance controls to manage access to select features (like web search or UI agent) and gate the sharing of flows across the organization by mandating approval reviews.

What are Flows?

Amazon Quick Flows is a capability within Amazon Quick Suite that empowers business users to transform their everyday tasks into workflows for individual and team productivity. Quick Flows makes it easy to design and manage workflows through a no-code interface, so you can build workflows without writing any code.

Using Amazon Quick Flows, you can create workflows that combine processes using AI-generated responses from their data in Quick suite, web search, direct model responses or generate images. You can write agent instructions to perform if/else statements, loops and conditions and leverage UI actions to interact with websites. You can also use integrations to connect to their external services and gather User Input as text or uploaded files.

When running a flow, users get the option of interacting with their workflows as a sequence of pre-defined steps, along with the flexibility of conversational UI to chat with it to refine outputs or ask follow-up questions. Each flow consists of the following steps:

AI-generated responses

- Quick Suite data step that retrieves responses from spaces and knowledge bases.
- Quick Sight step that gets insights from your dashboards and topics.
- Web step that generates results from internet search.
- General knowledge step that generates responses directly from models
- Image steps to generate AI images from text and image inputs.

Instruct agent

- Reasoning group step that process if, when, and how to run steps.
- UI agent step to perform tasks on public websites.

Integrations

- Action step that performs read or write operations in connected systems

User input

- Text step to gather free-form text input from users.
- Files step to accept files from the user.

Why Flows?

Organizations today face increasingly complex business processes that often require both human judgment and system interactions. Amazon Quick Flows bridges this gap by combining AI reasoning capabilities with direct business actions.

Consider a typical customer support scenario where a customer submits a technical issue. Without automation, this process involves multiple manual steps: reading the ticket, researching the issue, consulting knowledge bases, possibly escalating to specialists, and finally responding with a solution.

With Amazon Quick Flows, business users can streamline their routine tasks like generate responses for RFPs, review statements of work (SOW), or collate the latest industry trends and insights into a sales pitch. You can transform this process through:

1. *Automated information gathering*: Input steps collect and process customer information when a support ticket arrives.
2. *Intelligent analysis*: Amazon Bedrock models analyze the ticket content and identify potential solutions.
3. *Seamless system integration*: Action steps connect to your knowledge bases, CRM systems, and other tools to gather context and update records.
4. *Contextual decision making*: Reasoning defines how and when the flow runs, guiding multiple steps such as applying conditions, performing validations, or running loops to determine the appropriate next step.
5. *Flexible output delivery*: Output steps provide solutions to customers in the appropriate format, whether through chat interfaces or structured responses.

Creators can describe a pre-defined set of steps that accomplish tasks through a predictable, reusable workflow. These workflows can fetch information, take action, generate content, and handle process-specific requirements through creator-specified logic.

Amazon Quick Flows offers agentic runtime support enabling end users to chat with their workflow and request updates (like summarizing an output or skipping a step). This flow reduces response times from hours to minutes while ensuring consistent quality and 24/7 availability. As your business grows, Quick Flows scales effortlessly, handling increasing volumes without proportional increases in staffing.

Beyond customer support, Amazon Quick Flows can transform numerous business processes across departments:

- *Sales, marketing, and operations*: Qualifying leads, generating personalized proposals, creating marketing content, updating CRM records, and supporting processes like RFP responses
- *HR*: Processing employee requests, answering policy questions, and automating onboarding steps
- *Finance*: Analyzing expense reports, flagging anomalies, and processing routine approvals
- *IT*: Automating troubleshooting, system monitoring, and access management

By combining AI reasoning with direct business actions, Amazon Quick Flows unlocks new levels of efficiency and scalability across your organization. For more detailed information about how Flows work, see [the section called "Key concepts"](#).

Quick Flows features and capabilities

Enterprise customers can use Amazon Quick Flows, a feature within Amazon Quick Suite, to leverage intelligent workflow capabilities. Quick Flows allows business users to predefine a set of steps across the applications and services they use to get work done, to save time and effort on repetitive business processes. Creators can build Quick Flows by providing instructions on the user defined goal in natural language. Quick Flows will apply Quick Suite level governance like guardrails and determining permissions at a user level to restrict access to data sources and tools. Quick Flows provides an intuitive visual creation experience building experience with visual flow design and step configuration.

Quick Flows runtime provides on-demand execution where users can trigger flows interactively. Runtime chat mode enables conversational experiences within Flows, while runtime structured mode provides guided interfaces for standardized processes. Creators can configure core capabilities like name, description, and set instructions for how and when steps run (reasoning) for the flow and define a sequence of steps that generate AI-powered responses from data across Quick Suite and the web, take actions in third party applications, and apply agentic reasoning to control execution of steps or even navigate websites with human-like interactions. Creators can share their Quick Flows with individuals or groups as either co-owners (who can edit and manage the flow) or viewers (who can only run the flow), or with all users who have appropriate access permissions.

Quick Flows offers a comprehensive set of features designed to streamline your repetitive tasks. Here's what you can expect:

Administration

- User Roles integrated with Amazon Quick Suite identity management
- Ability to enable/disable flows for the Quick account
- Ability to enable/disable usage of Bedrock models for output refinement in Flows
- Ability to enable/disable access to internet to enhance responses
- Ability to enable/disable access to UI agent to perform browser tasks

Creation

- Use natural language prompts (NLP) to describe the flow to create
- Create Quick Flows manually using the visual UI
- Transform conversation with chat agent into a flow
- Duplicate shared Quick Flows
- Step-level editing capabilities
- Different step types supported including file upload (default file) and text input (static text)
- Reasoning instructions (loops, validations, conditions) to specify how one or more steps should execute within a group
- Leverage General knowledge (instead of Company knowledge) (powered by Amazon Bedrock) for text and image output steps with multi-modal inputs (including images and video)
- Actions to perform read/ write operations connected applications
- Amazon Quick Sight step that generates response from dashboards and topics
- Web search to enhance responses with latest updates from the internet
- UI agent (Preview) to navigate public websites (that don't require a login) and perform tasks

Runtime

- Conversational runtime support for follow-up questions, refining outputs, and error state handling
- Structured runtime with steps
- Agentic runtime validates inputs based on constraints defined in reasoning instructions
- Visual indication of flow execution progress
- Support to view and resume past flow runs through history

Collaboration and governance

Amazon Quick Flows includes comprehensive features for team collaboration and organizational governance, giving IT teams peace of mind with governance controls to manage access to select features (like web search or UI agent) and gate the sharing of flows across the organization by mandating approval reviews.

Sharing controls

Quick Flows provides flexible sharing options to enable collaboration while maintaining organizational control:

- Share with individual users, teams, or with all users using the Quick account
- Share access can be as viewer enabling other users to run the flow or as co-owners, allowing them to edit and further share a given flow
- Ability to unpublish any shared flows and manage ownership transfer of flows

Approval workflows

Administrators can implement approval processes to ensure proper governance before flows are shared across the organization:

- Ability to enable/disable need for approval review before sharing

Getting started with Quick Flows

This section provides essential information to help you get started with Amazon Quick Flows. It covers user personas to help you identify your role, key concepts and terminology you'll encounter when working with flows, prerequisites to ensure your environment is properly configured, and an overview of the Amazon Quick Flows landing page interface.

Begin by understanding the different user personas and identifying which role best describes your responsibilities within your organization. Familiarize yourself with the core concepts and terminology that form the foundation of Quick Flows. Review the prerequisites to ensure your environment is properly configured, and explore the Amazon Quick Flows landing page to understand how to navigate the interface.

These foundational topics will prepare you to start creating and using your own flows effectively.

User personas

Amazon Quick Flows is designed to serve different types of users, each with distinct roles, responsibilities, and skill sets. Understanding these user personas helps you identify how you and your team members can best leverage Quick Flows within your organization.

Creator (Business User)

The flow creator is a business user who possesses a deep understanding of process nuances and combines business expertise with the process knowledge needed to translate repetitive tasks into streamlined workflows.

A critical aspect of the Creator role involves ensuring compliance with industry standards and regulatory requirements while validating business rules and process changes. This responsibility requires them to stay current with both regulatory developments and business process evolution within their organization.

End-User (Business User)

The End-User is a business user who performs repetitive business processes and benefits directly from the automation that Quick Flows provides. This persona represents the primary consumers of Flows created by Creators and managed by Admins.

End-Users have experience in their specific business domain with some understanding of end-to-end processes, workflows, and business rules. While they may not create Flows themselves, their domain expertise is valuable for testing and validating automated processes to ensure they meet real-world business needs.

These users act as the primary contact for business teams, providing feedback on Flow performance and identifying areas where automation could be improved or expanded. Their hands-on experience with daily business processes makes them essential contributors to the continuous improvement of automated workflows.

Admin (Technical User)

The Admin is a technical user responsible for configuring and managing governance controls within the Quick Suite. Admins manage which users have access to which capabilities within Quick Flows.

Admins have experience configuring Amazon services and possess stakeholder management skills for requirement gathering, providing updates, and managing expectations across organizational levels.

Key administrative capabilities for Quick Flows include the ability to enable or disable Quick Flows for the entire Quick account, controlling whether users can create and run automated workflows, or restrict access to Quick Flows for select users (through [the section called "Using Custom Permissions page"](#)). Admins can also manage Bedrock model usage by enabling or disabling

specific models for output refinement in Quick Flows, ensuring appropriate AI capabilities are available to users.

Additional governance controls allow Admins to enable or disable approval requirements before users can share Flows, maintaining organizational oversight of shared automation. Admins also have the authority to unpublish any shared flows and manage ownership transfer of flows between users. This is especially useful to prevent the problem of orphaned assets (here, flows) when users leave the organization.

Terminology and key concepts

Amazon Quick Flows empowers business users to transform their everyday tasks into workflows for individual and team productivity. Understanding the core terminology and concepts helps you effectively create, run, share, and maintain Flows within your organization. Flows offer agentic runtime support enabling end users to chat with their workflow and request updates like summarizing an output or skipping a step.

User inputs

User inputs represent the entry points where users provide information to initiate a flow. These steps capture the necessary context from a user that is needed to execute the flow. Input steps can include text fields for user queries or file upload capabilities for document processing. The design of input steps directly impacts user experience and determines how effectively users can communicate their needs to the flow.

When configuring user inputs, consider the type of information your flow needs to process and choose appropriate input methods that align with user expectations. Well-designed user inputs provide clear guidance about what information is required through placeholder text and some default values, if applicable.

Note

Eligible steps can be referenced in subsequent steps through the @ reference notation. Type @ in prompt to see menu of eligible steps within that flow.

AI-generated responses

AI-generated responses are outputs that define how Amazon Quick Flows presents results and information back to users. These steps encompass various formats including text responses, AI-

generated images using Bedrock models, or structured data presentations as Amazon Quick Sight visuals. Adding a specific step for AI-generated response denotes what data will be used to produce output for that step, which will be available later in the flow for further processing. These outputs can be generated from connected knowledge bases, Amazon Quick Sight dashboards and topics, spaces defined in Quick Suite, or results from the web.

Image generation capabilities allow Flows to create visual content dynamically, creating customized, consistent images for ads, marketing assets, stock photos, social media, and e-commerce content. These AI-powered image outputs can be tailored to specific requirements and integrated seamlessly with other output formats within the same Flow.

Action steps

Action steps enable Amazon Quick Flows to interact with external systems and perform automated tasks beyond content generation. These steps connect your flows to third-party applications, Amazon services, and internal systems through a comprehensive library of pre-built connectors. Action steps can perform tasks such as sending notifications, updating databases, creating calendar events, or any custom action defined with MCP or an OpenAPI specification (learn more about Quick actions in [the section called "Action steps"](#)).

The extensive action connector library supports integration with popular business applications, development tools, and Amazon services. When implementing action steps, consider authentication requirements, data security implications, and the specific permissions needed for your intended integrations.

Reasoning groups

Reasoning groups organize related processing logic within Amazon Quick Flows, allowing creators to control how the flow executes, using natural language instructions. These groups help manage the flow of information and decision-making processes by grouping related steps that work together to accomplish specific objectives. Reasoning groups improve workflow organization, making Flows easier to understand, maintain, and troubleshoot.

Within reasoning groups, you can establish dependencies between steps, control execution order, and manage data flow between different components of your workflow. This organizational structure becomes particularly valuable when building flows that require multiple processing stages or conditional logic. Instructions within reasoning groups provide specific guidance to define how the flow should interpret inputs, process information, and generate appropriate responses.

Editor and Run mode

Amazon Quick Flows provides two distinct modes for working with flows: Editor mode for building and configuring flows, and Run mode for executing and testing them. You can switch between these modes using the mode toggle in the interface.

In Editor mode, you design your flow by adding and configuring steps, defining logic, and setting up the workflow structure. This is where you build the automation before making it available to users.

Run mode allows you to execute and test your flow. Within Run mode, you can interact with your flow using two different interaction patterns: chat mode for conversational, iterative interactions where users can ask follow-up questions and refine their requests, or structured mode for guided, step-by-step workflows that follow predetermined paths and ensure consistent data collection.

Getting started with terminology

Understanding these core concepts provides the foundation for effectively using Amazon Quick Flows. As you begin building flows, refer back to these definitions to ensure you're leveraging the appropriate components for your specific use cases and organizational requirements.

Prerequisites for Quick Flows

Before you can create and use Amazon Quick Flows, you need to ensure that your Amazon Quick Suite administrator has completed the following prerequisites.

Administrator setup requirements

Your Quick Suite administrator must complete the following tasks before you can create and use Quick Flows:

- Set up and configure Quick Suite for your organization.
- (Optional) Restrict access to flows for select users using custom permissions.

For information about how to set up Quick Suite, see [Setting up and signing in](#).

Required permissions

Permissions to create, run, share, and govern flows are a result of user subscriptions and any configured custom permissions. Learn more about Quick Suite subscriptions in [the section called "Managing subscriptions"](#).

For more information about user roles and permissions in Quick Suite, see [the section called “Managing user access”](#).

Amazon Bedrock model access

Quick Flows uses Amazon Bedrock models for AI reasoning capabilities in General knowledge step. Your administrator must:

- Enable access in custom permissions for output refinement in flows using Bedrock models.

For information about managing Amazon Bedrock model access for Quick Flows, see [the section called “Response preferences”](#).

Action connector prerequisites

To use action connectors in your flows, your administrator must:

- Enable the specific action connectors that your organization needs.
- Configure authentication for each connector (2-legged OAuth or 3-legged OAuth).
- Grant appropriate permissions to users who need to use specific connectors.

For more information about action connector prerequisites and configuration, see [the section called “Action steps”](#).

Quick Suite integration prerequisites

If you plan to use Amazon Quick Sight integration in your flows, your administrator must:

- Set up and configure Amazon Quick Sight for your organization.
- Enable Amazon Quick Sight integration for Quick Flows.
- Grant appropriate permissions to users who need to access Amazon Quick Sight dashboards and visualizations.

For more information about Amazon Quick Sight integration with Quick Flows, see [the section called “Amazon Quick Sight steps”](#).

Browser requirements

Amazon Quick Flows supports the following web browsers:

- Google Chrome (latest three versions)
- Mozilla Firefox (latest three versions)
- Microsoft Edge (latest three versions)
- Apple Safari (latest three versions)

For the best experience, we recommend using the latest version of Google Chrome or Mozilla Firefox.

Availability in Amazon Regions

Quick Flows is available in all four regions with Quick Suite:

- US East (N. Virginia)
- US West (Oregon)
- Europe (Dublin)
- Asia Pacific (Sydney)

Next steps

After ensuring that all prerequisites are met, you can:

- Learn about the key concepts in Quick Flows. See [the section called “Key concepts”](#).
- Create your first flow. See [the section called “Creating flows”](#).
- Explore advanced features and capabilities. See [the section called “Editing flows”](#).

Amazon Quick Flows landing page

When you access Amazon Quick Flows, you are directed to the landing page. This page serves as your central hub for discovering, accessing, and managing all your flows.

Flow views

The Amazon Quick Flows landing page offers multiple views to help you organize and find your flows efficiently. You can switch between these views using the tabs at the top of the page.

Recently used

The Recently used view displays flows you have recently accessed or modified. This view helps you quickly return to flows you are actively working with, without having to search through all available flows.

Flows in this view are sorted with the most recently accessed flow appearing first. This view automatically updates as you interact with different flows in your environment.

My Flows

The My Flows view shows all flows that you have created or that have been shared with you specifically. This view helps you focus on flows that are most relevant to your work.

You can further organize your flows in this view by using filters or sorting options available in the interface. This makes it easier to manage your personal collection of flows as it grows over time.

All Flows

The All Flows view provides a comprehensive list of all flows available in your organization. This includes flows created by you and those created by others that have been shared with the broader organization.

This view is particularly useful when you need to discover flows created by colleagues that might be relevant to your work. You can use sorting and filtering options to navigate through the complete collection of flows.

Searching for Flows

The Amazon Quick Flows landing page includes a powerful search functionality that allows you to quickly find specific flows. The search bar is prominently displayed at the top of the landing page.

You can search for flows using various criteria including:

- Flow name
- Flow description

As you type in the search bar, the results update dynamically to show matching flows. This allows you to quickly narrow down your search and find the exact flow you need.

The search functionality works across all views, allowing you to search within Recently used, My Flows, or All Flows depending on which tab you have selected.

Managing Flows from the landing page

The Amazon Quick Flows landing page provides several options for managing your flows directly from the interface. For each flow displayed on the landing page, you can:

- Open a flow to use it or edit it
- Duplicate a flow to create a new version
- Delete a flow you no longer need

These management options are accessible through the context menu for each flow, which you can open by clicking the menu icon next to the flow name.

Creating a new flow

The landing page includes a prominent **Create flow** button that allows you to quickly start building a new flow. Clicking this button will take you to the flow creation interface where you can define the structure and functionality of your new flow.

For more information about creating flows, see [the section called “Creating flows”](#).

Creating and managing flows

This section covers the core processes for creating and managing Amazon Quick Flows. You'll learn how to create new flows, edit existing ones, manage versions, and share your flows with others in your organization.

Creating a flow involves defining its structure, components, and logic. Once created, you can edit your flows to refine their functionality or add new capabilities. Sharing enables collaboration and allows others to benefit from the flows you've created.

These topics provide comprehensive guidance on the full lifecycle of flow development and management, from initial creation through ongoing maintenance and collaboration.

Creating flows

After enabling Amazon Quick Flows in the Amazon Quick Suite management console, users with creator permissions and an appropriate Amazon Quick Suite license can create flows using several different methods. This section describes the various ways to create a flow and the requirements for each method.

Prerequisites for creating flows

Before you can create a flow, ensure that you have:

- Creator permissions in Quick Suite.
- Access to the Amazon Quick Flows feature, which must be enabled by your administrator.
- A supported web browser. For more information about browser requirements, see [the section called “Browser requirements”](#).

Creating a flow using natural language prompt (NLP)

The natural language prompt method allows you to describe your desired flow in plain language, and Amazon Quick Flows will generate a flow based on your description.

1. Sign in to the Quick Suite console.
2. In the navigation pane, choose *Flows*.
3. Choose *Create Flow*. Once you select creating a flow, you see a page where you can enter a natural language prompt or use a sample prompt.
4. In the prompt field, describe the flow you want to create. Be as specific as possible about:
 - The purpose of the flow
 - The inputs it should accept
 - The processing it should perform
 - The outputs or actions it should produce
5. Choose *Generate Flow*.
6. Review the generated flow and make any necessary adjustments.
7. Preview the flow to test its functionality and ensure it works as expected. If the flow is ready for use, you can share or publish it with individuals, groups, or all users who have appropriate access permissions.

Note

The quality of the generated flow depends on the clarity and specificity of your natural language prompt. You may need to refine your prompt or make manual adjustments to the flow to achieve your desired outcome.

Creating a flow from scratch (blank flow)

Creating a flow from scratch gives you complete control over the design and functionality of your flow.

1. Sign in to the Quick Suite web experience.
2. In the navigation pane, choose *Flows*.
3. Choose *Create Flow*.
4. Create a blank flow.
5. Enter a name and optional description for your flow.
6. In the Flow builder, add and configure the steps you need:
 - Add input steps to collect information from users.
 - Add reasoning groups to process information using AI models.
 - Add output steps to get AI-generated responses from Quick Suite data, Quick Sight dashboards and topics, Web search, UI Agent, or General knowledge.
 - Add action steps to perform operations in connected systems.
 - Add image step to generate professional grade content from text and image inputs.
7. Choose *Save* to save your flow.

For more information about the different components you can use in a flow, see [the section called "Key concepts"](#).

Creating a flow by duplicating an existing flow

Duplicating a flow allows you to copy an existing flow. This helps users to understand how the flow was created and easily create flows based on others' ideas.

1. Sign in to the Quick Suite console.
2. In the navigation pane, choose *Flows*.
3. Find the flow you want to duplicate.
4. Choose the ellipsis (:) next to the flow name, then choose *Duplicate*.
5. Enter a name and optional description for the new flow.
6. Make any necessary modifications to the duplicated flow.
7. Choose *Save* to save your flow.

Note

When you duplicate a flow, all components, connections, and configurations are copied to the new flow. However, you may need to reconfigure certain settings, such as authentication for action connectors.

Creating a flow from an agent conversation

You can create a flow directly from a conversation with Amazon Q by using the Quick Flows icon in the chat bar.

1. Sign in to the Quick Suite console.
2. Start a conversation with Amazon Quick Agent in the chat interface.
3. Discuss the flow you want to create, including its purpose, inputs, processing, and outputs.
4. When you're ready to create the flow, click the Quick Flows icon in the chat bar.
5. Choose *Create Flow from conversation*.
6. Enter a name and optional description for the flow.
7. Choose *Create*.
8. Review the generated flow and make any necessary adjustments.
9. Choose *Save* to save your flow.

This method is particularly useful when you want to explore ideas and possibilities through conversation before creating a flow.

Flow creation requirements

When creating a flow, keep the following requirements in mind:

Flow name

Flow names must be unique within your Quick Suite account and can contain up to 128 characters.

Flow components

A valid flow must have at least one input step and one output step or action step.

Reasoning groups

Each reasoning group must have at least one instruction that defines what the AI model should do with the inputs.

Action steps

Action steps require proper authentication and configuration before they can be used in a flow.

Next steps

After creating a flow, you can:

- Test your flow to ensure it works as expected. .
- Publish your flow to make it available to users. See [the section called "Publishing your flow"](#).
- Share your flow with specific users or groups. See [the section called "Sharing flows"](#).
- Monitor your flow's usage and performance. See [the section called "What is a progress tracker?"](#).

Editing flows

Once you've created your flow, you can edit and configure it to meet your specific requirements. This section describes how to modify your flow's components, structure, and settings.

Accessing the Flow editor

To edit an existing flow:

1. Sign in to the Amazon Quick Suite console.
2. In the navigation pane, choose *Flows*.
3. Find the flow you want to edit.
4. Choose the flow name to open it in the Flow editor.

The Flows editor provides a visual interface where you can modify your flow's components, connections, and settings.

Configuring step types

Each type of step in a flow has specific configuration options. This section describes how to configure each step type.

Configuring input text steps

Input text steps collect text input from users. To configure an input text step:

1. Select the input text step in the Flow editor.
2. In the configuration panel, set the following options:
 - *Label*: The text that appears above the input field.
 - *Placeholder*: Optional text that appears inside the input field when it's empty.
 - *Default value*: Optional text that pre-fills the input field.
 - *Required*: Toggle to specify whether the input is required.
3. Choose *Save* to save your changes.

Configuring file upload steps

File upload steps allow users to upload files to your flow. To configure a file upload step:

1. Select the file upload step in the Flow editor.
2. In the configuration panel, set the following options:
 - *Label*: The text that appears above the upload control.
 - *Upload default file*: Upload a default file that can be used during flow run.
 - *Allow override of default files*: Enable runtime users to override default files.
3. Choose *Save* to save your changes.

Note

File uploads are processed according to the capabilities of the selected model. Some models have limitations on file types and sizes they can process.

Configuring general knowledge steps

General knowledge steps display text responses from models to users. To configure a general knowledge step:

1. Select the general knowledge step from the add step menu in the Flow editor.
2. In the configuration panel, set the following options:

- *Output preference*: Choose output response preference from Faster responses or Versatility and Performance.
- *Prompt*: Write the prompt that instructs the model what to generate. You can use @ references to include data from previous steps.
- *Advanced settings*: Configure model specific parameters such as creativity level to manage randomness of the LLM response.

3. Choose *Save* to save your changes.

Configuring Quick suite data steps

Quick suite data steps display text responses from internet search to users. To configure a Quick suite data:

1. Select the Quick suite data step from the add step menu in the Flow editor.
2. In the configuration panel, set the following options:
 - *Prompt*: Write the prompt that instructs what content to generate from the web. You can use @ references to include data from previous steps.
 - *Link specific resources*: Select spaces and knowledge bases that you would like to get insights from. By default, responses are generated from all knowledge sources user has access to.
3. Choose *Save* to save your changes.

Configuring web steps

Web steps display text responses from internet search to users. To configure a web step:

1. Select the web step from the add step menu in the Flow editor.
2. In the configuration panel, set the following options:
 - *Prompt*: Write the prompt that instructs what content to generate from the web. You can use @ references to include data from previous steps.
3. Choose *Save* to save your changes.

For more information about writing effective prompts, see [the section called “Prompt writing for output steps”](#).

Configuring output image steps

Output image steps generate and display images to users. To configure an output image step:

1. Select the output image step in the Flow editor.
2. In the configuration panel, set the following options:
 - *Prompt*: Write the prompt that describes the image to generate. You can use @ references to include data from previous steps.
 - *Advanced settings*: Configure mode-specific parameters such as creativity level (defines the randomness of the LLM response), Exclude (a parameter to define what not to include), and image seed (control the determinism of the images generated).
3. Choose *Save* to save your changes.

Configuring output Quick Sight steps

Output Quick Sight steps display Quick Sight visualizations to users. To configure an output Quick Sight step:

1. Select the output Quick Sight step in the Flow editor.
2. In the configuration panel, set the following options:
 - *Quick Sight source*: Choose from Dashboard or a Topic.
 - *Prompt*: Describe insights you want to get from your Quick Sight dashboard or topic. You can use @ references to include data from previous steps.
3. Choose *Apply* to save your changes.

For more information about integrating Quick visualizations in your flows, see [the section called "Amazon Quick Sight steps"](#).

Configuring action steps

Action steps perform operations in connected systems. To configure an action step:

1. Select the action step in the Flow editor.
2. In the configuration panel, set the following options:
 - *Action connector*: Select the connector to use (e.g., Salesforce, Jira, Slack).
 - *Action*: Select the specific action to perform.

- *Prompt*: Write prompt instructions to execute your actions. You can use @ references to include data from previous steps.
3. Choose *Apply* to save your changes.

For more information about configuring action steps, see [the section called "Action steps"](#).

Configuring reasoning groups

Reasoning groups process information using AI models. To configure a reasoning group:

1. Select the reasoning group in the Flow editor.
2. In the configuration panel, set the following options:
 - *Instructions*: Write the instructions that tell the model what to do with the inputs. You can use @ references to include data from previous steps.
3. Choose *Apply* to save your changes.

Prompt writing for output steps

Writing effective prompts is essential for getting the desired results from output steps. This section provides guidance on writing prompts and using @ references.

Prompt writing basics

When writing prompts for output steps, consider the following best practices:

- Be clear and specific about what you want the model to generate.
- Provide context to help the model understand the task.
- Specify the desired format, tone, and style of the output.
- Use examples to illustrate the expected output when appropriate.

Example prompt for a customer support response:

```
You are a helpful customer support agent for a software company.  
Write a response to the customer's inquiry below.  
Be professional, empathetic, and solution-oriented.  
Include specific steps the customer can follow to resolve their issue.
```

```
Customer inquiry: @{input_text}
```

Using @ references

@ references allow you to include data from previous steps in your prompts. To use an @ reference:

1. In the prompt field, type the @ symbol (@).
2. A dropdown menu will appear showing available references from previous steps.
3. Select the reference you want to include.
4. The reference will be inserted in the format @{step_id}.

You can use @ references in various ways:

- Include user input: @{input_text}
- Include file content: @{file_upload}
- Include action results: @{action_step}

Example Prompt with multiple @ references

```
Analyze the customer's message: @{customer_input}
```

```
Consider the customer's account information:
```

- Account type: @{account_info.type}
- Subscription status: @{account_info.status}
- Support level: @{account_info.support_level}

```
Based on this information, provide a personalized response that addresses the customer's concerns and offers appropriate solutions.
```

Adding and removing steps

You can add new steps to your flow or remove existing ones as needed.

Adding steps

To add a new step to your flow:

1. In the Flow editor, choose *Add step* from the toolbar.
2. Select the type of step you want to add from the dropdown menu.
3. Drag the step to the desired position in your flow.
4. Configure the step as needed.
5. Connect the step to other steps in your flow using @ reference of existing steps.

Removing steps

To remove a step from your flow:

1. Select the step you want to remove.
2. Choose *Delete* from the context menu.
3. Confirm the deletion when prompted.

Note

When you remove a step, any connections to and from that step are also removed. You may need to reconnect other steps to maintain the flow of your application.

Sequential flow

Quick Flows uses a sequential flow model, where steps are run in a specific order based on their connections.

To create a sequential flow:

1. Arrange your steps in the order you want them to run.
2. Connect each step to the next by dragging a connection from the output port of one step to the input port of the next.
3. Ensure that all steps are connected in a logical sequence, with no disconnected steps.

The sequential flow determines:

- The order in which steps are presented to users
- The data flow between steps

- The availability of @ references from previous steps

Quick Suite data vs general knowledge

When configuring reasoning groups and output steps, you can choose between using Quick Suite data and general knowledge.

Knowledge sources

Quick Suite data

Uses your organization's knowledge base to provide responses based on your company's specific information, documents, and data.

General knowledge

Uses the model's built-in knowledge to provide responses based on general information available during the model's training.

Output preference details

Instead of selecting specific models, you can choose output preferences that optimize the AI response for your specific needs. When configuring output preferences, consider:

- *Faster responses*: Optimized for speed, providing quicker results when time is critical for your workflow.
- *Versatility and Performance*: Balanced approach that handles a wide range of tasks effectively across different use cases.

The system automatically selects the most appropriate Amazon Bedrock model based on your chosen preference and the specific requirements of your flow. For more information about output preferences and model abstraction, see [the section called "Response preferences"](#).

Spaces details

Spaces are containers for company knowledge that can be used in your flows. You can create and configure spaces with specific knowledge sources (if you have Author Pro or Reader Pro tier access), such as:

- Document repositories

- Wikis and knowledge bases
- Databases and structured data
- Custom data sources

When configuring a reasoning group or output step, you can select which space to use as the knowledge source. This determines what information is available to the model when processing inputs and generating outputs.

Publishing changes

After making changes to your flow, you need to publish them to make them available to users.

1. In the Flow editor, choose *Save* to save your changes.
2. Choose *Publish* to publish your changes.
3. Choose *Publish* to confirm.

When you publish changes to a flow:

- The changes become immediately available to all users who have access to the flow.

Updating Flow details

You can update your flow's title, description, and view the original prompt used to create it (if applicable).

Updating title and description

To update your flow's title and description:

1. In the Flow editor, directly edit title and description of your flow in-line.
2. Choose *Save* to apply your changes.
3. Choose *Publish* to publish your changes.

Viewing the original prompt

If your flow was created using a natural language prompt, you can view the original prompt:

1. In the Flow editor, select the kebab menu in the header.

2. From the options displayed in the menu, choose *View prompt*.
3. The original prompt used to create the flow is displayed.

Note

The original prompt is read-only and cannot be modified. If you want to create a new Flow based on a modified prompt, you can create a new flow using the natural language prompt method.

Best practices for editing flows

Consider these best practices when editing your flows:

- Test your changes thoroughly before publishing them to ensure they work as expected.
- Use clear and descriptive names for steps to make your flow easier to understand and maintain.
- Write detailed prompts and instructions to get the best results from AI models.
- Use @ references to create dynamic flows that adapt to user inputs.
- Consider the user experience when designing the flow of steps.
- Document your changes in the publication description to maintain a clear history of updates.

Versioning

Versions help you manage your Amazon Quick Flows as you build, test, and share them with others. Think of versions like saving different copies of your work - you can have a working copy while keeping a clean, published version for others to use. Amazon Quick Flows keep it simple with no more than three versions for any given flow.

How versioning works

Your organization's setup determines which version types you'll see:

If your organization uses approval review

You'll work with three types of versions. For detailed information about how approval reviews work, see the [Amazon Quick Flows - Approval reviews documentation](#).

Draft

What it is: your working copy where you build and test your flow. You can continue to run Drafts as private flows without ever sharing them.

What you can do:

- Edit and modify your flow freely
- Add cards like Amazon Quick Sight cards, file upload cards, and output cards
- Test your changes before sharing
- Save your work automatically as you go
- Create multiple drafts to try different approaches

Pending Approval

What it is: this version of the flow is waiting for an eligible user to review and approve it for sharing.

What you can do:

- View your submitted flow (but can't edit it)
- Check the approval status
- Respond to reviewer feedback
- Wait for approval or address requested changes before re-submitting it for review

Published

What it is: Your approved flow that others can now use

What you can do:

- See your live, published flow
- Use it as a starting point for new drafts

If your organization doesn't use approval review

You'll work with two types of versions:

Draft

What it is: Your private workspace for building flows

What you can do:

- Edit and test your flow
- Save changes automatically
- Keep your work private until ready to share

Published

What it is: Your published flow that others can use immediately

What you can do:

- Publish directly when your flow is ready
- Make it available to users right away
- Use it as a baseline for future changes

Key versioning concepts

Understanding these core concepts will help you work effectively with Amazon Quick Flows versioning:

Multi-versioning for creators: As a flow creator, you can see and work with multiple versions of your flow (draft, pending approval, shared). This allows you to continue developing while users access the stable published version.

Single version for users: End users always see only one version of your flow - the currently published version. This prevents confusion and ensures everyone has the same experience.

Publish changes workflow: When you make changes to an existing shared flow, you must "publish changes" to make those updates available to users. This creates a new version and may trigger approval review.

Version replacement: Each time you publish changes, the new version replaces the previous published version for all users. Historical versions are maintained for creators but users always see the latest approved version.

App definition versioning: Changes to core flow elements like title, description, and structure are considered app definition changes and require going through the full publishing process.

Creating and editing flows

All new flows begin as drafts in your private workspace. You can build and test your flow, adding cards, configuring settings, and iterating as needed. Your work saves automatically, but you can

manually save important milestones. When you're ready to share your flow with others, you'll need to "publish changes" to make it available to your intended audience.

Publishing your flow

With approval review: When your draft is ready, you "publish changes" to submit it for review. Your flow moves to "Pending Approval" status where reviewers can evaluate it. If changes are requested, you'll need to create a new draft with updates and resubmit. Once approved, your flow becomes "Published" and available to users. Any future changes to a shared flow require republishing and going through the approval process again.

Without approval review: When your draft is ready, you "publish changes" to share it immediately. Your flow becomes "Published" and available to users right away. You can continue making changes in draft mode and publish updates as needed.

Making changes to shared flows

Once you've shared a flow, any changes you make require going through the "publish changes" process to make them available to users. This ensures that users always have access to a stable, tested version while you continue developing improvements.

The publish changes workflow

When you modify a shared flow, you're working in a new draft version. Users continue to see and use the current published version until you complete the publish changes process. This approach prevents users from seeing incomplete or untested changes.

Title and description changes: Updates to your flow's title or description are considered app definition changes. These require publishing changes and, if approval review is enabled, going through the approval process again.

Content changes: Modifications to cards, prompts, or flow logic also require publishing changes. Each published update creates a new version that replaces the previous version for all users.

Sharing permission changes: Adding new users or changing permissions (such as converting viewers to co-owners) may require approval depending on your organization's settings.

Unpublishing flows

You can unpublish a shared flow to remove it from general availability while keeping it accessible to co-owners. When you unpublish a flow:

- **Viewers lose access:** Users with viewer permissions can no longer see or use the flow
- **Co-owners retain access:** Co-owners can still access and work with the unpublished flow
- **Draft state:** The flow returns to draft state for further development
- **No approval required:** Unpublishing does not require going through approval review

This feature is useful when you need to make significant changes to a widely-shared flow or temporarily remove access while addressing issues.

Tips for success

Follow these best practices to ensure your flows work reliably and provide a great experience for your users.

Before you submit or publish

- # Test your flow thoroughly with different scenarios and inputs
- # Check that all steps work as expected
- # Verify data permissions are set up correctly for your intended users
- # Add clear descriptions and help text for users
- # Review any error handling or edge cases

Managing multiple versions

- Keep track of what changes you made in each version
- Use descriptive names or notes to identify different versions
- Don't delete important drafts until you're sure you won't need them

Working with your team

- Communicate with teammates about who's working on what
- Coordinate timing for publishing shared flows
- Let users know when important flows are being updated
- Ensure proper access permissions are in place before sharing

Version reference

Use this table as a quick reference to understand what you can do with each version type and what the next steps are in your workflow.

#	Version Type	Can Edit?	Visible to Users?	Next Step
1	Draft	# Yes	# No	Submit/Publish
2	Pending Approval	# No	# No	Wait for approval
3	Shared	# No	# Yes	Create new draft for changes

Sharing flows

Sharing allows you to make your flows available to other people in your organization. When you share a flow, you're giving others permission to view, run, and potentially collaborate on your work. Think of sharing as publishing your flow - it moves from being a private draft that only you can see to being accessible by the people or groups you specify.

Sharing is essential for collaboration, knowledge sharing, and making your flows useful to your broader team or organization.

How sharing works

Amazon Quick Flows offers flexible sharing options that let you control exactly who can access your flows and what they can do with them. Your sharing choices depend on your role and your organization's approval review settings.

Types of sharing

You can share flows in three different ways, each serving different collaboration needs:

Private sharing with individuals

What it is: Share your flow with specific people using their email addresses

Best for: Testing with colleagues, sharing with specific team members, or controlled distribution

Who can do this: Author, Admin, Author Pro, and Admin Pro users

Private sharing with groups

What it is: Share your flow with pre-defined groups in your organization

Best for: Department-wide sharing, team collaboration, or role-based access

Who can do this: Author, Admin, Author Pro, and Admin Pro users

Share with everyone

What it is: Make your flow available to all users in your Amazon Quick Flows instance

Best for: Organization-wide tools, widely useful flows, or public resources

Who can do this: Author, Admin, Author Pro, and Admin Pro users (may require approval)

Private sharing with individuals

Private sharing with individuals gives you precise control over who can access your flow. This is perfect for collaboration with specific colleagues or when you want to test your flow with a select group before broader distribution.

How to share with individuals

- **Complete your flow** - Ensure it's ready for others to use
- **Click the share button** - Access sharing options from your flow
- **Add email addresses** - Enter the full email addresses of people you want to share with
- **Choose permissions** - Select whether they should be viewers or co-owners
- **Add a message** (optional) - Include context about your flow
- **Share the flow** - Send invitations to your selected users

What happens when you share

- **Recipients get notified** - They'll receive an email notification about the shared flow
- **Flow appears in their library** - They can find it in their Amazon Quick Flows library
- **Access is immediate** - They can start using the flow right away (unless approval is required)
- **You maintain control** - You can modify sharing permissions or remove access at any time

Best practices for individual sharing

- **Use full email addresses** - Ensure you have the correct email for each person
- **Include context** - Add a brief message explaining what the flow does and why you're sharing it
- **Start small** - Test with a few people before sharing more broadly
- **Check permissions** - Verify that shared users have access to any required data sources

Private sharing with groups

Group sharing allows you to share flows with pre-defined groups in your organization, making it easy to collaborate with entire teams or departments without managing individual email addresses.

Understanding groups

Groups are collections of users that your organization's administrators have set up in your identity system. Common examples include:

- Department groups (Marketing, Sales, Engineering)
- Project teams (Project Alpha, Q4 Initiative)
- Role-based groups (Managers, Analysts, Executives)
- Location-based groups (Seattle Office, Remote Workers)

How to share with groups

- **Access sharing options** - Click the share button in your flow
- **Select group sharing** - Choose the option to share with groups
- **Find your group** - Search for or select the appropriate group from the list
- **Set permissions** - Choose viewer or co-owner access for the group
- **Add context** - Include a message about the flow's purpose
- **Share with the group** - Send the flow to all group members

Share with everyone

Sharing with everyone makes your flow available to all users in your Amazon Quick Flows instance. This is the broadest form of sharing and is ideal for flows that provide value to your entire organization.

When to share with everyone

Consider organization-wide sharing when your flow:

- **Solves common problems** - Addresses needs that many people have
- **Provides general utility** - Offers broadly useful functionality
- **Represents best practices** - Demonstrates good flow design for others to learn from
- **Supports organizational goals** - Helps achieve company-wide objectives

How to share with everyone

- **Prepare your flow** - Ensure it's polished and ready for wide distribution
- **Access sharing settings** - Click the share button in your flow
- **Select "Share with all"** - Choose the organization-wide sharing option
- **Add description** - Provide clear information about what the flow does
- **Submit for approval** (if required) - Wait for approval if your organization requires it
- **Monitor and maintain** - Keep the flow updated and respond to user feedback

Viewer permissions

Viewer permissions give people access to use your flow without the ability to modify it. This is the default permission level and is appropriate for most sharing scenarios.

What viewers can do

- **Run the flow** - Execute the flow with their own inputs
- **View results** - See outputs and generated content
- **Access the flow** - Find it in their library and use it anytime
- **See flow details** - View descriptions and help information

What viewers cannot do

- **Edit the flow** - Cannot modify cards, prompts, or flow structure
- **Change sharing** - Cannot add or remove other users
- **Delete the flow** - Cannot remove the flow from the system
- **See draft versions** - Only see the published version

Co-owner permissions

Co-owner permissions give people full access to your flow, including the ability to edit, share, and manage it. This enables true collaboration but requires trust and coordination.

What co-owners can do

- **Edit the flow** - Modify cards, prompts, and flow structure
- **Share with others** - Add new viewers or co-owners
- **Manage permissions** - Change or remove access for other users
- **Delete the flow** - Remove the flow entirely (except for the original owner)
- **See draft versions** - Access unpublished changes and work-in-progress
- **Publish changes** - Share updates with all flow users

What co-owners cannot do

- **Remove the original owner** - The person who created the flow always retains access
- **Override each other** - Only one person can edit at a time to prevent conflicts

Tips for success

Following these best practices will help you share flows effectively and maintain good collaboration with your colleagues.

Choosing the right sharing approach

- **Start narrow** - Begin with individual or group sharing before going organization-wide
- **Match audience to content** - Share broadly only when flows have wide applicability

- **Consider approval time** - Factor in approval delays when planning to share with everyone
- **Think about maintenance** - Ensure you can support the level of sharing you choose

Setting appropriate permissions

- **Default to viewer** - Most users only need to run flows, not edit them
- **Be selective with co-owners** - Only grant editing access when truly needed
- **Plan for coverage** - Have multiple co-owners for critical flows
- **Review regularly** - Audit permissions periodically to ensure they're still appropriate

Maintaining shared flows

- **Keep descriptions current** - Update help text when you change functionality
- **Test after changes** - Verify flows still work for all users after modifications
- **Communicate updates** - Let users know about significant changes or improvements
- **Gather feedback** - Ask users about their experience and suggestions for improvement

Sharing and permissions reference

Use these tables to understand sharing options and permissions:

Sharing Types

#	Sharing Type	Audience	Approval Required?	Best For
1	Individual	Specific people	Sometimes*	Testing, targeted collaboration
2	Group	Pre-defined groups	Sometimes*	Team/department sharing
3	Everyone	All users	Often*	Organization-wide tools

*Depends on your organization's approval review settings

Permission Levels

#	Permission Level	Can Run?	Can Edit?	Can Share?	Best For
1	Viewer	# Yes	# No	# No	Most users, tool consumption
2	Co-owner	# Yes	# Yes	# Yes	Collaboration, shared ownership

Approval review

Approval reviews provide administrators with governance control over how flows are shared within your organization. When approval review is enabled, all users must submit their flows for review and approval before they can be shared with others. This helps prevent flow sprawl, ensures quality control, and maintains oversight of what flows are available to your users.

Think of approval reviews as a quality gate - it allows your organization to review flows before they become available to everyone, ensuring that only appropriate, well-designed flows are shared broadly.

How approvals work

Your organization's administrators decide whether to enable approval review. Once enabled, it applies to all flows submitted for sharing, requiring an approval before any other users see them in their library.

If your organization uses approval reviews

When approval reviews are enabled, you'll work with a three-stage approval cycle:

Submitted

What it means: Your flow has been submitted for review and is waiting for approval

What you can do:

- View the submitted flow (but cannot edit it)

- Check the approval status
- Add comments about what changes were made
- Withdraw the request if needed
- Wait for reviewer feedback

Rejected

What it means: Your flow was reviewed but not approved for sharing

What you can do:

- See the reason for rejection (if provided by the reviewer)
- Make necessary changes to address feedback
- Resubmit the updated flow for approval
- View rejection comments from the reviewer

Approved

What it means: Your flow has been approved and is now shared with your intended audience

What you can do:

- See your flow available to users
- View approval details and who approved it
- Make new changes (which will require a new approval cycle)

Admin Pro and Author Pro view in the library

Administrators and Author Pro users have special access to manage the approval process:

Note

This review ability is not provided to users with promotional Enterprise access (Authors). Only users with full Admin Pro or Author Pro subscriptions can review and approve flows.

Review flows tab

Purpose: Central location to review all submitted flows

What you can see:

- All flows submitted for approval
- Flow details and what changes were made
- Comments from flow creators about their changes
- Approval history and status

Approval actions

- **Approve flows:** Accept flows for sharing with intended audiences
- **Reject flows:** Decline flows with optional feedback for improvement
- **Bulk actions:** Approve or reject multiple flows at once
- **View details:** See what specific changes were made or who is being added

Enabling and disabling approval review

Administrators have full control over when approval review is active in your organization.

Enabling approval review

When administrators enable approval review:

- All future Flow sharing requires approval
- Existing shared flows remain accessible
- Users can still create and edit flows in draft mode
- Only Admin Pro and Author Pro users can approve flows

Disabling approval review

When administrators disable approval review:

- All pending approval requests are automatically rejected
- Users can immediately share flows without approval
- Previously approved flows remain shared
- The approval process is bypassed for all future sharing

Configuration options

Administrators can configure approval review to require approval for:

- Sharing with individuals
- Sharing with groups
- Sharing with everyone in the organization
- Publishing changes to existing shared flows
- Adding new users to existing flows

Working with approval reviews - step by step

Understanding the approval process helps you navigate it efficiently and get your flows approved quickly.

For flow creators

Submitting for approval

- **Complete your flow** - Ensure it's fully built and tested
- **Add sharing details** - Specify who should have access
- **Include change notes** - Explain what you built or changed
- **Submit for review** - Send to the approval queue
- **Monitor status** - Check for updates and feedback

If your flow is rejected

- **Review feedback** - Read the reviewer's comments carefully
- **Make necessary changes** - Address the specific concerns raised
- **Test your updates** - Ensure changes work as expected
- **Resubmit with notes** - Explain how you addressed the feedback
- **Follow up if needed** - Contact reviewers for clarification

For reviewers (Admin Pro and Author Pro users)

Reviewing submitted flows

- **Access the review tab** - Find submitted flows in the library
- **Examine the flow** - Test functionality and review design
- **Check sharing scope** - Verify who will have access
- **Review change notes** - Understand what the creator built or modified
- **Make a decision** - Approve or reject with clear feedback

Best practices for reviewers

- **Be timely** - Review submissions promptly to avoid delays
- **Be specific** - Provide clear, actionable feedback for rejections
- **Be consistent** - Apply the same standards across all reviews
- **Be collaborative** - Work with creators to improve flows rather than just rejecting

Approval review states reference

Use this table to understand the approval review states and what actions are available:

#	Flow State	Can Edit?	Visible to Users?	Available Actions
1	Draft	# Yes	# No	Submit for approval
2	Submitted	# No	# No	View status, withdraw request
3	Rejected	# Yes (after changes)	# No	Make changes, resubmit
4	Approved	# No	# Yes	Create new version for changes

Flow components and features

This section details the specialized components and features that you can use to enhance your flows. These capabilities allow you to customize your flows with advanced AI models, data visualizations, external actions, web search capabilities, and UI agent functionality.

Amazon Bedrock models provide powerful AI reasoning capabilities that you can configure to meet your specific needs. Amazon Quick Suite integration enables you to incorporate data visualizations and analytics into your flows. Action steps allow your flows to connect with external systems and perform automated tasks. Web search capabilities enable your flows to retrieve information from the internet, and UI agent functionality provides specialized interface options.

Understanding these components and features will help you build more capable flows that can address complex business requirements.

Using response preferences in General knowledge step

This guide covers how to configure response preferences to refine and optimize your outputs in Amazon Quick Flows, providing flexibility in response optimization based on your specific use case requirements.

Key highlights

Simplified model selection

Flow builders get a benefit-based preference selection for their output refinement where they can choose from 2 modes - Faster responses or Versatility and performance. This reduces cognitive load for builders and creators can focus on their objectives rather than technical model comparisons.

Intelligent model selection in runtime

Depending on your output preference, flows service will automatically select the most appropriate model based on real-time context size, task and multi-modal requirements.

Modality supported for general knowledge step

Input: Text/document files, Image or Video, output: Text. Users can upload up to 50 MB of document files, 1GB of video files, and 4.5 MB of image files as inputs.

Getting started: Response preferences in flows

When building flows in Amazon Quick Flows, you can select response preferences to optimize performance for your specific use case. The response preference interface allows you to choose the most appropriate optimization based on your requirements for speed, versatility, and performance.

To select response preferences:

1. Navigate to your flow configuration
2. Add a General knowledge step
3. Access the response preference options
4. Choose from Faster responses or Versatility and Performance
5. Configure additional settings as needed

Configuring output types: Text vs Image

Different Amazon Bedrock models support various output formats. Configure your output type based on your application needs:

Text outputs

Text outputs are optimized for natural language generation and support both structured and unstructured text with variable length responses based on model capabilities.

Image outputs

Image outputs provide visual content generation capabilities with support for various image formats and resolutions, including integration with text prompts for image generation.

Advanced model settings: Creativity slider, Exclude, and Seed

Fine-tune model behavior using advanced configuration options:

Creativity slider

The creativity slider controls the randomness and creativity of model outputs. Lower values produce more deterministic results, while higher values increase variability and creative responses.

Exclude settings

Exclude settings allow you to specify content or patterns to exclude from image outputs, helping maintain content guidelines and restrictions with customizable filtering based on your requirements.

Seed configuration

Seed configuration enables reproducible outputs for testing and consistency. Use specific seed values to generate consistent results, which is useful for debugging and quality assurance workflows.

Multi-modality support using Amazon Bedrock models

Leverage models that support multiple input and output modalities:

- **Text-to-text:** Traditional language model interactions
- **Text-to-image:** Generate visual content from text descriptions
- **Image-to-text:** Extract information or descriptions from images
- **Multi-modal combinations:** Process both text and image inputs simultaneously

File uploads using general knowledge

Amazon Quick Flows supports various file types and processing capabilities with Amazon Bedrock models. Supported formats include documents, images, and structured data files with processing options to extract text, analyze content, or generate summaries. Integration workflows seamlessly incorporate file content into model prompts, though you should refer to model-specific file size restrictions.

Total context limit supported for Amazon Bedrock models

Understanding context limitations helps optimize your applications. Context window sizes vary by model type and version, so monitor input and output token usage. Use optimization strategies and techniques for working within context limits while balancing context size with response speed for performance considerations.

Note: If you don't see response preferences, contact admin

If the response preference options are not visible in your interface:

- Verify your user permissions and access levels
- Contact your system administrator to ensure "Enable bedrock model usage in General knowledge step for output refinement" is enabled
- Ensure you're using the latest version of the Amazon Quick Flows interface

For additional support and configuration details, administrators can refer to the comprehensive capabilities documentation.

Amazon Quick Sight steps in flows

The Amazon Quick Sight card enables flow creators to generate outputs using structured data that are already available within your customer's Amazon Quick Suite instance. Using the Quick Sight card in flows, users can leverage existing Dashboards and Topics to create multi-visual answers that provide comprehensive data insights.

Overview

The Quick Sight card integration brings the power of your existing Quick Sight resources directly into Amazon Quick Flows. This functionality allows you to seamlessly incorporate data visualizations and analytics into your flow-based applications. By connecting to your established Quick Sight environment, you can create rich, data-driven responses that combine multiple visual elements.

Configuring Amazon Quick Sight card

Configuration of the Quick Sight card requires establishing the connection between your flow and your Quick Suite instance. Begin by selecting the Quick Sight card from the available Flow components and specify the target Quick Suite resources you want to access.

Connection setup

Establish the connection to your Quick Sight instance by providing the necessary authentication credentials and specifying the Amazon region where your Quick Sight resources are located. The card will automatically discover available dashboards and topics within your authorized scope.

Resource selection

Choose from your existing Quick Sight resources including dashboards and topics. The card provides a browsable interface that displays your available resources, making it easy to select the specific data sources and visualizations you want to incorporate into your flow.

Data source configuration

Configure how the card will interact with your selected Quick Sight resources. Specify parameters such as filters, date ranges, and other criteria as natural language prompts in your Quick Sight card configuration that will determine which data gets included in your flow outputs. This configuration ensures that the generated responses are relevant and targeted to your specific use case.

Running Amazon Quick Sight card

Running the Quick Sight card executes queries against your configured Quick Sight resources and generates structured outputs based on the available data. The functionality provides intelligent responses that combine multiple data sources and visualizations.

Query execution

When the Quick Sight card runs, it processes the input parameters and executes the appropriate queries against your selected Quick Sight resources. The card intelligently determines which dashboards or topics are most relevant to the current Flow context and retrieves the corresponding data.

Multi-visual response generation

The card generates comprehensive responses that can include multiple visual elements from your Quick Sight environment. These responses combine charts, graphs, tables, and other visualizations to provide complete answers to user queries. The multi-visual approach ensures that users receive rich, contextual information that addresses their specific needs.

Data processing and formatting

Retrieved data undergoes processing to ensure compatibility with your flow's output requirements. The card handles data transformation, formatting, and presentation automatically, converting Amazon Quick Suite visualizations into formats suitable for your flow's target audience and use case.

Performance optimization

The Quick Sight card includes built-in performance optimization features that manage query execution and data retrieval efficiently. These optimizations ensure that your flows respond quickly while maintaining access to comprehensive Quick Sight data and visualizations.

Integration benefits

The Quick Sight card integration provides several key advantages for Flow creators. You can leverage existing investments in Quick Sight dashboards and Topics without recreating data connections or visualizations. This approach reduces development time and ensures consistency with your organization's established data presentation standards.

The card also enables dynamic data access, allowing your flows to provide up-to-date information from your Quick Sight environment. Users receive current data insights without manual intervention, ensuring that flow responses remain accurate and relevant over time.

Use cases

Common use cases for the Quick Sight card include executive reporting flows that combine multiple dashboard elements into comprehensive briefings, operational monitoring flows that surface key performance indicators from various dashboards, and customer-facing applications that provide data-driven insights using your organization's established Quick Sight resources.

Action steps in flows

Action connectors enable flows to integrate with external systems and services, extending the capabilities of your Amazon Quick Flows beyond data retrieval and analysis. These connectors allow your flows to perform actions such as creating records, sending notifications, updating databases, and triggering workflows in connected applications.

Amazon Quick Suite Actions supports hundreds of actions across multiple third-party systems, enabling comprehensive automation and integration capabilities.

Prerequisites for adding action connectors

Administrators must configure action connectors in the Amazon Quick Suite console before they become available to Amazon Quick Flows creators. Only plugins and associated actions configured by administrators will be accessible for use in Flow creation. This includes third-party connectors, MCP connectors, and custom API connectors.

System requirements

Your Amazon Quick Flows environment must have the appropriate permissions and network connectivity to communicate with target systems. Verify that firewall rules, security groups, and network access control lists allow outbound connections to the services you plan to integrate.

Additionally, confirm that your Amazon Quick Flows instance has sufficient resources to handle the additional processing load that action connectors may introduce.

Authentication credentials

Obtain the required authentication credentials for each service you plan to connect. This typically includes API keys, client IDs, client secrets, or other service-specific authentication tokens. Store these credentials securely using your organization's approved credential management system, ensuring they are accessible to your Amazon Quick Flows environment while maintaining security best practices.

Service account setup

Configure service accounts or application registrations in the target systems where your action connectors will operate. These accounts should have the minimum necessary permissions to perform the required actions while adhering to the principle of least privilege. Document the specific permissions granted to each service account for future reference and security audits.

Testing environment

Establish a testing environment that mirrors your production setup to validate action connector functionality before deployment. This environment should include access to test instances of your target systems, allowing you to verify connector behavior without affecting production data or processes.

Setting up Action steps in flows

To add action steps to your flow, follow these steps:

1. In the Flow builder, select the **+ Add step** button.
2. From the menu, choose **Application actions step**.
3. In the **Edit** panel, choose your connector and your action.
4. Write your prompt for the action execution.
5. Save your configuration.

Using action connectors in flows

Action connectors integrate seamlessly into your flow design, appearing as actionable steps that can be configured and chained with other Flow components. The implementation process involves

selecting the appropriate connector, configuring its parameters, and defining how it interacts with other elements in your flow. Amazon Quick Flows supports enhanced action flows that display eligible fields as dropdowns, reducing cognitive load and improving user experience.

Connector selection and configuration

Choose the appropriate action connector based on your integration requirements and the target system's capabilities. Each connector provides a configuration interface where you specify connection details, authentication information, and action-specific parameters. Flow creators can select from available plugins configured by administrators and choose specific actions from dropdown menus.

Data mapping and transformation

Configure how data flows between your Amazon Quick Flows and the connected systems. This includes mapping input parameters from your flow to the connector's expected format and defining how response data should be processed and passed to subsequent Flow steps. The mapping interface provides tools for data transformation, allowing you to modify data formats, apply filters, or perform calculations as needed.

Error handling and retry logic

Implement appropriate error handling mechanisms to manage potential connectivity issues, authentication failures, or service unavailability. Configure retry policies that define how the connector should respond to temporary failures, including retry intervals, maximum retry attempts, and escalation procedures for persistent issues.

Flow integration patterns

Design your flows to effectively utilize action connectors within the broader Flow logic. This includes determining the optimal placement of action steps, configuring conditional running based on previous step results, and implementing parallel processing where appropriate to optimize performance.

Authentication: 2 Legged OAuth vs 3 Legged OAuth

Action connectors support multiple authentication methods to accommodate different security requirements and integration scenarios. Understanding the differences between 2-legged and 3-legged OAuth helps you choose the appropriate authentication approach for your specific use case.

2 Legged OAuth

2-legged OAuth, also known as client credentials grant flow, provides server-to-server authentication without requiring user interaction. This method is ideal for automated processes where your flow needs to perform actions on behalf of the application rather than a specific user. The authentication process involves your application directly exchanging credentials with the target service to obtain an access token. This flow is suitable for service-to-service token authentication and is used when the identity of the end user does not matter, such as when accessing service-owned resources or for bot operations.

This approach offers simplified implementation and reliable automation since it doesn't depend on user presence or interaction. However, actions performed using 2-legged OAuth are typically associated with the application or service account rather than individual users, which may limit audit trails and personalization capabilities.

3 Legged OAuth

3-legged OAuth, also known as authorization code grant flow, involves the end user in the authentication process, requiring them to explicitly grant permission for your flow to access their account in the target system. This method is suitable when the identity of the end user matters, such as when accessing user-owned resources or when operations should be performed under the user's identity rather than a bot. This flow is supported by most ISVs who support OAuth and provides user-specific access with appropriate user context and permissions.

The 3-legged OAuth process requires users to authenticate with the target service and authorize your application's access during their first interaction with the flow. Subsequent runs can use stored refresh tokens to maintain access without repeated user intervention, provided the tokens remain valid and the user hasn't revoked access.

This authentication method provides better security and audit capabilities since actions are tied to specific user accounts. However, it requires more complex implementation and user interaction, which may not be suitable for fully automated processes.

Available action connectors

Amazon Quick Flows supports a comprehensive range of action connectors that enable integration with diverse systems and workflows. For the most current list of available connectors and their capabilities, see [Working with integrations](#) and [Action integrations](#).

Connector capabilities and limitations

Each connector category offers different levels of functionality, support, and maintenance. Second-party connectors offer reliable integration with established partner services, while third-party connectors provide broad compatibility with varying levels of feature completeness.

When selecting connectors, consider factors such as authentication requirements, rate limiting, data format compatibility, and long-term maintenance commitments. Review the specific documentation for each connector to understand its capabilities, limitations, and best practices for implementation.

Web search in flows

Amazon Quick Flows provides web search capabilities that allow you to retrieve current information from the internet. This feature enables your flows to access real-time data, news, and information beyond what's available in your organization's internal knowledge base.

Overview of web search in flows

Web search in Amazon Quick Flows allows you to incorporate internet search capabilities into your flows. This feature is particularly useful when you need to access current information, verify facts, or gather data from public sources. Web search uses a secure, managed search service that provides relevant results while maintaining privacy and security.

When you add a web search step to your flow, you can configure it to search for specific information based on user input or predefined queries. The search results can then be processed by subsequent steps in your flow, such as reasoning steps that analyze and summarize the information.

Adding web search to your flow

To add web search capabilities to your flow, follow these steps:

- Open the Flow editor and navigate to the point where you want to add web search.
- Click the "+" button to add a new step.
- Select "Web Search" from the available step types.
- Configure the search parameters as needed.

Once added, the web search step will appear in your flow's canvas. You can connect it to other steps to create a comprehensive workflow that leverages internet information.

Search results may vary over time as internet content changes. A Flow that works well today might return different results in the future.

Some content may not be accessible through web search due to various restrictions or the nature of the content.

Related topics

For more information about related features and capabilities, see the following topics:

- [the section called "Creating flows"](#)
- [the section called "Editing flows"](#)
- [the section called "Response preferences"](#)

UI agent in Amazon Quick Flows

The UI agent in Amazon Quick Flows (currently in Preview) allows you to automate interactions with web-based user interfaces and applications. This feature empowers business users to create automation workflows that can navigate websites, fill forms, extract data, and perform actions across various web applications without requiring technical API integration knowledge.

Overview of UI agent

The UI agent feature in Amazon Quick Flows represents a sophisticated web automation capability that transforms how users interact with websites and web applications. This feature functions as an AI agent running on a remote host that can perform human-like interactions with web interfaces, including navigating to specific pages, scrolling through content to locate information, clicking buttons, entering text into form fields, selecting options from dropdown menus, and extracting data from web pages. What makes this particularly powerful is its ability to navigate websites and perform tasks (like filling forms) across multiple websites without requiring users to have technical API integration knowledge or programming skills. Business users can simply add UI agent steps to their flows and for each such step, describe in natural language what they want done on which websites. The feature is currently in Preview as we look to extend possibilities with support for websites and applications that require user login and captcha resolution.

UI agent capabilities

The UI agent provides several key capabilities that make it powerful for business automation:

- Web navigation across multiple sites and applications
- Form filling and data entry automation
- Data extraction from web pages
- Conditional logic based on web content

These capabilities allow business users to create automation workflows without writing code or understanding API integration concepts.

Common use cases

The UI agent is particularly valuable for scenarios where API integrations are unavailable or would require significant development effort. Common use cases include:

Data entry automation allows you to transfer information from one system to another by extracting data from source applications and inputting it into destination applications. This eliminates manual copy-paste operations and reduces errors.

Report generation and extraction enables you to navigate to reporting interfaces, configure parameters, generate reports, and extract the resulting data for further processing or analysis.

Multi-system workflows help you create end-to-end processes that span multiple applications, such as retrieving customer information from a CRM system and using it to create invoices in a billing system.

Legacy system integration provides a way to automate interactions with older systems that lack modern API capabilities but still have web interfaces.

Setting up UI agent in your flow

To incorporate UI agent capabilities into your flow:

- Add a UI agent step to your flow from the **Add step** menu.
- Describe in natural language what you want done. Adding specific URLs, and one per step, will help improve accuracy and speed.

The UI agent provides a visual interface for configuring these actions, making it accessible to business users without technical expertise.

Note

UI agent is currently in Preview. Some websites implement anti-automation measures that may limit UI agent capabilities. These can include CAPTCHA challenges or other mechanisms designed to detect and block automated interactions.

Related topics

For more information about related features and concepts, see these topics:

- [the section called "Creating flows"](#)
- [the section called "Editing flows"](#)
- [the section called "Response preferences"](#)

Running and integrating flows

This section covers how to run Amazon Quick Flows and integrate them with other Amazon Quick Suite capabilities. You'll learn about executing flows in different runtime modes, using pre-built flows, and invoking flows from custom agents and My Assistant.

Running Flows in runtime allows end users to interact with your flows through structured or conversational interfaces. Pre-built flows provide ready-to-use templates for common business scenarios. Integration capabilities enable you to invoke Flows from custom agents and My Assistant, extending their functionality and reach within your organization.

These topics will help you maximize the value of your flows by ensuring they can be effectively executed and integrated with other systems and capabilities.

Interacting with flows in Runtime Mode

Runtime execution is how you interact with and run Amazon Quick Flows after they've been created and shared. When you open a flow from the library, you enter the runtime environment where you can provide inputs, run steps, and get results from the flow. Think of runtime as the "using" phase of flows - where the actual work gets done.

Runtime execution provides flexible ways to interact with flows, from structured step-by-step running to conversational chat-based interaction, giving you the control and experience that works best for your workflow.

Runtime modes

Amazon Quick Flows offers three different runtime modes to match your preferred way of working and the complexity of your flow.

Dual mode runtime

Dual mode gives you the best of both worlds by combining structured and chat interfaces in a single view. This is the default mode when you open most flows from the library.

What you see: The interface shows both a structured view on the left displaying all flow steps, inputs, and outputs, alongside a chat interface on the right where you can have conversations with the flow's AI agent.

How it works: You can interact with your flow in either interface - provide inputs through the structured forms or by chatting with the agent. The agent understands your flow's structure and can guide you through the process, request inputs, and run steps based on your conversational requests.

Best for: Most users and most flows, especially when you want flexibility to switch between structured interaction and natural conversation. This mode is particularly useful when you're learning a new flow or when the flow has complex logic.

Key features:

- Progress tracking visible in structured view
- Conversational guidance from the AI agent
- Ability to jump to any step or run steps out of order
- Real-time synchronization between chat and structured interfaces

Full structured mode

Full structured mode provides a traditional form-based interface where you can see and interact with all Flow components in an organized, step-by-step layout.

What you see: A clean, organized view showing all flow steps, input fields, output areas, and action buttons. Each step is clearly labeled and shows its current status (pending, running, completed).

How it works: You work through the flow by filling in input fields, clicking run buttons, and reviewing outputs in a structured sequence. You can see the entire flow layout and jump to any step you need to work on.

Best for: Users who prefer traditional form interfaces, complex flows with many steps, or situations where you need to see the complete flow structure at once. Also ideal for flows that require precise input formatting or when working with multiple related inputs.

Key features:

- Clear visual progress through flow steps
- Direct input into form fields
- Immediate visibility of all Flow components
- Easy navigation between steps

Full chat mode

Full chat mode provides a purely conversational interface where you interact with your flow entirely through natural language conversation with an AI agent.

What you see: A chat interface similar to messaging applications, where you have conversations with an AI agent that understands your flow and can run it based on your requests.

How it works: The AI agent guides you through the flow by asking for inputs, explaining what the flow can do, and running steps based on your conversational requests. You can ask the agent to run specific parts of the flow, regenerate outputs, or modify results through natural language.

Best for: Users who prefer conversational interfaces, simple flows with straightforward inputs, or when you want to focus on the task rather than the flow structure. Particularly useful for flows you use frequently.

Key features:

- Natural language interaction
- AI agent guidance and explanations
- Ability to request specific actions conversationally
- Simplified interface focused on conversation

What is a progress tracker?

The progress tracker is a visual indicator that shows you the current status and progress of your flow execution. It helps you understand where you are in the flow, what steps have been completed, and what's currently running or waiting for input.

How the progress tracker works

Step status indicators: Each step in your flow shows its current state - pending (not yet started), running (currently running), completed (finished successfully), or requiring input (waiting for you to provide information).

Active step highlighting: The step you're currently working on or that's currently running is highlighted and brought into view automatically. If multiple steps are running in parallel, all active steps are highlighted.

Progressive disclosure: Steps open and close automatically based on the flow's execution. When you start a flow, only the first step is open. As you progress, subsequent steps open when they become active, keeping your focus on the current work.

Visual progress indicators: Running steps show progress indicators (like spinning icons) to let you know something is happening, especially for steps that take time to complete.

Understanding step states

Closed steps: Steps that haven't been reached yet or aren't currently relevant remain closed to reduce visual clutter.

Open steps: Steps that are active, or require your input are automatically opened and visible.

Completed steps: Steps that have finished successfully show their results and cannot be closed.

Error states: If a step encounters an error, it's clearly marked and provides information about what went wrong and how to resolve it.

How can I start a new flow run?

Starting a new flow run gives you a fresh execution environment to work with your flow. There are several ways to begin a new run depending on your current context.

Starting from the library

When you open a flow from the Amazon Quick Flows library, you automatically start a new run. The flow opens in dual mode by default, with the AI agent providing a greeting message and overview of what the flow can accomplish.

What happens: The flow initializes with empty inputs and a clean state. The AI agent introduces the flow and asks for any initial inputs needed to get started. All previous execution history is separate from this new run.

Starting a new run from within a flow

If you're already working with a flow and want to start fresh, you can create a new run without leaving the current flow.

Using the New run button: Click the **New run** button in the chat interface or in the structured mode to start a completely new flow run. This creates a new conversation and execution context while preserving your previous run in the history.

Confirmation process: When you start a new run, you'll see a confirmation dialog to make sure you want to leave your current progress and begin fresh. This prevents accidental loss of work.

What gets reset vs. preserved

New run button:

- Creates entirely new execution context
- Clears all inputs and outputs
- Starts fresh chat conversation
- Previous run is saved to history

Ability to resume runs from flow history

Flow history allows you to return to previous executions of your flows, continuing where you left off or reviewing past results. This feature is essential for ongoing work and maintaining context across multiple sessions.

Accessing flow history

History icon: Each flow has a history icon in the chat interface that shows all your previous runs of that specific flow. Click this icon to see a chronological list of your past executions.

Recent runs first: Your flow history is organized with the most recent runs at the top, along with timestamps showing when each run was created or last modified.

Flow-specific history: Each flow maintains its own separate history. When you view history for a Customer Issue Resolver flow, you only see runs for that specific flow, not other flows you've used.

Resuming previous runs

Continuing where you left off: When you select a previous run from history, both the chat conversation and structured mode are restored to their last state. If you were in the middle of providing input for a specific step, that step will be highlighted and ready for you to continue.

Preserved context: All your previous inputs, generated outputs, and conversation history are maintained exactly as they were when you last worked on that run.

Seamless continuation: You can immediately continue working - provide additional inputs, run next steps, or modify previous results without having to start over.

Managing your run history

Automatic saving: Every interaction with a flow is automatically saved to your run history. You don't need to manually save your progress.

History persistence: flow runs are preserved in your history for an extended period of 30 days, allowing you to return to work even after days or weeks.

Run identification: Each run in your history is identified by either your first message to the agent or a smart summary of what you accomplished in that session, making it easy to find the run you're looking for.

Updating existing runs: When you resume and continue working on a previous run, it's updated in place and moved to the top of your history list as the most recent activity.

Working with runtime

Understanding how to effectively use the different runtime modes and features will help you get the most out of your flows.

Switching between modes

You can easily switch between runtime modes while working with a flow:

Dual to structured: Click the "X" button in the chat interface to close the chat and work in full structured mode.

Structured to dual: Click the sparkle icon to reopen the chat interface and return to dual mode.

Dual to chat: Expand the chat interface to focus entirely on conversational interaction.

Best practices for runtime

Choose the right mode: Use dual mode for learning new flows, structured mode for complex multi-step processes, and chat mode for familiar flows or quick tasks.

Leverage the AI agent: In chat and dual modes, the AI agent can explain flow capabilities, suggest next steps, and help troubleshoot issues. You can also ask the agent to customize output responses by specifying length requirements or adjusting tone. Request modifications like "summarize meeting notes in 200 words" or "make email draft more professional". Don't hesitate to ask questions about what the flow can do.

Use history effectively: Regularly check your flow history to resume important work or reference previous results. This is especially valuable for flows you use repeatedly with different inputs.

Monitor progress: Pay attention to the progress tracker to understand what's happening, especially for flows with long-running steps or complex logic.

Handling errors and interruptions

Error recovery: When errors occur, the flow provides clear information about what went wrong and suggested next steps for resolution or gives you an option to retry.

Input validation: The system validates your inputs based on the flow's requirements and provides helpful feedback if something needs to be corrected.

Amazon Quick Suite reference

Use this table to understand the different execution modes:

Runtime Modes

#	Runtime Mode	Interface	Best For	Key Features
1	Dual Mode	Chat + Structured	Most users, learning flows	Flexible interaction, AI guidance
2	Full Structured	Forms and buttons	Complex flows, precise control	Clear progress, direct input

#	Runtime Mode	Interface	Best For	Key Features
3	Full Chat	Conversation only	Simple flows, mobile use	Natural language, AI assistance

Common Actions

#	Action	How To	Result
1	Start new run	Click <i>New run</i> button	Fresh execution, new history entry
2	Resume previous run	History → Select run	Continue where you left off
3	Switch to structured	Click X in chat	Full structured mode
4	Switch to dual	Click sparkle icon	Chat + structured mode

Using Pre-built flows

Amazon Quick Flows provides pre-built, ready-to-use templates. These templates help you get started quickly with common business tasks without building flows from scratch. Pre-built flows are designed based on best practices and can be used as-is or customized to meet your specific requirements.

Accessing pre-built flows

You can access pre-built flows from the Quick Flows gallery. When creating a new flow, you'll see a section dedicated to pre-built flows that you can browse and select from.

To access pre-built flows:

1. Sign in to Amazon Quick Suite.
2. Navigate to the Quick Flows section.
3. Select *Create Flow*.

4. Browse the *Pre-built flows* section of the gallery.
5. Select a pre-built flow to view its details and purpose.

Available pre-built flows

Amazon Quick Flows offers several categories of pre-built flows to address different business needs. Each pre-built flow is designed for a specific purpose and includes appropriate input fields, processing steps, and output formatting.

Content creation flows

Content creation flows help you generate various types of written content for different purposes.

Blog post generator

The blog post generator helps you create engaging blog content based on a topic, target audience, and key points. This Flow guides you through the process of creating well-structured blog posts with appropriate headings, paragraphs, and calls to action.

Key features include:

- Topic and audience targeting
- Structured content with introduction, body, and conclusion
- SEO-friendly formatting suggestions
- Tone and style customization

Email composer

The email composer helps you draft professional emails for various business scenarios. This Flow provides templates for common email types and guides you through creating effective email content.

Key features include:

- Multiple email templates (announcement, request, follow-up)
- Tone adjustment based on recipient relationship
- Subject line optimization
- Call-to-action suggestions

Data analysis flows

Data analysis flows help you extract insights and create visualizations from your data.

Data summarizer

The data summarizer helps you analyze and extract key insights from structured data. This Flow can process CSV or tabular data and provide meaningful summaries and observations.

Key features include:

- Statistical analysis of numerical data
- Trend identification and highlighting
- Anomaly detection
- Visualization recommendations

Chart creator

The chart creator helps you generate appropriate visualizations based on your data and analysis goals. This Flow guides you through selecting the right chart type and configuring it for maximum clarity.

Key features include:

- Chart type recommendations based on data characteristics
- Customizable visualization parameters
- Integration with Amazon Quick Sight for advanced visualizations
- Exportable chart configurations

Business process flows

Business process flows help you streamline common operational tasks and workflows.

Meeting summarizer

The meeting summarizer helps you create concise and structured summaries of meeting discussions. This Flow processes meeting notes or transcripts and extracts key points, decisions, and action items.

Key features include:

- Automatic identification of discussion topics
- Extraction of decisions and action items
- Assignment tracking for action items
- Follow-up reminder creation

Project status reporter

The project status reporter helps you create comprehensive project status updates based on key metrics and milestones. This Flow guides you through collecting relevant project information and formatting it into a clear status report.

Key features include:

- Milestone tracking and reporting
- Risk and issue highlighting
- Resource allocation summary
- Next steps and recommendations

Customizing pre-built flows

While pre-built flows are designed to be used immediately, you can also customize them to better fit your specific requirements. Customization allows you to adapt the pre-built templates to your organization's processes, terminology, and branding.

To customize a pre-built flow:

1. Select the pre-built flow you want to customize.
2. Click *Duplicate* to create your own copy of the flow.
3. Open the duplicated flow in the Flow editor.
4. Modify the Flow components as needed, such as:
 - Adjusting input prompts and fields
 - Editing reasoning group instructions
 - Adding or removing steps

- Changing output formats
5. Save your customized Flow with a new name.

When customizing pre-built flows, consider preserving the core structure and best practices incorporated in the original template while adapting specific elements to your needs.

Best practices for using pre-built flows

To get the most value from pre-built flows, consider the following best practices:

- Start with the pre-built flow that most closely matches your use case, then customize as needed rather than building from scratch.
- Review the structure and design of pre-built flows to understand the best practices they incorporate.
- Use pre-built flows as learning tools to understand effective Flow design patterns.
- Share customized versions of pre-built flows with your team to ensure consistency.
- Periodically check for new pre-built flows that might address emerging use cases.

Related resources

For more information about working with flows, see the following topics:

- [the section called "Creating flows"](#)
- [the section called "Editing flows"](#)
- [the section called "Response preferences"](#)

Invoking flows in My Assistant and custom agents

Flow invocation in My Assistant and custom agents allows you to seamlessly access and run your Amazon Quick Flows directly within conversational AI experiences. Instead of switching between different interfaces, you can discover, run, and interact with flows while maintaining the context of your ongoing conversations.

This integration brings the power of structured workflows into natural conversation environments, making it easier to accomplish complex tasks without losing the context of your work or having to navigate between multiple applications.

How flow invocation works

Amazon Quick Flows can be invoked within two main conversational environments, each designed to provide seamless integration while maintaining the unique characteristics of the conversational experience.

Integration environments

My Assistant: The default conversational AI experience in Amazon Quick Suite that provides access to your organization's knowledge, data, and tools through natural conversation.

Custom Agents: Specialized conversational AI agents that your organization has configured for specific purposes, domains, or workflows.

Both environments allow you to access your flows while maintaining conversational context, but each has specific behaviors and capabilities tailored to their intended use.

Key benefits of flow invocation

Context preservation: Your conversation history and context are maintained when you invoke flows, allowing the flow to leverage information from your previous interactions.

Seamless transitions: Move between general conversation and structured flow running without losing your place or having to start over.

Unified experience: Access all your flows from within your regular conversational workflow, reducing the need to switch between different applications or interfaces.

Enhanced capabilities: Combine the flexibility of conversational AI with the structured power of flows to accomplish complex, multi-step tasks.

Choosing your invocation environment

The choice between My Assistant and Custom Agents depends on your specific needs and how your organization has configured these environments.

When to use My Assistant

My Assistant is ideal for general-purpose work where you need access to a broad range of organizational capabilities alongside your flows. It provides a comprehensive environment for knowledge work, data analysis, and task running.

Best for: Daily knowledge work, exploratory analysis, general task automation, and situations where you need access to multiple organizational resources and flows.

When to use Custom Agents

Custom Agents are designed for specific domains, roles, or workflows. They provide specialized conversational experiences tailored to particular use cases or organizational functions.

Best for: Domain-specific work, role-based tasks, specialized workflows, and situations where you need focused functionality within a particular area of expertise.

Getting started with flow invocation

To begin using flows within conversational agents, you'll need access to either My Assistant or Custom Agents in your Amazon Q Business environment. Your organization's administrators control access to these features and configure which flows are available in each environment.

Prerequisites

Flow access: You must have access to flows either as their creator or through sharing permissions. Only flows you can normally access will be available for invocation within agents.

Agent access: Your organization must have enabled Chat agents settings in custom permissions, and you must have appropriate permissions to use these conversational environments.

Integration enabled: Your administrators must have enabled flow invocation within the conversational environments you want to use.

Basic invocation process

Discovery: Find flows through search, recommendations, or browsing within the conversational interface.

Selection: Choose the flow you want to run explicitly.

Running: Interact with the flow conversationally or through structured interfaces, depending on your preference and the flow's design.

Completion: Return to general conversation or continue with related flows as needed.

Detailed guidance by environment

For specific information about using flows in each environment, see the detailed guides:

- **Invoking flows in My Assistant:** Comprehensive guide to using flows within the default Amazon Quick Suite conversational experience
- **Invoking flows in Custom Agents:** Detailed information about flow invocation within specialized custom agents

Each guide covers the specific behaviors, capabilities, and best practices for that environment, including getting started instructions, invocation methods, context sharing, and management features.

Common features across environments

While My Assistant and Custom Agents have unique characteristics, they share several common features for flow invocation.

Universal capabilities

Explicit invocation: Directly search for and select flows to run within the conversational interface.

Context sharing: Leverage information from your conversation history to provide inputs and context to flows.

Progress tracking: Monitor flow running progress through visual indicators and status updates.

Guardrails: Benefit from consistent security and compliance controls across both conversational and flow experiences.

History integration: Access your flow running history and conversation logs in a unified way.

Consistent user experience

Natural language interaction: Communicate with flows using natural language, just as you would in regular conversation.

Flexible running modes: Choose between conversational and structured interaction modes based on your preferences and the flow's complexity.

Error handling: Receive clear guidance when issues occur, with consistent error resolution patterns across environments.

File handling: Upload and work with files seamlessly within the conversational context when flows require file inputs.

Best practices for flow invocation

To get the most out of flow invocation in conversational agents, follow these recommended approaches.

Managing flow running

Monitor progress: Pay attention to progress indicators to understand what's happening, especially for complex or long-running flows.

Provide clear inputs: When flows request information, be specific and complete in your responses to avoid back-and-forth clarification.

Use appropriate modes: Choose between conversational and structured interaction based on the complexity of the flow and your comfort level.

Amazon Quick Suite reference

Environment Comparison

#	Environment	Best For	Key Features
1	My Assistant	General knowledge work	Broad access, comprehensive capabilities
2	Custom Agents	Specialized tasks	Domain focus, tailored experience

Common Actions

#	Action	How To	Available In
1	Find flows	Search or browse	Both environments
2	Invoke flow	Select from search	Both environments
3	Share context	Reference conversation history	Both environments

#	Action	How To	Available In
4	Track progress	Monitor visual indicators	Both environments
5	End flow	Complete workflow or exit explicitly	Both environments

Invoking flows in My Assistant

My Assistant is your default conversational AI experience in Amazon Quick Suite, designed to help you access organizational knowledge, analyze data, and accomplish tasks through natural conversation. When flows are enabled in My Assistant, you can seamlessly invoke and run your Amazon Quick Flows without leaving your conversational context.

Getting started: Invoking flows in My Assistant

What My Assistant provides

My Assistant serves as a comprehensive workspace where you can ask questions about company data, request analysis, generate content, and now run structured workflows through flows. It's designed to be your primary interface for knowledge work and task automation within your organization.

Unified experience: Access flows alongside all other My Assistant capabilities, including knowledge search, data analysis, content generation, and action running.

Organizational context: My Assistant understands your organization's data, processes, and available resources, making it an ideal environment for flow running that leverages company-specific information.

Broad access: Unlike specialized custom agents, My Assistant provides access to the full range of flows available to you, from simple automation to complex multi-step workflows.

Prerequisites for flow invocation

Before you can invoke flows in My Assistant, several conditions must be met:

My Assistant access: Your organization must have enabled My Assistant, and you must have appropriate permissions to use the conversational interface.

Flow permissions: You can only invoke flows that you have permission to access - either flows you've created or flows that have been shared with you through the standard Amazon Quick Flows sharing mechanisms.

Required capabilities: My Assistant must have access to the underlying services and data sources that your flows require, such as Amazon Quick Sight for visualization flows or action connectors for workflow automation.

Initial setup and access

When you first access My Assistant with flow capabilities enabled, you'll notice additional options for discovering and invoking flows. The interface remains familiar while providing new ways to access your automated workflows.

Flow discovery options: My Assistant provides multiple ways to find and access flows, including search and browsing.

Integration indicators: Visual cues in the interface indicate when flow capabilities are available and when you're interacting with flows versus standard My Assistant features.

Explicit flow invocation

Explicit flow invocation allows you to directly search for, select, and run specific flows within My Assistant. This approach gives you full control over which flows to use and when to use them.

Discovering available flows

My Assistant provides several methods for finding flows that are available to you:

Browse by category: Access flows organized by categories such as Recently Used, Popular, or All Available flows. This organization helps you quickly find flows based on usage patterns and relevance.

Selecting and launching flows

Once you've identified a flow you want to use, the invocation process is straightforward:

Direct selection: Choose a flow from search results or browse lists to begin running immediately.

Preview information: Before launching, you can view basic information about the flow, including its purpose, required inputs, and expected outputs.

Context preparation: If the flow can benefit from your current conversation context, My Assistant will prepare and share relevant information automatically.

Flow running modes in My Assistant

When you invoke a flow in My Assistant, you can choose how to interact with it:

Conversational mode: Interact with the flow entirely through natural language conversation with My Assistant. The assistant guides you through the flow steps, requests inputs, and presents outputs within the chat interface.

Structured mode: Access a more traditional form-based interface for flows that require precise inputs or complex configuration. This mode appears alongside your conversation and provides direct access to Flow components.

Hybrid approach: Switch between conversational and structured interaction as needed, allowing you to use the most appropriate method for each part of the flow.

Context sharing and log history

One of the key advantages of invoking flows within My Assistant is the ability to leverage your conversation history and context to enhance flow running.

How context sharing works

Automatic context detection: My Assistant analyzes your conversation history to identify information that might be relevant to the flow you're invoking. This includes data points, file references, previous analysis results, and stated goals or preferences.

Intelligent information extraction: When launching a flow, My Assistant can automatically extract relevant information from your conversation to pre-populate flow inputs, reducing the need for you to re-enter information you've already provided.

Seamless information flow: Information flows naturally between your general conversation and flow running, creating a unified experience where context is preserved and leveraged effectively.

Managing conversation history

Unified logging: Your flow invocations and runs are integrated into your My Assistant conversation history, creating a complete record of your work session.

Context preservation: When you return to previous conversations, the context of flow runs is preserved, allowing you to reference results or continue related work.

Cross-session continuity: Your conversation history with My Assistant persists across sessions, enabling you to build on previous work and maintain context over time.

Guardrails and tone honored

My Assistant maintains consistent security, compliance, and behavioral standards across both conversational interaction and flow running, ensuring a unified and safe experience.

For information about security and compliance guardrails, see [the section called “Managing user access”](#).

Behavioral consistency

Tone and style: My Assistant maintains its configured tone and communication style when guiding you through flow running, providing a consistent conversational experience.

Response patterns: The way My Assistant presents information, asks for clarification, and provides guidance remains consistent whether you're having a general conversation or running a flow.

Error handling: When issues occur during flow running, My Assistant handles them using the same helpful, clear communication patterns you're accustomed to in regular conversation.

User experience standards: The overall user experience standards that govern My Assistant interactions also apply to flow invocation, ensuring quality and consistency.

Customization and personalization

Organizational customization: If your organization has customized My Assistant's behavior, tone, or capabilities, these customizations are maintained during flow running.

Personal preferences: Your individual preferences for how My Assistant interacts with you are preserved when invoking and running flows.

Ending a flow in the conversation

Properly concluding flow running within My Assistant ensures a smooth transition back to general conversation and maintains the context of your work.

Manual flow termination

Explicit exit: You can end a flow at any time by explicitly stating your intention to stop, such as "Cancel this flow" or "I want to stop and return to our conversation."

Graceful interruption: If you need to interrupt a running flow, My Assistant handles the interruption gracefully, saving any progress where possible and providing options for resuming later.

Context preservation: When you manually end a flow, any progress made and information gathered is preserved in your conversation history for potential future use.

Clear transition: My Assistant clearly indicates when you've exited flow running and returned to general conversation mode, helping you understand the current context.

Post-flow conversation

Seamless continuation: After completing or ending a flow, you can immediately continue your conversation with My Assistant, referencing flow results or moving on to new topics.

Result discussion: You can discuss flow results with My Assistant, ask for interpretation or analysis, or request modifications or follow-up actions.

Flow chaining: You can easily invoke additional flows based on the results of previous flows, creating complex workflows through natural conversation.

Session management: My Assistant helps you manage your overall work session, keeping track of completed flows and helping you organize your activities.

Amazon Quick Suite reference

Common Flow Invocation Actions

#	Action	How To	Example
1	Launch a specific flow	Select from search or browse results	Click on "Expense Report Automation"
2	Use conversation context	Reference previous discussion	"Use the data we analyzed earlier"

#	Action	How To	Example
3	End a flow	Click on End button	"Select end button on chat"

Invoking flows in Custom Agents

Custom Agents are specialized conversational AI experiences designed for specific domains, roles, or organizational functions. Unlike the general-purpose My Assistant, Custom Agents are tailored to particular use cases and often have specialized knowledge, capabilities, and workflows. When flows are enabled in Custom Agents, you can access relevant automation within these focused conversational environments.

Getting started: Invoking flows in Custom Agents

What Custom Agents provide

Custom Agents are built to serve specific purposes within your organization, such as customer support, financial analysis, HR processes, or technical troubleshooting. Each Custom Agent has its own personality, knowledge base, and set of capabilities designed to excel in its particular domain.

Specialized focus: Custom Agents are optimized for specific tasks, roles, or domains, providing deep expertise and tailored workflows for their area of specialization.

Domain expertise: Each Custom Agent understands the terminology, processes, and requirements specific to its domain, making it more effective for specialized tasks than general-purpose assistants.

Curated capabilities: Custom Agents typically have access to a carefully selected set of flows that are relevant to their specific purpose, reducing complexity and improving task efficiency.

Consistent experience: The agent's personality, communication style, and approach remain consistent whether you're having general conversations or running flows within its domain.

Prerequisites for flow invocation

Before you can invoke flows in Custom Agents, several conditions must be met:

Custom Agent access: Your organization must have created Custom Agents, and you must have appropriate permissions to interact with the specific agents that have flow capabilities.

Relevant flow access: You can invoke any flow from a custom agent as long as you have access to the flow.

Understanding agent specialization

Each Custom Agent is designed with specific goals and capabilities:

Purpose-built knowledge: Custom Agents have specialized knowledge relevant to their domain, which they can leverage when helping you select and run appropriate flows.

Workflow optimization: The agent understands common workflows in its domain and can guide you toward the most effective flows and running patterns for your specific needs.

Context awareness: Custom Agents understand the typical context and requirements of their domain, allowing them to provide more targeted assistance and recommendations.

Explicit flow invocation

Explicit flow invocation in Custom Agents follows similar patterns to My Assistant but with domain-specific enhancements and curated selections.

Domain-optimized running

Flow running within Custom Agents benefits from specialized knowledge and capabilities:

Contextual guidance: Custom Agents provide domain-specific guidance during flow running, helping you understand inputs and outputs in the context of your specialized work.

Enhanced error handling: When issues occur, Custom Agents can provide domain-specific troubleshooting advice and alternative approaches based on their specialized knowledge.

Workflow continuity: Custom Agents understand how flows fit into broader domain workflows and can help you maintain continuity across complex, multi-step processes.

Results interpretation: Custom Agents can help interpret flow results in the context of domain-specific goals and requirements, providing more valuable insights than generic interpretation.

Context sharing and log history

Context sharing in Custom Agents is enhanced by domain specialization and focused conversation history.

Guardrails and tone honored

Custom Agents maintain their specialized personality and domain-specific guardrails while ensuring consistent security and compliance standards. For more information about agent guardrails, see [the section called “Managing chat customization”](#).

Ending a flow in the conversation

Concluding flow running within Custom Agents leverages domain expertise to provide specialized guidance and next steps.

Domain-focused post-flow conversation

Specialized result discussion: Custom Agents can engage in discussion about flow results using domain expertise, helping you understand implications and opportunities.

Expert interpretation: Custom Agents provide expert interpretation of flow results within the context of domain knowledge and best practices.

Workflow progression: Custom Agents guide you through logical next steps in domain workflows, leveraging their understanding of how different activities connect.

Amazon Quick Suite reference

Custom Agent Flow Invocation

#	Action	How To	Domain Benefits
1	Find relevant flows	Manual flow invocation through search	Curated, relevant results
2	Run with context	Leverage agent's domain knowledge	Enhanced running quality
3	Interpret results	Discuss outcomes with domain expert	Specialized insights
4	Plan next steps	Get domain-specific guidance	Expert workflow advice

Managing Quick Flows

As an administrator, you have comprehensive control over how Amazon Quick Flows operates within your organization. These administrative capabilities enable you to establish governance policies, manage user permissions, control feature availability, and oversee the Amazon Bedrock models available to flow creators.

Administrative management of Quick Flows encompasses two primary areas: permissions and asset management through the Custom Permissions and Asset Management pages, and Amazon Bedrock model allowlisting for controlling which AI models flow creators can use in their workflows. These controls ensure that Quick Flows deployment aligns with your organization's security, compliance, and operational requirements.

Administrative access requirements

To manage Quick Flows settings, you must have administrator privileges in your Amazon Quick Suite instance. Administrative controls are available through the Amazon Quick Suite console's management interface, where you can configure permissions, manage assets, and control model availability.

Governance capabilities overview

Quick Flows administrative controls provide multiple layers of governance:

- **Feature enablement** - Control whether flows are available to users in your account
- **Approval workflows** - Require administrative approval before flows can be shared with users
- **Asset management** - Oversee Flow sharing, ownership, and visibility across your organization

Managing Quick Flows permissions and assets

Administrative control over Quick Flows permissions and assets is managed through two primary interfaces in the Amazon Quick Suite console: the Custom Permissions page for controlling capabilities and features, and the Asset Management page for overseeing Flow sharing and ownership. These tools provide comprehensive governance over how Quick Flows operates within your organization.

Using Custom Permissions page

The Custom Permissions page provides centralized control over Quick Flows capabilities and features available to users in your account. Through this interface, you can enable or disable core functionality and manage approval workflows that govern how flows are shared within your organization.

Restrict capabilities

The restrict capabilities section allows you to control fundamental Quick Flows functionality at the account level. This provides a master switch for Quick Flows availability across your entire organization.

Enable or disable flows for all users

Administrators can enable or disable Quick Flows access for all users in the account through a single control. This setting affects both the creation and usage of flows across your organization.

When Quick Flows is enabled, users with appropriate permissions can create, edit, share, and run flows according to their assigned roles. When disabled, all Quick Flows functionality becomes unavailable to users, though existing flows remain preserved and can be re-enabled later.

By default, Quick Flows is enabled for new Amazon Quick Suite instances and for existing Amazon Quick Sight instances that upgrade to the Amazon Quick Suite experience. This opt-out approach ensures that organizations can immediately begin using Quick Flows capabilities while providing administrators the flexibility to disable the feature if needed for governance or compliance reasons.

Restrict features

The restrict features section provides granular control over specific Quick Flows capabilities, allowing administrators to selectively enable or disable advanced features based on organizational requirements.

Enable or disable approval review

Administrators can enable or disable approval review for Quick Flows sharing within the organization. When approval review is enabled, any flow that needs to be shared with users, groups, or made available to all users requires administrative approval before becoming accessible.

The approval review setting is disabled by default, allowing flow creators to share their work immediately upon creation. When you enable approval review, the system implements a review

process where administrators must explicitly approve each flow before it becomes available to the intended audience.

This feature is particularly valuable for organizations that require content review and approval processes, ensuring that all shared flows meet organizational standards and compliance requirements before reaching end users.

Enable or disable Amazon Bedrock model usage for General knowledge step in flows

Administrators can control whether Amazon Bedrock models are available for use in General knowledge steps within Flows. This setting provides organizational control over AI model usage in Flow workflows.

By default, Amazon Bedrock model usage is enabled for General knowledge steps in flows. Administrators can disable this capability to prevent flow creators from using Bedrock models in their General knowledge steps, ensuring alignment with organizational policies or compliance requirements.

When Bedrock model usage is disabled, flow creators will not be able to configure General knowledge steps for output refinement or generate images that rely on Amazon Bedrock models. Existing flows with General knowledge steps using Bedrock models may be affected and should be reviewed by their creators.

Enable or disable use of UI agent in flows

Admins can enable UI agents to perform browser tasks in flows.

Using Asset Management page

The Asset Management page provides comprehensive oversight of all flows within your organization, enabling administrators to manage Flow sharing, ownership, and visibility. This centralized interface allows you to maintain governance over your organization's flow ecosystem.

Ability to share flows

Through the Asset Management page, administrators can directly share flows with users or groups. This capability provides administrative oversight of flow distribution and ensures that valuable workflows reach their intended audiences.

Administrative sharing capabilities include the ability to share Flows on behalf of their creators, redistribute Flows to different user groups, and manage sharing permissions for Flows across the

organization. This is particularly useful for managing orphaned Flows or redistributing workflows when team structures change.

Administrators can also transfer ownership of flows between users, ensuring continuity when flow creators leave the organization or when workflows need to be reassigned to different teams or departments.

Ability to unlist flows from viewers

Administrators have the authority to unlist or remove Flows from the shared library, making them unavailable to viewers while preserving the underlying Flow for potential future use. This capability is essential for maintaining quality control and governance over shared content.

The unlisting capability allows administrators to quickly respond to issues with shared Flows, such as content that doesn't meet organizational standards, contains outdated information, or requires updates before continued use. Unlisted Flows remain accessible to their creators for editing and can be re-shared once any issues are resolved.

This feature supports organizational governance by providing administrators with the ability to curate the shared Flow library, ensuring that only approved, high-quality workflows remain visible to end users while maintaining the flexibility to restore access when appropriate.

When unlisting flows, administrators can view comprehensive details including the flow name, description, creation time, and creator information.

Quick Flows limits

Flow limits

The following limits apply to flow creation and management in your Quick Suite instance:

- **Maximum flows per instance:** 10,000 flows
- **Maximum flows per user:** 100 flows
- **Maximum steps per flow:** 35 steps

Input limits

Input limits vary depending on whether you're using General knowledge or Quick Suite data as your data source.

Text input limits

- **General knowledge:** 40,000 characters
- **Quick Suite data:** 15,000 characters

File upload limits

File upload limits depend on your knowledge source and file type:

File size limits:

- **General knowledge:**
 - Document files: Up to 50 MB (supported file types depend output preference)
 - Video files: Up to 1 GB
 - Image files: Up to 4.5 MB
- **Quick Suite data:**
 - Document files: Up to 50 MB (.pdf, .txt, .rtf, .doc, .docx, .ppt, .pptx, .csv, .xls)
 - Image files: Up to 10 MB (.png, .jpg, .jpeg)

Maximum number of files:

- **General knowledge:** Up to 5 files
- **Quick Suite data:** Up to 20 files

Character limits for file content:

- **General knowledge:**
 - Faster responses: Up to 1M characters total context limit
 - Versatility and Performance: Up to 1M characters total context limit

The total context limit includes characters from output prompts defined at build time and characters from input steps and file upload steps at runtime.

- **Quick Suite data:** 665,000 characters

Output limits

Output limits control the size of prompts and generated responses:

- **Maximum input prompt size:**
 - General knowledge: 5,000 characters
 - Quick Suite data: 5,000 characters
- **Maximum output characters rendered:**
 - General knowledge: Up to 40,000 characters
 - Quick Suite data: 8,000 characters

Using Amazon Quick Automate

Overview

Amazon Quick Automate is an AI-powered application that creates sophisticated automations using natural language or documentation. Amazon Quick Automate revolutionizes enterprise workflow by transforming complex processes into intelligent, adaptive automations.

Enterprises struggle to automate end-to-end business processes due to the growing complexity of their application ecosystems and the limitations of traditional automation tools. Amazon Quick Automate provides the easiest way for businesses to build, deploy, and continuously maintain workflow automations with the set of capabilities needed to support complex use cases that span multiple applications, data sources, and departments.

Customers can simply provide their process requirements by uploading process documents or simply describing their business process in a natural language. From the documents or description, Amazon Quick Automate creates a plan with steps and actions for your automation. You can expand upon the plan by adding additional steps or actions or modify the existing steps or actions. You can add additional steps by dragging actions or steps from an extensive Actions library.

Amazon Quick Automate's intelligent agents can seamlessly navigate and interact with web applications, make contextual decisions based on real-time data, execute API calls, and even generate custom code to handle complex scenarios. When human oversight is required, the system can create human-in-the-loop tasks for manual review or approval.

Key concepts

Automation groups

Automation groups create logical groupings for automations, allowing organizations to manage permissions and resources efficiently across different departments or projects. These groups provide a way to organize your automations and setup permissions based on business units, functional areas, or project teams.

By organizing automations into groups, you can apply consistent access controls, resource restrictions, and governance policies. This organizational structure supports scalable automation management as your organization's automation needs grow and evolve.

Project

A *project* is the primary container that encapsulates all metadata, assets, and contextual information related to an automation initiative. It serves as the foundational unit for organizing, authoring, testing, deploying, and maintaining automations across their lifecycle.

Lifecycle of a project

- **Create a project:** Define the automation's scope including its name, purpose, business case, and attach process documentation.
- **Build automation:** Within the project, build and test and automation.
- **Version and deploy:** Publish automation versions for testing or production. Projects maintain version history and deployment status.
- **Monitor and iterate:** Review logs, metrics, for each run and iterate on the automation. Each iteration creates traceable project-level change history.

Automations

An *automation* is a visual artifact that represents the executable logic of a process. It is created using the *Visual Designer* within , where you define the sequence of actions, decisions, and agent interactions needed to complete a business process. Each automation serves as the operational blueprint for how work is performed, combining AI-powered agents, system actions, and human inputs in a unified flow.

Actions

Actions are the building blocks for your automations. Each action performs a specific task, from reading data to interacting with external systems.

Agents

An *agent* is an AI-powered action that can understand natural language input, reason through scenarios, and perform complex tasks on your behalf. Agents bring intelligence and adaptability to automations by interpreting intent, handling variability, and executing actions dynamically.

- **UI agent** is a native agent that understands natural language instructions to perform complex browser actions. It can autonomously navigate websites, click, type, read data, and produce structured outputs optimized for downstream automation steps.
- **Custom agent** is an AI-powered action that processes natural language inputs to automate complex steps using integrated tool-calling capabilities. It primarily uses integrations as its tool interface, while offering extensibility to use Code as tool, and other native actions like human-in-the-loop task. The agent delivers structured, predictable outputs optimized for seamless integration into downstream automation steps.

Getting started

Amazon Quick Automate is a browser-based solution that requires no installation, has minimal IT dependencies, and is always up to date. Amazon Quick Automate is designed for **technical teams and power users** who build, test, and maintain intelligent automations across enterprise systems. These users typically have a strong understanding of business processes, system integrations, but may not necessarily be software developers. Amazon Quick Automate equips advanced users and technical teams with the flexibility, control, and depth needed to design and scale enterprise-grade automations. To get started with Amazon Quick Automate, simply access it through your Amazon Quick Suite and begin exploring the automation capabilities that best fit your needs.

Key capabilities

- Multi-agent architecture
 - Planning Agent designs automations from natural language inputs
 - UI Agent handles interface interactions
 - Specialized agents for process-specific tasks
 - Integration with
- Enterprise controls and governance
 - Granular access control management
 - Comprehensive activity logging

- Built-in version control
- Governance features for compliance
- Human-in-the-loop capabilities
- Case management
- Exception handling support
- Real-time monitoring
- Built in scalability

Setup tasks

Complete these setup tasks to prepare your Amazon Quick Automate environment for creating automations.

Sign up

Before you can use Amazon Quick Automate, you need to sign up for and create your account. The sign-up process establishes your initial workspace and provides access to Amazon Quick Automate.

Set up automation groups

Automation groups create logical groupings for automations, allowing organizations to manage permissions and resources efficiently across different departments or projects. These groups provide a way to organize your automations and setup permissions based on business units, functional areas, or project teams. By organizing automations into groups, you can apply consistent access controls, resource restrictions, and governance policies. This organizational structure supports scalable automation management as your organization's automation needs grow and evolve. After signing up, configure automation groups and permissions to organize your automations and control access across your organization. This step ensures that different departments and teams can work with automations appropriate to their roles and responsibilities.

Setting up proper automation groups and permissions from the beginning helps maintain security and compliance as your automation environment grows. You'll establish the organizational structure that supports efficient workflow management across multiple departments or projects.

Setup user roles

Amazon Quick Automate provides granular role based access: Viewer, Contributor, and Owner. Viewers can submit new project ideas, interact with human in the loop tasks, and access monitoring

dashboards. Contributors can build, test, and deploy automations. Owners can control automation access to systems through integrations and secure credentials as well as user access.

Configure integrations

If your automation interacts with external applications such as Salesforce or Jira via APIs, make sure all necessary integrations are configured. For more information, see [the section called “Action connectors”](#).

- Navigate to **Connections** → **Integrations** in the left panel to create new integrations.
- Currently, only integrations available under the **Actions** tab are supported in Amazon Quick Automate.
- Once an integration is created, associate it with the **Automation Group** where it will be used.
- The associated actions will then appear in the canvas. During deployment, you can select the appropriate connection to be used by the deployed automation.

Set up Virtual Private Cloud (VPC) connections

Amazon Quick Automate can access publicly hosted internet endpoints by default. For more complex setups, including accessing privately hosted websites, Amazon Quick Automate automation jobs can optionally be associated with your VPC. For more information, see [the section called “Setting up a VPC to use with Amazon Quick Suite”](#).

With this configuration in place, all outbound network traffic from the managed browser is routed through an ENI in your configured VPC. From there, you can send that traffic to privately hosted resources or use it to obtain a persistent IP for IP allow listing.

The VPC-connected Amazon Quick Automate setup provides:

- **Direct access** to privately hosted resources in the connected VPC
- **Extended access** to other privately hosted resources in peered, PrivateLinked, etc. VPCs
- **Internet access** through your managed Internet Gateway

Setup instructions

Step 1: Choose your VPC

Select an existing VPC or create a new one that includes private and public subnets in one or more Availability Zones, NAT and Internet Gateways.

Note

If creating a new VPC, the default selections in Amazon VPC creation wizards are sufficient.

Step 2: Setup connectivity to target resources

Bot traffic will egress through an ENI in your connected VPC. Choose the appropriate connectivity method:

For resources in the same VPC:

- No additional setup required

For resources in different VPCs, use one of these Amazon constructs:

- Amazon VPC Endpoint Services (PrivateLink) - Exposing individual service endpoints from one VPC to another
- Amazon Transit Gateway - Centralized management of a cluster of peered networks
- Amazon VPC Peering - Direct connection between a pair of VPCs

Step 3: Create Quick VPC Connection

Use Quick Admin tools to create a VPC Connection. This step creates two ENIs in the private subnet within your chosen VPC.

Step 4: Associate VPC Connection with Automation Group

Once the VPC Connection is created, associate it with your automation group from the VPC connection section in the automation group page. After this association is configured, all subsequent automations run within the specified group will use the VPC Connection for network access.

Create project

This section walks you through the process of creating a new automation project in Amazon Quick Automate.

To create a new project

1. Click the **Create Project** button to begin.
2. Fill in the project details in the Provide project details section:
 - **Name:** Enter a descriptive name for your automation project.
 - **Group:** Select the appropriate group for your automation.
 - **Description** (optional): Provide additional context about the project's purpose.
 - **Upload Documents** (optional): Attach existing documentation to jumpstart your automation.
3. Click **Next: Business case** to proceed to the business case section.

Defining the business case (Section is optional)

This section is optional but recommended for tracking ROI and prioritization. Complete the following fields:

- **Hours saved per case:** Estimate time savings per instance.
- **Cases per year:** Enter the expected annual volume.
- **Hours saved per year:** This may calculate automatically based on previous entries.
- **Project priority:** Set the relative importance of this automation.
- **Target launch date:** Define the target implementation date.

Click **Create** to finalize and generate your new automation project.

Note

Required fields are marked accordingly in the interface. Optional fields provide additional context but can be completed later if needed.

After creating your project, you can access a summary page with Overview, launch goals and document section will open for review. This comprehensive layout allows you to quickly grasp the essential information about your project or resource. You will have Summary (default selected), Versions and Deployments tabs.

In Summary, the Overview section provides a high-level description and launch Goals. The Documents section contains relevant document uploaded. On the right-hand side of the page, you'll notice a Status panel. This panel offers real-time updates on the current state of your project. It may include information such as deployment status, health checks, or any ongoing processes.

Building automations

Amazon Quick Automate provides multiple methods for creating automations to suit different needs and skill levels. This section describes the available creation methods and tools.

Prerequisites

Before creating an automation, you need to create a project within a specific automation group. The automation group will control what integrations and credentials are available to you when building that automation. Ensure the required integrations and credentials are set up for any external systems you plan to interact with in your automation.

Once you have created a project, click the **Start building** button to get started.

Note

You must have owner permissions to manage integrations and credentials for an automation group.

Creation methods

Choose from three approaches to build your automation. If you're new to Amazon Quick Automate, we recommend:

- Start by exploring sample automations
- Try creating a simple automation using natural language
- Experiment with manual editing once you're familiar with the basics

Using sample automations

Get started quickly by exploring pre-built samples that demonstrate common automation patterns. These samples provide hands-on examples of different automation capabilities and serve as

learning resources to help you become familiar with the capabilities of Amazon Quick Automate. To use a sample automation:

- Search to explore different samples by industry.
- Click on the sample to view an overview of the automation.
 - If any integrations are required, they will be listed as prerequisites.
- Click **Start with sample** to open the automation in the canvas.
- You will now be able to view the automation, edit it, and test it.

Creating with natural language

The Automation Assistant provides an AI-powered approach to creating automations through:

- Chat interface - Describe your process conversationally in natural language
- Document upload - Import any documentation that describes your process steps and requirements

When you provide input through either method, the Automation Assistant will generate your automation through a two-step guided experience:

- **High-level plan**
 - First, the Assistant analyzes your requirements and generates a high-level automation plan. The plan consists of **process steps** that logically group the related actions as part of the automation. Each process step initially contains natural language instructions.
 - Review and edit the instructions for each process step and then click **Generate** to create the low-level actions for that step.
- **Low-level actions**
 - Next, the Assistant will build the low-level actions and process logic based on the step instructions.
 - Review and edit the generated actions. Click **Run** or **Debug** to begin testing.

Best practices for providing natural language inputs:

- Specify the step by step process as if you were training a new hire
- Include exact details needed to perform each step:

- Email address if sending a message
- File name and location if uploading/downloading a document
- URL if navigating to a website

Creating from scratch (blank automation)

Build automations manually using the visual designer interface:

- Navigate to your project summary
- Click **Start building**
- Select **Skip** to access the canvas

The designer interface provides several key components:

- **Canvas** - Visualizes the automation. You can zoom in and out, expand and collapse process steps, and reorder actions on the canvas.
- **Actions panel** - The actions panel has a comprehensive list of all the automation actions made available within your automation group. You can search and filter to find actions organized by category (Agents, Process flow, Web browser, etc.). Drag-and-drop actions onto the canvas to build your automation. You can also click the plus sign found when hovering between actions to add a new action directly in place on the canvas.
- **Properties panel** - Once you add an action or click on an existing one in your process, the properties panel allows you to configure the input and output parameters that control how that action will behave.
- **Settings panel** - Allows you to create and edit runtime configurations and explore the credentials available to your automation.
- **Variables panel** - Displays all the variables used within your automation. Creating a variable is seamless by adding a new reference directly in your automation.

Runtime configurations and variables

Runtime configurations allow you to create easily editable settings that are used by your automation. Examples of runtime configurations include:

- Website URLs

- File names and locations
- Email configurations like subjects and sender lists

Variables are used to store and pass information between the actions in your automation when it runs. Add a new variable to your automation by simply providing a new reference name in an output property. Common variable types include:

- String - Plain text (e.g., name, description)
- Number - Numeric value (e.g., quantity, score)
- Boolean - True/False value
- Array - Collection or list of items
- Object - Key-value pairs
- File - Documents and media (e.g., PDFs, images)
- Data table - Spreadsheet-like data with rows and columns

Editing automations

Once you've created your automation, you can modify it using either the Automation Assistant or by directly editing on the canvas.

Using Automation Assistant

The Automation Assistant helps you make changes through natural language. You can:

- Ask for updates to the entire automation, specific steps or specific actions
 - Focus the chat on specific steps by selecting them first
- Ask questions and get in-product help and suggestions based on your actual automation
- Ask for help writing expressions and code based on your needs

To edit with the Assistant:

- Click **Build with Assistant** in the toolbar
- Describe the changes you want to make
- Review and confirm the suggested modifications

- The chat includes a summary of changes. Click **Reject changes** to revert to the previous version.
- Test the updated automation

Editing on canvas

Make changes directly in the visual designer:

- Select the step or action to modify
- Use the properties panel to adjust settings
- Rearrange steps using drag-and-drop

Note

Every action has a menu of options with common actions like duplicating the action or deleting it.

Best practices for editing:

- Make incremental changes
- Test after each significant modification

Note

Changes are saved automatically to the live version. If you want to deploy those changes, commit and deploy the updated version.

Tip

If you encounter issues while building:

- Use the Automation Assistant to get help
- Check the action documentation for specific requirements

Managing automation versions

Amazon Quick Automate provides version control capabilities to help you track and maintain the history of updates to your automations and easily restore previous versions.

Key concepts

- **Live version** - The current working copy of your automation that you can edit in the canvas. Changes are auto-saved to the live version as you edit.
- **Committed versions** - Read-only snapshots of your automation which can be deployed. Commit a new version of your automation to keep track of significant updates and be able to revert back if needed.
- **Deployed versions** - Committed versions can be deployed and activated to run on a schedule.

Note

Each version maintains its own runtime configuration values. Changes to the Live version's runtime settings don't affect previously committed or deployed versions.

Committing versions

To commit a version:

- Click **Commit** in the canvas toolbar
- Add a descriptive note explaining your changes
- Choose how to increment the version number:
 - Minor version (e.g., 1.05 → 1.06) for smaller updates (default)
 - Major version (e.g., 1.0 → 2.0) for significant changes
- Click **Commit**

Important

You cannot commit a version if the automation has any validation errors. Resolve all errors before attempting to commit.

When you commit a version:

- The committed version becomes a read-only snapshot
- The current runtime configuration is saved with the committed version
- A new Live version is created based on that snapshot to continue editing

Important

Only committed versions can be deployed. The Live version must be committed first before it can be made deployed.

Viewing versions

View versions in two places:

- **Version dropdown in canvas** - Defaulted to the Live version. Choose any previous version from the dropdown to view in read-only mode on the canvas.
- **Versions tab** - Found within Project details. Shows a complete version history including version number, when the version was committed, the user that committed the version, and any version notes. Click on the actions menu to view version details or deploy the version.

Restoring previous versions

To restore a previous version and continue editing it as the Live version:

- Select the version from the version dropdown above canvas
- Choose **Restore live version**
- The selected version replaces your current Live version
- Continue editing the Live version to make any needed adjustments
- Commit as a new version to take a snapshot or deploy the updates

Note

Restoring a version creates a new working copy but doesn't delete any version history.

Best practices

- Add clear version notes to track the purpose of changes
- Commit versions after significant updates
- Test changes before committing
- Review runtime configuration before committing

Automation actions

Actions are the building blocks for your automations. Each action performs a specific task, from reading data to interacting with external systems. This section describes all available actions and how to use them effectively.

Action Categories

- **Agents** - AI-powered actions that use natural language to perform complex tasks like web automation and multi-tool orchestration
- **Process flow** - Control the flow of your automation with steps, loops, and conditions
- **General** - Common utility actions
- **Web browser** - Interact with websites through clicking, typing, and extracting data
- **Microsoft Excel** - Read, write, and manipulate Excel spreadsheets
- **Data tables** - Work with structured data in memory for processing and analysis
- **Orchestration** - Manage work items through cases and coordinate human-in-the-loop tasks
- **Exception handling** - Define how your automation responds to different types of errors
- **Files** - Manage files used within your automation
- **PDF** - Create, modify, and extract content from PDF documents
- **Code** - Write custom Python code to extend automation capabilities
- **Integrations** - Connect with Amazon services and external applications like Salesforce, ServiceNow, and Microsoft 365

Agents

- **UI agent** - AI agent for web browser tasks. Used for dynamic and intelligent web automation. Simply write instructions to have it navigate websites, extract data, and generate structured outputs.
- **Custom agent** - AI agent for complex tasks. Create an agent that can understand instructions in natural language and take actions using available tools. Used for tasks that require reasoning, judgement, and dynamic planning.

UI agents

UI agent is a native agent that understands natural language instructions to perform complex browser actions. It can autonomously navigate websites, click, type, read data, and produce structured outputs optimized for downstream automation steps. Example use cases include summarizing products on a webpage or fetching data by navigating websites.

Properties

Title

Name of the step/UI agent

Instructions

In this field you write the prompt for the agent in natural language. Best practices while writing the prompt:

- Be clear and explicit about what you want.
- Structure the prompt. Start with mentioning the 'Task' or 'Role' first and then 'Instructions' to achieve the task with numbered steps
- Add constraints (e.g., only review the products section) and specify when to stop/end (e.g., stop when you find the relevant info)
- Provide positive and negative (don't do this) examples
- Specify length requirements (e.g., less than 100 words) or output format (e.g., date in MM/DD/YY format) clearly

Wrap the text in triple quotes (""") to write multiline prompts. For example:

```
"""Task: Locate the company's latest annual report.
```

- * Visit the provided URL.
- * Look for the annual report. The report may be titled 'Annual Report', 'Financial Report', 'Year in Review', or similar variations...""

Structured Output (optional)

Agent Response: Name of the variable to assign the output of this operation

How to configure structured output fields

Adding fields

- Click Add field to create a new output field
- Enter the Output name - this becomes the JSON property name
- Select the Type from the dropdown
- Check Required if the field must always be present
- Add a Description to guide the AI agent

Field types

- *String* - Text values (Names, descriptions, summaries)
- *Number* - Numeric values (Counts, scores, percentages)
- *Boolean* - True/false values (Status flags, yes/no questions)
- *Object* - Nested structure (Complex data groupings)
- *Array* - List of items (Tags, categories, multiple values)
- *File* - File references (Document attachments, images)
- *Data table* - Tabular data (Structured datasets, reports)

Working with complex types

Objects and Arrays can contain nested fields:

- Click the expand arrow (►) next to Object or Array fields
- Use Add field within the nested structure
- Keep nesting to 2-3 levels maximum for optimal performance

Example configuration

Here's a simple configuration for summarizing customer feedback:

```
{
  "orderId": "12345",
  "numberOfOrders": 3,
  "hasShipped": true,
  "orderDetails": {
    "quantity": 2,
    "productName": "ABC",
  },
  "tags": ["electronics", "urgent"]
}
```

This structure would be configured as:

- orderId (String, required)
- numberOfOrders (Number, required)
- hasShipped (Boolean, required)
- orderDetails (Object, required)
 - quantity (Number, required)
 - productName (String, required)
- tags (Array of Strings, Optional)

Best practices

- Use descriptive field names - Help the AI understand what data to extract
- Add clear descriptions - Provide context for complex fields
- Mark critical fields as required - Ensure essential data is always present
- Limit nesting depth - Keep structures simple for better performance
- Test your configuration - Verify the output matches your expectations by running the agent step and verifying the response.

Important notes

- JSON Knowledge: Unfamiliar with JSON? Learn the basics at json.org

- **No validation:** Currently, the system doesn't validate output structure - ensure your automation handles missing or malformed data

Custom agents

Custom agent is an intelligent action that processes natural language inputs to automate complex steps using integrated tool-calling capabilities. It primarily uses integrations as its tool interface, while offering extensibility to use Code as tool, and other native actions like human-in-the-loop task. The agent delivers structured, predictable outputs optimized for seamless integration into downstream automation steps.

Properties

Title

Name of the step/custom agent

Mode

A mode defines how the Agent operates based on your use case. The three available modes are: Fast, Pro, and Custom. Fast is best for simple tasks like summarization, classification, and high-volume automations, and Pro is ideal for complex tasks that involve reasoning and orchestration of multiple tools or actions. Fast and Pro are fully managed modes that require no pre-setup needed in advance. In Custom Mode, you'll need an Bedrock Converse connector and can select the model you want to use. This is ideal if you already have a prompt fine-tuned for a particular Bedrock model, specifically need a particular Bedrock model for the Agent, or want to include your own custom or fine-tuned model hosted on Bedrock. In Custom Mode, since you bring your own model from Bedrock via a connector, model inference is billed separately to the account associated with that Bedrock connector.

Instructions

In this field you write the prompt for the agent in natural language. Best practices while writing the prompt:

- Be clear and explicit about what you want.
- Structure the prompt. Start with mentioning the 'Task' or 'Role' first and then 'Instructions' to achieve the task with numbered steps
- To improve tool-call accuracy and guide the Agent, clearly specify in the prompt which tool to use at each step, if applicable.

- Specify length requirements (e.g., less than 100 words) or output format (e.g., date in MM/DD/YY format) clearly

Wrap the text in triple quotes (""") to write multiline prompts. For example:

```
""""You are content summarization agent.  
Summarize the last two paragraphs of the provided text, focusing only on the main  
conclusion.""""
```

Tools (Optional)

A tool enables the AI agent to interact with external systems or perform specific tasks

General tools

Create user task

If enabled, this tool allows the Agent to trigger a Human-in-the-Loop (HITL) task whenever it gets stuck and needs assistance during execution. The Agent will pause and wait for human input, then resume once the required information is provided. The HITL task will be visible in the task center. For best results, the author can specify in the prompt exactly when the Agent should invoke HITL.

Integrations

If you've added specific connectors—such as Salesforce, MS Exchange, or Bedrock—to your automation group, their corresponding actions will appear here. The author can then select the relevant actions to use as tools for the Agent. For optimal performance, it's recommended to limit the Agent to 3–5 tools.

Structured Output (optional)

Configure your AI agent to return structured JSON output that can be easily processed in subsequent steps. This feature is ideal for text summarization, report generation, data transformation, and extracting statistics from unstructured content. This is an optional field. If you don't define structured output, the agent will by default return output in natural language.

Agent response: Name of the variable to assign the output of this operation

Note

The structured output configuration for Custom agents follows the same format as UI agents. Refer to the UI agent structured output section for detailed configuration instructions.

Custom agent testing

Users can test the agent independently of the full automation to validate behavior, debug prompts, and iterate faster.

Start test

- Hover on the agent card, a separate run button will show up on top of the card
- Click on the button to unit test this specific agent
- A variable collection window will pop up and automatically detect any variables used in the prompt/instruction
 - A preview of the prompt of this agent is displayed and highlights all the auto detected variables
 - Input put values for each variables before kick off unit test. Similar to all other expression fields of Amazon Quick Automate, the value of a given variable has to be valid expression syntax. Otherwise, an error will show up on the screen and prevent user to start test,

Test running

Users can see the execution log feed in the audit panel on the right side. The experience is the same as running the whole automation.

After test run

- User can see the input variables and output result at the Watch Variables` tab below the log feed.
- User can see basic metric card above the log feed (total time used and tools used).

Process flow

- **Process step** - Groups logically related actions. Used for organizing your process.

- **If decision** - Checks a true/false condition. The process is split into two branches based on the result.
- **Loop through items** - Repeats actions for each item. Used for processing multiple items like emails, files, or data rows one at a time.
- **Loop while true** - Repeats while a condition is met. The condition is checked before each loop and must be True for the loop to continue.
- **End process** - Ends the current run successfully. Optionally, provide a custom message or value as the final outcome of the process.

General

- **Save value** - Assigns a value to a variable. Used to save information that you can reference later in your automation.
- **Write log message** - Adds a message to the logs. Useful for tracking progress, troubleshooting issues, or recording important events.
- **Wait** - Pauses the automation. After the specified amount of time, the automation will continue with the next action.

Web browser

- **Start browser session** - Opens a new browser session. Actions placed inside will interact with this browser. The browser closes automatically when finished.
- **Go to webpage** - Navigates to a specific URL. Used to go to a new webpage in the current browser tab.
- **Click** - Clicks on a webpage element. Used to interact with buttons, links, or other clickable elements.
- **Enter text** - Types text into an input field. Used to fill in forms, search boxes, or other text inputs on webpages.
- **Enter credential** - Types a username or password. Used to securely login to a website using saved credentials.
- **Select from dropdown** - Chooses a value from a dropdown. Used to make selections in forms, filters, or other dropdown menus on webpages.
- **Extract text** - Reads a value from a webpage. Used to capture the text from an individual field.

- **Extract table** - Reads data from a webpage table. Used to capture structured information found in tables.
- **Take screenshot** - Captures an image of the webpage. The screenshot is saved to a file to be used later in your process.
- **Visual Q&A** - Answer questions about a webpage. Uses AI to visually analyze the page and answer true/false questions about the content.
- **Click and upload files** - Uploads files to a webpage. Used to automate file pickers by clicking on a file upload button (like 'Choose file', 'Browse', or 'Upload') and choosing the files to upload.
- **Enter keystroke** - Sends a key command or shortcut. Used to automate typing individual keys or key combinations on a webpage.
- **Refresh webpage** - Reloads the current webpage. Some pages may redirect on refresh.
- **Restart browser** - Closes and reopens the browser. A new session will be started from a blank page.
- **Wait for element** - Checks for an element to appear. Used to wait for changes or validate conditions on a webpage. Returns true if the element is found within the maximum wait time, false if not.
- **Find matching elements** - Gets elements matching a pattern. Used to get a list of similar elements on a webpage that you can process one by one - for example, reading text from each search result or clicking each hyperlink.
- **Wait for new window** - Checks for a new window to open. Used after clicking a link or button that opens a new child window in the current browser session. Outputs a window identifier you can use to switch to or close the window in subsequent steps.
- **Switch window** - Changes browser windows. After switching, subsequent browser actions will interact with that window.
- **Close window** - Closes a browser window. Once closed, the automation returns to the main window of the current browser session.
- **Save to clipboard** - Copies text to the clipboard. Used to save text that you can paste later on. Existing clipboard text will be replaced.
- **Paste from clipboard** - Inserts text from the clipboard. Used to paste copied text into webpage fields.

Excel

- **Create new workbook** - Creates a blank Excel file. Outputs a workbook identifier you can use to update the spreadsheet in subsequent steps.
- **Open existing workbook** - Loads an existing Excel file. Outputs a workbook identifier you can use to update the spreadsheet in subsequent steps. Supports only Excel (.xlsx) files.
- **Save workbook** - Saves updates to an Excel file. Allows you to update the file name and file type for the saved workbook.
- **Read sheet** - Gets data from a range of cells. The output is stored in a data table variable containing rows and columns.
- **Read cell** - Gets the value from a cell. Used to read an individual cell value from a worksheet.
- **Write to sheet** - Outputs a data table to a sheet. Used to write a range of rows and columns to a worksheet.
- **Write to cell** - Outputs a value to a cell. Used to update individual cells in a worksheet.
- **Write new row** - Adds a row of data to the sheet. The new row will be appended to the end of the existing data.
- **Create new sheet** - Adds a blank worksheet. The new sheet is added to the end of the workbook.
- **Copy sheet** - Creates a copy of the worksheet. The new sheet will be created within the same workbook.
- **Rename sheet** - Updates the name of a worksheet.
- **Delete sheet** - Removes a sheet from a workbook. Cannot delete the last remaining sheet of a workbook.
- **Set cell color** - Updates cell background color. Used to highlight a specific range of cells.
- **Get cell color** - Reads the cell background color. Outputs the color in RGB hex code format (e.g., "FF0000" for red).
- **Hide rows** - Makes rows hidden in the sheet. The data remains intact but will not be visible.
- **Unhide rows** - Makes hidden rows visible. Used to show previously hidden rows.

Data tables

- **Sort table** - Orders a table by a column. Used to organize your data in ascending or descending order.

- **Filter table** - Keeps rows matching a criteria. Used to extract relevant rows from a larger data set.
- **Lookup value** - Searches a value in a table. Used to lookup a value in one column and get the corresponding value from another column in the same row.
- **Add columns** - Creates new columns in a table. Newly added columns are appended to the end of the existing table.
- **Remove columns** - Deletes columns from a table. Outputs a table with all of the remaining columns.
- **Keep columns** - Drops extra columns from a table. Used to select a specific subset of columns you want to keep.
- **Add new row** - Appends a new row to a table. The new row can be created with specific values or as a blank row and will be added to the bottom of the table.
- **Remove rows** - Deletes rows from a table. Outputs a table with all of the remaining rows.
- **Remove duplicates** - Deletes duplicate rows. Used to create a dataset of unique rows based on specific columns.
- **Append tables** - Combines the rows of two tables. Used to add data from one table to another.
- **Create new table** - Creates an empty table. Used to set up a table with required columns to add rows to in subsequent steps. The new table will have no rows.
- **Convert text to table** - Creates a table from text. Used to convert delimited text such as comma-separated values into a table.
- **Convert table to HTML** - Creates an HTML formatted table. Used to output your table as formatted text for a document, email, and more.

Orchestration

Orchestration enables you to automate complex end-to-end processes, coordinating automated tasks with human decisions where needed. The orchestration capabilities in Amazon Quick Automate include case management for tracking work items and human-in-the-loop integration for incorporating human judgment.

Case management

Cases track individual items of work for your automation to process such as a new claim, an IT support ticket, a customer complaint, or a new invoice. Each case has:

- A unique reference name (e.g., "Order-12345")
- A case type for grouping related work (e.g., "Sales orders")
- Custom data stored as key-value pairs
- A defined lifecycle with status tracking and an audit trail
- Associated human-in-the-loop tasks if applicable

Case status lifecycle

To track the progress of work through your automation, cases follow a defined lifecycle with specific status transitions:

- **Ready** - The case has been created and is waiting to be picked up for processing. This is the initial state of all new cases.
- **In Progress** - The case is actively being processed by the automation. You can monitor real-time progress through the logs.
- **Pending Resolution** - The case is waiting for a human-in-the-loop task to be completed before processing can continue. Cases automatically move back to Ready status after tasks are resolved.
- **Successful** - The case has completed all processing steps without any exceptions.
- **Failed** - The case encountered one of the following exception types:
 - **Business Exception** - The case encountered a handled business rule violation and stopped processing. Detailed exception information is available in the logs.
 - **System Exception** - The case experienced a technical error and stopped processing. Detailed error information is available in the logs.

Actions

Amazon Quick Automate provides the following actions to create, process, and manage cases throughout their lifecycle.

Create new case

Creates a single case to track work. Each case has a reference name (e.g., "Order-12345") and stores custom data as key-value pairs. Cases are grouped by type (e.g., "Sales orders") for organization and reporting. Properties:

- **Title (optional)** - Action name shown in the process visualization (e.g., "Create sales order case")
- **Reference name (required)** - Business identifier used for searching and tracking cases in monitoring dashboards (e.g., "Order-12345")
- **Case type (required)** - Label that groups related cases by business purpose. Provides organization and metric tracking (e.g., "Sales orders", "Support tickets")
- **Custom data (required)** - Additional case information stored as key-value pairs. This data is visible to users and can be retrieved, updated, or appended to during case processing (e.g., {"Product": "ABC-12345", "Qty": "100", "Priority": "High"})

Create multiple cases

Bulk creates cases from a data table. Each row becomes a new case, inheriting the specified case type and using a designated column for reference names. All other columns become custom data fields. Properties:

- **Title (optional)** - Action name shown in the process visualization (e.g., "Create cases from invoice batch")
- **Case type (required)** - Label that groups related cases by business purpose (e.g., "Invoice processing")
- **Cases data table (required)** - Data table where each row creates a new case. All columns are stored as custom data (e.g., cases_table)
- **Reference column name (required)** - Name of the column containing the reference name identifier for each case (e.g., "Invoice number")

Process cases

Loops through cases in 'Ready' status. Automatically manages case status by moving cases to 'In Progress' during processing and marking them as successful or failed based on completion. When adding actions inside the Process cases loop, uncaught exceptions are handled for you, marking the case as failed with the associated error details and moving to the next case. The action also provides access to case data during processing. Properties:

- **Title (optional)** - Action name shown in the process visualization (e.g., "Process pending invoices")
- **Case type (required)** - Type of cases to process (e.g., "Invoice processing")
- **Current case (output)** - Variable containing the case being processed in each loop. Access case data using `current_case["custom_data"]["key_name"]`

Update case data

Saves changes to case information during processing. Can add new fields or modify existing data. Only available while a case is 'In Progress'. Used to track processing results and manage case lifecycle. Properties:

- `Title` (optional) - Action name shown in the process visualization (e.g., "Update case status")
- `Case ID` (required) - Unique identifier of the case to update (e.g., `current_case["case_id"]`)
- `Data to update` (required) - New or modified information to store as key-value pairs (e.g., `{"Status": "Approved", "ProcessedDate": "2024-01-20"}`)
- `Updated case` (output) - Variable containing the modified case. Access updated data using `updated_case["custom_data"]["key_name"]`

Search cases

Retrieves cases for reporting and analysis. Filter by case type, status, or reference name. Returns results as a data table containing case information and custom data. Properties:

- `Title` (optional) - Action name shown in the process visualization (e.g., "Find overdue cases")
- `Case type` (required) - Type of cases to search (e.g., "Support tickets")
- `Filter by status` (optional) - Filter for cases based on their current status (e.g., "Pending Resolution")
- `Filter by reference name` (optional) - Filter for cases with a specific reference name (e.g., "Ticket-12345")
- `Maximum results` (optional) - Limits the number of cases to return. Leave empty to return all cases up to 10,000 (e.g., 100)
- `Search results` (output) - Variable name that will store the data table of cases retrieved by the search (e.g., `cases_table`)

Human-in-the-Loop (HITL)

While many decisions can be automated, some scenarios require human judgment. Amazon Quick Automate's HITL capabilities allow you to seamlessly incorporate human decision-making into your automated processes. There are two ways to create HITL tasks:

- Using the Create user task action
- Using HITL chat through a Custom agent

Create user task

Creates a task for human input. Tasks appear in a list for users to review, make decisions, or provide information. Tasks can only be created when in a case processing loop while the case is in progress. Creating a task moves the case to 'Pending task' status and the automation continues with the next case. Once the task is resolved, the case returns to 'Ready' status for further processing. Properties:

- **Title (required)** - Subject line that appears in the user's task list. Should clearly describe what needs to be done (e.g., "Review loan application", "Approve expense report")
- **Description (required)** - Detailed instructions and context provided to users completing the task (e.g., "Please review the attached application details and either approve or reject based on the following criteria...")
- **Form (required)** - Interface for collecting structured input from users
- **Severity (required)** - Priority level to help users triage tasks:
 - **Critical** - Urgent attention required, highest priority
 - **Moderate** - Normal priority (default)
 - **Low** - Non-urgent task
- **Due in (required):**
 - **Number** - Quantity of time units (e.g., 24)
 - **Unit** - Hours or Days
- **Case ID (required)** - Unique identifier of the case this task is for (e.g., `current_case["case_id"]`)

Important

When creating a HITL task:

- The Process cases loop automatically updates the case status to Pending Resolution and moves on to the next case
- Do not add actions after creating a task inside the Process cases loop as they won't execute for the current case

Building forms

When creating a HITL task, you can build custom forms to collect exactly the information needed from human reviewers. The following components are available for form design:

- Text field - Single or multi-line text input for collecting user inputs (e.g., "Please provide reason for rejection in detail")
- Dropdown - Selection from predefined options (e.g., ["Approve", "Reject", "Escalate"])
- Date/time picker - Calendar-based input
- Media - Display images or documents for review
- Display text - Read-only information to provide context (e.g., case details, policy references)

HITL chat

For scenarios requiring more dynamic interaction, Amazon Quick Automate also supports HITL chat through Custom agents. This enables natural language conversation between the agent and human reviewer. The agent can create a HITL task when it needs human input. The case will remain In progress while the agent chats with the human reviewer. Once sufficient input is provided, the agent will continue.

```
# Example agent prompt for HITL chat
""
Review the data and if any values are outside normal ranges,
create a HITL task for human review. Include the specific
values and acceptable ranges in the task description.
""
```

Note

HITL chat tasks will be given a default severity (Moderate) and due in (24 hours).

Task resolution and case re-entry

After a human completes a HITL task, your automation needs to handle the case's return to processing. Understanding how to work with task resolution data is key to implementing effective human-in-the-loop workflows. When a task is resolved:

- The associated case automatically returns to Ready status
- The case becomes available for processing again through Process cases
- The `latest_task_resolution` field contains all form inputs and decisions

Example checking for HITL resolution and handling the case accordingly:

```
# Example handling HITL resolution
if current_case["latest_task_resolution"]:
    # Handle case returning from HITL
    resolution_data = current_case["latest_task_resolution"]["form_data"]
    if resolution_data["decision"] == "Approve":
        # Process approved case
    else:
        # Handle rejected case
else:
    # First time processing this case
```

Best practices for handling HITL resolution:

- Always check latest_task_resolution when processing cases
- Structure your logic to handle both new cases and those returning from HITL
- Consider implementing different processing paths based on human decisions
- Include error handling for missing or invalid resolution data

Exception handling

Exception handling allows you to manage errors and unexpected conditions in your automations. Amazon Quick Automate provides capabilities for creating, raising, and handling exceptions appropriately to ensure robust error management and process continuity.

Types of exceptions

Amazon Quick Automate supports two types of exceptions:

- Business exceptions - Used for expected errors from business rule violations such as:
 - Missing required data
 - Invalid input values
 - Policy violations
- System exceptions - Used for technical or infrastructure errors such as:
 - Connection failures
 - Service timeouts
 - Authentication errors

Actions

Raise exception

Interrupts the process flow with an exception. Used to trigger or throw an exception to either end the process or to take specific actions to handle the exception scenario. If you do not want to end the process immediately, add an Exception flow to your Process step to handle the exception. Properties:

- **Title** (optional) - Action name shown in the process visualization (e.g., "Raise validation error")
- **Create custom exception** (required) - When selected on, create a new exception with custom details. When off, raise or re-raise an existing exception stored in a variable
- **Exception type** (required) - Category of the exception (e.g., "Business exception", "System exception")
- **Exception reason** (required) - Short identifier used to categorize similar exceptions for reporting (e.g., "REQUIRED_FIELD_MISSING", "INVALID_FORMAT")
- **Exception message** (required) - Detailed description that will help with troubleshooting (e.g., "Customer email is required but not provided")
- **Exception to raise** (required when Create custom exception is off) - The exception you want to raise or re-raise, typically stored as a variable (e.g., new_exception)

Create custom exception

Creates an exception without raising it. Used to prepare an exception variable that can be raised later in subsequent steps. Properties:

- **Title** (optional) - Action name shown in the process visualization (e.g., "Create validation exception")
- **Exception type** (required) - Category of the exception (e.g., "Business exception", "System exception")
- **Exception reason** (required) - Short identifier used to categorize similar exceptions for reporting (e.g., "REQUIRED_FIELD_MISSING")
- **Exception message** (required) - Detailed description that will help with troubleshooting (e.g., "Customer email is required but not provided")
- **Error code** (optional) - Numerical code used to distinguish different types of errors (e.g., 404 for not found, 503 for service unavailable)
- **Caused by** (optional) - Reference to another exception that triggered this one, typically stored as a variable (e.g., original_exception)

- **New exception** (output) - Variable that will store the newly created exception (e.g., `new_exception`)

Exception flow

Defines how to handle exceptions that occur within a Process step. Used to create an alternative path to follow when exceptions occur in your process. After an exception is handled, the process continues with the next Step. Raise the exception again if you want to end the process. Add an Exception flow by clicking the **Exception flow** plus button found in the Process step. Properties:

- **Title** (optional) - Action name shown in the process visualization (e.g., "Handle validation errors")
- **Exception reference** (required) - Variable name to refer to the exception within the Exception flow. Access exception details using `variable["property"]` where property can be: type, reason, message, code, or `caused_by` (e.g., `error`)

Note

Exception flows can only be added to Process steps. They are not found otherwise in the Actions panel.

Working with exceptions

Exception properties

When handling exceptions, you can access these properties:

- `type` - Category of the exception ("Business exception" or "System exception")
- `reason` - Short identifier for the exception
- `message` - Detailed description of the error
- `code` - Numerical error code if provided
- `caused_by` - Original exception that triggered this one if applicable

Best practices

- Use descriptive exception reasons for easy categorization and reporting
- Include detailed messages for troubleshooting

- Implement retry logic for temporary system exceptions
- Consider human intervention for business exceptions
- Consider the impact on case status when handling exceptions
- Ensure critical cleanup operations occur even after exceptions such as restarting a browser
- Monitor exception frequencies and patterns
- Review exception logs regularly
- Re-raise exceptions once handled if you do not want to proceed with the remaining steps

Files

- **Duplicate file** - Creates a copy of a file. Used to create a new version of the same file.
- **Zip files** - Creates a compressed zip file. Used to combine multiple related files into a single package.
- **Unzip file** - Extracts files from a zip. Used to unpack individual files from a single zip file.

PDF

- **Create PDF file** - Generates a PDF file from text. Used to create a new PDF document with formatted contents.
- **Read PDF text** - Extracts text from a PDF file. Used to get the content of text-based PDF documents. Does not extract text from scanned documents or preserve the structure of the extracted text.
- **Split PDF file** - Extracts pages from a PDF. Used to create a new PDF from selected pages of an original file.
- **Merge PDF files** - Combines multiple PDF files. Used to create a single PDF from several documents by appending their pages.
- **Get page count** - Gets the count of pages in a PDF. The number of pages can be used in subsequent steps to extract text from specific pages or split the PDF.

Code

- **Single-line expression** - Runs a one-line Python expression. Used to perform simple operations that require no output such as adding items to a list (`list.append`) or updating dictionary values (`dictionary.update`).

- **Custom code block** - Runs a multi-line Python function. Used for more complex operations or custom logic that can't be achieved with standard actions.

Integrations

Amazon Services

Amazon S3

actions allow you to store and retrieve files and data.

Available Actions

- **Download File** - Downloads a file from Amazon S3 and returns a file object handle.
- **Upload File** - Uploads a file object handle to Amazon S3.
- **Copy Object** - Creates a copy of an object that is already stored in Amazon S3 from one source file to a destination location.
- **Delete Object** - Deletes a specified object from an S3 bucket.
- **List Objects** - Lists all of the objects in the given bucket. Can be filtered using a prefix.

actions allow you to connect to foundation models (FMs) from leading AI companies through a single API.

- **Basic model converse** - Sends a prompt to an AI model. Use natural language to generate text, analyze content, and get AI-powered responses.
- **Advanced model converse** - Prompts model with more options. Advanced options include custom message formats, tool integrations, and granular control over model behavior.

Agents

Agent actions allows you to invoke AI assistants that can connect to your data sources, APIs, and backend systems.

- **Basic agent invoke** - Sends a request to an AI agent. Use for simple interactions that don't require file attachments, user input, or custom configurations.
- **Advanced agent invoke** - Calls an agent with more options. Used for advanced interactions that require file attachments, user input, or custom configurations.

Data Automation

Bedrock Data Automation (BDA) simplifies the process of extracting valuable insights from unstructured content using generative AI.

Available Actions

- **Invoke Data Automation Async** - Process data with a data automation project or blueprints. Starts asynchronous data processing and returns an invocation ID.
- **Get Data Automation Status** - Gets details about a data automation invocation status and output location.

Amazon Textract

Amazon Textract helps you add document text detection and analysis to your automations.

- **Basic document query** - Extracts data based on questions. Each query contains the question you want to ask and the alias you want to associate. Textract provides the text answer to each question and a confidence score.
- **Basic document analysis** - Extracts data from documents. Automatically detects and extracts forms (key-value pairs), tables (structured data), and text while preserving document layout and data relationships.

is a natural language processing (NLP) service for gaining insight into document content. It helps analyze text to extract key phrases, identify sentiment, and classify documents.

Available Actions

- **Classify Document** - Creates a classification request to analyze a single document in real-time using a custom model endpoint. Supports text input or document files (PDF, Word, or image).
- **Detect Key Phrases** - Detects the key noun phrases found in the text, identifying important concepts and topics within the content.

External Applications

Microsoft Outlook

The Microsoft Outlook actions allow your automations to manage emails and calendars.

- Email
 - **Send User Email** - Send an email message on behalf of a specific user.
 - **Reply to Email** - Reply to the sender of a message.
 - **List Folder Messages** - Get the messages in a specific folder of the signed-in user's mailbox.
 - **List User Mails** - Get the messages in the signed-in user's mailbox (including Deleted Items and Clutter folders).
 - **View Email** - Retrieve the properties and relationships of a specific message object.
 - **Update Email** - Update the properties of a message object.
 - **Delete Email** - Delete a message in the specified user's mailbox.
 - **List Email Attachments** - Retrieve a list of attachment objects attached to a message.
 - **Get Attachment** - Read the properties, relationships, or raw contents of an attachment.
- Folders
 - **List Mail Folders** - Get the mail folder collection directly under the root folder of the signed-in user.
 - **Move Email to Folder** - Move a message to another folder within the specified user's mailbox.
 - **Get Mailbox Settings** - Get the user's mailbox settings, such as automatic replies, date format, and time zone.
- Calendar
 - **List Calendar Events** - Fetch a list of calendar events for the specified user or resource from Microsoft Outlook.
 - **Create Calendar Event** - Create an event in the user's default calendar or specified calendar.
 - **Update Calendar Event** - Update the properties of an existing event object.
- Other
 - **List Places** - Get a collection of place objects (rooms or room lists) defined in the tenant.
 - **List Users** - Retrieve a list of users in the organization from Microsoft Outlook with their basic information.
 - **List Contacts** - Get a contact collection from the default contacts folder of the signed-in user.

Microsoft Teams

The Microsoft Teams actions allow your automations to communicate with team members through channels and chat messages, as well as manage meetings.

- Messages
 - **List Chats** - View all chat conversations
 - **Create Chat** - Start a new chat conversation
 - **Get Individual Chat** - View details of a specific chat
 - **List All Chat Messages** - View messages in a chat conversation
 - **Send Chat Message** - Send a new message in a chat
- Teams
 - **List Teams** - View all teams accessible to the authenticated user
 - **Get Team** - View detailed information about a specific team
 - **List All Team Members** - View members of a team
 - **Invite User To Team** - Add a new member to a team
- Channels
 - **List All Channels** - View all channels within a specific team
 - **Create Channel** - Create a new channel within a team
 - **Get Channel** - View detailed information about a specific channel
 - **List All Channel Messages** - View messages posted in a channel
 - **Send Channel Message** - Post a new message to a channel
 - **Reply To Channel Message** - Add a reply to an existing channel message
 - **Invite Channel Member** - Add a member to a specific channel
- Meetings
 - **Create Online Teams Meeting** - Schedule a new online meeting
 - **Get Online Teams Meeting** - View details of a specific meeting
 - **List All Recordings** - View recordings from a meeting
 - **List All Transcripts** - View transcripts from a meeting
- Users
 - **List Users** - View list of users in the organization

Microsoft SharePoint

The Microsoft SharePoint allows you to centrally store, manage, and share documents and information through your automations.

- Lists and items
 - **Get List** - Get a list of operations associated with a list
 - **Create Item** - Create a new list Item in a list
 - **Update Item** - Update the properties on a listItem
 - **Delete Item** - Removes an item from a list
 - **Get Item** - Returns the metadata for an item in a list
 - **List Item** - Get the collection of items in a list
- Sheets and files
 - **Add Sheet** - Add a new worksheet to the workbook
 - **Delete Sheet** - Remove worksheet from workbook
 - **Update Sheet** - Edit worksheet properties (rename, visibility)
 - **List Sheets** - Retrieve a list of worksheet objects
 - **Read Sheet** - Retrieve properties of a worksheet
 - **Upload File** - Creates or updates a workbook file
- Ranges and cells
 - **Read Range** - Retrieves values from a specified range
 - **Delete Range** - Deletes cells in a specified range
 - **Write Range** - Updates values in a specified range
 - **Get Used Range** - Returns the used range of the given worksheet
 - **Clear Range** - Clear range values including format, fill, and border
 - **Read Cell** - Gets the value from a specific cell
 - **Write Cell** - Sets the value of a specific cell

Microsoft OneDrive

Microsoft OneDrive actions allow you to store, sync, and share files. You can interact directly with Microsoft Excel files saved to OneDrive for spreadsheet actions.

- File management
 - **Get Item** - View details of a specific file or folder
 - **Get Drive** - View OneDrive account details and storage information
 - **List Items** - View all files and folders in a drive

- **Create Item / Create Folder** - Create a new folder or item in OneDrive
- **List Children** - View items contained within a specific folder
- **Delete Item** - Remove a file or folder
- **Update Item / Move Item** - Update metadata or move files between folders
- **Copy Item** - Create a copy of a file or folder in a different location
- **Add Permissions** - Configure sharing and access permissions for files
- **Upload File** - Add new files to OneDrive
- **Excel**
 - **Add Sheet** - Create a new worksheet in an Excel workbook
 - **Clear Range** - Remove content and formatting from a range of cells
 - **Delete Range** - Remove cells, rows, or columns from a worksheet
 - **Delete Sheet** - Remove a worksheet from a workbook
 - **List Sheets** - View all worksheets in a workbook
 - **Read Cell** - Get the value of a specific cell
 - **Read Range** - Get values from a range of cells
 - **Read Sheet** - View contents of an entire worksheet
 - **Read Used Range** - Get the range of cells that contain data
 - **Update Sheet** - Modify worksheet properties
 - **Write Cell** - Set the value of a specific cell
 - **Write Range** - Set values for a range of cells

Salesforce

Salesforce actions allow you to automate the process of managing customer relationships across cases, opportunities, leads, campaigns, contacts and more.

- **Cases**
 - **Get Case List** - Gets a list of customer support cases from Salesforce, including their metadata and recent items.
 - **Get Case** - View details of a specific case
 - **Create Case** - Create a new support case
 - **Delete Case** - Remove an existing case

- **Update Case** - Edit case details and information
- Opportunities, Leads, Campaigns, Feed Items, Contacts, Users, System

Jira

Jira actions allow you to automate managing issues, projects, sprints and more.

- Issues, Comments, Projects, Sprints, Users

ServiceNow

ServiceNow actions allow you to automate IT service management of incidents, change requests, and more.

- Incidents, Change Requests, Knowledge Base Articles, Problem Records, Attachments, System

SAP

SAP actions allow you to automate processes involving business partner data, material stocks, bills of material, product masters, and physical inventory.

- SAP Business Partners, SAP Material Stock, SAP Bill Of Material, SAP Product Master, SAP Physical Inventory

Custom REST APIs

Actions for working with any REST API endpoint.

- **Get request** - Retrieves data from a REST API. Used to fetch information from web services and APIs.
- **Post request** - Sends data to a REST API. Used to create or submit information to web services and APIs.
- **Put request** - Replaces data in a REST API endpoint. Used to update existing information through web services and APIs.

Testing automations

Amazon Quick Automate provides comprehensive capabilities for testing your automations. This section describes the tools and features available for validating your automation behavior.

Test modes

Choose from two modes to test your automation:

- **Run mode** - Executes the automation from start to finish without stopping. Any breakpoints in the automation will be ignored. This mode is best for testing the complete flow of your automation.
- **Debug mode** - Executes the automation and pauses at any breakpoints you've set or if an exception occurs. This allows you to inspect the automation state, monitor variables, and validate behavior at specific points in your process.

Test pane

The Test pane includes includes two main tabs that help you monitor and control your automation test:

Setup tab

The Setup tab provides tools for managing your test configuration:

- **Breakpoint manager** - Lists all breakpoints currently set in your automation. Remove breakpoints directly from this view without returning to the canvas.
- **Test navigation** - Quick access links to view:
 - **View runs** - See all automation runs from testing
 - **View cases** - View cases created during testing
 - **View tasks** - Review any human-in-the-loop tasks generated

Monitor tab

The Monitor tab displays real-time information about your running automation:

- **UI streaming** - Watch browser automation actions execute in real-time
- **Variable watcher** - Monitor and modify variable values

- Logs - Track detailed automation progress

Test statuses

During testing, your automation will be in one of these statuses:

- Running - The automation is actively executing actions
- Paused (breakpoint) - Execution has stopped at a user-defined breakpoint
- Paused (exception) - Execution has stopped due to an error condition
- Completed - The automation has finished all steps successfully
- Failed - The automation encountered an unrecoverable error and stopped

Running and debugging

Setting breakpoints

Breakpoints allow you to pause your automation at specific points:

To add a breakpoint:

- Hover over an action or step on the canvas
- Click **Add breakpoint**
- The breakpoint icon appears on the action or step

To remove a breakpoint:

- Click **Remove breakpoint** on the action
- Or use the breakpoint manager in the Setup tab

Note

Set up breakpoints before starting debug mode.

Debug controls

When your automation is paused, you have several control options:

- Next - Run the next action in your automation and pause again
- Continue - Resume the run until the next breakpoint is reached
- Stop - End the current test run
- Ignore - If paused on an exception, this will clear the exception and continue as if that action was successful. This is useful if the exception is not blocking the rest of your test and you want to continue.
- Retry - If paused on an exception, this will attempt to run the current action again. This is useful if the issue is temporary or if you want to change the value of a variable and retry with the new value

Deactivating steps

To temporarily disable any actions or steps without removing them, you can deactivate them before starting your test. Deactivated actions or steps will be skipped during testing.

- Select the step to deactivate
- Click the **Menu options**
- Choose **Deactivate**
- The step appears grayed out

To reactivate:

- Select the deactivated step
- Click the **Menu options**
- Choose **Activate**

Variable watcher

The Variable watcher in the Monitor tab provides visibility to inspect data being used by your automation during the test. To inspect a variable:

- Open the Variable watcher in the Monitor tab while the test is paused
- Click **Add** to search and select variables to inspect
- View the current value for the variable
- Modify values to test different scenarios:

- Click the **pencil icon** next to any variable
- Enter a new value to test
- Continue the run to proceed with the new value

This feature is particularly useful for:

- Testing different data conditions
- Validating error handling
- Simulating edge cases

Troubleshooting

Fix with Assistant

When your automation encounters an exception, the Automation Assistant can help identify and resolve issues:

- Click **Fix with Assistant** when paused at an error
- The Assistant will:
 - Analyze the error message
 - Review the run logs
 - Check its troubleshooting knowledge base
 - Provide a clear explanation of the issue
 - Suggest potential solutions
- Chat with the Assistant to:
 - Ask questions about the error
 - Update the automation based on any of the recommended changes

Test metrics

After your test completes, Amazon Quick Automate provides comprehensive test metrics to help understand automation performance and validate test coverage. View key metrics about your test run:

- Total running time - How long the test took from start to finish

- Number of completed cases - Count of cases that finished processing (excluding those created or pending)
- Average time per case - The processing average time for each completed case
- Case success rate - Percentage of cases that completed without exceptions
- Count of tasks created - Number of human-in-the-loop tasks generated during the test
- Action coverage - Percentage of actions in your automation that executed during the test, helping identify untested paths

Deploying automations

After you create and test your automations in Amazon Quick Automate, the next step is to **deploy** them so they can run regularly on a trigger. Deployment makes the automation operational and ready for execution. You can add a trigger to a deployed automation to run it at a predefined schedule. Deployment involves configuring runtime settings, assigning users for human-in-the-loop tasks, verifying credentials and integrations.

This section explains each step of the deployment process and how to configure your automation for reliability, security, and optimal performance.

Prerequisites

Before deploying your automation, ensure that the following steps are complete:

- **Automation is tested thoroughly** - Validate your automation through end-to-end testing to confirm that all logic, actions, and agent interactions work as expected.
- **Version committed for deployment** - Only committed automation versions can be deployed. Review your changes, finalize the version, and commit it before proceeding.
- **Integrations configured** - If your automation interacts with external applications such as Salesforce or Jira via APIs, make sure all necessary integrations are configured.
 - Navigate to **Connections** → **Integrations** in the left panel to create new integrations.
 - Currently, only integrations available under the **Actions** tab are supported in Amazon Quick Automate.
 - Once an integration action is created, associate it with the **Automation Group** where it will be used.
 - The associated actions will then appear in the canvas. During deployment, you can select the appropriate connection to be used by the deployed automation.

- **Credentials configured** - Verify that all credentials required by your automation are correctly set up.

Deploying an automation

You can deploy an automation directly from the Canvas by clicking Deploy, or by navigating to the Deployment tab on the automation's landing page. Once you initiate deployment, the system guides you through a series of steps to complete the configuration and release process.

Release details

On the **Release Details** page, select the version of the automation you want to deploy. Only **committed versions** are available for deployment and will appear in the dropdown list.

Additional settings

Additional settings include:

- **Runtime configuration**
- **Tasks (for HITL assignments)**
- **Access**

Runtime configuration

Runtime configurations are parameters that may differ between environments such as development, testing, and production.

For example, an automation step that sends an email may use your personal email address during testing but should switch to a shared team address in production. Such environment-dependent values can be defined as **runtime configurations** when authoring the automation.

At deployment time, you can review and override these configurations to ensure the automation runs correctly in the intended environment. Runtime parameters can include:

- Email addresses or notification recipients
- File paths or URLs specific to the environment

This flexibility helps maintain a single automation definition across environments while adapting key parameters as needed.

Tasks

Select the resolver users or user groups for each of the tasks that require human-in-the-loop.

Access (Credentials and Connections)

Automations often need to connect to external systems, databases, or services. Amazon Quick Automate provides secure methods for managing credentials and connections without embedding sensitive information within your automation logic.

Connection and credential data are securely stored and encrypted, and made available to workflows at runtime without exposing them to authors or end users. This design ensures strong separation of secrets from automation definitions, improving both security and maintainability.

You can store and use two primary types of credentials:

- **Website credentials** - Used for UI automation steps that require website logins (username and password).
- **Action credentials (Integrations)** - Used for connecting Amazon services (e.g. S3) or external systems through configured integrations (e.g., Salesforce, Jira).

At deployment, ensure that the correct credentials and connections are selected so the automation can access all required systems securely and successfully.

Setting up triggers

Triggers determine when and how your automations run. You can configure automations to start based on a predefined schedule. To setup a trigger:

- On the Deployment page, click "Create Trigger" and configure the rules.
- Select the frequency
- Select the start date and time (Note that the actual execution will initiate within 15 minutes of the start time selected)
- Select the end date and time
- Select the timezone
- Amazon Quick Automate provided built-in scalability. Select the number of parallel executions of the automation (you can select a maximum of 10 parallel executions per trigger and 50 across all automations within an account. Please reach out to Amazon)

- For complex scheduling needs, you can use cron expressions to define precise running patterns. For example, to run an automation at 2:30 AM every Monday, Wednesday, and Friday, you would use the cron expression: `30 2 * * 1,3,5`.

Monitoring performance

Amazon Quick Automate provides comprehensive monitoring capabilities that help you track performance, audit and troubleshoot your automations. This section describes how to use the performance monitoring features in Amazon Quick Automate.

Prerequisites

Before using the Performance dashboard, you need:

- Active automations - At least one automation must be run to see metrics
- Automation group access - Permission to view the automation groups you want to monitor

Automation summary dashboard

The Automations tab provides high-level metrics and key performance indicators across your automation portfolio.

Automation summary charts

The dashboard includes three main performance visualizations:

- **Automation Distribution** - Shows the number of deployed automations organized by automation group, helping you understand your automation footprint across the organization
- **Success Rate** - Displays the percentage of cases that completed successfully versus those that encountered exceptions, providing insight into automation reliability
- **Time Savings** - Calculates total hours saved based on successful case completions multiplied by the estimated time savings configured for each automation project

Automation details table

Each automation is listed with key metrics and status information:

- **Automation name** - Name and version of the deployed automation

- **Latest run** - Most recent execution status and timestamp
- **Case metrics** - Count of cases processed, broken down by:
 - Successfully completed cases
 - Cases with business exceptions
 - Cases with system exceptions
- **Tasks** - Number of human-in-the-loop tasks created by the automation

Runs dashboard

The Runs tab provides detailed visibility into individual automation executions and their outcomes.

Run statuses

Automations can be in one of these run statuses:

- **Waiting** - The automation is queued and awaiting available system resources to begin execution. This is typically a brief transitional state.
- **In Progress** - The automation is actively executing its configured steps. You can monitor real-time progress through the logs.
- **Completed** - The automation has finished all steps successfully.
- **Failed** - The automation encountered an error it could not recover from and stopped execution. Detailed error information is available in the logs.
- **Stopped** - A user manually interrupted the automation execution using either the Exit or Terminate action.

Note

Count of incomplete runs includes Waiting and In Progress.

Run details table

Each run provides comprehensive execution information through the following columns:

- **Status** - Current run status (as defined in Run Status Types)

- **Automation** - Name of the automation being run
- **Version** - Specific version number of the automation which ran
- **Start time** - When the automation run began
- **End time** - When the automation completed or stopped
- **Duration** - Total running time of the automation
- **Case metrics** - Count of cases processed, broken down by:
 - Successfully completed cases
 - Cases with business exceptions
 - Cases with system exceptions
- **Tasks** - Number of human-in-the-loop tasks created in the run

Note

Data is refreshed when first navigating to the page. Click the refresh button to pull the latest data at any point.

Cases dashboard

The Cases tab enables tracking of individual case records through their lifecycle.

Case charts

The dashboard includes two main performance visualizations:

- **Total cases created** - Shows the number of cases created over time, helping you understand your automation volumes.
- **Top exceptions** - Displays the top 5 exception reasons and the relative count of each, helping you prioritize your optimization efforts to improve success rates.

Case statuses

Cases can be in one of these processing statuses:

- **Ready** - The case has been created and is waiting to be picked up for processing. This is the initial state of all new cases.

- **In Progress** - The case is actively being processed by the automation. You can monitor real-time progress through the logs.
- **Pending Resolution** - The case is waiting for a human-in-the-loop task to be completed before processing can continue. Cases move back to Ready status after tasks are resolved.
- **Successful** - The case has completed without any exceptions.
- **Failed** - The case encountered one of the following exception types:
 - **Business Exception** - The case encountered a handled business rule violation and stopped processing. Detailed exception information is available in the logs.
 - **System Exception** - The case experienced a technical error and stopped processing. Detailed error information is available in the logs.

For detailed information about case handling and orchestration patterns, see [Orchestration actions](#).

Search and filtering

The Performance dashboard includes search and filtering capabilities available across all monitoring views to help you find specific information. Filter by:

- Time range
- Status
- Automation group
- Automation

Note

You can also search on the cases page by reference name or exception reason.

Environment selection

Switch between viewing metrics from:

- Test - Data from automation testing in development environment
- Deployed - Data from deployed automations

 Note

Your selected environment applies across all dashboard tabs until changed.

Human-in-the-loop task center

Human-in-the-loop (HITL) tasks enable you to incorporate human judgment at critical points in your automated processes. The Task Center provides a centralized interface for managing HITL tasks. Users can review, resolve, and track tasks that require human intervention.

Overview

The Task center displays tasks requiring human attention and provides metrics on task processing.

Task metrics

The dashboard shows key performance indicators:

- **Open tasks** - Total number of tasks awaiting completion
- **Critical tasks** - Count of high-priority tasks requiring attention
- **Overdue tasks** - Number of tasks past their due date

Task details table

Each task provides comprehensive information through the following columns:

- **Title** - Brief description of what needs to be done
- **Status** - Current task state (Open, Completed, Archived)
- **Due date** - When the task should be completed (includes overdue warning)
- **Severity** - Task priority (Low, Moderate, Critical)
- **Assignee** - User or group responsible for the task
- **Creation time** - When the task was created
- **Completion time** - When the task resolution was submitted by a user

Prerequisites

Before working with tasks, you need:

- Task permissions - Access granted through automation deployment settings
- Active automations - At least one automation with HITL tasks must be run

Types of tasks

Chat tasks

Interactive conversations where AI agents gather input through natural language. To participate in a chat task:

- Select the task
- Review the conversation history
- Respond to agent questions in the chat interface
- The agent will close the task once it has sufficient information

Note

Chat tasks are closed automatically by the AI agent. You can archive unresolvable tasks, which the agent will treat as exceptions.

Form tasks

Tasks requiring structured input through form fields. To complete a form task:

- Select the task
- Fill in the required form fields
- Click Submit when finished

Note

Enable "Automatically open next task" to streamline processing multiple tasks. Use "Save as Draft" to preserve partial progress.

Managing tasks

Task statuses

Tasks can be in one of these states:

- Open - Awaiting human input
- Completed - Successfully submitted
- Archived - Removed from active view

Task actions

Manage tasks through these available actions:

- Submit - Complete a form task
- Save as Draft - Preserve partial form progress
- Archive - Remove task from active view
- Unarchive - Restore archived task

Note

Only Open tasks can be archived.

Task assignment

Distribute work by assigning tasks:

- Select a task
- Click the assignee field
- Search for and select users or groups
- Click Remove assignee to clear assignment

Note

Users must have appropriate permissions to be assigned tasks.

Related information

Access connected resources through the Sources menu:

- Case - View related case details (if applicable)
- Automation - See the automation that created the task

Note

Source access requires appropriate permissions.

Search and filtering

Locate specific tasks using:

- Search - Find tasks by title
- Filters:
 - Status
 - Due date
 - Severity
 - Assignee
 - Creation date

Note

Click refresh to update task data at any time.

Environment selection

Switch between viewing tasks from:

- Test - Tasks from automation testing
- Deployed - Tasks from deployed automations

Work with integrations in Amazon Quick Suite

Integrations in Amazon Quick Suite connect you to external applications and services for AI-powered analysis and automation. With integrations, you can perform actions in external applications and bring data from external applications to enhance your AI agents and chat experiences.

You can create three types of integrations:

Action connectors

Initiate real-time calls to perform actions in external services. Use action connectors to address user natural language chat requests, automate workflows, send notifications, and/or trigger processes in connected applications.

Data access integrations

Connect Quick Suite with external applications to access data. Data access integrations serve as the foundation for creating knowledge bases. You cannot use these integrations directly for analysis.

Knowledge bases

Connect to data access integrations to make specific content accessible for AI analysis. Choose which data or knowledge to include from external applications. Add knowledge bases to spaces or use them directly with chat agents.

Use the Integrations tab in the Amazon Quick Suite console to create these connections. You can connect to Model Context Protocol (MCP) servers, Google Drive, Amazon S3, web crawlers, OneDrive, and other supported applications.

Note

Not all applications support all integration types. Available options depend on each application's capabilities.

Topics

- [Supported integrations](#)

- [Set up integrations in the console](#)
- [Integration workflows](#)
- [Knowledge bases](#)
- [Action connectors](#)
- [Integration-specific guides](#)
- [Creating Amazon service action connectors using the admin console](#)
- [Bring Your Own Amazon Q Business Index \(BYOI\)](#)
- [Cross-region inference in Amazon Quick Suite integrations](#)

Supported integrations

Amazon Quick Suite supports integrations with various third-party applications and services. Each integration supports different combinations of actions and knowledge base creation capabilities. The following table shows the supported integrations and their capabilities.

Supported Integration Capabilities

Integration	Actions	Knowledge Base
Amazon S3	✓	✓
Asana	✓	—
Atlassian Confluence Cloud	✓	✓
BambooHR	✓	—
Google Drive	—	✓
Atlassian Jira Cloud	✓	—
Microsoft Outlook	✓	—
Microsoft OneDrive	✓	✓
Microsoft SharePoint Cloud	✓	✓
Microsoft Teams	✓	—

Integration	Actions	Knowledge Base
Model Context Protocol (MCP)	✓	—
OpenAPI Specification	✓	—
PagerDuty	✓	—
REST API	✓	—
Salesforce	✓	—
SAP Bill of Materials	✓	—
SAP Business Partner	✓	—
SAP Material Stock	✓	—
SAP Physical Inventory Docs	✓	—
SAP Product Master	✓	—
ServiceNow	✓	—
Slack	✓	—
Smartsheet	✓	—
Web Crawler	—	✓
Zendesk Suite	✓	—

 Note

Not all applications support all integration types. The available options depend on the capabilities of each specific application and your user role.

Integration capability definitions

Actions

Call APIs and perform actions in external applications directly from Amazon Quick Suite. You can share action connectors with other users and use them in automated workflows.

Knowledge base

Create searchable repositories of information from external sources. Knowledge bases are children of data access integrations. Add them to spaces or use them directly in chat agents.

Authentication method definitions

User auth

Custom user-based OAuth authentication requiring base URL, client ID, client secret, token URL, auth URL, and redirect URL.

Service auth

Service-to-service authentication using either API key (with base URL and email) or service-to-service OAuth (with base URL, client ID, client secret, and token URL).

Managed OAuth 2.0

Managed OAuth 2.0 authentication flow with provider-specific sign-in interface.

Amazon credentials

Amazon-specific authentication using Amazon access keys and permissions.

Basic auth

Username and password authentication.

Form/SAML auth

Form-based or SAML authentication with configurable field selectors.

JSON schema

Schema-based authentication for OpenAPI specifications.

Set up integrations in the console

The Integrations console provides a streamlined interface for creating and managing integrations. When you access the Integrations tab, you see all available applications in a unified grid. The console setup process adapts based on the integration you select, your user tier, and existing integrations.

Choose integration options

When you select an integration from the main grid, the console branches based on several factors:

- **Integration capabilities** - Each application supports different combinations of actions and knowledge base creation. For example, Google Drive supports data ingestion not actions. OneDrive supports both data ingestion and actions.
- **User tier** - Professional subscriptions enable you to create knowledge bases out of Google Drive, OneDrive, Confluence, and SharePoint. In order to set up actions and bring data from S3 and Web Crawler, you need to have the Enterprise subscriptions.
- **Existing integrations** - The console first shows existing integrations before offering to create new ones.

View setup process examples

The following examples show how different integrations guide you through different console setup processes:

Google Drive - Bring data for Q&A

Google Drive supports data access integrations and knowledge base creation.

1. In the console, choose **Integrations**.
2. From the integration grid, choose **Google Drive**, and select the add icon.
3. The OAuth popup appears. Complete the OAuth process and grant it all the necessary permissions. Choose **Continue**.
4. You are taken to the Create knowledge base page. Select the files you want to assign to your knowledge base, and choose **Select**.
5. Enter a name and description. Choose **Create**.

OneDrive - Perform actions and bring data for Q&A

OneDrive supports the capabilities to both perform actions and bring data for Q&A. The console provides options for both capabilities during setup.

1. In the console, choose **Integrations**.
2. From the integration grid, choose **OneDrive**.
3. In the Create Integration dialog, select your integration type:
 - **Perform actions in OneDrive** - Enables actions like creating, updating, or managing OneDrive files.
 - **Bring data from OneDrive** - Enables you to ask questions about and get insights from your OneDrive content.
 - **Both** - Combines task and data ingestion capabilities.
4. Enter integration details and configure Microsoft 365 authentication.
5. If you selected "Bring data from OneDrive", select the files you want to add to your knowledge base using the file picker interface.
6. Choose **Create** to complete the setup.

Working with existing integrations

When you select an integration type that you've used before, the console first displays your existing integrations.

1. In the console, choose **Integrations**.
2. Choose an integration type you've used before (for example, **OneDrive**).
3. The console displays a table of existing integrations with columns for:
 - **Name** - Integration name
 - **Status** - The status of the integration (ex. Available, Completed with issues)
 - **Visibiliy** - The visibility level of the integration (ex. Personal, Shared)
 - **Owner** - The integration creator
 - **Last Modified** - Most recent update
 - **Actions** - A dropdown menu to perform actions on existing integrations

4. Choose an existing integration to view details, or choose add + to set up an additional connection.

Integration management options

After creating integrations, you can manage them through several console options:

- **Edit integration** - Modify integration settings, authentication details, and configuration options
- **Delete integration** - Remove integrations with confirmation dialogs to prevent accidental deletion
- **Knowledge base management** - Separate flows for creating, editing, and deleting knowledge bases associated with your integrations

Integration workflows

The following procedures describe the general workflows for creating and managing different types of integrations in Amazon Quick Suite.

Creating a knowledge base from scratch

Data access integrations establish the connection to external systems creating knowledge bases from external data sources.

1. In the Amazon Quick Suite console, choose **Integrations**.
2. From the integration grid, choose the application you want to connect to (for example, **Google Drive**, **OneDrive**, or **S3**).
3. In the Integration details section, select the "Add" option (+). If required, complete the authentication process in the popup that occurs.
4. Fill in the appropriate details, depending on your chosen integration. For example, for Amazon S3, select your Amazon account and your Amazon S3 bucket url.
5. Enter a **Name** for your integration.
6. Enter the required connection details for your chosen application.
7. If required, choose **Create and continue** to continue to knowledge base creation.
8. Specify a name for your knowledge base.

9. Specify the files you want to include in your knowledge base using the file picker or appropriate sync options (for example, **S3** allows you to choose to add all content of specific content).

10. Choose **Create**.

Syncing of your content will automatically begin after creation of the knowledge base.

Creating an action connector

Action Connectors enable you to perform actions in external applications directly from Amazon Quick Suite.

1. In the Amazon Quick Suite console, choose **Integrations**.
2. From the integration grid, choose an application that supports action connectors (for example, **OneDrive**, **Confluence**, or **Slack**).
3. In the Integration details section, select the "Add" option (+).
4. Enter a **Name** for your action connector.
5. Configure the task-specific settings for your chosen application.
6. Choose **Next** to complete the authentication and setup process.

After successful creation, your action connector is available for use in Amazon Quick Suite workflows and can be triggered from analyses, dashboards, or automated processes.

Creating a knowledge base integration from an existing data access integration

Knowledge base integrations allow you to create searchable repositories of information from external sources.

Follow the procedure below to quickly create a knowledge base for one of the six supported data sources.

1. Open the Amazon Quick Suite console and choose **Integrations**. Then click on the **Data** tab.
2. Select your data source from the available options in the **Data** tab,
3. Selecting the three-dot menu option under **Actions**, and then select **Create knowledge base**.
4. Add a name and enter an optional description.
5. Select the content to include in your knowledge base. Choose **Create**.

Data synchronization begins automatically. Sync jobs can take a few mins to hours depending upon the size of your data, but you can see what documents are available for querying from the Refresh Reports tab. The knowledge base status changes from *Syncing* to *Available* when the sync is complete.

Managing existing integrations

You can edit, delete, and manage existing integrations from the Integrations console.

1. In the Amazon Quick Suite console, choose **Integrations**.
2. From the integrations table, select the integration you want to modify.
3. Select the three-dot drop-down menu under "Actions" and then select "Edit integration".
4. Modify the integration settings as needed.
5. Choose **Save changes** to apply your modifications.

Your integration settings are updated and the changes take effect immediately for new operations using the integration.

To delete an integration

1. In the Amazon Quick Suite console, choose **Integrations**.
2. From the integrations table, select the integration you want to remove.
3. Select the three-dot drop-down menu under "Actions" and then select "Delete integration".
4. In the confirmation dialog, review the integration details and any dependent resources that will be affected.
5. Choose **Delete** to confirm the removal.

The integration is permanently removed from your account. Any dependent resources (such as knowledge bases) that rely on this integration will be impacted.

Knowledge bases

A knowledge base is an organized, indexed collection of documents or content from data sources optimized for generative AI-powered retrieval and question answering. Whether your team stores documentation in Confluence, collaborates through SharePoint, or manages files in cloud storage, you can bring all this information into one unified search experience by creating knowledge bases.

The built-in integrations can be set up with just a few clicks to sync your data in Quick Suite and make it effortless to tap into your organization's knowledge across Google Drive, OneDrive, Confluence, SharePoint, S3, and Web Crawler. Whether your team stores documentation in Confluence, collaborates through SharePoint, or manages files in cloud storage, you can bring all this information into one unified search experience with the help of knowledge bases.

How knowledge bases work

Knowledge base is an indexed collection of documents or content from data sources such as Google Drive, optimized for generative AI-powered retrieval and question answering. Multiple knowledge bases can be created from the same source, and all can reside within a shared Quick Index. For example, if you sync two folders from Google Drive and create two knowledge bases — one for “Policy Documents” to answer queries such as *“What’s our refund policy”* and one for “Customer feedback” to answer queries such as *“What are the common customer complaints”* — both can be part of the same index. Quick distinguishes between them using the knowledge base id, so queries can be filtered to retrieve only the relevant documents from the desired knowledge base. This allows users to organize, secure, and retrieve information relevant to different domains or use cases, even though the underlying data is indexed together.

Your knowledge bases can be used individually or shared with team members through Amazon Quick Suite spaces. Our coarse-grained access control enables security at the knowledge base level, ensuring that users only receive information from knowledge bases they're authorized to access.

Creation process

You can create knowledge bases while setting up a new data access integration and use existing integrations to create additional knowledge bases:

1. **Data access integration setup** - Connect to your external data source
2. **Content selection** - Choose which content to include through filters and scope settings
3. **Indexing** - Amazon Quick Suite processes and indexes the selected content
4. **Availability** - The knowledge base becomes available for use in spaces and by AI agents

Capabilities

Each knowledge base provides the following capabilities:

- **Content indexing** - Processes text, documents, and structured data from external sources

- **Semantic search** - Enables AI-powered search across indexed content
- **Automatic synchronization** - Keeps content up-to-date with configurable sync schedules
- **Coarse-grained access control** - Ensures that users only receive information from knowledge bases they're authorized to access.
- **Multi-space usage** - Can be used across multiple spaces and by different AI agents

General workflow

The typical workflow for working with knowledge bases follows these steps:

1. **Set up data source integration** - Connect to your external application (such as SharePoint, Google Drive, or Confluence) with appropriate authentication. For more information, see [the section called "Integration-specific guides"](#).
2. **Create a knowledge base** - You can create a knowledge base while configuring your new integration. Configure your content filters by setting up include filters, file type restrictions, and folder selections to focus on relevant content.
3. **Set sync schedule** - Data refresh frequency is set to daily by default. You can edit the sync frequency to configure how often the knowledge base should be updated with new content from the source.
4. **Monitor and manage** - Review sync status, manage access permissions.

Common configuration settings

Knowledge bases share common configuration patterns across different data source integrations. Understanding these settings helps you optimize content indexing and manage sync behavior effectively.

Note

While these configuration options are available across most integrations, specific settings and available options may vary depending on your chosen data source integration.

File size and content limits

Configure file size limits to optimize processing performance and manage storage costs. The specific limits vary by content type and are displayed in the console when you configure your knowledge base.

Standard text documents

Applies to documents like PDFs, Word files, and text files. File size limit is 50 MB.

Video files

Available when video processing is enabled. Supported formats include .mp4, .mov, .m4v. Limit is 10 GB (10240 MB) for video files.

Audio files

Available when audio processing is enabled. Supported formats include .mp3, .wav, .m4a, .flac, and .ogg. Limit is 2 GB (2048 MB) for audio files.

Note

Files with extracted text that exceeds the 30 MB system limit are not indexed, regardless of the original file size. The maximum amount of text that can be extracted from a single document is 30 MB.

Sync schedule and safeguards

Configure how often your knowledge base updates and protect against unintended content deletion:

Sync frequency

Data refresh frequency is set to daily by default. You can edit the sync frequency to configure how often the knowledge base should update with new content from the source

Document deletion safeguard

Protect your indexed content from accidental mass deletion by setting a maximum deletion percentage threshold. If a sync job would delete more documents than your threshold allows, the deletion phase is skipped, preserving your existing indexed content.

This safeguard protects against temporary network issues, permission changes, or source system problems that might make content temporarily unavailable.

Troubleshooting knowledge bases

When you encounter issues with your Quick Suite knowledge base, you can use this troubleshooting guide to identify and resolve common problems. Knowledge base issues typically involve document synchronization, refresh job failures, or access permissions.

Documents don't appear in your knowledge base

When documents you expect to see don't appear in your knowledge base, several factors might cause this issue.

Common causes:

- **Sync in progress** – Documents might still be processing. Check the refresh status to confirm the refresh is complete.
- **Unsupported file format** – Verify your documents are in supported formats: Word, Excel, PowerPoint, PDF, CSV, TXT, RTF, JSON, XML, HTML
- **File size too large** – Each file must be less than 50 MB.
- **Insufficient access permissions** – Confirm the knowledge base has proper permissions to access the document source.
- **Document filtering** – Check if filters or exclusion rules prevent certain documents from being indexed.

To troubleshoot:

1. Review the refresh history for error messages related to specific documents that failed to sync.
2. Verify your document formats and file sizes meet requirements.
3. Check your access permissions and connection settings.

Refresh job fails

A refresh job typically fails when there's a configuration error in the knowledge base or data source connection.

Common causes:

- **Permission issues** – The integration lacks sufficient permissions to access the data source.
- **Configuration errors** – Incorrect URLs or data source connection settings.
- **Resource limitations** – Rate limiting from the source system.

To resolve:

1. Check the refresh history details for specific error messages.
2. Verify all connection settings and permissions are correctly configured.
3. Take the recommended action based on the error message.

Refresh job completes with issues

When a refresh job completes with issues, the job processed successfully but encountered problems with some documents.

What this means:

- **Partial success** – Some documents synced successfully while others failed.
- **Document-level errors** – Individual files might have formatting issues, corruption, or access problems.
- **Metadata issues** – Problems with document metadata or associated information.
- **Size or format violations** – Some files might exceed size limits or be in unsupported formats.

To resolve:

1. Review the detailed refresh reports to identify which documents encountered issues.
2. Address the individual document problems.
3. Run another refresh after resolving the issues.

Refresh job succeeds but no documents appear

If a refresh job shows as successful but no documents appear in your knowledge base, check these potential causes.

Common causes:

- **Empty source** – The configured data source location contains no documents.
- **Incorrect path configuration** – The source path or connection settings don't point to the correct location.
- **Document filters** – Inclusion or exclusion criteria might filter out all documents.
- **Read permissions missing** – The job connected successfully but lacked permissions to read the actual documents.

To resolve:

1. Verify your data source configuration points to the correct location.
2. Confirm documents are present in the specified location.
3. Check that appropriate access permissions are configured.
4. Review any document filters that might exclude content.

File format issues during refresh

Quick Suite knowledge bases support specific file formats. Files must meet format, size, and character limit requirements.

Requirements:

- **Supported formats:** Word, Excel, PowerPoint, PDF, CSV, TXT, RTF, JSON, XML, HTML
- **File size limit:** 50 MB per file
- **File condition:** Not corrupted or password-protected

To resolve format issues:

1. Verify your files meet the format and size requirements.
2. Convert unsupported formats to supported ones.
3. Remove password protection from files.
4. Check that files aren't corrupted.

Access denied errors

Access denied errors typically occur due to authentication or authorization issues.

Common causes:

- **Invalid credentials** – Authentication tokens or passwords might have expired.
- **Insufficient permissions** – The account used in the integration lacks read access to the data source.
- **Network restrictions** – Firewall or security policies block access.
- **SSL/TLS issues** – Certificate problems with secure connections.

To resolve:

1. **Verify authentication credentials** – Confirm that authentication credentials are current and valid. Edit the integration to re-authenticate and generate a new token.
2. **For web crawler data sources** – Verify that secure connections are properly configured and SSL certificates are properly configured and trusted.
3. **Contact your system administrator** – If you continue experiencing access issues, contact your system administrator. They might need to adjust permissions or security settings.

Action connectors

Action connectors use secure connections to external services and execute actions based on your authentication level and permissions.

How action connectors work

Action connectors in Amazon Quick Suite create secure connections between Amazon Quick Suite and external services. When you configure these integrations, you can perform actions based on your authentication level and permissions.

Core components

Action connectors

The foundational resources that integrate with external services. Amazon Quick Suite supports 15 third-party integrations and 5 Amazon service integrations. For information about setting up Amazon built-in service action connectors, see [Amazon service action connectors](#).

Authentication methods

Action connectors support multiple authentication methods including managed (3LO), custom user-based, API key, and 2LO. For detailed information about each authentication method, see [the section called “Authentication methods”](#).

Implementation types

- **On-demand actions for immediate, user-triggered operations** - Real-time operations that execute immediately when you trigger them. You can initiate actions through chat interfaces, dashboards, or Amazon Q Apps. Examples include creating tickets, sending messages, or querying data.
- **Automated workflows for scheduled or system-triggered tasks** - System-managed operations that execute based on schedules or triggers. They run in the background without user intervention. Examples include data synchronization, report generation, or system maintenance.

Permission models

- **Personal access permissions through 3LO** - You can grant specific permissions to Amazon Quick Suite through Three-Legged OAuth, maintaining control over your service access. Permissions are tied to your identity and credentials in the target service.
- **Service-level permissions for automated workflows** - Applied to automated workflows, these permissions support system-to-system interactions without user involvement. They're configured at the service level and typically use API keys or service account credentials.
- **Entity-level access controls** - Govern access to actions within Amazon Quick Suite, determining which users or groups can execute specific actions. These controls integrate with Amazon Quick Suite's broader permission system for consistent access management across the platform.

Types of actions

Amazon Quick Suite supports two methods of invoking actions, each serving different use cases and authentication models.

On-demand actions

On-demand actions execute immediately when you trigger them. These actions support interactive operations that require real-time response.

Key characteristics:

- User-initiated execution - You trigger actions through natural language in the chat interface.
- Interactive form completion - You fill out forms with required parameters before the action executes.
- Immediate response - Actions execute in real-time and provide instant feedback on success or failure.
- Personal authentication (3LO) - Uses your individual credentials and permissions from the target service.

Common use cases:

- Creating tickets in Jira.
- Sending messages in Slack.
- Updating Salesforce records.
- Retrieving information from SharePoint.

Automated workflows

Automated workflows execute actions on a schedule or in response to specific triggers. These are useful for background and system-level operations.

Key characteristics:

- System-level execution - Actions run automatically without user intervention based on predefined triggers.
- Scheduled or event-triggered - Execute on time-based schedules or in response to specific system events.
- Non-interactive operation - Run in the background without requiring user input or form completion.
- Service-level authentication - Use system credentials rather than individual user authentication.

Common use cases:

- Regular data synchronization.
- Scheduled report generation.

- Automated ticket updates.
- System health checks.

Bounded and unbounded agents

Amazon Quick Suite offers two types of agents that provide different levels of access and functionality: bounded and unbounded agents. Understanding the differences between these agent types helps you implement the right solution for your use case.

Bounded agents

Bounded agents operate within defined parameters, specifically linked to one or more spaces within Amazon Quick Suite. These agents can only access and perform actions on resources that are explicitly connected to their assigned spaces. For example, a bounded agent configured for the HR space can only access HR-related documents, datasets, and execute HR-related actions.

Use bounded agents for:

- Department-specific workflows (HR, Finance, IT).
- Project team collaborations.
- Sensitive data handling.
- Compliance-focused operations.

The bounded nature provides enhanced security by ensuring the agent can't access resources outside its designated spaces. This makes it ideal for scenarios where data isolation is important.

Unbounded agents

Unbounded agents have broader access capabilities and can work across all configured actions and resources within the Amazon Quick Suite environment. These agents aren't restricted to specific spaces and can access any properly configured action connector available in the system.

Use unbounded agents for:

- Organization-wide assistance.
- Cross-departmental workflows.
- General-purpose actions.

- Scenarios requiring access to multiple systems.

Prerequisites

Before using actions in Amazon Quick Suite, ensure you have the following:

License requirements

One of the following Amazon Quick Suite licenses:

- Reader Pro - Provides read access to data and the ability to execute actions in connected services.
- Author - Includes Reader Pro capabilities plus the ability to create and modify content and configurations.
- Author Pro - Full feature access including advanced action configuration and administrative capabilities.

Service requirements

For third-party services (such as Jira or Salesforce), ensure that you have:

- Appropriate permissions in the target services.
- Authentication credentials for each service.

For Amazon action connectors, you need admin access to the relevant services.

Amazon account requirements

- Active Amazon account - A valid Amazon account with billing enabled and in good standing.
- Appropriate IAM permissions - IAM roles and policies that allow Amazon Quick Suite to access the required Amazon services.
- Required service quotas - Sufficient service limits for the Amazon services you plan to integrate with your actions.

Supported action connector types and available actions

Amazon Quick Suite supports multiple connector types, each with specific actions available:

External service connectors

- **Salesforce** - Create records, update opportunities, search accounts, manage leads.
- **JIRA** - Create issues, update tickets, search projects, manage workflows.
- **Microsoft Outlook** - Send emails, manage calendar events, access contacts.
- **Slack** - Send messages, create channels, manage notifications.
- **ServiceNow** - Create incidents, update requests, manage workflows.
- **Zendesk** - Create tickets, update cases, search knowledge base.
- **PagerDuty** - Create incidents, manage escalations, update on-call schedules.
- **Asana** - Create actions, update projects, manage team workflows.
- **BambooHR** - Access employee data, manage time-off requests.
- **Smartsheet** - Update sheets, manage project data.
- **Factset** - Access financial data, generate reports.
- **Confluence** - Create, update, and manage pages, spaces, and other Confluence objects.
- **SharePoint** - Perform actions on SharePoint lists, items, and Excel files with 19 available actions for creating, updating, deleting, and retrieving SharePoint content.
- **OneDrive** - Create, update, delete, and manage OneDrive files and folders.
- **SAP** - Access SAP S/4HANA systems to perform Read only operation on enterprise data.

Amazon service connectors

- **Amazon S3** - Upload files, manage buckets, retrieve objects.
- **Amazon Bedrock** - Generate content, analyze data, process requests.
- **Amazon Textract** - Extract text and data from documents.
- **Amazon Comprehend** - Natural language processing and sentiment analysis.
- **Amazon Comprehend Medical** - Medical text analysis and entity extraction.

Action connector compatibility matrix

The following table shows which Amazon Quick Suite features each action connector type supports:

Action Connector Feature Compatibility

Action Connector	Chat Agents	Flows	Dashboard Visuals	Dashboard Alerts	Automations	Companions
Amazon Built-in Services						
Amazon Bedrock Agent Runtime	—	—	—	—	✓	—
Amazon Bedrock Data Automation Runtime	—	—	—	—	✓	—
Amazon Bedrock Runtime	—	—	—	—	✓	—
Amazon Comprehend	—	—	—	—	✓	—
Amazon Comprehend Medical	—	—	—	—	—	—
Amazon S3	—	—	—	—	✓	—
Amazon Textract	—	—	—	—	✓	—
External Service Connectors						
Asana	✓	✓	—	—	—	✓
Atlassian Confluence Cloud	✓	✓	—	—	✓	✓
Atlassian Jira Cloud	✓	✓	✓	✓	✓	✓
BambooHR	✓	✓	—	—	—	✓

Action Connector	Chat Agents	Flows	Dashboard Visuals	Dashboard Alerts	Automations	Companions
Microsoft OneDrive	✓	✓	—	—	✓	✓
Microsoft Outlook	✓	✓	✓	✓	✓	✓
Microsoft SharePoint	✓	✓	—	—	✓	✓
Microsoft Teams	✓	✓	✓	✓	✓	✓
PagerDuty	✓	✓	—	—	✓	✓
Salesforce	✓	✓	✓	✓	✓	✓
SAP	✓	—	—	—	✓	✓
ServiceNow	✓	✓	✓	✓	✓	✓
Slack	✓	✓	✓	—	—	✓
Smartsheet	✓	✓	—	—	—	✓
Zendesk	✓	✓	—	—	—	✓
Custom Connector Types						
Model Context Protocol (MCP)	✓	✓	—	—	✓	—
OpenAPI	✓	✓	—	—	—	—
REST API	—	—	—	—	✓	—

Authentication Support:

- **Chat Agents and Companions** - Support user authentication (3LO, Basic)

- **Dashboard Visuals** - Support user authentication (3LO)
- **Dashboard Alerts** - Support system authentication (2LO or API Key)
- **Automations** - Support system authentication (2LO)

Action connector APIs

Action connector APIs let you programmatically create and manage connections between Amazon Quick Suite and external services. These APIs support the action integration functionality that allows users to perform actions in third-party applications directly from Amazon Quick Suite chat interfaces and automated workflows.

What are action connector APIs?

Action connectors serve as the foundational resources that enable integration with first and third party applications. Through these APIs, you can authenticate to applications, manage permissions, and control which actions are available to users within your Amazon Quick Suite applications.

How action connector APIs support action integrations

Action connector APIs provide the backend infrastructure for Amazon Quick Suite action integrations. When you create an action connector through the API, you establish a secure connection that lets you:

- Execute actions in external services through chat interfaces.
- Perform automated workflows in background processes.
- Integrate third-party services with Amazon Quick Suite applications.
- Manage authentication and permissions for service access.

The APIs handle the complex authentication flows, credential management, and permission controls needed to securely connect Amazon Quick Suite with external services.

Authentication methods

Action connector APIs support multiple authentication methods to accommodate different use cases and security requirements:

Managed authentication (3LO)

Three-Legged OAuth provides the simplest setup for personal access to third-party services:

- No initial configuration required.
- User-specific authentication through service provider login.
- Automatic token refresh with 90-day lifecycle.
- Secure credential storage managed by Amazon Quick Suite.

Service-to-service authentication (2LO)

For complex enterprise integrations:

- Supports client credentials OAuth flow.
- Enables system-to-system interactions.
- Requires client ID, client secret, and token URL configuration.
- Suitable for automated workflows requiring sophisticated security.
- OAuth - Dynamic Client Registration (DCR - applicable only for select MCP servers).

API key authentication

Simplified authentication for automated workflows:

- Single token-based authentication.
- Service-level permissions.
- Ideal for background processes and scheduled actions.
- Requires valid API key from target service.

Basic Auth

Basic authentication provides a simple username/password authentication method:

- Uses standard HTTP Basic Authentication headers.
- Credentials are base64 encoded.
- Suitable for services that don't support OAuth or API keys.
- Requires secure HTTPS connection.
- Not recommended for public-facing services.

None

No authentication required:

- Used for public APIs and services.
- No credentials or tokens needed.
- Limited to read-only or public operations.
- Typically used for public data feeds and documentation.
- Should not be used for sensitive operations.

Permissions and access control

Action connector APIs implement comprehensive permission controls through Access Control Lists (ACLs):

Resource-level permissions

- **Owner** - Full control including delete and permission management.
- **Contributor** - Can use and modify connector settings.
- **Viewer** - Can view connector details and use enabled actions.

API operations for permission management

- `DescribeActionConnectorPermissions` - Retrieve current permission settings.
- `UpdateActionConnectorPermissions` - Grant or revoke user permissions.

Supported connector categories

Dual-purpose connectors

These connectors support both action integrations and knowledge base creation:

- **Amazon S3** - Use the Admin Console to create Actions for file operations, use the webapp to create knowledge bases from S3 content.
- **Microsoft SharePoint** - Document management actions, content indexing.
- **OneDrive** - File operations, document search capabilities.

- **Confluence** - Content creation actions, knowledge base integration.

Action-only connectors

Specialized for action execution without knowledge base capabilities:

- **Salesforce** - Enterprise CRM integration supporting account and contact operations, custom object CRUD operations, Sales process automation.
- **JIRA** - Issue tracking and project management.
- **Microsoft Outlook** - Send emails, manage calendar events, access contacts.
- **Slack** - Communication and notification workflows.
- **ServiceNow** - IT service management operations.
- **Zendesk** - Create tickets, update cases, search knowledge base.
- **PagerDuty** - Create incidents, manage escalations, update on-call schedules.
- **Asana** - Create actions, update projects, manage team workflows.
- **BambooHR** - Access employee data, manage time-off requests.
- **Smartsheet** - Update sheets, manage project data.
- **FactSet** - Access financial data, generate reports.
- **SAP** - Access SAP systems, execute business functions, and manage enterprise data.

Knowledge base-only connectors

Focused on knowledge base integration without action capabilities:

- **Google Drive** - Document indexing and search.
- **Web Crawler** - Content discovery and indexing.

API lifecycle management

Credential management

- Automatic refresh token handling for OAuth action connectors.
- Secure storage of authentication credentials using Amazon KMS.
- Support for credential rotation and updates.

- Cross-account access for Amazon S3 connectors.

Connection updates

Use the `UpdateActionConnector` API to:

- Modify authentication credentials.
- Update service configuration parameters.
- Change action connector metadata.

Monitoring and troubleshooting

- Track API usage through CloudWatch metrics.
- Monitor connection health and authentication status.
- Implement error handling for common failure scenarios.
- Use validation APIs to diagnose configuration issues.

Rate limiting and quotas

Action connector APIs implement standard Amazon API rate limiting:

- Standard Amazon API throttling applies to all operations.
- Connection validation may have additional limits.
- Action execution rates depend on target service capabilities.
- Implement exponential backoff for retry logic.

Cross-account support

For Amazon S3 connectors, the APIs support cross-account access:

- Specify different Amazon account IDs during connector creation.
- Configure appropriate IAM permissions for cross-account access.
- Use Amazon KMS for secure credential management across accounts.
- Validate permissions before enabling cross-account connections.

Error handling and troubleshooting

Action connector APIs return standard Amazon error responses:

Common error types

- `AccessDeniedException` - Insufficient permissions for the operation.
- `InvalidParameterValueException` - One or more parameter values are invalid for the operation.
- Invalid configuration parameters - Service-specific configuration values are incorrect or missing.
- `ResourceNotFoundException` - Connector or resource not found.
- `ThrottlingException` - Rate limit exceeded.
- `ConflictException` - Resource conflict or duplicate names.
- `InternalFailureException` - Internal service error occurred during request processing.
- `ResourceExistsException` - Attempt to create a resource that already exists.
- `InvalidNextTokenException` - The pagination token provided is invalid or expired.
- `AccessTokenNotFoundException` - User needs to authorize the connection (that is, sign-button). This exception is used by UX to ask users for authorization.
- `TokenResponseException` - Action setup is not valid.

Implement proper error handling in your applications to manage these scenarios gracefully and provide meaningful feedback to users.

Using Action Connector APIs with Amazon CLI

You can use the Amazon CLI to manage action connectors programmatically. The following examples demonstrate common operations using generic placeholder values.

Creating an action connector

Use the `create-action-connector` command to create a new action connector for integrating with external services.

```
aws quicksight create-action-connector \
  --aws-account-id "123456789012" \
  --name "MyS3Connector" \
```

```
--action-connector-id "my-s3-connector-id" \  
--type "AMAZON_S3" \  
--authentication-config '{  
  "AuthenticationType": "IAM",  
  "AuthenticationMetadata": {  
    "IamConnectionMetadata": {  
      "RoleArn": "arn:aws:iam::123456789012:role/MyConnectorRole"  
    }  
  }  
}' \  
--enabled-actions "CreateBucket" "ListBuckets" \  
--description "S3 connector for automation workflows" \  
--region "us-east-1"
```

Listing action connectors

Use the `list-action-connectors` command to retrieve all action connectors in your account.

```
aws quicksight list-action-connectors \  
  --aws-account-id "123456789012" \  
  --max-results 10 \  
  --region "us-east-1"
```

Describing an action connector

Use the `describe-action-connector` command to get detailed information about a specific action connector.

```
aws quicksight describe-action-connector \  
  --aws-account-id "123456789012" \  
  --action-connector-id "my-s3-connector-id" \  
  --region "us-east-1"
```

Updating an action connector

Use the `update-action-connector` command to modify an existing action connector's configuration.

```
aws quicksight update-action-connector \  
  --aws-account-id "123456789012" \  
  --action-connector-id "my-s3-connector-id" \  
  --name "UpdatedS3Connector" \  
  --authentication-config '{  
    "AuthenticationType": "IAM",  
    "AuthenticationMetadata": {  
      "IamConnectionMetadata": {  
        "RoleArn": "arn:aws:iam::123456789012:role/UpdatedConnectorRole"  
      }  
    }  
  }' \  
  --enabled-actions "CreateBucket" "ListBuckets" "DeleteBucket" \  
  --region "us-east-1"
```

Searching action connectors

Use the `search-action-connectors` command to find action connectors based on specific criteria.

```
aws quicksight search-action-connectors \  
  --aws-account-id "123456789012" \  
  --max-results 5 \  
  --filters ' [{  
    "Name": "ACTION_CONNECTOR_NAME",  
    "Operator": "StringLike",  
    "Value": "S3"  
  } ]' \  
  --region "us-east-1"
```

Managing action connector permissions

Use the `update-action-connector-permissions` command to grant or revoke permissions for an action connector.

```
aws quicksight update-action-connector-permissions \  
  --aws-account-id "123456789012" \  
  --region "us-east-1"
```

```
--action-connector-id "my-s3-connector-id" \  
--grant-permissions '[{  
  "Actions": [  
    "quicksight:DescribeActionConnector",  
    "quicksight:UpdateActionConnector",  
    "quicksight:DeleteActionConnector"  
  ],  
  "Principal": "arn:aws:quicksight:us-east-1:123456789012:user/default/myuser"  
}]' \  
--region "us-east-1"
```

Viewing action connector permissions

Use the `describe-action-connector-permissions` command to view current permissions for an action connector.

```
aws quicksight describe-action-connector-permissions \  
  --aws-account-id "123456789012" \  
  --action-connector-id "my-s3-connector-id" \  
  --region "us-east-1"
```

Deleting an action connector

Use the `delete-action-connector` command to remove an action connector from your account.

```
aws quicksight delete-action-connector \  
  --aws-account-id "123456789012" \  
  --action-connector-id "my-s3-connector-id" \  
  --region "us-east-1"
```

Next steps

After understanding action connector APIs, you can:

- Review the complete API reference documentation for detailed parameter specifications.

- Explore specific connector setup guides for your target services.
- Implement authentication flows appropriate for your use case.
- Set up monitoring and error handling for production deployments.
- Configure permissions and access controls for your organization.

Authentication methods

Amazon Quick Suite supports multiple authentication methods, each designed for specific use cases and security requirements.

Managed authentication (3LO)

Three-Legged OAuth (3LO) is the recommended authentication method for personal access to third-party services.

Key features of 3LO:

- No initial configuration required.
- User-specific authentication.
- Secure credential storage.
- Automatic token refresh.
- 90-day refresh token lifecycle.

3LO setup process:

1. Select connector.
2. Choose managed authentication.
3. Complete service provider login.
4. Grant requested permissions.
5. Confirm connection.

Custom user-based authentication

For scenarios that require specific organizational control or custom configuration.

Required information:

- Client ID.
- Client Secret.
- Domain URL.
- Authorization URL.
- Token URL.
- Redirect URL.

Configuration steps:

1. Obtain credentials from service provider.
2. Configure authentication settings.
3. Validate connection.
4. Test access permissions.

When configuring user-based authentication in the Amazon Quick Suite console, obtain the proper credentials from your service provider and configure your authentication settings. Then validate the connection and test your access permissions.

API key authentication

Used primarily for automated workflows and system-level access.

Key features:

- Simple token-based authentication.
- Single credential management.
- Service-level permissions.
- Suitable for automated processes.

Setup requirements:

When setting up API Key authentication, ensure that you have the following:

- Valid API key from service.

- Appropriate service permissions.
- Secret storage configuration.

Service-to-service authentication

For automated workflows that require complex authentication.

Configuration requirements:

- Client ID.
- Client Secret.
- Domain URL.
- Token URL.
- Service-specific parameters.

Action execution methods

Amazon Quick Suite provides multiple ways to execute actions, accommodating different use cases and interaction preferences.

Chat interface

You can execute implicit actions in the Amazon Quick Suite chat.

Implicit actions

Amazon Quick Suite also supports implicit action execution through natural conversation with agents. Using advanced natural language processing, the system can identify when your conversation indicates the need for specific actions. Conversations are analyzed to determine which actions are required to fulfill your request.

A single request might require multiple actions to complete. When this happens, the system handles these actions sequentially, guiding you through each step. For each identified action, the system presents the appropriate form for you to complete. After each action completes, you receive a confirmation before moving on to the next action in the sequence.

For example, if you ask "Create a Jira ticket for this issue and notify the team in Slack," the system would:

1. First present the Jira ticket creation form.
2. After completing the ticket creation, show the Slack message form.
3. Complete both actions in sequence.

Throughout the process, you can track your progress through multiple actions. When all actions complete, the system provides a comprehensive summary showing all executed actions and their outcomes. You can access related documentation if needed and review any error states that may have occurred during the process.

Monitoring and maintenance

Monitoring your action connectors helps ensure reliable performance and identify issues before they impact users. Regular monitoring allows you to track usage patterns, optimize performance, and maintain healthy connections to external services.

Performance monitoring

You can assess action connector performance using the following metrics and analytics.

CloudWatch metrics

- Action execution success rates - Track the percentage of successful action executions to identify reliability issues.
- Response times - Monitor how long actions take to complete and identify performance bottlenecks.
- Error frequencies - Track error patterns to identify common failure points and areas for improvement.
- API quota usage - Monitor usage against service limits to prevent throttling and plan for capacity.

Usage analytics

The following usage analytics are collected for action connectors:

- Active users - Track how many users are actively using action connectors to understand adoption and usage patterns.

- Popular actions - Identify which actions are used most frequently to prioritize optimization efforts.
- Execution patterns - Analyze when and how often actions are executed to optimize resource allocation.
- Error trends - Monitor error patterns over time to identify systemic issues and improvement opportunities.

Connection health

You can assess action connector health using the following connection health tools:

Status monitoring

- Connection state - Monitor whether connectors are actively connected and functioning properly.
- Authentication validity - Track the status of authentication tokens and credentials to prevent access failures.
- Token expiration tracking - Monitor when authentication tokens will expire and need renewal.
- Service availability - Track the availability and response status of connected external services.

Automated maintenance

- Token refresh handling.
- Connection recovery.
- Error retry logic.
- Performance optimization.

CloudWatch metrics reference

Available CloudWatch metrics

Metric	Description	Unit
ActionSuccess	Successful executions	Count
ActionLatency	Execution time	Milliseconds

Metric	Description	Unit
AuthFailures	Failed authentications	Count
APIThrottling	API throttling events	Count

Best practices

Following best practices for action connectors helps ensure secure, reliable, and efficient operations. These practices help you maintain optimal performance, protect sensitive data, and minimize operational issues.

Security

Authentication management

- Regular credential rotation - Update API keys and OAuth tokens on a scheduled basis to maintain security.
- Periodic permission reviews - Audit user and service permissions quarterly to ensure least-privilege access.
- Token lifecycle monitoring - Track token expiration dates and set up alerts before credentials expire.
- Access audit logging - Enable comprehensive logging to track who accessed which services and when.

Access control

- Implement least-privilege access - Grant only the minimum permissions necessary for each action to function properly.
- Regular permission audits - Review and validate that current permissions align with actual usage patterns and business needs.
- Document access patterns - Maintain clear documentation of who has access to which connectors and why.
- Monitor usage anomalies - Set up alerts for unusual access patterns that might indicate security issues.

Performance

Action configuration

- Optimize form defaults - Pre-populate commonly used values to reduce user input time and errors.
- Configure appropriate timeouts - Set realistic timeout values based on typical response times for each service.
- Set up error handling - Implement robust error handling with clear user messages and retry logic where appropriate.
- Document dependencies - Clearly document any prerequisites or dependencies between different actions.

Resource management

- Monitor API quotas.
- Track usage patterns.
- Optimize refresh schedules.
- Regular cleanup of unused connectors.

Maintenance

Regular actions

- Review connector status.
- Update configurations.
- Validate connections.
- Document changes.

Troubleshooting

- Monitor error patterns.
- Review CloudWatch logs.
- Track resolution times.
- Document solutions.

Troubleshooting

When action connectors encounter issues, systematic troubleshooting helps you quickly identify and resolve problems. This guidance covers common issues and their solutions to minimize downtime and restore functionality.

Common issues and solutions

Authentication problems

Token expiration

Symptom: "Authentication token expired" error
Resolution:

1. Choose "Reconnect" in console.
2. Complete authentication flow.
3. Retry action.

Permission errors

Symptom: "Insufficient permissions" message
Resolution:

1. Verify service permissions.
2. Check connector configuration.
3. Review action requirements.

Connection failures

Symptom: "Unable to connect to service" error
Resolution:

1. Verify service availability.
2. Check network connectivity.
3. Validate credentials.
4. Review service quotas.

Action-specific issues

Form submission failures

Validation errors

- Check required fields.
- Verify data formats.
- Review field limitations.
- Check for special characters.

Timeout issues

- Reduce form complexity.
- Check network latency.
- Review service response times.
- Consider breaking into multiple actions.

Sync and performance issues

Slow response times

Resolution:

1. Check API rate limits.
2. Review concurrent executions.
3. Monitor service health.
4. Optimize action configuration.

Failed executions

Resolution:

1. Review CloudWatch logs.
2. Check error messages.
3. Verify service status.

4. Test connection health.

Common error messages

Error codes and resolutions

Error code	Description	Resolution
AUTH_001	Authentication failed	Verify credentials and retry
CONN_002	Connection timeout	Check network and service status
PERM_003	Insufficient permissions	Review required permissions
TOKEN_004	Token expired	Reinitiate authentication

Integration-specific guides

Use these guides to configure integrations with specific applications. Each guide includes setup instructions for both action and data ingestion capabilities when the integration supports them.

Amazon S3 integration

With Amazon S3 integration in Amazon Quick Suite, you can create knowledge bases from documents stored in S3 buckets. This integration supports data ingestion capabilities for indexing and searching S3 content. Amazon S3 actions are only supported for Quick Automate.

Amazon Quick Suite supports source attribution with citations. If you specify the `_source_uri` metadata field when you add metadata to your Amazon S3 bucket, the source attribution links returned by Amazon Quick Suite in the chat results will direct users to the configured URL. If you don't specify a `_source_uri`, users can still access the source documents through clickable citation links that will download the file at query time. This allows users to verify information even when no source URI is configured. To learn how to add metadata for your Amazon S3 connector, see [the section called "Adding document metadata in Amazon S3"](#).

What you can do

Amazon S3 users can ask questions about content stored in their Amazon S3 buckets. For example, users can inquire about key findings from documents, search for specific information across

multiple file types, or analyze data patterns. The integration enables users to quickly access and understand information from their Amazon S3 content, regardless of file location or type, while providing contextual details such as modification dates, and file metadata—all contributing to more efficient information discovery and better-informed decision making.

Note

This guide covers Amazon S3 data ingestion integration for knowledge base creation. For Amazon S3 action connectors that perform Amazon S3 operations (upload, download, delete files), these must be created through the admin console. For more information, see [the section called “Creating action connectors using the admin console”](#).

Before you begin

Before you set up Amazon S3 integration, make sure you have the following:

- Amazon account with Amazon S3 access.
- Amazon S3 bucket with documents to index.
- Amazon Quick Suite Enterprise subscription.
- Necessary permissions to create Amazon S3 integrations.

If you need to access Amazon S3 buckets in a different Amazon account, verify that cross-account access has been enabled by your administrator.

Note

Cross-account Amazon S3 access is only supported within the same Amazon region.

Enable cross-account access (administrators only)

If you need to enable cross-account Amazon S3 access for your organization, complete the following steps.

To enable cross-account Amazon S3 access

1. Open the Amazon Quick Suite Admin console.

2. Choose **Amazon resource page**, and then choose **Amazon S3 configuration**.
3. Select **Choose accessible buckets from other Amazon accounts**.

Prepare IAM role and policy setup

Before setting up the integration in Amazon Quick Suite, prepare your IAM role and policy configuration. Amazon S3 integration uses Amazon authentication to access your Amazon S3 buckets.

Required IAM permissions

Make sure your Amazon account has the following minimum permissions for the Amazon S3 bucket:

- `s3:GetObject` - Read objects from the bucket.
- `s3:ListBucket` - List bucket contents.
- `s3:GetBucketLocation` - Get bucket region information.
- `s3:GetObjectVersion` - Get object versions.
- `s3:ListBucketVersions` - List bucket versions.

Configure Amazon S3 bucket permissions for cross-account access

If you're accessing Amazon S3 buckets in a different Amazon account, you must configure IAM policies in the source Amazon account.

To configure Amazon S3 bucket permissions for cross-account access

1. Sign in to the Amazon Management Console for the account that contains the Amazon S3 bucket.
2. Open the Amazon S3 console at <https://console.aws.amazon.com/s3/>.
3. Choose the bucket that you want to grant access to.
4. Choose **Permissions**, and then choose **Bucket Policy**.
5. Add a bucket policy with the following elements:
 - **Version** – Set to "2012-10-17"
 - **Statement** – Array containing policy statements with:
 - **Sid** – "AllowQuickSuiteS3Access"

- **Effect** – "Allow"
- **Principal** – Amazon ARN for the Amazon Quick Suite service role in your account.
For example, the principal should look like this: `"Principal": { "AWS": "arn:aws:iam::<quick_account_id>:role/service-role/aws-quicksight-service-role-v0" }`
- **Action** – Array of Amazon S3 permissions: `s3:GetObject`, `s3:ListBucket`, `s3:GetBucketLocation`, `s3:GetObjectVersion`, `s3:ListBucketVersions`
- **Resource** – "*" (applies to the current key), the Amazon S3 bucket path should look like the following: `"Resource": [ "arn:aws:s3:::bucket_name" ]`

6. Choose **Save changes**.

Configure KMS key permissions (if your bucket uses encryption)

If your Amazon S3 bucket uses Amazon KMS encryption, complete the following steps.

1. Open the Amazon Key Management Service (Amazon KMS) console at <https://console.aws.amazon.com/kms>.
2. Choose the KMS key that is used to encrypt your Amazon S3 bucket.
3. Choose **Key policy**, and then choose **Edit**.
4. Add a statement to the key policy with the following structural elements:
 - **Sid** – "AllowQuickSuiteKMSAccess"
 - **Effect** – "Allow"
 - **Principal** – Amazon ARN for the Amazon Quick Suite service role in your account.
For example, the principal should look like this: `"Principal": { "AWS": "arn:aws:iam::<quick_account_id>:role/service-role/aws-quicksight-service-role-v0" }`
 - **Action** – Array of KMS permissions: `kms:Decrypt`, `kms:DescribeKey`
 - **Resource** – "*" (applies to the current key), the Amazon S3 bucket path should look like the following: `"Resource": [ "arn:aws:s3:::bucket_name" ]`
5. Choose **Save changes**.
6. Wait 2-3 minutes for the policy changes to propagate.

During the integration setup, you will need to:

- Verify the connection and bucket access.

Set up Amazon S3 integration

After preparing your IAM role and policy configuration, follow these steps to create your Amazon S3 integration:

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Choose **Add** (the plus + button).
3. Choose your Amazon account option:
 - **Default Account** – Use this option to access Amazon S3 buckets in the same Amazon account where Amazon Quick Suite is enabled. This option is selected by default.
 - **Other Amazon Account** – Use this option to access Amazon S3 buckets in a different Amazon account.

Note

If order for **Other Amazon Account** to be used, your admin must enable the feature.

Important

Cross-account Amazon S3 access is only supported within the same Amazon region.

4. If you selected **Other Amazon Account**, enter the Amazon account ID that contains the Amazon S3 bucket.
5. Fill in the integration details:
 - **Name** - Descriptive name for your Amazon S3 integration.
 - **Amazon S3 bucket URL** - The Amazon S3 bucket path containing your documents. Your Amazon S3 bucket should be in the same region as your Amazon Quick Suite region.
 - **Metadata Folder** - Specific folder within the bucket (optional).
6. Select **Create and continue**. The system validates your configuration. If errors occur, review the error message for specific remediation steps. You can use the Amazon S3 policy details specified above and copy them into your S3 bucket.

7. On the **Add files or folders** page, select the files to place in your knowledge base. You select the files you want Amazon Quick Suite to sync with using a point-and-click experience.
8. Complete the knowledge base details:
 - **Name** - Enter a descriptive name for your knowledge base.
 - **Description** - Describe the purpose of this knowledge base.
9. Select **Create**.

After clicking create, the data sync is started automatically.

Manage knowledge bases

After setting up your Amazon S3 integration, you can create and manage knowledge bases from your Amazon S3 content.

Edit existing knowledge bases

You can modify your existing Amazon S3 knowledge bases:

1. In the Amazon Quick Suite console, choose **Knowledge bases**.
2. Select your Amazon S3 knowledge base from the list.
3. Choose the three-dot icon under **Actions**, then choose **Edit knowledge base**.
4. Update your configuration settings as needed and choose **Save**.

Create additional knowledge bases

You can create multiple knowledge bases from the same Amazon S3 integration:

1. In the Amazon Quick Suite console, choose **Integrations**, and then select the **Data** tab.
2. Choose your existing Amazon S3 integration from the list.
3. Choose the three-dot icon under **Actions**, then choose **Create knowledge base**.
4. Configure your knowledge base settings and choose **Create**.

For detailed information about knowledge base configuration options, see [the section called "Common configuration settings"](#).

Note

Amazon Quick Suite doesn't sync ACLs from data sources. When you create a knowledge base in Amazon Quick Suite, by default only you can get insights from the knowledge base. For shared content, you can provide access to different users and groups by updating the knowledge base permissions.

Adding document metadata in Amazon S3

To customize chat results for your end users, you can add metadata or document attributes to documents in an Amazon S3 bucket by using a metadata file. Metadata is additional information about a document, such as its title and the date and time it was created.

Amazon Quick Suite supports source attribution with citations. If you specify the `_source_uri` metadata field when you add metadata to your Amazon S3 bucket, the source attribution links returned by Amazon Quick Suite in the chat results will direct users to the configured URL. If you don't specify a `_source_uri`, users can still access the source documents through clickable citation links that will download the file at query time. This allows users to verify information even when no source URI is configured.

Document metadata location

In Amazon S3, each metadata file can be associated with an indexed document. Your metadata files must be stored in the same Amazon S3 bucket as your indexed files. You can specify a location within the Amazon S3 bucket for your metadata files when configuring your Amazon S3 integration in Amazon Quick Suite.

If you don't specify an Amazon S3 prefix, your metadata files must be stored in the same location as your indexed documents. If you specify an Amazon S3 prefix for your metadata files, they must be in a directory structure parallel to your indexed documents. Amazon Quick Suite looks only in the specified directory for your metadata. If the metadata isn't read, check that the directory location matches the location of your metadata.

The following examples show how the indexed document location maps to the metadata file location. The document's Amazon S3 key is appended to the metadata's Amazon S3 prefix and then suffixed with `.metadata.json` to form the metadata file's Amazon S3 path.

Note

The combined Amazon S3 key, the metadata's Amazon S3 prefix, and the `.metadata.json` suffix must be no more than a total of 1,024 characters. We recommend that your Amazon S3 key is less than 1,000 characters to account for additional characters when combining your key with the prefix and suffix.

Example Example 1: No metadata path specified

```
Bucket name:
    s3://bucketName
Document path:
    documents
Metadata path:
    none
File mapping
    s3://bucketName/documents/file.txt ->
    s3://bucketName/documents/file.txt.metadata.json
```

Example Example 2: Metadata path specified

```
Bucket name:
    s3://bucketName
Document path:
    documents/legal
Metadata path:
    metadata
File mapping
    s3://bucketName/documents/legal/file.txt ->
    s3://bucketName/metadata/documents/legal/file.txt.metadata.json
```

Document metadata structure

You define your document metadata itself in a JSON file. The file must be a UTF-8 text file without a BOM marker. The file name of the JSON file must be

<document>.<extension>.metadata.json. In this example, document is the name of the document that the metadata applies to and extension is the file extension for the document. The document ID must be unique in <document>.<extension>.metadata.json.

The content of the JSON file uses the following template:

```
{
  "DocumentId": "document ID",
  "Attributes": {
    "_authors": ["author of the document"],
    "_category": "document category",
    "_created_at": "ISO 8601 encoded string",
    "_last_updated_at": "ISO 8601 encoded string",
    "_source_uri": "document URI",
    "_version": "file version",
    "_view_count": number of times document has been viewed
  },
  "Title": "document title",
  "ContentType": "PDF | HTML | MS_WORD | PLAIN_TEXT | PPT | RTF | XML | XSLT |
MS_EXCEL | CSV | JSON | MD"
}
```

If you provide a metadata path, make sure that directory structure inside the metadata directory exactly matches the directory structure of data file.

For example, if the data file location is at `s3://bucketName/documents/legal/file.txt`, the metadata file location should be at `s3://bucketName/metadata/documents/legal/file.txt.metadata.json`.

All of the attributes and fields are optional, so it's not necessary to include all attributes. However, you must provide a value for each attribute that you want to include; the value can't be empty.

The `_created_at` and `_last_updated_at` metadata fields are ISO 8601 encoded dates. For example, `2012-03-25T12:30:10+01:00` is the ISO 8601 date-time format for March 25, 2012, at 12:30PM (plus 10 seconds) in the Central European Time time zone.

Troubleshooting Amazon S3 integration issues

If you encounter issues connecting to your Amazon S3 bucket, review the following common causes and solutions.

Cross-account access not configured

Issue: Your administrator hasn't granted access to use Amazon S3 buckets from other Amazon accounts in Amazon Quick Suite.

Solution for administrators:

1. Navigate to Amazon Quick Suite Admin settings → Manage Amazon Quick Suite → Manage Amazon resources.
2. Under **Allow access and autodiscovery for these resources**, choose **Select Amazon S3 Bucket** and enable **Amazon S3 Buckets You Can Access Across Amazon**.
3. Select **Choose accessible buckets from Amazon account**.

Contact your administrator to confirm cross-account access has been properly configured.

Bucket not in approved list

If your admin is only allowing certain buckets to be accessed, check with your admin to make sure your bucket is in the approved list.

Issue: The bucket you're trying to access hasn't been authorized by your administrator.

Solution:

- Confirm the bucket name is spelled correctly.
- Verify with your administrator that the bucket is included in the approved list.
- Request your administrator to add the bucket to the authorized buckets list if needed.

Insufficient IAM permissions

Issue: Your IAM role or user lacks the necessary permissions to access the Amazon S3 bucket.

Solution:

- Verify your IAM policy includes the required Amazon S3 permissions:
 - `s3:GetObject`
 - `s3:ListBucket`
 - `s3:GetBucketLocation`

- `s3:GetObjectVersion`
- `s3:ListBucketVersions`
- Check your own buckets for any explicit Deny statements that might be blocking access.

Note

The ARN `arn:aws:iam::account-id:role/service-role/aws-quicksight-service-role-v0` is the default service role used when no custom role has been created. If a custom service role exists, contact your administrator to obtain the custom service role ARN and use it instead of the default.

Cross-region restrictions

Issue: The Amazon S3 bucket is located in a different Amazon region than your Amazon Quick Suite account or service.

Solution:

- Verify the bucket region matches your Amazon Quick Suite service region.
- Check bucket region using Amazon CLI: `aws s3api get-bucket-location --bucket bucket-name`
- Use a bucket in the same region as your service.

Additional troubleshooting steps

- **Test bucket accessibility** using Amazon CLI:

```
aws s3 ls s3://bucket-name --profile your-profile
```

- **Review CloudTrail logs** for AccessDenied errors to identify the specific permission issue.
- **Check Amazon S3 Block Public Access settings** - while these typically don't affect authenticated access, verify they're not interfering with your specific use case.
- **Verify bucket ownership** - ensure the bucket exists and you have the correct bucket name.

Limitations

When using Amazon S3 integrations in Amazon Quick Suite, be aware of the following limitations:

- The Amazon S3 bucket must be in the same Amazon Region as your Amazon Quick Suite application.

Asana integration

Connect Amazon Quick Suite to your Asana workspace to manage projects, tasks, and team collaboration. You can create, update, and manage Asana content without leaving your Amazon Quick Suite environment. This integration requires Amazon Quick Suite Pro tier or higher.

What you can do

With Asana integration, you can perform actions within your Asana workspaces through the Asana API.

Action connector

Create, update, and manage projects, tasks, and team assignments through the Asana API.

Set up Asana integration

Follow these steps to connect Amazon Quick Suite to your Asana workspace.

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Click **Add** (plus "+" button).
3. Fill in the following details:
 - **Name** - Enter a descriptive name for your Asana integration.
 - **Description** - Describe the purpose of this integration.
4. Choose the connection type and configure the network type settings.
5. Configure the authentication settings based on your chosen authentication method.
6. Select **Create and continue**.
7. Add users to share the integration with.
8. Click **Next**.

Configure authentication

Asana integration uses custom user-based OAuth authentication. Configure the following authentication fields:

- **Base URL** - Asana API base URL.
- **Client ID** - Your Asana OAuth app client ID.
- **Client Secret** - Your Asana OAuth app client secret.

Required OAuth scopes

When you create your Asana OAuth application, configure these scopes:

- `tasks:write` - Create and modify tasks.
- `tasks:read` - Read task information.
- `workspaces:read` - Access workspace information.
- `workspaces.typeahead:read` - Search within workspaces.
- `stories:read` - Read task stories and comments.
- `users:read` - Access user information.
- `projects:read` - Read project information.
- `project_templates:read` - Access project templates.

Manage Asana integrations

You can perform these management tasks for your Asana integrations:

- **Edit integration settings** - Update authentication settings or Asana configuration.
- **Share integration access** - Make the integration available to other users.
- **Delete integration** - Remove the integration and revoke authentication.

BambooHR integration

Connect Amazon Quick Suite to your BambooHR system to manage employee data, time-off requests, and HR processes. You can create, update, and manage HR content without leaving your Amazon Quick Suite environment. This integration requires Amazon Quick Suite Pro tier or higher.

What you can do

With BambooHR integration, you can perform actions within your BambooHR systems through the BambooHR API.

Action connector

Create, update, and manage employee records, time-off requests, and other HR processes through the BambooHR API.

Set up BambooHR integration

Follow these steps to connect Amazon Quick Suite to your BambooHR system.

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Click **Add** (plus "+" button).
3. Complete the integration details:
 - **Name** - Enter a descriptive name for your BambooHR integration.
 - **Description** (Optional) - Describe the purpose of this integration.
4. Choose connection type (user or service authentication).
5. Complete the connection settings based on authentication method:
 - **For User authentication (OAuth):**
 - **Base URL** - Your BambooHR instance URL.
 - **Client ID** - OAuth client identifier.
 - **Client Secret** - OAuth client secret.
 - **Token URL** - OAuth token endpoint.
 - **Auth URL** - OAuth authorization endpoint.
 - **Redirect URL** - OAuth callback URL.

Required OAuth scopes: When you create your BambooHR OAuth application, configure the scopes you need for your use case. Common scopes include:

- `employee` - Access employee information.
- `employee.write` - Write employee information.
- `time_off` - Access time off information.

- `time_off.write` - Write time off information.
- `company:info` - Access company information.
- `payroll` - Access payroll data.

 Note

Additional scopes may be required depending on the specific BambooHR actions you plan to use. Consult your BambooHR administrator for the complete list of available scopes.

- **For Service authentication (API Key):**
 - **Base URL** - Your BambooHR instance URL.
 - **API Key** - Your BambooHR API key.
 - **Email** - Email address associated with the API key.
6. Select **Create and continue**.
 7. Select users to share the integration with.
 8. Click **Next**.

Manage BambooHR integration

You can perform these management tasks for your BambooHR integration:

Edit integration settings

To modify your BambooHR integration settings:

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Choose **BambooHR** from the integration grid.
3. Select your integration from the list and choose **Edit**.
4. Modify your integration settings as needed.
5. Choose **Save changes**.

Share integration

You can share your BambooHR action connector with other users in your organization.

1. From the BambooHR integration details page, choose **Share**.
2. Configure your sharing options and permission levels.
3. Choose **Share integration**.

Delete integration

To permanently remove your BambooHR integration:

1. From the BambooHR integration details page, choose **Delete**.
2. Review the deletion impact, including any workflows or automations using this integration.
3. Type the integration name to confirm deletion.
4. Choose **Delete integration**.

What happens next

After you complete the setup, your BambooHR integration appears in the integrations list. You can use it in Amazon Quick Suite workflows, automations, and AI agents to perform HR-related tasks.

Troubleshoot BambooHR integration

If you encounter issues with your BambooHR integration, try these solutions:

Authentication failures

Verify that your BambooHR credentials are correct and that the API key or OAuth application has the necessary permissions.

Connection timeouts

Check that your BambooHR instance URL is correct and accessible from Amazon Quick Suite.

Permission errors

Ensure that the authenticated user or API key has the required permissions to perform the requested HR operations in BambooHR.

Action execution failures

Review the action parameters and ensure they match the expected format for BambooHR API calls.

Atlassian Confluence integration

With Atlassian Confluence integration in Amazon Quick Suite, you can perform actions on Confluence content and create knowledge bases from Confluence spaces, pages, and blog posts. This integration supports both action capabilities and data ingestion capabilities.

What you can do

Confluence users can ask questions about content stored in their Confluence spaces and pages. For example, users can inquire about project documentation, team wikis, or search for specific information across multiple spaces. The integration enables users to quickly access and understand information from their Confluence content, regardless of location or type, while providing contextual details such as publication dates, modification history, and page ownership—all contributing to more efficient information discovery and better-informed decision making.

With Confluence actions, you can create, update, and manage pages, spaces, and other Confluence objects directly through Amazon Quick Suite. These action capabilities allow you to streamline content management tasks without leaving your Amazon Quick Suite environment.

Before you begin

Before you set up Confluence integration, make sure you have the following:

- Atlassian Confluence Cloud
- For action connectors: Amazon Quick Suite Enterprise subscription.
- For data access: Amazon Quick Suite Professional subscription or higher.

Prepare Atlassian app setup and authentication

Before setting up the integration in Amazon Quick Suite, prepare your Atlassian app and authentication credentials. Confluence integration supports different authentication methods depending on your integration type.

Action connector authentication setup

For action connectors, gather authentication credentials using one of these methods:

User authentication (3LO)

Gather the following information from your Atlassian app setup:

- **Base URL** - Your Confluence instance URL. This is the URL used for API calls, and not the same URL that users of Confluence log into. It will look something like the following example: <https://api.atlassian.com/ex/confluence/yourInstanceId>
- **Client ID** - Atlassian app client ID.
- **Client Secret** - Atlassian app client secret.
- **Token URL** - Atlassian OAuth token endpoint.
- **Auth URL** - Atlassian OAuth authorization endpoint.
- **Redirect URL** - OAuth redirect URI.

Required OAuth scopes:

- `search:confluence` - Search Confluence content. This is a classic scope in Confluence.
- `read:page:confluence` - Read page content. This is a granular scope in Confluence.
- `write:page:confluence` - Create and modify pages. This is a granular scope in Confluence.
- `read:space:confluence` - Access space information. This is a granular scope in Confluence.

Service authentication (API Key)

Gather the following information from your Confluence administrator:

- **API Key** - Confluence API token.
- **Base URL** - Your Confluence instance URL, used for API calls.
- **Email** - Associated user account email.

Data access authentication setup

For **Bring data from Atlassian Confluence** integrations, prepare for Confluence authentication. During the integration setup, you will need to:

1. Complete the Confluence authentication popup that appears.
2. Grant permissions for Amazon Quick Suite to access your Confluence content.
3. Review and complete the authentication process.

Set up Confluence integration

After preparing your Atlassian app and authentication credentials, follow these steps to create your Confluence integration:

Set up knowledge base integration

To create a Confluence integration for knowledge base creation:

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Click **Add** (plus "+" button).
3. Select **Bring data from Atlassian Confluence Cloud**.
4. Select **Next**.
5. Complete the Name and Confluence URL fields:
 - **Name** - Descriptive name for your integration.
 - **Confluence URL** - The URL of your Atlassian site, will resemble the following example:
your-site.atlassian.net
6. Select **Next**.
7. Add Integration details and complete authentication.
8. Select **Next**.
9. Complete the Knowledge base details:
 - **Knowledge base name** - Name for your knowledge base.
 - **Description** - Purpose of the knowledge base.
 - **Content selection** - Choose spaces and pages to include.

Set up action connector integration

To create a Confluence integration for performing actions:

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Click **Add** (plus "+" button).
3. Select **Perform actions**.
4. Select **Next**.
5. Complete the integration details:

- **Name** - Descriptive name for your integration.
 - **Description** - Purpose of the integration.
6. Choose connection type (user or service authentication).
 7. Configure authentication settings based on your chosen method:
 - **For User authentication (3LO):**
 - **Base URL** - Your Confluence instance URL used for API calls.
 - **Client ID** - Atlassian app client ID.
 - **Client Secret** - Atlassian app client secret.
 - **Token URL** - Atlassian OAuth token endpoint.
 - **Auth URL** - Atlassian OAuth authorization endpoint.
 - **Redirect URL** - OAuth redirect URI.
 - **For Service authentication (API Key):**
 - **API Key** - Confluence API token.
 - **Base URL** - Your Confluence instance URL, which is the URL used for API calls.
 - **Email** - Associated user account email.
 8. Select **Create and continue**.
 9. Select users to share the integration with.
 10. Select **Next**.

Manage knowledge bases

After setting up your Confluence integration, you can create and manage knowledge bases from your Confluence content.

Edit existing knowledge bases

You can modify your existing Confluence knowledge bases:

1. In the Amazon Quick Suite console, choose **Knowledge bases**.
2. Select your Confluence knowledge base from the list.
3. Choose the three-dot icon under **Actions**, then choose **Edit knowledge base**.
4. Update your configuration settings as needed and choose **Save**.

Create additional knowledge bases

You can create multiple knowledge bases from the same Confluence integration:

1. In the Amazon Quick Suite console, choose **Integrations**, and then select the **Data** tab.
2. Choose your existing Confluence integration from the list.
3. Choose the three-dot icon under **Actions**, then choose **Create knowledge base**.
4. Configure your knowledge base settings and choose "Create".

For detailed information about knowledge base configuration options, see [the section called "Common configuration settings"](#).

Supported content types

You can include these content types in your knowledge base:

- Confluence pages and blog posts.
- Spaces content
- Page and blog attachments

Note

Amazon Quick Suite doesn't sync ACLs from data sources. When you create a knowledge base in Amazon Quick Suite, by default only you can get insights from the knowledge base. For shared content, you can provide access to different users and groups by updating the knowledge base permissions.

Google Drive integration

With Google Drive integration in Amazon Quick Suite, you can create knowledge bases from documents stored in Google Drive. This integration supports data access capabilities for indexing and searching Google Drive content including Google Docs, Sheets, and Slides.

What you can do

Google Drive users can ask questions about content stored in their Google Drive. For example, users can inquire about key findings from Google Docs, presentation highlights from Google Slides,

or search for specific information across multiple document types. The integration enables users to quickly access and understand information from their Google Drive content, regardless of file location or type, while providing contextual details such as publication dates, modification history, and document ownership—all contributing to more efficient information discovery and better-informed decision making.

Note

Google Drive integration supports data ingestion only. It doesn't provide action capabilities for managing Google Drive files through APIs.

Before you begin

Before you set up Google Drive integration, make sure you have the following:

- Google account with Google Drive access.
- Google Drive files and folders to index
- Amazon Quick Suite Professional subscription or higher.

Prepare Google Drive authentication

Google Drive integration uses managed three-legged OAuth (3LO) authentication. Before setting up the integration in Amazon Quick Suite, understand the authentication process and required permissions:

Authentication methods

Managed three-legged OAuth (3LO)

Personal access method for individual Google Drive content:

- Requires user authentication and consent
- Provides access to personal and shared content

Required permissions and scopes

Amazon Quick Suite requests the following Google Drive permissions during authentication:

- <https://www.googleapis.com/auth/drive.readonly> - Read-only access to Google Drive files and folders
- <https://www.googleapis.com/auth/drive.metadata.readonly> - Read-only access to file metadata, sharing settings, and properties

OAuth authentication process

During the integration setup, you will be prompted to:

1. Select your Google account from the list or sign in with your credentials.
2. Review the permissions that Amazon Quick Suite is requesting:
 - See and download all your Google Drive files.
 - See information about your Google Drive files.
3. Select both permissions and click **Continue** to grant permissions and complete the authentication process.

Set up Google Drive integration

After understanding the authentication requirements, follow these steps to create your Google Drive integration:

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Click **Add** (plus "+" button).
3. Select your account and complete the authentication process.
4. Select **Continue**.
5. On the "Add files or folders" section, select the files to place in your knowledge base. You select the files you want Amazon Quick Suite to sync with using a point-and-click experience.
6. Complete the knowledge base details:
 - **Name** - Enter a descriptive name for your knowledge base.
 - **Description** (optional) - Describe the purpose of this knowledge base.
7. Select **Create**.

After clicking create, the data sync is started automatically.

Manage knowledge bases

After setting up your Google Drive integration, you can create and manage knowledge bases from your Google Drive content.

Edit existing knowledge bases

You can modify your existing Google Drive knowledge bases:

1. In the Amazon Quick Suite console, choose **Knowledge bases**.
2. Select your Google Drive knowledge base from the list.
3. Choose the three-dot icon under **Actions**, then choose **Edit knowledge base**.
4. Update your configuration settings as needed and choose **Save**.

Create additional knowledge bases

You can create multiple knowledge bases from the same Google Drive integration:

1. In the Amazon Quick Suite console, choose **Integrations**, and then select the **Data** tab.
2. Choose your existing Google Drive integration from the list.
3. Choose the three-dot icon under **Actions**, then choose **Create knowledge base**.
4. Configure your knowledge base settings and choose **Create**.

For detailed information about knowledge base configuration options, see [the section called "Common configuration settings"](#).

Note

Amazon Quick Suite doesn't sync ACLs from data sources. When you create a knowledge base in Amazon Quick Suite, by default only you can get insights from the knowledge base. For shared content, you can provide access to different users and groups by updating the knowledge base permissions.

Limitations

When using Google Drive integrations in Amazon Quick Suite, be aware of the following limitations:

- File comments synchronization is not supported..

Jira Cloud integration

Connect Amazon Quick Suite to your Jira Cloud instance to manage issues, projects, and other Jira objects. You can create, update, and query Jira content without leaving your Amazon Quick Suite environment. This integration requires Amazon Quick Suite Pro tier or higher.

What you can do

With Jira Cloud integration, you can perform actions within your Jira instances through the Jira REST API.

Action connector

Create, update, and query Jira issues, projects, and other objects through the Jira REST API.

Note

This integration supports action execution only. You can't use it to create knowledge bases or access Jira data for search purposes.

Before you begin

You need the following to set up Jira integration:

- Jira Cloud instance with appropriate permissions.
- Jira user account or API token credentials.
- Amazon Quick Suite Author or higher.
- Administrative access to configure OAuth applications (for user authentication).

Create a Jira API token

To use API key authentication, create an API token in your Atlassian account.

1. Go to your Atlassian account settings and choose **Security**.

2. Create a new API token for Amazon Quick Suite integration.
3. Copy the generated token to use in your Amazon Quick Suite integration setup.
4. Note your Jira Cloud instance URL and user email address. Your Jira Cloud instance URL is a special value for API calls that can be found by navigating to `https://your-domain.atlassian.net/_edge/tenant_info`.

Important

Jira Cloud may return HTTP 200 success responses even when API tokens are revoked or improperly configured. For more information see [JRACLOUD-82932](#). If your integration appears to connect successfully but actions fail unexpectedly, verify that your API token is valid and has not been revoked.

Set up Jira integration

Follow these steps to connect Amazon Quick Suite to your Jira Cloud instance.

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Click **Add** (plus "+" button).
3. Complete the integration details:
 - **Name** - Enter a descriptive name for your Jira integration.
 - **Description** (Optional) - Describe the purpose of this integration.
4. Choose connection type (user or service authentication).
5. Complete the connection settings based on authentication method:
 - **For User authentication (OAuth):**
 - **Base URL** - Your Jira Cloud instance URL. This must be provided in the following format: `https://api.atlassian.com/ex/jira/instanceid`. Your *instanceid* can be found by navigating to `https://your-domain.atlassian.net/_edge/tenant_info`.
 - **Client ID** - Your Atlassian OAuth app client ID.
 - **Client Secret** - Your Atlassian OAuth app client secret.
 - **Token URL** - Atlassian OAuth token endpoint.

- **Auth URL** - Atlassian OAuth authorization endpoint.
 - **Redirect URL** - OAuth redirect URI.
 - **For Service authentication (API Key):**
 - **Base URL** - Your Jira Cloud instance URL.
 - **API Key** - Your Jira API token.
 - **Email** - Email address associated with your Jira user account.
6. Select **Create and continue**.
 7. Select users to share the integration with.
 8. Click **Next**.

Configure authentication

Jira Cloud integration supports two authentication methods. Choose the method that works best for your organization.

User authentication (OAuth)

Use OAuth when you want actions to be performed under individual user identities. Configure these fields:

- **Base URL** - Your Jira Cloud instance URL.
- **Client ID** - Your Atlassian OAuth app client ID.
- **Client Secret** - Your Atlassian OAuth app client secret.
- **Token URL** - Atlassian OAuth token endpoint.
- **Auth URL** - Atlassian OAuth authorization endpoint.
- **Redirect URL** - OAuth redirect URI.

Required OAuth scopes:

- `read:jira-user` - Read user information.
- `read:jira-work` - Read issues, projects, and work items.
- `write:jira-work` - Create and modify issues and work items.
- `manage:jira-project` - Manage project settings.
- `manage:jira-configuration` - Manage Jira configuration.
- `manage:jira-webhook` - Manage webhooks.

- `read:sprint:jira-software` - Read sprint information.
- `write:sprint:jira-software` - Create and modify sprints.
- `delete:sprint:jira-software` - Delete sprints.
- `write:board-scope:jira-software` - Manage board configurations.

Service authentication (API Key)

Use API key authentication for service-to-service connections. Configure these fields:

- **Base URL** - Your Jira Cloud instance URL.
- **API Key** - Your Jira API token.
- **Email** - Email address associated with your Jira user account.

Note

Due to a known Jira Cloud behavior, the integration may appear to connect successfully even with revoked or invalid API tokens. If actions fail unexpectedly after setup, verify your API token is still valid in your Atlassian account settings.

Available actions

After you create your Jira integration, you can use these actions to interact with Jira Cloud:

- Create, read, update, and delete Jira issues.
- Manage issue transitions and workflows.
- Query issues using JQL (Jira Query Language).
- Manage projects, components, and versions.
- Add comments and attachments to issues.
- Manage user assignments and watchers.
- Access project and issue metadata.

Note

The actions you can use depend on the permissions configured in your Jira Cloud instance and your authentication method.

Share integrations

You can share your Jira action connector with other users in your organization.

1. After you create the integration, choose **Share integration**.
2. Select the users or groups you want to share the integration with.
3. Set appropriate permissions for shared access.
4. Confirm your sharing settings.

Manage Jira integrations

You can perform these management tasks for your Jira integrations:

- **Edit integration settings** - Update authentication settings or Jira instance configuration.
- **Share integration access** - Make the integration available to other users in your organization.
- **Monitor usage metrics** - View integration activity and API usage metrics.
- **Review available actions** - See the complete list of available Jira actions.
- **Delete integration** - Remove the integration and revoke associated authentication.

Model Context Protocol (MCP) integration

With Model Context Protocol (MCP) integration in Amazon Quick Suite, you can connect to MCP servers for both task execution and data access capabilities. MCP provides a standardized way to connect AI systems with external tools and data sources.

What you can do

With MCP integration, you can connect to MCP servers for both action execution and data access through standardized protocols.

Action connector

Connect to MCP servers to perform actions and execute tasks through standardized MCP protocols.

Data access integration

Access data sources through MCP servers to create knowledge bases and retrieve information.

Before you begin

Before you set up MCP integration, make sure you have the following:

- MCP server endpoint with appropriate access.
- Authentication credentials for the MCP server.
- Amazon Quick Suite Author or higher for action connectors.
- Amazon Quick Suite Professional subscription

 Note

MCP integration supports remote servers only. HTTP streaming is preferred over Server-Sent Events (SSE). Local stdio connections are not supported. VPC connectivity is not supported.

Popular MCP servers

The following are examples of popular MCP servers being used with Amazon Quick Suite and the remote MCP server URLs:

:

Popular MCP servers

Provider	MCP Server URL
Asana	https://mcp.asana.com/sse
Atlassian	https://mcp.atlassian.com/v1/sse
Amazon Knowledge MCP Server	https://knowledge-mcp.global.api.aws
Box	http://mcp.box.com/
HubSpot	https://mcp.hubspot.com
HuggingFace	https://huggingface.co/mcp

Provider	MCP Server URL
Intercom	https://mcp.intercom.com/mcp
Linear	https://mcp.linear.app/mcp
Monday	https://mcp.monday.com/sse
Notion	https://mcp.notion.com/mcp
PagerDuty	https://identity.pagerduty.com/global/oauth/anonymous/.well-known/openid-configuration
Workato	https://MCP_ID.apim.mcp.workato.com
Zapier	https://mcp.zapier.com/MCP_ID

The proceeding table shows just some of the many MCPs supported for use with Amazon Quick Suite.

 Note

These servers require appropriate authentication credentials. Refer to each provider's documentation for specific authentication requirements.

Prepare MCP server setup and authentication

Before setting up the integration in Amazon Quick Suite, prepare your MCP server configuration and authentication credentials. MCP integration supports multiple authentication methods and configuration approaches. Choose the method that matches your MCP server requirements:

User authentication (OAuth)

MCP servers support two configuration approaches for OAuth authentication:

Dynamic Client Registration (DCR)

When your MCP server advertises Dynamic Client Registration support, no manual configuration is required. The MCP client automatically registers itself with the server and receives the necessary credentials during the connection process. This streamlined approach eliminates the need to manually gather client IDs, secrets, and endpoint URLs.

Manual configuration

For MCP servers that don't support DCR, gather the following information from your MCP server configuration:

- **Client ID** - OAuth client ID.
- **Client Secret** - OAuth client secret.
- **Token URL** - OAuth token endpoint.
- **Auth URL** - OAuth authorization endpoint.
- **Redirect URL** - OAuth redirect URI.

Service authentication (Service-to-Service)

Gather the following information from your MCP server configuration:

- **Client ID** - Service client ID.
- **Client Secret** - Service client secret.
- **Token URL** - Service token endpoint.

API Key authentication

Gather the following information from your MCP server administrator:

- **API Key** - MCP server API key.

Username/Password authentication

Gather the following information from your MCP server administrator:

- **Username** - MCP server username.
- **Password** - MCP server password.

Set up MCP integration

After preparing your MCP server configuration and authentication credentials, follow these steps to create your MCP integration:

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Click **Add** (plus "+" button).
3. On the Create Integration page, enter integration details:
 - **Name** - Descriptive name for your MCP integration.
 - **Description** (Optional) - Purpose of the integration.
 - **MCP server endpoint** - URL of the MCP server.
4. Click **Next**.
5. Select the authentication method (user or service).
6. Provide the appropriate configuration details.
7. Select **Create and continue**.
8. Review the integration details.
9. Select **Next**.
10. Share the integration with other users if needed.

After creating your MCP integration, a knowledge base is automatically created on completion of integration creation.

Manage knowledge bases

After setting up your MCP integration, you can create and manage knowledge bases from your MCP data sources.

Edit existing knowledge bases

You can modify your existing MCP knowledge bases:

1. In the Amazon Quick Suite console, choose **Knowledge bases**.
2. Select your MCP knowledge base from the list.
3. Choose **Actions**, then choose **Edit knowledge base**.
4. Update your configuration settings as needed and choose **Save**.

Create additional knowledge bases

You can create multiple knowledge bases from the same MCP integration:

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Choose your existing MCP integration from the list.
3. Choose **Actions**, then choose **Create knowledge base**.
4. Configure your knowledge base settings and choose **Create**.

For detailed information about knowledge base configuration options, see [the section called "Common configuration settings"](#).

Review integration

After you configure authentication, review the MCP integration capabilities:

1. The system connects to the MCP server and discovers available capabilities.
2. Review the list of available actions and tasks provided by the MCP server.
3. Verify data access capabilities if the MCP server provides data sources.
4. Confirm the integration configuration and capabilities.

Capability discovery

MCP integration automatically discovers and lists:

- Available tools and actions.
- Data sources and resources.
- Supported protocols and methods.
- Server metadata and capabilities.

Manage MCP integrations

After you create your MCP integration, you can manage it using these options:

- **Share integration** - Make the integration available to other users in your organization.
- **Review tools** - Review tools that are enabled.

Note

MCP integrations depend on the availability and configuration of the target MCP server. Changes to the server or authentication requirements may affect integration functionality.

Limitations

When using MCP integrations in Amazon Quick Suite, be aware of the following limitations:

- MCP operations have a fixed 60-second timeout. Operations that exceed this limit automatically fail with an HTTP 424 error.
- Custom HTTP headers are not supported in MCP operations. Only standard system headers are transmitted.
- Tool lists remain static after initial registration. You must manually refresh actions to detect server-side changes.
- Connector creation may fail if the Amazon Quick Suite callback URI is not allow-listed by third-party providers.
- Server connectivity issues result in immediate failure without retry attempts.

Microsoft OneDrive integration

With Microsoft OneDrive integration, you can perform actions on OneDrive files and folders. You can also create knowledge bases from OneDrive content including documents, spreadsheets, and presentations.

What you can do

OneDrive integration provides two types of capabilities to help you work with your OneDrive content.

Action connector

Create, update, delete, and manage OneDrive content through API calls.

Knowledge base

Connect to OneDrive to ask questions and get insights from your documents, spreadsheets, presentations, and other file types.

Before you begin

Make sure you have the following before you set up OneDrive integration.

- Microsoft OneDrive account.
- Appropriate OneDrive permissions for the content you want to access.
- For action connectors: Amazon Quick Suite Enterprise subscription
- For data access: Amazon Quick Suite Professional subscription –

As a one-time step from the admin, your Microsoft admin may need to grant organizational consent first before users can create a OneDrive integration to bring data. This is because Microsoft requires admin approval when using new applications (in our case the managed OAuth application), unless your user consent setting allows bypassing the admin approval. Admins can grant organization-wide consent by signing in and checking "Consent on behalf of your organization" during an integration creation for OneDrive.

Prepare Microsoft App Registration and authentication

Before setting up the integration in Amazon Quick Suite, prepare your Microsoft App Registration and authentication credentials. OneDrive integration supports different authentication methods depending on your integration type.

Action connector authentication setup

For action connectors, gather authentication credentials using one of these methods:

User authentication (OAuth)

Gather the following information from your Microsoft App Registration:

- **Base URL** - Microsoft Graph API base URL.
- **Client ID** - Microsoft 365 application client ID.
- **Client Secret** - Microsoft 365 application client secret.
- **Token URL** - Microsoft OAuth token endpoint.
- **Auth URL** - Microsoft OAuth authorization endpoint.
- **Redirect URL** - OAuth redirect URI, this is the Amazon Quick Suite URL.

Required OAuth scopes (Delegated):

- Files.Read.All
- Notes.Read.All
- User.Read
- Sites.Read.All
- offline_access
- AllSites.Read (sharepoint)

Service authentication (OAuth)

Gather the following information from your Microsoft App Registration:

- **Base URL** - Microsoft Graph API base URL.
- **Client ID** - Service application client ID.
- **Client Secret** - Service application client secret.
- **Token URL** - Microsoft OAuth token endpoint.

Required scope for token generation:

- `.default` - Default application permissions scope.

Data access authentication setup

Amazon Quick Suite supports creating OneDrive knowledge base integrations using three-legged OAuth (3LO). For this approach, Amazon Quick Suite registers a multi-tenant app with delegated permissions (*Sites.Read.All*) in their Microsoft Entra account. The delegated permissions model keeps risk low, despite being a multi-tenant application. Amazon Quick Suite can only access data through authenticated users' permissions. The security model ensures Quick is limited to what your users can access.

When an admin grants the organizational consent to allow users to bring data to Amazon Quick Suite, Azure automatically creates a Service Principal (Enterprise Application) in the customer tenant. You can disable or delete this Service Principal anytime from their Enterprise Applications, immediately revoking all access.

For data access integrations, prepare for Microsoft sign-in authentication. During the integration setup, you will need to:

1. Sign in with your Microsoft account that has access to the OneDrive content.
2. Grant the requested permissions to allow Amazon Quick Suite to access your OneDrive data.

3. Complete the authentication process.

Set up OneDrive integration

After preparing your Microsoft App Registration and authentication credentials, use the Integrations tab in the Amazon Quick Suite console to set up OneDrive integration. The setup process varies based on whether you want to perform actions, create knowledge bases, or both.

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Click the Add (plus "+") button.
3. Choose **Microsoft OneDrive** from the integration options.
4. Select your integration type:
 - **Bring data from Microsoft OneDrive** - Sets up data connections for knowledge base creation.
 - **Perform actions in Microsoft OneDrive** - Enables actions like creating, updating, or managing OneDrive files.
5. For knowledge base creation (Bring data from Microsoft OneDrive):
 - a. Select **Next**.
 - b. Complete authentication.
 - c. Select which files from your drive you want to add.
 - d. Add Name and Description.
 - e. Select **Create**.
6. For actions (Perform actions in Microsoft OneDrive):
 - a. Select **Next**.
 - b. Add Name and description.
 - c. Choose connection type:
 - **User authentication** - OAuth-based authentication for individual user access.
 - **Service authentication** - Service-to-service authentication for application access.
 - d. Fill in connection settings based on your selected authentication method (either user or service) using the credentials you prepared earlier.
 - e. Select **Create and continue**.

- f. Choose users to share the integration with.
- g. Click **Next**.

Manage knowledge bases

After setting up your OneDrive integration, you can create and manage knowledge bases from your OneDrive content.

Edit existing knowledge bases

You can modify your existing OneDrive knowledge bases:

1. In the Amazon Quick Suite console, choose **Knowledge bases**.
2. Select your OneDrive knowledge base from the list.
3. Choose **Actions**, then choose **Edit knowledge base**.
4. Update your configuration settings as needed and choose **Save**.

Create additional knowledge bases

You can create multiple knowledge bases from the same OneDrive integration:

1. In the Amazon Quick Suite console, choose **Integrations**, and then select the **Data** tab.
2. Choose your existing OneDrive integration from the list.
3. Choose the three-dot icon under **Actions**, then choose **Create knowledge base**.
4. Configure your knowledge base settings and choose **Create**.

For detailed information about knowledge base configuration options, see [the section called "Common configuration settings"](#).

You can create knowledge bases from these OneDrive content types.

- Microsoft Office documents (Word, Excel, PowerPoint).
- PDF files.
- Text files and rich text documents.
- Text documents with embedded images
- Audio, video, and other common document formats.

Note

Amazon Quick Suite doesn't sync ACLs from data sources. When you create a knowledge base in Amazon Quick Suite, by default only you can get insights from the knowledge base. For shared content, you can provide access to different users and groups by updating the knowledge base permissions.

Limitations

When using One Drive integrations in Amazon Quick Suite, be aware of the following limitations:

- File comments synchronization is not supported.

Microsoft Outlook action connector

With Microsoft Outlook action connector in Amazon Quick Suite, you can perform actions within Outlook systems, including managing emails, calendar events, contacts, and interacting with Microsoft Graph APIs. This action connector supports task execution only and requires Amazon Quick Suite Pro tier or higher.

What you can do

With Microsoft Outlook integration, you can perform actions within your Outlook systems through the Outlook API.

Action connector

Create, update, and manage emails, calendar events, contacts, and other messaging operations through the Outlook API.

Before you begin

Before you set up Microsoft Outlook action connector integration, make sure you have the following:

- Administrative access to Microsoft Azure Active Directory.
- Microsoft Outlook or Exchange Online account with appropriate permissions.
- Amazon Quick Suite Author or higher.

Step 1: Register an application in Azure

As an administrator, create a new OAuth 2.0 Microsoft Exchange application in the Microsoft Azure developer console with scoped permissions for performing actions in Amazon Quick Suite.

1. Sign in to the Azure portal and navigate to Azure Active Directory.
2. Choose **App registrations**, then choose **New registration**.
3. Enter a name for your application and configure the registration settings.
4. After registration, copy the **Application (client) ID** and **Directory (tenant) ID** from the application overview page.

Step 2: Configure API permissions

Add the required Microsoft Graph API permissions to your registered application.

1. In your Azure application, choose **API permissions**.
2. Choose **Add a permission** and select **Microsoft Graph**.
3. Add the following required permissions:
 - `Mail.ReadWrite` - Read and write access to user mail.
 - `Mail.Send` - Send mail as a user.
 - `Calendars.ReadWrite` - Read and write access to user calendars.
 - `User.Read.All` - Read all users' basic profiles.
 - `Contacts.Read` - Read user contacts.
 - `Place.Read.All` - Read all place information.
 - `MailboxSettings.Read` - Read user mailbox settings.
4. Choose **Grant admin consent** to approve the permissions.

Step 3: Generate a client secret

Create a client secret for your Azure application to authenticate with Amazon Quick Suite.

1. In your Azure application, navigate to **Manage > Certificates & secrets**.
2. Choose **New client secret**.
3. Enter a description and choose an expiration period.

4. Choose **Add**.
5. Copy the **Secret ID** value immediately. This value is only displayed once.

Step 4: Set up Microsoft Outlook action connector

After preparing your Azure application credentials, create the Microsoft Outlook action connector in Amazon Quick Suite.

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Click the Add (plus "+") button.
3. Choose **Microsoft Outlook** from the integration options.

Note

Microsoft Outlook integration supports action execution only. Data access and knowledge base creation are not available for Outlook systems.

4. Fill in the integration details:
 - **Name** - Enter a descriptive name for your Outlook action connector.
 - **Description** (Optional) - Describe the purpose of this action connector.
5. Choose your connection type:
 - **User authentication** - OAuth for individual user access to Outlook.
 - **Service authentication** - Service-to-service authentication for application access.
6. Fill in the connection settings based on your selected authentication method (either user or service) using the Azure application details you prepared earlier.
7. Select **Create and continue**.
8. Choose users to share the integration with.
9. Click **Next**.

Step 5: Associate action connector to automation groups

To use Microsoft Outlook actions in automations, you must associate the action connector with your automation groups.

1. Navigate to your automation group settings.
2. Associate the Microsoft Outlook action connector with the automation group that will use these actions.
3. Create a new automation for the automation group to access Microsoft Outlook actions in your workflows.

Manage Microsoft Outlook action connectors

You can perform these management tasks for your Microsoft Outlook action connectors once they have been created. Select them in the list of "Actions":

- **Edit action connector** - Update authentication settings or Outlook configuration.
- **Share action connector** - Make the action connector available to other users in your organization.
- **Delete action connector** - Remove the action connector and revoke authentication.

What happens next

After you complete the setup, your Microsoft Outlook action connector appears in the actions list. You can use it in Amazon Quick Suite automations and workflows to manage emails, calendar events, and contacts.

Microsoft SharePoint integration

With Microsoft SharePoint integration, you can perform actions on SharePoint lists and items. You can also create knowledge bases from SharePoint content including document libraries, sites, and pages.

What you can do

SharePoint users can ask questions about content stored in their SharePoint sites and document libraries. For example, users can inquire about project documents, team sites, or search for specific information across multiple libraries. The integration enables users to quickly access and understand information from their SharePoint content, regardless of location or type, while providing contextual details such as publication dates, modification history, and document ownership—all contributing to more efficient information discovery and better-informed decision making.

Before you begin

Make sure you have the following before you set up SharePoint integration.

- Microsoft SharePoint Online
- For action connectors: Amazon Quick Suite Enterprise subscription
- For data access: Amazon Quick Suite Professional subscription or higher – As a one-time step from the admin, your Microsoft admin may need to grant organizational consent first before users can create a SharePoint integration to bring data. This is because Microsoft requires admin approval when using new applications (in our case the managed OAuth application), unless your user consent setting allows bypassing the admin approval. Admins can grant organization-wide consent by signing in and checking "Consent on behalf of your organization" during an integration creation for SharePoint.

Prepare Microsoft App Registration and authentication

Before setting up the integration in Amazon Quick Suite, prepare your Microsoft App Registration and authentication credentials. SharePoint integration supports different authentication methods depending on your integration type and user tier.

Action connector authentication setup

For action connectors, prepare authentication credentials using one of these methods:

User authentication (3LO)

Gather the following information from your Microsoft Azure AD app registration:

- **Base URL** - Your SharePoint site URL.
- **Client ID** - Amazon application client ID.
- **Client Secret** - Amazon application client secret.
- **Token URL** - OAuth token endpoint.
- **Auth URL** - OAuth authorization endpoint.
- **Redirect URL** - OAuth redirect URI.

Required OAuth scopes:

- offline_access

- `Notes.Read.All`
- `Sites.Read.All`
- `AllSites.Read` (sharepoint)
- `User.Read`

Service authentication (API Key)

Gather the following information from your SharePoint administrator:

- **Base URL** - Your SharePoint site URL.
- **API Key** - SharePoint API access key.
- **Email** - Associated service account email.

Required scope for token generation:

- `.default` - Default application permissions scope.

Data access authentication setup

Amazon Quick Suite supports creating SharePoint knowledge base integrations using three-legged OAuth (3LO). For this approach, Amazon Quick Suite registers a multi-tenant app with delegated permissions (*Sites.Read.All*) in their Microsoft Entra account. The delegated permissions model keeps risk low, despite being a multi-tenant application. Quick can only access SharePoint through authenticated users' permissions. The security model ensures Quick is limited to what your users can access.

When an admin grants the organizational consent to allow users to bring data to Amazon Quick Suite, Azure automatically creates a Service Principal (Enterprise Application) in the customer tenant. You can disable or delete this Service Principal anytime from their Enterprise Applications, immediately revoking all access.

For data access integrations, prepare for Microsoft 3LO authentication. During the integration setup, you will need to:

1. Sign in with your Microsoft account that has access to the SharePoint content.
2. Grant the requested permissions to allow Amazon Quick Suite to access your SharePoint data.
3. Complete the authentication process.

Set up SharePoint integration

After preparing your Microsoft App Registration and authentication credentials, use the Integrations tab in the Amazon Quick Suite console to set up SharePoint integration. The setup process varies based on whether you want to perform actions, access data, or both.

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Choose **Microsoft SharePoint** from the integration options, and click the Add (plus "+") button.
3. Select your integration type:
 - **Bring data from Microsoft SharePoint** - Sets up data connections for knowledge base creation.
 - **Perform actions in Microsoft SharePoint** - Enables actions on SharePoint lists and items.
4. For knowledge base creation (Bring data from Microsoft SharePoint):
 - a. Select **Next**.
 - b. Fill in Name and complete authentication.
 - c. Select which pages/files/folders you want to add.
 - d. Fill in details like Name, Description, and content URLs.
 - e. Select **Create and continue**.
5. For actions (Perform actions in Microsoft SharePoint):
 - a. Select **Next**.
 - b. Add Name and description.
 - c. Choose connection type:
 - **User authentication** - 3LO-based authentication for individual user access.
 - **Service authentication** - API key-based authentication for service access.
 - d. Fill in connection settings based on your selected authentication method (either user or service) using the credentials you prepared earlier.
 - e. Select **Create and continue**.
 - f. Choose users to share the integration with.
 - g. Click **Next**.

After clicking **Create**, the data sync is started automatically.

Review available task actions

After setting up your SharePoint integration with action capabilities, review the available task actions to understand what operations you can perform. See [the section called “Available task actions”](#) for the complete list of available actions.

Configure knowledge bases

After setting up your SharePoint integration with data access capabilities, you can create additional knowledge bases. See [the section called “Manage knowledge bases”](#) for detailed configuration options.

Available task actions

SharePoint action connector provides 19 actions for managing SharePoint content.

SharePoint Actions

Action Name	Description	Type
Create Item	Create a new list item	Write
Update Item	Update an existing list item	Write
Delete Item	Delete a list item	Write
Get Item	Retrieve a specific list item	Read
List Items	List items in a SharePoint list	Read
Get List	Retrieve details of a SharePoint list	Read

Note

Additional actions are available for Excel file management, site operations, and advanced list management. Review the complete action list after you create your integration.

Manage knowledge bases

After setting up your SharePoint integration, you can create and manage knowledge bases from your SharePoint content.

Edit existing knowledge bases

You can modify your existing SharePoint knowledge bases:

1. In the Amazon Quick Suite console, choose **Knowledge bases**.
2. Select your SharePoint knowledge base from the list.
3. Choose the three-dot icon under **Actions**, then choose **Edit knowledge base**.
4. Update your configuration settings as needed and choose **Save**.

Create additional knowledge bases

You can create multiple knowledge bases from the same SharePoint integration:

1. In the Amazon Quick Suite console, choose **Integrations**, and then select the **Data** tab.
2. Choose your existing SharePoint integration from the list.
3. Choose the three-dot icon under **Actions**, then choose **Create knowledge base**.
4. Configure your knowledge base settings and choose **Create**.

For detailed information about knowledge base configuration options, see [the section called "Common configuration settings"](#).

Supported content types

You can create knowledge bases from these SharePoint content types.

- **Document libraries:** Word, Excel, PowerPoint, PDF, OneNote (.one)
- **Media files:** MP3, MP4, MOV, WMV
- **Site pages and wiki pages**

Manage SharePoint integrations

Note

Amazon Quick Suite doesn't sync ACLs from data sources. When you create a knowledge base in Amazon Quick Suite, by default only you can get insights from the knowledge base. For shared content, you can provide access to different users and groups by updating the knowledge base permissions.

Troubleshooting

Use this section to resolve common issues with SharePoint integration.

Authentication issues

Azure AD problems

Symptoms: Authentication fails with Azure AD errors, token refresh failures, or permission denied messages.

Causes and resolutions:

- **Incorrect app registration:** Verify the Azure AD app registration includes the required API permissions and OAuth scopes.
- **Expired client secret:** Check if the client secret has expired and generate a new one if needed.
- **Insufficient permissions:** Ensure the app registration has been granted admin consent for the required permissions.
- **Incorrect redirect URI:** Verify the redirect URI in Azure AD matches the one configured in Amazon Quick Suite.

Permission sync failures

Symptoms: Users cannot access content they should have permissions for, or see content they shouldn't access.

Causes and resolutions:

- **Permission propagation delay:** SharePoint permissions may take time to propagate. Wait 15-30 minutes and retry.

- **Nested group permissions:** Check if the user is part of nested security groups that may affect permissions.
- **Broken permission inheritance:** Verify that SharePoint items haven't broken permission inheritance unexpectedly.

Sync performance issues

Slow sync

Symptoms: Knowledge base sync takes longer than expected or appears to hang.

Causes and resolutions:

- **Large content volume:** Reduce the scope of content being synced by applying more restrictive filters.
- **Network connectivity:** Check network connection stability and bandwidth availability.
- **SharePoint throttling:** SharePoint may be throttling requests. Retry the sync during off-peak hours.

Content discovery problems

Symptoms: Expected SharePoint content is not appearing in the knowledge base.

Causes and resolutions:

- **Content filters too restrictive:** Review and adjust content filtering settings to include the missing content.
- **Unsupported content types:** Verify that the content type is supported by SharePoint integration.
- **Permission restrictions:** Ensure the integration has appropriate permissions to access the content.

Common error messages

AADSTS50020: User account from identity provider does not exist in tenant

Cause: The user account is not properly configured in Azure AD.

Resolution: Verify the user account exists in the correct Azure AD tenant and has appropriate licenses.

Access denied. You do not have permission to perform this action

Cause: Insufficient SharePoint permissions for the requested operation.

Resolution: Contact your SharePoint administrator to verify and grant appropriate permissions.

The request is throttled

Cause: SharePoint is limiting the number of requests due to high usage.

Resolution: Wait and retry the operation. Consider reducing the frequency of requests or syncing during off-peak hours.

Microsoft Teams integration

With Microsoft Teams integration, you can create actions that interact with Microsoft Teams for messaging, channel management, meeting operations, and team collaboration. This integration supports action capabilities for managing Teams communications and workflows.

What you can do

Microsoft Teams integration provides team collaboration capabilities to help you work with your Microsoft Teams environment.

Action integration

Execute Microsoft Teams actions such as sending messages, managing channels, creating meetings, and accessing team information through Microsoft Graph API calls.

Before you begin

Make sure you have the following before you set up Microsoft Teams integration.

- Microsoft 365 account with Teams access.
- Azure Active Directory application registration.
- Appropriate Teams permissions and roles.
- Author or higher.
- Administrative access to configure OAuth applications.

Set up Microsoft Teams integration

Follow these steps to create your Microsoft Teams integration.

1. In the console, choose **Integrations**.
2. Choose **Microsoft Teams** from the integration options, and click the Add (plus "+") button.
3. Fill in the integration details:
 - **Name** - Descriptive name for your Microsoft Teams integration.
 - **Description** (Optional) - Purpose of the integration.
4. Choose your connection type:
 - **User authentication** - OAuth-based authentication for individual user access.
 - **Service authentication** - Application permissions for automated access.
5. Fill in the connection settings based on your selected authentication method (either user or service):
 - a. For **User authentication (OAuth)**, configure the following fields:
 - **Client ID** - Azure AD application client ID.
 - **Client Secret** - Azure AD application client secret.
 - **Token URL** - Microsoft OAuth token endpoint.
 - **Auth URL** - Microsoft OAuth authorization endpoint.
 - **Redirect URL** - OAuth redirect URI.

Required OAuth scopes:

- `ChatMessage.Send` - Send chat messages.
- `ChannelMessage.Read.All` - Read channel messages.
- `ChannelMessage.Send` - Send channel messages.
- `Chat.ReadWrite` - Read and write chat information.
- `Team.ReadBasic.All` - Read basic team information.
- `TeamMember.ReadWrite.All` - Read and write team member information.
- `Channel.ReadBasic.All` - Read basic channel information.
- `ChannelMember.ReadWrite.All` - Read and write channel member information.

- `OnlineMeetingRecording.Read.All` - Read meeting recordings.
 - `OnlineMeetingTranscript.Read.All` - Read meeting transcripts.
 - `User.Read.All` - Read user profiles.
 - `offline_access` - Maintain access when user is offline.
- b. For **Service authentication (Application permissions)**, configure the following fields:
- **Client ID** - Service application client ID.
 - **Client Secret** - Service application client secret.
 - **Token URL** - Microsoft OAuth token endpoint.

Required scope for token generation:

- `.default` - Default application permissions scope.

Set the following scopes in your Azure AD application:

- `Chat.Send` - Send chat messages.
- `ChannelMessage.Read.All` - Read channel messages.
- `ChannelMessage.Send` - Send channel messages.
- `Chat.Read.All` - Read chat information.
- `Team.ReadBasic.All` - Read basic team information.
- `TeamMember.ReadWrite.All` - Read and write team member information.
- `Channel.ReadBasic.All` - Read basic channel information.
- `ChannelMember.ReadWrite.All` - Read and write channel member information.
- `OnlineMeetingRecording.Read.All` - Read meeting recordings.
- `OnlineMeetingTranscript.Read.All` - Read meeting transcripts.
- `User.Read.All` - Read user profiles.

6. Select **Create and continue**.
7. Choose users to share the integration with.
8. Click **Next**.

Available actions

Microsoft Teams integration provides the following action capabilities.

- Send messages to channels and chats.
- Create and manage team channels.
- Manage team and channel memberships.
- Read channel and chat messages.
- Access meeting recordings and transcripts.
- Retrieve team and user information.
- Manage team settings and configurations.

Note

The specific actions available depend on the OAuth scopes configured and the permissions granted during authentication.

Manage Microsoft Teams integrations

After you create your Microsoft Teams integration, you can manage it through several options.

- **Edit integration** - Update OAuth settings or Teams configuration.
- **Share integration** - Make the integration available to other users.
- **Monitor usage** - View integration activity and API usage.
- **Review actions** - See available Microsoft Teams actions.
- **Delete integration** - Remove the integration and revoke Teams access.

OpenAPI Specification integration

With OpenAPI Specification integration, you can create custom integrations based on OpenAPI schemas. This allows you to connect to any API that provides an OpenAPI specification. This integration supports action execution only and requires Amazon Quick Suite Pro tier or higher.

What you can do

OpenAPI Specification integration provides schema-based connectivity to help you work with custom APIs.

Action connector

Perform actions based on OpenAPI specifications. Execute API calls, manage resources, and interact with custom services through dynamically generated actions based on the provided schema.

Schema-based configuration

Import OpenAPI specifications to automatically generate available actions and actions. Support for JSON schema validation and parameter mapping.

Before you begin

Before you set up OpenAPI Specification integration, make sure you have the following:

- OpenAPI specification (version 3.0.0 or higher) for your target API.
- API authentication credentials (API key, OAuth, or other).
- Amazon Quick Suite Author or higher.
- API documentation and endpoint access.

Prepare OpenAPI specification and authentication

Before setting up the integration in Amazon Quick Suite, prepare your OpenAPI specification and authentication credentials. Your OpenAPI specification must meet specific requirements for successful integration.

Basic requirements

Your OpenAPI specification must meet these basic requirements.

- **OpenAPI version** - Version 3.0.0 or higher is required.
- **File format** - JSON format only (application/json).
- **File size limit** - Maximum 1MB for the entire OpenAPI specification.
- **Operation limit** - Minimum 1 and maximum 100 API operations per connector.

Required schema fields

Your OpenAPI specification must include these required sections.

openapi

The OpenAPI version being used (must be "3.0.0" or higher).

info

Service information including title, description, and version.

servers

Base URL and server information for API connectivity.

paths

API endpoint definitions with HTTP methods and operations.

Supported authentication schemes

Prepare authentication credentials based on the supported authentication methods in OpenAPI specifications.

OAuth 2.0

Supports both Authorization Code Grant and Client Credentials Grant flows. Requires authorizationUrl, tokenUrl, and scopes definitions in the security scheme.

API Key

API key authentication passed in query parameters or headers.

No authentication

Default option when no security schemes are defined in the specification.

Note

Unsupported authentication schemes in the OpenAPI specification will be ignored, and the connector will default to no authentication.

Set up OpenAPI Specification integration

After preparing your OpenAPI specification and authentication credentials, follow these steps to create your OpenAPI Specification integration.

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Click **Add** (plus "+" button).
3. On the Add Schema page, select **Import Schema** and choose a schema to import.
4. Select **Next**.
5. Fill in the integration details including name and description.
6. Review the available actions generated from your OpenAPI specification.
7. Select **Create and continue**.
8. Choose users to share the integration with.
9. Click **Next**.

Expected results

After successful setup, your OpenAPI Specification integration appears in the integrations list with dynamically generated actions based on your schema. The integration is available for use in Amazon Quick Suite workflows, automations, and AI agents, with tasks corresponding to the endpoints defined in your OpenAPI specification.

Format and pattern requirements

Your OpenAPI specification must follow these format requirements.

- **Path patterns** - Must follow the pattern: `^/[^\/*]*(/[^\/*]*)*$` and begin with a forward slash (/).
- **Operation IDs** - Must follow the pattern: `^[a-zA-Z0-9_]{1,256}$`.
- **Descriptions** - Required for all operations and parameters, maximum 5000 characters.
- **Content type** - Only application/json is supported for request and response bodies.

Unsupported features

The following OpenAPI features are not supported and will cause validation errors.

- **Array types** - Schemas with array types (for example, {"type": "array", "items": {"string"}}) are not supported.
- **Composition keywords** - allOf, oneOf, anyOf, and not keywords are not supported.
- **Circular references** - Circular or cyclic references (\$ref inside of \$ref) are not supported.
- **Complex nested structures** - Avoid deeply nested object structures when possible.

Schema best practices

Follow these best practices when creating your OpenAPI specification.

Write effective descriptions

Use these guidelines to write clear and helpful descriptions.

- **Operation descriptions** - Describe what the operation does, when to use it, and how it behaves.
- **Parameter descriptions** - Concisely explain the parameter's purpose and impact on operation behavior.
- **Self-contained content** - Ensure descriptions provide sufficient guidance without external references.
- **Dependency clarity** - Make dependencies between operations explicit in descriptions.
- **Operation references** - Use operationId when referring to other operations, not API paths.

Parameter structure recommendations

Structure your parameters for optimal usability.

- **Flatten complex objects** - Instead of nested objects, use separate parameters (for example, start_date, start_time, end_date, end_time).
- **Use standard formats** - Use standard format values like "date-time" or "date" for ISO-8601 dates and times.
- **Clear parameter names** - Use descriptive parameter names that clearly indicate their purpose.
- **Required field marking** - Properly mark parameters as required or optional based on API behavior.

Supported extensions

We support custom OpenAPI extensions to enhance the user experience.

x-amzn-form-display-name

Use at the parameter level to override the default field name displayed in confirmation forms. Currently supported for request body parameters only.

```
"x-amzn-form-display-name": "User Name"
```

x-amzn-operation-type

Define whether an operation should be treated as "read" or "write". By default, GET methods are "read" operations and all other HTTP methods are "write" operations.

```
"x-amzn-operation-type": "read"
```

Schema validation and error handling

When you upload an OpenAPI specification, we perform comprehensive validation.

- **File size validation** - Ensures the specification doesn't exceed 1MB.
- **Operation count validation** - Verifies between 1-100 operations are defined.
- **Schema structure validation** - Checks for required fields and proper formatting.
- **Security scheme validation** - Validates supported authentication methods.
- **Content type validation** - Ensures only application/json is used.

Validation errors appear below the schema editor and must be resolved before the integration can be created. Common validation issues include:

- Missing required fields (openapi, info, servers, paths).
- Invalid path patterns or operation IDs.
- Unsupported schema features (arrays, composition keywords).
- Missing or overly long descriptions.
- Circular references in schema definitions.

Manage OpenAPI Specification integration

After you create your OpenAPI Specification integration, you can manage it through several options.

Share integration

You can share OpenAPI Specification action connectors with other users in your organization.

1. From the OpenAPI integration details page, choose **Share**.
2. Configure sharing options:
 - **Share with specific users** - Enter user names or email addresses.
 - **Share with organization** - Make available to all users in your organization.
3. Set permission levels for shared access.
4. Choose **Share integration** to apply the sharing settings.

Delete integration

Follow these steps to permanently remove your OpenAPI integration.

1. From the OpenAPI integration details page, choose **Delete**.
2. Review the deletion impact, including any workflows or automations using this integration.
3. Type the integration name to confirm deletion.
4. Choose **Delete integration** to permanently remove it.

Troubleshoot OpenAPI Specification integration

Use these troubleshooting tips to resolve common OpenAPI Specification integration issues.

Schema validation errors

Ensure your OpenAPI specification follows the correct format and includes all required fields.
Use an OpenAPI validator to check your schema before importing.

No tasks generated

Verify that your OpenAPI specification includes path definitions with HTTP methods and operation IDs. Actions are generated based on the operations defined in your schema.

Authentication failures

Check that the authentication scheme defined in your OpenAPI specification matches the actual API requirements and that credentials are properly configured.

Action execution failures

Review the action parameters and ensure they match the parameter definitions in your OpenAPI specification, including required fields and data types.

PagerDuty integration

Connect Amazon Quick Suite to your PagerDuty system to manage incidents, alerts, schedules, and on-call rotations. You can perform these actions without leaving your Amazon Quick Suite environment. This integration requires Amazon Quick Suite Pro tier or higher.

What you can do

With PagerDuty integration, you can perform actions within your PagerDuty systems through the PagerDuty API.

Action connector

Create, update, and manage incidents, alerts, schedules, and escalation policies through the PagerDuty API.

Set up PagerDuty integration

Follow these steps to connect Amazon Quick Suite to your PagerDuty system.

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Click the Add (plus "+") button.
3. Choose **PagerDuty** from the integration options.
4. Fill in the integration details:
 - **Name** - Enter a descriptive name for your PagerDuty integration.
 - **Description** (Optional) - Describe the purpose of this integration.
5. Choose your connection type:

- **User authentication** - OAuth for individual user access to PagerDuty.
 - **Service authentication** - API key-based authentication for service-to-service access.
6. Fill in the connection settings based on your selected authentication method (either user or service).
 7. Select **Create and continue**.
 8. Choose users to share the integration with.
 9. Click **Next**.

What happens next

After you complete the setup, your PagerDuty integration appears in the integrations list. You can use it in Amazon Quick Suite workflows, automations, and AI agents.

Manage PagerDuty integrations

You can perform these management tasks for your PagerDuty integrations:

- **Edit integration** - Update authentication settings or PagerDuty configuration.
- **Share integration** - Make the integration available to other users in your organization.
- **Delete integration** - Remove the integration and revoke authentication.

REST API Connection integration

With REST API Connection integration in Amazon Quick Suite, you can perform actions with custom REST APIs and web services. This integration supports action execution only.

What you can do

With REST API Connection integration, you can perform actions with custom REST APIs and web services through the action connector.

Action connector

Perform HTTP requests, retrieve data, and interact with APIs using flexible authentication options.

Note

REST API Connection integration doesn't support data access or knowledge base creation. It's designed specifically for task execution and API interactions with custom web services.

Before you begin

Before you set up REST API integration, make sure you have the following:

- REST API endpoint with appropriate access permissions.
- API authentication credentials (OAuth, API key, or other).
- API documentation for the target web service.

Prepare API endpoint and authentication

Before configuring the integration in Amazon Quick Suite, prepare your REST API endpoint and authentication credentials. REST API Connection integration supports multiple authentication methods. Choose the method that matches your API requirements:

User authentication (OAuth)

Gather the following information from your API provider:

- **Base URL** - REST API base URL.
- **Client ID** - OAuth application client ID.
- **Client Secret** - OAuth application client secret.
- **Token URL** - OAuth token endpoint.
- **Auth URL** - OAuth authorization endpoint.
- **Redirect URL** - OAuth redirect URI.

Service authentication (API Key)

Gather the following information from your API provider:

- **API Key** - REST API access key or token.
- **Base URL** - REST API base URL.
- **Email** - Associated user account email (if required).

Custom headers and parameters

You can use custom headers and parameters for flexible authentication and API interaction:

- Custom authentication headers.
- API version headers.
- Content-Type specifications.
- Custom query parameters.

Set up REST API integration

After you have prepared your API endpoint and authentication credentials, follow these steps to create your REST API integration:

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Choose **REST API Connection** from the integration options, and click the Add (plus "+") button.
3. Fill in the integration details:
 - **Name** - Descriptive name for your REST API integration.
 - **Description** (Optional) - Purpose of the integration.
4. Choose your connection type:
 - **User authentication** - OAuth-based authentication for individual user access.
 - **Service authentication** - API key-based authentication for service access.
5. Fill in the connection settings based on your selected authentication method (either user or service).
6. Select **Next**.
7. Review the actions that are available.
8. Select **Create and continue**.

Available task actions

After you create your REST API integration, you can review the available actions for interacting with the REST API. Common REST API actions include:

- HTTP GET requests for data retrieval.

- HTTP POST requests for data creation.
- HTTP PUT/PATCH requests for data updates.
- HTTP DELETE requests for data removal.
- Custom endpoint interactions.
- JSON and XML data processing.
- Query parameter and header management.

Note

The specific actions available depend on the REST API endpoints and the authentication permissions configured for your integration.

API configuration options

You can configure various aspects of your REST API integration to match your specific requirements.

Endpoint configuration

Configure these endpoint settings:

- Base URL and endpoint paths.
- HTTP method specifications.
- Request and response format handling.
- Error handling and retry logic.

Data handling

Configure how your integration processes different data formats:

- JSON request and response processing.
- XML data transformation.
- Form data and multipart uploads.
- Binary data handling.

Manage REST API integrations

After you create your REST API integration, you can manage it using these options:

- **Edit integration** - Update authentication settings, base URL, or API configuration.
- **Share integration** - Make the integration available to other users in your organization.
- **Monitor usage** - View integration activity and API call metrics.
- **Review actions** - See the complete list of available REST API actions.
- **Test endpoints** - Validate API connectivity and authentication.
- **Delete integration** - Remove the integration and revoke associated authentication.

Important

REST API integrations depend on the availability and configuration of the target web service. Changes to the API or authentication requirements may affect integration functionality.

Salesforce integration

With Salesforce action connector in Amazon Quick Suite, you can perform actions within Salesforce organizations, including managing records, querying data, and interacting with Salesforce APIs. This action connector supports task execution only and requires Amazon Quick Suite Pro tier or higher.

What you can do

With Salesforce integration, you can perform actions within your Salesforce organizations through the action connector.

Action connector

Create, update, and query Salesforce objects such as leads, accounts, contacts, and opportunities.

 Note

Salesforce integration doesn't support data access or knowledge base creation. It's designed specifically for task execution and API interactions with Salesforce objects.

Before you begin

Before you set up Salesforce integration, make sure you have the following:

- Salesforce organization with appropriate permissions.
- Salesforce connected app or API access credentials.
- Amazon Quick Suite Author or higher.
- Administrative access to configure OAuth applications (if using user authentication).

Step 1: Set up Salesforce connected app

Create a connected app in Salesforce to enable OAuth authentication with Amazon Quick Suite.

1. Sign in to your Salesforce account and navigate to the Setup page using the Setup icon in the top right.
2. In the Quick Find bar, enter **Apps** and select **App Manager**.
3. Choose **New Connected App**.
4. Choose **Create a connected app**.
5. In the Basic Information section, enter the following required information:
 - **Connected App Name** - A descriptive name for your connected app.
 - **API Name** - A unique API name for your application.
 - **Contact Email** - Your contact email address.
6. In the OAuth Settings section, select the following checkboxes:
 - **Enable OAuth Settings**
 - **Require Secret for Web Server Flow**
 - **Require Secret for Refresh Token Flow**
 - **Enable Client Credentials Flow**

- **Enable Authorization Code and Credential Flow**
 - **Enable Token Exchange Flow**
 - **Require Secret for Token Exchange Flow**
7. Add the following required OAuth scopes:
 - `api` - Access Salesforce APIs
 - `refresh_token` - Maintain access when user is offline
 - `offline_access` - Perform requests at any time
 - `full` - Full access to all data
 - `web` - Web-based access
 - `visualforce` - Access Visualforce pages
 - `custom_permissions` - Access custom permissions
 - `chatter_api` - Access Chatter API
 - `wave_api` - Access Analytics API
 - `eclair_api` - Access Einstein Analytics API
 - `pardot_api` - Access Pardot API
 - `id` - Access identity information
 - `email` - Access email address
 - `profile` - Access basic profile information
 - `address` - Access address information
 - `phone` - Access phone number
 - `open_id` - Access OpenID Connect
 8. Enter the callback URL in the format: `<quicksuite-url>/sn/oauthcallback`
 9. Choose **Save**.

Step 2: Configure consumer details and execution user

Configure the consumer details and set up an execution user for the client credentials flow.

1. From the Manage Connected Apps page, choose **Manage Consumer Details**. You might need to verify your identity.
2. Copy the **Consumer Key (Client ID)** and **Consumer Secret (Client Secret)**.

3. Choose **Apply**.
4. Choose **Initial Access Token**, then choose **OK**.
5. Configure the execution user:
 - a. From the connected app detail page, choose **Edit** under the Action column.
 - b. Under OAuth Policies > Refresh Token Policy, select **Immediately expire refresh token**.
 - c. Under Client Credentials Flow, for Run As, choose the user to assign the client credentials flow.
 - d. Choose **Save**.

Step 3: Set up Salesforce action connector

After preparing your Salesforce connected app credentials, create the Salesforce action connector in Amazon Quick Suite.

Salesforce integration supports action execution only - data access and knowledge base creation are not available for Salesforce systems.

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Choose **Salesforce** from the integration options, and click the Add (plus "+") button.
3. Fill in the integration details:
 - **Name** - Descriptive name for your Salesforce action connector.
 - **Description** (Optional) - Purpose of the action connector.
4. Choose your connection type:
 - **User authentication** - OAuth-based authentication for individual user access.
 - **Service authentication** - API key-based authentication for service access.
5. Fill in the connection settings based on your selected authentication method (either user or service):
 - a. For **User authentication (OAuth)**, configure the following fields:
 - **Base URL** - Salesforce instance URL (for example, <https://your-domain.salesforce.com>).
 - **Client ID** - Salesforce connected app consumer key.
 - **Client Secret** - Salesforce connected app consumer secret.

- **Token URL** - Salesforce OAuth token endpoint.
 - **Auth URL** - Salesforce OAuth authorization endpoint.
 - **Redirect URL** - OAuth redirect URI configured in your connected app.
- b. For **Service authentication (API Key)**, configure the following fields:
- **API Key** - Salesforce API access token or session ID.
 - **Base URL** - Salesforce instance URL.
 - **Email** - Associated Salesforce user email.
6. Select **Create and continue**.
 7. Choose users to share the integration with.
 8. Click **Next**.

Step 4: Associate action connector to automation groups

To use Salesforce actions in automations, you must associate the action connector with your automation groups.

1. Navigate to your automation group settings.
2. Associate the Salesforce action connector with the automation group that will use these actions.
3. Create a new automation for the automation group to access Salesforce actions in your workflows.

Available task actions

After you create your Salesforce integration, you can review the available actions for interacting with Salesforce objects. Common Salesforce actions include:

- Create, read, update, and delete (CRUD) operations on standard and custom objects.
- Query Salesforce data using SOQL (Salesforce Object Query Language).
- Manage leads, accounts, contacts, and opportunities.
- Execute Apex methods and custom logic.
- Manage cases, tasks, and activities.
- Access reports and dashboards.

Share integrations

You can share Salesforce action connectors with other users in your organization. Follow these steps:

1. After you create the integration, choose **Share integration**.
2. Select users or groups to share the integration with.
3. Set appropriate permissions for shared access.
4. Confirm sharing settings.

Shared users can use the Salesforce integration to perform actions within the connected Salesforce organization, subject to the permissions configured in the original authentication setup.

Manage Salesforce action connectors

After you create your Salesforce action connector, you can manage it using these options:

- **Edit action connector** - Update authentication settings or Salesforce instance configuration.
- **Share action connector** - Make the action connector available to other users in your organization.
- **Monitor usage** - View action connector activity and API usage metrics.
- **Review actions** - See the complete list of available Salesforce actions.
- **Delete action connector** - Remove the action connector and revoke associated authentication.

SAP workload integrations

With SAP workload integrations, you can perform actions within various SAP systems. Manage business data, inventory, materials, and business processes. These integrations support action execution only and require Amazon Quick Suite Pro tier or higher.

What you can do

SAP workload integrations provide enterprise-level connectivity to help you work with your SAP systems.

Action connector

Perform actions within SAP systems. Create, update, and manage business data, inventory records, material information, and other enterprise operations through SAP APIs.

Multiple SAP modules

Support for five different integration types: Bill of Materials, Business Partner, Material Stock, Physical Inventory Documents, and Product Master.

Before you begin

Before you set up SAP workload integrations, make sure you have the following:

- SAP system with appropriate modules installed and configured.
- SAP user account with necessary permissions and authorizations.
- Amazon Quick Suite Author or higher.
- Network connectivity between Amazon Quick Suite and your SAP system.

Prepare SAP system configuration and authentication

Before setting up the integration in Amazon Quick Suite, prepare your SAP system configuration and user authentication. SAP workload integrations support multiple authentication methods and require proper system setup.

Authentication methods

SAP workload integrations support two authentication methods:

OAuth 2.0 (Recommended)

Secure authentication method for automated workflows. Requires OAuth configuration in your SAP system.

Required parameters:

- **Client ID** - SAP OAuth Client ID
- **Client Secret** - SAP OAuth Client Secret
- **Token URL** - OAuth token endpoint (e.g., `https://hostname:port/sap/bc/sec/oauth2/token?sap-client=100`)

- **Domain URL** - SAP system API endpoint (e.g., `https://hostname:port/sap/opu/odata/sap/API_BUSINESS_PARTNER`)

Basic Authentication

Username and password authentication for direct SAP system access.

Required parameters:

- **Username** - SAP system username
- **Password** - SAP system password
- **Domain URL** - SAP system API endpoint

SAP system configuration requirements

Before using SAP workload integrations, ensure your SAP system is properly configured:

OAuth 2.0 setup

For OAuth 2.0 authentication:

- Configure the OAuth authorization server with appropriate scopes
- Verify OAuth is enabled for the service using transaction code `/IWFND/MAINT_SERVICE`
- Refer to SAP documentation for detailed OAuth 2.0 configuration: [OAuth 2.0 Configuration Guide](#)

SAP API activation

Ensure the required SAP API services are active:

- Activate the specific API service for your chosen SAP connector
- Verify API service status in your SAP system
- Refer to SAP documentation for API activation: [SAP API Service Activation Guide](#)

Available SAP workload integration types

Choose the SAP integration type that matches your business needs and ensure the corresponding SAP module is available in your system.

SAP Bill of Materials

Manage bill of materials data, including component lists, quantities, and manufacturing specifications.

Required scope: ZAPI_BILL_OF_MATERIAL_SRV_0002

SAP Business Partner

Handle business partner information, including customer and vendor data, contact details, and relationship management.

Required scope: ZAPI_BUSINESS_PARTNER_0001

SAP Material Stock

Access and manage material stock levels, inventory movements, and warehouse information.

Required scope: ZAPI_MATERIAL_STOCK_SRV_0001

SAP Physical Inventory Documents

Create and manage physical inventory documents, stock counts, and inventory reconciliation processes.

Required scope: ZAPI_MATERIAL_STOCK_SRV_0001

SAP Product Master

Maintain product master data, including material specifications, classifications, and product hierarchies.

Required scope: ZAPI_PRODUCT_SRV_0001

SAP workload authentication setup

Prepare SAP user authentication credentials and ensure proper system access permissions.

- **SAP system URL** - Obtain the base URL or server address for your SAP system.
- **User credentials** - Create or identify a SAP user account with appropriate permissions.
- **System permissions** - Ensure the user account has the necessary SAP authorizations and transaction codes for the specific SAP modules you plan to integrate.
- **Network access** - Verify that your SAP system is accessible from external connections and that firewall rules allow the integration.

Set up SAP integrations

After preparing your SAP system configuration and authentication credentials, the setup process is similar for all SAP integration types. Choose between OAuth 2.0 or Basic Authentication based on your security requirements.

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Select one of the available SAP connectors:
 - **SAP Bill of Materials**
 - **SAP Business Partner**
 - **SAP Material Stock**
 - **SAP Physical Inventory Documents**
 - **SAP Product Master**
3. Click **Add** (plus "+" button).
4. Fill in the name and description for your SAP integration.
5. Choose the connection type for your integration.
6. Configure authentication using one of the supported methods:

- a. **OAuth 2.0 configuration**

For OAuth 2.0 authentication, provide:

- **Client ID** - Your SAP OAuth Client ID
- **Client Secret** - Your SAP OAuth Client Secret
- **Token URL** - OAuth token endpoint URL
- **Domain URL** - SAP system API endpoint URL

- b. **Basic Authentication configuration**

For Basic Authentication, provide:

- **Username** - Your SAP system username
- **Password** - Your SAP system password
- **Domain URL** - SAP system API endpoint URL

7. Select **Create and continue**.
8. Choose users to share the integration with.

9. Click **Next**.

Expected results

After successful setup, your SAP workload integration appears in the integrations list and is available for use in Amazon Quick Suite workflows, automations, and AI agents. You can perform SAP-specific actions directly from Amazon Quick Suite using the configured authentication credentials.

Available operations by connector type

Each SAP workload connector provides specific operations tailored to its business domain. Review the available operations for your chosen connector type.

SAP Bill of Materials operations

Available operations for managing bill of materials data:

- **Get Material BOM Item** - Retrieves bill of material details for specified materials

Operation ID: getMaterialBOMItem

Endpoint: GET /MaterialBOMItem

SAP Business Partner operations

Available operations for managing business partner data:

- **Get Business Partner** - Retrieves business partner general data

Operation ID: getBusinessPartner

- **Get Business Partner Addresses** - Retrieves business partner address data

Operation ID: getBusinessPartnerAddress

- **Get Business Partner Roles** - Retrieves business partner role data

Operation ID: getBusinessPartnerRole

- **Get Business Partner by ID** - Retrieves business partner data by business partner number

Operation ID: getBusinessPartnerByID

Required parameter: BusinessPartner (string, max 10 characters)

- **Get Business Partner Role by ID** - Retrieves business partner role data using key fields

Operation ID: getBusinessPartnerRoleByID

Required parameters: BusinessPartner (string, max 10 characters), BusinessPartnerRole (string, max 6 characters)

SAP Material Stock operations

Available operations for managing material stock data:

- **Get Material Stock In Account** - Retrieves material stock information posted in the account model

Operation ID: getMaterialStockInAccount

SAP Physical Inventory Documents operations

Available operations for managing physical inventory documents:

- **Get PhysInventory Doc Item** - Reads information on physical inventory items

Operation ID: getPhysInventoryDocItem

SAP Product Master operations

Available operations for managing product master data:

- **Get Product Master Item** - Returns product master records

Operation ID: getProductMaster

- **Get Plant Data By Material** - Returns plant data of product master record

Operation ID: getPlantDataByMaterial

Required parameter: Product (string, max 40 characters)

- **Get Supply Planning Data By Material** - Returns supply planning data by product number and plant

Operation ID: getSupplyPlanningDataByMaterial

Required parameters: Product (string, max 40 characters), Plant (string, max 4 characters)

Query parameters

SAP connectors support standard query parameters to filter, sort, and format API responses. Use these parameters to optimize data retrieval and processing.

Supported Query Parameters

#	Parameter	Description	Type
1	\$top	Limits the number of returned items	integer
2	\$skip	Skips the specified number of items	integer
3	\$filter	Filters results based on specified criteria	string
4	\$orderby	Orders results by specified fields	array
5	\$select	Selects specific properties to return	array
6	\$expand	Expands related entities	array
7	\$inlinecount	Includes count of items in response	string

Manage SAP workload integrations

After you create your SAP workload integration, you can manage it through several options.

Edit integration settings

Follow these steps to modify your SAP workload integration settings.

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Choose your SAP workload integration type from the integration grid.
3. Select your integration from the list and choose **Edit**.
4. Modify integration settings:

- Update authentication credentials (username and password).
 - Change SAP system URL or connection settings.
 - Modify integration name or description.
5. Choose **Save changes** to apply your modifications.

Share integration

You can share SAP workload action connectors with other users in your organization.

1. From the SAP integration details page, choose **Share**.
2. Configure sharing options:
 - **Share with specific users** - Enter user names or email addresses.
 - **Share with organization** - Make available to all users in your organization.
3. Set permission levels for shared access.
4. Choose **Share integration** to apply the sharing settings.

Delete integration

Follow these steps to permanently remove your SAP integration.

1. From the SAP workload integration details page, choose **Delete**.
2. Review the deletion impact, including any workflows or automations using this integration.
3. Type the integration name to confirm deletion.
4. Choose **Delete integration** to permanently remove it.

Troubleshoot SAP workload integrations

Use these troubleshooting tips to resolve common SAP workload integration issues.

Authentication issues

OAuth 2.0 authentication failures

Symptoms: Token generation fails, invalid client credentials, or OAuth scope errors.

Resolution:

- Verify OAuth client ID and client secret are correct
- Check that OAuth is properly configured in SAP using transaction /IWFND/MAINT_SERVICE
- Ensure the required scopes are properly configured for your SAP workload connector type
- Verify the token URL format matches your SAP system configuration

Basic authentication failures

Symptoms: Login failures, invalid credentials, or access denied errors.

Resolution:

- Verify SAP username and password are correct
- Check that the user account has necessary SAP authorizations
- Ensure the user account is not locked or expired
- Verify the domain URL is accessible and correctly formatted

SAP system configuration issues

API service not activated

Symptoms: Service unavailable errors, API endpoint not found, or HTTP 404 responses.

Resolution:

- Verify the required SAP API service is activated in your system
- Check API service status using SAP transaction codes
- Ensure the API service corresponds to your chosen SAP connector type
- Contact your SAP administrator to activate missing API services

Connection timeouts

Symptoms: Request timeouts, network connectivity errors, or slow response times.

Resolution:

- Check that your SAP system URL is correct and accessible
- Verify network connectivity allows connections to SAP system
- Ensure firewall rules permit the integration traffic

- Check SAP system performance and availability

Permission and authorization errors

Insufficient SAP authorizations

Symptoms: Access denied errors, missing authorization messages, or restricted operation failures.

Resolution:

- Ensure the authenticated user has required SAP authorizations for the specific module
- Verify the user has access to necessary transaction codes
- Check that the user account has appropriate role assignments
- Contact your SAP administrator to grant missing permissions

Scope permission errors

Symptoms: OAuth scope errors, insufficient permissions for API operations, or restricted access messages.

Resolution:

- Verify the OAuth configuration includes the required scope for your connector type
- Check that the scope permissions are properly granted in SAP system
- Ensure the OAuth client has been granted the necessary API access rights

Data format and parameter errors

Invalid parameter formats

Symptoms: Data validation errors, invalid field length messages, or parameter format exceptions.

Resolution:

- Review action parameters and ensure they match expected SAP data formats
- Check field lengths match SAP system requirements (e.g., BusinessPartner max 10 characters)
- Verify data types are correct for the specific SAP module
- Ensure all required fields are provided for the operation

Query parameter errors

Symptoms: Query syntax errors, unsupported parameter messages, or malformed request errors.

Resolution:

- Verify parameters use correct syntax (e.g., `$filter`, `$top`, `$skip`)
- Check that the parameter values are properly formatted
- Ensure the SAP API supports the specific parameters being used
- Refer to SAP API documentation for supported query options

SAP system availability issues

SAP system unavailability

Symptoms: Connection refused errors, system not responding, or service unavailable messages.

Resolution:

- Check SAP system status and availability with your SAP administrator
- Verify if there are scheduled maintenance windows affecting the system
- Check for any SAP system alerts or known issues
- Retry the operation after confirming system availability

ServiceNow integration

With ServiceNow integration in Amazon Quick Suite, you can perform actions within ServiceNow instances, including managing incidents, requests, and other ServiceNow records. This integration supports action execution only and requires Amazon Quick Suite Pro tier or higher.

What you can do

With ServiceNow integration, you can perform actions within your ServiceNow instances through the action connector.

Action connector

Create, update, and query ServiceNow records such as incidents, requests, changes, and custom tables.

Note

ServiceNow integration doesn't support data access or knowledge base creation. It's designed specifically for action execution and API interactions with ServiceNow instances.

Before you begin

Before you set up ServiceNow integration, make sure you have the following:

- ServiceNow instance with appropriate permissions.
- ServiceNow user account or service account credentials.
- Amazon Quick Suite Author or higher.
- Administrative access to configure OAuth applications (if using user authentication).

Set up ServiceNow integration

Follow these steps to create your ServiceNow integration:

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Choose **ServiceNow** from the integration options, click the Add (plus "+") button.
3. Fill in the integration details:
 - **Name** - Descriptive name for your ServiceNow integration.
 - **Description** (Optional) - Purpose of the integration.
4. Choose your connection type:
 - **User authentication** - OAuth-based authentication for individual user access.
 - **Service authentication** - API key-based authentication for service access.
5. Fill in the connection settings based on your selected authentication method (either user or service):
 - a. For **User authentication (OAuth)**, configure the following fields:
 - **Base URL** - ServiceNow instance URL (for example, <https://your-instance.service-now.com>).

- **Client ID** - ServiceNow OAuth application client ID.
- **Client Secret** - ServiceNow OAuth application client secret.
- **Token URL** - ServiceNow OAuth token endpoint.
- **Auth URL** - ServiceNow OAuth authorization endpoint.
- **Redirect URL** - OAuth redirect URI.

ServiceNow OAuth authentication doesn't require specific OAuth scopes - access is controlled through ServiceNow's role-based permissions system.

b. For **Service authentication (API Key)**, configure the following fields:

- **API Key** - ServiceNow API access token.
- **Base URL** - ServiceNow instance URL.
- **Email** - Associated ServiceNow user email.

6. Select **Create and continue**.
7. Choose users to share the integration with.
8. Click **Next**.

Available actions

After you create your ServiceNow integration, you can review the available actions for interacting with ServiceNow instances. Common ServiceNow actions include:

- Create, read, update, and delete (CRUD) operations on ServiceNow tables.
- Manage incidents, requests, and change records.
- Query ServiceNow data using REST API.
- Execute ServiceNow workflows and business rules.
- Manage user accounts and group memberships.
- Access ServiceNow reports and analytics.

Note

The specific actions available depend on the permissions configured in your ServiceNow instance and the authentication method used.

Manage ServiceNow integrations

After you create your ServiceNow integration, you can manage it using these options:

- **Edit integration** - Update authentication settings or ServiceNow instance configuration.
- **Share integration** - Make the integration available to other users in your organization.
- **Monitor usage** - View integration activity and API usage metrics.
- **Review actions** - See the complete list of available ServiceNow actions.
- **Delete integration** - Remove the integration and revoke associated authentication.

Slack integration

With Slack integration in Amazon Quick Suite, you can perform actions within Slack workspaces, including sending messages, managing channels, and interacting with Slack APIs. This integration supports action execution only and requires Amazon Quick Suite Pro tier or higher.

What you can do

With Slack integration, you can perform actions through the action connector.

Action connector

Send messages, manage channels, and access Slack APIs through authenticated connections.

Before you begin

Before you set up Slack integration, make sure you have the following:

- Slack workspace with appropriate permissions.
- Slack app or bot token with required scopes.
- Amazon Quick Suite Author or higher.
- Administrative access to configure OAuth applications (if using user authentication).

Set up Slack OAuth app

Before you configure the integration, you need to create a Slack app with OAuth capabilities. Follow these steps:

1. Go to the Slack API website and create a new Slack app.
2. Configure OAuth scopes based on the actions you want to perform.
3. Set up redirect URLs to match your Amazon Quick Suite integration configuration.
4. Note the client ID and client secret for use in Amazon Quick Suite integration setup.

Set up Slack integration

Use the unified Integrations tab in the Amazon Quick Suite console to set up Slack integration for task execution. Follow these steps:

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Choose **Slack** from the integration options, click the Add (plus "+") button.
3. Fill in the integration details:
 - **Name** - Descriptive name for your Slack integration.
 - **Description** (Optional) - Purpose of the integration.
4. Choose your connection type (user authentication for Slack).
5. Fill in the connection settings for user authentication:
 - Configure the following OAuth fields:
 - **Base URL** - Slack API base URL (typically https://slack.com/api).
 - **Client ID** - Slack app client ID.
 - **Client Secret** - Slack app client secret.
 - **Token URL** - Slack OAuth token endpoint.
 - **Auth URL** - Slack OAuth authorization endpoint.
 - **Redirect URL** - OAuth redirect URI configured in your Slack app.
6. Select **Create and continue**.
7. Choose users to share the integration with.
8. Click **Next**.

Available task actions

After you create your Slack integration, you can review the available actions for interacting with Slack workspaces. Common Slack actions include:

- Send messages to channels or direct messages.
- Create and manage channels.
- Retrieve channel information and member lists.
- Upload and share files.
- Manage user presence and status.
- Access workspace and team information.

Note

The specific actions available depend on the OAuth scopes configured in your Slack app and the permissions granted during authentication.

Share integrations

You can share Slack action connectors with other users in your organization. Follow these steps:

1. After you create the integration, choose **Share integration**.
2. Select users or groups to share the integration with.
3. Set appropriate permissions for shared access.
4. Confirm sharing settings.

Shared users can use the Slack integration to perform actions within the connected Slack workspace, subject to the permissions configured in the original OAuth setup.

Manage Slack integrations

After you create your Slack integration, you can manage it through the Integrations console using these options:

- **Edit integration** - Update authentication settings or OAuth configuration.
- **Share integration** - Make the integration available to other users in your organization.
- **Monitor usage** - View integration activity and API usage metrics.
- **Review actions** - See the complete list of available Slack actions.

- **Delete integration** - Remove the integration and revoke associated OAuth tokens.

Important

Deleting a Slack integration will revoke the OAuth tokens and prevent any shared users from accessing the Slack workspace through Amazon Quick Suite.

Smartsheet integration

With Smartsheet integration in Amazon Quick Suite, you can perform actions within Smartsheet workspaces, including managing sheets, rows, and collaborative work. This integration supports action execution only and requires Amazon Quick Suite Pro tier or higher.

What you can do

With Smartsheet integration, you can perform actions within your Smartsheet workspaces through the action connector.

Action connector

Create, update, and manage sheets, rows, columns, and collaborative features through the Smartsheet API.

Set up Smartsheet integration

Follow these steps to create your Smartsheet integration:

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Choose **Smartsheet** from the integration options, click the Add (plus "+") button.
3. Fill in the integration details:
 - **Name** - Descriptive name for your Smartsheet integration.
 - **Description** (Optional) - Purpose of the integration.
4. Choose your connection type:
 - **User authentication** - OAuth-based authentication for individual user access.
 - **Service authentication** - API key-based authentication for service access.

5. Fill in the connection settings based on your selected authentication method (either user or service):
 - a. For **User authentication (OAuth)**, configure the following fields:
 - **Base URL** - Smartsheet API base URL.
 - **Client ID** - Smartsheet OAuth app client ID.
 - **Client Secret** - Smartsheet OAuth app client secret.
 - **Token URL** - Smartsheet OAuth token endpoint.
 - **Auth URL** - Smartsheet OAuth authorization endpoint.
 - **Redirect URL** - OAuth redirect URI.
 - Required OAuth scopes are READ_SHEETS and WRITE_SHEETS.
 - b. For **Service authentication (API Key)**, configure the following fields:
 - **API Key** - Smartsheet API access token.
 - **Base URL** - Smartsheet API base URL.
 - **Email** - Associated Smartsheet user email.
6. Select **Create and continue**.
7. Choose users to share the integration with.
8. Click **Next**.

Manage Smartsheet integrations

After you create your Smartsheet integration, you can manage it using these options:

- **Edit integration** - Update authentication settings or Smartsheet configuration.
- **Share integration** - Make the integration available to other users.
- **Delete integration** - Remove the integration and revoke authentication.

Web Crawler integration

With Web Crawler integration in Amazon Quick Suite, you can create knowledge bases from website content by crawling and indexing web pages. This integration supports data ingestion capabilities with different authentication options based on your user tier.

What you can do

Web Crawler users can ask questions about content stored on websites and web pages. For example, users can inquire about documentation sites, knowledge bases, or search for specific information across multiple web pages. The integration enables users to quickly access and understand information from web content, regardless of location or type, while providing contextual details such as publication dates, modification history, and page ownership—all contributing to more efficient information discovery and better-informed decision making.

Note

Web Crawler integration supports data ingestion only. It doesn't provide action capabilities for managing websites or web services.

Before you begin

Before you set up Web Crawler integration, make sure you have the following:

- Website URLs to crawl and index.
- Amazon Quick Suite Enterprise subscription
- The website you want to crawl needs to be public and can't be behind a firewall or require special browser plugins to connect.

Prepare website access and authentication

Before setting up the integration in Amazon Quick Suite, prepare your website access credentials. Web Crawler integration supports different authentication methods based on your user role:

No authentication

Available for all users. Use for crawling public websites that don't require authentication.

Basic authentication

Standard HTTP Basic Authentication for secured websites. HTTP Basic Authentication is a simple way to protect web resources by requiring a username and password. When you visit a protected site using Basic Authentication, your browser will show a pop-up dialog box asking for your credentials.

Required credentials:

- **Login page URL** - The URL of the login page
- **Username** - Basic auth username
- **Password** - Basic auth password

Form authentication

For websites that use HTML form-based login pages.

The form is set up for you to specify XPath. XPath (XML Path Language) is a query language used to navigate through elements and attributes in an HTML or XML document. To identify an XPath for a web page element, a user can utilize their browser's developer tools, typically accessed by right-clicking on the desired element and selecting "Inspect" or pressing F12. Once the element is highlighted in the developer tools, the user can right-click on the corresponding HTML code, select "Copy," and then choose "Copy XPath" from the submenu. This generates a unique path that identifies the element's exact location in the document structure. The resulting XPath might look something like `//input[@id='username']` or `//button[@type='submit']`, where the double forward slashes (`//`) indicate the path can start anywhere in the document, and the square brackets contain attributes that help identify the specific element.

Required information:

- **Login page URL** - URL of the login form (e.g., `https://example.com/login`)
- **Username** - Login username
- **Password** - Login password
- **Username field XPath** - XPath to username input field (e.g., `//input[@id='username']`)
- **Username button XPath** (Optional) - XPath to username button field (e.g., `//input[@id='username_button']`)
- **Password field XPath** - XPath to password input field (e.g., `//input[@id='password']`)
- **Password button XPath** - XPath to password button (e.g., `//button[@type='password']`)

SAML authentication

For websites that use SAML-based single sign-on authentication.

SAML (Security Assertion Markup Language) authentication is a federated identity standard that enables single sign-on (SSO) by allowing users to authenticate through a centralized identity provider rather than entering credentials directly into each application. Unlike

traditional form authentication where users type their username and password into fields on the application's login page, SAML redirects users to their organization's identity provider (like Microsoft Azure AD or Okta) to authenticate, then passes a secure token back to the application to grant access. This approach provides a seamless user experience across multiple applications, centralized user management for IT administrators, and enhanced security through features like multi-factor authentication, while form authentication requires separate credential management for each individual application

Required information:

- **Login page URL** - URL of the SAML login page
- **Username** - SAML username
- **Password** - SAML password
- **Username field XPath** - XPath to username input field (e.g., `//input[@id='username']`)
- **Username button XPath** (Optional) - XPath to username button field (e.g., `//input[@id='username_button']`)
- **Password field XPath** - XPath to password input field (e.g., `//input[@id='password']`)
- **Password button XPath** - XPath to password button (e.g., `//button[@type='password']`)

XPath configuration examples

Use these XPath examples to configure form and SAML authentication:

Username field examples:

```
//input[@id='username']  
//input[@name='user']  
//input[@class='username-field']
```

Password field examples:

```
//input[@id='password']  
//input[@name='pass']  
//input[@type='password']
```

Submit button examples:

```
//button[@type='submit']  
//input[@type='submit']  
//button[contains(text(), 'Login')]
```

Set up Web Crawler integration

After preparing your website access requirements, create the Web Crawler integration in Amazon Quick Suite.

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Choose **Web Crawler** from the integration options, and click the **Add** button (plus "+" button).
3. Choose **Access data from Web Crawler**. Web Crawler integration supports data access only - action execution is not available for web crawling.
4. Configure integration details and authentication method, then create knowledge bases as needed.
 1. Select the authentication type for your web crawler integration.
 2. Fill in the required details based on your selected authentication method.
 3. Select **Create and continue**.
 4. Fill in the Name and description for your knowledge base.
 5. Add the content URLs you want to crawl.
 6. Select **Create**.

After clicking create, the data sync is started automatically.

Configure crawling

You can configure which websites and pages to crawl and how to filter the content.

Configure URLs and content sources

Configure which websites and pages to crawl:

Direct URLs

Specify individual URLs to crawl:

```
https://example.com/docs  
https://example.com/blog  
https://example.com/support
```

Limit: Maximum 10 URLs per dataset

Content filters and crawl settings

Crawl scope settings

To view these settings, you must first set up a knowledge base and then examine the advanced settings option.

Crawl depth

- Range: 0-10 (default: 1)
- 0 = crawl only specified URLs
- 1 = include linked pages one level deep
- Higher values follow links deeper into the site

Maximum links per page

- Default: 1000
- Maximum: 1,000
- Controls how many links to follow from each page

Wait time

- Default: 1
- The amount of time the web crawler will wait for each page after the page reaches "page ready" state. This is useful for pages that have dynamic javascript load characteristics where the page has content blocks that loads after the main template is loaded. Increase wait time if you have visually rich content or anticipate high load times.

Manage knowledge bases

After setting up your Web Crawler integration, you can create and manage knowledge bases from your crawled website content.

Edit existing knowledge bases

You can modify your existing Web Crawler knowledge bases:

1. In the Amazon Quick Suite console, choose **Knowledge bases**.
2. Select your Web Crawler knowledge base from the list.
3. Choose the three-dot icon under **Actions**, then choose **Edit knowledge base**.

4. Update your configuration settings as needed and choose **Save**.

Create additional knowledge bases

You can create multiple knowledge bases from the same Web Crawler integration:

1. In the Amazon Quick Suite console, choose **Integrations**, and then select the **Data** tab.
2. Choose your existing Web Crawler integration from the list.
3. Choose the three-dot icon under **Actions**, then choose **Create knowledge base**.
4. Configure your knowledge base settings and choose **Create**.

For detailed information about knowledge base configuration options, see [the section called "Common configuration settings"](#).

Attachments and file crawling

Control whether the system processes files and attachments linked from web pages:

- **Enable file attachment crawling** - Select this option to crawl and index files and attachments found on web pages, such as PDFs, documents, and media files.

Crawling behavior and sync configuration

Your Web Crawler integration follows these crawling practices:

- **Incremental sync model:** First sync performs full crawl, subsequent syncs capture changes only
- **Automatic retry:** Built-in retry logic for failed requests
- **Duplicate handling:** Automatic detection and handling of URLs
- **Crawler identification:** Identifies itself with user-agent string "aws-quick-on-behalf-of-<UUID>" in request headers

Robots.txt compliance

Web Crawler respects the robots.txt protocol and honors user-agent and allow/disallow directives. This enables you to control how the crawler accesses your site.

How robots.txt checking works

- **Host-level checking:** Web Crawler reads robots.txt files at the host level (for example, example.com/robots.txt)
- **Multiple host support:** For domains with multiple hosts, Web Crawler honors robots rules for each host separately
- **Fallback behavior:** If Web Crawler cannot fetch robots.txt due to blocking, parsing errors, or timeouts, it will behave as if robots.txt does not exist and will crawl the site

Supported robots.txt fields

Web Crawler recognizes these robots.txt fields (field names are case-insensitive, values are case-sensitive):

user-agent

Identifies which crawler the rules apply to

allow

A URL path that may be crawled

disallow

A URL path that may not be crawled

sitemap

The complete URL of a sitemap

crawl-delay

Specified amount of time (in seconds) to wait between requests to your website

Meta tag support

Web Crawler supports page-level robots meta tags that you can use to control how your data is used. You can specify page-level settings by including a meta tag on HTML pages or in an HTTP header.

Supported meta tags

`noindex`

Do not index the page. If you don't specify this rule, the page may be indexed and eligible to appear in experiences

`nofollow`

Do not follow the links on this page. If you don't specify this rule, Web Crawler may use the links on the page to discover those linked pages

You can combine multiple values using a comma (for example, "noindex, nofollow").

Note

To detect meta tags, Web Crawler needs to access your page, so do not block your page with robots.txt which will prevent it from being recrawled.

Troubleshooting

Use this section to resolve common issues with Web Crawler integration.

Authentication failures

Symptoms:

- "Unable to authenticate" error messages
- 401/403 HTTP responses
- Login page redirect loops
- Session timeout errors

Resolution steps:

1. Verify the site is reachable from the Amazon region the Amazon Quick Suite instance is setu
2. Verify credentials accuracy and ensure they haven't expired
3. Check authentication endpoint availability and accessibility
4. Validate XPath configurations by testing them in browser developer tools

5. Review browser network logs to understand the authentication flow
6. Ensure login page URL is correct and accessible
7. Test authentication manually using the same credentials

Access and connectivity issues

Symptoms:

- Connection timeouts and network errors
- Network unreachable errors
- DNS resolution failures

Resolution steps:

1. Verify network connectivity to target websites
2. Validate site accessibility:
 - Check DNS resolution for target domains
 - Verify SSL/TLS configuration and certificates
 - Test access from different networks if possible

Crawl and content issues

Symptoms:

- Missing or incomplete content
- Incomplete crawls or early termination
- Rate limiting errors (429 responses)
- Content not being indexed properly

Resolution steps:

1. Review robots.txt restrictions:
 - Check robots.txt file for crawl restrictions
 - Verify crawler is allowed to access target paths

- Ensure robots.txt compliance isn't blocking content
2. Check rate limiting and throttling:
 - Monitor response headers for rate limit information
 - Implement appropriate crawl delays
 3. Verify URL patterns and filters:
 - Test regex patterns for accuracy
 - Check URL formatting and structure
 - Validate include/exclude pattern logic
 4. Review content restrictions:
 - Check for noindex meta tags on pages
 - Verify content type support
 - Ensure content size is within limits
 5. Update the Wait time to an appropriate value so the content loads on the page before the crawler tries crawling

Known limitations

Web Crawler integration has the following limitations:

- **URL limits:** Maximum of 10 URLs, sitemap not supported
- **Crawl depth:** Maximum crawl depth of 10 levels
- **Security requirements:** HTTPS required for web proxy configurations

Zendesk Suite integration

With Zendesk Suite integration in Amazon Quick Suite, you can perform actions within Zendesk instances, including managing tickets, users, and customer support workflows. This integration supports action execution only and requires Amazon Quick Suite Pro tier or higher.

What you can do

With Zendesk Suite integration, you can perform actions within your Zendesk instances through the action connector.

Action connector

Create, update, and manage tickets, users, and customer support processes through the Zendesk API.

Set up Zendesk integration

Follow these steps to create your Zendesk integration:

1. In the Amazon Quick Suite console, choose **Integrations**.
2. Choose **Zendesk Suite** from the integration options, click the Add (plus "+") button.
3. Fill in the integration details:
 - **Name** - Descriptive name for your Zendesk integration.
 - **Description** (Optional) - Purpose of the integration.
4. Choose your connection type:
 - **User authentication** - OAuth-based authentication for individual user access.
 - **Service authentication** - API key-based authentication for service access.
5. Fill in the connection settings based on your selected authentication method (either user or service):
 - a. For **User authentication (OAuth)**, configure the following fields:
 - **Base URL** - Zendesk instance URL.
 - **Client ID** - Zendesk OAuth app client ID.
 - **Client Secret** - Zendesk OAuth app client secret.
 - **Token URL** - Zendesk OAuth token endpoint.
 - **Auth URL** - Zendesk OAuth authorization endpoint.
 - **Redirect URL** - OAuth redirect URI.

Required OAuth scopes are `tickets:read`, `tickets:write`, and `read`.

- b. For **Service authentication (API Key)**, configure the following fields:
 - **API Key** - Zendesk API token.
 - **Base URL** - Zendesk instance URL.

- **Email** - Associated Zendesk user email.
6. Select **Create and continue**.
 7. Choose users to share the integration with.
 8. Click **Next**.

Manage Zendesk integrations

After you create your Zendesk integration, you can manage it using these options:

- **Edit integration** - Update authentication settings or Zendesk configuration.
- **Share integration** - Make the integration available to other users.
- **Delete integration** - Remove the integration and revoke authentication.

Support escalation path

If you cannot resolve an issue for an integration using the troubleshooting steps, follow this escalation path:

1. Gather diagnostic information:
 - Integration configuration details
 - Error messages and timestamps
 - Steps to reproduce the issue
 - Environment details if applicable (Online vs. Server, version)
2. Check Amazon Quick Suite service health and known issues in the console.
3. Contact Amazon Support through the Amazon Quick Suite console or Amazon Support Center.
4. For complex authentication or permission issues, be prepared to involve your administrator.

Creating Amazon service action connectors using the admin console

Action connectors enable you to integrate Amazon services with your automation workflows. You can create these connectors through the Amazon Quick Suite admin console, which provides a user-friendly interface for configuring and managing connector permissions.

This procedure describes how to create first-party (Amazon) action connectors that can be used in Amazon Quick Automate workflows to perform actions on Amazon resources.

Prerequisites

Before you create action connectors, ensure that you have the following:

- Administrative access to your Amazon Quick Suite account
- An IAM role with the necessary permissions for the connector you want to create
- Access to the Amazon Quick Suite admin console

Note

For information about creating customer IAM roles with Amazon service permissions, see the Amazon Quick Automate documentation for creating Amazon action connectors.

Accessing the admin console

To create action connectors, you need to access the Amazon Quick Suite admin console.

1. Sign in to the Amazon using your administrative credentials.
2. Navigate to the Amazon Quick Suite service.
3. In the Amazon Quick Suite console, choose Admin to access the admin console.
4. In the left navigation pane, choose Permissions.
5. Choose Amazon action connectors to view the action connectors page.

Creating a new action connector

After accessing the action connectors page, you can create a new connector.

1. On the action connectors page, choose New action in the upper right corner.
2. Select the first-party connector type that you want to create from the available options.
3. Fill in the connector details:
 - a. Provide a descriptive name for your connector.

- b. In the connection details section, enter the Role ARN for the IAM role that grants the connector permissions to perform actions on resources.
 - c. Configure any additional settings specific to the connector type you selected.
 4. Configure sharing permissions for the connector:
 - a. Choose Share to add users or groups who can access the connector.
 - b. Select the appropriate permission level:
 - **Owner** - Grants full permissions to share, edit, and delete the connector. Recommended if you are granting the connector to your own user account.
 - **User** - Grants permission to use the connector in automation workflows.
 - c. Choose Add to apply the sharing settings.
 5. Review your configuration and choose Create to create the connector.
 6. Verify that the new connector appears in the list of actions on the admin console page.

Using the new connector in Amazon Quick Suite

After creating the action connector, you can use it in Amazon Quick Suite automation workflows.

1. In a new browser tab, sign in to Amazon Quick Suite as a regular user using the same account that you used for the admin console.
2. In the left navigation pane, choose Add Connections, then choose Actions.
3. Verify that you can see the newly created action connector in the list.

Note

If you have Owner permissions, you can also view, edit, share, and delete the connector from this interface.

4. To use the connector in an automation:
 - a. Add the connector to an Automation Group.
 - b. Create an automation that uses the connector's capabilities.
 - c. Verify that the connector can be attached to the Automation Group and used in the workflow.

Next steps

After creating your action connector, you can:

- Create automation workflows that use the connector to perform actions on Amazon resources
- Share the connector with additional users or groups as needed
- Monitor the connector's usage and performance through the admin console
- Update the connector's configuration or permissions as requirements change

Bring Your Own Amazon Q Business Index (BYOI)

Amazon Quick Suite enables you to use your existing Amazon Q Business indexes as data sources. You can leverage your enterprise data without having to recreate indexes. This capability, known as Bring Your Own Index (BYOI), lets you connect your Amazon Q Business indexes to Amazon Quick Suite and use them alongside other data sources for comprehensive analytics and intelligent responses.

Topics

- [Overview of Amazon Q Business indexes in Amazon Quick Suite](#)
- [Prerequisites](#)
- [Setting up permissions](#)
- [Creating knowledge bases from Amazon Q Business indexes](#)
- [Sharing Amazon Q Business index knowledge bases](#)
- [Using Amazon Q Business index knowledge bases](#)
- [Limitations](#)
- [Billing](#)
- [User types and capabilities](#)
- [Common use cases](#)
- [Interacting with your Amazon Q Business indexes](#)

Overview of Amazon Q Business indexes in Amazon Quick Suite

Amazon Q Business indexes contain indexed enterprise data that you can now use directly in Amazon Quick Suite. When you connect an Amazon Q Business index to Amazon Quick Suite, it

becomes available as a knowledge base. You can use it in spaces, agents, and automations, just like any other knowledge base in Amazon Quick Suite.

Key benefits of using Amazon Q Business indexes in Amazon Quick Suite include:

Leverage existing data investments

Use your existing Amazon Q Business indexes and indexed data directly in Amazon Quick Suite. You don't need to recreate indexes or upload the same data again.

Unified analytics experience

Query across multiple data sources including Amazon Q Business indexes, uploaded documents, and structured data. Use Amazon Quick Suite's unified analytical environment for all your data.

Consistent security and permissions

Amazon Q Business index knowledge bases in Amazon Quick Suite maintain the same security and access controls as in Amazon Q Business. Users only see content they have permission to access.

Enhanced collaboration

Share Amazon Q Business index knowledge bases with other users in Amazon Quick Suite, add them to spaces, and use them with agents and automations to create comprehensive business solutions.

Prerequisites

Before you can use Amazon Q Business indexes in Amazon Quick Suite, ensure that you meet the following prerequisites:

- You have an existing Amazon Q Business index with indexed data.
- Both the Q Business instance and the Amazon Quick Suite instance are:
 - Authenticating through IAM Identity Center
 - Are in the same Amazon account
 - Are in the same region
- You have administrator permissions in both Amazon Q Business and Amazon Quick Suite.

 Note

Currently, BYOI requires both Amazon Q Business and Amazon Quick Suite to authenticate through IAM Identity Center. Support for Amazon Q Business instances using other authentication methods is coming in the next few weeks.

Setting up permissions

To use Amazon Q Business indexes in Amazon Quick Suite, you need to set up the appropriate permissions:

1. Sign in to the Amazon Quick Suite console as an administrator.
2. Navigate to the **Admin** section.
3. Select **Amazon Resources**.
4. Choose **Amazon Q Business** from the list of available data sources.
5. Enable permissions for Amazon Quick Suite to access Amazon Q Business indexes in your Amazon account.

Once permissions are set up, Amazon Q Business becomes available as a data source in Amazon Quick Suite, and Author users can create knowledge bases from Amazon Q Business indexes.

Creating knowledge bases from Amazon Q Business indexes

After setting up permissions, Author users can create knowledge bases from Amazon Q Business indexes:

1. Sign in to the Amazon Quick Suite console as an Author PRO user.
2. Navigate to **Knowledge bases**.
3. Choose **Create knowledge base**.
4. Select **Amazon Q Business** as the data source.
5. Choose the Amazon Q Business index you want to use from the list of available indexes.
6. Provide a name and description for the knowledge base.
7. Choose **Create** to create the knowledge base.

Note

Each Amazon Q Business index creates exactly one knowledge base in Amazon Quick Suite. If multiple Author users create knowledge bases from the same Amazon Q Business index, these knowledge bases will functionally be identical.

Sharing Amazon Q Business index knowledge bases

After creating a knowledge base from an Amazon Q Business index, Author users can share it with other users in Amazon Quick Suite:

1. Navigate to **Knowledge bases**.
2. Select the Amazon Q Business index knowledge base you want to share.
3. Choose **Share**.
4. Select the users or groups you want to share the knowledge base with.
5. Choose **Share** to confirm.

Users who have access to an Amazon Q Business index in Amazon Q Business automatically have access to the corresponding knowledge base in Amazon Quick Suite.

Using Amazon Q Business index knowledge bases

Once created, Amazon Q Business index knowledge bases can be used in Amazon Quick Suite just like any other knowledge base:

Using in spaces

Author users can add Amazon Q Business index knowledge bases to spaces:

1. Navigate to the space where you want to add the knowledge base.
2. Choose **Add resources**.
3. Select **Knowledge bases**.
4. Choose the Amazon Q Business index knowledge base from the list.
5. Choose **Add** to confirm.

Using in output cards

Author can use Amazon Q Business index knowledge bases via Spaces in output cards.

Using with agents

Authors can add Amazon Q Business index knowledge bases to custom agents:

1. Navigate to **Agents**.
2. Select an existing agent or create a new one.
3. In the agent configuration, choose **Add knowledge bases**.
4. Select the Amazon Q Business index knowledge base from the list.
5. Choose **Add** to confirm.

Using with automations

Authors can add Amazon Q Business index knowledge bases to automations:

1. Navigate to **Automations**.
2. Select an existing automation or create a new one.
3. In the automation configuration, add a step that uses knowledge bases.
4. Select the Amazon Q Business index knowledge base from the list.
5. Configure the step and save the automation.

Querying Amazon Q Business index knowledge bases

Readers can query Amazon Q Business index knowledge bases through the Amazon Quick Suite web application:

1. Navigate to the Amazon Quick Suite web application.
2. Select a space that contains the Amazon Q Business index knowledge base or use the default agent.
3. Enter your query in the chat interface.
4. View the response, which includes citations and clickable links to the source documents from the Amazon Q Business index.

Limitations

When using Amazon Q Business indexes in Amazon Quick Suite, be aware of the following limitations:

- You can connect up to 2 Amazon Q Business indexes to Amazon Quick Suite in the current release.
- Amazon Q Business index knowledge bases cannot be changed like other knowledge bases in Amazon Quick Suite.
- Amazon Q Business index knowledge bases are currently limited to text files and tabular data.
- Amazon Q Business guardrails are not automatically applied in Amazon Quick Suite. You need to configure similar guardrails in Amazon Quick Suite separately.
- QApps, Actions, and Amazon Q Business chat guardrails are not included in the BYOI capability.
- Amazon Q Business indexes must be in the same Amazon account and region as Amazon Quick Suite, and use the same instance of IAM Identity Center.

Billing

When using Amazon Q Business indexes in Amazon Quick Suite, billing works as follows:

- You are billed for index and subscription costs in Amazon Q Business (recommended at least 1 user).
- You are billed for subscription costs in Amazon Quick Suite based on your Amazon Quick Suite user subscriptions.
- There are no additional charges for connecting Amazon Q Business indexes to Amazon Quick Suite.

For more information about Amazon Quick Suite pricing, see [Amazon Quick Suite pricing](#).

User types and capabilities

Different user types have different capabilities when working with Amazon Q Business indexes in Amazon Quick Suite:

Amazon Q Business Administrator

Manages Amazon Q Business indexes in the Amazon Q Business console.

Amazon Quick Suite Administrator

Enables permissions for Amazon Quick Suite to access Amazon Q Business indexes in the same Amazon account and region.

Amazon Quick Suite Author

Creates knowledge bases from Amazon Q Business indexes, shares them with other users, and adds them to spaces.

Amazon Quick Suite Author

Creates agents and automations that leverage Amazon Q Business index knowledge bases.

Amazon Quick Suite Reader

Uses agents and automations that leverage Amazon Q Business index knowledge bases, and queries these knowledge bases through the Amazon Quick Suite web application.

Common use cases

Here are some common use cases for using Amazon Q Business indexes in Amazon Quick Suite:

Finding relevant information

Query about specific topics across your enterprise data, such as HR policies, support documentation, or compliance guidelines.

Example: "What is the maternity leave policy for my company in Washington state?" where national guidelines are in an Amazon Q Business index and state-specific information is in Amazon Quick Suite.

Business analytics and insights

Extract key findings from status updates across multiple groups or analyze the health of your product portfolio.

Example: "Analyze the health of my product portfolio" where product sales data is in an Amazon Q Business index, and product engineering and support data is in Amazon Quick Suite.

Content creation and summarization

Create summaries or reports based on documents from multiple sources.

Example: "Summarize last week's Service WBR and Sales update into a one-pager reflecting key initiatives" where the Service WBR is indexed in an Amazon Q Business index and the latest sales update is uploaded to Amazon Quick Suite.

Interacting with your Amazon Q Business indexes

Once you've connected your Amazon Q Business index to Amazon Quick Suite, you can interact with it through chat using agents.

1. From the left navigation menu in the console, select **Agents**.
2. From the **Actions** column for the agent you want to share, select the menu icon, and then select **Chat**.
3. You can then narrow down the resource to just the Knowledge Base created from the Amazon Q Business index by selecting the All Resources option and clicking on your Amazon Q Business index from the Knowledge Bases sub-tab.

When you ask questions, the agent analyzes data from your Amazon Q Business index alongside other data sources to provide comprehensive responses.

You can ask complex questions in plain language about your enterprise data and get detailed responses that draw from your Amazon Q Business index's indexed content. The chat interface allows you to:

- Query your enterprise data using natural language
- Get contextual responses that combine data from your Amazon Q Business index with other Amazon Quick Suite data sources
- View source citations with clickable references to see where information originated
- Continue conversations and build on previous queries

For detailed information about using the chat interface and working with agents, see [the section called "Use a chat agent"](#).

Cross-region inference in Amazon Quick Suite integrations

With cross-region inference, Amazon Quick Suite automatically selects the optimal region within your geography to process your inference requests, maximizing available compute resources and

model availability, and providing the best customer experience. With cross-region inference, you get:

- Complete access to the most advanced Amazon Quick Suite AI capabilities and features
- Access to a variety of models suitable for different tasks
- Improved performance for all your applications and integrations

Cross-region inference requests are kept within the Amazon Regions that are part of the geography where the data originally resides. For example, a request made within the US is kept within the Amazon Regions in the US. Although the data remains stored only in the primary region, when using cross-region inference, your input prompts and output results may move outside of your primary region. All data will be transmitted encrypted across Amazon's secure network.

 **Note**

There's no additional cost for using cross-region inference.

Amazon CloudWatch and Amazon CloudTrail logs won't specify the Amazon Region in which data inference occurs.

Supported regions for Amazon Quick Suite cross-region inference

Supported Amazon Quick Suite Cross-Region Inference Regions

Supported Amazon Quick Suite geography	Inference regions
United States	US East (N. Virginia) (us-east-1)
	US East (Ohio) (us-east-2)
	US West (Oregon) (us-west-2)
Europe	Europe (Frankfurt) (eu-central-1)
	Europe (Stockholm) (eu-north-1)
	Europe (Milan) (eu-south-1)

Supported Amazon Quick Suite geography	Inference regions
	Europe (Spain) (eu-south-2)
	Europe (Ireland) (eu-west-1)
	Europe (Paris) (eu-west-3)
Australia	Asia Pacific (Tokyo) (ap-northeast-1)
	Asia Pacific (Seoul) (ap-northeast-2)
	Asia Pacific (Osaka) (ap-northeast-3)
	Asia Pacific (Mumbai) (ap-south-1)
	Asia Pacific (Singapore) (ap-southeast-1)
	Asia Pacific (Singapore) (ap-southeast-1)
	Asia Pacific (Sydney) (ap-southeast-2)
	Asia Pacific (Melbourne) (ap-southeast-2)

Connect Amazon Quick Suite to your existing tools with extensions

Note

Before you install an extension as an author, your Quick Suite admin must [configure Amazon Quick Suite access](#) to the extension.

Amazon Quick Suite extensions enable you to optimize your daily tasks with assistance from an AI tool added to the tools you already use. Whether you're browsing the web, collaborating in team channels, managing emails, or creating documents, Quick Suite extensions bring AI-powered assistance directly into your existing workflows without requiring you to switch between applications.

The following extension types are available:

- **Browser Extension** - Enhance your web browsing experience with AI-powered page summarization, content analysis, and knowledge access across Google Chrome, Firefox, and Microsoft Edge browsers. No admin setup required.
- **Slack extension** - Access Quick Suite capabilities directly within Slack channels and direct messages, enabling seamless collaboration and knowledge sharing without leaving your team communication platform. Requires admin to configure extension access, uses OAuth installation.
- **Microsoft Outlook extension** - Streamline email management with AI assistance for drafting responses, summarizing email threads, and accessing organizational knowledge while composing or reviewing emails. Requires admin to configure extension access, uses manifest-only installation.
- **Microsoft Teams extension** - Integrate Quick Suite into your Teams workspace for real-time collaboration, thread analysis, and knowledge retrieval during meetings and conversations. Requires admin to configure extension access, uses OAuth installation.
- **Microsoft Word extension** - Enhance document creation and editing with AI-powered writing assistance, content review, and document analysis capabilities built into your Word environment. Requires admin to configure extension access, uses manifest-only installation.

Each extension maintains enterprise-grade security and privacy standards while providing contextual AI assistance tailored to the specific application environment. All extensions integrate with your organization's Quick Suite knowledge bases, custom agents, and configured data sources to provide relevant, accurate responses based on your permissions and access levels.

Topics

- [Amazon Quick Suite user interaction with extensions](#)
- [Amazon Quick Suite browser extension](#)
- [Amazon Quick Suite Microsoft Outlook extension](#)
- [Amazon Quick Suite Slack extension](#)
- [Amazon Quick Suite Microsoft Word extension](#)
- [Amazon Quick Suite Microsoft Teams extension](#)

Amazon Quick Suite user interaction with extensions

Amazon Quick Suite extensions enable different user roles to collaborate effectively in creating, managing, and utilizing AI-powered analytics tools. Each role—administrators, authors, and end users—has distinct responsibilities and capabilities when working with extensions. Understanding these role-based interactions ensures proper governance, security, and optimal user experience across your organization's Amazon Quick Suite deployment.

Topics

- [User capabilities and permissions](#)
- [Amazon Quick Suite administrator responsibilities](#)
- [Amazon Quick Suite author capabilities](#)
- [Amazon Quick Suite end user access](#)
- [Extension access troubleshooting](#)

User capabilities and permissions

Amazon Quick Suite implements a hierarchical permission model for extensions that balances organizational control with user productivity. Administrators establish the foundational policies and infrastructure, authors manage the deployment and configuration of specific extension

instances, and readers benefit from the AI-powered assistance within their authorized scope. The following table shows how each user type in Amazon Quick Suite interacts with extensions.

User Type	Capabilities
Administrators	• Enable/remove extension access to their organization's workspaces and organizations • Create extensions for Slack, Teams, Word, and Outlook • Assign owners for extensions created • Control whether extensions are available for users in Amazon Quick Suite • Use all enabled extensions
Authors	• Create and install extensions using administrator enabled extension access • Edit/delete extensions for which they are owners • Use all enabled extensions
Readers (Reader Pro only)	• Use all enabled extensions

Beyond the broad user type categories, Amazon Quick Suite implements granular permission controls that determine specific actions users can perform with individual extensions. These permissions operate independently of user types, allowing administrators to fine-tune access based on organizational needs and security requirements. The following table outlines how user permissions determine what you can do with a extension:

Permissions Type	Permissions
Owners	• Edit extension name and description • Install the extension in a workspace or organization • Grant ownership and editing permissions to specific users and groups
Viewers	• View extension name and description • Access links to open in Slack, M365, and browser

Permissions Type	Permissions
	• Use the extensions

 Important

The Amazon Quick Suite browser extension extension doesn't need any admin setup to install and use.

These permission structures ensure that extension access aligns with your organization's security requirements while enabling appropriate levels of functionality for each user role. Administrators maintain control over extension availability and initial setup, authors handle deployment and configuration, and all users can benefit from the AI-powered assistance once extensions are properly configured.

Amazon Quick Suite administrator responsibilities

Amazon Quick Suite administrators, or users granted admin privileges, establish the foundational infrastructure that enables extension functionality across the organization. They handle the critical setup and security configurations that allow Amazon Quick Suite authors to deploy extensions and end users to access them safely.

 Note

Creating and adjusting extension access requires IAM administrator privileges. All administrators will see the extension links, but if not [signed in with IAM credentials](#), you will need to sign in with appropriate IAM permissions to manage extension access. This must be an IAM user who is an Amazon Quick Suite administrator and not non-IAM admin users.

Responsibility area	Administrator capabilities
Infrastructure setup	• Connect extension applications to organizational systems (Slack workspaces, Microsoft 365 tenants) • Handle OAuth authorization and tenant-level permissions

Responsibility area	Administrator capabilities
	• Manage platform-specific installation methods: manifest-only for Word/Outlook, OAuth-only for Slack/Teams
Security and compliance	• Manage enterprise security and compliance settings • Control which extension types are available to the organization

Amazon Quick Suite author capabilities

Authors bridge the gap between administrative infrastructure and end-user access by creating, configuring, and deploying specific extension instances. They manage the day-to-day operational aspects of extension deployment and maintenance.

Author capabilities depend on the permission level granted by administrators. There are two permission types that determine what authors can do:

- **Viewers:** Can view extension name and description, access links to open extensions, and use the extensions.
- **Owners:** Can edit extension name and description, install extensions in workspaces or organizations, and share ownership permissions with users and groups.

This translates to two main setup flows:

- **Limited permissions (view, share, delete only):** Authors can only use the basic landing page functionality to view, share, and delete extensions after admin completes all setup.
- **Full permissions (deploy, view, share, delete, edit):** Authors can download manifests or use OAuth deployments to complete app setup, rename extensions, and access editing features.

Capability area	Author functions
Configuration and deployment	• Download and manage installation manifests for deployment (with full permissions) • Deploy extension instances using OAuth or manifest-based methods (with full permissions)

Capability area	Author functions
	• Complete app setup and installation (with full permissions)
Management and sharing	• Manage extension sharing and access permissions • View, share, and delete deployed extensions • Edit extensions and access creation features (with full permissions)

Amazon Quick Suite end user access

End users represent the primary beneficiaries of extension functionality, accessing AI-powered assistance directly within their existing workflow applications. Their interaction with extensions focuses on daily productivity and seamless integration with organizational knowledge.

Access level	User capabilities
Daily usage and productivity	• Access extensions directly within existing workflow applications • Use extensions for knowledge access, document analysis, and action taking • Benefit from seamless integration without context switching between applications
Knowledge and actions	• Access organizational knowledge bases and custom agents through extension interfaces • Perform external actions in third-party applications using configured action connectors • Receive permissions-aware responses based on individual access levels

The following table shows how each user type in Amazon Quick Suite interacts with extensions.

User Type	Capabilities
Administrators	• Enable/remove extension access to their organization's workspaces and organizations

User Type	Capabilities
	• Create extensions for Slack, Teams, Word, and Outlook • Assign owners for extensions created • Control whether extensions are available for users in Amazon Quick Suite • Use all enabled extensions
Authors	• Create and install extensions using administrator enabled extension access • Edit/delete extensions for which they are owners • Use all enabled extensions
Readers (Reader Pro only)	• Use all enabled extensions

The following table outlines how user permissions determine what you can do with a extension:

Permissions Type	Permissions
Owners	• Edit extension name and description • Install the extension in a workspace or organization • Grant ownership and editing permissions to specific users and groups
Viewers	• View extension name and description • Access links to open in Slack, M365, and browser • Use the extensions

 Important

The Amazon Quick Suite browser extension extension doesn't need any admin setup to install and use.

Extension access troubleshooting

If you encounter issues accessing or creating extensions:

- **Cannot create extensions:** Your IAM user who is a Amazon Quick Suite administrator must first configure extension access for the extension type you want to use. Other administrators will not be able to create/edit/delete extension access.
- **Cannot find my extensions:** Check that extension access has been configured by your administrator and that you have the appropriate permissions to view extensions.
- **Extensions appear but cannot edit:** You have view-only access. Ask your administrator or extension owner to share edit permissions with you.
- **No extensions visible:** Contact your administrator to set up extension access and create default extensions for your organization.

Amazon Quick Suite browser extension

Important

The Amazon Quick Suite browser extension extension doesn't need any admin setup to install and use.

The Amazon Quick Suite browser extension integrates AI-powered assistance directly into your web browsing experience. Users can access Quick Suite capabilities within their browser to enhance productivity, streamline research, and get instant help without switching between applications.

Key capabilities include:

- Summarize web pages for quick understanding.
- Ask questions about multiple web pages and access organizational knowledge from spaces.
- Use actions from [action connectors](#) configured in Quick Suite.
- Analyze files uploaded to Quick Suite during browsing sessions.

The AI-powered assistance enables you to maximize reading productivity, streamline research and analysis of complex information, and get instant help when creating content without context switching during web browsing.

Topics

- [Supported browsers](#)
- [Browser extension benefits](#)
- [Browser extension use cases](#)
- [Install and configure the browser extension](#)
- [Deploy the browser extension in your organization](#)
- [Browser extension usage guidelines](#)
- [Browser extension known limitations](#)

Supported browsers

The Quick Suite browser extension is compatible with the following browsers:

- Google Chrome
- Mozilla Firefox
- Microsoft Edge

Browser extension benefits

The Quick Suite browser extension provides powerful AI-driven capabilities that transform how you interact with web content. These key benefits help you work more efficiently and make better use of information you encounter while browsing, bringing enterprise knowledge directly into your web research workflows without requiring application switching.

Embedded productivity features

- Summarize a snapshot of any web page for quick understanding and analysis.
- Ask questions about one or more web pages using your organizational knowledge.
- Analyze both files you upload to Quick Suite and specific spaces it can access.
- Execute actions from configured action connectors without leaving your browser.

Browser extension use cases

The following are common use cases that help you make the best use of your Quick Suite browser extension:

Summarize web pages

Use this procedure to quickly understand web page content without reading the entire page.

1. Open the Quick Suite browser extension.
2. Log in and navigate to the web page you want to summarize.
3. From the chat interface, select **Summarize** for a summary of a snapshot of that web page.

Your conversation now contains a snapshot of this web page. You can continue to chat about the web page and ask follow-up questions.

Add individual pages to your conversation scope

Use this procedure to include multiple web pages in your conversation context for comprehensive analysis.

1. Navigate to the web page you want to add to your conversation.
2. Select the **+** icon to add the page to your scope.
3. Navigate from tab to tab and select the **+** icon for each page you want to chat about.
4. You can now chat about all the pages you've added to your scope.

This allows you to analyze and compare information across multiple web pages in a single conversation.

Add files to a conversation

Upload files to enhance your conversation with document analysis capabilities.

1. Select the **paperclip** icon.
2. Select the files to add to the conversation.
3. Select **Add to chat**.

The files are now available to use in the Quick Suite conversation.

Add webpages as context to a conversation

Include multiple browser tabs in your conversation for comprehensive web content analysis.

1. Select **No tabs selected**.
2. From **Add tabs**, choose the tabs you want to include in your Quick Suite conversation.
3. Select **Confirm**.

The tabs are now available in your Quick Suite conversation.

Reset conversation context

Start fresh when you need to clear your current conversation context.

1. Select the **+ icon inside the bubble** to start a new conversation.
2. Alternatively, you can reset your conversation by selecting the **X** icon to end the current conversation and then open the Quick Suite extension to start a new conversation.

This clears all previous context and allows you to begin with a clean slate.

Use actions in chat

Execute external actions directly from your browser extension conversation.

1. Ask the chat to perform a specific action.
2. Follow the prompts in the chat and on-screen to complete the action.

This enables you to perform tasks in external applications without leaving your browser.

Install and configure the browser extension

As a user, you can install and configure the browser extensions with the following steps.

- Mozilla based browsers — [Mozilla Firefox Add-ons web store](#)
- Chromium based browsers (including Microsoft Edge) — [Chrome Web Store](#)

Once you have successfully logged on, you can use your Quick Suite browser extension.

Note

You can also **pin** your Quick Suite browser extension to have it readily accessible while using your browser. Instructions for this are specific to your browser of choice. The following third-party information about pinning extensions might be helpful.

- Google Chrome — <https://www.howtogeek.com/683099/how-to-pin-and-unpin-extensions-from-the-chrome-toolbar>
- Mozilla Firefox — https://support.mozilla.org/en-US/kb/extensions-button#w_manage-pinned-extensions
- Microsoft Edge — <https://www.microsoft.com/en-us/edge/features/pin-to-taskbar>

Deploy the browser extension in your organization

As an author, after enabling the browser extension in your Quick Suite application, you can deploy it across your organization using enterprise management tools and policies.

For Google Chrome Enterprise environments, you can use organizational unit policies to manage extension deployment. The Chrome Web Store for Enterprise provides additional deployment options specifically designed for organizational use. Mozilla Firefox Enterprise offers similar capabilities through policy templates and enterprise distribution methods that allow for automated extension deployment across the organization. Microsoft Edge Enterprise provides policy settings for managing extensions through mobile device management (MDM) software.

Policy settings from browser vendors: [Firefox](#), [Chrome](#), and [Edge](#).

Deploy to Google Chrome Enterprise

Follow these steps to deploy the browser extension across your Chrome Enterprise environment.

1. Access your Google Admin console and navigate to the Chrome management section.
2. Configure extension policies using organizational units to target specific user groups or departments.
3. Use the [Chrome extension management guide](#) to set up automated installation policies.
4. Reference the [Chrome Web Store Enterprise documentation](#) for advanced deployment configurations.

The extension will now be automatically installed for all users in your Chrome Enterprise environment.

Deploy to Mozilla Firefox Enterprise

Use these steps to deploy the extension across your Firefox Enterprise environment.

1. Configure Firefox deployment policies using the enterprise policy framework.
2. Follow the [Firefox extension deployment guide](#) for step-by-step instructions.
3. Use the [enterprise distribution documentation](#) to configure organization-wide deployment.
4. Apply extension settings using [Mozilla policy templates](#) for centralized configuration management.

The extension is now deployed across your Firefox Enterprise environment with centralized management.

Deploy to Microsoft Edge Enterprise

Deploy the extension to Microsoft Edge Enterprise using these organizational deployment steps.

1. Install the browser extension for all users using the software deployment processes of your organization.
2. Configure extension settings using [Edge policy settings](#) through mobile device management (MDM) software.
3. Follow the [Microsoft Edge extension management guide](#) for detailed deployment instructions.

Your Microsoft Edge Enterprise deployment is complete with centralized policy management.

Browser extension usage guidelines

As a user, you are responsible for keeping company information safe. The following guidelines help you use Quick Suite browser extensions securely while maintaining data privacy and compliance.

Using the browser extension

When using the Quick Suite browser extension, users can review past conversations and attachments to maintain context across browsing sessions.

Conversation retention and data privacy

Amazon Quick Suite maintains strict data privacy through automated retention policies:

- Amazon Quick Suite automatically deletes conversations and related files after 30 days of inactivity.
- Browser extension conversation retention follows your company's specific history rules.
- Browser extension retention periods may exceed Amazon Quick Suite's standard 30-day limit based on organizational policies.

Browser extension known limitations

The following are known limitations of the Amazon Quick Suite browser extension:

- The browser extension does not support [flows](#).
- The browser extension supports up to a total of 20 web pages and uploaded files at a time. For more information, see [Upload files and chat](#).

Amazon Quick Suite Microsoft Outlook extension

Note

Before you deploy a Microsoft Outlook extension as an author, your Amazon Quick Suite admin must [configure extension access to Microsoft Outlook](#).

The Amazon Quick Suite Microsoft Outlook extension brings AI-powered assistance directly into email workflows, enabling access to organizational knowledge and document analysis without leaving Microsoft Outlook.

Key capabilities include:

- Extension access through Microsoft Outlook add-ins for continuous availability.
- Email thread analysis and summarization for quick understanding of lengthy discussions.
- Draft email responses based on conversation context and organizational knowledge.
- External actions in third-party applications using [action connectors](#) configured by your organization.

- Unlimited usage for enhanced email productivity and communication workflows.
- General knowledge capabilities.
- Company-specific information search when enabled.
- Source-attributed responses.
- File upload and processing of up to 20 files simultaneously.
- Contextual summaries and insights generation.

Topics

- [Amazon Quick Suite Microsoft Outlook extension author guide](#)
- [Amazon Quick Suite Microsoft Outlook extension user guide](#)

Amazon Quick Suite Microsoft Outlook extension author guide

As an Amazon Quick Suite author, you can deploy Quick Suite Microsoft Outlook extensions after your admin establishes the foundational connection to your organization's Microsoft 365 tenant. Your capabilities depend on the permission level granted by your administrator.

With **limited permissions** (view, share, delete only), you can manage basic extension operations through the landing page after admin completes all setup. With **full permissions** (deploy, view, share, delete, edit), you can download manifests for deployment in the M365 admin center, rename extensions, and access all editing features.

Author capabilities for Microsoft Outlook extensions include:

- Configure extension connections (requires access to M365 Admin Center portal).
- Deploy extensions to your organization's Microsoft 365 tenant (requires full permissions).
- Manage sharing and access permissions (available with limited or full permissions).
- Customize extension settings - names, descriptions (requires full permissions).
- Download manifest files for advanced Microsoft 365 deployment (requires full permissions).

Note

Before you deploy a Microsoft Outlook extension as an author, your Quick Suite admin must [configure Amazon Quick Suite access to Microsoft Outlook](#).

Topics

- [Deploy Microsoft Outlook extension](#)
- [Edit Microsoft Outlook extension](#)
- [Share Microsoft Outlook extension](#)
- [Delete Microsoft Outlook extension](#)

Deploy Microsoft Outlook extension

Deploy a new Microsoft Outlook extension instance in the Quick Suite console. This process establishes the foundation for connecting AI-powered assistance to your Microsoft Outlook environment.

Note

This action requires full author permissions.

1. Sign in to the Amazon Quick Suite console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select **Create extension**.
4. Select **Microsoft Outlook**. Then, select **Next**.
5. Configure the following fields:
 - **Name** - A name for your extension is pre-filled for you. You can edit this and enter a descriptive name for the Microsoft Outlook extension.
 - **Description** (optional) - A description for your extension is pre-filled for you. You can edit this and enter a new description to provide additional context about this extension configuration.
 - **Installation** type - Your Microsoft Outlook extension uses manifest-only installation by default.
6. Select **Next** to save your configuration.
7. From the **Extension** summary page, navigate to the extension you just configured.
8. Then, from the **Actions** menu, navigate to the extension you just configured.
9. Select **Download manifest**. Then, from the **Complete installation for Outlook assistant** dialog box that opens, select **Download**.

The manifest file will be downloaded to your computer.

10. From the success message, select **Install extension** to finish downloading the manifest for your extension.

 Note

You can also navigate to the extensions summary page and download the manifest for your extension from the **Actions** menu.

11. In the screen asking for permissions to allow your Amazon Quick Suite Outlook extension to access Outlook, select **Allow**.

 Note

You will now continue the remainder of this procedure within the Microsoft 365 admin center.

12. In the Microsoft 365 admin center, choose **Integrated apps** from the left navigation and choose **Upload custom apps**. This will open the **Deploy New App** page.
13. Choose **Office Add-in** as your App type.
14. Paste the manifest URL link you copied in the **Provide link to manifest file** and choose **Validate**.
15. Choose the users you want to add in the **Add users** section.
16. Choose **Accept permissions** in the **Accept permissions requests** section and deploy the Add-in. Once deployment is completed, your users will be able to install the Amazon Quick Suite Add-in in their Microsoft Outlook.

Your Outlook extension has now been successfully deployed and is available for users.

Edit Microsoft Outlook extension

As an author, you can edit the extensions you deploy to your users. Modify extension settings to update names, descriptions, or configuration options. Changes take effect immediately and apply to all users with access to the extension.

1. Sign in to the Amazon Quick Suite console.

2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Microsoft Outlook extension you need to edit.
4. Select **Edit**.
5. Edit the configuration as required and select **Save** to confirm the changes.

Share Microsoft Outlook extension

Share ownership and management permissions with specific users and groups, enabling multiple users to manage extensions and assist with installation. You can assign different permission levels and manage access as needed.

1. Sign in to the Amazon Quick Suite console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Microsoft Outlook extension you need to share.
4. Select **Share**.
5. Enter the users and groups you would like to share the extension with.
6. Select **Share** to send the access email to each group and user.
7. From the drop-down next to each name, you can edit their access levels (Viewer or Owner).
8. **Optional:** You could select **Remove access** to delete the access for the selected group or user.

Delete Microsoft Outlook extension

As an author, you can delete the extensions you deploy to your users. Permanently remove a extension from your Quick Suite console and revoke access for all users. This action cannot be undone and requires confirmation.

1. Sign in to Amazon Quick Suite console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Microsoft Outlook extension you need to delete.
4. Select **Delete**.
5. Enter the word, "confirm", and select **DELETE**.

Amazon Quick Suite Microsoft Outlook extension user guide

As a user, you can access Amazon Quick Suite directly within Microsoft Outlook to get AI-powered assistance without leaving your email environment. The extension integrates seamlessly into your Microsoft Outlook workflow, providing instant access to knowledge bases and intelligent responses.

Topics

- [Add a Microsoft Outlook Extension to your Microsoft Outlook](#)
- [Access Quick Suite in your Microsoft Outlook navigation pane at any time](#)
- [Microsoft Outlook extension use cases](#)
- [Amazon Quick Suite Microsoft Outlook extension usage guidelines](#)
- [Microsoft Outlook extension known limitations](#)

Add a Microsoft Outlook Extension to your Microsoft Outlook

Install the Quick Suite app in your Microsoft Outlook environment to enable AI-powered assistance. This one-time setup makes Quick Suite accessible throughout your Microsoft Outlook workflow.

Add a Microsoft Outlook extension to your Microsoft Outlook

1. Open your Microsoft Outlook.
2. Go to **Add-Ins** and search for **Quick Suite** and choose **Quick Suite**.
3. You will now see the Amazon Quick Suite logo on the top bar on the right where you can access the chat assistant.
4. Once the add-in is installed, you can use your Quick Suite **Add-In**.

Access Quick Suite in your Microsoft Outlook navigation pane at any time

Once you've added Quick Suite to your Microsoft Outlook environment, you can access it conveniently through the navigation pane for quick assistance without interrupting your email workflow. The navigation pane provides persistent access to Quick Suite while you work with emails and conversations.

Microsoft Outlook extension use cases

With Quick Suite integrated into Microsoft Outlook, you can search and access your organization's knowledge bases directly from your email conversations. This allows you to find relevant information, get answers to questions, and enhance your communications without leaving Microsoft Outlook.

Get summaries and action items from lengthy email threads

1. Locate the Quick Suite icon under **add-ins** in your Microsoft ribbon and select the **summarize** shortcut.
2. Alternatively, open the Quick Suite side panel and ask any question of your choice about the email thread in focus.

Draft speedy replies to emails

1. Navigate to an email thread of your choice and locate the Quick Suite icon under **add-ins** in your Microsoft ribbon. Select the **reply** or **reply all** shortcut.
2. Quick Suite will create a draft response and open a reply window that is ready to send.
3. Iterate on the draft with Quick Suite and click send.

Incorporate Quick Suite's company knowledge and general knowledge into your emails

1. Open an email thread of your choice and ask Quick Suite about its knowledge.
2. Use the **insert text** or **replace selection** buttons to insert Quick Suite's response into your email drafts.

Perform actions in external applications

1. Start a direct message (DM) with Quick Suite.
2. Ask it to perform an action of your choice in an external application using [action connectors](#) configured by your organization.

Amazon Quick Suite Microsoft Outlook extension usage guidelines

As a user, you are responsible for keeping company information safe. The following guidance helps you use Amazon Quick Suite apps securely while maintaining data privacy and compliance.

Conversation retention policy and accuracy

Amazon Quick Suite automatically deletes conversations after 30 days of inactivity. Microsoft Outlook conversation retention follows your company's specific history rules, and Microsoft Outlook retention periods may exceed Amazon Quick Suite's 30-day limit.

Users can review past conversations and attachments. Closing the Quick Suite chat panel ends the current conversation. Reopen the panel to start a new conversation. However, users need to visit the Amazon Quick Suite chat instance to manage their conversations.

Security considerations

When you use Amazon Quick Suite with email content, remember that responses reflect individual user permissions and content may include information not accessible to all email recipients. It may generate responses that contain sensitive data that can't be publicly shared.

To protect your organization's data, carefully evaluate email usage and plan deployment with data privacy in mind. Amazon Quick Suite maintains strict data privacy by not using customer data for service improvements, not using customer data to enhance language models, and not indexing Microsoft Outlook conversations into your company's Amazon Quick Suite instance.

Microsoft Outlook extension known limitations

The following are known limitations of the Amazon Quick Suite Microsoft Outlook extension:

- The Microsoft Outlook extension doesn't support Amazon Quick Suite accounts that use IAM Identity Center for authentication at this time.
- The Microsoft Outlook extension is only able to access your current selected email (which can include your email thread). It is unable to answer questions across all emails/calendar events in your inbox. To enable access to your entire inbox and calendar, configure the Microsoft Outlook [actions integration](#).
- The Microsoft Outlook extension can't generate visuals from structured data.
- Web search isn't supported in the Microsoft Outlook extension at this time.

- You cannot review your conversation history in Microsoft Outlook at this time. To review and manage your conversation history, use the Amazon Quick Suite chat instance.
- The Microsoft Outlook extension doesn't support [flows](#).
- The Microsoft Outlook extension supports up to 19 uploaded files and your current email thread. For more information on file upload support, see [Upload files and chat](#).
- The Microsoft Outlook extension doesn't support filtering during chat based on **Recently used** agents. Users will have access to the full list of available agents and can pick agents they want to use.
- The Microsoft Outlook extension doesn't support the **Actions** menu, so you can't invoke actions explicitly. However, implicit actions are supported. For more information about explicit and implicit actions, see [Using Actions in Chat](#).

Amazon Quick Suite Slack extension

Note

Before you deploy a Slack extension as an author, your Amazon Quick Suite admin must [configure Amazon Quick Suite access to Slack](#).

The Amazon Quick Suite Slack extension brings AI-powered assistance directly into team conversations, enabling access to organizational knowledge and document analysis without leaving Slack.

Key capabilities include:

- Extension access through Slack's AI assistant side panel for continuous availability.
- Engagement with Quick Suite through **@Amazon Quick Suite** mentions in channels for contextual team assistance. Contextual responses with relevant content from the Slack channel are not supported at this time.
- Direct interaction through private messages for individual productivity.
- External actions in third-party applications from Slack using [action connectors](#) configured by your organization.
- Unlimited usage for enhanced team collaboration and knowledge sharing.
- General knowledge capabilities.

- Company-specific information search when enabled.
- Source-attributed responses.
- File upload and processing of up to 20 files simultaneously.
- Contextual summaries and insights generation.

Topics

- [Amazon Quick Suite Slack extension author guide](#)
- [Amazon Quick Suite Slack extension user guide](#)

Amazon Quick Suite Slack extension author guide

As an Amazon Quick Suite author, you can deploy Quick Suite Slack extensions after your admin establishes the foundational connection to your organization's Slack workspace. Your capabilities depend on the permission level granted by your administrator.

With **limited permissions** (view, share, delete only), you can manage basic extension operations through the landing page after admin completes all setup. With **full permissions** (deploy, view, share, delete, edit), you can complete the Slack OAuth application deployment, rename extensions, and access all editing features.

Author capabilities for Slack extensions:

- Deploy extensions to your organization's Slack workspace using OAuth (requires full permissions).
- Customize extension settings - names, descriptions (requires full permissions).
- Manage sharing and access permissions (available with limited or full permissions).

Note

Before you deploy a Slack extension as an author, your Quick Suite admin must [configure Amazon Quick Suite access to Slack](#).

Topics

- [Deploy a Slack extension](#)

- [Edit Slack extension](#)
- [Share Slack extension](#)
- [Delete Slack extension](#)

Deploy a Slack extension

Deploy a new Slack extension instance in the Quick Suite console. This process establishes the foundation for connecting AI-powered assistance to your Slack workspace.

Note

This action requires full author permissions.

1. Sign in to the Amazon Quick Suite console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select **Create extension**.
4. Select **Slack**. Then, select **Next**.
5. Configure the following fields:
 - **Name** - A name for your extension is pre-filled for you. You can edit this and enter a descriptive name for the Slack extension.
 - **Description** (optional) - A description for your extension is pre-filled for you. You can edit this and enter a new description to provide additional context about this extension configuration.
 - **Installation** type - Your Slack extension supports OAuth installation by default.
6. Select **Next** to save your configuration.
7. From the **Extension** summary page, navigate to the extension you just configured.
8. Then, from the **Actions** menu, navigate to the extension you just configured.
9. Select **Install**. Then, from the **Complete installation for Slack assistant** dialog box that opens, select **Install**.

A success message will open up on the top right of your screen.

10. From the success message, select **Install extension** to finish installing your extension.

Note

You can also navigate to the extensions summary page and deploy your extension from the **Actions** menu.

11. The link will open to a Slack login page where after you login (as admin) you will need to find and be asked to install the Amazon Quick Suite app within your Slack workspace.

Note

You may see a "This app is not approved by Slack banner." This message can be ignored.

12. Choose **Allow** to install your Amazon Quick Suite App for Slack.
13. Once the installation is complete, you will see the page confirming that the **Congratulations! Your Slack App has been successfully installed.**
14. Choose **Open the Amazon Quick Suite App in Slack.**
15. This will open your Slack workspace where all users will be required to sign-in.

Your Slack extension has now been successfully deployed and is available for users.

Edit Slack extension

Note

This action requires full author permissions.

As an author, you can edit the extensions you deploy to your users. Modify extension settings to update names, descriptions, or configuration options. Changes take effect immediately and apply to all users with access to the extension.

1. Sign in to the Amazon Quick Suite console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Slack extension you need to edit.
4. Select **Edit**.

5. Edit the configuration as required and select **Save** to confirm the changes.

Share Slack extension

Share ownership and management permissions with specific users and groups, enabling multiple users to manage extensions and assist with installation. You can assign different permission levels and manage access as needed.

1. Sign in to the Amazon Quick Suite console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Slack extension you need to share.
4. Select **Share**.
5. Enter the users and groups you would like to share the extension with.
6. Select **Share** to send the access email to each group and user.
7. From the drop-down next to each name, you can edit their access levels (**Viewer** or **Owner**).
8. **Optional:** You could select **Remove access** to delete the access for the selected group or user.

Delete Slack extension

As an author, you can delete the extensions you deploy to your users. Permanently remove a extension from your Quick Suite console and revoke access for all users. This action cannot be undone and requires confirmation.

1. Sign in to Amazon Quick Suite console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Slack extension you need to delete.
4. Select **Delete**.
5. Enter the word, "confirm", and select **DELETE**.

Amazon Quick Suite Slack extension user guide

As a user, you can access Amazon Quick Suite directly within Slack to get AI-powered assistance without leaving your workspace. The extension integrates seamlessly into your Slack environment, providing instant access to knowledge bases and intelligent responses.

Topics

- [Add a Slack Extension to your Slack](#)
- [Access Quick Suite in your Slack side panel at any time](#)
- [Slack extension use cases](#)
- [Amazon Quick Suite Slack extension usage guidelines](#)
- [Slack extension known limitations](#)

Add a Slack Extension to your Slack

Install the Amazon Quick Suite app in your Slack workspace to enable AI-powered assistance. This one-time setup makes Quick Suite accessible throughout your Slack environment.

Add a Slack Extension to your Slack

1. Open and login to the Slack workspace for your company.
2. Choose **More** in the left navigation, then select **Automations**.
3. Choose **Apps**.
4. Choose **+ Add apps**.
5. Search for **Quick Suite** and choose Quick Suite.
6. Select your Slack profile picture in the bottom left of Slack.
7. Navigate to **Preferences > Navigation > AI Apps and Agents > select Show app agents > and select Quick Suite**.
8. Quick Suite is now available to access anywhere in Slack through an icon in the top right of your Slack window.

Note

For more information, see [Understand AI apps in Slack](#) in the Slack help center.

Access Quick Suite in your Slack side panel at any time

Once you've added Quick Suite to your Slack workspace, you can access it conveniently through the side panel for quick assistance without interrupting your workflow. The side panel provides persistent access to Quick Suite while you work in Slack channels and direct messages.

Access Quick Suite in your Slack side panel at any time

1. Click on your Slack profile picture in the bottom left of Slack.
2. Navigate to **Preferences > Navigation > AI Apps and Agents** and select **Amazon Quick Suite**.
3. Quick Suite will now be available to access anywhere in Slack through an icon in the top right of your Slack window.

Slack extension use cases

With Quick Suite integrated into Slack, you can search and access your organization's knowledge bases directly from your conversations. This allows you to find relevant information, get answers to questions, and enhance your communications without leaving Slack.

Search your organization's knowledge wherever you are

1. Open the Quick Suite side panel from the top right of your Slack window or start a DM with Quick Suite under **Apps**.
2. Ask Quick Suite a question about its knowledge from your spaces and connected knowledge bases.
3. To personalize your conversation using your data, use the gear icon (visible after sending your first message to Quick Suite) within your conversation to select an agent or a space to respond from.

Get Quick Suite's in-context help in Slack conversation threads

1. Navigate to a Slack conversation thread of your choice.
2. Type **@Amazon Quick Suite** and ask any question about the messages in that thread or its knowledge from your spaces and Quick Suite knowledge bases.
3. If Quick Suite is not already added to the channel, you will be prompted to add it before receiving a response.

Perform actions in external applications

1. Start a direct message (DM) with Quick Suite.
2. Ask it to perform an action of your choice in an external application using [action connectors](#) configured by your organization.

Upload and analyze images

1. In any Slack channel or direct message, upload an image by dragging and dropping it or using the attachment button.
2. Mention **@Amazon Quick Suite** and ask questions about the uploaded image, such as requesting analysis, descriptions, or insights.
3. Quick Suite will analyze the image and provide relevant information based on your request.

Amazon Quick Suite Slack extension usage guidelines

As a user, you are responsible for keeping company information safe. The following guidance helps you use Amazon Quick Suite apps securely while maintaining data privacy and compliance.

When using the Slack extension

To start fresh, begin a new conversation by selecting **New chat** in direct messages. Closing the side panel ends your current conversation. You can access conversation history through Slack message history or the Amazon Quick Suite chat instance, which includes conversations from all channels. You can view complete conversation details, including attachment names, and access the same knowledge base available in your Amazon Quick Suite chat instance.

Note

Users can access conversation history in Slack. Deleting the conversation in Slack will not delete from conversation history in Amazon Quick Suite. However, users need to visit the Amazon Quick Suite chat instance to manage their conversations.

Conversation retention policy and accuracy

Amazon Quick Suite automatically deletes conversations after 30 days of inactivity. Slack conversation retention follows your company's specific history rules, and Slack retention periods may exceed Amazon Quick Suite's 30-day limit.

Users can access conversation history in Slack through the AI App History tab. Deleting the conversation in Slack will not delete from conversation history in Amazon Quick Suite. However, users need to visit the Amazon Quick Suite chat instance to manage their conversations.

Security considerations

When you use Amazon Quick Suite in public channels, remember that responses reflect individual user permissions and content may include information not accessible to all channel members. It may generate responses that contain sensitive data that can't be publicly shared.

To protect your organization's data, carefully evaluate public channel usage, consider the security implications, and plan deployment with data privacy in mind. Amazon Quick Suite maintains strict data privacy by not using customer data for service improvements, not using customer data to enhance language models, and not indexing Slack conversations into your company's Amazon Quick Suite instance.

Slack extension known limitations

The following are known limitations of the Amazon Quick Suite Slack extension:

- The Slack extension isn't supported for Amazon Quick Suite accounts that use IAM Identity Center for authentication.
- The Slack extension doesn't have access to all your Slack conversations. It only has access to the messages in Slack conversation threads where it is mentioned.
- The Slack extension doesn't have access to previous messages when mentioned in a Slack channel or in group message.
- The Slack extension isn't accessible in group direct messages or in direct messages with teammates unless it is explicitly added to the conversation as a member.
- You can filter Amazon Quick Suite agents and spaces only in direct messages with the Slack extension. The Slack extension will use the default assistant and all spaces available to you when invoked in public such as in channels or group messages.
- The Slack extension can't render visuals for prompts about your structured data.
- Web search isn't supported in the Slack extension.
- The Slack extension can't be configured to automatically reply to messages in channels or posts.
- The Slack extension will not respond when invoked by Slack's workflow bot as it is not able to authenticate the user who made the query in these situations.
- The Slack extension isn't available in huddles.
- [Actions](#) (using your action integrations) can only be performed in direct messages with the Slack extension.
- Actions that require file uploads as inputs are not supported by the Slack extension.

- [Flows](#) are not supported in the Slack extension at this time.
- File upload limitations are the same within Amazon Quick Suite chat. For more information, see [Upload files and chat](#).

Amazon Quick Suite Microsoft Word extension

Note

Before you deploy a Microsoft Word extension as an author, your Amazon Quick Suite admin must [configure Amazon Quick Suite access to Microsoft Word](#).

The Amazon Quick Suite Microsoft Word extension brings AI-powered assistance directly into document workflows, enabling access to organizational knowledge and document analysis without leaving Microsoft Word.

Key capabilities include:

- Extension access through Microsoft Word add-ins for continuous availability.
- Document review and content suggestions for improving clarity, grammar, and style.
- Document summarization and content generation capabilities.
- External actions in third-party applications using [action connectors](#) configured by your organization.
- Unlimited usage for enhanced document productivity and writing workflows.
- General knowledge capabilities.
- Company-specific information search when enabled.
- Source-attributed responses.
- File upload and processing of up to 20 files simultaneously.
- Contextual summaries and insights generation.

Topics

- [Amazon Quick Suite Microsoft Word extension author guide](#)
- [Amazon Quick Suite Microsoft Word extension user guide](#)

Amazon Quick Suite Microsoft Word extension author guide

As an Amazon Quick Suite author, you can deploy Quick Suite Microsoft Word extensions after your admin establishes the foundational connection to your organization's Microsoft 365 tenant. Your capabilities depend on the permission level granted by your administrator.

With **limited permissions** (view, share, delete only), you can manage basic extension operations through the landing page after admin completes all setup. With **full permissions** (deploy, view, share, delete, edit), you can download manifests for deployment in the M365 admin center, rename extensions, and access all editing features.

Author capabilities for Microsoft Word extensions include:

- Configure extension connections (requires access to M365 Admin Center portal).
- Deploy extensions to your organization's Microsoft 365 tenant (requires full permissions).
- Customize extension settings - names, descriptions (requires full permissions).
- Download manifest files for advanced Microsoft 365 deployment (requires full permissions).
- Manage sharing and access permissions (available with limited or full permissions).

Note

Before you deploy a Microsoft Word extension as an author, your Quick Suite admin must [configure Amazon Quick Suite access to Microsoft Word](#).

Topics

- [Deploy Microsoft Word extension](#)
- [Edit Microsoft Word extension](#)
- [Share Microsoft Word extensions](#)
- [Delete Microsoft Word Extensions](#)

Deploy Microsoft Word extension

Deploy a new Microsoft Word extension instance in the Amazon Quick Suite console. This process establishes the foundation for connecting AI-powered assistance to your Microsoft Word environment.

 Note

This action requires full author permissions.

Use this procedure to create and configure a new Microsoft Word extension for your organization.

1. Sign in to the Amazon Quick Suite console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select **Create extension**.
4. Select **Microsoft Word**. Then, select **Next**.
5. Configure the following fields:
 - **Name** - A name for your extension is pre-filled for you. You can edit this and enter a descriptive name for the Microsoft Word extension.
 - **Description** (optional) - A description for your extension is pre-filled for you. You can edit this and enter a new description to provide additional context about this extension configuration.
 - **Installation** type - Your Microsoft Word extension uses manifest-only installation by default.
6. Select **Next** to save the details and download the manifest file to your computer.

 Note

You will now continue the remainder of this procedure within the Microsoft 365 admin center.

1. In the Microsoft 365 admin center, choose **Integrated apps** from the left navigation and choose **Upload custom apps**. This will open the **Deploy New App** page.
2. Choose **Office Add-in** as your App type.
3. Choose **Upload manifest file (.xml) from device**, select **Choose from file**, select the downloaded manifest xml file and choose **Next**.
4. Choose the users you want to add in the **Add users** section.

5. Choose **Accept permissions** in the **Accept permissions requests** section and deploy the Add-in. Once deployment is completed, your users will be able to install the Amazon Quick Suite Add-in in their Microsoft Word.

Your Microsoft Word extension is now created and ready for deployment to users in your organization.

Edit Microsoft Word extension

Note

This action requires full author permissions.

As an author, you can edit the extensions you deploy to your users. Modify existing Microsoft Word extension configurations to update settings, change descriptions, or adjust deployment parameters as your organizational needs evolve.

Use this procedure to modify settings and configuration for an existing Microsoft Word extension.

1. Sign in to the Amazon Quick Suite console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Microsoft Word extension you need to edit.
4. Select **Edit**.
5. Enter your changes and select **Save** to confirm the new configuration.

Your changes are now applied and will be reflected in the extension configuration for all users.

Share Microsoft Word extensions

Share ownership and management permissions with specific users and groups, enabling multiple users to manage extensions and assist with installation. You can assign different permission levels and manage access as needed.

Use this procedure to share your Microsoft Word extension with other users and manage their access permissions.

1. Sign in to Amazon Quick Suite console.

2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Microsoft Word extension you need to share.
4. Select **Share**.
5. Enter the users and groups you would like to share the extension with.
6. Select **Share** to send the access email to each group and user.
7. From the drop-down next to each name, you can edit their access levels (**Viewer** or **Owner**).
8. **Optional** You could select, **Remove access** to delete the access for the selected group or user.

The specified users and groups now have access to your Microsoft Word extension with the permissions you assigned.

Delete Microsoft Word Extensions

As an author, you can delete the extensions you deploy to your users.

Use this procedure to permanently remove a Microsoft Word extension from your organization.

1. Sign in to Amazon Quick Suite console
2. In the left navigation, under **CONNECTIONS**, select **Extensions**
3. Select the three dot menu icon for the Microsoft Word extension you need to delete
4. Select **Delete**
5. Enter, confirm and **DELETE** to delete the Extension

The Microsoft Word extension has been permanently removed and is no longer accessible to users in your organization.

Amazon Quick Suite Microsoft Word extension user guide

As a user, you can access Amazon Quick Suite directly within Microsoft Word to get AI-powered assistance without leaving your workspace. The extension integrates seamlessly into your Microsoft Word environment, providing instant access to knowledge bases and intelligent responses.

Topics

- [Add the Quick Suite Microsoft Word extension](#)
- [Access Quick Suite in your Microsoft Word home menu at any time](#)
- [Microsoft Word extension use cases](#)

- [Amazon Quick Suite Microsoft Word extension usage guidelines](#)

Add the Quick Suite Microsoft Word extension

Install the Amazon Quick Suite app in your Microsoft Word workspace to enable AI-powered assistance. This one-time setup makes Amazon Quick Suite accessible throughout your Microsoft Word environment.

1. Open Microsoft Word for the web or the desktop application.
2. Navigate to **Insert > Get Add-ins**.
3. Search for **Amazon Quick Suite** in the **Office Add-ins** store.
4. Select **Add** to install the Quick Suite add-in.
5. Once the add-in is installed, you can use your Quick Suite **Add-In**.

Access Quick Suite in your Microsoft Word home menu at any time

Once you've added Quick Suite to your Microsoft Word environment, you can access it conveniently through the home menu for quick assistance without interrupting your document workflow. The home menu provides persistent access to Quick Suite while you work with documents.

Microsoft Word extension use cases

With Quick Suite integrated into Microsoft Word, you can search and access your organization's knowledge bases directly from your documents. This allows you to find relevant information, get answers to questions, and enhance your content without leaving Microsoft Word.

Get summaries and insights from your documents

1. Open the Quick Suite add-in from the ribbon or side panel in your Microsoft Word document.
2. Select the **summarize** shortcut from the ribbon, or ask any question about the document content in the side panel.
3. Quick Suite will analyze your document and provide relevant summaries or insights based on your request.

Generate and improve document content

1. Navigate to the location in your document where you want to add or improve content.

2. Open the Quick Suite add-in and select the **generate content** or **improve writing** shortcut.
3. Describe what you need, and Quick Suite will generate suggestions that match your document's style and context.

Review and enhance selected text

1. Select the text you want to review or improve in your document.
2. Open the Quick Suite side panel and use quick prompts such as **explain in simple terms** or **simplify writing**.
3. Quick Suite will provide suggestions for improving clarity, grammar, and style based on your selected content.

Incorporate Quick Suite's company knowledge and general knowledge into your docs

1. Open the Quick Suite side panel while working on your document.
2. Ask questions about your organization's knowledge bases and spaces to find relevant information for your document.
3. Use the **insert text** or **replace selection** buttons to insert Quick Suite's response into your document.

Perform actions in external applications

1. Open the Quick Suite side panel in your Microsoft Word document.
2. Ask Quick Suite to perform an action in an external application using [action connectors](#) configured by your organization.

Amazon Quick Suite Microsoft Word extension usage guidelines

As a user, you are responsible for keeping company information safe. The following guidance helps you use Amazon Quick Suite apps securely while maintaining data privacy and compliance.

Conversation retention policy and accuracy

Quick Suite automatically deletes conversations and related files after 30 days of inactivity. Microsoft Word conversation retention follows your company's specific history rules, and Microsoft Word retention periods may exceed Quick Suite's 30-day limit.

Users can review past conversations and attachments. To start fresh, begin a new conversation by selecting **New chat**. However, users need to visit the Amazon Quick Suite chat instance to manage their conversations.

 Note

Conversation history is not currently available in Microsoft Word. To manage your conversations, use the Quick Suite chat instance.

Security considerations

When you use Amazon Quick Suite with document content, remember that responses reflect individual user permissions and content may include information not accessible to all document recipients. It may generate responses that contain sensitive data that can't be publicly shared.

To protect your organization's data, carefully evaluate document usage and plan deployment with data privacy in mind. Amazon Quick Suite maintains strict data privacy by not using customer data for service improvements, not using customer data to enhance language models, and not indexing Microsoft Word documents into your company's Quick Suite instance.

No document is read when the side panel is closed, and no data or document content is sent to Quick Suite unless you explicitly send a prompt to the extension.

Quick Suite uses generative AI, so double check responses before using them. Usage of the Quick Suite extension for Microsoft Word is subject to the Amazon Responsible AI Policy. [Amazon Responsible AI Policy](#)

Microsoft Word extension known limitations

The following are known limitations of the Amazon Quick Suite Microsoft Word extension:

- The Microsoft Word extension doesn't support Amazon Quick Suite accounts that use IAM Identity Center for authentication.
- The Microsoft Word extension can't generate visuals from structured data.
- Web search isn't supported in the Microsoft Word extension.
- You can't review your conversation history in the Microsoft Word extension at this time. To review and manage your conversation history, visit your Amazon Quick Suite chat instance.

- The Microsoft Word extension doesn't support [flows](#).
- The Microsoft Word extension supports up to 19 uploaded files and your Microsoft Word document. For more information on file upload limits, see [Upload files and chat](#).
- The Microsoft Word extension doesn't support filtering during chat based on **Recently used** agents. Users will have access to the full list of available agents and can pick agents they want to use.
- The Microsoft Word extension doesn't support the **Actions** menu, so you can't invoke actions explicitly. However, implicit actions are supported. For more information about explicit and implicit actions, see [Using Actions in Chat](#).

Amazon Quick Suite Microsoft Teams extension

Note

Before you deploy a Microsoft Teams extension as an author, your Quick Suite admin must [configure Amazon Quick Suite access to Microsoft Teams](#).

The Amazon Quick Suite Microsoft Teams extension brings AI-powered assistance directly into team conversations, enabling access to organizational knowledge and document analysis without leaving Microsoft Teams.

Key capabilities include:

- Seamless integration within Microsoft Teams allowing you to access information while staying in your conversations.
- Collaborative assistance through **@Amazon Quick Suite** mentions in thread discussions for thread-specific information.
- Action item extraction and summary generation from conversations.
- File upload and processing capabilities for enhanced productivity.

Topics

- [Amazon Quick Suite Microsoft Teams extension author guide](#)
- [Amazon Quick Suite Microsoft Teams extension user guide](#)

Amazon Quick Suite Microsoft Teams extension author guide

As an Amazon Quick Suite author, you can deploy Quick Suite Microsoft Teams extensions after your admin establishes the foundational connection to your organization's Microsoft Teams workspace. Your capabilities depend on the permission level granted by your administrator.

With **limited permissions** (view, share, delete only), you can manage basic extension operations through the landing page after admin completes all setup. With **full permissions** (deploy, view, share, delete, edit), you can complete OAuth deployments, download manifests, rename extensions, and access all editing features.

Author capabilities for Microsoft Teams extensions include:

- Configure extension connections (requires admin privileges).
- Deploy extensions to your organization's Microsoft Teams workspace (requires full permissions).
- Manage sharing and access permissions (available with limited or full permissions).
- Customize extension settings - names, descriptions (requires full permissions).
- Download manifest files for advanced Microsoft Teams customization (requires full permissions).

Note

Before you deploy a Microsoft Teams extension as an author, your Quick Suite admin must [configure Amazon Quick Suite access to Microsoft Teams](#).

Topics

- [Deploy Microsoft Teams extension](#)
- [Edit Microsoft Teams extension](#)
- [Share Microsoft Teams extension](#)
- [Delete Microsoft Teams extension](#)

Deploy Microsoft Teams extension

Deploy a new Microsoft Teams extension instance in the Amazon Quick Suite console. This process establishes the foundation for connecting AI-powered assistance to your Microsoft Teams workspace.

 Note

This action requires full author permissions.

Use this procedure to create and configure a new Microsoft Teams extension for your organization.

1. Sign in to the Amazon Quick Suite console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select **Create extension**.
4. Select **Microsoft Teams**. Then, select **Next**.
5. Configure the following fields:
 - **Name** - A name for your extension is pre-filled for you. You can edit this and enter a descriptive name for the Microsoft Teams extension.
 - **Description** (optional) - A description for your extension is pre-filled for you. You can edit this and enter a new description to provide additional context about this extension configuration.
 - **Installation** type - Your Microsoft Teams extension supports OAuth installation by default.
6. Select **Next** to save your configuration.
7. From the **Extension** summary page, navigate to the extension you just configured.
8. Then, from the **Actions** menu, navigate to the extension you just configured.
9. Select **Install**. Then, from the **Complete installation for Teams assistant** dialog box that opens, select **Install**.

A success message will open up on the top right of your screen.

10. From the success message, select **Install extension** to finish installing your extension.

 Note

You can also navigate to the extensions summary page and deploy your extension from the **Actions** menu.

11. Open the link and login as Global Admin or someone with administrative permissions can add the Amazon Quick Suite App to the Microsoft Teams admin center for your organization.

12. Choose **Teams apps** in the left navigation.
13. Choose **Amazon Quick Suite** from the list of available apps.
14. Review and grant admin consent by choosing the **Permissions** tab and reviewing the permissions and choose **Grant admin consent**.

 Note

If permissions are already granted, proceed to the end of the procedure, there is no further action required.

15. Authenticate and choose **Accept** for Amazon Quick Suite app.
16. Confirm that an app titled Amazon Quick Suite **Permissions** tab now says **Admin consent granted for all required permissions**.
17. All users assigned to the app from the Teams admin center can now find the app in the **Built for your org** section of the **Apps** page of their Teams app.

Your Microsoft Teams extension has now been successfully deployed and is available for users.

Edit Microsoft Teams extension

 Note

This action requires full author permissions.

As an author, you can edit the extensions you deploy to your users. Modify extension settings to update names, descriptions, or configuration options. Changes take effect immediately and apply to all users with access to the extension.

Use this procedure to modify settings and configuration for an existing Microsoft Teams extension.

1. Sign in to the Amazon Quick Suite console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Microsoft Teams extension you need to edit.
4. Select **Edit**.
5. Edit the configuration as required and select **Save** to confirm the changes.

Your changes are now applied and will be reflected in the extension configuration for all users.

Share Microsoft Teams extension

Share ownership and management permissions with specific users and groups, enabling multiple users to manage extensions and assist with installation. You can assign different permission levels and manage access as needed.

Use this procedure to share your Microsoft Teams extension with other users and manage their access permissions.

1. Sign in to the Amazon Quick Suite console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Microsoft Teams extension you need to share.
4. Select **Share**.
5. Enter the users and groups you would like to share the extension with.
6. Select **Share** to send the access email to each group and user.
7. From the drop-down next to each name, you can edit their access levels (**Viewer** or **Owner**).
8. **Optional:** You could select **Remove access** to delete the access for the selected group or user.

The specified users and groups now have access to your Microsoft Teams extension with the permissions you assigned.

Delete Microsoft Teams extension

As an author, you can delete the extensions you deploy to your users. Permanently remove a extension from your Quick Suite console and revoke access for all users. This action cannot be undone and requires confirmation.

Use this procedure to permanently remove a Microsoft Teams extension from your organization.

1. Sign in to Amazon Quick Suite console.
2. In the left navigation, under **CONNECTIONS**, select **Extensions**.
3. Select the three dot menu icon for the Microsoft Teams extension you need to delete.
4. Select **Delete**.
5. Enter the word, "confirm", and select **DELETE**.

The Microsoft Teams extension has been permanently removed and is no longer accessible to users in your organization.

Amazon Quick Suite Microsoft Teams extension user guide

As a user, you can access Quick Suite directly within Microsoft Teams to get AI-powered assistance without leaving your workspace. The extension integrates seamlessly into your Microsoft Teams environment, providing instant access to knowledge bases and intelligent responses.

Topics

- [Add a Microsoft Teams Extension to your Microsoft Teams](#)
- [Microsoft Teams extension use cases](#)
- [Amazon Quick Suite Microsoft Teams extension usage guidelines](#)
- [Microsoft Teams extension known limitations](#)

Add a Microsoft Teams Extension to your Microsoft Teams

Install the Quick Suite app in your Microsoft Teams workspace to enable AI-powered assistance. This one-time setup makes Quick Suite accessible throughout your Microsoft Teams environment.

Add a Microsoft Teams Extension to your Microsoft Teams

1. Open and login to the Microsoft Teams workspace for your company.
2. Choose **Apps** in the left navigation and search for **Quick Suite**.
3. Choose **Quick Suite**.
4. On the top bar, Quick Suite icon is now available. You can select it to access and chat with Quick Suite.

Note

For more information, see [Chat with a bot in Microsoft Teams](#) from Microsoft support.

Microsoft Teams extension use cases

With Amazon Quick Suite integrated into Microsoft Teams, you can search and access your organization's knowledge bases directly from your conversations. This allows you to find relevant information, get answers to questions, and enhance your communications without leaving Team.

Search your organization's knowledge without leaving Teams

1. Start a chat with Quick Suite after adding it from the **Apps** section.
2. Ask Quick Suite a question about its knowledge from your spaces and connected knowledge bases.
3. To personalize your conversation using your data, use the gear icon (visible after sending your first message to Quick Suite) within your conversation to select an agent or a space to respond from.

Get Quick Suite's in-context help in Teams posts

1. Navigate to a Teams post of your choice.
2. Type **@Amazon Quick Suite** and ask any question about the messages in that thread or its knowledge from your spaces and your organization's knowledge bases.
3. If Quick Suite is not already added to the channel, you will be prompted to add it before receiving a response.

Perform actions in external applications

1. Start a direct message (DM) with Quick Suite.
2. Ask it to perform an action of your choice in an external application using [action connectors](#) configured by your organization.

Amazon Quick Suite Microsoft Teams extension usage guidelines

As a user, you are responsible for keeping company information safe. The following guidance helps you use Amazon Quick Suite apps securely while maintaining data privacy and compliance.

Conversation retention policy and accuracy

Amazon Quick Suite automatically deletes conversations after 30 days of inactivity. Teams conversation retention follows your company's specific history rules. Teams retention periods may exceed Amazon Quick Suite's 30-day limit. To start fresh, begin a new conversation by using the `/new_conv` command

Users can review past conversations in Teams or all conversations from all channels (Teams, browser extensions, etc.) in your Amazon Quick Suite chat instance.

Note

Deleting conversations or messages in Teams does not remove them from Amazon Quick Suite. To manage your conversation history, use the Amazon Quick Suite chat instance.

Security considerations

When Amazon Quick Suite is invoked by a user in a public Teams channel, it generates responses based on the invoking user's permissions, which may include content that other channel members aren't authorized to access. To prevent unintended exposure of sensitive information, carefully evaluate the use of Amazon Quick Suite in public channels.

File upload behavior in Amazon Quick Suite Teams Extension:

- Responses will be limited to the uploaded file content.
- General knowledge access depends on admin settings.
- Start a new chat to access company knowledge again.

Data privacy and conversation behavior

Amazon Quick Suite does not use customer data for service improvement or for improving its underlying large language models (LLMs). Also, none of the data you include in your Teams conversations will be indexed into your company's Amazon Quick Suite instance.

The Amazon Quick Suite Teams app will have access to the same knowledge available in the corresponding Amazon Quick Suite web experience.

Closing the Amazon Quick Suite bot for Teams side panel will end the current conversation. Users can review past conversations in Teams or all conversations from all channels (Teams, browser

extensions, etc.) in your Amazon Quick Suite web experience. You can access all the history of previous conversations including, the names of the attachments in those conversations.

Microsoft Teams extension known limitations

The following are known limitations of the Amazon Quick Suite Microsoft Teams extension:

- The Microsoft Teams extension isn't supported for Amazon Quick Suite accounts that use IAM Identity Center for authentication.
- The Microsoft Teams extension doesn't have access to all your Microsoft Teams conversations. It only has access to the messages in Microsoft Teams posts where it is mentioned.
- The Microsoft Teams extension isn't accessible in group direct messages or in direct messages with other teammates.
- You can filter Amazon Quick Suite agents and spaces only in direct messages with the Microsoft Teams extension. The Microsoft Teams extension will use the default assistant and all spaces available to you when invoked in public, such as in channels and Microsoft Teams posts.
- The Microsoft Teams extension can't render visuals for prompts about your structured data.
- Web search is not supported in the Microsoft Teams extension.
- The Microsoft Teams extension can't be configured to automatically reply to messages in channels or posts.
- The Microsoft Teams extension isn't available for your Microsoft Teams meetings.
- You cannot review your conversation history in Microsoft Teams. To review and manage your conversation history, go to your Amazon Quick Suite chat instance.
- [Actions](#) (using your action integrations) can only be performed in direct messages with the Microsoft Teams extension.
- Actions that require file uploads as inputs are not supported by the Microsoft Teams extension at this time.
- [Flows](#) are not supported in the Microsoft Teams extension at this time.
- File upload limitations are the same as those in the web experience. For more information see [Upload files and chat](#).

Document history for the Amazon Quick Suite User Guide

This page describes changes to the *Quick Suite User Guide*. For notifications about these documentation changes, subscribe to the RSS feed using the link near the top of this page.

To find out what's new in Quick Suite, see [What's New with Analytics](#). To find out about changes to the Quick Suite service, sign up for the [Quick Suite newsletter](#).

New Amazon Quick Suite releases appear in different Amazon Web Services Regions over a period of time, beginning with the first Region on the initial release date. Documentation is released in coordination with this process. If you have questions, contact [Amazon Support](#) or your technical account manager.

 **Note**

The following table describes important changes in each *Amazon Quick Suite User Guide* release since March 4, 2019.

Change	Description	Date
New Amazon Quick Suite capabilities	Amazon Quick Suite provides an integrated suite of AI-powered analytics, automation, and research tools that transcend traditional business intelligence. Chat and AI capabilities: Amazon Quick Suite chat - Natural language interface for data exploration with contextual awareness, file uploads, web search, and artifact generation	October 9, 2025

- Quick Suite chat agents - Customizable AI assistants with domain-specific expertise, configurable personas, and organizational context

Workflow automation:

- Amazon Quick Flows - Intelligent workflows combining generative AI with automated actions through no-code interface, supporting conversational and structured runtime modes
- Amazon Quick Automate - AI-powered automation platform that creates sophisticated workflows from natural language descriptions or process documents

Research and analysis:

- Amazon Quick Research - Comprehensive research capability that analyzes multiple data sources and generates detailed reports with citations and source tracing for market research, competitive analysis, and business intelligence

Data and knowledge management:

- Amazon Quick Index - Data discovery and cataloging capabilities
- Spaces - Collaborative workspaces for organizing and sharing data, dashboards, topics, and files with team-specific knowledge bases

Integration and connectivity:

- Integrations - Connect to external applications for unstructured data access, including Google Drive, Amazon S3, OneDrive, and MCP servers
- Action connectors - Perform actions in third-party services like Microsoft Outlook, PagerDuty, and Slack with 15+ supported integrations
- Knowledge bases - Create AI-accessible content from third-party applications and data sources
- Extensions - Browser, Slack, Microsoft Teams, Outlook, and Word extensions for AI assistance within existing tools

- Bring Your Own Index (BYOI) - Connect existing Amazon Q Business indexes as data sources

[Amazon QuickSight rebrands to Amazon Quick Suite](#)

Amazon QuickSight has been rebranded to Amazon Quick Suite, expanding from a standalone business intelligence service to a comprehensive analytics and AI platform. QuickSight continues as Amazon Quick Sight, a core component within the Quick Suite ecosystem.

October 9, 2025

Amazon Quick Sight retains existing functionality including interactive data visualization, SPICE in-memory analytics, embedded analytics, and dashboard sharing. All existing QuickSight APIs, SDKs, and integrations continue to work without changes.

[Font customization](#)

Authors can now customize fonts for data labels and axes in supported charts, in addition to the previously supported font customizations for visual titles, subtitles, and legends, as well as table and pivot table headers.

October 6, 2025

Google Sheets data source	You can now create datasets using a Google Sheets data source. Admins will need to perform a one-time setup to enable Google Sheets as a data source.	September 27, 2025
Q&A null support	Q&A now supports comprehensive support for null value handling, enabling users to create more sophisticated analyses and answer complex business questions.	September 18, 2025
Interactive Highchart features	Highchart visualizations now support custom actions, highlighting, and custom field color consistencies, allowing you to create interactive and visually cohesive charts that integrate seamlessly with other visuals.	September 2, 2025
Calculated fields limitation increase	The number of limits for calculated fields in an analysis has increased from 200 to 500. For more information, see Item limits for analyses .	August 15, 2025
Apache Impala data source	You can now create datasets using an Apache Impala data source.	August 1, 2025

[Improve Q&A accuracy with custom instructions](#)

Authors can now curate responses to questions by adding domain-specific knowledge that cannot be captured through a topic's metadata settings. This ensures more accurate, relevant, and tailored answers that are better suited for business needs.

July 18, 2025

[New topic: Trusted identity propagation support with Athena data sources](#)

Added information about using trusted identity propagation with QuickSight and Athena. For more information, see [Using trusted identity propagation with Athena](#).

June 30, 2025

[2 billion row SPICE datasets](#)

QuickSight now supports SPICE datasets that contain up to 2 billion (2,000,000,000) rows. For more information, see [SPICE quotas for imported data](#).

June 30, 2025

[Topic Permissions](#)

Authors can now manage permissions for dashboards and their linked topics from a single location. Control topic viewer access directly from a dashboard's sharing preferences, eliminating the need to manage permissions in multiple locations.

April 18, 2025

[Sheet Highlighting](#)

Added the ability to highlight specific data points across visuals in a sheet. When you select or hover over a data point on a visual, related data across other visuals will stand out, while unrelated data is dimmed.

April 11, 2025

[Updated topic: Opting out of Amazon Q in Quick Suite](#)

Added additional information and guidance for opting out of Amazon Q in Quick Suite. For more information, see [Opting out of Amazon Q in Quick Suite](#).

April 8, 2025

[Enable Generative BI features in embedded dashboards and consoles](#)

Added documentation for adding executive summaries to embedded dashboards, and adding executive summaries, data stories, authoring, and Q&A to embedded consoles. For more information, see [Embedding with the Amazon QuickSight APIs](#).

March 31, 2025

[Customize alert emails](#)

Previously, the custom email template in your QuickSight account only affected dashboard report and paginated report emails. Now, it also customizes alert emails. For more information, see [Customizing report and alert emails](#).

March 31, 2025

Rules Dataset flag	A Rules Dataset flag has been implemented to distinguish permission datasets used for row-level security from regular datasets.	March 28, 2025
Publish a previous version of a dashboard	You can now revert back to a previous version of a dashboard by searching for it under the dashboard's Version History and publishing it. Each dashboard can store up to 1000 versions that are never deleted.	March 28, 2025
Scenarios general availability	Updated the scenarios documentation to support its general availability release. For more information, see Working with scenarios in Amazon QuickSight .	March 25, 2025
New Region	Amazon QuickSight is now available in the Europe (Spain) (eu-south-2) Region. For more information, see Supported Amazon Regions for Amazon QuickSight .	March 11, 2025

[Dashboard Q&A experience](#)

Amazon QuickSight authors can now to turn on the Q&A experience directly from their dashboards in one click without the need to create a topic in QuickSight. For more information, see [Turn on the Dashboard Q&A experience in Amazon QuickSight.](#)

January 31, 2025

[Unique key configuration](#)

Amazon QuickSight authors can now configure a unique key column to a QuickSight dataset during data preparation. For more information, see [Adding a unique key to an Amazon QuickSight dataset.](#)

December 19, 2024

[Amazon QuickSight scenarios](#)

Amazon QuickSight Pro users can now use Amazon Q in Quick Suite to create scenarios to analyze complex business problems with simple natural language. For more information, see [Working with scenarios in Amazon QuickSight.](#)

December 4, 2024

[Augmenting Amazon QuickSight insights with Amazon Q Business](#)

Amazon QuickSight account admins can now connect their QuickSight account to Amazon Q Business to augment insights with unstructured data sources. For more information, see [Augmenting Amazon QuickSight insights with Amazon Q Business](#).

December 3, 2024

[Prompted reports](#)

Amazon QuickSight authors can now create prompts on pixel-perfect reports to allow dashboard users to filter data in on-demand and scheduled reports. For more information, see [Setting up prompts for paginated reports](#).

November 22, 2024

[New visual type: layer map](#)

Amazon QuickSight now supports layer map visuals. For more information, see [Using layer maps](#).

November 22, 2024

[New visual type: image component](#)

You can now add images to QuickSight dashboards and analyses with image component visuals in Amazon QuickSight. For more information, see [Using image components](#).

November 22, 2024

New visual type: Highcharts visual	Amazon QuickSight authors can now configure Highcharts visuals in QuickSight. For more information, see Using Highcharts .	November 22, 2024
Import visual public preview	(Public preview) Amazon QuickSight authors can now import QuickSight visuals from one dashboard or analysis to a new analysis that has access permissions. For more information, see Importing Amazon QuickSight visuals to an analysis .	November 22, 2024
Font customization	Amazon QuickSight dashboard viewers can now schedule their own dashboard reports for themselves. For more information, see Creating a reader generated report in Amazon QuickSight .	November 22, 2024
Font customization	Amazon QuickSight authors can now customize the font, font size, and font color of visual titles, subtitles, legends, headers, and cells. For more information, see Titles and subtitles on visual types in QuickSight .	November 22, 2024

[Cross-Region inference with Amazon Q in Quick Suite](#)

Amazon Q in Quick Suite now supports cross-Region inferencing. For more information, see [Cross-Region inference with Amazon Q in Quick Suite](#).

November 21, 2024

[Create a custom brand in Amazon QuickSight](#)

Amazon QuickSight account admins can now customize their application's branding and visual theme to align with their organization's guidelines. For more information, see [Amazon QuickSight brand customization](#).

November 14, 2024

[Connect to Starburst with OAuth client credentials](#)

You can now use OAuth client credentials to connect your Amazon QuickSight account with Starburst through the QuickSight APIs. For more information, see [Creating an Amazon QuickSight data source connection to Starburst with OAuth client credentials](#).

November 8, 2024

[Connect to Snowflake with OAuth client credentials](#)

You can now use OAuth client credentials to connect your Amazon QuickSight account with Snowflake through the QuickSight APIs. For more information, see [Creating an Amazon QuickSight data source connection to Snowflake with OAuth client credentials](#).

November 8, 2024

[Subfolders for restricted shared folders](#)

Assets that are located in a QuickSight restricted folder can now be moved within the restricted folder tree into one or more subfolders. For more information, see [Organizing assets into folders for Amazon QuickSight](#).

October 15, 2024

[Show an apply button](#)

QuickSight authors can now add an **Apply** button to a Calendar, Relative date picker, or Multiselect dropdown control that delays visual reload until the user chooses **Apply**. For more information, see [Adding filter controls to analysis sheets](#).

September 27, 2024

[Data story personalization](#)

Amazon Q in Quick Suite now leverages user location and job-related information from your IAM Identity Center instance to generate personalized data stories that are more relevant to authors and readers. For more information, see [Personalize data stories in Amazon QuickSight](#).

September 27, 2024

[Shared view for embedded dashboards](#)

QuickSight developers can now use the Amazon QuickSight Embedding SDK (version 2.8.0 and higher) to allow readers of embedded dashboards to receive and distribute shareable links to their view of an embedded dashboard. For more information, see [Share an embedded view](#).

August 23, 2024

[5 new regions](#)

Amazon Q in Quick Suite is now available in 5 new opt-in Regions: Asia Pacific (Mumbai), Canada (Central), Europe (Ireland), Europe (London), and South America (São Paulo). For more information, see [Supported Amazon Regions for Amazon Q in Quick Suite](#).

August 9, 2024

[Nested filters](#)

You can now apply nested filters to QuickSight analyses. For more information, see [Adding nested filters](#).

August 1, 2024

[Joined datasets in SPICE](#)

The size limit for secondary tables of joins for datasets that are stored in SPICE have been updated. For more information, see [Joining data](#).

July 10, 2024

[Advanced options for paginated reports](#)

You can now create repeating sections for paginated reports in Amazon QuickSight. For more information, see [Create repeating sections](#).

June 27, 2024

[4 New Regions](#)

QuickSight is now available in 4 new opt-in Regions: Europe (Milan), Europe (Zurich), Africa (Cape Town), and Asia Pacific (Jakarta). For more information, see [Amazon Web Services Regions, web sites, IP address ranges, and endpoints](#).

May 28, 2024

[New sorting options for tables in Amazon QuickSight](#)

QuickSight Authors can now configure single column, multi column, and off-visual sorts for tables in Amazon QuickSight. For more information, see [Sorting tables](#).

May 20, 2024

[Key management in QuickSight: Updated IAM policy conditions](#)

The required IAM policy permissions that grant access to Amazon KMS keys in the QuickSight console and the QuickSight APIs have been updated. For more information, see [IAM identity-based policies for Amazon QuickSight: using the admin key management console](#).

May 16, 2024

[???](#)

QuickSight Admins can now update and manage CMKs that are registered to a QuickSight account with the QuickSight APIs. For more information about creating and managing CMKs with the QuickSight APIs, For more information, see [Key management operations](#).

May 16, 2024

[Hide SPICE capacity labels](#)

QuickSight Admins can now choose to hide the account-wide SPICE capacity usage and remaining size labels from QuickSight authors. For more information, see [Hiding SPICE capacity labels](#).

May 2, 2024

[Embed the Generative Q&A experience](#)

You can now set up an embedded Q&A that uses enhanced NLQ capabilities powered by LLMs. The Q&A is the recommended replacement for the embedded Q Search Bar and provides an updated BI experience for users. For more information, see [Embed the Amazon Q in Quick Suite Generative Q&A experience](#).

April 30, 2024

[Amazon Q in Quick Suite: New Pro roles](#)

Amazon Q integrates with Amazon QuickSight to give QuickSight users access to a suite of new generative BI capabilities. With Amazon Q in Quick Suite, you can utilize the generative BI authoring experience, create executive summaries of your data, ask and answer questions of data, and generate data stories. To get started with Amazon Q in QuickSight Generative BI capabilities, upgrade your account's users to Admin Pro, Author Pro, or Reader Pro roles. For more information, see [Get started with Generative BI](#).

April 30, 2024

[VPC endpoint restrictions](#)

You can now create VPC endpoint restrictions in Amazon QuickSight to limit access to your QuickSight account. For more information, see [Turning on IP and VPC endpoint restrictions in Amazon QuickSight](#).

April 9, 2024

[Account instances of IAM Identity Center for account management](#)

Amazon QuickSight now supports account instances of IAM Identity Center for identity management. For more information, see [Managing user access inside Amazon QuickSight](#).

April 9, 2024

Managed policy update	Amazon QuickSight added a new managed policy to simplify Asset Bundle export operations. For more information, see Amazon managed policy: AWSQuickSightAssetBundleExportPolicy .	March 27, 2024
Managed policy update	Amazon QuickSight added a new managed policy to simplify Asset Bundle import operations. For more information, see Amazon managed policy: AWSQuickSightAssetBundleImportPolicy .	March 27, 2024
New Q pane	Amazon QuickSight users that are opted in to the Generative BI preview can now use the Q pane to access all relevant Generative BI features. For more information, see Using Generative BI with Amazon Q in Quick Suite .	March 25, 2024
New Region: China (north)	Amazon QuickSight is now available in China (Beijing) (cn-north-1). For more information, see Amazon Web Services Regions, web sites, IP address ranges, and endpoints .	January 25, 2024

[Amazon Q Business in Amazon QuickSight](#)

Amazon Q Business integrates with Amazon QuickSight to give QuickSight users access to a suite of new generative BI capabilities. With Amazon Q in Quick Suite, you can utilize the generative BI authoring experience, create executive summaries of your data, ask and answer questions of data, and generate data stories. For more information, see [Using generative BI with Amazon Q in Quick Suite](#).

November 28, 2023

[Trusted identity propagation in Amazon QuickSight](#)

QuickSight authors can now create data sources with trusted identity propagation, which allows database administrators to centrally manage data security in Amazon Redshift and automatically apply all data security rules to data consumers in QuickSight. For more information, see [Enabling trusted identity propagation with Amazon Redshift](#).

November 26, 2023

[Runtime filtering and theming for embedded dashboards and visuals](#)

QuickSight authors can now customize filters and themes for embedded dashboards and visuals. For more information, see [Customizing embedded assets](#).

November 17, 2023

[Custom time zones and week start days in Amazon QuickSight](#)

QuickSight authors can now restrict the functionality that people can access in Amazon QuickSight. For more information, see [Customizing access to the Amazon QuickSight console](#).

November 17, 2023

[Custom time zones and week start days in Amazon QuickSight](#)

QuickSight authors can now turn on SPICE auto capacity purchasing to allow Amazon QuickSight to automatically manage your QuickSight account's SPICE capacity. For more information, see [Turning SPICE auto capacity purchasing on](#).

November 17, 2023

[Amazon EventBridge integration](#)

You can now integrate your Amazon QuickSight account with Amazon EventBridge. For more information, see [Amazon QuickSight EventBridge integration](#).

November 17, 2023

[New look for the QuickSight analysis workspace](#)

We've redesigned the Amazon QuickSight analysis workspace. You might encounter screenshots or procedural text that doesn't reflect the new look in the QuickSight console. We're in the process of updating screenshots and procedural text. For more information on QuickSight's new look, see [Introducing new analysis experience on Amazon QuickSight](#).

November 16, 2023

[Custom time zones and week start days in Amazon QuickSight](#)

QuickSight authors can now set custom time zones and week start days for their analyses. For more information, see [Customize date and time values of an analysis](#).

November 14, 2023

[Float decimal type](#)

Decimal values in calculated fields of datasets that are stored in SPICE can now be assigned the Fixed or Float decimal data type. For more information, see [Supported data types and values](#).

November 8, 2023

[Custom total values in Amazon QuickSight](#)

QuickSight authors can now define the total and subtotal aggregations for their table or pivot table visuals from the field wells. For more information, see [Custom total values](#).

October 25, 2023

Restricted and unrestricted shared folders in Amazon QuickSight	QuickSight authors can now create restricted and unrestricted shared folders with the QuickSight CLI to help organize and manage the permissions of important QuickSight assets. For more information, see Organizing assets into folders for Amazon QuickSight .	October 24, 2023
New data source	QuickSight now supports connecting to a Trino data source. For more information, see Using Trino with Amazon QuickSight .	October 20, 2023
New data source	QuickSight now supports connecting to a Starburst data source. For more information, see Using Starburst with Amazon QuickSight .	October 20, 2023
SageMaker AI Canvas integration	QuickSight authors can now export data into SageMaker AI Canvas to build ML models that can be sent back to QuickSight and applied to analyses and dashboards. For more information, see Build predictive models with SageMaker AI Canvas .	October 6, 2023

Managed policy update	The AWSQuickSightSageMakerPolicy is updated to reflect the Amazon QuickSight–SageMaker AI integration. For more information, see Amazon managed policy: AWSQuickSightSageMakerPolicy .	October 6, 2023
Amazon Redshift run-as role	QuickSight authors can now connect Amazon Redshift data with run as role to enhance data security with fine grained access policies. For more information, see Run queries as an IAM role in Amazon QuickSight .	October 6, 2023
Generative BI public preview	The public preview for Generative BI is now available to users who have a QuickSight Q add-on subscription. For more information, see Working with Generative BI in QuickSight Q .	September 28, 2023
Freeze table columns in Amazon QuickSight	You can now freeze single or groups of table columns in Amazon QuickSight. For more information, see Freeze columns to table visuals .	September 26, 2023
Enhanced KPI options in Amazon QuickSight	New layouts and formatting options for KPIs are now available in Amazon QuickSight. For more information, see Using KPIs .	September 15, 2023

Scaled folders in Amazon QuickSight	You can create scaled folders that can be shared with thousands of namespaces at once with the Amazon QuickSight APIs. For more information, see Creating scaled folders with the Amazon CLI .	August 30, 2023
Excel format reports in Amazon QuickSight	You can now create and schedule Excel snapshot reports of a dashboard in Amazon QuickSight. For more information, see Scheduling and sending reports by email .	August 24, 2023
Add contextual row subtitles to pivot tables in Amazon QuickSight	You can now add contextual row subtitles to pivot tables in Amazon QuickSight. For more information, see Totals and subtotals .	August 16, 2023
New layouts for pivot tables in Amazon QuickSight	Amazon QuickSight now supports two types of layout for pivot tables: tabular and hierarchy. For more information, see Choosing a layout .	August 11, 2023
Default column width in Amazon QuickSight	You can now set the default column width of your pivot table visuals in Amazon QuickSight. For more information, see Row and column size in tables and pivot tables in Amazon QuickSight .	August 11, 2023

[Amazon QuickSight integration with IAM Identity Center](#)

You can now integrate your Amazon QuickSight Enterprise edition account with IAM Identity Center. For more information, see [Configure your Amazon QuickSight account with IAM Identity Center](#).

August 11, 2023

[Embedded callback actions at runtime](#)

You can now use embedded data point callback actions to build tighter integrations between your SaaS application and your Amazon QuickSight embedded dashboards and visuals. For more information, see [Add embedded callback actions to at runtime in Amazon QuickSight](#).

August 9, 2023

[Working with field level coloring in Amazon QuickSight](#)

You can now assign specific colors to specific field values across all visuals in a QuickSight analysis or dashboard. For more information, see [Working with field level coloring in Amazon QuickSight](#).

July 13, 2023

[Small multiples axis options in Amazon QuickSight](#)

You can now configure the x and y axes for each individual panel of a small multiples visual in Amazon QuickSight. For more information, see [Small multiples axis options](#).

July 13, 2023

[Use the analysis menu in Amazon QuickSight](#)

You can now use menu options to efficiently perform tasks without needing to manually navigate through your analysis in Amazon QuickSight. For more information, see [The analysis menu](#).

July 7, 2023

[Advanced report scheduling options in Amazon QuickSight](#)

You can now access more advanced PDF report scheduling options in Amazon QuickSight. For more information, see [Configuring email report settings for a dashboard in QuickSight Enterprise edition](#).

June 30, 2023

[Format visual language data based on language settings in Amazon QuickSight](#)

You can now choose how your numeric data values appear in visuals so that they align with the regional language that you have chosen in Amazon QuickSight. For more information, see [Formatting visual numeric data based on language settings in Amazon QuickSight](#).

May 26, 2023

[Create geospatial heatmaps in Amazon QuickSight](#)

You can now create geospatial heatmaps in Amazon QuickSight. For more information, see [Geospatial heatmaps in Amazon QuickSight](#).

May 26, 2023

[Work with Q topics using the QuickSight CLI in Amazon QuickSight](#)

You can now work with QuickSight Q topics using the Amazon QuickSight command line interface (CLI). For more information, see [Work with QuickSight Q topics using the Amazon QuickSight CLI](#).

May 4, 2023

[Use dataset parameters in Amazon QuickSight](#)

You can now use dataset parameters in direct query to dynamically customize their datasets and apply reusable logic to your datasets. For more information, see [Using dataset parameters in Amazon QuickSight](#).

May 4, 2023

[Enhanced options for scatter plots in Amazon QuickSight](#)

You can now plot unaggregated values on scatter plots in Amazon QuickSight. For more information, see [Using scatter plots](#).

May 4, 2023

[Create and update VPC connections Amazon QuickSight](#)

You can now create and update VPC connections in Amazon QuickSight. For more information, see [Using dataset parameters in Amazon QuickSight](#).

May 4, 2023

[Create tag-based RLS rules using the OR condition in Amazon QuickSight](#)

You can now add the OR condition to your tag-based rules to further customize the way data is presented to your QuickSight account users. For more information, see [Using row-level security \(RLS\) with tag-based rules to restrict access to a dataset when embedding dashboards for anonymous users](#).

April 7, 2023

[Threshold alerts on Table and Pivot table visuals in Amazon QuickSight](#)

You can now create threshold alerts for Table and Pivot table visuals in Amazon QuickSight. For more information, see [Working with threshold alerts in Amazon QuickSight](#).

March 17, 2023

[Hide collapsed columns in Amazon QuickSight table visuals](#)

You can now customize the way Amazon QuickSight readers view pivot tables to make them easier to read and understand at a glance. For more information, see [Pivot table layout options](#).

March 9, 2023

[Enhanced embedding capabilities in Amazon QuickSight](#)

You can now use QuickSight's Embedding SDK (v2.0) to improve developer experience when embedding QuickSight in your application or website. For more information, see [Embedding with the QuickSight APIs](#).

March 9, 2023

[Run queries with Amazon S3 data sources in Amazon QuickSight](#)

You can now enhance data security by using fine-grained access policies rather than broader permissions for data sources connected to Amazon S3 in Amazon QuickSight. For more information, see [Run queries with Amazon S3 data sources](#).

February 21, 2023

[Radar charts in Amazon QuickSight](#)

You can now create radar charts in Amazon QuickSight. For more information, see [Using radar charts in Amazon QuickSight](#).

January 30, 2023

[Data bars for tables and pivot tables in Amazon QuickSight](#)

You can now use data bars to add visual context to your table visuals in Amazon QuickSight. For more information, see [Adding data bars to tables in Amazon QuickSight](#).

January 24, 2023

[New question types: Boolean, Forecast, and Why](#)

You can now enter boolean, forecast, and why questions into the QuickSight Q search bar. For more information, see [Asking questions with Amazon QuickSight Q](#).

November 29, 2022

[Automated data prep for QuickSight Q](#)

Amazon QuickSight Q now uses AI-enhanced data preparation to automatically create topics that are relevant to your end users. For more information, see [Working with Amazon QuickSight Q topics](#).

November 29, 2022

[Paginated reporting](#)

You can now create, schedule, and share highly formatted multi-page reports in Amazon QuickSight. For more information, see [Working with paginated reports in Amazon QuickSight](#).

November 28, 2022

[Embedding the Q search bar for anonymous \(unregistered\) users](#)

You can now embed the QuickSight Q search bar for anonymous (unregistered) users. For more information, see [Embedding the Amazon QuickSight Q search bar for anonymous \(unregistered\) users](#).

November 19, 2022

[Asset management](#)

You can now manage all of the assets on your Amazon QuickSight account in one unified view. For more information, see [Amazon QuickSight asset management](#).

November 19, 2022

Text boxes	You can now add static and dynamic textual content using text boxes in Amazon QuickSight. For more information, see Using text boxes .	November 18, 2022
Small multiples	You can now create small multiples in line, bar, and pie chart visuals. For more information, see Using small multiples .	November 18, 2022
Set a reserved value in QuickSight datasets	You can now set a reserved value to determine the value of the Select all value of a dataset in Amazon QuickSight. For more information, see Setting up parameters in Amazon QuickSight .	November 18, 2022
Run queries as an IAM role in Amazon QuickSight	You can now enhance data security by using fine-grained access policies rather than broader permissions for data sources connected to Athena. For more information, see Run queries as an IAM role in Amazon QuickSight .	November 18, 2022
Pin Amazon QuickSight Q answers to your pinboard	You can now pin visuals from Q for easy access to frequently asked questions. For more information, see Pinning visuals in Amazon QuickSight Q .	November 18, 2022

[Marker clustering on geospatial map visuals](#)

You can now use marker clustering to improve readability of co-located points on a map. For more information, see [Marker clustering on geospatial point maps in QuickSight](#).

November 18, 2022

[Line chart customization](#)

You can now add customizations to line chart visuals to emphasize what you want readers to focus on: color, line style, and markers. For more information, see [Line and marker styling on line charts in QuickSight](#).

November 18, 2022

[Connecting Amazon QuickSight accounts to Databricks](#)

You can now use Amazon QuickSight to connect to Databricks on Amazon. For more information, see [Using Databricks in QuickSight](#).

November 18, 2022

[Account termination protection](#)

You can now use account termination protection to prevent someone from deleting a QuickSight user account accidentally or maliciously. For more information, see [Deleting your Amazon QuickSight subscription and closing your account](#).

November 18, 2022

[1 billion row SPICE datasets](#)

Amazon QuickSight now supports SPICE datasets that contain up to 1 billion (1,000,000,000) rows or 1 terabyte (TB) of data. For more information, see [SPICE quotas for imported data](#).

November 18, 2022

[Monitor SPICE consumption metrics in Amazon CloudWatch](#)

You can now monitor your QuickSight account's SPICE consumption metrics in Amazon CloudWatch. For more information, see [Aggregate SPICE metrics](#).

November 8, 2022

[Use Amazon KMS to encrypt your SPICE datasets in Amazon QuickSight](#)

You can now encrypt your SPICE datasets using the keys that you have stored in Amazon KMS. This provides you with the tools to audit access to data and satisfy regulatory security requirements. For more information, see [Key management](#).

October 27, 2022

[Add datasets that contain row-level security \(RLS\) to Q topics](#)

QuickSight Q now supports questions for access restricted datasets that use Row level Security (RLS) with user based rules. Readers can now ask questions about Topics that contain restricted access datasets and instantly receive accurate and pertinent answers based on access control rules defined by authors in RLS settings. For more information, see [Adding datasets with row-level security \(RLS\) to a topic.](#)

October 10, 2022

[Use Amazon Secrets Manager secrets instead of database credentials in Amazon QuickSight](#)

Amazon QuickSight administrators can grant QuickSight read-only access to secrets they create in Amazon Secrets Manager. These secrets can be used in place of database credentials when creating and editing data sources using the QuickSight API. For more information, see [Using Amazon Secrets Manager secrets instead of database credentials in Amazon QuickSight.](#)

October 6, 2022

[Visual Embedding](#)

You can now embed visuals using a 1-click embed code or with the Amazon QuickSight API. For more information, see [Embedding visuals with the Amazon QuickSight API](#).

August 25, 2022

[Undo and redo changes in Q answers](#)

You can now undo or redo any changes you make to a Q answer by choosing the undo or redo arrows in the Q search bar. For more information, see step six in [Asking questions with Q](#).

August 22, 2022

[Updates to the sign in experience](#)

The sign in experience in QuickSight has been updated. For more information about how to sign in, see [Signing in to QuickSight](#).

July 21, 2022

[Bookmarks](#)

You can now bookmark views of a dashboard. For more information, see [Bookmarking views of a dashboard](#).

July 21, 2022

[Base Maps](#)

When creating map visuals in QuickSight, you can now change map base. For more information, see [Changing base maps](#).

July 21, 2022

[Q updates](#)

You can now try QuickSight Q out for free before getting the Q add-on. For more information, see [Trying QuickSight Q](#).

July 7, 2022

[Rolling dates in filters and parameters](#)

You can now set a rolling date on date range filters and relative date parameters in analyses. For more information, see [Creating range date filters in analyses](#) and [Setting up parameters](#).

July 5, 2022

[Calculations update](#)

You can now create level-aware calculations in QuickSight. You can specify at what level to group computations using level-aware calculation - aggregation (LAC-A) functions. You can also specify the window or partition to compute a calculation using level-aware calculation - window (LAC-W) functions (previously known as level-aware aggregations). For more information, see [Using level-aware calculations](#).

July 5, 2022

[Allowed domains](#)

Developers can now add a domain at run-time to an allow list with the `AllowedDomains` parameter of a `GenerateEmbedUrlForAnonymousUser` or a `GenerateEmbedUrlForRegisteredUser` API call. For more information, see [Allow listing domains at run-time with the QuickSight API](#).

July 5, 2022

[Account provisioning](#)

You can now sign up for an Amazon QuickSight account with the Amazon QuickSight API. For more information, see [CreateAccountSubscription](#).

July 5, 2022

[Monitoring data with Amazon CloudWatch](#)

You can now monitor metrics for Amazon QuickSight dashboards, visuals, and dataset ingestions on Amazon CloudWatch to provide your readers with a consistent, high-performing, and uninterrupted experience on QuickSight. For more information, see [Monitoring data in Amazon QuickSight](#).

June 22, 2022

[Forecast boundaries](#)

Authors can now set a minimum and maximum for forecasted values. For more information, see [Forecasting and creating what-if scenarios](#).

June 16, 2022

[Table and pivot table enhancements](#)

Authors and Readers can now resize rows and columns in a table or pivot table visual. For more information, see [Resizing rows and columns in tables and pivot tables](#).

June 14, 2022

Hiding fields in pivot tables	You can now hide fields in pivot table visuals. For more information, see Showing and hiding pivot table columns .	June 14, 2022
ML insights and forecasting for Q	Readers can now see key data insights and add forecasts to Q answers. For more information, see Asking questions with QuickSight Q .	June 2, 2022
What's New in QuickSight	To find out what's new in QuickSight, see the What's New page. What's New posts provide a brief overview of all QuickSight feature announcements as they are released.	June 1, 2022
Line chart updates	QuickSight now supports 10,000 data points for line charts. For more information, see Using line charts .	May 22, 2022
1-click public embedding	You can now embed dashboards for anyone on the internet in your public applications with a 1-click static embed code. For more information, see Enabling public access to dashboards with a 1-click embed code .	May 18, 2022

Creating and managing groups	You can now create and manage groups in the QuickSight console. For more information, see Creating and managing groups in Amazon QuickSight .	March 21, 2022
Automatic refresh controls for direct query	In QuickSight, filter controls are now refreshed every 24 hours for direct query. For more information, see Refreshing data .	March 8, 2022
Updating files in a dataset	You can now update files in a dataset to get the latest version of those files. For more information, see Updating files in a dataset .	January 27, 2022
Rich text formatting for titles and subtitles	QuickSight now offers rich text formatting for titles and subtitles, and the ability to add hyperlinks and parameters in titles. For more information, see Formatting a visual title and subtitle .	January 27, 2022
Comparative and cumulative date/time calculations	Comparative and cumulative period functions, such as year-over-year and year-to-date, are now supported in QuickSight. For more information, see Add comparative and cumulative date/time calculations in QuickSight .	January 27, 2022

Sharing dashboards	You now have the option to grant everyone in your QuickSight account access to a dashboard. You can also share a link to the dashboard with anyone who has access to it. For more information, see Sharing dashboards .	November 23, 2021
Customizing email report templates	You can now customize how dashboard email reports appear and behave for QuickSight account users. For more information, see Customizing email report templates .	November 23, 2021
1-click enterprise embedding	You can now embed dashboards for registered users in your internal applications with a 1-click static embed code. For more information, see Embedding dashboards for registered users with a 1-Click embed code .	November 23, 2021
New data source	QuickSight now supports connecting to an Exasol data source. For more information, see Supported data sources .	November 22, 2021
Dataset versioning	QuickSight now supports dataset versioning. For more information, see Reverting datasets back to previous published versions .	November 22, 2021

[Adding datasets with RLS to Q topics](#)

QuickSight Q topics now support datasets with row-level security (RLS). For more information, see [Adding datasets with RLS to a Q topic](#).

November 19, 2021

[Sheet change performance operations](#)

QuickSight now only refreshes visuals when switching sheets if required. For more information, see [Refreshing visuals](#).

November 12, 2021

[Incrementally refresh a SQL-based dataset](#)

In QuickSight Enterprise Edition, you can now refresh your SQL-based SPICE datasets incrementally within a look-back window of time. For more information, see [Refreshing data](#).

October 25, 2021

[Embedding the Q search bar](#)

You can now embed the Q search bar in your application for registered users of QuickSight. For more information, see [Embedding the QuickSight Q search bar for registered users](#).

October 22, 2021

[Email syncing for federated users](#)

In Enterprise Edition, QuickSight administrators can now restrict new users from using personal email addresses when provisioning through their identity provider (IdP) to QuickSight. For more information, see [Configuring email syncing for federated users](#).

October 22, 2021

[Adding RLS tags to a dataset](#)

You can now add tag-based rules to a dataset in QuickSight when applying row-level security (RLS). For more information, see [Using row-level security \(RLS\) with tag-based rules to restrict access to a dataset when embedding dashboards for anonymous users](#).

October 19, 2021

[IP restrictions](#)

You can now limit access to your organization's QuickSight account to a predefined list of Internet Protocol (IP) ranges. For more information, see [Turning on Internet Protocol \(IP\) restrictions in Amazon QuickSight](#).

October 18, 2021

[New table and pivot table styling options](#)

You can now customize the look of tables and pivot tables. For more information, see [Formatting tables and pivot tables](#).

October 12, 2021

New SPICE quota for imported data	In QuickSight Enterprise edition, you can now import up to 500 million rows into a SPICE dataset. For more information, see Data source quotas .	October 12, 2021
Links and images in tables	You can now add links to tables. You can also render links as images in tables. For more information, see Field styling .	October 12, 2021
Custom date formats	You can now customize how dates are formatted in your filter and parameter controls. For more information, see Customizing date formats .	October 8, 2021
QuickSight q	You can now ask questions about your data in the Q bar and get answers in the form of visuals. For more information, see Working with QuickSight q .	September 24, 2021
New geospatial region	QuickSight now supports creating geospatial charts for India. For more information about creating geospatial charts, see Using geospatial charts (maps) .	September 22, 2021

Free-form layouts	You can now place visuals anywhere in your dashboard using precise coordinates. For more information, see Customizing dashboards and visuals .	September 22, 2021
Passing IAM roles to QuickSight	You can now pass IAM roles to QuickSight. For more information, see Using QuickSight with IAM .	September 15, 2021
Create datasets from existing datasets	You can now create datasets from existing datasets. For more information, see Creating a dataset using an existing dataset .	September 15, 2021
Updated Amazon managed policy	QuickSight updated an existing Amazon managed policy. For more information, see Amazon QuickSight updates to Amazon managed policies .	September 8, 2021
New Amazon managed policy	QuickSight added a new Amazon managed policy. For more information, see Amazon QuickSight updates to Amazon managed policies .	September 8, 2021
Forecast multiple measures	You can now forecast up to three measures in the same chart. For more information, see Forecasting and creating what-if scenarios .	August 30, 2021

[Share your view of a dashboard](#)

You can now share your view of a dashboard. For more information, see [Sharing dashboards](#).

August 20, 2021

[Hide the data zoom on a chart](#)

You can now choose to show or hide the data zoom for your charts in Amazon QuickSight. For more information, see [Formatting axes and Grid lines](#).

August 16, 2021

[Row-level security with tags](#)

You can now use row-level security (RLS) with tags to specify which data your (unauthenticated) users can see in an embedded Amazon QuickSight dashboard depending on who they are. For more information, see [Using RLS with tags to restrict access to a dataset when embedding dashboards for anonymous users](#).

July 29, 2021

[New embedding API operations](#)

Amazon QuickSight has the following new API operations for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` API operations to embed dashboards and the QuickSight console, but they don't contain the latest embedding capabilities. For more information about using the new API operations, see [Embedding overview](#). For more information on using the old API operations, see [Embedding analytics using the old API operations](#).

July 29, 2021

[Customizing charts](#)

You can now hide the sort icon and the axis or field title in charts in Amazon QuickSight. For more information, see [Formatting axes and Grid lines](#).

July 16, 2021

[Threshold alert updates](#)

You can now choose to be notified when there is no data to check your alert rule against in Amazon QuickSight. For more information, see [Working with threshold alerts](#).

July 9, 2021

New chart type: Custom visual content	You can now embed webpages and online videos, forms, and images in your QuickSight dashboards using the custom visual content chart type. For more information, see Using custom visual content in QuickSight .	June 29, 2021
Feature improvements	You can now show totals for stacked bar charts. For more information, see Using bar charts .	June 29, 2021
Threshold alert updates	Dataset owners can now set a threshold alert evaluation schedule for a dataset in Amazon QuickSight. For more information, see Working with threshold alerts .	June 17, 2021
Duplicating analyses	You can now duplicate an analysis in Amazon QuickSight. For more information, see Duplicating an analysis .	June 17, 2021
Threshold alert updates	You can now choose how often you want to receive threshold alerts in Amazon QuickSight. For more information, see Working with threshold alerts .	May 25, 2021
Wildcard search	Amazon QuickSight now supports wildcard search. For more information, see Using wildcard search .	May 1, 2021

[Threshold alerts](#)

Amazon QuickSight now supports creating threshold alerts using KPI and Gauge visuals in dashboards. For more information, see [Working with threshold alerts](#).

May 1, 2021

[Service control policies](#)

Amazon QuickSight now supports using service control policies (SCPs) to restrict how individuals in your organization can sign up for QuickSight. For more information, see [Using service control policies to restrict QuickSight signup options](#).

April 26, 2021

[Parameter improvements](#)

Amazon QuickSight now supports dynamically displaying parameter values in titles and descriptions throughout charts and analyses. For more information, see [Using parameters in titles and descriptions](#).

April 12, 2021

[Feature improvements](#)

QuickSight now supports selecting a time granularity for date and time range filters and datetime parameters. For more information, see [Adding a date filter](#) and [Setting up parameters](#).

April 1, 2021

[Export to PDF](#)

QuickSight now supports PDF exports of the current sheet in dashboards and analyses. In Enterprise edition, you can now attach a PDF to email reports. For more information, see [Exporting an analysis or dashboard as a PDF](#).

April 1, 2021

[Seconds support for date fields](#)

QuickSight now supports aggregating date fields at the second level. For more information, see [Changing date field granularity](#).

March 30, 2021

[Pivot table sorting improvements](#)

QuickSight now supports sorting values in pivots tables by fields in the **Rows** and **Columns** field wells or by column headers. For more information, see [Sorting pivot tables](#).

March 30, 2021

[New feature](#)

Amazon QuickSight now supports customizing tooltips in visuals. For more information, see [Customizing tooltips in a visual](#).

March 16, 2021

Feature improvements

In Enterprise edition, QuickSight now supports computing anomalies for a time and measure field only, and computing the exact combination of fields in the Category field well. For more information, see [Adding an ML insight to detect outliers and key drivers](#).

March 16, 2021

Line chart improvements

Amazon QuickSight now supports creating line charts with dual axes. For more information, see [Creating a dual-axis line chart](#).

March 1, 2021

New formatting options for charts

Amazon QuickSight now supports hiding axis lines, axis labels, and grid lines, and also customizing the how axis labels appear in a chart. For more information, see [Formatting axis lines, axis labels, and Grid lines](#).

February 5, 2021

New chart type: Sankey diagrams

Amazon QuickSight now supports creating Sankey diagrams. For more information, see [Using sankey diagrams](#).

February 5, 2021

Field mapping improvements	Amazon QuickSight now supports updating field mapping between mismatched and missing fields when replacing a dataset in an analysis. For more information, see Replacing datasets .	February 5, 2021
Data prep enhancements	Amazon QuickSight now supports organizing fields into folders and subfolders when preparing or editing datasets. For more information see, Organizing fields into folders .	February 5, 2021
New features	In Enterprise edition, QuickSight now supports displaying anomaly ranges and multiple anomalies when exploring anomalies. Line charts now display notifications when QuickSight detects an anomaly, key driver, or forecasting opportunity. For more information, see Detecting outliers with ML-powered anomaly detection .	February 3, 2021
New region: South America (São Paulo)	Amazon QuickSight is now available in South America (São Paulo) (sa-east-1). For more information, see Amazon Web Services Regions, websites, IP address ranges, and endpoints .	December 23, 2020

New region: Canada (central)	Amazon QuickSight is now available in Canada (Central) (ca-central-1). For more information, see Amazon Web Services Regions, websites, IP address ranges, and endpoints .	December 23, 2020
Define an action to open a different sheet	QuickSight now supports navigation actions to enable you to open a different sheet with the included parameter values. For more information, see Using custom actions for filtering and navigating .	December 21, 2020
Data source enhancement: Athena federated query	QuickSight now supports connecting to Athena Federated Query. For more information, see Creating a dataset using Amazon Athena data .	December 21, 2020
New list items sheet control type	Amazon QuickSight now supports single and multi-select list control on dashboards. For more information, see Parameter controls .	December 18, 2020
New formatting options for tables	In QuickSight, you can now reorder columns in table charts. For more information, see Using tables as visuals .	December 17, 2020

[Row-level security \(RLS\) now supports GRANT option only](#)

To streamline and simplify RLS setup, QuickSight no longer supports using RLS in deny-access mode. To create a new RLS configurations, use the explicit grant-access model. Current RLS datasets and deny-access configurations should continue to work as expected. For more information, see [Using row-level security \(RLS\) to restrict access to a dataset](#).

December 6, 2020

[New percentile functions](#)

QuickSight now supports variations on percentile calculations, including aggregate functions `percentileCont` and `percentileDisc` and `OVER` functions `percentileContOver` and `percentileDiscOver`. You can use all of these functions in the calculations editor in analysis mode. For more information, see [Functions by category](#).

December 6, 2020

[Dynamic queries, plus dataset editor improvements](#)

QuickSight now enables you to improve performance for visuals based on direct queries by specifying unique keys for joined tables. There are also multiple interface improvements, such as settings for optional autopreviews, zoom in and out of data diagrams, join recommendations, and more. For more information, see [Joining data](#).

December 6, 2020

[Developer portal for embedding](#)

The Dev Portal helps you learn how to use embedding by example in your web site or application. For more information, see [Using the dev portal](#).

November 30, 2020

[New data source: Amazon OpenSearch](#)

QuickSight now supports connecting to Amazon OpenSearch. For more information, see [Using Amazon OpenSearch with QuickSight](#).

November 25, 2020

[Format metrics on pie charts](#)

QuickSight now supports formatting metrics on pie charts to display values, percentages, or both. For more information, see [Customizing data labels](#).

November 25, 2020

[New formatting options for tables](#)

In QuickSight, you can now vertically align and wrap the text for headers in table charts. For more information, see [Using tables as visuals](#).

November 24, 2020

[New chart type: Box plots](#)

QuickSight now supports using box plots so you can visualize how your data is distributed across an axis or over time. For more information, see [Box plots](#).

November 24, 2020

[Filled maps](#)

QuickSight now supports filled maps so you can visualize your data over a geographical area. For more information, see [Filled maps](#).

November 24, 2020

[Favorite folders](#)

In QuickSight, you can now favorite your folders for easy access. For more information, see [Organizing assets into folders](#).

November 24, 2020

[Change in display limits](#)

For parameter controls and filters, Amazon QuickSight now displays up to 1,000 sample values. When you have more than that, you use the search box to locate a value. For more information, see [Using a control with a parameter](#).

November 16, 2020

Null rendering	QuickSight now allows null values to be updated to a custom string. For more information, see Changing a field format .	November 9, 2020
New chart type: Waterfall charts	QuickSight now supports using waterfall charts so you can view your data sequentially. For more information, see Waterfall charts .	November 9, 2020
Embedding for everyone	QuickSight Enterprise edition now supports embedding for everyone. When you purchase session packs with capacity pricing, you can enable visitors to use embedded dashboards without registering them as QuickSight users. For more information, see Embedded analytics .	November 9, 2020
Column-level security	QuickSight now supports using column-level security to restrict access to a dataset. For more information, see Using column-level security .	November 9, 2020

Service provider initiated federation	QuickSight Enterprise edition now supports service provider initiated federation, so you can sign on directly to QuickSight using your IAM Identity Center login and password. For more information, see Setting up service provider–Initiated federation with QuickSight enterprise edition .	October 30, 2020
New chart type: Funnel charts	QuickSight now supports using funnel charts so you can display your data in a linear process. For more information, see Funnel charts .	October 29, 2020
Custom sorting	QuickSight now enables you to apply sort options on fields outside the field wells. This feature is available for all charts that support sorting. For more information, see Describing data .	October 29, 2020
New region: Amazon GovCloud (US-West)	Amazon QuickSight is now available in Amazon GovCloud (US-West) (gov-west-1). For more information, see QuickSight - Amazon GovCloud (US) .	October 28, 2020
New data source: Oracle	QuickSight now supports connecting to Oracle. For more information, see Supported data sources .	October 23, 2020

Filter across datasets	In QuickSight, you can now create filters that apply to multiple datasets in an analysis. For more information, see Filtering data .	October 23, 2020
Add field descriptions to datasets	QuickSight now enables you to add column or field descriptions to datasets. This metadata is visible in both datasets and analyses, helping you make your data self-explanatory. For more information, see Describing data .	October 23, 2020
Filter controls on analysis sheets	QuickSight now offers a filter control that you can add to your analysis with a single click. You can put filters beside dashboard visuals and resize them to fit. For more information, see Using filter controls .	October 5, 2020
New data source: Amazon Timestream	QuickSight now supports Amazon Timestream as a data source. For more information, see Using Amazon Timestream data with QuickSight .	October 1, 2020
Export to excel	QuickSight now supports exporting data from pivot tables and table charts to Microsoft Excel (.xlsx) format. For more information, see Exporting data .	September 14, 2020

Duplicate sheets	QuickSight now supports duplicating sheets. For more information, see Working with multiple sheets in an QuickSight analysis .	September 14, 2020
Customize labels of totals and subtotals	In QuickSight, you can now add custom text to display beside totals and subtotals for pivot tables and table charts. For more information, see Table and pivot table formatting options in QuickSight .	September 14, 2020
Skipped row details	When you import data into QuickSight, you now get better error messages and row-by-row information on what caused rows to be skipped during ingestion into SPICE. For more information, see Troubleshooting skipped rows .	September 12, 2020
New configuration option for multivalue lists	QuickSight now supports a configuration option for multivalue lists to start with no values selected. For more information, see Setting up parameters in Amazon QuickSight .	September 11, 2020

Reference lines	QuickSight now supports using reference lines in bar, line, and combo charts. For more information, see Reference lines .	September 3, 2020
New table calculations	QuickSight now supports using window functions <code>firstValue</code> and <code>lastValue</code> in analyses. For more information, see Table calculations .	August 25, 2020
Multivalue default parameters	QuickSight now supports dynamic default values for multi-valued parameters. For more information, see Creating parameter defaults in QuickSight .	August 25, 2020
Last active status for users	QuickSight administrators can now see the last date and time that people accessed QuickSight. For more information, see Viewing user details .	August 25, 2020
Join SQL queries, tables, and files	QuickSight now supports using custom SQL queries in joins with other queries, tables, and files. For more information, see Joining data .	August 25, 2020

[Calculated expressions editor](#)

The enhanced editor adds ease-of-use features to help you more easily create calculated expressions in QuickSight. With the new full-screen editor, you can add functions and values directly from the menu. For more information, see [Adding a calculated field to an analysis](#).

August 25, 2020

[Personal and shared folders](#)

QuickSight Enterprise edition now offers personal and shared folders to make it easier to discover, organize, share, and explore your available assets. For more information, see [Organizing QuickSight assets into folders](#).

August 5, 2020

[QuickSight support for multitenancy](#)

In Enterprise edition, QuickSight now offers an additional layer of security to support multitenancy. By creating users and groups in a QuickSight namespace, you can isolate them from users and groups in other namespaces. For more information, see [Supporting multitenancy with isolated namespaces](#).

July 23, 2020

[Embed the QuickSight console](#)

In QuickSight Enterprise edition, you can now provide the full dashboard authoring experience of the QuickSight console in your own custom-branded authoring portal. For more information, see [Embedding the QuickSight console](#).

July 23, 2020

[Datasets with 2,000 columns](#)

QuickSight now supports datasets that have up to 2,000 columns. For more information, see [Data source limits](#).

July 23, 2020

[Customized permissions](#)

In QuickSight Enterprise edition, you can now create custom roles to restrict a person from using specific functionality in the QuickSight console, for example to control who can manage data sources and datasets or who can manage or subscribe to email reports. For more information, see [Customizing access to the QuickSight console](#).

July 23, 2020

[Customize the console](#)

In QuickSight Enterprise edition, you can now choose whether to show or hide QuickSight startup samples and videos. You can also create and specify a default theme to add branding to your experience of the QuickSight console. For more information, see [Customizing the QuickSight console](#).

July 23, 2020

[Larger SPICE datasets](#)

QuickSight Enterprise edition now supports SPICE datasets with up to 250 million (250,000,000) rows or 500 GB. For more information, see [Data source limits](#).

July 9, 2020

[Display width settings](#)

In QuickSight, you can now choose between responsive and fixed layouts for your analyses and dashboards. You can also temporarily change your current view from the menu bar. For more information, see [Display settings](#).

July 9, 2020

[Data refresh notifications](#)

QuickSight now supports sending email to notify SPICE dataset owners when data refresh fails. For more information, see [Refreshing data](#).

July 9, 2020

[Color settings for heat maps and tree maps](#)

You can now customize gradient colors for your heat maps and tree maps in QuickSight. For more information, see [Changing colors on heat maps and tree maps](#).

July 9, 2020

[Font choice](#)

You can now customize your theme with a font. Choose from a selection of fonts supported by QuickSight. For more information, see [Using themes in QuickSight](#).

July 8, 2020

[Amazon Lake Formation](#)

If you use QuickSight Enterprise edition to query Amazon Athena, you can simplify data access management by using Amazon Lake Formation to control permissions to Athena. For more information, see [Authorizing connections through Amazon Lake Formation](#).

June 29, 2020

[New chart type](#)

QuickSight now supports using histograms so you can display the distribution of values in your data. For more information, see [Histogram charts](#).

June 12, 2020

Gauge chart update

You can now add conditional formatting to gauge charts in QuickSight. For more information, see [Adding conditional formatting to visuals](#).

June 12, 2020

Five new languages

QuickSight now supports Danish, Dutch, Finnish, Norwegian, and Swedish. These languages expand upon the existing 15 languages already available in QuickSight. For more information, see [Choosing a language in Amazon QuickSight](#).

June 12, 2020

New region: Mumbai

Amazon QuickSight is now available in Asia Pacific (Mumbai) (ap-south-1). For more information, see [Amazon Web Services Regions, websites, IP address ranges, and endpoints](#).

June 3, 2020

[New SQL editor](#)

QuickSight enhanced the SQL editor that you use to create datasets from SQL queries. The new SQL editor supports syntax highlighting, basic autocomplete, autoindent, and line numbering. Also, you can use the new schema explorer to interactively explore schemas, tables, fields, and data types. For more information, see [Using a SQL query](#).

May 20, 2020

[Amazon SageMaker AI integration now generally available](#)

The integration of SageMaker AI with QuickSight launched in preview just before Amazon re:Invent 2019. Now, this feature is generally available in all Amazon Web Services Regions that are supported by QuickSight. Amazon QuickSight and SageMaker AI together make it faster, easier, and more cost-effective for customers to make use of their machine learning models for visualization and predictions. For more information, see [Amazon QuickSight integration with SageMaker](#).

May 19, 2020

Update to presto data sources	You can now use nonauthenticated private Presto as a data source. For more information, see Creating a data source using presto .	May 18, 2020
Rewrite of Amazon QuickSight VPC connection section	Based on your feedback, we completely rewrote this documentation to better explain using Amazon VPC with QuickSight. It includes examples and Amazon CLI commands to help you establish the path from Amazon QuickSight to your data source, even if you're using on-premises sources. For more information, see QuickSight VPC connections .	May 18, 2020
Use logarithmic scale with QuickSight	You can now format visuals using log scale. For more information, see Changing the visual scale .	May 1, 2020
Severity levels for QuickSight anomaly detection insights	You can now choose the level of severity for anomalies that display in your insight widgets. For more information, see Using anomaly detection .	April 17, 2020

[QuickSight adds a new visual type](#)

The new visual type, stacked area line charts, is similar to area charts except that the stacked values indicate the relationship that each value contributes to the whole. Using stacked area charts, you can display cumulative totals over time while displaying the breakdown by the chosen category. For more information, see [Line charts](#).

April 17, 2020

[QuickSight activates legends by adding menus for quick feature access](#)

QuickSight now supports clicking on legend items to quickly access features like filtering, drilling, chart colors, and custom actions for filtering and URLs. For more information, see [Focusing on visual elements](#) and [Custom actions in QuickSight](#).

April 17, 2020

[Improvements to visual controls in QuickSight](#)

The on-visual menu can now float at the upper-right corner of the visual, moving inside or outside the border of the visual. Plus, you can make your visuals smaller and have more of them, as many as 36 across. For smaller visuals, the new visual controls have more viewing area for charts and insights. To see the new menu, choose one of your visuals. For more information, see [Working with visuals](#).

April 17, 2020

[QuickSight adds cascading filters](#)

You can now create cascading filters by adding them to custom actions in your analyses and dashboards. Each custom filter action can target one or more visuals in the same sheet, creating a cascading filter effect. For more information, see [Custom actions in QuickSight](#).

April 2, 2020

[Modulo operation available in Amazon QuickSight](#)

You can now use modulo operation to find the remainder after dividing one number into another. For more information, see [mod](#).

March 8, 2020

[Images in Amazon QuickSight narratives](#)

You can now insert images from URLs to be rendered as part of a narrative. Images can be resized, conditionally shown within an IF block, and also hyperlinked with the URL feature. For more information, see [Expression editor screen and menus](#).

March 8, 2020

[Amazon QuickSight adds minimum and maximum dates](#)

You can now use minimum and maximum date aggregations in tables and pivot tables. For more information, see [max](#) and [min](#).

March 8, 2020

[Now you can customize minimum delta values to enhance QuickSight ML-powered anomaly detection](#)

Now you can control the minimum delta value that QuickSight uses to detect anomalies (also known as outliers). You can change this setting when you are configuring or exploring anomalies. For more information, see [Detecting outliers with ML-powered anomaly detection](#).

February 4, 2020

[Improvements in the QuickSight narrative expression editor](#)

QuickSight now provides a much more spacious interface for editing narrative expressions. The interface now has instant feedback in a preview of your evaluated narrative before you save any changes. Plus, narratives now support both static and dynamic URLs in your custom narratives for links to searches, apps, other dashboards, other sheets in the same dashboard, and more. For more information, see [Using autonarratives](#).

February 4, 2020

[SageMaker AI integration now available in more Amazon Web Services Regions](#)

The public preview for integrating SageMaker AI with QuickSight is now available in all of the Amazon Regions that are supported by QuickSight. For more information, see [Amazon QuickSight integration with SageMaker](#).

January 10, 2020

[QuickSight now enables you to choose an Amazon Athena workgroup](#)

By choosing an Athena workgroup, you can better manage your Athena datasets. Using this option, you can also allocate Athena costs to the workgroup that QuickSight is using, for better cost reporting. For more information, see [Creating a dataset using Amazon Athena](#).

December 10, 2019

[Private VPC support for presto in QuickSight](#)

Now Presto can use a VPC connection. For more information, see [Supported data sources for VPC](#).

December 10, 2019

[New mathematical functions in QuickSight](#)

Amazon QuickSight now supports functions for base 10 logarithms (log), natural logarithm (ln), absolute value (abs), square root (sqrt), and base of natural log e raised to the power of (exp). For more information, see [Functions and operators](#).

December 10, 2019

[New LAA functions in QuickSight](#)

Now you can use rank, denseRank , and percentileRank for level-aware aggregations. For more information, see [Using level-aware aggregations](#).

December 10, 2019

[Integrate SageMaker AI ML models into your QuickSight analyses and dashboards](#)

In this public preview, QuickSight launches the ability to integrate machine learning (ML) models created and trained in SageMaker AI. This new feature makes it easier to augment your business data with ML predictions. Add your data scientists' prebuilt inferences and predictions to your analysts' dashboards, and let the collaboration and decision-making begin. For more information, see [Amazon QuickSight integration with SageMaker](#).

November 26, 2019

[New formatting options for visuals in QuickSight](#)

Amazon QuickSight launches more formatting for visuals and visuals titles, including font sizes from extra small to extra large. In key performance indicators (KPIs), you can set font sizes for primary and comparison values. In pivot tables and tables, you can see font sizes for table headers, cells, totals, and subtotals. Now your visuals render better even when smaller and with less padding between charts. For combo charts, you can now synchronize the Y-axes for both bars and lines into a single axis. For more information, see [Formatting a visual](#).

November 23, 2019

[QuickSight supports the now function in SPICE](#)

For more information, see [now](#).

November 22, 2019

[QuickSight now supports seconds in SPICE](#)

Seconds are now supported in SPICE datasets. Datetime fields are no longer truncated to minutes. Now you can use the SS option for the period parameter in date functions including `addDateTime`, `dateDiff`, `extract`, and `truncDate`. For more information, see [Date functions](#).

November 22, 2019

[QuickSight now has themes](#)

You can now create a collection of themes, and apply a theme to an analysis and all its dashboards. For more information, see [Using themes in QuickSight](#).

November 22, 2019

[New QuickSight API operations](#)

Amazon QuickSight launches new API operations to programmatically manage your data, dashboards, and fine-grained access control capabilities linked with Amazon Identity and Access Management (IAM). With new data API operations, you can create, update, and delete Amazon Quick Sight data sources and datasets programmatically. You can also manage data refreshes on your SPICE datasets with API operations. Amazon QuickSight also introduces templates, which store the visual configuration and data schemas required for a dashboard. You can transfer templates across accounts or use them to instantiate dashboards with the same visual presentation but different data. For more information, see the [Amazon QuickSight API reference](#).

November 22, 2019

[QuickSight now supports conditional formatting for tables, pivot tables, and key performance indicators \(KPIs\).](#)

For tables and pivot tables, you can set multiple conditions for fields or supported aggregations, along with format options to apply to a target cell. For KPIs, you can format the primary value based on conditions that are applied to any dimension in the dataset. The conditional formatting options now supported are text color, background color, and placement of supported icons. You can use icons from the provided set, or you can use Unicode icons instead. For more information, see [Adding conditional formatting to visuals](#).

November 18, 2019

[View history for SPICE ingestion on your Amazon QuickSight datasets](#)

You can now view the ingestion history for SPICE datasets in Amazon QuickSight. See information like when the latest ingestion started, how long it took, and what its status is. For more information, see [View SPICE ingestion history](#).

November 7, 2019

[Add your own functionality to visuals with Amazon QuickSight actions](#)

Amazon QuickSight enables you to add to the basic functionality for visuals by creating your own custom actions for filtering or opening URLs. For more information, see [Custom actions in Amazon QuickSight](#).

November 7, 2019

[Amazon QuickSight now supports dynamic sheet titles](#)

You can now use parameters in sheet titles to make the context of a dashboard clearer to the reader. For more information, see [Formatting a visual](#).

November 6, 2019

[Use Amazon QuickSight to join data from multiple data sources](#)

Amazon QuickSight now supports creating datasets that join multiple data sources. For more information, see [Joining data](#).

November 5, 2019

[Amazon QuickSight mobile iOS update and a new Android app](#)

Get access to insights from your data on the fly. Download the updated iOS app or the new Android app. You can browse, favorite, and interact with your dashboards and explore your data with drilldowns and filters. You can stay ahead of the curve by using forecasting. You can get email alerts when unexpected changes happen in your data, and you can share those insights with colleagues. For more information, see [Amazon QuickSight mobile](#).

November 5, 2019

[Amazon QuickSight supports printing](#)

Now you can print a dashboard or an analysis. For more information, see [Printing](#).

October 17, 2019

[Amazon QuickSight supports a new Amazon Web Services Region](#)

Amazon QuickSight is now available in Asia Pacific (Seoul) (ap-northeast-2). For more information, see [Amazon Web Services Regions, websites, IP address ranges, and endpoints](#).

October 17, 2019

[Amazon QuickSight now supports visual-level formatting](#)

Now, any formatting you apply from the field wells is applied only to the selected visual. For more information, see [Formatting a visual](#).

October 17, 2019

Transpose tables	In table visuals, you can now transpose columns and rows. For more information, see Tables .	October 10, 2019
Shared data sources	Share data sources with other users and groups. For more information, see Shared data sources .	October 10, 2019
SPICE enhancements	SPICE dashboards now support new wildcard filters on dimensions. You can now filter your data using any of the available wildcard filters: 'contains', 'starts with', 'end with', 'equals'. Additionally, SPICE dashboards now support new string functions (toString and parseDecimal) and two new date functions (parseDate and formatDate). For more information, see Calculated fields .	October 10, 2019
New publishing options for dashboards	Control your user experience with more publishing options for your dashboards. These options include ability to toggle the filter pane, tooltips, drill up/down, and more. For more information, see Publish a dashboard .	October 10, 2019

New median function	Amazon QuickSight supports median as a new aggregation. For more information, see median .	October 10, 2019
New filter functionality	Amazon QuickSight now supports additional relative date filter options. For more information, see Date filters .	October 10, 2019
New aggregations for fields	You can now use Nth percentile, median, standard deviation, and variance in field wells, filters, and ML Insights. For more information, see Field aggregations .	October 10, 2019
Filter for context	You can now filter text using the following wildcard comparison types: contains, does not contain, begins with, and ends with. For more information, see Text filters .	October 10, 2019
Use the power operator (^) in SPICE	QuickSight now supports using the power operator (^) on SPICE datasets. You can use the power operator with any numeric field, with any valid exponent. For example, "revenue ^ 0.5 " computes the square root of the revenue field. For more information, see Arithmetic and comparison operators .	September 11, 2019

[Use string functions in SPICE](#)

Now you can use string functions for calculated fields in your analyses based on SPICE datasets. For more information, see [Functions by category](#).

September 11, 2019

[Use level-aware aggregations to control your calculations](#)

QuickSight now supports level-aware aggregation calculations. Using this feature, you can perform aggregations at prefilter and preaggregation levels, before aggregations that happen in the display. This helps you explore more advanced query-building strategies than ever before. For more information, see [Level-aware aggregations](#).

September 11, 2019

[Sort your anomalies](#)

Now you can choose how to sort anomalies on the insight widget and in the anomaly exploration page. Choosing how to prioritize anomalies can help you identify the anomalies that are the most important to you. For more information, see [Using ML-powered anomaly detection](#).

September 11, 2019

[Put 100 million rows into a SPICE dataset](#)

We changed the SPICE dataset limit from 25GB to 100 million rows for Enterprise edition (and 25 million rows for Standard edition). For more information, see [SPICE data source limits](#).

September 11, 2019

[Rename your dashboard, keep your settings](#)

Amazon QuickSight now supports renaming your published dashboards in place. You can change the name and still keep your settings and subscribers. For more information, see [Publishing a dashboards](#).

September 6, 2019

[New visual type: Word clouds](#)

Amazon QuickSight supports word cloud visuals that are based on different aggregations over any dimension in your dataset. You can drill down, focus on specific data points, and apply specific colors to data points on the word cloud. For more information, see [Word clouds](#).

September 5, 2019

[Filter out the last \$n\$ time periods](#)

Amazon QuickSight lets you exclude a specific number and type of time periods from a time range (after) filter. For more information, see [Adding a date filter](#).

September 5, 2019

[Amazon QuickSight new feature to mark favorites](#)

Mark your favorite dashboards and analyses so you can come back to them quickly. For more information, see [Using the QuickSight start page](#).

September 5, 2019

[Amazon QuickSight Enterprise edition adds anomaly alerts](#)

Your readers can subscribe themselves to anomaly alerts on dashboards to get the latest ML-powered anomaly detection and contribution analysis by email. For more information, see [Using data dashboards](#).

September 5, 2019

[QuickSight adds custom colors for charts](#)

With the QuickSight enhanced color picker, you can choose custom colors for charts where color customization is supported. The enhanced color picker also retains the eight most recently used custom colors for easy selection across multiple charts. For more information, see [Changing visual colors](#).

August 1, 2019

[Detect anomalies with less data](#)

In Amazon QuickSight Enterprise edition, we lowered the required minimum data points to 15 for training anomaly detection. For more information, see [Dataset requirements for using ML insights with QuickSight](#).

August 1, 2019

[Amazon QuickSight supports new Amazon Web Services Regions](#)

Amazon QuickSight is available in Europe (London) and Europe (Frankfurt). Also, you can now connect to Active Directory in any Amazon Web Services Region supported by Amazon QuickSight, except Asia Pacific (Singapore) and Asia Pacific (Sydney). For more information, see [Amazon Web Services Regions and IP address ranges](#).

August 1, 2019

[Embedding for Active Directory](#)

In Amazon QuickSight Enterprise edition, you can now embed dashboards for users authenticated through Active Directory. For more information, see [Embedding dashboards](#).

July 11, 2019

[Granular access control for using Amazon services from Amazon QuickSight](#)

In Amazon QuickSight Enterprise edition, you can scope down your security policies to allow specific people access to specific Amazon resources. You can attach and detach IAM policies in the Amazon QuickSight interface, simplifying finer control over who can access your Amazon data sources. For more information, see [Controlling access to Amazon resources](#).

June 14, 2019

[Tabs for dashboards and analyses](#)

Amazon QuickSight now supports multiple sheets inside of analyses and dashboards. You can add multiple tabs to your dashboards, utilize URL actions and on-screen controls to simplify navigation, and filter across all of your sheets. For more information, see [Multiple sheets](#).

June 11, 2019

[New visual types](#)

You can now use Amazon QuickSight to chart your data on a customizable gauge and donut charts. For more information, see [Visual types](#).

June 11, 2019

[New functions](#)

Amazon QuickSight now supports functions Variance and Standard Deviations as both aggregation and table calculation. You can create Sample and Population variants for both the functions. You can create these calculations through the calculations editor in analysis mode. [Functions by category](#)

June 11, 2019

[Hide/Show columns in table visualizations](#)

In Amazon QuickSight, you can hide or show columns in visualizations that use the *table* visual type. For more information, see [Customizing a visual](#).

June 11, 2019

Conditional string functions on SPICE	We added support for <code>ifelse</code> in analyses that use SPICE. For more information, see ifelse	June 11, 2019
Choose period granularity for date differences	We added an optional parameter to <code>dateDiff</code> so you can show differences between dates in the period you choose (years, months, and so on). For more information, see dateDiff .	June 11, 2019
Duplicate visuals	In Amazon QuickSight, you can now duplicate visuals. For more information, see Duplicating a visual .	May 21, 2019
Angled data labels	In Amazon QuickSight, data labels that are too long are now angled by default on vertical bar, combo, and line charts. For more information, see Customizing data labels .	May 21, 2019
New languages added	Amazon QuickSight is now available in 10 languages: English, German, Spanish, French, Portuguese, Italian, Japanese, Korean, Simplified Chinese, and Traditional Chinese. For more information, see Choosing a language in Amazon QuickSight .	April 8, 2019

[New aggregation function](#)

Amazon QuickSight supports aggregating by percentile. This function helps you understand the distribution of your data. For more information, see [percentile](#).

April 8, 2019

[Customize how many data points to display](#)

You can now format your visuals to display a custom number of data points or groups before showing the "other" category. This feature is available for bar charts, combo charts, line charts, pie charts, heat maps, and tree maps. For more information, see [Customizing a visual](#).

April 8, 2019

[ML-powered forecasting](#)

With forecasting powered by machine learning and what-if analyses in Amazon QuickSight, nontechnical users can now easily forecast their key business metrics. No ML expertise or Microsoft Excel data modeling is required. The built-in ML algorithm in Amazon QuickSight is designed to handle complex real-world scenarios. Amazon QuickSight uses ML to provide more reliable forecasts than traditional means. For more information, see [Forecasts and what-ifs](#).

March 14, 2019

[ML-powered anomaly detection](#)

Amazon QuickSight uses proven Amazon technology to continuously run ML-powered anomaly detection on millions of metrics and billions of data points. This anomaly detection enables you to get deep insights that are often buried in the aggregates, not visible in plain sight, and not scalable with manual analysis.

With ML-powered anomaly detection, there's no need for manual analysis, custom development, or ML domain expertise. For more information, see [Anomaly detection](#).

March 14, 2019

[Automatic narratives](#)

Automatic narratives provide key insights in everyday language, embedded contextually in your dashboard, saving hours on manual analysis. With automatic narratives, Amazon QuickSight interprets the charts and tables in your dashboard and provides a number of suggested insights in natural language. Depending on the shape and form of your data, you might get different suggestions. For example, you might see what the day-over-day changes look like, what was the highest sales date, or what the growth rate is. Or you might see what the forecast looks like for the next seven days. As the author of the dashboard, you can customize the computations and business language for your needs. You can use automatic narratives to effectively tell the story of your data in plain language. For more information, see [Narrative insights](#).

March 14, 2019

Previous updates

The following table describes the important changes in each release of the *Amazon Quick Suite User Guide* before March 4, 2019.

Change	Description	Date changed
New features	Using custom window functions, you can calculate any aggregation for a defined window at the point of selection. You can define the window interval before and after any point in time. You can also do calculations for that window using the sum, average, minimum, maximum, and count functions. Doing this enables a moving window aggregation as you progress through your data points. For more information, see Table calculation functions. Send email reports with data tailored to each of your users and groups. You can now create email reports for datasets that use row-level security. Amazon QuickSight generates a custom email snapshot for each user or group based on their data permission that is defined in the dashboard. RLS for email reports works for both scheduled and ad hoc emails. For more information, see Scheduling and sending QuickSight reports by email. Amazon QuickSight now supports unbounded cardinality. That means you can have more than 10,000 values in your control or filter. For more information, see Using a control with a parameter in Amazon QuickSight. You can create a datetime parameter that has no static default value. Filters on these parameters become active after you choose a value. For more information, see Adding date filters.	March 4, 2019

Change	Description	Date changed
New features	There are new table calculations available: • percentileOver • runningAvg • runningCount • runningMax • runningMin	February 7, 2019
New features	Amazon QuickSight supports emailing reports that use row-level security, using controls with unbounded cardinality, and creating datetime parameters with no default values. For more information, see the following: • Scheduling and sending QuickSight reports by email • Using a control with a parameter in Amazon QuickSight • Adding date filters	January 22, 2019

Change	Description	Date changed
New features	We added conditional aggregations, including <code>sumIf</code>, <code>countIf</code>, <code>minIf</code>, <code>maxIf</code>, <code>avgIf</code>, and <code>distinct_countIf</code> . Pivot tables now support infinite scrolling through millions of rows. You can add up to 20 fields for columns, and 20 fields for rows. Plus, you can add subtotals and totals to rows and columns. For more information, see the following: • Aggregate functions • Using pivot tables	January 10, 2019
New feature	We enhanced the join editor to increase usability and functionality. You can now add tables from one or more schemas on the same data source, or add the same table twice. For more information, see Joining data.	January 3, 2019
New SDK	You can embed dashboards and manage users or groups by using the Amazon QuickSight SDKs. For more information, see Developing with QuickSight.	November 27, 2018
New features	You can use groups with row-level security (RLS), and you can add cascading controls for parameters. For more information, see the following: • Using row-level security with user-based rules to restrict access to a dataset • Using a control with a parameter in Amazon QuickSight	November 20, 2018

Change	Description	Date changed
Function renamed	We renamed the percentile function to make it more intuitive. Its new name is <code>percentileRank</code> . For more information, see percentileRank .	November 12, 2018
New features	Use a top and bottom filter to show the top or bottom n for the field you choose, based on values in another field. For example, you could choose to show the top five sales people based on revenue. For more information, see Adding text filters. Use cascading controls to limit the values displayed in the controls, so they only show values that are relevant to what is selected in other controls. For more information, see Setting up parameters in Amazon QuickSight.	November 1, 2018
New feature	Use JSON native data types with <code>parseJson</code> . For more information see, parseJson .	October 30, 2018
New features	Use date functions to find out what quarter a date is in: • addDateTime • Extract • truncDate Add parameters to URLs. For more information see, Using parameters in a URL.	September 10, 2018
New features	You can sort strings in SPICE datasets. For more information on these features, see Sorting visual data in Amazon QuickSight	August 20, 2018

Change	Description	Date changed
New features	You can schedule emailed reports, and add data labels to your visuals. For more information on these features, see the following sections: • Scheduling and sending QuickSight reports by email • Subscribing to email reports in QuickSight • Data labels on visual types in QuickSight	August 15, 2018
New features	You can create table calculations, using aggregated measures to discover how dimensions influence measures or each other. Also, you can visualize time data at granularities as low as one minute. For more information on these features, see the following sections: • Table calculations • Changing date field granularity	August 8, 2018

Change	Description	Date changed
New features	You can replace datasets, customize labels, and format dimensions that are aggregated with count and count distinct. Also, new visuals start out smaller in size. For more information on these features, see the following sections: • Replacing datasets • Labels on visual types in QuickSight • Customizing a field format • Visualizing data in QuickSight	June 21, 2018

Change	Description	Date changed
New features and a new Amazon Web Services Region	You can upgrade your Amazon QuickSight subscription from Standard edition to Enterprise edition. In Enterprise edition, Amazon QuickSight supports usage-based pricing for users in the reader role, sharing dashboards with all users in the reader role, and hourly refresh of datasets. Amazon QuickSight also supports private connections to data in a VPC with a private subnet in Enterprise edition. In both editions, Amazon QuickSight supports parameters with on-sheet controls, dashboard co-ownership, custom URL actions, and 25-GB SPICE datasets. Also, Amazon QuickSight is available in Asia Pacific (Tokyo). For more information on these features, see the following sections: • Upgrading your Amazon QuickSight subscription from Standard edition to Enterprise edition • Self-provisioning an Amazon QuickSight read-only user • Inviting users to access Amazon QuickSight • Sharing QuickSight dashboards • Connecting to a VPC with QuickSight • Refreshing a dataset on a schedule • Parameters in Amazon QuickSight •	May 30, 2018

Change	Description	Date changed
	• Exploring interactive dashboards in QuickSight (newly updated for read-only users) • Using custom actions for filtering and navigating • Data source quotas • Amazon Web Services Regions, websites, IP address ranges, and endpoints	
New feature	Amazon QuickSight supports quick sorting from axis labels, duplicating datasets, and showing or hiding totals on tabular reports. Also, you can add custom SQL to a dataset earlier in the creation workflow. For more information on these features, see the following sections: • Sorting visual data in Amazon QuickSight • Duplicating datasets • Totals and subtotals • Creating a basic SQL query	May 25, 2018
New features	You can use Amazon QuickSight to connect to Software as a Service (SaaS) providers. For more information, see Software as a service (SaaS) data. You can import JSON files to Amazon QuickSight. For more information, see JSON data. You can also parse JSON fields in a CSV file. For more information, see parseJson.	April 9, 2018

Change	Description	Date changed
New feature	You can use Amazon QuickSight with Amazon S3 files that are in a different Amazon account. For more information, see Datasets using S3 files in another Amazon account .	November 20, 2017
New visual types	You can create visuals using maps. You can also view tabular data in a visual. For more information on preparing geospatial data for use in a visual, see Adding geospatial data . For more information on creating a geospatial visual, see Creating maps and geospatial charts . For more information on using tabular reports, see Using tables as visuals .	November 20, 2017
New features	Amazon QuickSight can support 1000 columns in a dataset. For more information, see Data source quotas . Calculated fields are supported in SPICE datasets. For more information, see Adding calculated fields . Also, high cardinality values, which often display as a long tail on a visual, are placed into a category called other . For more information, see Visual types in Amazon QuickSight .	November 20, 2017
New feature	In Enterprise edition, you can restrict access to a dataset by adding row-level security. To learn more, see Using row-level security with user-based rules to restrict access to a dataset .	October 20, 2017
New visual type	You can create visuals using combo charts. To learn more about combo charts, see Using combo charts .	October 20, 2017

Change	Description	Date changed
New features	Amazon QuickSight supports creating custom aggregations for calculated fields in analyses, custom date formats, and copies of dashboards. For more information on aggregating calculated fields, see Aggregate functions . For more information on using unsupported dates by creating a custom date format, see Using unsupported or custom dates . For more information on duplicating dashboards, see Copying an QuickSight dashboard .	September 25, 2017
New feature	You can combine multiple filters using the And/Or operators. To learn more about filter groups, see Adding filter conditions (group filters) with AND and OR operators .	August 31, 2017
New data source	Amazon QuickSight supports Amazon S3 Analytics.	August 31, 2017
New Features	Amazon QuickSight supports importing ZIP files from Amazon S3. There is also a new search feature, to simplify finding analyses, datasets, and dashboards. For more information on the search feature, see Using the QuickSight console .	August 31, 2017
New Amazon Web Services Regions	Amazon QuickSight is now available in Asia Pacific (Singapore) and Asia Pacific (Sydney).	August 8, 2017
New feature	Amazon QuickSight supports Snowflake cloud data warehouse.	July 31, 2017
New feature	Amazon QuickSight added a new aggregation: Count Distinct. To learn more, see Changing field aggregation .	July 19, 2017

Change	Description	Date changed
New feature	Amazon QuickSight supports exploring Amazon S3 analytics data from an Amazon QuickSight dashboard that you reach from the Amazon Management Console. To learn more, see Exploring your Amazon data in Amazon QuickSight .	July 5, 2017
New feature	Amazon QuickSight supports Federated Single Sign-On (IAM Identity Center) in Enterprise Edition. To learn more, see Using external identity federation and single sign-on with Amazon QuickSight .	May 25, 2017
New feature	Amazon QuickSight supports Amazon Redshift Spectrum. To learn more, see Enabling access to Amazon Redshift Spectrum .	May 25, 2017
New feature	Amazon QuickSight supports Federated Single Sign-On (IAM Identity Center) in Standard Edition. To learn more, see Using external identity federation and single sign-on with Amazon QuickSight .	May 25, 2017
New feature	Amazon QuickSight supports just-in-time (JIT) user provisioning through the following policy actions: <code>quicksight:CreateUser</code> and <code>quicksight:CreateAdmin</code> . To learn more, see IAM policy examples for QuickSight .	May 25, 2017
New feature	Amazon QuickSight supports direct connections to Teradata 14.0 and later.	May 25, 2017
New feature	Amazon QuickSight added relative date filters for datasets and visuals. To learn more, see Adding date filters .	May 25, 2017

Change	Description	Date changed
New feature	Amazon QuickSight supports connecting to Apache Spark and Presto. To learn more, see Creating a data source using Apache Spark and Creating a data source using Presto .	May 3, 2017
New feature	Amazon QuickSight supports operational logging with Amazon CloudTrail. To learn more, see Logging operations with Amazon CloudTrail .	April 28, 2017
New feature	Amazon QuickSight is available in US East (Ohio). To learn more about Amazon Web Services Regions, see Amazon Web Services Regions, websites, IP address ranges, and endpoints .	April 11, 2017
New feature	- You can use the AD Connector with Amazon QuickSight. To learn more about managing Enterprise users, see Managing user access inside Amazon QuickSight. - You can add key performance indicators (KPIs) to your visualizations. To learn more, see Using KPIs. - You can import .xlsx files with headers and footers, comments, formatting, filter header, sort, frozen panel/header, hidden columns, groups, and formulas / references. You can also limit your import to a specific range. To learn more about importing ranges, see Choosing file upload settings.	April 6, 2017
New feature	You can now export to a comma-separated value (CSV) format file using Amazon QuickSight. To learn more, see Exporting data from visuals .	March 21, 2017

Change	Description	Date changed
New feature	You can now schedule data refreshes for SPICE data sets. To learn more, see Refreshing SPICE data .	February 14, 2017
New feature	You can now connect to Amazon Athena databases and use them as data sources in Amazon QuickSight. To learn more, see Creating a dataset using Amazon Athena data .	December 22, 2016
New edition	Amazon QuickSight now offers an Enterprise edition as well as a Standard edition. Both editions offer a full set of features for creating and sharing data visualizations, and Enterprise edition additionally offers encryption at rest and Active Directory integration. When you choose to use Enterprise edition, you select a Microsoft Active Directory directory in Amazon Directory Service and use that active directory to identify and manage your Amazon QuickSight users and administrators. To learn more, see Different editions of Amazon QuickSight .	December 15, 2016
New guide	This is the first release of <i>QuickSight User Guide</i> .	November 15, 2016

Attributions

Following, you can find out about attributions and third-party licenses for material used in Quick Suite documentation.

** Mapzen Who's On First; version latest -- <https://tinyurl.com/y8p87rwl>

Data from Who's On First. <http://whosonfirst.mapzen.com#License>

CC0 1.0 Universal

CREATIVE COMMONS CORPORATION IS NOT A LAW FIRM AND DOES NOT PROVIDE LEGAL SERVICES. DISTRIBUTION OF THIS DOCUMENT DOES NOT CREATE AN ATTORNEY-CLIENT RELATIONSHIP. CREATIVE COMMONS PROVIDES THIS INFORMATION ON AN "AS-IS" BASIS. CREATIVE COMMONS MAKES NO WARRANTIES REGARDING THE USE OF THIS DOCUMENT OR THE INFORMATION OR WORKS PROVIDED HEREUNDER, AND DISCLAIMS LIABILITY FOR DAMAGES RESULTING FROM THE USE OF THIS DOCUMENT OR THE INFORMATION OR WORKS PROVIDED HEREUNDER.

Statement of Purpose

The laws of most jurisdictions throughout the world automatically confer exclusive Copyright and Related Rights (defined below) upon the creator and subsequent owner(s) (each and all, an "owner") of an original work of authorship and/or a database (each, a "Work").

Certain owners wish to permanently relinquish those rights to a Work for the purpose of contributing to a commons of creative, cultural and scientific works ("Commons") that the public can reliably and without fear of later claims of infringement build upon, modify, incorporate in other works, reuse and redistribute as freely as possible in any form whatsoever and for any purposes, including without limitation commercial purposes. These owners may contribute to the Commons to promote the ideal of a free culture and the further production of creative, cultural and scientific works, or to gain reputation or greater distribution for their Work in part through the use and efforts of others.

For these and/or other purposes and motivations, and without any expectation of additional consideration or compensation, the person associating CC0 with a Work (the "Affirmer"), to the extent that he or she is an owner of Copyright and Related Rights in the Work, voluntarily elects to

apply CC0 to the Work and publicly distribute the Work under its terms, with knowledge of his or her Copyright and Related Rights in the Work and the meaning and intended legal effect of CC0 on those rights.

1. Copyright and Related Rights. A Work made available under CC0 may be protected by copyright and related or neighboring rights ("Copyright and Related Rights"). Copyright and Related Rights include, but are not limited to, the following:

- i. the right to reproduce, adapt, distribute, perform, display, communicate, and translate a Work;
- ii. moral rights retained by the original author(s) and/or performer(s);
- iii. publicity and privacy rights pertaining to a person's image or likeness depicted in a Work;
- iv. rights protecting against unfair competition in regards to a Work, subject to the limitations in paragraph 4(a), below;
- v. rights protecting the extraction, dissemination, use and reuse of data in a Work;
- vi. database rights (such as those arising under Directive 96/9/EC of the European Parliament and of the Council of 11 March 1996 on the legal protection of databases, and under any national implementation thereof, including any amended or successor version of such directive); and
- vii. other similar, equivalent or corresponding rights throughout the world based on applicable law or treaty, and any national implementations thereof.

2. Waiver. To the greatest extent permitted by, but not in contravention of, applicable law, Affirmer hereby overtly, fully, permanently, irrevocably and unconditionally waives, abandons, and surrenders all of Affirmer's Copyright and Related Rights and associated claims and causes of action, whether now known or unknown (including existing as well as future claims and causes of action), in the Work (i) in all territories worldwide, (ii) for the maximum duration provided by applicable law or treaty (including future time extensions), (iii) in any current or future medium and for any number of copies, and (iv) for any purpose whatsoever, including without limitation commercial, advertising or promotional purposes (the "Waiver"). Affirmer makes the Waiver for the benefit of each member of the public at large and to the detriment of Affirmer's heirs and successors, fully intending that such Waiver shall not be subject to revocation, rescission, cancellation, termination, or any other legal or equitable action to disrupt the quiet enjoyment of the Work by the public as contemplated by Affirmer's express Statement of Purpose.

3. Public License Fallback. Should any part of the Waiver for any reason be judged legally invalid or ineffective under applicable law, then the Waiver shall be preserved to the maximum extent

permitted taking into account Affirmer's express Statement of Purpose. In addition, to the extent the Waiver is so judged Affirmer hereby grants to each affected person a royalty-free, non transferable, non sublicensable, non exclusive, irrevocable and unconditional license to exercise Affirmer's Copyright and Related Rights in the Work (i) in all territories worldwide, (ii) for the maximum duration provided by applicable law or treaty (including future time extensions), (iii) in any current or future medium and for any number of copies, and (iv) for any purpose whatsoever, including without limitation commercial, advertising or promotional purposes (the "License"). The License shall be deemed effective as of the date CC0 was applied by Affirmer to the Work. Should any part of the License for any reason be judged legally invalid or ineffective under applicable law, such partial invalidity or ineffectiveness shall not invalidate the remainder of the License, and in such case Affirmer hereby affirms that he or she will not (i) exercise any of his or her remaining Copyright and Related Rights in the Work or (ii) assert any associated claims and causes of action with respect to the Work, in either case contrary to Affirmer's express Statement of Purpose.

4. Limitations and Disclaimers.

a. No trademark or patent rights held by Affirmer are waived, abandoned, surrendered, licensed or otherwise affected by this document.

b. Affirmer offers the Work as-is and makes no representations or warranties of any kind concerning the Work, express, implied, statutory or otherwise, including without limitation warranties of title, merchantability, fitness for a particular purpose, non infringement, or the absence of latent or other defects, accuracy, or the present or absence of errors, whether or not discoverable, all to the greatest extent permissible under applicable law.

c. Affirmer disclaims responsibility for clearing rights of other persons that may apply to the Work or any use thereof, including without limitation any person's Copyright and Related Rights in the Work. Further, Affirmer disclaims responsibility for obtaining any necessary consents, permissions or other rights required for any use of the Work.

d. Affirmer understands and acknowledges that Creative Commons is not a party to this document and has no duty or obligation with respect to this CC0 or use of the Work.

** Quattroshapes; version latest -- <https://tinyurl.com/y72sd8u8>

Includes data from foursquare quattroshapes. <https://github.com/foursquare/quattroshapes/>.

This map database contains open data from government and other sources. Consider including the following copyright statements when using the data.

Europe

EuroGeoGraphics data copyright is held by European National Mapping Agencies.

Austria © Bundesamt für Eich- und Vermessungswesen

Belgium © Institut Géographique National – Belgique

Bulgaria © Cadastre Agency, MRDPW

Croatia © Drzavna Geodetska Uprava

Cyprus © Lands and Surveys, Survey and Cartogr. Br.

Czech Republic © Český úřad zeměměřický a katastrální

Denmark © Kort og Matrikelstyrelsen

Estonia © Maaamet

Faroe Islands © Kort og Matrikelstyrelsen

Finland © Maanmittauslaitos

France © Institut National de l'Information Géographique et Forestière – France

Georgia © National Agency of Public Registry

Germany © Bundesamt für Kartographie und Geodäsie

Great Britain © Ordnance Survey

Greece © Hellenic Military Geographical Service

Greenland © Kort og Matrikelstyrelsen

Hungary © Földmérési és Távérzékelési Intézet

Iceland © Landmælingar Íslands

Ireland © Ordnance Survey Ireland

Italy © Istituto Geografico Militare Italiano

Latvia © Latvijas Republikas Valsts zemes dienests

Liechtenstein © Bundesamt für Landestopographie (Switzerland)

Lithuania © Nacionalinė žemės tarnyba

Luxembourg © Administration du Cadastre et de la Topographie

Malta © Malta Environment and Planning Authority (MEPA)

Moldova © State Agency for Land Relations and Cadastre

Netherlands © Topografische Dienst Nederland

Northern Ireland © Ordnance Survey of Northern Ireland

Norway © Statens Kartverk

Poland © Główny Urząd Geodezji i Kartografii

Portugal © Instituto Geográfico Português

Rep. of Slovakia © Geodetick a kartografick ústav

Romania © CNGCFT (National Center of Geodesie, Cartography, Fotogrametry and Remote Sensing)

Serbia © Republički geodetski zavod

Slovenia © Geodetska Uprava Republike Slovenije

Spain © Centro Nacional de Informacion Geografica – Instituto Geográfico Nacional

Sweden © Lantmäteriet

Switzerland © Bundesamt für Landestopographie

Ukraine © Research Institute of Geodesy and Cartography

Additional European data

United Kingdom: Contains Ordnance Survey data © Crown copyright and database right [2012]

Netherlands: Kadaster

Spain: Instituto Geográfico Nacional

France: Institut Géographique National

Switzerland: swisstopo

Europe-wide: European Environment Agency (EEA) urban morphological zones 2006

Americas

United States: US Census Bureau (Census 2010 geography files).

Canada: © Department of Natural Resources Canada. All rights reserved., Statistics Canada, and BC Stats

Brazil: IBGE

Mexico: INEGI

Chile: Global Map of Chile © International Steering Committee for Global Mapping / Instituto Geografico Militar de Chile

Asia

Indonesia: Global Map of Indonesia @ ISCGM/Indonesia

Australian: Geoscience Australia and Australian Bureau of Statistics

New Zealand: Land Information New Zealand

South Africa: Global Map of the Republic of South Africa © ISCGM/State Copyright, Department of Rural Development and Land Reform, Chief Directorate of National Geo-spatial Information, Website: <http://www.ngi.gov.za/>

Palestine: "Global Map of PALESTINE © ISCGM/Palestinian National Authority, Ministry of Planning. Geographic Center and Technical Support www.mop.gov.ps

Ordnance Survey

Contains OS data © Crown copyright and database right 2016

Contains Royal Mail data © Royal Mail copyright and Database right 2016

Contains National Statistics data © Crown copyright and database right 2016

Attribution 2.0

CREATIVE COMMONS CORPORATION IS NOT A LAW FIRM AND DOES NOT PROVIDE LEGAL SERVICES. DISTRIBUTION OF THIS LICENSE DOES NOT CREATE AN ATTORNEY-CLIENT RELATIONSHIP. CREATIVE COMMONS PROVIDES THIS INFORMATION ON AN "AS-IS" BASIS. CREATIVE COMMONS MAKES NO WARRANTIES REGARDING THE INFORMATION PROVIDED, AND DISCLAIMS LIABILITY FOR DAMAGES RESULTING FROM ITS USE.

License

THE WORK (AS DEFINED BELOW) IS PROVIDED UNDER THE TERMS OF THIS CREATIVE COMMONS PUBLIC LICENSE ("CCPL" OR "LICENSE"). THE WORK IS PROTECTED BY COPYRIGHT AND/OR OTHER APPLICABLE LAW. ANY USE OF THE WORK OTHER THAN AS AUTHORIZED UNDER THIS LICENSE OR COPYRIGHT LAW IS PROHIBITED.

BY EXERCISING ANY RIGHTS TO THE WORK PROVIDED HERE, YOU ACCEPT AND AGREE TO BE BOUND BY THE TERMS OF THIS LICENSE. THE LICENSOR GRANTS YOU THE RIGHTS CONTAINED HERE IN CONSIDERATION OF YOUR ACCEPTANCE OF SUCH TERMS AND CONDITIONS.

1. Definitions

a. "Collective Work" means a work, such as a periodical issue, anthology or encyclopedia, in which the Work in its entirety in unmodified form, along with a number of other contributions, constituting separate and independent works in themselves, are assembled into a collective whole. A work that constitutes a Collective Work will not be considered a Derivative Work (as defined below) for the purposes of this License.

b. "Derivative Work" means a work based upon the Work or upon the Work and other pre-existing works, such as a translation, musical arrangement, dramatization, fictionalization, motion picture version, sound recording, art reproduction, abridgment, condensation, or any other form in which the Work may be recast, transformed, or adapted, except that a work that constitutes a Collective Work will not be considered a Derivative Work for the purpose of this License. For the avoidance of doubt, where the Work is a musical composition or sound recording, the synchronization of the

Work in timed-relation with a moving image ("synching") will be considered a Derivative Work for the purpose of this License.

c. "Licensor" means the individual or entity that offers the Work under the terms of this License.

d. "Original Author" means the individual or entity who created the Work.

e. "Work" means the copyrightable work of authorship offered under the terms of this License.

f. "You" means an individual or entity exercising rights under this License who has not previously violated the terms of this License with respect to the Work, or who has received express permission from the Licensor to exercise rights under this License despite a previous violation.

2. Fair Use Rights. Nothing in this license is intended to reduce, limit, or restrict any rights arising from fair use, first sale or other limitations on the exclusive rights of the copyright owner under copyright law or other applicable laws.

3. License Grant. Subject to the terms and conditions of this License, Licensor hereby grants You a worldwide, royalty-free, non-exclusive, perpetual (for the duration of the applicable copyright) license to exercise the rights in the Work as stated below:

a. to reproduce the Work, to incorporate the Work into one or more Collective Works, and to reproduce the Work as incorporated in the Collective Works;

b. to create and reproduce Derivative Works;

c. to distribute copies or phonorecords of, display publicly, perform publicly, and perform publicly by means of a digital audio transmission the Work including as incorporated in Collective Works;

d. to distribute copies or phonorecords of, display publicly, perform publicly, and perform publicly by means of a digital audio transmission Derivative Works.

e. For the avoidance of doubt, where the work is a musical composition:

i. Performance Royalties Under Blanket Licenses. Licensor waives the exclusive right to collect, whether individually or via a performance rights society (e.g. ASCAP, BMI, SESAC), royalties for the public performance or public digital performance (e.g. webcast) of the Work.

ii. Mechanical Rights and Statutory Royalties. Licensor waives the exclusive right to collect, whether individually or via a music rights agency or designated agent (e.g. Harry Fox Agency), royalties for any phonorecord You create from the Work ("cover version") and distribute, subject to the

compulsory license created by 17 USC Section 115 of the US Copyright Act (or the equivalent in other jurisdictions).

f. **Webcasting Rights and Statutory Royalties.** For the avoidance of doubt, where the Work is a sound recording, Licensors waive the exclusive right to collect, whether individually or via a performance-rights society (e.g. SoundExchange), royalties for the public digital performance (e.g. webcast) of the Work, subject to the compulsory license created by 17 USC Section 114 of the US Copyright Act (or the equivalent in other jurisdictions).

The above rights may be exercised in all media and formats whether now known or hereafter devised. The above rights include the right to make such modifications as are technically necessary to exercise the rights in other media and formats. All rights not expressly granted by Licensors are hereby reserved.

4. **Restrictions.** The license granted in Section 3 above is expressly made subject to and limited by the following restrictions:

a. You may distribute, publicly display, publicly perform, or publicly digitally perform the Work only under the terms of this License, and You must include a copy of, or the Uniform Resource Identifier for, this License with every copy or phonorecord of the Work You distribute, publicly display, publicly perform, or publicly digitally perform. You may not offer or impose any terms on the Work that alter or restrict the terms of this License or the recipients' exercise of the rights granted hereunder. You may not sublicense the Work. You must keep intact all notices that refer to this License and to the disclaimer of warranties. You may not distribute, publicly display, publicly perform, or publicly digitally perform the Work with any technological measures that control access or use of the Work in a manner inconsistent with the terms of this License Agreement. The above applies to the Work as incorporated in a Collective Work, but this does not require the Collective Work apart from the Work itself to be made subject to the terms of this License. If You create a Collective Work, upon notice from any Licensors You must, to the extent practicable, remove from the Collective Work any reference to such Licensors or the Original Author, as requested. If You create a Derivative Work, upon notice from any Licensors You must, to the extent practicable, remove from the Derivative Work any reference to such Licensors or the Original Author, as requested.

b. If you distribute, publicly display, publicly perform, or publicly digitally perform the Work or any Derivative Works or Collective Works, You must keep intact all copyright notices for the Work and give the Original Author credit reasonable to the medium or means You are utilizing by conveying the name (or pseudonym if applicable) of the Original Author if supplied; the title of the Work if supplied; to the extent reasonably practicable, the Uniform Resource Identifier, if any, that Licensors

specifies to be associated with the Work, unless such URI does not refer to the copyright notice or licensing information for the Work; and in the case of a Derivative Work, a credit identifying the use of the Work in the Derivative Work (e.g., "French translation of the Work by Original Author," or "Screenplay based on original Work by Original Author"). Such credit may be implemented in any reasonable manner; provided, however, that in the case of a Derivative Work or Collective Work, at a minimum such credit will appear where any other comparable authorship credit appears and in a manner at least as prominent as such other comparable authorship credit.

5. Representations, Warranties and Disclaimer

UNLESS OTHERWISE MUTUALLY AGREED TO BY THE PARTIES IN WRITING, LICENSOR OFFERS THE WORK AS-IS AND MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND CONCERNING THE WORK, EXPRESS, IMPLIED, STATUTORY OR OTHERWISE, INCLUDING, WITHOUT LIMITATION, WARRANTIES OF TITLE, MERCHANTIBILITY, FITNESS FOR A PARTICULAR PURPOSE, NONINFRINGEMENT, OR THE ABSENCE OF LATENT OR OTHER DEFECTS, ACCURACY, OR THE PRESENCE OF ABSENCE OF ERRORS, WHETHER OR NOT DISCOVERABLE. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OF IMPLIED WARRANTIES, SO SUCH EXCLUSION MAY NOT APPLY TO YOU.

6. Limitation on Liability. EXCEPT TO THE EXTENT REQUIRED BY APPLICABLE LAW, IN NO EVENT WILL LICENSOR BE LIABLE TO YOU ON ANY LEGAL THEORY FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR EXEMPLARY DAMAGES ARISING OUT OF THIS LICENSE OR THE USE OF THE WORK, EVEN IF LICENSOR HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. Termination

a. This License and the rights granted hereunder will terminate automatically upon any breach by You of the terms of this License. Individuals or entities who have received Derivative Works or Collective Works from You under this License, however, will not have their licenses terminated provided such individuals or entities remain in full compliance with those licenses. Sections 1, 2, 5, 6, 7, and 8 will survive any termination of this License.

b. Subject to the above terms and conditions, the license granted here is perpetual (for the duration of the applicable copyright in the Work). Notwithstanding the above, Licensor reserves the right to release the Work under different license terms or to stop distributing the Work at any time; provided, however that any such election will not serve to withdraw this License (or any other license that has been, or is required to be, granted under the terms of this License), and this License will continue in full force and effect unless terminated as stated above.

8. Miscellaneous

- a. Each time You distribute or publicly digitally perform the Work or a Collective Work, the Licensor offers to the recipient a license to the Work on the same terms and conditions as the license granted to You under this License.
- b. Each time You distribute or publicly digitally perform a Derivative Work, Licensor offers to the recipient a license to the original Work on the same terms and conditions as the license granted to You under this License.
- c. If any provision of this License is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this License, and without further action by the parties to this agreement, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.
- d. No term or provision of this License shall be deemed waived and no breach consented to unless such waiver or consent shall be in writing and signed by the party to be charged with such waiver or consent.
- e. This License constitutes the entire agreement between the parties with respect to the Work licensed here. There are no understandings, agreements or representations with respect to the Work not specified here. Licensor shall not be bound by any additional provisions that may appear in any communication from You. This License may not be modified without the mutual written agreement of the Licensor and You.

** Yahoo! Geoplanet; version latest -- <https://tinyurl.com/y8p87rwl>

Content licensed under Creative Commons Attribution 3.0 United States

License. Provide attribution to "Yahoo! GeoPlanet"

License

THE WORK (AS DEFINED BELOW) IS PROVIDED UNDER THE TERMS OF THIS CREATIVE COMMONS PUBLIC LICENSE ("CCPL" OR "LICENSE"). THE WORK IS PROTECTED BY COPYRIGHT AND/OR OTHER

APPLICABLE LAW. ANY USE OF THE WORK OTHER THAN AS AUTHORIZED UNDER THIS LICENSE OR COPYRIGHT LAW IS PROHIBITED.

BY EXERCISING ANY RIGHTS TO THE WORK PROVIDED HERE, YOU ACCEPT AND AGREE TO BE BOUND BY THE TERMS OF THIS LICENSE. TO THE EXTENT THIS LICENSE MAY BE CONSIDERED TO BE A CONTRACT, THE LICENSOR GRANTS YOU THE RIGHTS CONTAINED HERE IN CONSIDERATION OF YOUR ACCEPTANCE OF SUCH TERMS AND CONDITIONS.

1. Definitions

"Adaptation" means a work based upon the Work, or upon the Work and other pre-existing works, such as a translation, adaptation, derivative work, arrangement of music or other alterations of a literary or artistic work, or phonogram or performance and includes cinematographic adaptations or any other form in which the Work may be recast, transformed, or adapted including in any form recognizably derived from the original, except that a work that constitutes a Collection will not be considered an Adaptation for the purpose of this License. For the avoidance of doubt, where the Work is a musical work, performance or phonogram, the synchronization of the Work in timed-relation with a moving image ("synching") will be considered an Adaptation for the purpose of this License.

"Collection" means a collection of literary or artistic works, such as encyclopedias and anthologies, or performances, phonograms or broadcasts, or other works or subject matter other than works listed in Section 1(f) below, which, by reason of the selection and arrangement of their contents, constitute intellectual creations, in which the Work is included in its entirety in unmodified form along with one or more other contributions, each constituting separate and independent works in themselves, which together are assembled into a collective whole. A work that constitutes a Collection will not be considered an Adaptation (as defined above) for the purposes of this License.

"Distribute" means to make available to the public the original and copies of the Work or Adaptation, as appropriate, through sale or other transfer of ownership.

"Licensor" means the individual, individuals, entity or entities that offer(s) the Work under the terms of this License.

"Original Author" means, in the case of a literary or artistic work, the individual, individuals, entity or entities who created the Work or if no individual or entity can be identified, the publisher; and in addition (i) in the case of a performance the actors, singers, musicians, dancers, and other persons who act, sing, deliver, declaim, play in, interpret or otherwise perform literary or artistic

works or expressions of folklore; (ii) in the case of a phonogram the producer being the person or legal entity who first fixes the sounds of a performance or other sounds; and, (iii) in the case of broadcasts, the organization that transmits the broadcast.

"Work" means the literary and/or artistic work offered under the terms of this License including without limitation any production in the literary, scientific and artistic domain, whatever may be the mode or form of its expression including digital form, such as a book, pamphlet and other writing; a lecture, address, sermon or other work of the same nature; a dramatic or dramatico-musical work; a choreographic work or entertainment in dumb show; a musical composition with or without words; a cinematographic work to which are assimilated works expressed by a process analogous to cinematography; a work of drawing, painting, architecture, sculpture, engraving or lithography; a photographic work to which are assimilated works expressed by a process analogous to photography; a work of applied art; an illustration, map, plan, sketch or three-dimensional work relative to geography, topography, architecture or science; a performance; a broadcast; a phonogram; a compilation of data to the extent it is protected as a copyrightable work; or a work performed by a variety or circus performer to the extent it is not otherwise considered a literary or artistic work.

"You" means an individual or entity exercising rights under this License who has not previously violated the terms of this License with respect to the Work, or who has received express permission from the Licensor to exercise rights under this License despite a previous violation.

"Publicly Perform" means to perform public recitations of the Work and to communicate to the public those public recitations, by any means or process, including by wire or wireless means or public digital performances; to make available to the public Works in such a way that members of the public may access these Works from a place and at a place individually chosen by them; to perform the Work to the public by any means or process and the communication to the public of the performances of the Work, including by public digital performance; to broadcast and rebroadcast the Work by any means including signs, sounds or images.

"Reproduce" means to make copies of the Work by any means including without limitation by sound or visual recordings and the right of fixation and reproducing fixations of the Work, including storage of a protected performance or phonogram in digital form or other electronic medium.

2. Fair Dealing Rights. Nothing in this License is intended to reduce, limit, or restrict any uses free from copyright or rights arising from limitations or exceptions that are provided for in connection with the copyright protection under copyright law or other applicable laws.

3. License Grant. Subject to the terms and conditions of this License, Licensor hereby grants You a worldwide, royalty-free, non-exclusive, perpetual (for the duration of the applicable copyright) license to exercise the rights in the Work as stated below:

to Reproduce the Work, to incorporate the Work into one or more Collections, and to Reproduce the Work as incorporated in the Collections;

to create and Reproduce Adaptations provided that any such Adaptation, including any translation in any medium, takes reasonable steps to clearly label, demarcate or otherwise identify that changes were made to the original Work. For example, a translation could be marked "The original work was translated from English to Spanish," or a modification could indicate "The original work has been modified.";

to Distribute and Publicly Perform the Work including as incorporated in Collections; and,

to Distribute and Publicly Perform Adaptations.

For the avoidance of doubt:

Non-waivable Compulsory License Schemes. In those jurisdictions in which the right to collect royalties through any statutory or compulsory licensing scheme cannot be waived, the Licensor reserves the exclusive right to collect such royalties for any exercise by You of the rights granted under this License;

Waivable Compulsory License Schemes. In those jurisdictions in which the right to collect royalties through any statutory or compulsory licensing scheme can be waived, the Licensor waives the exclusive right to collect such royalties for any exercise by You of the rights granted under this License; and,

Voluntary License Schemes. The Licensor waives the right to collect royalties, whether individually or, in the event that the Licensor is a member of a collecting society that administers voluntary licensing schemes, via that society, from any exercise by You of the rights granted under this License.

The above rights may be exercised in all media and formats whether now known or hereafter devised. The above rights include the right to make such modifications as are technically necessary

to exercise the rights in other media and formats. Subject to Section 8(f), all rights not expressly granted by Licensor are hereby reserved.

4. Restrictions. The license granted in Section 3 above is expressly made subject to and limited by the following restrictions:

You may Distribute or Publicly Perform the Work only under the terms of this License. You must include a copy of, or the Uniform Resource Identifier (URI) for, this License with every copy of the Work You Distribute or Publicly Perform. You may not offer or impose any terms on the Work that restrict the terms of this License or the ability of the recipient of the Work to exercise the rights granted to that recipient under the terms of the License. You may not sublicense the Work. You must keep intact all notices that refer to this License and to the disclaimer of warranties with every copy of the Work You Distribute or Publicly Perform. When You Distribute or Publicly Perform the Work, You may not impose any effective technological measures on the Work that restrict the ability of a recipient of the Work from You to exercise the rights granted to that recipient under the terms of the License. This Section 4(a) applies to the Work as incorporated in a Collection, but this does not require the Collection apart from the Work itself to be made subject to the terms of this License. If You create a Collection, upon notice from any Licensor You must, to the extent practicable, remove from the Collection any credit as required by Section 4(b), as requested. If You create an Adaptation, upon notice from any Licensor You must, to the extent practicable, remove from the Adaptation any credit as required by Section 4(b), as requested.

If You Distribute, or Publicly Perform the Work or any Adaptations or Collections, You must, unless a request has been made pursuant to Section 4(a), keep intact all copyright notices for the Work and provide, reasonable to the medium or means You are utilizing: (i) the name of the Original Author (or pseudonym, if applicable) if supplied, and/or if the Original Author and/or Licensor designate another party or parties (e.g., a sponsor institute, publishing entity, journal) for attribution ("Attribution Parties") in Licensor's copyright notice, terms of service or by other reasonable means, the name of such party or parties; (ii) the title of the Work if supplied; (iii) to the extent reasonably practicable, the URI, if any, that Licensor specifies to be associated with the Work, unless such URI does not refer to the copyright notice or licensing information for the Work; and (iv) , consistent with Section 3(b), in the case of an Adaptation, a credit identifying the use of the Work in the Adaptation (e.g., "French translation of the Work by Original Author," or "Screenplay based on original Work by Original Author"). The credit required by this Section 4 (b) may be implemented in any reasonable manner; provided, however, that in the case of a Adaptation or Collection, at a minimum such credit will appear, if a credit for all contributing authors of the Adaptation or Collection appears, then as part of these credits and in a manner at

least as prominent as the credits for the other contributing authors. For the avoidance of doubt, You may only use the credit required by this Section for the purpose of attribution in the manner set out above and, by exercising Your rights under this License, You may not implicitly or explicitly assert or imply any connection with, sponsorship or endorsement by the Original Author, Licensor and/or Attribution Parties, as appropriate, of You or Your use of the Work, without the separate, express prior written permission of the Original Author, Licensor and/or Attribution Parties.

Except as otherwise agreed in writing by the Licensor or as may be otherwise permitted by applicable law, if You Reproduce, Distribute or Publicly Perform the Work either by itself or as part of any Adaptations or Collections, You must not distort, mutilate, modify or take other derogatory action in relation to the Work which would be prejudicial to the Original Author's honor or reputation. Licensor agrees that in those jurisdictions (e.g. Japan), in which any exercise of the right granted in Section 3(b) of this License (the right to make Adaptations) would be deemed to be a distortion, mutilation, modification or other derogatory action prejudicial to the Original Author's honor and reputation, the Licensor will waive or not assert, as appropriate, this Section, to the fullest extent permitted by the applicable national law, to enable You to reasonably exercise Your right under Section 3(b) of this License (right to make Adaptations) but not otherwise.

5. Representations, Warranties and Disclaimer

UNLESS OTHERWISE MUTUALLY AGREED TO BY THE PARTIES IN WRITING, LICENSOR OFFERS THE WORK AS-IS AND MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND CONCERNING THE WORK, EXPRESS, IMPLIED, STATUTORY OR OTHERWISE, INCLUDING, WITHOUT LIMITATION, WARRANTIES OF TITLE, MERCHANTIBILITY, FITNESS FOR A PARTICULAR PURPOSE, NONINFRINGEMENT, OR THE ABSENCE OF LATENT OR OTHER DEFECTS, ACCURACY, OR THE PRESENCE OF ABSENCE OF ERRORS, WHETHER OR NOT DISCOVERABLE. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OF IMPLIED WARRANTIES, SO SUCH EXCLUSION MAY NOT APPLY TO YOU.

6. Limitation on Liability. EXCEPT TO THE EXTENT REQUIRED BY APPLICABLE LAW, IN NO EVENT WILL LICENSOR BE LIABLE TO YOU ON ANY LEGAL THEORY FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR EXEMPLARY DAMAGES ARISING OUT OF THIS LICENSE OR THE USE OF THE WORK, EVEN IF LICENSOR HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. Termination

This License and the rights granted hereunder will terminate automatically upon any breach by You of the terms of this License. Individuals or entities who have received Adaptations or Collections from You under this License, however, will not have their licenses terminated provided such individuals or entities remain in full compliance with those licenses. Sections 1, 2, 5, 6, 7, and 8 will survive any termination of this License.

Subject to the above terms and conditions, the license granted here is perpetual (for the duration of the applicable copyright in the Work). Notwithstanding the above, Licensor reserves the right to release the Work under different license terms or to stop distributing the Work at any time; provided, however that any such election will not serve to withdraw this License (or any other license that has been, or is required to be, granted under the terms of this License), and this License will continue in full force and effect unless terminated as stated above.

8. Miscellaneous

Each time You Distribute or Publicly Perform the Work or a Collection, the Licensor offers to the recipient a license to the Work on the same terms and conditions as the license granted to You under this License.

Each time You Distribute or Publicly Perform an Adaptation, Licensor offers to the recipient a license to the original Work on the same terms and conditions as the license granted to You under this License.

If any provision of this License is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this License, and without further action by the parties to this agreement, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

No term or provision of this License shall be deemed waived and no breach consented to unless such waiver or consent shall be in writing and signed by the party to be charged with such waiver or consent.

This License constitutes the entire agreement between the parties with respect to the Work licensed here. There are no understandings, agreements or representations with respect to the Work not specified here. Licensor shall not be bound by any additional provisions that may appear in any communication from You. This License may not be modified without the mutual written agreement of the Licensor and You.

The rights granted under, and the subject matter referenced, in this License were drafted utilizing the terminology of the Berne Convention for the Protection of Literary and Artistic Works (as amended on September 28, 1979), the Rome Convention of 1961, the WIPO Copyright Treaty of 1996, the WIPO Performances and Phonograms Treaty of 1996 and the Universal Copyright Convention (as revised on July 24, 1971). These rights and subject matter take effect in the relevant jurisdiction in which the License terms are sought to be enforced according to the corresponding provisions of the implementation of those treaty provisions in the applicable national law. If the standard suite of rights granted under applicable copyright law includes additional rights not granted under this License, such additional rights are deemed to be included in the License; this License is not intended to restrict the license of any rights under applicable law.

** geonames.org; version latest -- <https://www.geonames.org>

This work is licensed under a Creative Commons Attribution 4.0 License, see <http://creativecommons.org/licenses/by/4.0/>

The Data is provided "as is" without warranty or any representation of accuracy, timeliness or completeness.

** Zetashapes; version latest -- <https://tinyurl.com/y8p87rwl>

The basic source data is from US TIGER/Line Census Data which is public domain (Q10). This site also makes use of data scraped from the flickr api.

Creative Commons Attribution 4.0 International Public License

By exercising the Licensed Rights (defined below), You accept and agree to be bound by the terms and conditions of this Creative Commons Attribution 4.0 International Public License ("Public License"). To the extent this Public License may be interpreted as a contract, You are granted the Licensed Rights in consideration of Your acceptance of these terms and conditions, and the Licensor grants You such rights in consideration of benefits the Licensor receives from making the Licensed Material available under these terms and conditions.

Section 1 – Definitions.

a. Adapted Material means material subject to Copyright and Similar Rights that is derived from or based upon the Licensed Material and in which the Licensed Material is translated, altered,

arranged, transformed, or otherwise modified in a manner requiring permission under the Copyright and Similar Rights held by the Licensor. For purposes of this Public License, where the Licensed Material is a musical work, performance, or sound recording, Adapted Material is always produced where the Licensed Material is synched in timed relation with a moving image.

b. Adapter's License means the license You apply to Your Copyright and Similar Rights in Your contributions to Adapted Material in accordance with the terms and conditions of this Public License.

c. Copyright and Similar Rights means copyright and/or similar rights closely related to copyright including, without limitation, performance, broadcast, sound recording, and Sui Generis Database Rights, without regard to how the rights are labeled or categorized. For purposes of this Public License, the rights specified in Section 2(b)(1)-(2) are not Copyright and Similar Rights.

d. Effective Technological Measures means those measures that, in the absence of proper authority, may not be circumvented under laws fulfilling obligations under Article 11 of the WIPO Copyright Treaty adopted on December 20, 1996, and/or similar international agreements.

e. Exceptions and Limitations means fair use, fair dealing, and/or any other exception or limitation to Copyright and Similar Rights that applies to Your use of the Licensed Material.

f. Licensed Material means the artistic or literary work, database, or other material to which the Licensor applied this Public License.

g. Licensed Rights means the rights granted to You subject to the terms and conditions of this Public License, which are limited to all Copyright and Similar Rights that apply to Your use of the Licensed Material and that the Licensor has authority to license.

h. Licensor means the individual(s) or entity(ies) granting rights under this Public License.

i. Share means to provide material to the public by any means or process that requires permission under the Licensed Rights, such as reproduction, public display, public performance, distribution, dissemination, communication, or importation, and to make material available to the public including in ways that members of the public may access the material from a place and at a time individually chosen by them.

j. Sui Generis Database Rights means rights other than copyright resulting from Directive 96/9/EC of the European Parliament and of the Council of 11 March 1996 on the legal protection of databases, as amended and/or succeeded, as well as other essentially equivalent rights anywhere in the world.

k. You means the individual or entity exercising the Licensed Rights under this Public License. Your has a corresponding meaning.

Section 2 – Scope.

a. License grant.

1. Subject to the terms and conditions of this Public License, the Licensors hereby grants You a worldwide, royalty-free, non-sublicensable, non-exclusive, irrevocable license to exercise the Licensed Rights in the Licensed Material to:

A. reproduce and Share the Licensed Material, in whole or in part; and

B. produce, reproduce, and Share Adapted Material.

2. Exceptions and Limitations. For the avoidance of doubt, where Exceptions and Limitations apply to Your use, this Public License does not apply, and You do not need to comply with its terms and conditions.

3. Term. The term of this Public License is specified in Section 6(a).

4. Media and formats; technical modifications allowed. The Licensors authorizes You to exercise the Licensed Rights in all media and formats whether now known or hereafter created, and to make technical modifications necessary to do so. The Licensors waives and/or agrees not to assert any right or authority to forbid You from making technical modifications necessary to exercise the Licensed Rights, including technical modifications necessary to circumvent Effective Technological Measures. For purposes of this Public License, simply making modifications authorized by this Section 2(a)(4) never produces Adapted Material.

5. Downstream recipients.

A. Offer from the Licensors – Licensed Material. Every recipient of the Licensed Material automatically receives an offer from the Licensors to exercise the Licensed Rights under the terms and conditions of this Public License.

B. No downstream restrictions. You may not offer or impose any additional or different terms or conditions on, or apply any Effective Technological Measures to, the Licensed Material if doing so restricts exercise of the Licensed Rights by any recipient of the Licensed Material.

6. No endorsement. Nothing in this Public License constitutes or may be construed as permission to assert or imply that You are, or that Your use of the Licensed Material is, connected with, or

sponsored, endorsed, or granted official status by, the Licensor or others designated to receive attribution as provided in Section 3(a)(1)(A)(i).

b. Other rights.

1. Moral rights, such as the right of integrity, are not licensed under this Public License, nor are publicity, privacy, and/or other similar personality rights; however, to the extent possible, the Licensor waives and/or agrees not to assert any such rights held by the Licensor to the limited extent necessary to allow You to exercise the Licensed Rights, but not otherwise.

2. Patent and trademark rights are not licensed under this Public License.

3. To the extent possible, the Licensor waives any right to collect royalties from You for the exercise of the Licensed Rights, whether directly or through a collecting society under any voluntary or waivable statutory or compulsory licensing scheme. In all other cases the Licensor expressly reserves any right to collect such royalties.

Section 3 – License Conditions.

Your exercise of the Licensed Rights is expressly made subject to the following conditions.

a. Attribution.

1. If You Share the Licensed Material (including in modified form), You must:

A. retain the following if it is supplied by the Licensor with the Licensed Material:

i. identification of the creator(s) of the Licensed Material and any others designated to receive attribution, in any reasonable manner requested by the Licensor (including by pseudonym if designated);

ii. a copyright notice;

iii. a notice that refers to this Public License;

iv. a notice that refers to the disclaimer of warranties;

v. a URI or hyperlink to the Licensed Material to the extent reasonably practicable;

B. indicate if You modified the Licensed Material and retain an indication of any previous modifications; and

C. indicate the Licensed Material is licensed under this Public License, and include the text of, or the URI or hyperlink to, this Public License.

2. You may satisfy the conditions in Section 3(a)(1) in any reasonable manner based on the medium, means, and context in which You Share the Licensed Material. For example, it may be reasonable to satisfy the conditions by providing a URI or hyperlink to a resource that includes the required information.

3. If requested by the Licensor, You must remove any of the information required by Section 3(a)(1) (A) to the extent reasonably practicable.

4. If You Share Adapted Material You produce, the Adapter's License You apply must not prevent recipients of the Adapted Material from complying with this Public License.

Section 4 – Sui Generis Database Rights.

Where the Licensed Rights include Sui Generis Database Rights that apply to Your use of the Licensed Material:

- a. for the avoidance of doubt, Section 2(a)(1) grants You the right to extract, reuse, reproduce, and Share all or a substantial portion of the contents of the database;
- b. if You include all or a substantial portion of the database contents in a database in which You have Sui Generis Database Rights, then the database in which You have Sui Generis Database Rights (but not its individual contents) is Adapted Material; and
- c. You must comply with the conditions in Section 3(a) if You Share all or a substantial portion of the contents of the database.

For the avoidance of doubt, this Section 4 supplements and does not replace Your obligations under this Public License where the Licensed Rights include other Copyright and Similar Rights.

Section 5 – Disclaimer of Warranties and Limitation of Liability.

a. Unless otherwise separately undertaken by the Licensor, to the extent possible, the Licensor offers the Licensed Material as-is and as-available, and makes no representations or warranties of any kind concerning the Licensed Material, whether express, implied, statutory, or other. This includes, without limitation, warranties of title, merchantability, fitness for a particular purpose, non-infringement, absence of latent or other defects, accuracy, or the presence or absence of errors, whether or not known or discoverable. Where disclaimers of warranties are not allowed in full or in part, this disclaimer may not apply to You.

b. To the extent possible, in no event will the Licensor be liable to You on any legal theory (including, without limitation, negligence) or otherwise for any direct, special, indirect, incidental, consequential, punitive, exemplary, or other losses, costs, expenses, or damages arising out of this Public License or use of the Licensed Material, even if the Licensor has been advised of the possibility of such losses, costs, expenses, or damages. Where a limitation of liability is not allowed in full or in part, this limitation may not apply to You.

c. The disclaimer of warranties and limitation of liability provided above shall be interpreted in a manner that, to the extent possible, most closely approximates an absolute disclaimer and waiver of all liability.

Section 6 – Term and Termination.

a. This Public License applies for the term of the Copyright and Similar Rights licensed here. However, if You fail to comply with this Public License, then Your rights under this Public License terminate automatically.

b. Where Your right to use the Licensed Material has terminated under Section 6(a), it reinstates:

1. automatically as of the date the violation is cured, provided it is cured within 30 days of Your discovery of the violation; or
2. upon express reinstatement by the Licensor.

For the avoidance of doubt, this Section 6(b) does not affect any right the Licensor may have to seek remedies for Your violations of this Public License.

c. For the avoidance of doubt, the Licensor may also offer the Licensed Material under separate terms or conditions or stop distributing the Licensed Material at any time; however, doing so will not terminate this Public License.

d. Sections 1, 5, 6, 7, and 8 survive termination of this Public License.

Section 7 – Other Terms and Conditions.

a. The Licensor shall not be bound by any additional or different terms or conditions communicated by You unless expressly agreed.

b. Any arrangements, understandings, or agreements regarding the Licensed Material not stated herein are separate from and independent of the terms and conditions of this Public License.

Section 8 – Interpretation.

- a. For the avoidance of doubt, this Public License does not, and shall not be interpreted to, reduce, limit, restrict, or impose conditions on any use of the Licensed Material that could lawfully be made without permission under this Public License.
- b. To the extent possible, if any provision of this Public License is deemed unenforceable, it shall be automatically reformed to the minimum extent necessary to make it enforceable. If the provision cannot be reformed, it shall be severed from this Public License without affecting the enforceability of the remaining terms and conditions.
- c. No term or condition of this Public License will be waived and no failure to comply consented to unless expressly agreed to by the Licensor.
- d. Nothing in this Public License constitutes or may be interpreted as a limitation upon, or waiver of, any privileges and immunities that apply to the Licensor or You, including from the legal processes of any jurisdiction or authority.

** Microsoft Sample Databases: World Wide Importers; version latest -- <https://tinyurl.com/ybat5d8k>

Microsoft SQL Server Sample Code

Copyright (c) Microsoft Corporation

All rights reserved.

MIT License.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

**** Microsoft Sample Databases: AdventureWorks; version last -- <https://github.com/Microsoft/sql-server-samples/releases/tag/adventureworks2014>**

Microsoft SQL Server Sample Code

Copyright (c) Microsoft Corporation

All rights reserved.

Microsoft Public License (Ms-PL)

This license governs use of the accompanying software. If you use the software, you accept this license. If you do not accept the license, do not use the software.

1. Definitions

The terms "reproduce," "reproduction," "derivative works," and "distribution" have the same meaning here as under U.S. copyright law.

A "contribution" is the original software, or any additions or changes to the software.

A "contributor" is any person that distributes its contribution under this license.

"Licensed patents" are a contributor's patent claims that read directly on its contribution.

2. Grant of Rights

(A) Copyright Grant- Subject to the terms of this license, including the license conditions and limitations in section 3, each contributor grants you a non-exclusive, worldwide, royalty-free copyright license to reproduce its contribution, prepare derivative works of its contribution, and distribute its contribution or any derivative works that you create.

(B) Patent Grant- Subject to the terms of this license, including the license conditions and limitations in section 3, each contributor grants you a non-exclusive, worldwide, royalty-free license under its licensed patents to make, have made, use, sell, offer for sale, import, and/or otherwise dispose of its contribution in the software or derivative works of the contribution in the software.

3. Conditions and Limitations

(A) No Trademark License- This license does not grant you rights to use any contributors' name, logo, or trademarks.

(B) If you bring a patent claim against any contributor over patents that you claim are infringed by the software, your patent license from such contributor to the software ends automatically.

(C) If you distribute any portion of the software, you must retain all copyright, patent, trademark, and attribution notices that are present in the software.

(D) If you distribute any portion of the software in source code form, you may do so only under this license by including a complete copy of this license with your distribution. If you distribute any portion of the software in compiled or object code form, you may only do so under a license that complies with this license.

(E) The software is licensed "as-is." You bear the risk of using it. The contributors give no express warranties, guarantees or conditions. You may have additional consumer rights under your local laws which this license cannot change. To the extent permitted under your local laws, the contributors exclude the implied warranties of merchantability, fitness for a particular purpose and non-infringement.

** OurAirports; version latest -- <https://tinyurl.com/y8p87rwl>

Made with OurAirports. <http://ourairports.com/data/>

We hereby release all of these files into the Public Domain, with no warranty of any kind — By downloading any of these files, you agree that OurAirports.com, Megginson Technologies Ltd., and anyone involved with the web site or company hold no liability for anything that happens when you use the data, including (but not limited to) computer damage, lost revenue, flying into cliffs, or a general feeling of drowsiness that persists more than two days.

** Natural Earth; version latest -- <https://tinyurl.com/y8p87rwl>

Made with Natural Earth. Free vector and raster map data @ naturalearthdata.com.

All versions of Natural Earth raster + vector map data found on this

website are in the public domain. You may use the maps in any manner, including modifying the content and design, electronic dissemination, and offset printing. The primary authors, Tom Patterson and Nathaniel Vaughn Kelso, and all other contributors renounce all financial claim to the maps and invites you to use them for personal, educational, and commercial purposes.

No permission is needed to use Natural Earth. Crediting the authors is unnecessary.