

Amazon QuickSight



Amazon QuickSight: User Guide

Table of Contents

.....	xiv
What is Amazon QuickSight?	1
Why QuickSight?	1
Starting work with QuickSight	3
How it works	4
Terminology	4
Using sample data	5
Reader experience: exploring dashboards	6
Interacting with dashboards	7
Using filters	8
Filtering data during your session	10
Using dashboard elements	14
Sorting data	16
Exporting and printing interactive dashboard reports	17
Generate an executive summary of a dashboard	22
Interacting with paginated reports	23
Exporting and printing reports	24
Subscribing to dashboard emails and alerts	28
Getting email reports	28
Sign up for anomaly alerts	29
Bookmarking views of a dashboard	30
Creating bookmarks	30
Updating bookmarks	31
Renaming bookmarks	32
Making a bookmark the default view	33
Sharing bookmarks	34
Deleting bookmarks	35
Setting up	37
Sign up for Amazon	37
Create an IAM user	37
Sign up for an Amazon Web Services account	37
Secure IAM users	38
Integrating with IAM Identity Center	38
Signing up for a subscription	38

Getting started	42
Signing in to QuickSight	42
How to sign in to QuickSight	43
Quick start: Create an analysis using sample data	48
Create a dashboard using sample data	52
Tutorial: Create a prepared dataset	53
Tutorial: Create an analysis	59
Tutorial: Modify visuals	63
Tutorial: Create a dashboard	72
Using the console	73
Using the Amazon QuickSight menu and landing page	74
Using the Amazon QuickSight start page	76
Choosing a language in Amazon QuickSight	80
Using the Amazon QuickSight mobile app	82
Connecting to data	83
Supported data sources	85
Connecting to relational data	85
Importing file data	87
Software as a service (SaaS) data	88
.....	89
Data source quotas	89
SPICE quotas for imported data	90
Quotas for direct SQL queries	90
Supported data types and values	91
String and text data	91
Date and time data	92
Numeric data	92
Supported data types from other data sources	94
Connection examples	105
Amazon Athena	105
Amazon OpenSearch Service	107
Amazon S3 files	110
Apache Spark	122
Databricks	123
Google BigQuery	128
Microsoft Excel files	130

Presto	131
Starburst	133
SaaS sources	138
Salesforce	140
Trino	141
Text files	145
Timestream data	145
Creating datasets	149
From new data sources	149
From existing data sources	156
From existing datasets	160
Editing datasets	164
Editing a dataset from the Datasets page	164
Editing a dataset in an analysis	164
Reverting datasets	165
Troubleshooting	166
Duplicating datasets	167
Changing datasets	167
Sharing datasets	169
Sharing a dataset	169
Viewing and editing the permissions of users that a dataset is shared with	170
Revoking access to a dataset	170
Tracking dataset assets	170
Dataset parameters	171
Dataset parameter limitations	172
Create dataset parameters	172
Insert dataset parameters into custom SQL	173
Add dataset parameters to calculated fields	176
Add dataset parameters to filters	176
Working with dataset parameters in QuickSight analyses	179
Advanced use	182
Using row-level security (RLS)	188
Using user-based rules	188
Using tag-based rules	196
Column-level security (CLS)	208
Run-as role	210

Athena data sources	211
Amazon Redshift data sources	213
Amazon S3 data sources	216
Deleting datasets	219
Adding a dataset to an analysis	220
Add or edit a dataset	221
Replacing datasets	222
Remove a dataset from an analysis	224
Working with data sources	225
Creating a data source	226
Editing a data source	228
Deleting a data source	230
Refreshing data	232
Importing data into SPICE	233
Estimating the size of SPICE datasets	234
Refreshing SPICE data	235
Refreshing a dataset	235
Incrementally refreshing a dataset	236
Refreshing a dataset during data preparation	238
Refreshing a dataset on a schedule	239
Incrementally refreshing a dataset on a schedule	240
Using SPICE data in an analysis	242
View SPICE ingestion history	243
Troubleshooting skipped row errors	245
SPICE ingestion error codes	248
Row import errors	248
Data import errors	249
Updating files in a dataset	252
Preparing data	254
Describing data	254
Choosing file upload settings	256
Changing text file upload settings	256
Changing Microsoft Excel file upload settings	256
Preparing data fields	257
Editing field names and descriptions	257
Setting fields as a dimensions or measures	259

Changing a field data type	260
Adding drill-downs	262
Selecting fields	266
Organizing fields into folders	267
Mapping and joining fields	271
Adding calculations	273
Adding calculated fields	273
Order of evaluation	283
Level-aware calculations	286
Functions and operators	294
Previewing tables in a dataset	525
Joining data	526
Types of joined datasets	526
Facts about joining datasets	527
Creating a join	528
Join types	533
Filtering data	536
Viewing existing filters	537
Adding filters	541
Cross-sheet filters and controls	546
Filter types	557
Adding filter controls	571
Editing filters	580
Enabling or disabling filters	584
Deleting filters	588
Using SQL to customize data	592
Creating a basic SQL query	593
Adding geospatial data	594
Changing a geospatial grouping	600
Geospatial troubleshooting	600
Using unsupported or custom dates	605
Integrating SageMaker models	607
How SageMaker integration works	608
Costs incurred (no additional costs with integration itself)	609
Usage guidelines	610
Defining the schema file	611

Adding a SageMaker model to your QuickSight dataset	614
SageMaker Canvas	615
Preparing dataset examples	617
Preparing a dataset based on file data	618
Preparing a dataset based on Salesforce data	621
Preparing a dataset based on database data	622
Visualizing data	625
Working with an analysis	625
Starting an analysis	626
Adding titles and descriptions to an analysis	629
Renaming an analysis	629
Duplicating analyses	630
Viewing analysis details	630
Date and time settings	630
The analysis menu	634
Saving changes to analyses	635
Exporting data from analyses	636
Deleting an analysis	639
Adding sheets	639
Working with interactive sheets in Amazon QuickSight	641
Arranging visuals in an interactive dashboard	641
Working with paginated reports in Amazon QuickSight	702
Getting started	703
Creating reports from an analysis in Amazon QuickSight	704
Formatting reports in Amazon QuickSight	707
Consuming paginated reports in Amazon QuickSight	716
Working with items on sheets	718
Adding visuals	718
Using Q Topics	732
Visual types	733
Formatting	888
Customizing data presentation	973
Using themes in Amazon QuickSight	1008
Keyboard shortcuts	1015
Using shortcuts within a visual	1016
Gaining insights with ML	1018

Understanding the ML algorithm	1019
What's the difference between anomaly detection and forecasting?	1020
What RCF is and what it does	1021
How RCF is applied to detect anomalies	1022
How RCF is applied to generate forecasts	1022
References for machine learning and RCF	1023
Dataset requirements	1024
Adding insights	1025
Adding suggested insights	1025
Adding custom insights to your analysis	1028
Autonarratives	1029
Insights that include autonarratives	1030
Working with the expression editor screen and menus	1031
Adding URLs	1034
Walkthrough: Use the narrative expression editor	1035
Computations	1039
Detecting outliers	1066
Concepts for anomaly or outlier detection	1067
Setting up ML-powered anomaly detection for outlier analysis	1068
Exploring outliers and key drivers	1074
ML-powered forecasts and what-ifs	1081
Using forecasts and what-if scenarios	1081
Sharing data	1086
Sharing Amazon QuickSight analyses	1087
Sharing an analysis	1087
Viewing the users that an analysis is shared with	1089
Revoking access to an analysis	1089
Publishing dashboards	1089
Copying a dashboard	1094
Deleting dashboards	1095
Sharing dashboards	1096
Granting access to a dashboard	1097
Sharing a link to the dashboard	1097
Viewing who has access	1098
Revoking access	1099
With individual users and groups	1099

With everyone in your account	1102
With anyone on the internet	1104
With everyone in your account with the API	1111
With anyone on the internet using the API	1113
Share your view of a dashboard	1115
Sending reports	1116
Report billing	1117
Configuring email reports	1117
Subscribing to reports	1125
Threshold alerts	1126
Creating Alerts	1127
Managing Threshold Alerts	1131
Investigating Alert Failures	1133
Alert Scheduling	1134
Alert Permissions	1134
Printing	1134
Exporting as PDFs	1135
PDF Error codes	1136
Organizing assets into folders	1137
Overview of QuickSight folders	1138
Permissions overview for shared folders	1142
Create and manage permissions for shared folders	1143
Considerations	1144
Creating scaled folders with the Amazon CLI	1145
Monitoring data	1149
Accessing metrics in CloudWatch	1149
Metrics	1149
Per-dashboard metrics	1149
Per-dataset ingestion metrics	1151
Per-visual metrics	1152
Aggregate metrics	1154
Aggregate dashboard metrics	1154
Aggregate ingestion metrics	1155
Aggregate visual metrics	1157
Aggregate SPICE metrics	1158
Dimensions	1158

Graphing metrics by using the CloudWatch console	1159
Creating alarms by using the CloudWatch console	1159
Developing with Amazon QuickSight	1160
Required knowledge	1160
Available API operations for Amazon QuickSight	1160
Terminology and concepts	1162
QuickSight dev portal	1164
Coding with the QuickSight APIs	1164
Events integration	1172
Supported events	1173
Example event payload	1184
Creating rules to send events to Amazon CloudWatch	1184
Creating rules to send events to Amazon Lambda	1185
Embedded analytics	1188
Embedding overview	1190
Customize embedded assets	1191
1-click embedding and public embedding	1205
Embedding with the APIs	1217
Troubleshooting	1432
Resolving Amazon QuickSight issues and error messages	1432
Athena issues	1433
Athena column not found	1433
Athena invalid data	1434
Athena query timeout	1434
Athena staging bucket missing	1435
Amazon Glue table incompatible with Athena	1435
Athena Table not found	1442
Workgroup and output errors when using Athena with Amazon QuickSight	1442
Data source connectivity issues	1443
I can't connect although my data source connection options look right (SSL)	1444
I can't connect to Amazon Athena	1446
I can't connect to Amazon S3	1450
I can't create or refresh a dataset from an existing Adobe Analytics data source	1452
I need to validate the connection to my data source, or change data source settings	1453
I can't connect to MySQL (issues with SSL and authorization)	1453
I can't connect to RDS	1455

Login issues	1455
Insufficient permissions with Athena	1456
Amazon QuickSight isn't working in my browser	1457
How do I delete my Amazon QuickSight account?	1457
Individuals in my organization get "External Login is Unauthorized"	1458
My email sign-in stopped working	1463
Visual issues	1463
I can't see my visuals	1464
I get a feedback bar across my printed documents	1464
My map charts don't show locations	1465
My pivot table stops working	1465
My visual can't find missing columns	1465
My visual can't find the query table	1466
My visual doesn't update after I change a calculated field	1466
Values with scientific notation don't format correctly	1467
Amazon QuickSight Administration	1468
Different editions of Amazon QuickSight	1468
Availability of editions	1468
User management between editions	1469
Permissions for the different editions	1470
Regions and IP ranges	1471
Supported browsers	1474
Managing QuickSight	1475
Managing assets	1475
Managing your subscriptions	1477
SPICE capacity	1485
Manage account settings	1490
Domains and Embedding	1495
Multitenancy and namespaces	1498
Account customizations	1502
Welcome content	1503
Email report templates	1504
Default analysis theme	1513
Tracking cost and usage data	1515
Amazon security	1516
Data protection	1516

Data encryption	1518
Inter-network traffic privacy	1526
Accessing data sources	1528
Identity and access management	1567
Service control policies (SCP)	1568
IAM	1571
Identity management	1602
Managing user access	1634
Turning on IP and VPC endpoint restrictions	1649
Customizing access to the QuickSight console	1652
Incident response, logging, and monitoring	1656
Logging operations with Amazon CloudTrail	1657
Compliance validation	1663
Resilience	1664
Infrastructure security	1665
Network and database configuration requirements	1666
Connecting to a VPC with QuickSight	1672
Best practices	1704
Amazon managed policies	1705
AWSQuickSightElasticsearchPolicy	1706
AWSQuickSightOpenSearchPolicy	1707
AWSQuickSightSageMakerPolicy	1708
AWSQuickSightAssetBundleExportPolicy	1709
AWSQuickSightAssetBundleImportPolicy	1710
Policy updates	1710
Amazon Glossary	1712
Document history	1713
Previous updates	1788
Attributions	1801

Important: 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.

To find a feature or item, use the [Quick search bar](#).

For more information on QuickSight's new look, see [Introducing new analysis experience on Amazon QuickSight](#).

What is Amazon QuickSight?

Amazon QuickSight is a cloud-scale business intelligence (BI) service that you can use to deliver easy-to-understand insights to the people who you work with, wherever they are. Amazon QuickSight connects to your data in the cloud and combines data from many different sources. In a single data dashboard, QuickSight can include Amazon data, third-party data, big data, spreadsheet data, SaaS data, B2B data, and more. As a fully managed cloud-based service, Amazon QuickSight provides enterprise-grade security, global availability, and built-in redundancy. It also provides the user-management tools that you need to scale from 10 users to 10,000, all with no infrastructure to deploy or manage.

QuickSight gives decision-makers the opportunity to explore and interpret information in an interactive visual environment. They have secure access to dashboards from any device on your network and from mobile devices.

To learn more about the major components and processes of Amazon QuickSight and the typical workflow for creating data visualizations, see the following sections. Get started today to unlock the potential of your data and make the best decisions that you can.

Topics

- [Why QuickSight?](#)
- [Starting work with QuickSight](#)

Why QuickSight?

Every day, the people in your organization make decisions that affect your business. When they have the right information at the right time, they can make the choices that move your company in the right direction.

Here are some of the benefits of using Amazon QuickSight for analytics, data visualization, and reporting:

- The in-memory engine, called SPICE, responds with blazing speed.
- No upfront costs for licenses and a low total cost of ownership (TCO).
- Collaborative analytics with no need to install an application.
- Combine a variety of data into one analysis.

- Publish and share your analysis as a dashboard.
- Control features available in a dashboard.
- No need to manage granular database permissions—dashboard viewers can see only what you share.

For advanced users, QuickSight Enterprise edition offers even more features:

- Saves you time and money with automated and customizable data insights, powered by machine learning (ML). This enables your organization to do the following, without requiring any knowledge of machine learning:
 - Automatically make reliable forecasts.
 - Automatically identify outliers.
 - Find hidden trends.
 - Act on key business drivers.
 - Translate data into easy-to-read narratives, like headline tiles for your dashboard.
- Provides extra Enterprise security features, including the following:
 - Federated users, groups, and single sign-on (IAM Identity Center) with Amazon Identity and Access Management (IAM) Federation, SAML, OpenID Connect, or Amazon Directory Service for Microsoft Active Directory.
 - Granular permissions for Amazon data access.
 - Row level security.
 - Highly secure data encryption at rest.
 - Access to Amazon data and on-premises data in Amazon Virtual Private Cloud
- Offers pay-per-session pricing for the users that you place in the "reader" security role—*readers* are dashboard subscribers, people who view reports but don't create them.
- Empowers you to make QuickSight part of your own websites and applications by deploying embedded console analytics and dashboard sessions.
- Makes our business your business with multitenancy features for value-added resellers (VARs) of analytical services.
- Enables you to programmatically script dashboard templates that can be transferred to other Amazon accounts.
- Simplifies access management and organization with shared and personal folders for analytical assets.

- Enables larger data import quotas for SPICE data ingestion and more frequently scheduled data refreshes.

To discover the power of end-to-end BI from Amazon, sign up at <http://www.amazonaws.cn/QuickSight>.

Starting work with QuickSight

To start work with QuickSight, we recommend that you read the following sections:

- [How Amazon QuickSight works](#)– Learn essential terminology and how QuickSight components work together.
- [Getting started with Amazon QuickSight data analysis](#) – Complete important setup tasks and learn how to use a dashboard, create an analysis, and publish a dashboard.
- [Amazon security in Amazon QuickSight](#) – Understand how you can help to secure access to data in QuickSight.

How Amazon QuickSight works

Using Amazon QuickSight, 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 looks like this:

1. Create a new analysis.
2. Add new or existing datasets.
3. Choose fields to create the first chart. QuickSight 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.

The following illustration shows the basic workflow.

Terminology

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 (Super-fast, Parallel, In-memory Calculation Engine) is the robust in-memory engine that QuickSight 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.

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.

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 (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.

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.

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

Using sample data

To get a first look at how QuickSight works, you can explore Amazon QuickSight 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 QuickSight, for example the [Amazon public datasets](#). These datasets come in a variety of formats.

Reader experience: Exploring interactive dashboards in Amazon QuickSight

 **Intended audience:** Amazon QuickSight Dashboard subscribers or viewership

In Amazon QuickSight, 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 QuickSight, 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 Amazon QuickSight 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.

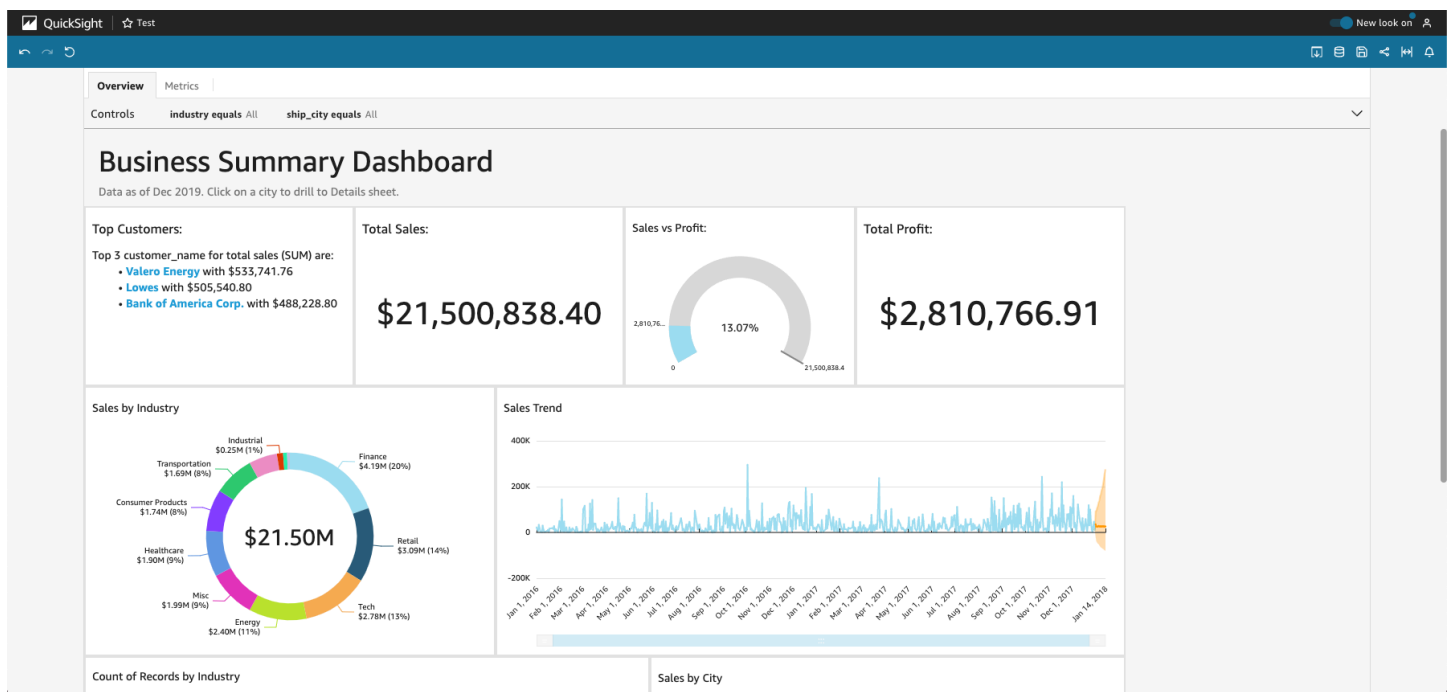
To learn to set up dashboards, see [Sharing and subscribing to data in Amazon QuickSight](#).

Topics

- [Interacting with Amazon QuickSight dashboards](#)
- [Interacting with paginated reports in Amazon QuickSight](#)
- [Subscribing to dashboard emails and alerts](#)
- [Bookmarking views of a dashboard](#)

Interacting with Amazon QuickSight 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. When you open the dashboard, the screen should look something like the following example.



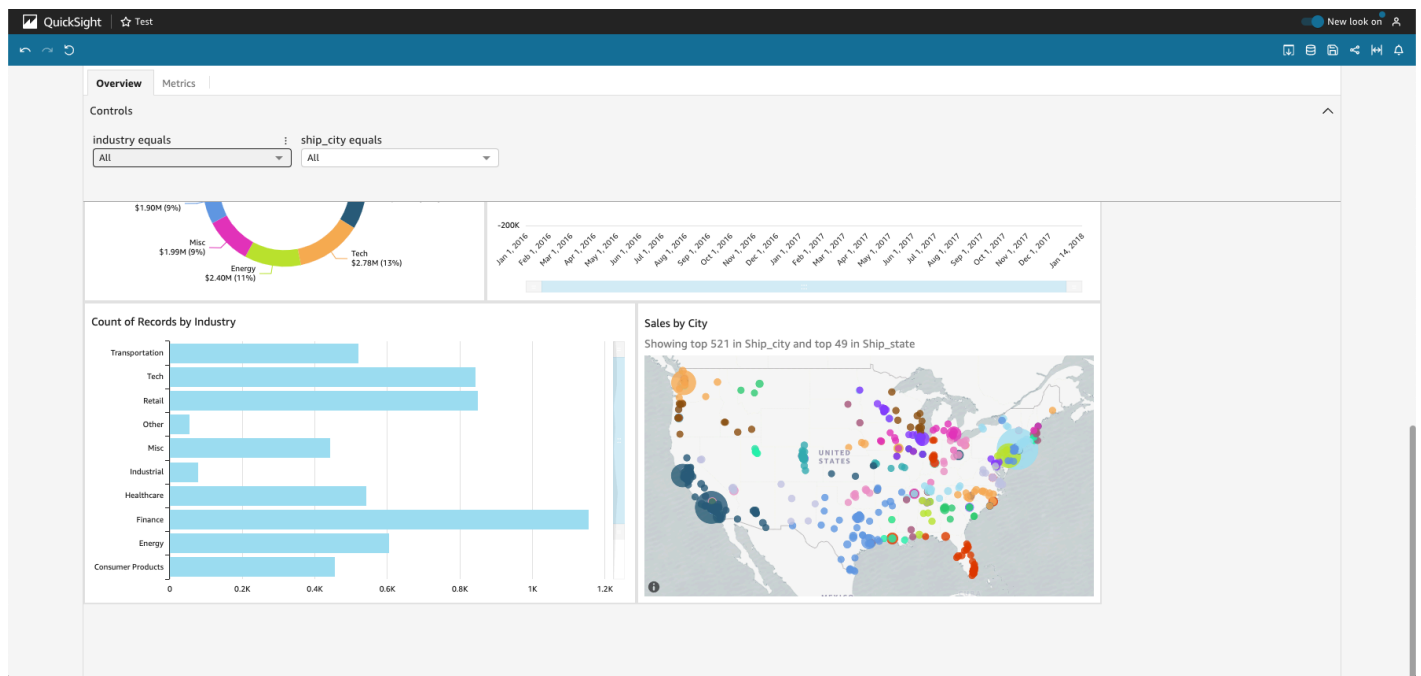
To fit the dashboard to your screen, open the **View** menu at upper right and select **Fit to window**.

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 QuickSight displays. It also has links to the Amazon QuickSight **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. The following screenshot shows more of the previous example dashboard:




Using filters on 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 QuickSight 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

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.



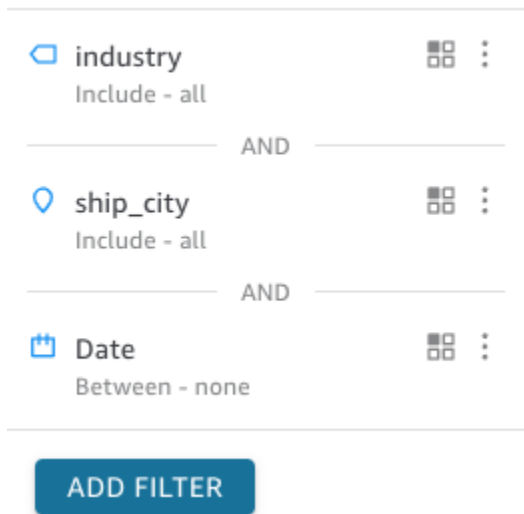
The scope icons match the ones that display on the filter menu when you are choosing the scope for the filter.

Viewing filter details

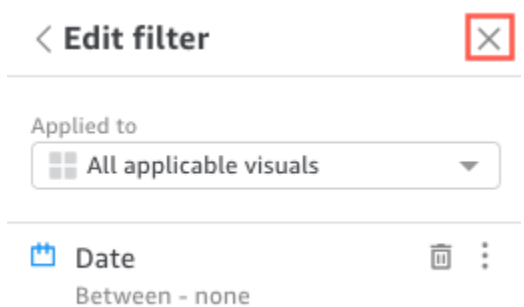
To see filter details, choose **Filter** at left. The filter view retains your last selection. So when you open **Filter**, you see either the **Applied filters** or the **Edit filter** view.

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.

Filters




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

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.

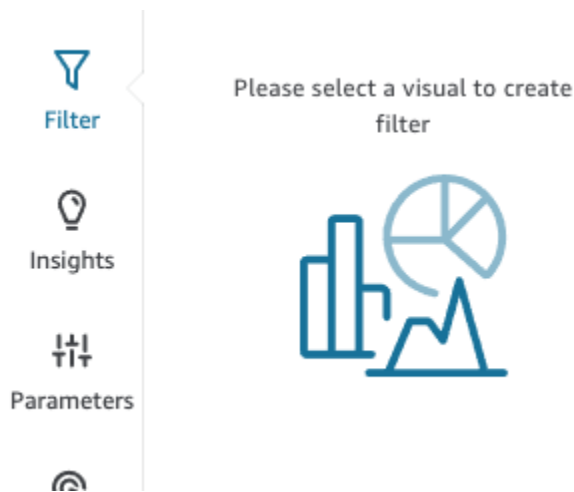


3. You can create your own filters by using the filter panel on the left side of the page. The filter icon looks like the following.

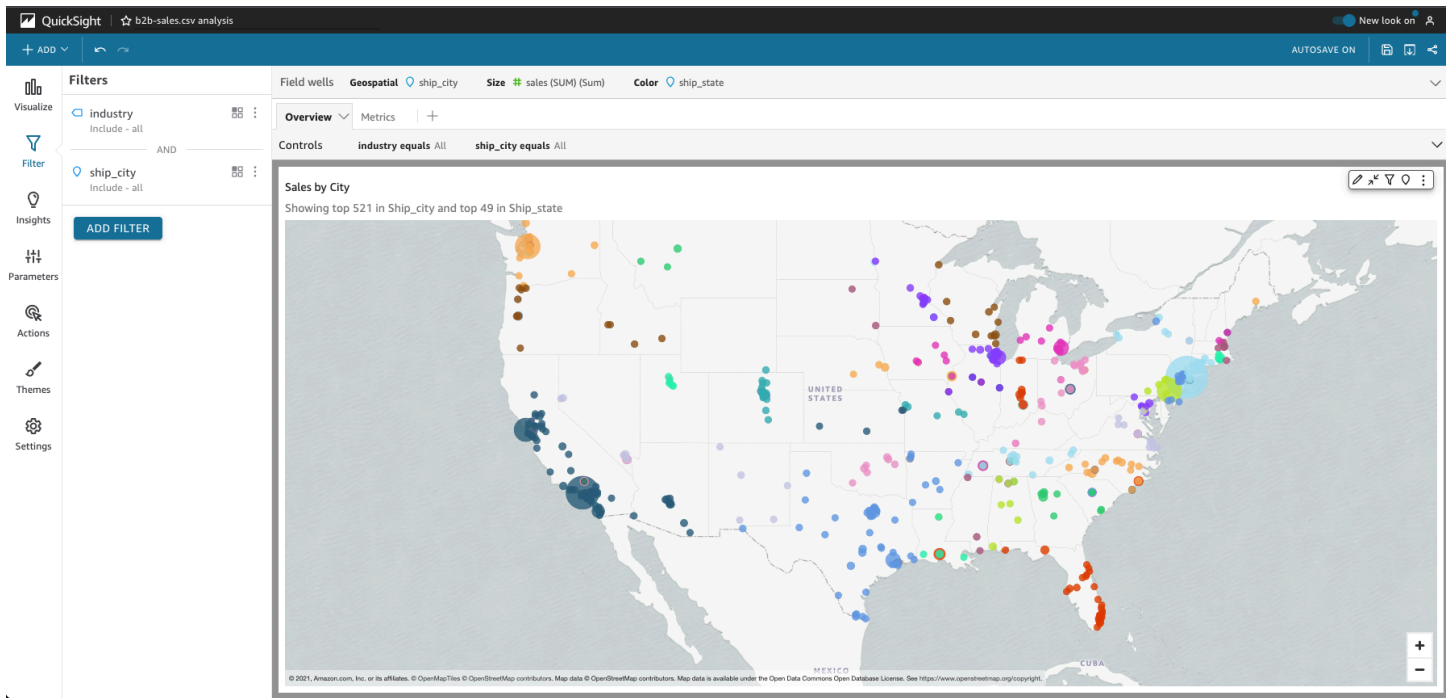


To create a filter, choose the **Filter** icon at left.

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.

< **Edit filter** ✕

Applied to
📄 Only this visual ▼

📅 **Date** 🗑️ ⋮
Between - none

Filter type
Date & time range ▼

Condition
Between ▼

Use parameters

Time granularity
Day ▼

Start date
YYYY/MM/DD

Include start date

End date
YYYY/MM/DD

Include end date

Null options
Exclude nulls ▼

OR

ADD FILTER CONDITION ▼

Note: There are limitations on how you can group filters.
[Learn more](#)


APPLY DELETE 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 QuickSight](#).


Using the elements on the 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. The following screenshot shows a list of actions available on most chart types.

Sales Trend



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

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 dashboard reports

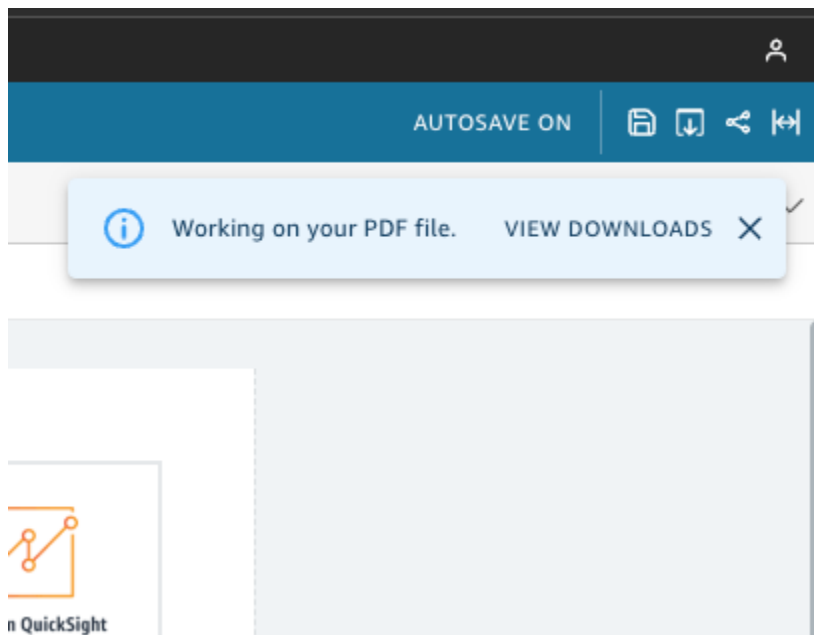
You can export or print a PDF version of an interactive dashboard. You can also export some visuals in a dashboard to a CSV. Exporting an entire dashboard to a CSV is not currently supported for interactive dashboards.

To export an interactive dashboard report as a PDF

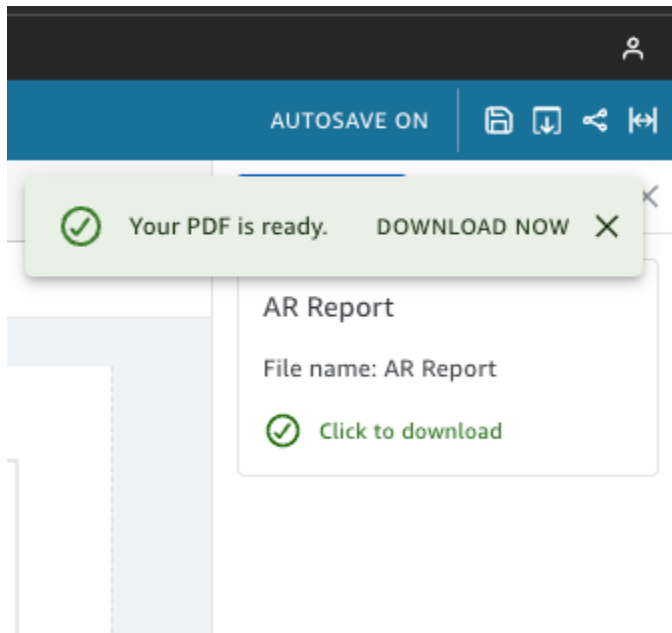
1. From the dashboard report that you want to export, choose the **Export** icon at the top right.
2. Choose **Generate PDF**.




3. When you choose **Generate PDF**, QuickSight 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.

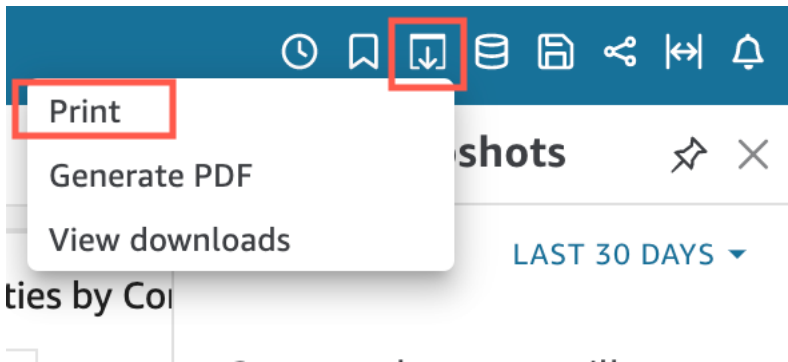


- Choose the **Export** icon at the top right, and then choose **View downloads** to view and download every report that is ready to download.

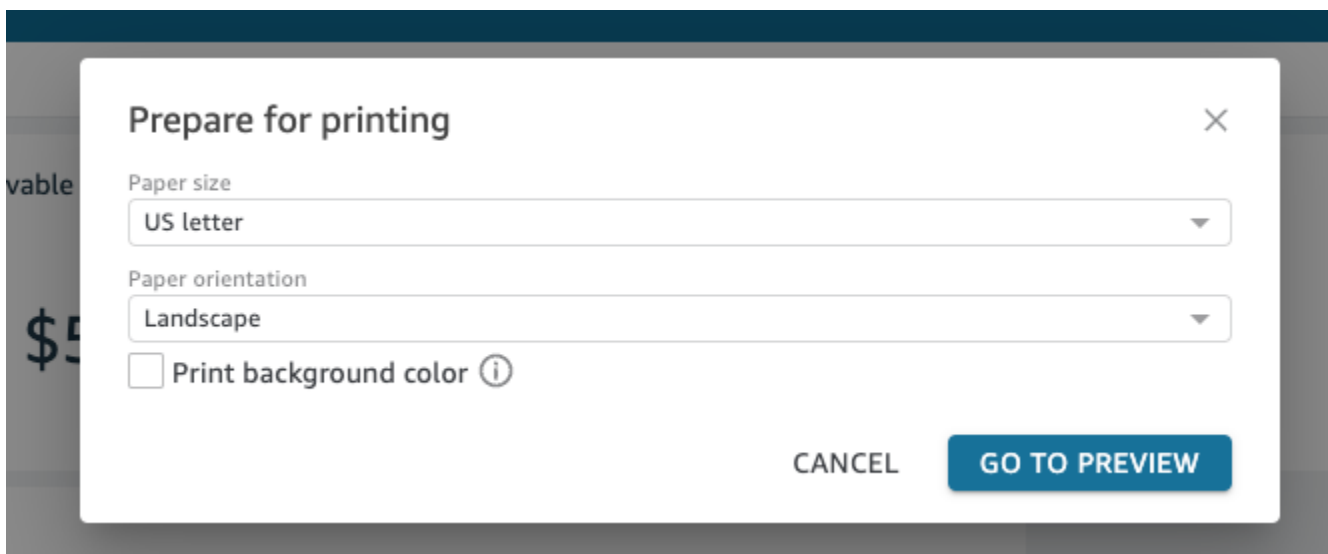


To print an interactive dashboard report

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**.


Channel	Segment	Cost	Cost
API	Enterprise	697,877.37	756,282.96
	SMB	182,391.93	147,228.72
	Startup	1,478,997.31	441,312.46
Mobile	Enterprise	1,548,957.28	1,475,086.85
	SMB	419,653.32	369,625.7
	Startup	3,322,570.56	1,088,844.05
Web	Enterprise	2,198,689.04	1,823,214.96
	SMB	589,733.39	444,692.8

Generate an executive summary of a dashboard

This feature for Amazon QuickSight is a public preview release and is subject to change. We recommend that you use this tool only for testing, and not in production environments.

Dashboard readers can generate executive summaries that provide a summary of all insights that QuickSight 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 does not 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, QuickSight readers can copy the summary to their clipboard in order to share with others, or include in a story. For more information about QuickSight stories, see [???](#).

Interacting with paginated reports in Amazon QuickSight

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. When you open the report, the screen should look something like the following example.

QuickSight | AWS Statistics / Original dashboard (Modified)

Software Sales Type a question about your data

Dashboard | Report | Table

80% Page 1 of 6 Nov 18, 2022 4:23PM PST

Attendance Statistics

Between 29-JUN-2022 and 03-JUL-2022

PROPERTY: AWS
USER: RS
CURRENCY: EUR

NET OR GROSS: Net
GROUP BY: Room Type

History

Date	Room	Avail Rooms	OOD Rms	Rms Sold	Rms to Sell	Occ	Cmp Rms	No Show Rms	Day Use Rms	Trans Rms	Grp Pickup Rms	Grp Rms	Centr Pickup Rms	Centr Rms	Room Revenue	ADR	RevPAR
6/29/2021 Tue	J51KT	5	1	4	0	5.33	0	0	1	2	2	0	0	0	€7,672.72	€1,518.18	€1,518.18
	RO1KT	2	0	0	2	0	0	0	1	0	0	0	0	0	€0.00	€0.00	€0.00
	RO2KT	8	0	5	3	6.67	0	0	1	3	2	0	0	0	€4,009.09	€801.82	€501.14
	RO3KT	4	0	1	3	1.33	0	0	1	0	1	0	0	0	€790.92	€790.92	€197.73
	RO4KT	7	0	7	0	9.33	0	0	1	6	1	0	0	0	€7,472.73	€1,067.53	€1,067.53
	RO5KT	2	0	1	1	1.33	0	0	1	1	0	0	0	0	€1,018.18	€1,018.18	€509.09
	SP1KT	2	0	2	0	2.67	0	0	1	1	1	0	0	0	€5,109.10	€2,554.55	€2,554.55
	SP2KT	1	0	0	1	0	0	0	1	0	0	0	0	0	€0.00	€0.00	€0.00
	ST1KT	3	0	2	1	2.67	0	0	1	2	0	0	0	0	€2,890.91	€1,445.46	€963.64
	ST2KT	6	0	4	2	5.33	0	0	1	0	4	0	0	0	€4,545.45	€1,136.37	€757.58
	ST3KT	3	0	3	0	4	0	0	1	3	0	0	0	0	€6,763.63	€2,254.54	€2,254.54
	ST4KT	11	0	9	2	12	0	0	1	3	6	0	0	0	€14,245.47	€1,582.83	€1,295.04
	ST5KT	14	1	10	3	13.33	0	0	1	5	5	0	0	0	€24,397.56	€2,439.76	€1,742.68
	ST6KT	1	0	1	0	1.33	0	0	1	1	0	0	0	0	€1,400.00	€1,400.00	€1,400.00
TS1KT	1	0	1	0	1.33	0	0	1	1	0	0	0	0	€3,490.91	€3,490.91	€3,490.91	
TS2KT	1	0	1	0	1.33	0	0	1	0	1	0	0	0	€72.72	€72.72	€72.72	
TS3KT	1	0	1	0	1.33	0	0	1	1	0	0	0	0	€2,409.09	€2,409.09	€2,409.09	
VL1KT	1	0	1	0	1.33	0	0	1	1	0	0	0	0	€2,177.27	€2,177.27	€2,177.27	
VL2KT	1	0	0	1	0	0	0	1	0	0	0	0	0	€0.00	€0.00	€0.00	

Nov 19 2022 12:23 Private and Confidential Page 1

Attendance Statistics

Date	Room	Avail Rooms	OOD Rms	Rms Sold	Rms to Sell	Occ	Cmp Rms	No Show Rms	Day Use Rms	Trans Rms	Grp Pickup Rms	Grp Rms	Centr Pickup Rms	Centr Rms	Room Revenue	ADR	RevPAR
6/29/2021 Tue	ZSTRK	4	0	0	0	0	0	0	1	0	0	0	0	0	€0.00	€0.00	€0.00
	ZSTRK	5	0	0	0	0	0	0	1	0	0	0	0	0	€0.00	€0.00	€0.00
	Total	93	2	53	19	70.64	0	0	21	30	23	0	0	0	€88,320.54	€26,414.69	€22,792...
6/30/2021 Wed	J51KT	5	1	4	0	5.33	0	0	1	2	2	0	0	0	€6,809.09	€1,702.27	€1,361.82
	RO1KT	2	0	0	2	0	0	0	1	0	0	0	0	0	€0.00	€0.00	€0.00
	RO2KT	8	0	4	4	5.33	0	0	1	3	1	0	0	0	€2,909.09	€727.27	€363.64

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.


— + 100% Page 1 of 6 Nov 18, 2022 9:50AM PST

Exporting and printing reports

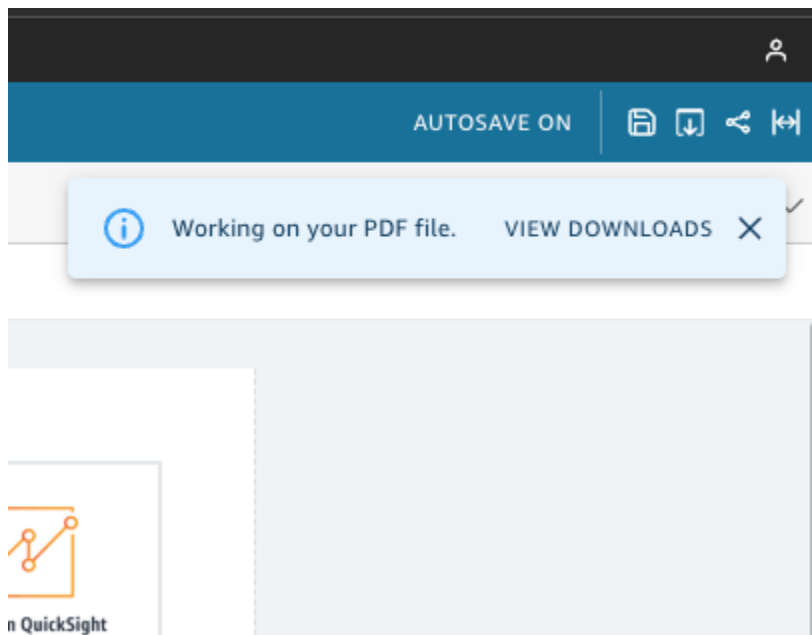
Paginated reports are designed to be viewed from a specific point of time. These reports, or snapshots, can be printed or downloaded as a PDF or CSV.

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**, QuickSight 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.

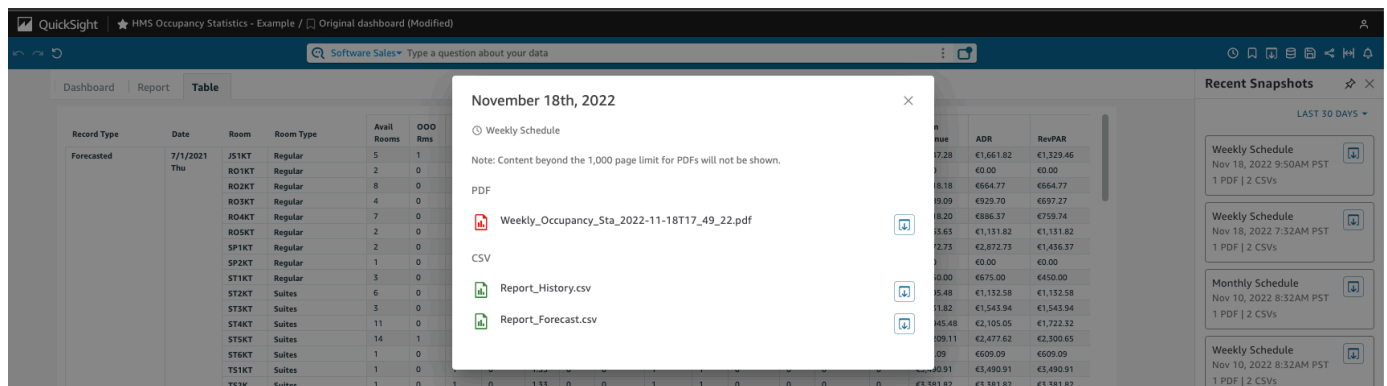


To export a paginated report as a CSV

1. From the report that you want to export, choose the **Scheduling** icon at the top right, and then choose **Recent snapshots**.

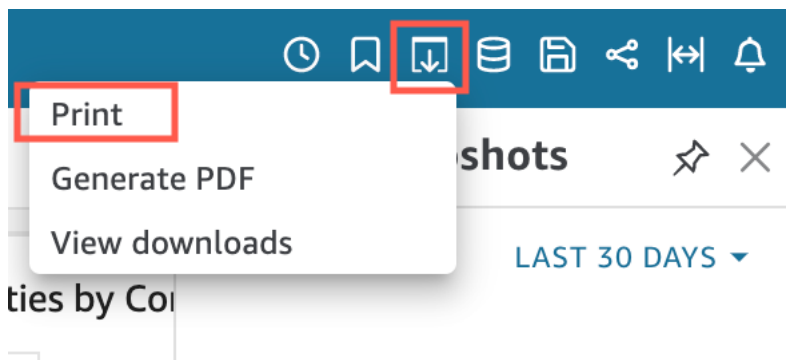


- In the **Recent snapshots** menu that appears on the right, snapshots are sorted from most recently generated to the oldest. Snapshots are stored for up to 1 year. Find the report that you want to download and choose the download icon to the right of the report.
- In the report pop-up that appears, 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

- From the report that you want to print, choose the **Export** icon at the top right, and then choose **Print**.



- 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.

The screenshot shows an Amazon QuickSight invoice report. The report includes contact information for Acme Corp, bill and ship to addresses, and invoice details. A table of late invoice details is also present. A print dialog is open on the right side of the screen, showing options for destination, pages, and print settings.

Invoice

Acme Corp (321) 123 - 4567
123 Rainier Way
Seattle, WA 98009
Email: ar@acmecorp.com
Website: www.acmecorp.com

Bill To
1080-NDGAE
111 Office Ct.
Bellevue, WA 98001

Ship To
1080-NDGAE
111 Office Ct.
Bellevue, WA 98001

Invoice Details
Invoice #: 12345
Invoice Date: 11/30/2022
Terms: Net 30
Due: 12/30/2022

Accounts Receivable: **\$2,646.81**
Overdue Receivable: **\$1,609.40**

Late Invoice Details

InvoiceNumber	InvoiceDate	DueDate	DaysLate	InvoiceAmount
9,632,048,192	Jul 9, 2012	Aug 8, 2012	15	\$128.28
857,712,918	Feb 22, 2013	Mar 24, 2013	11	\$93.39
9,390,786,866	Feb 24, 2013	Mar 26, 2013	11	\$74.62
915,652,542	Jan 5, 2012	Feb 4, 2012	9	\$78.29
246,081,324	Sep 6, 2012	Oct 6, 2012	7	\$92.53
106,360,977	Feb 19, 2012	Mar 20, 2012	6	\$93.48
2,121,660,618	Jan 25, 2013	Feb 24, 2013	6	\$79.79
3,954,057,080	May 10, 2013	Jun 9, 2013	6	\$74.50
4,701,158,835	Nov 18, 2013	Dec 18, 2013	5	\$83.59
8,193,630,211	Mar 10, 2012	Apr 9, 2012	5	\$85.82
2,912,484,665	Aug 5, 2013	Sep 4, 2013	4	\$92.16
3,289,097,967	Aug 30, 2013	Sep 29, 2013	4	\$82.60
2,592,238,538	May 27, 2012	Jun 26, 2012	3	\$92.25
5,144,461,624	Nov 13, 2013	Dec 13, 2013	3	\$101.93
5,536,610,902	Jun 29, 2013	Jul 29, 2013	3	\$89.97
641,300,165	Sep 12, 2013	Oct 12, 2013	2	\$94.57

Page 1 / 3

Print 3 pages

Destination: Save as PDF

Pages: All

Pages per sheet: 1

Print using system dialog... (Q-⌘P)

Open PDF in Preview

Cancel Save

Subscribing to dashboard emails and alerts

Using Amazon QuickSight, 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.

Bookmarking views of a dashboard

When you load a dashboard as an Amazon QuickSight 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 multiple bookmarks for a dashboard and share them by a URL link with other subscribers of that dashboard.

Dashboard bookmarks are available on the Amazon QuickSight console. Bookmarks for embedded dashboards are currently not supported.

Dashboard bookmarks for paginated reports are currently not supported. For more information on paginated reports, see [Working with paginated reports in Amazon QuickSight](#).


Use the following procedures to learn how to use bookmarks.

Creating bookmarks

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).



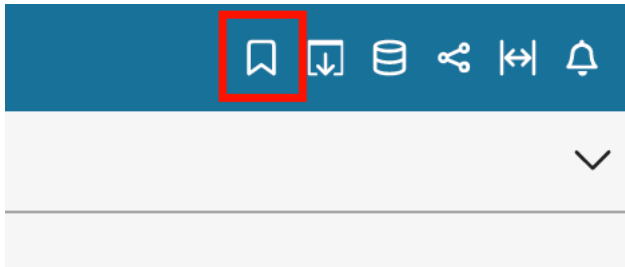
You can return to the original dashboard view that the author published at any time by selecting **Original dashboard** in the **Bookmarks** pane at right.

Updating bookmarks

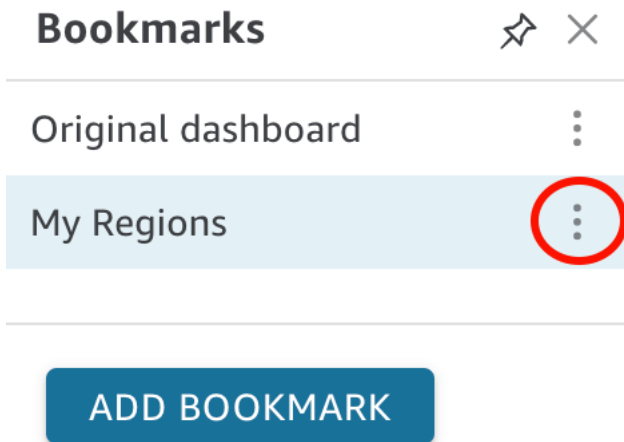
At any time, you can change a bookmark dashboard view and update the bookmark to always reflect those changes.

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

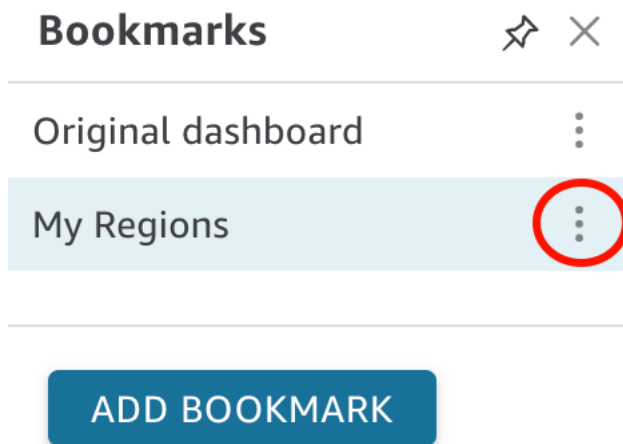
Use the following procedure to rename a bookmark.

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**.



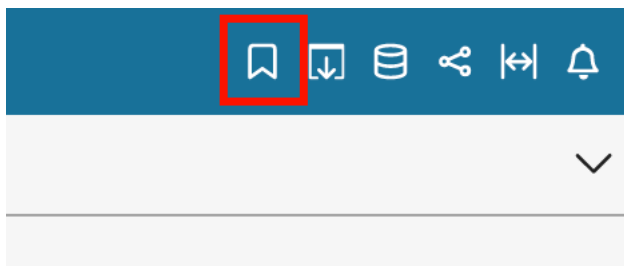
3. In the **Rename bookmark** pane, enter a name for the bookmark, and then choose **Save**.

Making a bookmark the default view


By default, when you update a dashboard, QuickSight 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**.

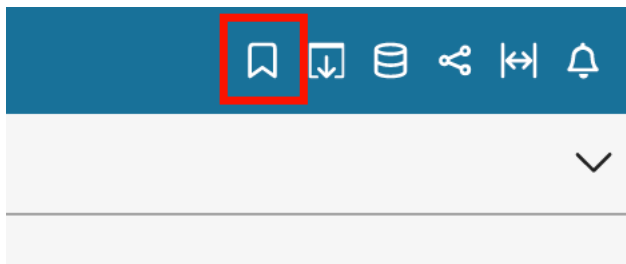


Sharing bookmarks

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 QuickSight 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 QuickSight dashboard](#).




Deleting bookmarks

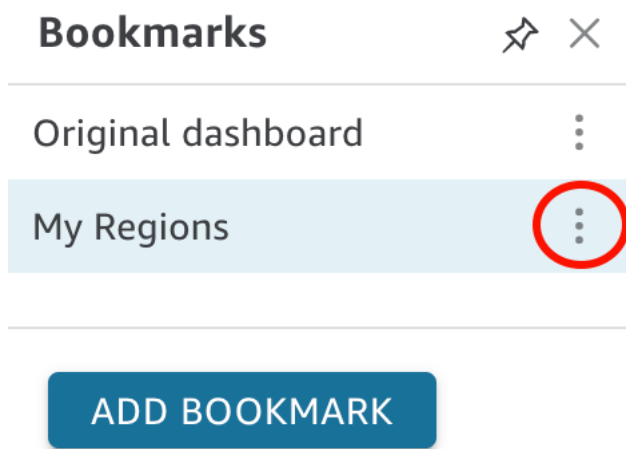
Use the following procedure to delete a bookmark.

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**.

Setting up for Amazon QuickSight

In this section, you sign up for an Amazon Web Services account, create an IAM user, and sign up for Amazon QuickSight.

Topics

- [Sign up for Amazon](#)
- [Create an IAM user](#)
- [Integrating with IAM Identity Center](#)
- [Signing up for an Amazon QuickSight subscription](#)

Sign up for Amazon

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.

Create an IAM user

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](#)

Integrating with IAM Identity Center

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 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.

Signing up for an Amazon QuickSight subscription

When you first sign up for Amazon QuickSight, you get a free trial subscription for four users for 30 days. During the process of signing up, you choose which edition of QuickSight to use and 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 [Sign up for Amazon](#). The person who signs up for QuickSight needs to have the correct Amazon Identity and Access Management (IAM) permissions. For more information, see [IAM policy examples for Amazon QuickSight](#).

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?](#)

To subscribe to QuickSight

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

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

2. Choose **Sign up for QuickSight**.
3. Choose **Standard** or **Enterprise**.
 - a. If you choose **Standard**, choose the method that you want to connect with. Choose one of the following:
 - **Use IAM federated identities and QuickSight-managed users.**
 - **Use IAM federated identities only.**
 - b. If you choose **Enterprise**, you can also choose to add QuickSight Q to your subscription. To add QuickSight Q to your subscription, choose **Enterprise + Q**. For more information about QuickSight Q, see [???](#).

If you don't want to add QuickSight Q to your account, choose **Enterprise**.

Select **Continue**. Then, choose the method that you want to connect with.

Choose one of the following:

- (Recommended) **Use IAM Identity Center enabled application**. This option is only available for Enterprise Edition accounts.
- **Use Active Directory**
- **Use IAM federated identities and QuickSight-managed users**
- **Use IAM federated identities only**

To sign up for a QuickSight Enterprise Edition account with an IAM Identity center enabled application, you need the correct permissions. For more information on the permissions

needed to use this method, see [IAM identity-based policies for Amazon QuickSight: All access for Enterprise edition with IAM Identity Center](#).

To sign up for QuickSight with federated users, you need the correct IAM permissions, defined as follows:

- To use role-based federation (that is, single sign-on, or IAM Identity Center) with QuickSight Standard Edition or QuickSight Enterprise Edition, see [IAM identity-based policies for Amazon QuickSight: All access for Standard edition](#).
- To use Microsoft Active Directory with QuickSight Enterprise Edition, see [IAM identity-based policies for Amazon QuickSight: all access for Enterprise edition with Active Directory](#). QuickSight Standard Edition doesn't work with Active Directory.

After you finish creating an Enterprise Edition account in Amazon QuickSight, you can add a subscription to Paginated Reports from the **Manage subscriptions** page of the **Manage QuickSight** menu. For more information on paginated reports, see [Working with paginated reports in Amazon QuickSight](#).

4. For both Standard and Enterprise editions, make choices for the following items:
 - Enter a unique account name for QuickSight. Your account name can only contain characters (A–Z and a–z), digits (0–9), and hyphens (-). Note that if your account begins with the characters D- or d-, an error occurs. If you use Microsoft AD, and it has a default alias, this alias is used for the account name.
 - Enter a notification email address for the QuickSight account owner or group. This email address receives service and usage notifications.
 - (Optional) Choose the Amazon Web Services Region that you want to use for your initial data storage capacity, called SPICE.
 - (Optional) Choose whether to allow autodiscovery of your Amazon resources. You can change these options later in **Manage Account**. For more information, see [Allowing autodiscovery of Amazon resources](#).
 - (Optional) For **IAM Role**, choose **Use an existing role**, and then from the list choose a role that you want to use. Or enter the IAM Amazon Resource Name (ARN) in the following format: `arn:aws-cn:iam::account-id:role/path/role-name`.

Note

Make sure to have your administrator give you permissions to pass any existing IAM roles in QuickSight. If you don't have permissions, or if you don't know if you have permissions, choose **QuickSight-managed role**. This is the default role. You can always switch to using a different role later if you have the correct permissions. For more information, see [Using an existing IAM role in QuickSight](#).

5. Review the choices that you made, then choose **Finish**.
6. (Optional) To open QuickSight, choose **Go to QuickSight**.

If you're using Enterprise edition, you can manage user groups by choosing **Manage access to QuickSight**. Otherwise, close the browser and notify your users how to connect.

7. (Optional) If you're using IAM Identity Center or federation, choose the users and groups that are going to use QuickSight.

Getting started with Amazon QuickSight data analysis

Use the topics in this section to create your first analysis. You can use sample data to create either a simple or a more advanced analysis. Or you can connect to your own data to create an analysis.

Before you create your first analysis, make sure to complete the steps in [Setting up for Amazon QuickSight](#).

Topics

- [Signing in to Amazon QuickSight](#)
- [Quick start: Create an Amazon QuickSight analysis with a single visual using sample data](#)
- [Tutorial: Create an Amazon QuickSight dashboard using sample data](#)
- [Using the Amazon QuickSight console](#)

Signing in to Amazon QuickSight

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

If you're a QuickSight 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 QuickSight 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*.

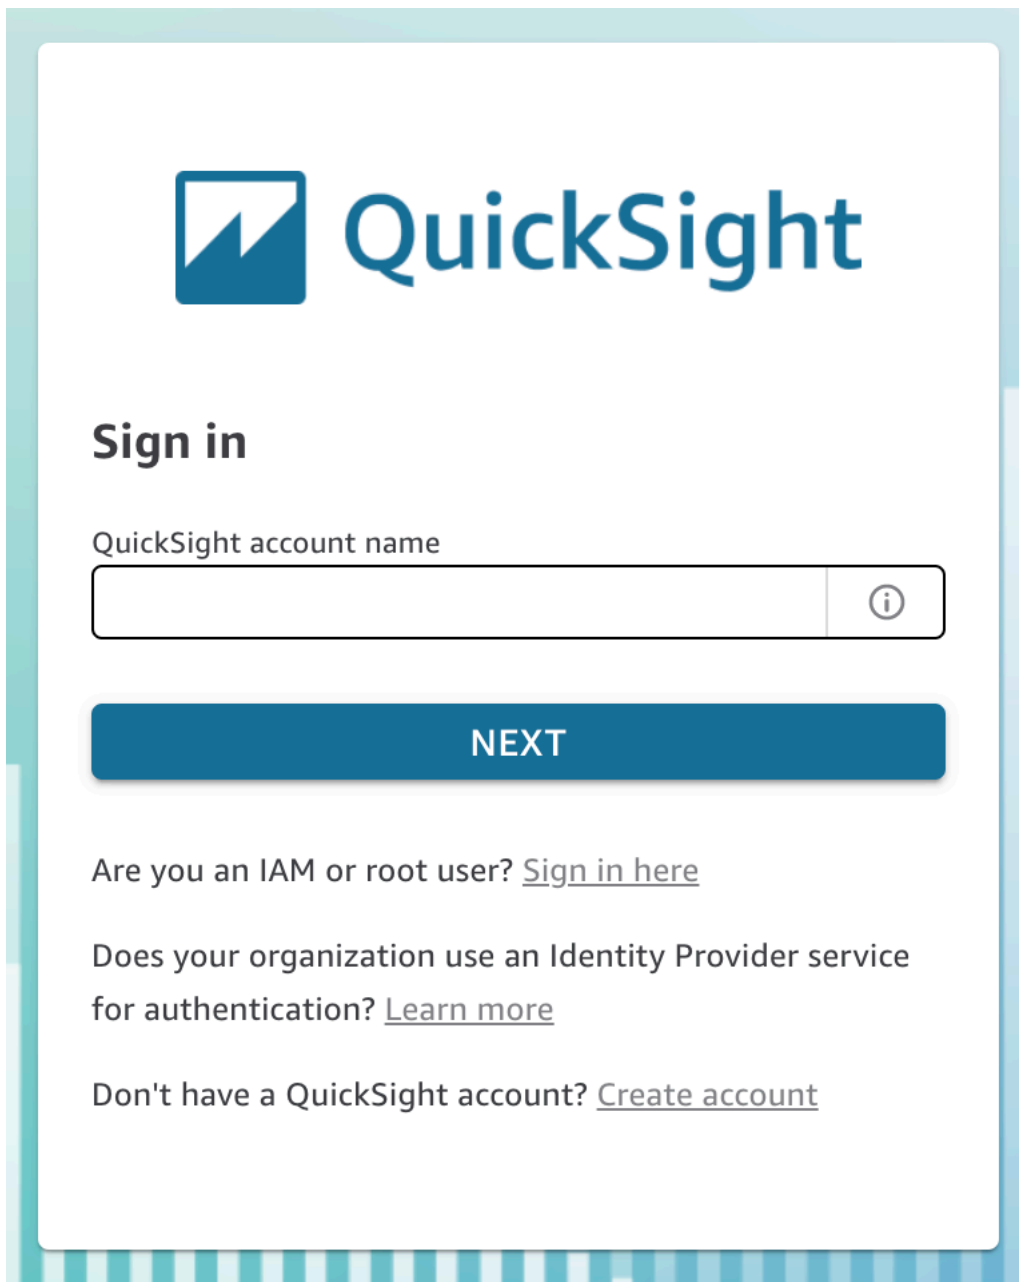
How to sign in to Amazon QuickSight


Use the following procedure to sign in to QuickSight.

To sign in to QuickSight

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

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



 **QuickSight**

Sign in

QuickSight account name


NEXT

Are you an IAM or root user? [Sign in here](#)

Does your organization use an Identity Provider service for authentication? [Learn more](#)

Don't have a QuickSight account? [Create account](#)

You can also find your QuickSight account name at the top of the menu at upper-right on the QuickSight console. In some cases, you might not have access to your QuickSight 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.



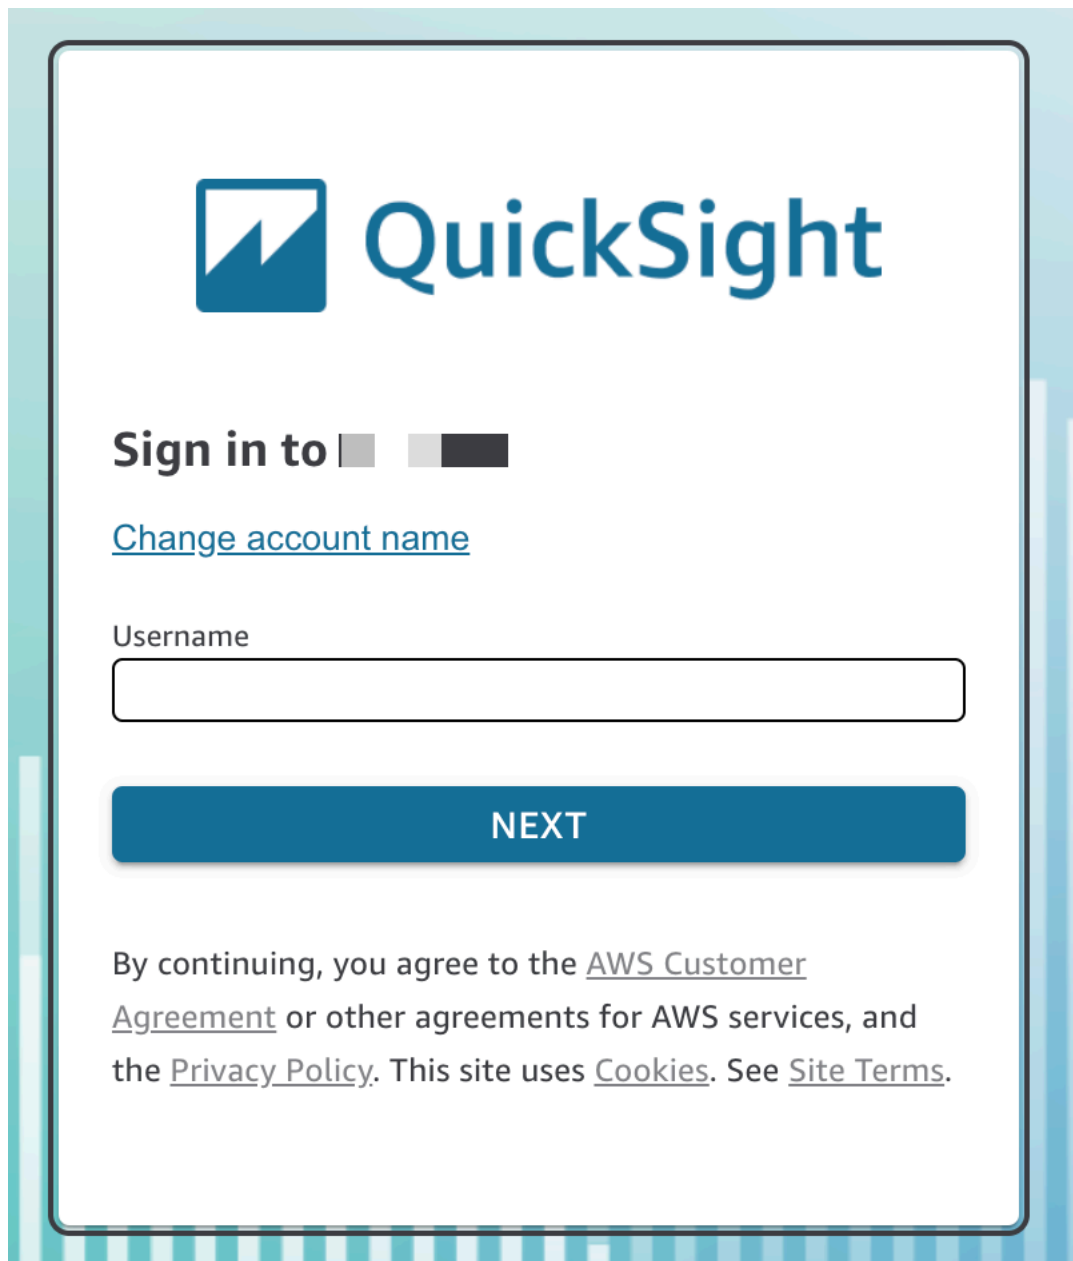
3. For **Username**, enter your QuickSight user name. User names that contain a semicolon (;) aren't supported. Choose one of the following:

- 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 QuickSight account. If you received an invitation email from another Amazon QuickSight 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 image shows a sign-in page for Amazon QuickSight. At the top left is the QuickSight logo, which consists of a blue square with a white diagonal line. To the right of the logo is the text "QuickSight" in a large, blue, sans-serif font. Below the logo and text, there is a "Sign in to" label followed by three small colored squares (gray, light gray, and dark gray). Underneath this is a blue hyperlink that says "Change account name". Below the link is a text input field labeled "Username". At the bottom of the input field is a large, blue button with the word "NEXT" in white, uppercase letters. Below the button is a paragraph of text: "By continuing, you agree to the [AWS Customer Agreement](#) or other agreements for AWS services, and the [Privacy Policy](#). This site uses [Cookies](#). See [Site Terms](#)."

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

Finishing QuickSight sign-in as a QuickSight or Active Directory user

If you're signing in directly through QuickSight 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 QuickSight if you sign in directly through QuickSight 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 QuickSight 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 QuickSight 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.

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

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

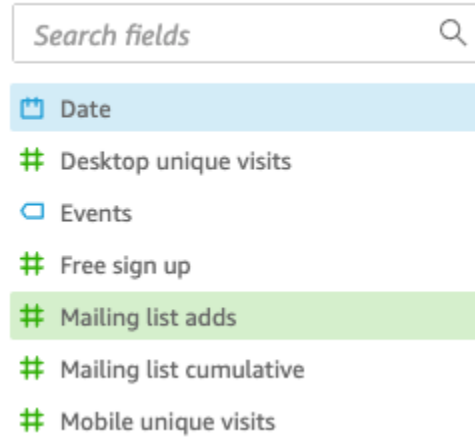
1. On the Amazon QuickSight start page, choose **New analysis**. If you don't have the sample data, you can download it from [web-and-social-analytics.csv.zip](#). Unzip the file so you can use the .csv file.

To upload the sample data, do the following:

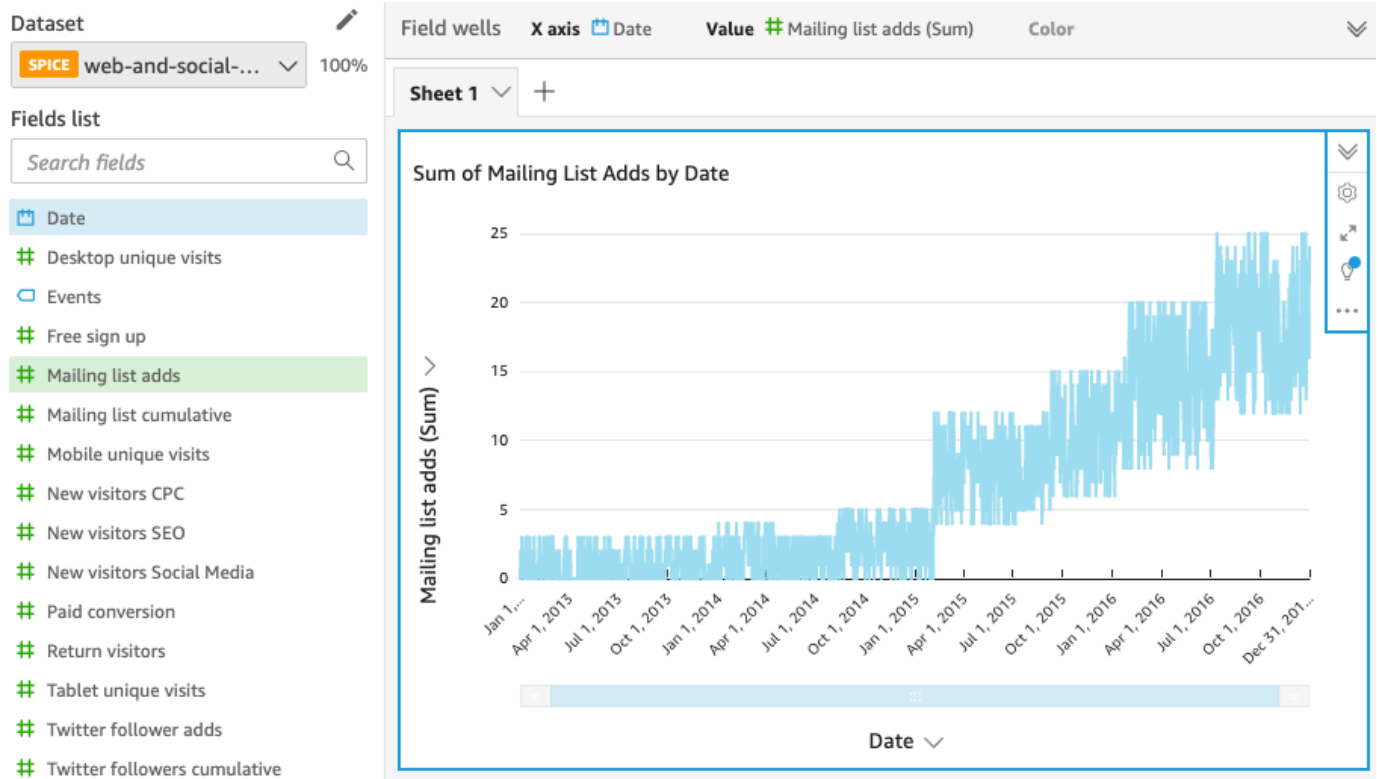
- a. Choose **New dataset** from the **New analysis** screen. (Or choose **Datasets** at left, and then choose **New dataset**.)
- b. Choose **Upload a file**.
- c. Choose the sample file, `web-and-social-analytics.csv`, from your drive. If you don't see it, check that you unzipped the `web-and-social-analytics.csv.zip` 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 **Fields list** pane, choose **Date**, and then choose **Mailing list adds**.

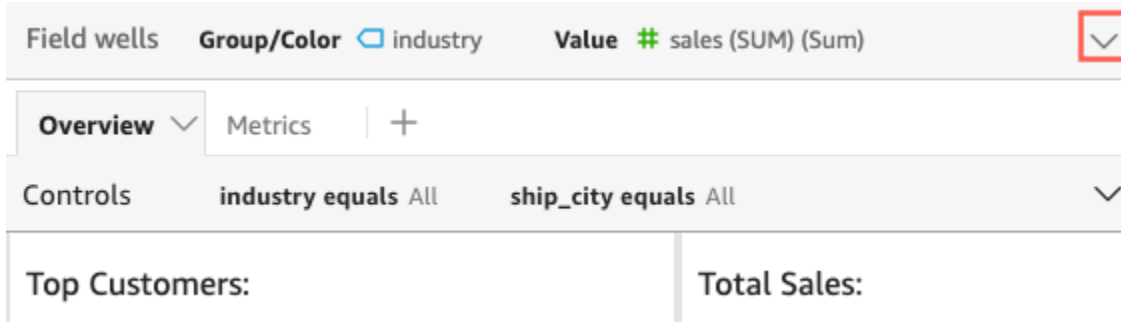
Fields list



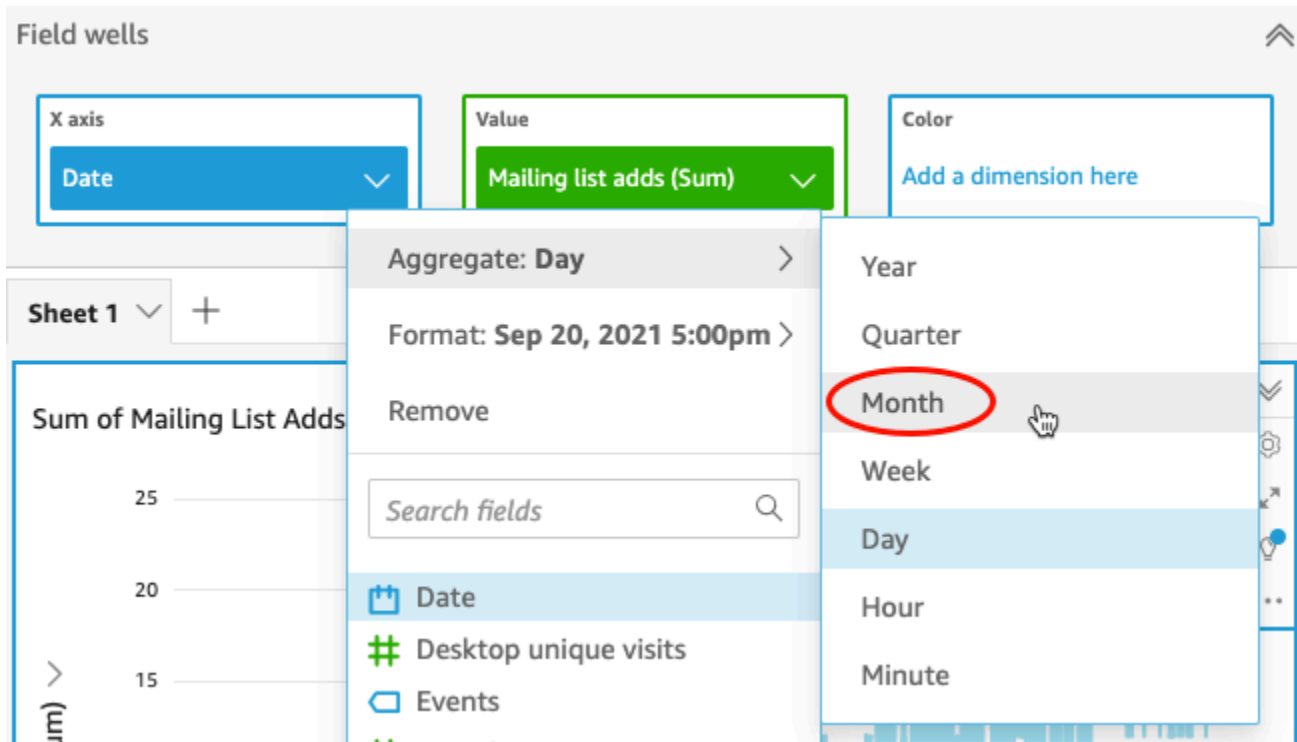
Amazon QuickSight uses AutoGraph to create the visual, selecting the visual type that it determines is most compatible with those fields. In this case, it selects a line chart that shows mailing list adds by day, which is the date granularity default.



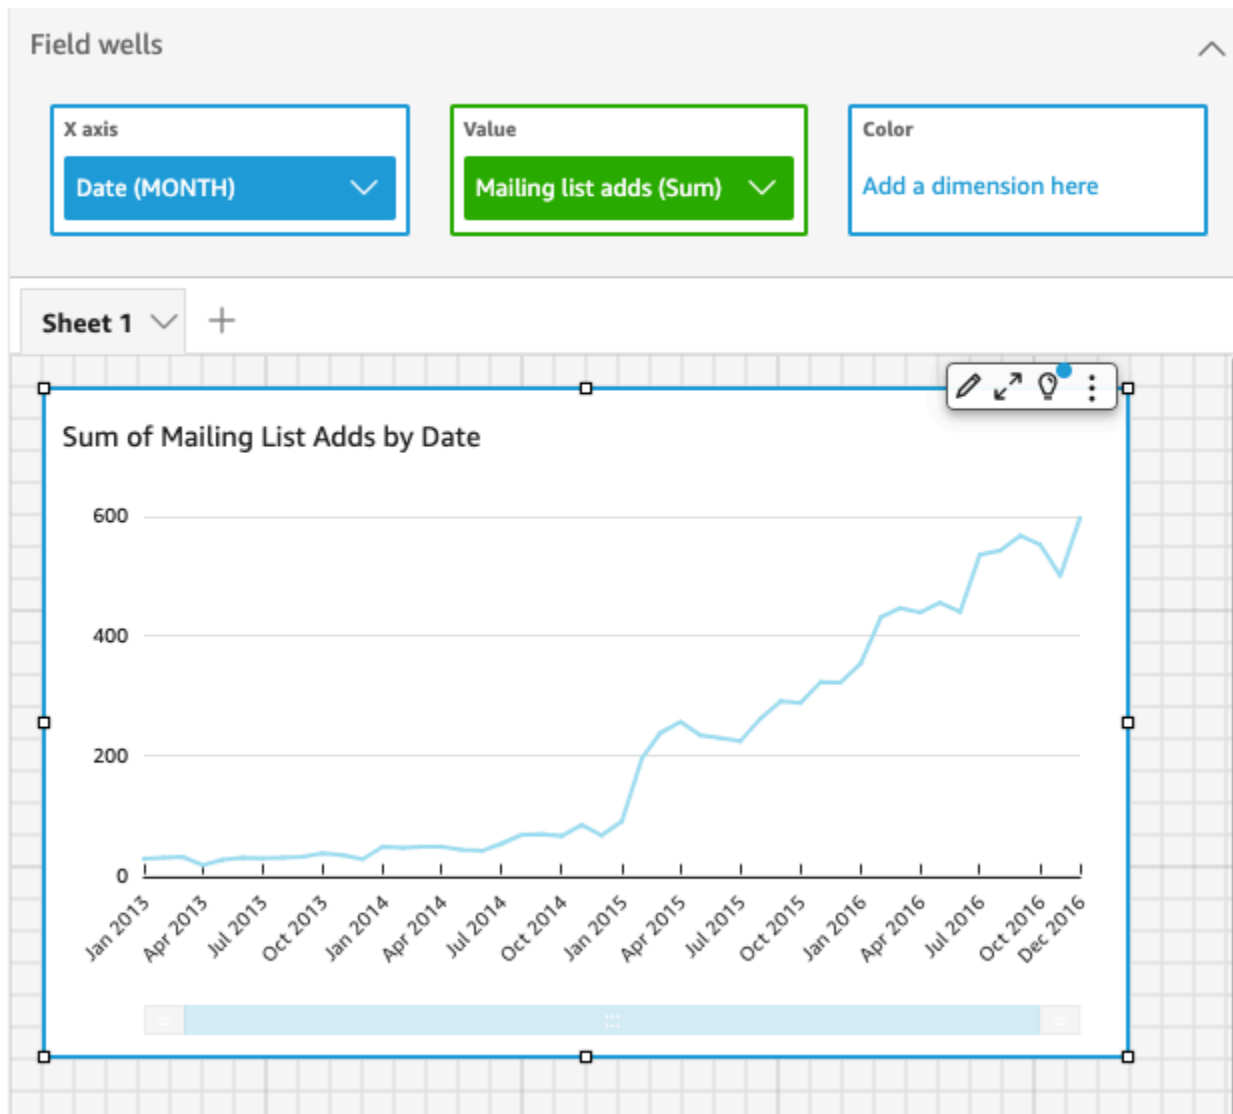
- Expand the **Field wells** pane by choosing the expand icon.



- Choose the **X axis** field well, 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 QuickSight 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 QuickSight dataset](#)
- [Tutorial: Create an Amazon QuickSight analysis](#)
- [Tutorial: Modify Amazon QuickSight visuals](#)
- [Tutorial: Create an Amazon QuickSight dashboard](#)

Tutorial: Create a prepared Amazon QuickSight 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 QuickSight, you can download it: web-and-social-analytics.csv.zip.

To prepare the Marketing dataset and create an analysis

1. On the Amazon QuickSight start page, choose **Datasets** at left.

Find analyses & more

- ★ Favorites
- 🕒 Recent
- 📁 My folders
- 📁 Shared folders

- Dashboards
- Analyses**
- Datasets**
- Topics

Analyses

Analysis

Business Review analysis
Updated 11 minutes ago

Analysis

Web and Social Media Anal...
Updated 2 months ago

Analysis

photos.csv analysis
Updated 4 months ago

- On the **Datasets** page, choose **New dataset**.

Find analyses & more

Datasets

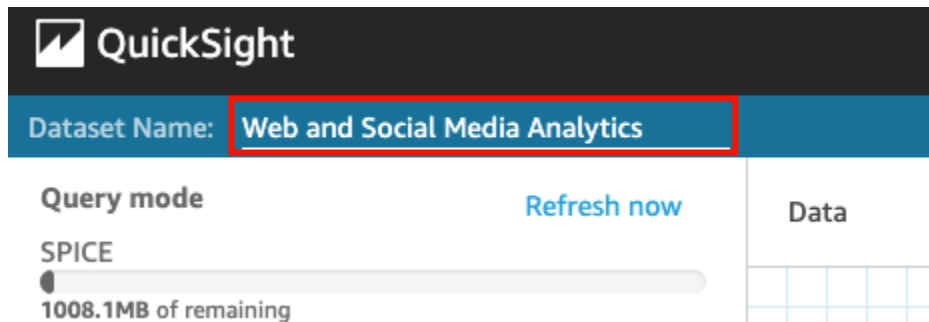
[New dataset](#)

Name	Owner	Last Modified
[Redacted]	Me	an hour ago
[Redacted]	Me	16 days ago
[Redacted]	Me	16 days ago

- In the **FROM EXISTING DATA SOURCES** section of the **Create a Data Set** page, choose the **Web and Social Media Analytics** Amazon S3 data source and then choose **Edit dataset**.

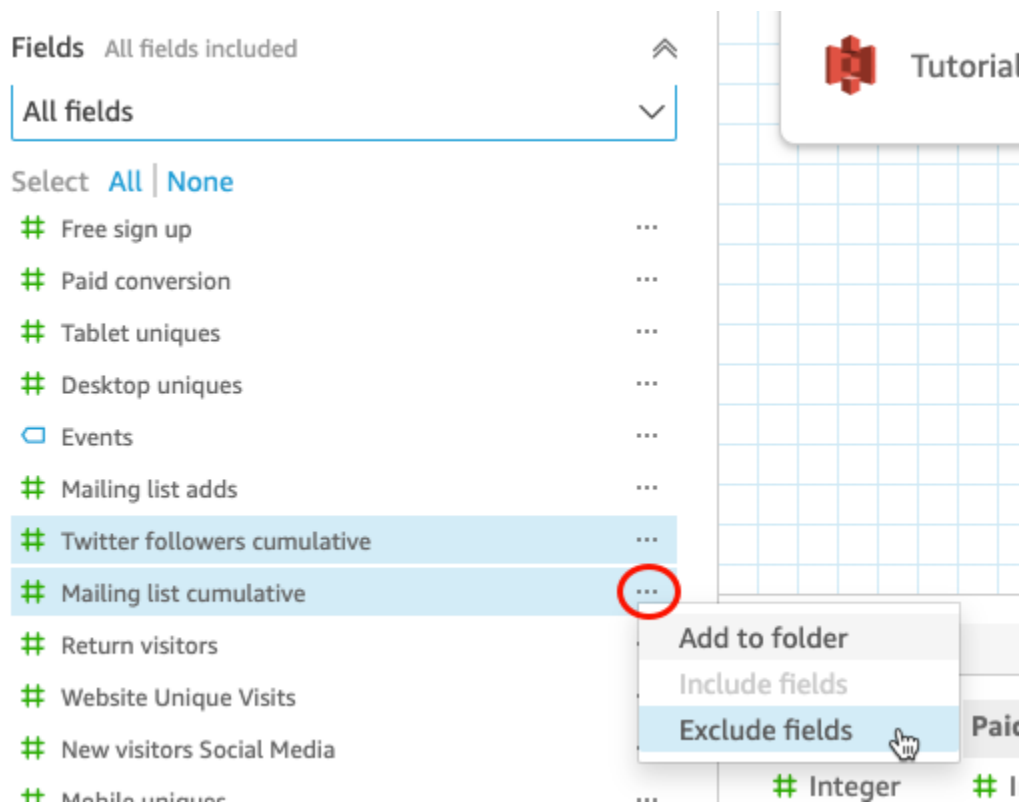
Amazon QuickSight opens the data preparation page.

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



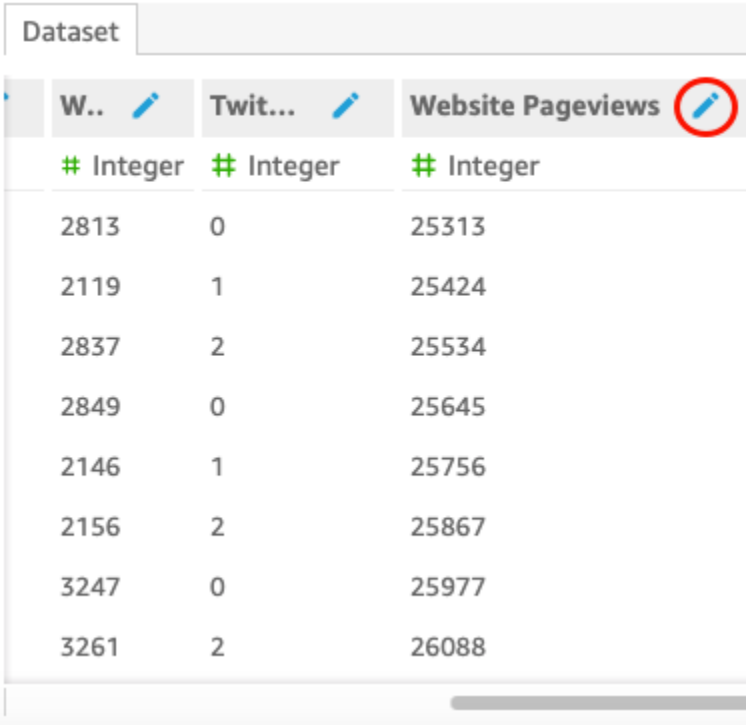
5. Exclude some fields from the dataset.

In the **Fields** pane, choose the field 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.

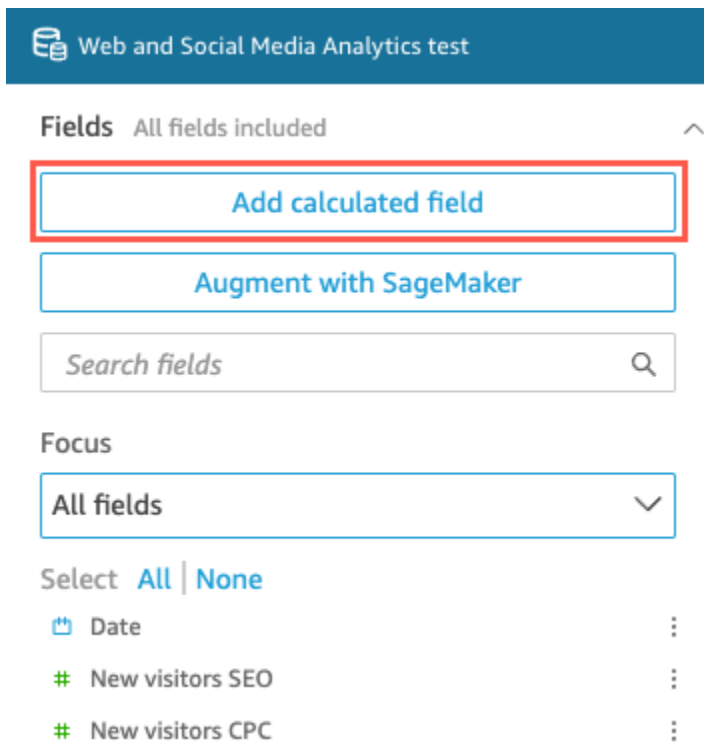
In the **Dataset** preview pane, scroll to the **Website Pageviews** field and choose the edit icon.



W..	Twit...	Website Pageviews
# Integer	# Integer	# Integer
2813	0	25313
2119	1	25424
2837	2	25534
2849	0	25645
2146	1	25756
2156	2	25867
3247	0	25977
3261	2	26088

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**.



Web and Social Media Analytics test

Fields All fields included ^

Add calculated field

Augment with SageMaker

Search fields

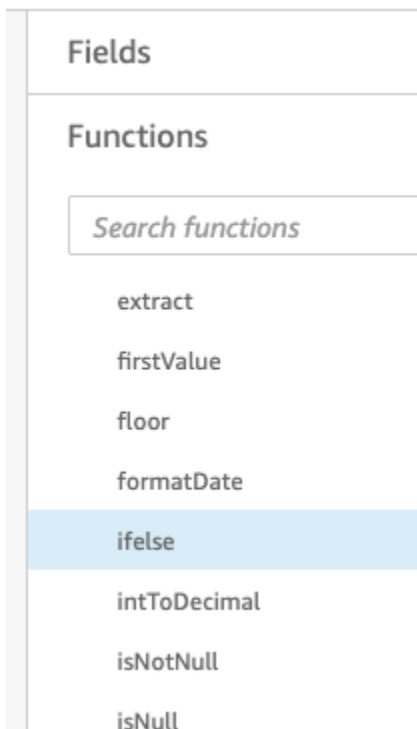
Focus

All fields

Select All | None

- Date
- New visitors SEO
- New visitors CPC

- In the **Add calculated field** page that opens, for **Add name**, enter **populated_event**.
- 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.




Fields

Functions

Search functions

- extract
- firstValue
- floor
- formatDate
- ifelse
- intToDecimal
- isNotNull
- isNull

- 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})`.


Add calculated field

populated_event 

```
1 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 QuickSight analysis](#).

Tutorial: Create an Amazon QuickSight 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 QuickSight dataset](#).

Create your analysis

Use the following procedure to create your analysis.

To create your analysis

1. On the Amazon QuickSight start page, choose **New analysis**.




2. On the **Datasets** page, choose the **Business Review** sample dataset, and then choose **Create Analysis**.

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 QuickSight creates a line chart using this data.




Create a scatter plot visual

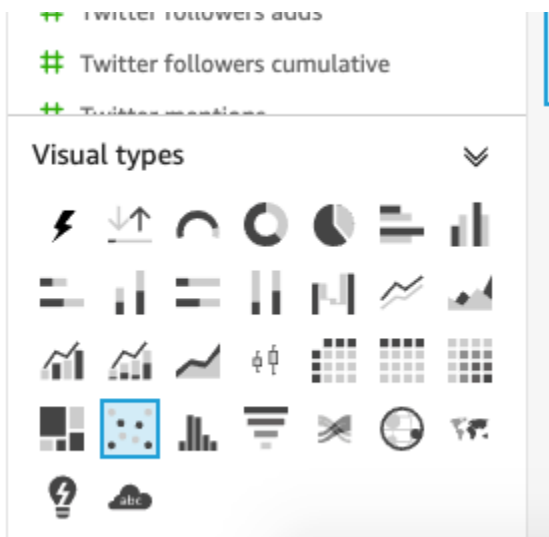
Create a visual by choosing a visual type and dragging fields to the field wells.

To create a scatter plot visual

1. On the analysis page, choose **Add** 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 QuickSight visuals](#).

Tutorial: Modify Amazon QuickSight visuals

Use the following procedures to modify the visuals that you created using the procedures in [Tutorial: Create an Amazon QuickSight 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.

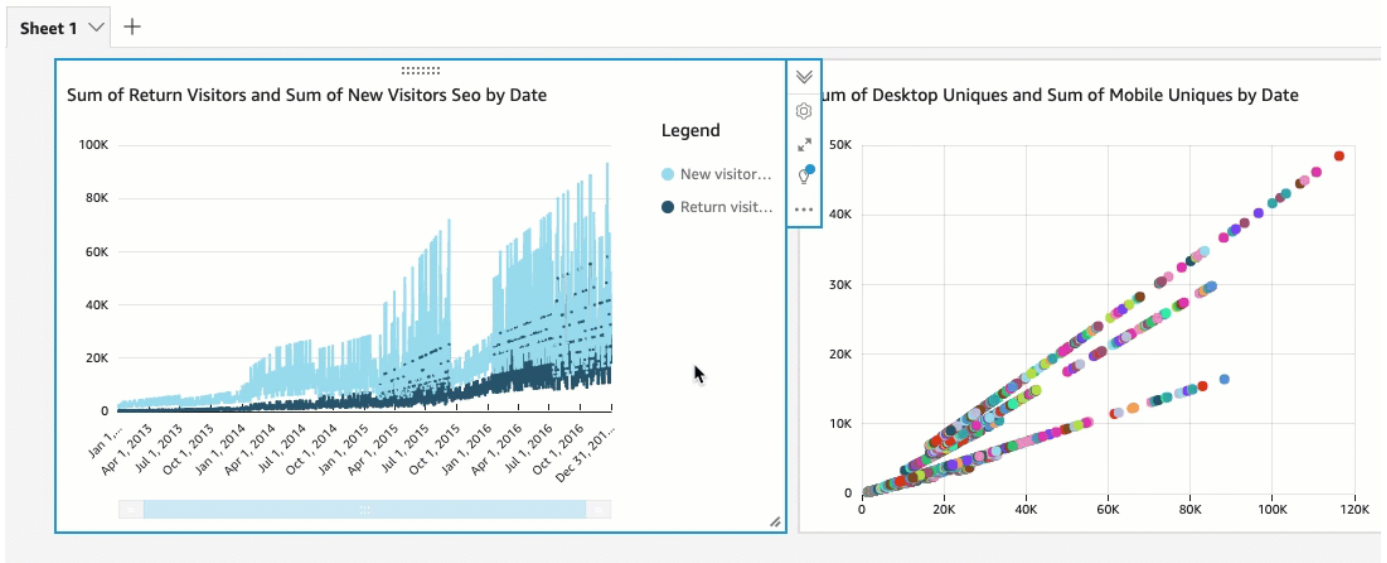
Select the **New visitors SEO** field in the **Fields list** pane. This measure is added to the **Value** field well, and the line chart updates with a line to represent it. The visual title also updates.



- 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**.


The screenshot shows the "Field wells" section of the Amazon QuickSight interface. The "X axis" well contains the "Date" field, which has a dropdown arrow highlighted with a red circle. The "Value" well contains "Return visitors (Sum)". A dropdown menu is open for the "Date" field, showing options for aggregation: "Aggregate: Day", "Format: Sep 20, 2021 5:00pm", and "Remove". A secondary dropdown menu is open, showing time intervals: "Year", "Quarter", "Month" (circled in red), "Week", "Day", "Hour", and "Minute". Below the field wells, a chart titled "Sum of Return Visitors and Su..." is visible, showing a bar chart with a y-axis from 0 to 100K and an x-axis with dates from 2013 to 2016.

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

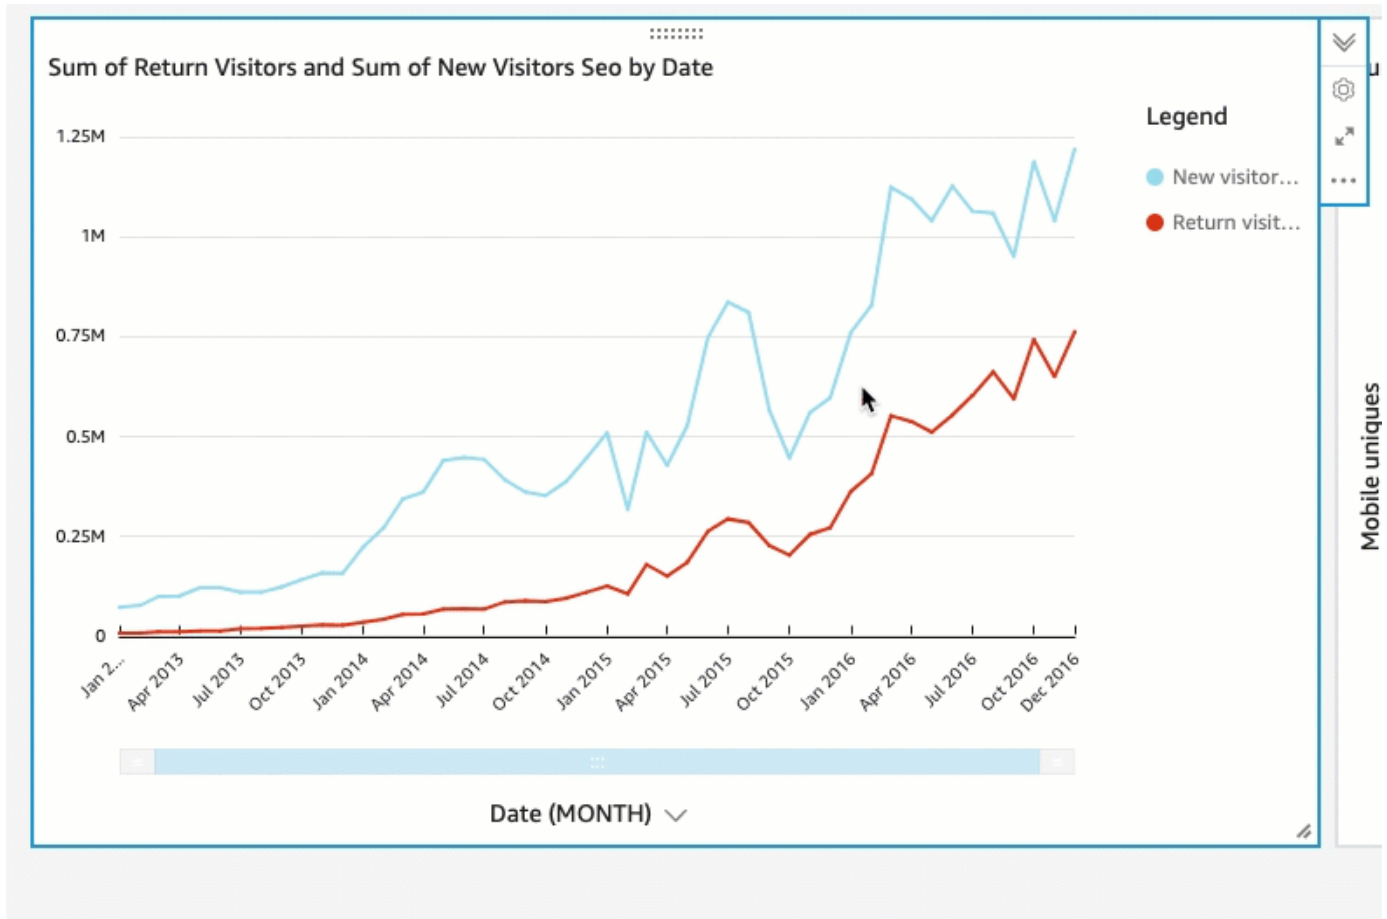
Modify both visuals by changing visual size and location, and by adding a filter and applying it to both of them.

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.



- Repeat this procedure on the scatter plot visual.
- 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

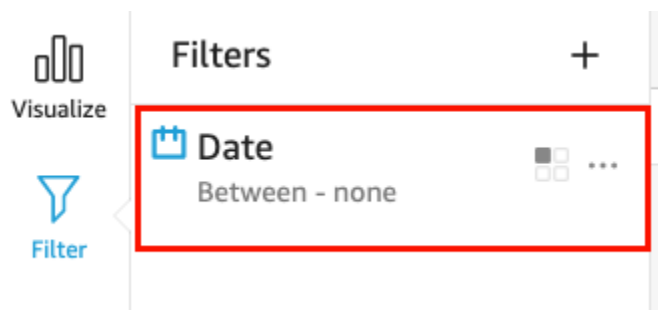
Modify both visuals by adding a filter and applying it to both of them.

To add a filter to both visuals

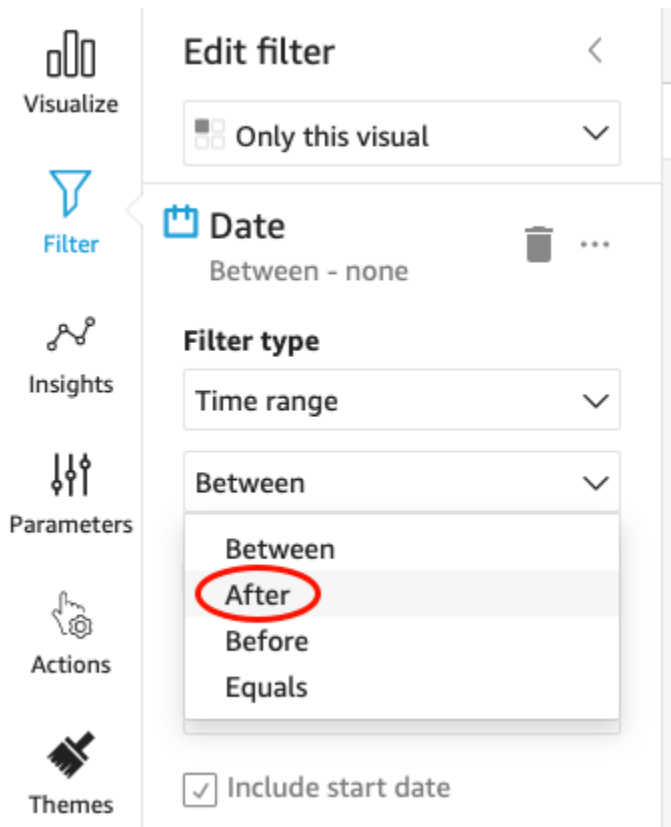
1. In the analysis, choose the scatter plot visual.
2. Choose **Filter** at left.
3. On the **Filters** pane, choose the plus icon, and then choose the **Date** field to filter on.



4. Choose the new filter to expand it.



5. In the **Edit filter** pane, for **Filter type**, 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.

Filter type

Time range ▼

After ▼

Use parameters

Date


2014-01-01 00:00

< **January 2014** ▶

Sun	Mon	Tue	Wed	Thu	Fri	Sat
29	30	31	1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18


- 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.



- 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 QuickSight dashboard](#).

Tutorial: Create an Amazon QuickSight dashboard

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

To create a dashboard from your analysis

1. In your analysis, choose **Share** in the application bar at upper right, and then choose **Publish dashboard**.




2. In the **Publish dashboard** page that opens, choose **Publish new dashboard as**, and enter the name **Marketing Dashboard**.

Publish a dashboard ✕

Publish new dashboard as

Marketing dashboard

Replace an existing dashboard

Advanced publish options 

Publish 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.

Using the Amazon QuickSight console

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

Topics

- [Using the Amazon QuickSight menu and landing page](#)
- [Using the Amazon QuickSight start page](#)
- [Choosing a language in Amazon QuickSight](#)
- [Using the Amazon QuickSight mobile app](#)


Using the Amazon QuickSight menu and landing page

After you sign in to Amazon QuickSight, you see the Amazon QuickSight landing page. This page provides tabs for your analyses, your dashboards, and our tutorial videos. It also provides a menu bar at the top, with options for the following:


- Searching Amazon QuickSight
- Choosing the Amazon Region that you want to work in
- Accessing your user profile (community, language selection, and help)
- Creating a new analysis
- Managing data

Note

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



To view videos about Amazon QuickSight, choose your user name at the upper-right of any page, and then choose **Tutorial videos**. Choose a video to play it.



To access the user profile menu, choose your user icon at the upper right of any page in Amazon QuickSight. Use this menu to manage Amazon QuickSight features, visit the community, send product feedback, choose a language, get help from the documentation, or sign out of Amazon QuickSight.

Username

Account name

Manage QuickSight

Community

Send feedback

🌐 English >

📍 N. Virginia >

Tutorial videos


Help

Sign out

The following options are available from the user profile menu:

- **Manage QuickSight** – If you have appropriate permissions, you can access administrative functions such as managing users, subscriptions, [SPICE](#) capacity, and account settings.

- **Community** – Choose this option to visit the Amazon QuickSight online community.
- **Send feedback** – This is your direct connection to the product team. Use this simple form to report problems, request features, or tell us how you are using Amazon QuickSight.
- **What's new** – Find out what new features are available in Amazon QuickSight.
- **Language setting** – Choose the language you want to use in the Amazon QuickSight 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 QuickSight administrator. Changing the Amazon Region changes where your work is stored.

- **Tutorial videos** – This will open the Tutorial videos page where you can watch videos about Amazon QuickSight.
- **Help** – This will open the official Amazon documentation, which you can view online, in Kindle, or as a PDF.
- **Sign out** – Choose this option to sign out of Amazon QuickSight and your Amazon session.

Using the Amazon QuickSight start page

To see available dashboards, choose **Dashboards** at left. Choose any dashboard on the page to open it.

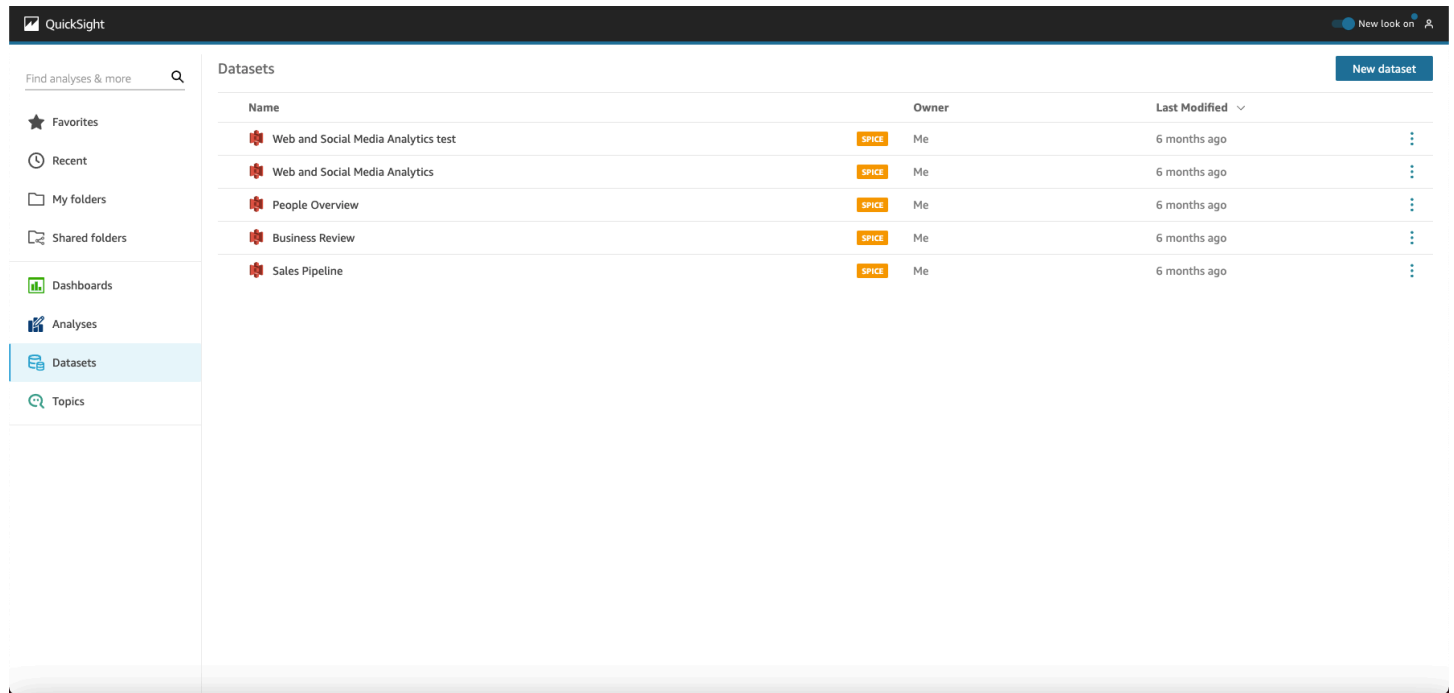
To see available analyses, choose **Analyses** at left. This is the default page when Amazon QuickSight opens. Choose any analysis to open it.



To see your list of favorite dashboards and analyses, choose **Favorites**. You can add items to your favorites by selecting the star near the title of the dashboard or analysis, so that the star is filled in. Clear the star to remove the item from your favorites.

To create a new analysis, choose **New analysis**, near the top left. This takes you to **Datasets**. Choose one to start analyzing it.

To see current datasets, or to create a new dataset, choose **Datasets**. This takes you to the **Datasets** page, which displays the datasets that you have access to. (If they don't all fit on one page, you can navigate between pages.) From here, you can choose a dataset to analyze.



The screenshot displays the Amazon QuickSight interface. On the left, there is a navigation sidebar with options: Favorites, Recent, My folders, Shared folders, Dashboards, Analyses, Datasets (highlighted), and Topics. The main area is titled 'Datasets' and features a 'New dataset' button in the top right. Below the button is a table listing datasets:

Name	Owner	Last Modified
Web and Social Media Analytics test	Me	6 months ago
Web and Social Media Analytics	Me	6 months ago
People Overview	Me	6 months ago
Business Review	Me	6 months ago
Sales Pipeline	Me	6 months ago

To create a new dataset, choose **New dataset**. From here, you can upload a file, or you can create a new dataset based on a data source (a connection to external data). Icons for new data sources are at the top of the screen under **From new data sources**. Icons for existing data sources are displayed below them, under **From existing data sources**.

Data Sets SPICE capacity for this region: 15.9MB of 1GB

Create a Data Set
FROM NEW DATA SOURCES

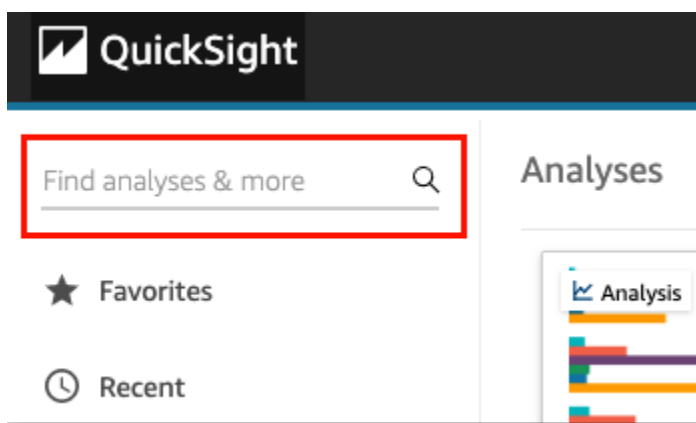
Upload a file (.csv, .tsv, .clf, .elf, .xlsx, .json)	Salesforce Connect to Salesforce	S3 Analytics	S3	Athena
RDS	Redshift Auto-discovered	Redshift Manual connect	MySQL	PostgreSQL
ORACLE	SQL Server	Aurora	MariaDB	Presto
Spark	Teradata Provided by Teradata	Snowflake	AWS IoT Analytics	Amazon Elasticsearch Se...
Timestream	GitHub	Twitter	Jira	ServiceNow
Adobe Analytics				

FROM EXISTING DATA SOURCES

covid-19-ds1 Updated a month ago	covid-19-ds1 Updated a month ago	covid-19-ds1 Updated a month ago	Tutorial Test Updated 2 months ago	Business Review Updated 3 months ago
Sales Pipeline Updated 3 months ago	People Overview Updated 3 months ago			

Searching Amazon QuickSight

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 data set, create an analysis from a data set, or access an analysis or dashboard. Choose an item from the search results to open it.

Choosing a language in Amazon QuickSight

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

Amazon QuickSight supports the following languages:

Languages available in the Amazon QuickSight user interface

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


Official name	Language code	Localized name
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 QuickSight 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 QuickSight 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.

Using the Amazon QuickSight mobile app

The Amazon QuickSight mobile app enables you to securely get insights from your data from anywhere; favorite, browse, and interact with your dashboards; explore your data with drilldowns and filters; stay ahead of the curve via forecasting; get email alerts when unexpected changes happen in your data; and share those insights with colleagues.

For a quick tour of the app, see [Amazon QuickSight announces the all-new QuickSight mobile app](#) on the Amazon Big Data Blog.

QuickSight accounts that are integrated with IAM Identity Center are not supported by the Amazon QuickSight mobile app.

To begin using the QuickSight Mobile app, do one of the following:

- [Download the iOS version from the iOS app store](#)
- [Download the Android version from google play](#)

Connecting to data in Amazon QuickSight

People in many different roles use Amazon QuickSight 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 QuickSight. QuickSight supports direct connections and uploads from a variety of sources.

After your data is available in QuickSight 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.

After your data is available in QuickSight 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 QuickSight 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.

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.

If you perform administrative functions in QuickSight, 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 QuickSight 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 QuickSight 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 QuickSight reader security role so they can securely interact with published data dashboards on a pay-per-session basis.

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 **Datasets** page, which you reach by choosing **Manage data** on the Amazon QuickSight 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 **New data set** on the **Datasets** page.

Topics

- [Supported data sources](#)
- [Data source quotas](#)
- [Supported data types and values](#)
- [Amazon QuickSight Connection examples](#)
- [Creating datasets](#)
- [Editing datasets](#)
- [Reverting datasets back to previous published versions](#)
- [Duplicating datasets](#)
- [Changing datasets](#)
- [Sharing datasets](#)
- [Tracking dashboards and analyses that use a dataset](#)
- [Using dataset parameters in Amazon QuickSight](#)
- [Using row-level security \(RLS\) in Amazon QuickSight](#)
- [Using column-level security \(CLS\) to restrict access to a dataset](#)
- [Run queries as an IAM role in Amazon QuickSight](#)
- [Deleting datasets](#)
- [Adding a dataset to an analysis](#)
- [Working with data sources in Amazon QuickSight](#)

Supported data sources

Amazon QuickSight 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 QuickSight:

- Amazon Athena

- Amazon Aurora
- Amazon OpenSearch Service
- Amazon Redshift
- Amazon Redshift Spectrum
- Amazon S3
- Amazon S3 Analytics
- 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 5.1 or later
- Oracle 12c or later
- PostgreSQL 9.3.1 or later
- 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 QuickSight:

- 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 QuickSight](#).

Importing file data

You can use files in Amazon S3 or on your local (on-premises) network as data sources. QuickSight 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

QuickSight 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 QuickSight 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 QuickSight 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

QuickSight 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, QuickSight 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 QuickSight:

- Enterprise Edition
- Unlimited Edition
- Developer Edition

To connect to on premises data sources, you need to add your data sources and a QuickSight-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 QuickSight](#).

Data source quotas

Data sources that you use with Amazon QuickSight 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 QuickSight, [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 QuickSight 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 QuickSight 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. QuickSight 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 QuickSight 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. QuickSight doesn't support binary and character large objects (BLOBs) in dataset columns. In the QuickSight documentation, the term "text" always means "string data".

The first time you query or import the data, QuickSight 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, QuickSight uses a field length of 8 bytes plus the UTF-8 encoded character length. Amazon QuickSight 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. QuickSight supports dates and times that use [supported date formats](#).

QuickSight uses UTC time for querying, filtering, and displaying date data. When date data doesn't specify a time zone, QuickSight assumes UTC values. When date data does specify a time zone, QuickSight 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 QuickSight as **2015-11-01T15:30:00**.

For each DATE field in imported data, QuickSight uses a field length of 8 bytes. QuickSight 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. QuickSight 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 QuickSight 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. The image below shows the **Update fields** pop up.

Update fields from decimal fixed to decimal float ✕

SPICE now supports decimal-float data type. We detected fields in this dataset that can be updated to decimal-float to match the data source type. If you update, fields will be stored as decimal-float in SPICE. If you do not update these fields they will continue to be stored as decimal-fixed in SPICE. Fields can still be changed individually while editing. [Learn more](#)

Available fields 2 Fields ^

Field	Source
# precision_column	datatypes_numeric
# real_column	datatypes_numeric

i Changing data types may update the values in dependent datasets, analyses and dashboards. You can revert dataset back to previous published versions if needed. [Learn more](#)

DO NOT UPDATE FIELDS
UPDATE FIELDS

Supported data types from other data sources

The following table lists data types that are supported when using the following data sources with Amazon QuickSight.

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
Amazon Athena, Presto, Starburst, Trino	<ul style="list-style-type: none"> • bigint • decimal • double • integer • real • smallint • tinyint 	<ul style="list-style-type: none"> • char • varchar 	<ul style="list-style-type: none"> • date • timestamp 	<ul style="list-style-type: none"> • boolean
Amazon Aurora, MariaDB, and MySQL	<ul style="list-style-type: none"> • bigint • decimal • double • int • integer • mediumint • numeric • smallint • tinyint 	<ul style="list-style-type: none"> • char • enum • set • text • varchar 	<ul style="list-style-type: none"> • date • datetime • timestamp • year 	

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
Amazon OpenSearch Service	<ul style="list-style-type: none"> • byte • integer • long • float • double 	<ul style="list-style-type: none"> • string (keyword string field type in OpenSearch Service) • ip 	<ul style="list-style-type: none"> • timestamp 	<ul style="list-style-type: none"> • boolean • binary
Oracle	<ul style="list-style-type: none"> • bigint • decimal • decimal • int • money • numeric • real • smallint • smallmoney • tinyint 	<ul style="list-style-type: none"> • char • nchar • nvarchar • text • varchar 	<ul style="list-style-type: none"> • date • datetime • datetime2 • datetimeoffset • smalldatetime 	bit

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
PostgreSQL	<ul style="list-style-type: none"> • bigint • decimal • double • integer • numeric • precision • real • smallint 	<ul style="list-style-type: none"> • char • character • text • varchar • varying character 	<ul style="list-style-type: none"> • date • timestamp 	<ul style="list-style-type: none"> • boolean
Apache Spark	<ul style="list-style-type: none"> • bigint • decimal • double • integer • real • smallint • tinyint 	<ul style="list-style-type: none"> • varchar 	<ul style="list-style-type: none"> • date • timestamp 	<ul style="list-style-type: none"> • boolean

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
Snowflake	<ul style="list-style-type: none"> • bigint • byteint • decimal • double • doubleprecision • float • float4 • float8 • int • integer • number • numeric • real • smallint • tinyint 	<ul style="list-style-type: none"> • char • character • string • text • varchar 	<ul style="list-style-type: none"> • date • datetime • time • timestamp • timestamp_* 	<ul style="list-style-type: none"> • boolean

Database engine or source	Numeric data types	String data types	Datetime data types	Boolean data types
Microsoft SQL Server	<ul style="list-style-type: none"> • bigint • bit • decimal • int • money • numeric • real • smallint • smallmoney • tinyint 	<ul style="list-style-type: none"> • char • nchar • nvarchar • text • varchar 	<ul style="list-style-type: none"> • date • datetime • datetime2 • smalldatetime 	<ul style="list-style-type: none"> • bit

Supported date formats

Amazon QuickSight supports the date and time formats described in this section. Before you add data to Amazon QuickSight, 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 Amazon S3 sources	Both 24-hour and	Supported date and time formats are described in the Joda API documentation.

Data source	Clocks	Date formats
Athena Salesforce	12-hour clocks	<p>For a complete list of Joda date formats, see Class DateTimeFormat on the Joda website.</p> <p>For datasets stored in memory (SPICE), Amazon QuickSight supports dates in the following range: Jan 1, 1400 00:00:00 UTC through Dec 31, 9999, 23:59:59 UTC .</p>

Data source	Clocks	Date formats
Relational databases sources	24-hour clock only	<p>The following data and time formats:</p> <ol style="list-style-type: none"> 1. dd/MM/yyyy HH:mm:ss , for example 31/12/2016 15:30:00. 2. dd/MM/yyyy , for example 31/12/2016. 3. dd/MMM/yyyy HH:mm:ss , for example 31/DEC/2016 15:30:00. 4. dd/MMM/yyyy , for example 31/DEC/2016. 5. dd-MMM-yyyy HH:mm:ss , for example 31-DEC-2016 15:30:00. 6. dd-MMM-yyyy , for example 31-DEC-2016. 7. dd-MM-yyyy HH:mm:ss , for example 31-12-2016 15:30:00. 8. dd-MM-yyyy , for example 31-12-2016. 9. MM/dd/yyyy HH:mm:ss , for example 12/31/2016 15:30:00. 10. MM/dd/yyyy , for example 12/31/2016. 11. MM-dd-yyyy HH:mm:ss , for example 12-31-2016 15:30:00. 12. MM-dd-yyyy , for example 12-31-2016. 13. MMM/dd/yyyy HH:mm:ss , for example DEC/31/2016 15:30:00. 14. MMM/dd/yyyy , for example DEC/31/2016.

Data source	Clocks	Date formats
		<p>15. MMM-dd-yyyy HH:mm:ss , for example DEC-31-2016 15:30:00.</p> <p>16. MMM-dd-yyyy , for example DEC-31-2016.</p> <p>17. yyyy/MM/dd HH:mm:ss , for example 2016/12/31 15:30:00.</p> <p>18. yyyy/MM/dd , for example 2016/12/31.</p> <p>19. yyyy/MMM/dd HH:mm:ss , for example 2016/DEC/31 15:30:00.</p> <p>20. yyyy/MMM/dd , for example 2016/DEC/31.</p> <p>21. yyyy-MM-dd HH:mm:ss , for example 2016-12-31 15:30:00.</p> <p>22. yyyy-MM-dd , for example 2016-12-31.</p> <p>23. yyyy-MMM-dd HH:mm:ss , for example 2016-DEC-31 15:30:00.</p> <p>24. yyyy-MMM-dd , for example 2016-DEC-31.</p> <p>25. yyyyMMdd'T'HHmmss , for example 20161231T153000.</p> <p>26. yyyy-MM-dd'T'HH:mm:ss , for example 2016-12-31T15:30:00.</p> <p>27. yyyyMMdd'T'HHmmss.SSS , for example 20161231T153000.123.</p> <p>28. MM/dd/yyyy HH:mm:ss.SSS , for example 12/31/2016 15:30:00.123.</p>

Data source	Clocks	Date formats
		<p>29. dd/MM/yyyy HH:mm:ss.SSS , for example 31/12/2016 15:30:00.123.</p> <p>30. yyyy/MM/dd HH:mm:ss.SSS , for example 2016/12/31 15:30:00.123.</p> <p>31. MMM/dd/yyyy HH:mm:ss.SSS , for example DEC/31/2016 15:30:00.123.</p> <p>32. dd/MMM/yyyy HH:mm:ss.SSS , for example 31/DEC/2016 15:30:00.123.</p> <p>33. yyyy/MMM/dd HH:mm:ss.SSS , for example 2016/DEC/31 15:30:00.123.</p> <p>34. yyyy-MM-dd'T'HH:mm:ss.SSS , for example 2016-12-31T15:30:00.123.</p> <p>35. MM-dd-yyyy HH:mm:ss.SSS , for example 12-31-2016 15:30:00.123.</p> <p>36. dd-MM-yyyy HH:mm:ss.SSS , for example 31-12-2016 15:30:00.123.</p> <p>37. yyyy-MM-dd HH:mm:ss.SSS , for example 2016-12-31 15:30:00.123.</p> <p>38. MMM-dd-yyyy HH:mm:ss.SSS , for example DEC-31-2016 15:30:00.123.</p> <p>39. dd-MMM-yyyy HH:mm:ss.SSS , for example 31-DEC-2016 15:30:00.123.</p> <p>40. yyyy-MMM-dd HH:mm:ss.SSS , for example 2016-DEC-31 15:30:00.123.</p>

Unsupported values in data

If a field contains values that don't conform with the data type that Amazon QuickSight 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 QuickSight 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

Amazon QuickSight Connection examples

The following connections are supported in Amazon QuickSight. 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 QuickSight](#)
- [Creating a dataset using Amazon S3 files](#)
- [Creating a data source using Apache Spark](#)
- [Using Databricks in QuickSight](#)
- [Creating a dataset using Google BigQuery](#)
- [Creating a dataset using a Microsoft Excel file](#)
- [Creating a data source using Presto](#)
- [Using Starburst with Amazon QuickSight](#)
- [Creating a data source and data set from SaaS sources](#)
- [Creating a dataset from Salesforce](#)
- [Using Trino with Amazon QuickSight](#)
- [Creating a dataset using a local text file](#)
- [Using Amazon Timestream data with Amazon QuickSight](#)

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 **Datasets** from the navigation pane at left, then choose **New dataset**.

2. a. To use an existing Athena connection profile (common), scroll down to the **FROM EXISTING DATA SOURCES** section, and choose the card for the existing data source that you want to use. Choose **Create dataset**.

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. In the **FROM NEW DATA SOURCES** section, choose the **Athena** data source card.
 2. For **Data source name**, enter a descriptive name.
 3. For **Athena workgroup**, choose your workgroup.
 4. Choose **Validate connection** to test the connection.
 5. Choose **Create data source**.
 6. (Optional) Select an IAM role ARN for queries to run as.
3. 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 QuickSight administrator for assistance. For more information, see [Authorizing connections to Amazon Athena](#).

4. Choose **Edit/preview data**.
5. Create a dataset and analyze the data using the table by choosing **Visualize**. For more information, see [Visualizing data in Amazon QuickSight](#).

Using Amazon OpenSearch Service with Amazon QuickSight

Following, you can find how to connect to your Amazon OpenSearch Service data using Amazon QuickSight.

Creating a new QuickSight data source connection for OpenSearch Service

Following, you can find how to connect to OpenSearch Service

Before you can proceed, Amazon QuickSight 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 QuickSight administrator can authorize connections to Amazon resources.

To authorize QuickSight to initiate a connection to OpenSearch Service

1. Open the menu by clicking on your profile icon at top right, then choose **Manage QuickSight**. If you don't see the **Manage QuickSight** option on your profile menu, ask your QuickSight 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 **Datasets** from the navigation pane at left, then choose **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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight before you can add them.

To edit permissions on a data source

1. Choose **Datasets** 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 QuickSight 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 QuickSight 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 **Datasets, 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 QuickSight 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 Amazon QuickSight. Amazon QuickSight 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 Amazon QuickSight access to any Amazon S3 buckets that you want to read files from. For information about granting Amazon QuickSight access to Amazon resources, see [Accessing data sources](#).

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 Amazon QuickSight start page, choose **Datasets**.
5. On the **Datasets** page, choose **New dataset**.
6. In the **FROM NEW DATA SOURCES** section of the **Create a Data Set** page, choose the Amazon S3 icon.
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 Amazon QuickSight.
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 QuickSight](#).

Datasets based on multiple Amazon S3 files

You can use one of several methods to merge or combine files from Amazon S3 inside Amazon QuickSight:

- **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. Amazon QuickSight 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 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 Amazon QuickSight 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 Amazon QuickSight, 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 Amazon QuickSight account](#)
- [Setting up Amazon QuickSight to access Amazon S3 files in another Amazon account](#)

Setting up Amazon S3 to allow access from a different Amazon QuickSight account

Use this section to learn how to set permissions in Amazon S3 files so they can be accessed by Amazon QuickSight in another Amazon account.

For information on accessing another account's Amazon S3 files from your Amazon QuickSight account, see [Setting up Amazon QuickSight 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 QuickSight 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 Amazon QuickSight. 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 Amazon QuickSight needs access to.
9. Notify the Amazon QuickSight user that the files are now available for use.

Setting up Amazon QuickSight to access Amazon S3 files in another Amazon account

Use this section to learn how to set up Amazon QuickSight 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 Amazon QuickSight account, see [Setting up Amazon S3 to allow access from a different Amazon QuickSight account](#).

Use the following procedure to access another account's Amazon S3 files from Amazon QuickSight. 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 QuickSight

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 Amazon QuickSight**.
3. Choose **Security & permissions**.
4. Under **QuickSight 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 Amazon QuickSight account**. It shows all the buckets your Amazon QuickSight 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 Amazon QuickSight account." This error message appears if you don't have either account permissions or Amazon QuickSight permissions.

Note

To use Amazon Athena, Amazon QuickSight 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](#).

Supported formats for Amazon S3 manifest files

You use JSON manifest files to specify files in Amazon S3 to import into Amazon QuickSight. These JSON manifest files can use either the Amazon QuickSight 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 Amazon QuickSight 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, Amazon QuickSight processes the optional `mandatory` option as Amazon Redshift does. If the associated file isn't found, Amazon QuickSight 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. Amazon QuickSight supports UTF-8 file encoding, but not UTF-8 (with ap-south-1). 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 Amazon QuickSight access to. For information about granting Amazon QuickSight access to Amazon resources, see [Accessing data sources](#).

Manifest file format for Amazon QuickSight

Amazon QuickSight 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.

Amazon QuickSight 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 Amazon QuickSight 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/awsexamplebucket/data.csv</code>	
<code>s3://<bucket name>/<file name></code>	<code>s3://awsexamplebucket/data.csv</code>	
<code>https://<bucket name>.s3.amazonaws.com/<file name></code>	<code>https://awsexamplebucket.s3.amazonaws.com/data.csv</code>	
<code>https://s3-<region name>.amazonaws.com/<bucket name>/<file name></code>	<code>https://s3-us-east-1.amazonaws.com/awsexamplebucket/data.csv</code>	This URI type identifies the Amazon Web Services Region for the Amazon S3 bucket.
<code>https://<bucket name>.s3-<region name>.amazonaws.com/<file name></code>	<code>https://awsexamplebucket.s3-us-</code>	This URI type identifies

URI format	Example	Comments
	<i>east-1 .amazonaws .com/data.csv</i>	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. Amazon QuickSight recursively retrieves files from child folders.

QuickSight 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 QuickSight 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/awsexamplebucket/</code>	
<code>https://s3.amazonaws.com/<bucket name>/<folder name1>/(<folder name2>/etc.)</code>	<code>https://s3.amazonaws.com/awsexamplebucket/folder1/</code>	
<code>s3://<bucket name></code>	<code>s3://awsexamplebucket/</code>	
<code>s3://<bucket name>/<folder name1>/(<folder name2>/etc.)</code>	<code>s3://awsexamplebucket/folder1/</code>	
<code>https://<bucket name>.s3.amazonaws.com</code>	<code>https://awsexamplebucket.s3.amazonaws.com</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 /awsexamplebucket /</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 /awsexamplebucket /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://awsexamplebucket .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, Amazon QuickSight 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 Amazon QuickSight

The following are some examples of completed Amazon QuickSight 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/awsexamplebucket/data.tsv"
      ]
    },
    {
      "URIPrefixes": [
        "https://s3-us-east-1.amazonaws.com/awsexamplebucket/"
      ]
    }
  ],
  "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 Amazon QuickSight 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://awsexamplebucket.your-s3-url.com",
        "s3://awsexamplebucket2/"
      ]
    }
  ],
  "globalUploadSettings": {

```

```

    "format": "CLF"
  }
}

```

The following example uses the Amazon Redshift format to identify a .csv file for import.

```

{
  "entries": [
    {
      "url": "https://awsexamplebucket.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 a data source using Apache Spark

You can connect directly to Apache Spark using Amazon QuickSight, 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 Amazon QuickSight. You can either directly query your data through Spark, or you can import the results of your query into [SPICE](#).

Before you use Amazon QuickSight with Spark products, you must configure Spark for Amazon QuickSight.

Amazon QuickSight 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, Amazon QuickSight refuses the connection to the server. To use Amazon QuickSight 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 Amazon QuickSight access, follow the instructions in [Network and database configuration requirements](#).

Using Databricks in QuickSight

Use this section to learn how to connect from QuickSight to Databricks.

To connect to Databricks

1. Begin by creating a new dataset. Choose **Datasets** from the navigation pane at left, then choose **New Dataset**.
2. Choose the **Databricks** data source card.
3. 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.

The following screenshot shows the connection screen for Databricks.

New Databricks data source ×

Data source name

Connection type

Public network ∨

Database server

HTTP Path

Port

Username

Password

SSL is enabled

4. 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.

5. For **Database server**, enter the **Hostname of workspace** specified in your Databricks connection details.
6. For **HTTP Path**, enter the **Partial URL for the spark instance** specified in your Databricks connection details.
7. For **Port**, enter the **port** specified in your Databricks connection details.
8. For **Username** and **Password**, enter your connection credentials.
9. To verify the connection is working, click **Validate connection**.
10. To finish and create the data source, click **Create data source**.

Adding a new QuickSight 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 **Datasets** 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**. The following popup displays:

Choose your table ×

t1

Catalog: contain sets of schemas.

hive_metastore ∨

Schema: contain sets of tables.

default ∨

Tables: contain the data you can visualize.

all_flights

singletuplestocktradedtable

[Edit/Preview data](#) [Use custom SQL](#) [Select](#)

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, choose **Augment with SageMaker**. This option is only available in QuickSight 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.

Amazon QuickSight Administrator's guide to connecting Databricks

You can use Amazon QuickSight 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 QuickSight 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 QuickSight, see [Connecting to a VPC with Amazon QuickSight](#) in the QuickSight 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 QuickSight 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:

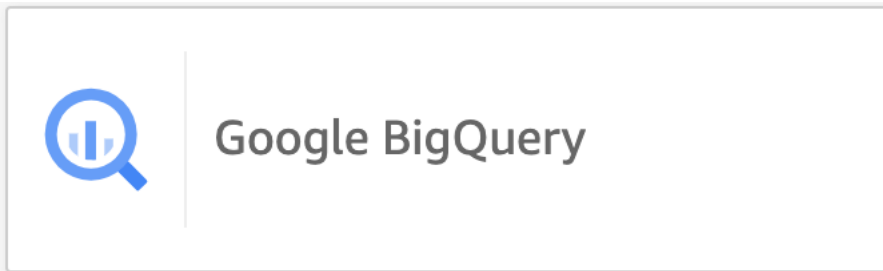
- **Project ID** – The project ID that is associated with your Google account. To find this, navigate to the Google Cloud console and choose the name of the project that you want to connect to QuickSight. Copy the project ID that appears in the new window and record it for later use.
- **Dataset Region** – 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 QuickSight, 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 Amazon QuickSight](#).

Use the following procedure to connect your QuickSight account with your Google BigQuery data source.

To create a new connection to a Google BigQuery data source from Amazon QuickSight

1. Open the [QuickSight console](#).
2. From the left navigation pane, choose **Datasets**, and then choose **New Dataset**.
3. Choose the **Google BigQuery** tile.



4. 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.
5. Choose **Sign in**.
6. In the new window that opens, enter the login credentials for the Google account that you want to connect to.
7. Choose **Continue** to grant QuickSight access to Google BigQuery.
8. After you create the new data source connection, continue to [Step 4](#) in the following procedure.

Adding a new QuickSight 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 [QuickSight console](#).
2. From the start page, choose **Datasets**, and then choose **New Dataset**.

3. On the **Create a dataset** page that opens, choose the **Google BigQuery** tile, and then choose **Create dataset**.
4. 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 QuickSight, see [Using SQL to customize data](#).
5. Choose **Edit/Preview**.
6. (Optional) In the **Data prep** page that opens, you can add customizations to your data with calculated fields, filters, and joins.
7. When you are finished making changes, choose **Save** to save and close the dataset.

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 Amazon QuickSight start page, choose **Datasets**.
3. On the **Datasets** page, choose **New dataset**.
4. In the **FROM NEW DATA SOURCES** section of the **Create a Data Set** page, 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 Amazon QuickSight.

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 QuickSight](#).

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 QuickSight. For more information about possible implications of changing datasets, see [Changing 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 Amazon QuickSight datasets. Presto processes the analytic queries on the backend databases. Then it returns results to the Amazon QuickSight client. You can directly query your data through Presto, or you can import the results of your query into SPICE.

Before you use Amazon QuickSight as a Presto client to run queries, make sure that you configure data source profiles. You need a data source profile in Amazon QuickSight 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 QuickSight (console)

1. On the Amazon QuickSight start page, choose **Datasets** at top right. Then choose **New dataset**.
2. 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.

3. 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 [Connecting to a VPC with Amazon QuickSight](#).

- **Database server** – The name of the database server.
- **Port** – The port that the server using to accept incoming connections from Amazon QuickSight
- **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. Amazon QuickSight applies the same user name and password to all connections that use this data source profile. If you want to monitor Amazon QuickSight separately from other accounts, create a Presto account for each Amazon QuickSight 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 QuickSight encrypts all credentials that you use in data source profile. For more information, see [Data encryption in Amazon QuickSight](#).
 - **Enable SSL** – SSL is enabled by default.
4. Choose **Validate connection** to test your settings.
 5. After you validate your settings, choose **Create data source** to complete the connection.


Using Starburst with Amazon QuickSight

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 QuickSight to

Starburst. All traffic between QuickSight 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 **QuickSight**. You should be able to see information, such as hostname and port. Amazon QuickSight supports basic username and password authentication to Starburst.

Creating a data source connection for Starburst

1. Begin by creating a new dataset. From the left navigation pane, choose **Datasets**, 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.

You should see the following data source creation modal.

New Starburst data source ×

Product Type

Starburst Enterprise Starburst Galaxy

Data source name

Connection type

Public network ▾

Database server

Port

Catalog

Username

Password

SSL is enabled

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 QuickSight, see [Configuring the VPC connection in Amazon QuickSight](#). 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 QuickSight and Starburst was validated using Starburst version 420.

Adding a new Amazon QuickSight dataset for Starburst

After you go through the [data source creation process](#) for Starburst, you can create Starburst datasets to use for analysis. You can create new datasets from a new or an existing Starburst data source. When you are creating a new data source, Amazon QuickSight 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 Starburst data source, see the following steps.

1. From the start page, choose **Datasets** and then choose **New dataset** at the top right.
2. Scroll down to the section **FROM EXISTING DATA SOURCES** and choose the Starburst data source you just created in the section above.
3. Choose **Create data set**.
4. To specify the table you want to connect to, choose a schema. The screenshot below shows a chosen sample schema. If you don't want to choose a schema, you can also use your own SQL statement.

Choose your table ×

Starburst Data Source

Schema: contain sets of tables.

simpledb ▾

Tables: contain the data you can visualize.

all_flights

all_flights_snapi

single_tuple_stock_trade_table

[Edit/Preview data](#) [Use custom SQL](#) [Select](#)

5. To specify the table you want to connect to, first select the **Schema** you want to use. For **Tables**, select 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** at 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 checkbox 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 QuickSight and Starburst was validated using Starburst version 420.

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 Amazon QuickSight. 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 Amazon QuickSight can access the application's Domain Name System (DNS) name over the network. If Amazon QuickSight 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 Amazon QuickSight 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 Amazon QuickSight start page, choose **Datasets**.
2. On the **Datasets** page, choose **New dataset**.

3. In the **FROM NEW DATA SOURCES** section of the **Create a Data Set** page, 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 Amazon QuickSight.

The following example shows the authorization for Amazon QuickSight to access the GitHub account for the Amazon QuickSight documentation.

 **Note**

Amazon QuickSight 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 Amazon QuickSight. If you want to do this, follow the prompts on your SaaS screen, then choose to authorize Amazon QuickSight.

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 Amazon QuickSight 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 Amazon QuickSight start page, choose **Datasets**.
3. On the **Datasets** page, choose **New dataset**.
4. In the **FROM NEW DATA SOURCES** section of the **Create a Data Set** page, 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 Amazon QuickSight 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 QuickSight

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 QuickSight to Trino. All traffic between Amazon QuickSight and Trino is enabled by SSL. Amazon QuickSight 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 **Datasets**, then choose **New Dataset**.
2. Choose the **Trino** data source card.



You should see the following data source creation modal.

New Trino data source ✕

Data source name

Connection type

Public network ▾

Database server

Port

Catalog

Username

Password

SSL is enabled

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 QuickSight, see [Configuring the VPC connection in Amazon QuickSight](#).
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 QuickSight 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 QuickSight 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 **Datasets** and then choose **New dataset** in the top right.
2. Scroll down to the section that says **FROM EXISTING DATA SOURCES** and 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. The screenshot below shows a chosen sample schema. If you don't want to choose a schema, you can also use your own SQL statement.

Choose your table ×

Trino Data Source

Schema: contain sets of tables.

simpleddb ▾

Tables: contain the data you can visualize.

all_flights

all_flights_snapi

single_tuple_stock_trade_table

[Edit/Preview data](#) [Use custom SQL](#) [Select](#)

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 QuickSight 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 Amazon QuickSight start page, choose **Datasets**.
3. On the **Datasets** page, choose **New dataset**.
4. In the **FROM NEW DATA SOURCES** section of the **Create a Data Set** page, 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 Amazon QuickSight.

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 QuickSight](#).

Using Amazon Timestream data with Amazon QuickSight

Following, you can find how to connect to your Amazon Timestream data using Amazon QuickSight. For a brief overview, see the [Getting started with Amazon Timestream and Amazon QuickSight](#) video tutorial on YouTube.

Creating a new Amazon QuickSight data source connection for a Timestream database

Following, you can find how to connect to Amazon Timestream from Amazon QuickSight.

Before you can proceed, Amazon QuickSight needs to be authorized to connect to Amazon Timestream. If connections aren't enabled, you get an error when you try to connect. A QuickSight 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 [Accessing data sources](#).

To connect to Amazon Timestream

1. Begin by creating a new dataset. Choose **Datasets** from the navigation pane at left, then choose **New Dataset**.
2. Choose the Timestream data source card.
3. 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.
4. Choose **Validate connection** to check that you can successfully connect to Timestream.
5. Choose **Create data source** to proceed.
6. For **Database**, choose **Select** to view the list of available options.
7. Choose the one you want to use, then choose **Select** to continue.
8. Do one of the following:
 - To import your data into QuickSight's in-memory engine (called SPICE), choose **Import to SPICE for quicker analytics**.
 - To allow QuickSight 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.

9. Choose **Edit/Preview** and then **Save** to save your dataset and close it.
10. 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 QuickSight

dataset. Currently, QuickSight 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 QuickSight.

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 QuickSight before you can add them.

To edit permissions on a dataset

1. Choose **Datasets** 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 QuickSight 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 QuickSight 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 QuickSight data source connection for each table.

To create a dataset using Amazon Timestream

1. Choose **Datasets** 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, choose **Augment with SageMaker**. This option is only available in QuickSight 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 QuickSight 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](#).

Creating datasets

To create a dataset, choose **New data set** on the **Datasets** 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.

Topics

- [Creating datasets using new data sources](#)
- [Creating a dataset using an existing data source](#)
- [Creating a dataset using an existing dataset in Amazon QuickSight](#)

Creating datasets using new data sources

To create a dataset from a new data source, 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 datasets from new database data sources

You can use a variety of database data sources to provide data to Amazon QuickSight. 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.

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](#).

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 QuickSight account. For details on pricing, see the pricing page for the service in question.

Topics

- [Creating a dataset from a database](#)

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 QuickSight 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 QuickSight access by following the instructions in [Network and database configuration requirements](#).
4. On the Amazon QuickSight start page, choose **Datasets**.
5. On the **Datasets** page, choose **New dataset**.
6. In the **FROM NEW DATA SOURCES** section of the **Create a Data Set** page, 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 QuickSight 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 QuickSight access by following the instructions in [Network and database configuration requirements](#).
4. On the Amazon QuickSight start page, choose **Manage data**.
5. On the **Datasets** page, choose **New data set**.
6. In the **FROM NEW DATA SOURCES** section of the **Create a Data Set** page, 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 [QuickSight 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 QuickSight doesn't accept self-signed SSL certificates as valid. For a list of accepted certificates, see [QuickSight SSL and CA certificates](#).

Amazon QuickSight 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 QuickSight 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 QuickSight 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 QuickSight start page, choose **Datasets**.
2. On the **Datasets** page, choose **New dataset**.
3. In the **FROM EXISTING DATA SOURCES** section of the **Create a Data Set** page, 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 QuickSight start page, choose **Manage data**.
2. On the **Datasets** page, choose **New data set**.

In the **FROM EXISTING DATA SOURCES** section of the **Create a Data Set** page, 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 QuickSight 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 QuickSight start page, choose **Manage data**.
2. On the **Datasets** page, choose **New data set**.
3. In the **FROM EXISTING DATA SOURCES** section of the **Create a Data Set** page, 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 QuickSight start page, choose **Manage data**.
2. On the **Datasets** page, choose **New data set**.
3. In the **FROM EXISTING DATA SOURCES** section of the **Create a Data Set** page, 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 QuickSight

After you create a dataset in Amazon QuickSight, 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 QuickSight 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 QuickSight 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 QuickSight 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 = product quantity x 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 QuickSight 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 QuickSight start page, choose **Datasets** in the pane at left.
2. On the **Datasets** page, 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 QuickSight](#).
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 \(RLS\) with user-based rules to restrict access to a dataset](#). For more information about setting up CLS, see [Using column-level security \(CLS\) 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 QuickSight, 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 QuickSight start page, choose **Datasets** in the pane at left.
2. On the **Datasets** page, choose the dataset that you want to restrict creating new datasets from.
3. On the page that opens for that dataset, choose **Edit dataset**.



DELETE DATASET

EDIT DATASET

USE IN ANALYSIS

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 Amazon QuickSight data preparation functionality, see [Preparing data in Amazon QuickSight](#).

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.

Editing a dataset from the Datasets page

1. From the QuickSight start page, choose **Datasets** at left.
2. On the **Datasets** 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 QuickSight](#).

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 QuickSight](#).

Reverting datasets back to previous published versions

When you save and publish changes to a dataset in Amazon QuickSight, 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 QuickSight start page, choose **Datasets**.
2. On the **Datasets** 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, QuickSight 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 QuickSight start page, choose **Datasets** at left.
2. On the Datasets page, choose the dataset that you want to duplicate.
3. On the dataset details page that opens, choose the drop-down for **Edit datasource**, 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.

Changing 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 QuickSight 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 QuickSight](#).

Sharing datasets


You can give other Amazon QuickSight 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 QuickSight start page, choose **Datasets** at left.
2. On the **Datasets** 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 QuickSight 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 QuickSight start page, choose **Datasets** at left.
2. On the **Datasets** 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 QuickSight start page, choose **Datasets** at left.
2. On the **Datasets** 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 Amazon QuickSight, 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 QuickSight start page, choose **Datasets** in the pane at left.
2. On the **Datasets** 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 QuickSight

In Amazon QuickSight, 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 QuickSight](#). 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 QuickSight 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** – QuickSight 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 QuickSight](#)
- [Insert dataset parameters into custom SQL](#)

- [Add dataset parameters to calculated fields](#)
- [Add dataset parameters to filters](#)
- [Working with dataset parameters in QuickSight 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 QuickSight.

- 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.

Creating dataset parameters in Amazon QuickSight

Use the following procedures to get started using dataset parameters.

To create a new dataset parameter

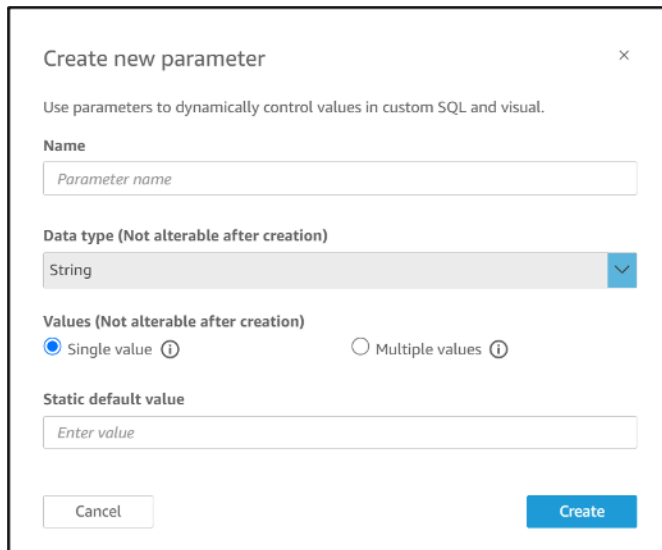
1. From the QuickSight start page, choose **Datasets** on the left, choose the ellipsis (three dots) next to the dataset that you want to change, and then choose **Edit**.
2. On the **Dataset** page that opens, choose **Parameters** on the left, and then choose the (+) icon to create a new dataset parameter.
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.

- 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.

- 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.
- When you are finished configuring the new parameter, choose **Create** to create the parameter.



The screenshot shows a dialog box titled "Create new parameter" with a close button (X) in the top right corner. Below the title is a subtitle: "Use parameters to dynamically control values in custom SQL and visual." The dialog contains the following fields and controls:

- Name:** A text input field with the placeholder text "Parameter name".
- Data type (Not alterable after creation):** A dropdown menu currently set to "String".
- Values (Not alterable after creation):** Two radio button options: "Single value" (which is selected) and "Multiple values". Both options have an information icon (i) next to them.
- Static default value:** A text input field with the placeholder text "Enter value".
- At the bottom, there are two buttons: "Cancel" on the left and "Create" on the right.

Insert 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, QuickSight 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**. In the following screenshot, the default value of the State parameter is IL, or Illinois. The SQL query filters the data from the dataset and returns all entries in the dataset where the origin state is IL.

The screenshot shows the Amazon QuickSight interface. On the left, there are navigation options: Fields, Filters (with 'State' selected), Parameters, and Schema. The main area is titled 'Data' and contains a 'Custom SQL name' field with 'New custom SQL' and a 'Custom SQL' text area containing the query: `1 SELECT * FROM all_flights` and `2 WHERE (origin_state_abr = <<$State>>)`. Below the query, the 'Data source' is set to 'Redshift'. At the bottom, there is a 'Dataset' tab for 'New custom SQL' showing a table of results. The table has columns: origin_city..., origin_state_abr, origin_stat..., origin_wac, dest_airpo..., dest_city..., dest_city..., and dest_state. The first row shows 'Bloomington...' for origin_city, 'IL' for origin_state_abr, 'Illinois' for origin_stat..., and so on. The 'Query mode' at the bottom left is set to 'Direct query' and shows '6.7TB of remaining'.

Add 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 QuickSight](#).

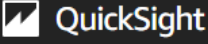
Add dataset parameters to filters

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**.



+ ADD ▾

Fields

Filters

Parameters

Community

Edit filter No filters applied ×

= [blurred] Include - all

Filter type
Custom filter ▾

Filter condition
Contains ▾

Use parameters

Select a parameter
[blurred] ▾

APPLY DELETE FILTER

Data

[blurred] Custom SQL ▾

Dataset [blurred]

[blurred]

String

Query mode

- SPICE
- Direct query

Working with dataset parameters in QuickSight 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 QuickSight 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 **Parameters** pane in the analysis menu on the left.
2. 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.

The following screenshot shows the configuration of a new analysis parameter that is mapped to a dataset parameter.

For more information on creating analysis parameters, see [Setting up parameters in Amazon QuickSight](#).

Create new parameter ×

Use parameters to dynamically control values in your fields, filters, and sheet.

Name
State

Data type (Not alterable after creation)
String

Values (Not alterable after creation)
Single value ⓘ

Static default value
IL

Dynamic default
[Set a dynamic default](#)

Dataset Parameters	Dataset	Actions
State	Flights by Arrival	Remove

ADD DATASET PARAMETER ▾ REMOVE ALL 1-1 of 1 < >

Advanced settings ▾

Cancel Create

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**.

Edit parameter ×

Use parameters to dynamically control values in your fields, filters, and sheet.

Name
State

Data type (Not alterable after creation)
String

Values (Not alterable after creation)
Single value ⓘ

Static default value
IL

Dynamic default
[Set a dynamic default](#)

Mapped Dataset Parameters ^

Dataset Parameters	Dataset	Actions
<input type="checkbox"/> State	Flights by Arrival	Remove
<input type="checkbox"/> State	Flights By Destination	Remove

ADD DATASET PARAMETER ▾ **REMOVE ALL** 1-2 of 2 < >

Advanced settings ▾

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.

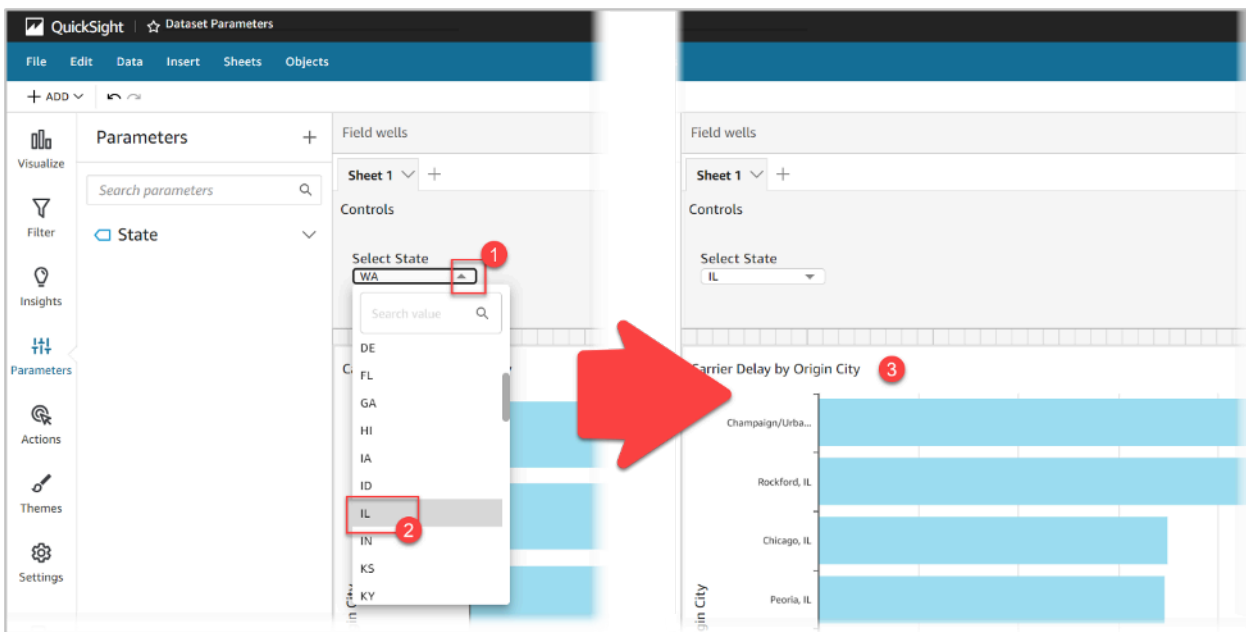
Add filter controls to mapped analysis parameters

After you map a dataset parameter to an analysis parameter in QuickSight, 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.

In the following screenshot, a dropdown control called **Select State** is created and applied to a sheet. When a user chooses **IL** from the control dropdown, the dataset parameter filters the data and only returns cities located in Illinois.



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.

The screenshot shows the Amazon QuickSight interface with the 'Format control' panel open for a 'State' parameter. The 'Link to a dataset field' option is selected, and the 'State Codes' dataset is chosen. The 'Field' dropdown is set to 'origin_state_abr'. The main visualization shows a large number '12.9' and a table of average arrival delays by origin state.

origin_state_abr	dest_state_abr	arr_delay
IL	MO	9.16
IL	NM	9.15
IL	IA	9.11
IL	IN	9.05
IL	NC	9
IL	MA	8.96
IL	KY	8.88

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.

The screenshot shows the Amazon QuickSight interface. On the left, the 'Dataset' is 'Flights by Arrival'. Below it is a 'Fields list' with a search bar and a list of fields including 'actual_elapsed_time', 'air_time', 'arr_del15', 'arr_delay', 'arr_delay_group', 'arr_time', 'cancellation_code', 'cancelled', 'carrier_delay', 'day_of_month', 'day_of_week', 'dep_del15', 'dep_delay', 'dep_delay_group', 'dep_time', 'dest_airport_id', 'dest_city_market_id', 'dest_city_name', and 'dest_state_abr'. On the right, the 'Field wells' section shows 'Sheet 1' and 'Controls'. A 'State' control is visible with a dropdown menu open, showing a search bar and a list of state abbreviations: AK, AL, AR, AZ, CA, CO, CT, DE, FL, GA, IL, NC, MA, and KY. The 'Select all' option is also visible in the dropdown menu. Below the dropdown, a table shows the 'arr_delay' values for each state: IL (9.18), NC (9.15), MA (9.11), GA (9.03), IL (9.03), NC (9), IL (8.96), and KY (8.88).

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.

The screenshot shows a dialog box titled "Edit parameter" with a close button (X) in the top right corner. Below the title is a subtitle: "Use parameters to dynamically control values in custom SQL and visual." The dialog contains several sections:

- Name:** A text input field containing "States".
- Data type (Not alterable after creation):** A label followed by "String".
- Values (Not alterable after creation):** A label followed by "Multiple values" and a help icon (i).
- Static multiple default values:** A large text area containing "All States" with a small 'X' icon in the bottom right corner.

 At the bottom of the dialog are two buttons: "Cancel" on the left and "Update" on the right.

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 QuickSight analyses](#). For more information on creating controls in an analysis that is using dataset parameters, see [Add 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 (RLS) in Amazon QuickSight

Applies to: Enterprise Edition

In the Enterprise edition of Amazon QuickSight, 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 QuickSight dashboards in your application for unregistered users of QuickSight, 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 QuickSight. 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 QuickSight. To do so, define tags on columns to secure data. Values to tags must be passed when embedding dashboards.

Topics

- [Using row-level security \(RLS\) with user-based rules to restrict access to a dataset](#)
- [Using row-level security \(RLS\) with tag-based rules to restrict access to a dataset when embedding dashboards for anonymous users](#)

Using row-level security (RLS) with user-based rules to restrict access to a dataset

Applies to: Enterprise Edition

In the Enterprise edition of Amazon QuickSight, 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 QuickSight 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 QuickSight 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 QuickSight 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 **Datasets** 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 QuickSight 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	

GroupName	Region	Segment	
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	
NikhiJayashankar	US	SMB, Startup	
PauloSantos	US	Startup	
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

UserARN	GroupARN	Region
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**.

Creating 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 **Datasets** page, choose the dataset
3. On the dataset details page that opens, for **Row-level security**, choose **Set up**.

The screenshot shows the 'About' section of a dataset. It includes a 'SPICE' badge and a size of 4.4MB. Below this is a 'REFRESH' section with a green checkmark indicating 'Status: Completed' and '9994 rows imported (100% success)'. The last successful refresh is dated July 26, 2022 at 2:11 PM PDT. The 'ACCESS SETTINGS' section is visible, with 'Sharing' set to 'Owners (1) Viewers (0)'. Under 'Row-level security', the current setting is 'No restrictions', and a red box highlights the 'Set up' link next to it.

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 QuickSight skips empty RLS rules that grant access to everyone. An *empty RLS rule* occurs when all columns of a row have no value. QuickSight 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.

Rules for UserName, GroupName, Region, Segment	Granted access	
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 QuickSight. 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 (RLS) with tag-based rules to restrict access to a dataset when embedding dashboards for anonymous users

Applies to: Enterprise Edition

Intended audience: Amazon QuickSight Administrators and Amazon QuickSight developers

When you embed Amazon QuickSight dashboards in your application for users who are not provisioned (registered) in QuickSight, 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 QuickSight, 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 QuickSight data 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 QuickSight 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](#).

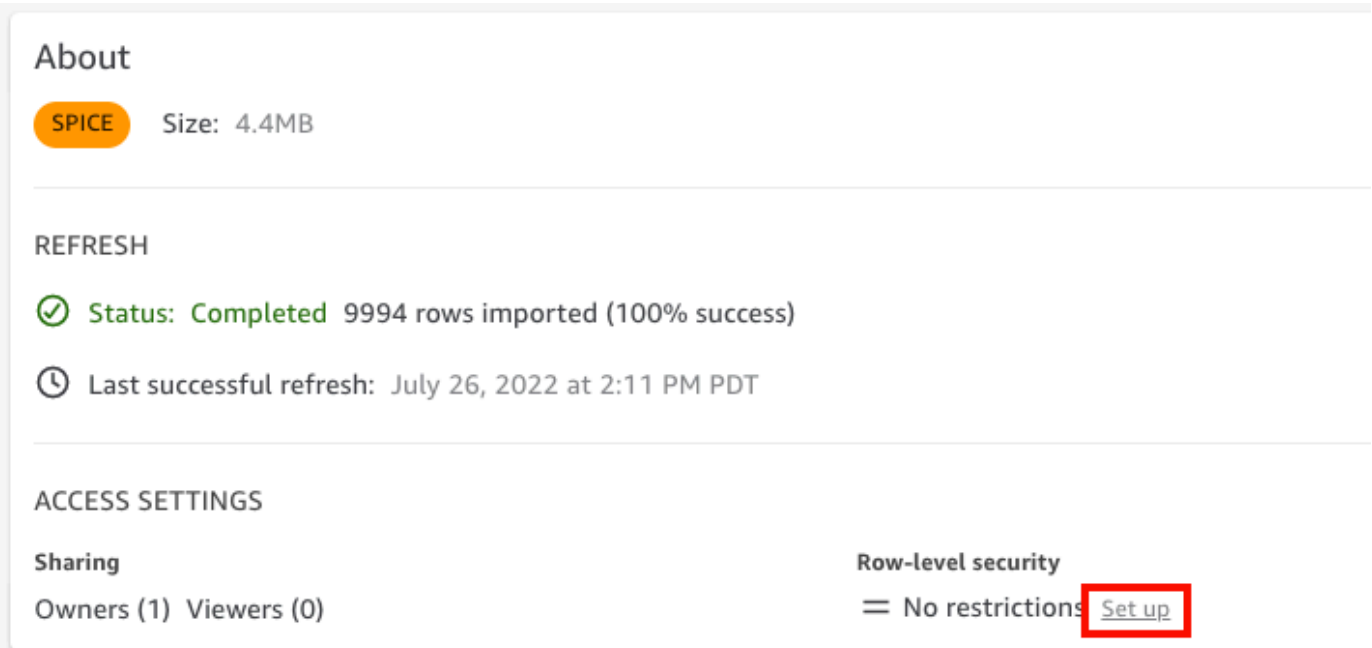
Step 1: Add RLS tags to a dataset

You can add tag-based rules to a dataset in Amazon QuickSight. 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 QuickSight.

To add RLS tags to a dataset

1. From the QuickSight start page, choose **Datasets** at left.
2. On the **Datasets** page, 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.

- 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.

- (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.

- (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

Column ⓘ	Tag	Delimiter ⓘ	Match all ⓘ
Select column ▼	Tag key	▼	▼

[Add](#)

- 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.

- When you're ready to apply the tag rules to the dataset, choose **Apply rules**.

Tag-based rules ⓘ

Choose this when restricting access to data in embedded dashboards for viewers who are not provisioned in QuickSight. Pass tag values to the QuickSight embedding API to define who sees the data. [Learn more](#)

There are unapplied changes to 1 rule. Unapplied rules don't affect readers.

[Apply rules](#)


- 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.

Tag-based rules ON ^

Choose this when restricting access to data in embedded dashboards for viewers who are not provisioned in QuickSight. Pass tag values to the QuickSight embedding API to define who sees the data. [Learn more](#)

 To enable dashboard and analysis authors see the data, add them to user-based rules and enable user-based rules setting above.

On the **Datasets** 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 QuickSight readers when active.

Important

After tags are assigned and enabled on the dataset, make sure to give QuickSight authors permissions to see any of the data in the dataset when authoring a dashboard. To give QuickSight 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**.

Manage tags

Search by column or tag 🔍

Column ⓘ	Tag	Delimiter ⓘ	Match all ⓘ	
Select column ▼	<input type="text" value="Tag key"/>	<input type="text" value=""/> ▼	<input type="text" value=""/> ▼	Add
segment	segment_tag			✎ 🗑
customer_id	customer_tag			✎ 🗑
industry	industry_tag			✎ 🗑

Manage rules

[Add OR Condition](#)

segment_tag AND
customer_tag AND
industry_tag
✎ 🗑

Cancel
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 QuickSight account users. When you use the OR condition with your tag-based rules, visuals in QuickSight 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**.

Manage tags

Search by column or tag 🔍

Column ⓘ	Tag	Delimiter ⓘ	Match all ⓘ	
Select column ▼	<input type="text" value="Tag key"/>	<input type="text" value=""/> ▼	<input type="text" value=""/> ▼	Add
segment	segment_tag			✎ 🗑
customer_id	customer_tag			✎ 🗑
industry	industry_tag			✎ 🗑

Manage rules

[Add OR Condition](#)

segment_tag AND
customer_tag AND
industry_tag
✎ 🗑

Cancel
Apply

- 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.

The screenshot shows the 'Manage rules' interface. At the top right is an 'Add OR Condition' button. Below it, a rule is displayed: 'segment_tag AND customer_tag AND industry_tag' followed by an 'OR' operator. A dropdown menu is open, showing 'Select tag' with a downward arrow and three options: 'segment_tag', 'customer_tag', and 'industry_tag'. To the right of the dropdown is an 'Update' button with a close icon. At the bottom right are 'Cancel' and 'Apply' buttons.

- Choose **Update** to add the condition to your rule, then choose **Apply** to apply the updated rule to your dataset.

The screenshot shows the 'Manage rules' interface after the update. The rule now consists of two parts: 'segment_tag AND customer_tag AND industry_tag' followed by an 'OR' operator, and 'customer_tag AND product_tag'. The 'Apply' button at the bottom right is highlighted with a red box.

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 QuickSight 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-2:111122223333: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 QuickSight authors permissions to see any of the data in the dataset when authoring a dashboard.

To give QuickSight 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 QuickSight 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 QuickSight 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 QuickSight data dashboards for anonymous \(unregistered\) users](#), or [GenerateEmbedUrlForAnonymousUser](#) in the *Amazon QuickSight API Reference*

Using column-level security (CLS) to restrict access to a dataset

In the Enterprise edition of Amazon QuickSight, 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 Amazon QuickSight 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 Amazon QuickSight start page, choose the **Datasets** tab.
2. On the Datasets page, open your dataset
3. On the dataset details page that opens, for **Column-level security**, choose **Set up**.

About

SPICE Size: 4.4MB

REFRESH

✔ Status: **Completed** 9994 rows imported (100% success)

🕒 Last successful refresh: July 26, 2022 at 2:11 PM PDT

ACCESS SETTINGS

Sharing

Owners (1) Viewers (0)

Row-level security

☰ No restrictions [Set up](#)

Column-level security

|| No restrictions [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 Amazon QuickSight navigation pane, choose the **Analyses** tab.
2. Choose the analysis that you want to create a dashboard of.
3. At upper right, choose **Share**, and then choose **Publish dashboard**.
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 QuickSight dashboards](#).

Run queries as an IAM role in Amazon QuickSight

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 Amazon QuickSight 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 QuickSight 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 QuickSight instance that uses it.

- The IAM role must have a trust relationship allowing QuickSight to assume the role.
- The identity that calls QuickSight'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 QuickSight connects to another service in Amazon, it uses an IAM role. By default, this less granular version of the role is created by QuickSight 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 QuickSight.

```
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 QuickSight 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 \  
  --aws-region us-east-1 \  
  --aws-iam-role TestAthenaRoleForQuickSight \  
  --aws-iam-policy-arn arn:aws-cn:iam::aws:policy/AmazonS3Access1
```

```
--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 QuickSight.

```
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"  
    },  
    {  
      "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`.

```
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": [
        "*"
      ]
    }
  ]
}
```

After you verify your permissions, you can use the role in QuickSight 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"
    ]
  }
]
}

```

Run queries with Amazon S3 data sources

Amazon S3 data sources contain a manifest file that QuickSight uses to find and parse your data. You can upload a JSON manifest file through the QuickSight console, or you can provide a URL that points to a JSON file in an S3 bucket. If you choose to provide a URL, QuickSight must be granted

permission to access the file in Amazon S3. Use the QuickSight 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.

New S3 data source



Data source name

Enter a name for the data source

Upload a [manifest file](#)

URL Upload

Enter URL of your JSON manifest file

Connect

To get started, create an Amazon S3 manifest file. Then, you can either upload it to Amazon QuickSight 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 QuickSight, 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 QuickSight 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 Amazon QuickSight 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 **Datasets** 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

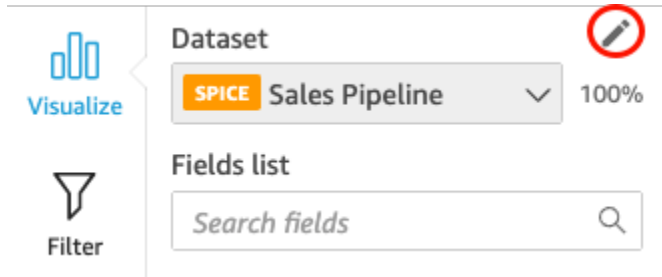
Topics

- [Add or edit a dataset](#)

- [Replacing datasets](#)
- [Remove a dataset from 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 **Fields list** 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 **Fields list** 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** on the application bar and then **Add Visual** to create a new visual using the selected dataset.



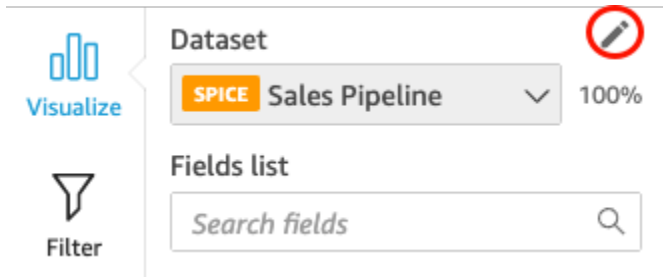
If you choose **Suggested** on the tool bar to see suggested visuals, you'll see visuals based on the currently selected dataset.

Only filters for the currently selected dataset are shown in the **Filter** pane, and you can only create filters on the currently selected dataset.

Add or edit a dataset


Use the following procedure to add a dataset to an analysis or edit a dataset used by an analysis.

1. On the analysis page, choose the pencil-shaped edit icon at the top of the **Fields list** pane.



2. Choose **Add dataset** to add a dataset. Or, choose the pencil-shaped edit icon to edit a dataset. For more information about editing a dataset, see [Editing datasets](#).

Datasets in this analysis

Dataset	Status
 Web and Social Media Analytics test SPICE	Available ⋮

3. A list of your datasets appears. Choose a dataset and then choose **Select**. To cancel, choose **Cancel**.

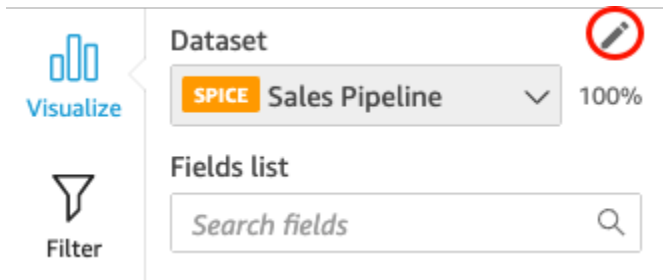
Replacing datasets

In an analysis, you can add, edit, replace, or remove datasets. Use this section to learn how to replace your dataset.

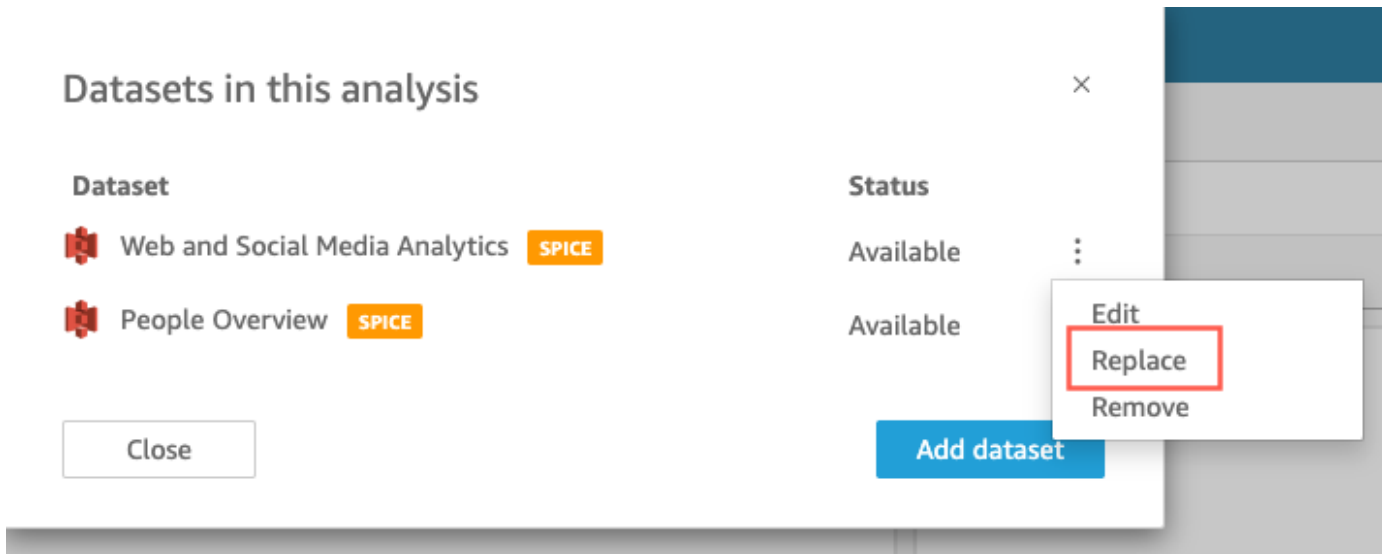
When you replace a dataset, the new dataset should have similar columns, if you expect the visual to work the way you designed it. Replacing the dataset also clears the undo and redo history for the analysis. This means you can't use the undo and redo buttons on the application bar to navigate your changes. 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 screen, choose the pencil icon above the dataset drop-down menu.



2. In the **Data sets in this analysis** page that opens, choose the ellipses next to the dataset that you want to replace, and then choose **Replace**.



3. 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.

4. In the **Replacing dataset** page, choose **Replace**.

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 Amazon QuickSight 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

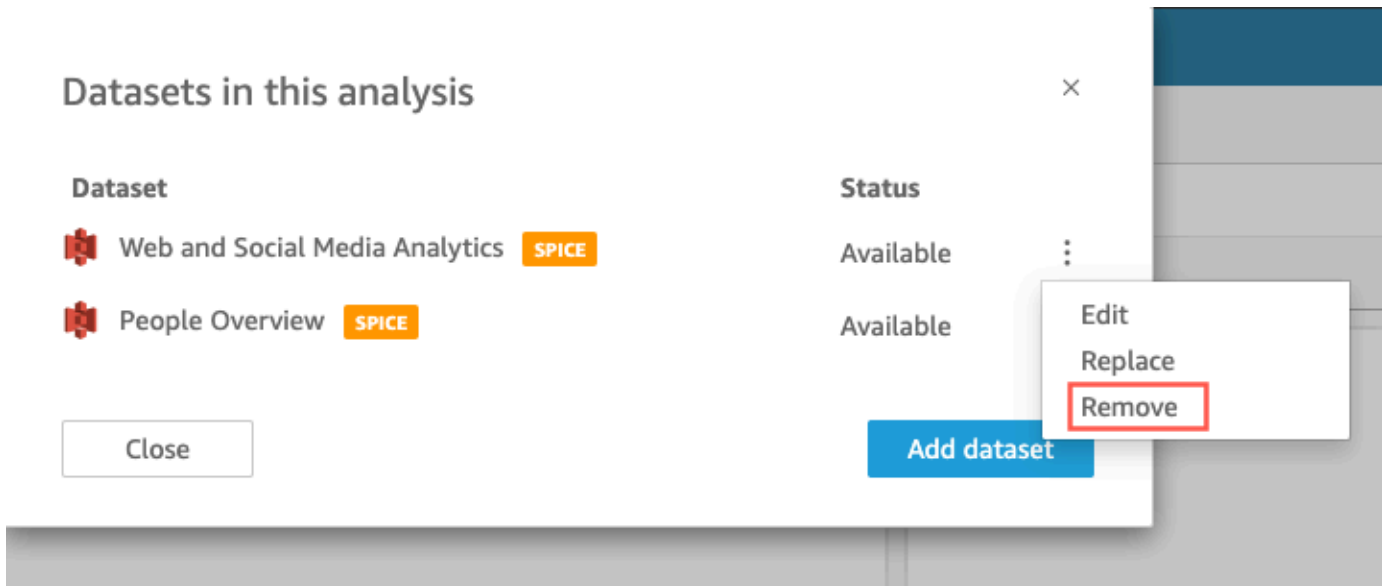
Use the following procedure to delete a dataset from an analysis.

To delete a dataset from an analysis

1. On the analysis page, choose the dataset list at the top of the **Fields list** pane, and then choose **Edit analysis datasets**.



- In the **Data sets in this analysis** dialog box, choose the dataset that you want to delete, and then choose the delete icon. You can't delete a dataset if it's the only one in the analysis.



- Choose **Close** to close the dialog box.

Working with data sources in Amazon QuickSight

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 QuickSight authors

As an analysis author in Amazon QuickSight, 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 Amazon QuickSight 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 QuickSight 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 QuickSight, 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 QuickSight. 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 QuickSight 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 QuickSight 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 QuickSight to use.

Every time Amazon QuickSight 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 QuickSight 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 QuickSight.

To view your existing connection profiles, open the Amazon QuickSight start page, choose **Datasets**, choose **New Dataset**, and then scroll to the heading **FROM EXISTING DATA SOURCES**.

For more information, see the following topics:

- [Creating a dataset from a database](#)
- [Creating datasets from new database data sources](#)
- [Creating a dataset using Amazon S3 files](#)
- [Creating a dataset using Amazon Athena data](#)
- [Creating a dataset from Salesforce](#)
- [Creating a data source using Presto](#)
- [Creating a data source using Apache Spark](#)
- [Creating a data source and data set from SaaS sources](#)

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 QuickSight start page, choose **Datasets** at left, and then choose **New dataset**.
2. Scroll down to the **FROM EXISTING DATA SOURCES** section and 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 QuickSight start page, choose **Datasets** at left, and then choose **New dataset**.
2. Scroll down to the **FROM EXISTING DATA SOURCES** section, and then 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. In the **FROM EXISTING DATA SOURCES** section of the **Create a Data Set** page, choose the data source that you want to delete.
2. Choose **Delete**.

Refreshing data in Amazon QuickSight

When refreshing data, Amazon QuickSight handles datasets differently depending on the connection properties and the storage location of the data.

If QuickSight 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, QuickSight must independently authenticate using stored credentials to connect to the data. QuickSight can't refresh manually uploaded data—even from S3 buckets, even though it's stored in SPICE—because QuickSight 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 QuickSight Standard Edition, you can do a full refresh of your SPICE data at any time. For Amazon QuickSight 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 QuickSight 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 QuickSight in a single Amazon Web Services Region. The other Amazon Web Services Regions have no SPICE capacity unless you choose to purchase some. QuickSight 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 QuickSight administrator can purchase more SPICE capacity or release unused SPICE capacity as needed. For more information, see [Managing 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 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. does not recommend that you use this formula to estimate the total SPICE capacity that your 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 on the **Datasets** page.

To refresh SPICE data from the datasets page

1. On the **Datasets** page, 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 QuickSight, 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. QuickSight queries the database asking for all data since June 24 (7 days ago), which is 7,000 records. QuickSight then deletes the data currently in SPICE from June 24 and after, and appends the newly queried data. The next day (July 2), QuickSight 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 on the **Datasets** page.

To incrementally refresh a SQL-based SPICE dataset

1. On the Datasets page, 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. On the **Datasets** page, 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 QuickSight, 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. On the Datasets page, 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 Amazon QuickSight 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 on the **Datasets** page.

To set an incremental refresh schedule for a SQL-based SPICE dataset

1. On the **Datasets** page, 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.

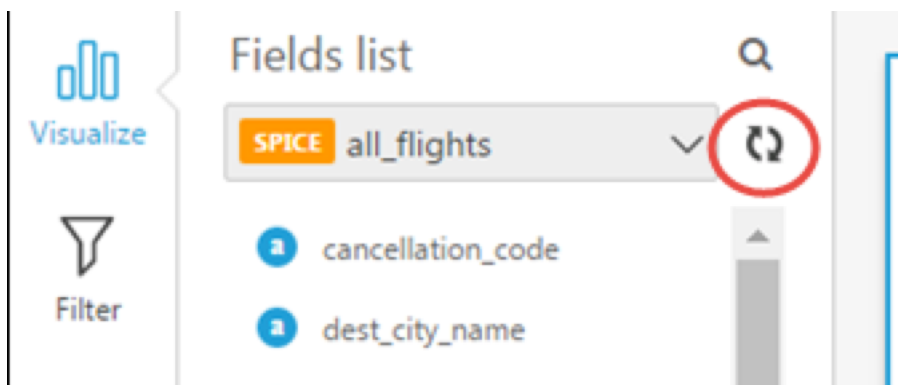
Scheduled dataset ingestions take place within 10 minutes of the scheduled date and time.

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, this icon appears as a spinner.



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 start screen, choose **Datasets** at left.
2. On the **Datasets** page, 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.

Q - SaaS Sales ADD NEW SCHEDULE REFRESH NOW

Summary **Refresh** Permissions Usage

Email owners when a refresh fails

Schedules

Refresh type	Occurrence	Start time	Timezone	Actions
No rows				

History

Show times within Last 90 days with status of All

Refresh start	St...	Duration	S...	In...	D...	Refresh type
July 27, 2022 at 11:41...	Completed	16 seconds	0	9994	9994	Manual, Full refresh
June 30, 2022 at 3:18 ...	Completed	16 seconds	0	9994	9994	Manual, Initial

1-2 of 2 < >


4. (Optional) Choose a time frame to filter the entries from the last hour to the last 90 days.

Show times within Last 90 days with status of All

- Last 90 days
- Last 30 days
- Last 2 weeks
- Last week
- Last 24 hours
- Last 6 hours
- Last hour

1-2 of 2 < >

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 QuickSight previews a portion of your data. If it can't interpret a row for any reason, it skips this row.

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](#), QuickSight skips rows where that field doesn't contain a valid date.

QuickSight 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. On the **Datasets** page, 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.

i 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 QuickSight.

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 QuickSight.
3. Import the data by following these steps:

- a. Choose **Datasets, 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.

Topics

- [Error codes for skipped rows](#)
- [Error codes during data import](#)

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 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 nondate 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 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 schema doesn't match the data type received from SageMaker.

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, QuickSight 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 QuickSight 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 QuickSight 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 QuickSight data sources.

DATA_TOLERANCE_EXCEPTION – There are too many invalid rows. Amazon QuickSight 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 QuickSight couldn't assume the correct Amazon Identity and Access Management (IAM) role. Verify the policies for Amazon QuickSight-service-role in the IAM console.

FAILURE_TO_PROCESS_JSON_FILE – Amazon QuickSight couldn't parse a manifest file as valid JSON.

IAM_ROLE_NOT_AVAILABLE – Amazon QuickSight doesn't have permission to access the data source. To manage Amazon QuickSight permissions on Amazon resources, go to the **Security and Permissions** page under the **Manage Amazon QuickSight** option as an administrator.

INGESTION_CANCELED – The ingestion was canceled by the user.

INGESTION_SUPERSEDED – This ingestion has been superseded by another workflow.

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 QuickSight 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 QuickSight 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 QuickSight and the Amazon QuickSight 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 QuickSight data source or dataset was deleted or became unavailable during ingestion. Check your dataset in Amazon QuickSight 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 QuickSight 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 QuickSight dataset. Update your Amazon QuickSight dataset definition.

SQL_TABLE_NOT_FOUND – Amazon QuickSight 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 QuickSight 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 QuickSight can't resolve the host name of the data source. Verify the host name of the data source and try again.

UNROUTABLE_HOST – Amazon QuickSight 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 QuickSight 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 QuickSight, choose **Datasets** at left.
2. On the **Datasets** page, 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.

QuickSight 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 QuickSight

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 QuickSight, 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 QuickSight](#)
- [Adding calculations](#)
- [Previewing tables in a dataset](#)
- [Joining data](#)
- [Filtering data in Amazon QuickSight](#)
- [Using SQL to customize data](#)
- [Adding geospatial data](#)
- [Using unsupported or custom dates](#)
- [Integrating Amazon SageMaker models with Amazon QuickSight](#)
- [Preparing dataset examples](#)


Describing data

Using Amazon QuickSight, 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. On the QuickSight start page, choose **Datasets** at left.
2. On the **Datasets** page, 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 QuickSight 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 QuickSight

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 Amazon QuickSight](#)
- [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**.

- # New visitors SEO
- # New visitors CPC
- # New visitors Social Media
- # Return visitors
- # Twitter mentions
- # Twitter followers adds
- # Twitter followers cumulative
- # Mailing list adds
- # Mailing list cumulative



2. Enter the new name or description that you want to change, and choose **Apply**.







Edit field ×

Name

Description

500 characters left

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.

Dataset					
Date	New visito...	New visito...	New visito...	Return visi...	Tw...
 Date	 Integer	 Integer	 Integer	 Integer	 I
2013-01-01...	2194	338	56	225	1
2013-01-02...	1653	254	42	169	0
2013-01-03...	2213	340	57	227	2
2013-01-04...	2223	342	57	228	0
2013-01-05...	1674	258	43	172	0
2013-01-06...	1681	259	43	172	0

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, Amazon QuickSight 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 Amazon QuickSight 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 Amazon QuickSight 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 QuickSight. 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.

QuickSight 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 QuickSight) 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. Amazon QuickSight 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 Amazon QuickSight, 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 Amazon QuickSight, see [Supported data types from other 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 Amazon QuickSight understand that they are not useful for performing mathematical calculations and can only be aggregated with the Count function. For more information about how Amazon QuickSight 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 QuickSight start page, choose **Datasets**, 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 Amazon QuickSight

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
Pivot table	Drill-down not supported
Scatter plot	Group/Color
Tabular Reports	Drill-down not supported
Tree map	Group by

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.

Note

You can't add drill-downs to pivot tables.

2. Click anywhere on the **Field wells** to expand them.
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.

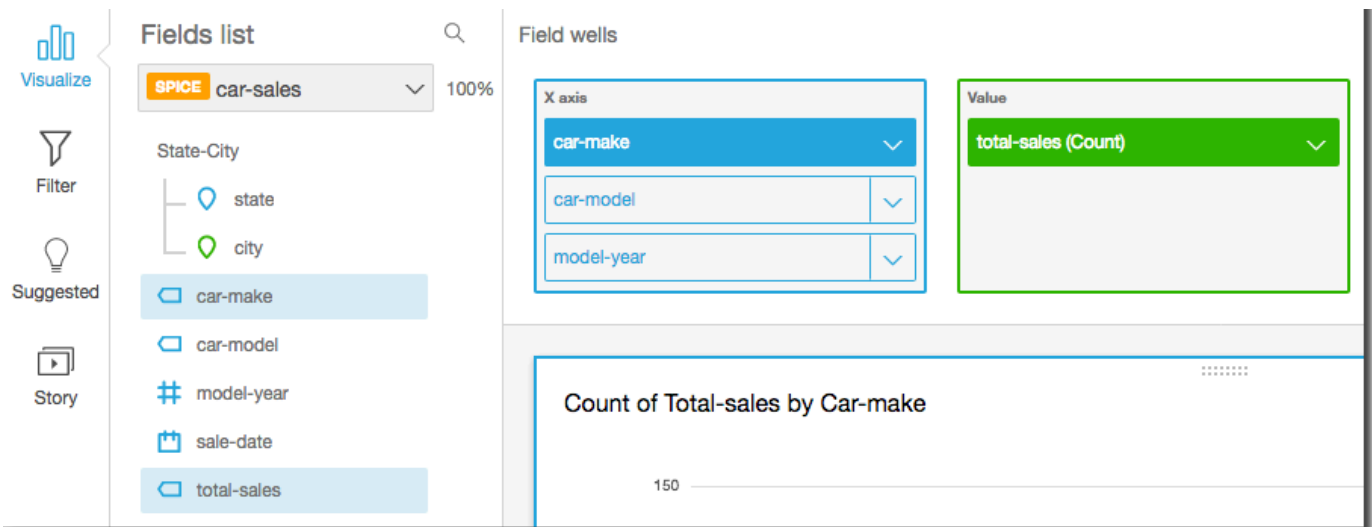
4. Drag a field that you want to use in the drill-down hierarchy to an appropriate field well, depending on the visual type. Make sure that the label for the dragged field says **Add drill-down layer**. Position the dragged field above or below the existing field based on where you want it to be in the hierarchy you're creating.

The screenshot shows the Amazon QuickSight interface. On the left is a navigation sidebar with icons for Visualize, Filter, Suggested, and Story. The main area is divided into three sections: Fields list, Field wells, and a visualization. The Fields list shows a search for 'car-sales' and a list of fields including 'State-City', 'state', 'city', 'car-make', 'car-model', 'model-year', 'sale-date', and 'total-sales'. The Field wells section shows 'car-make' in the X axis well and 'total-sales (Count)' in the Value well. The visualization is a bar chart titled 'Count of Total-sales by Car-make' with a single bar representing a value of 150.

This screenshot shows the same interface as the first, but with a tooltip for 'car-model' appearing over the 'car-make' field in the X axis well. The tooltip contains the text 'car-model' and 'Add drill-down layer'. The visualization remains the same, showing a bar chart with a value of 150.

This screenshot shows the 'car-model' field added as a second layer to the X axis well, below 'car-make'. A red circle highlights the 'car-model' field in the X axis well. The visualization area is currently blank, indicating that the chart has not yet been updated to reflect the new drill-down layer.

- 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**.



- 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 Amazon QuickSight, 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.

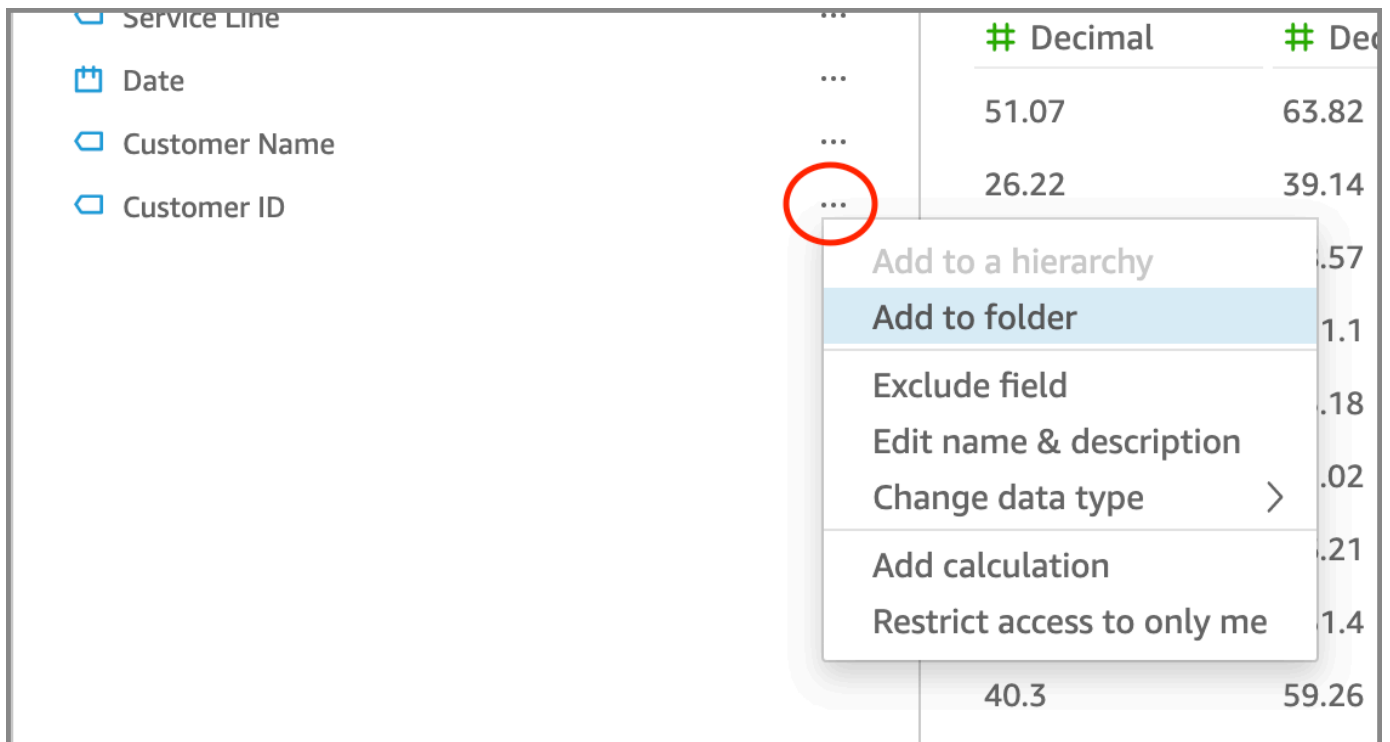
Creating a folder

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 field menu for the fields that you want to place in a folder 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 Amazon QuickSight, 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.

Topics

- [Mapping fields](#)
- [Joining fields](#)

Mapping fields

Amazon QuickSight can automatically map fields between datasets in the same analysis. The following tips can help make it easier for Amazon QuickSight 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 QuickSight](#).

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 QuickSight](#)
- [Using level-aware calculations in Amazon QuickSight](#)
- [Calculated field function and operator reference for Amazon QuickSight](#)

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. In your analysis, choose **Add** at top left, and then choose **Add 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.
2. When finished, choose **Save**.

For more information about how to create formulas using the available functions in QuickSight, see [Calculated field function and operator reference for Amazon QuickSight](#).


Adding calculated fields to a dataset

Amazon QuickSight 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, QuickSight computes and materializes the result in SPICE. If the calculated field relies on an aggregation function, QuickSight 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, QuickSight 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

Operator	Input type	Output type
round	Decimal-fixed	Decimal-fixed
	Int	Decimal-fixed
	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

	Integer	Decimal-fixed	Decimal-float
Decimal-fixed	Decimal-float	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 QuickSight

When you open or update an analysis, before displaying it Amazon QuickSight evaluates everything that is configured in the analysis in a specific sequence. Amazon QuickSight 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 QuickSight 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 Amazon QuickSight](#).



The following list shows the sequence in which Amazon QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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 Amazon QuickSight

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}])
```

Field wells

Group by

- Region
- Country
- Product

Value

- sum({Sales}) (Custom)
- sum({Sales}, [{Country}]) (Sum)

Sheet 1

Sum({sales}), Sum of Sum({sales} and [{Country}]) by Region, Country, and Product

Regi...	Country	Product	sum({Sales})	sum({Sales}, [{Country}])
AMER	Argentina	Big Ol Database	9,899.85	35,764.31
AMER	Argentina	ChatBot Plugin	742.8	35,764.31
AMER	Argentina	ContactMatcher	3,947.81	35,764.31
AMER	Argentina	Data Smasher	1,023.56	35,764.31
AMER	Argentina	FinanceHub	2,728.24	35,764.31
AMER	Argentina	Marketing Suite	2,275.88	35,764.31
AMER	Argentina	Marketing Suite - Gold	4,669.08	35,764.31
AMER	Argentina	OneView	4,204.36	35,764.31
AMER	Argentina	SaaS Connector Pack	950.97	35,764.31
AMER	Argentina	SaaS Connector Pack - Gold	153.7	35,764.31
AMER	Argentina	Site Analytics	3,577.75	35,764.31
AMER	Argentina	Storage	54.12	35,764.31
AMER	Argentina	Support	1,536.19	35,764.31
APJ	Australia	Alchemy		
APJ	Australia	Big Ol Database		
APJ	Australia	ChatBot Plugin	930.03	80,166.1

sum(sum({Sales}, [{Country}])

Page size 500

- 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).

Field wells

Group by

Region

Value

Sales (Average)

sum({Sales}, [{Country}]) (Average)

Sheet 1 +

Average of Sum({sales}, [{Country}]) and Average of Sales by Region

Region	Sales	sum({Sales}, [{Country}])
AMER	228.61	104,731.2
APJ	196.9	41,546.42
EMEA	247.43	34,796.23


avg({Sales}) avg(sum({Sales}, [{Country}]))

- 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 QuickSight 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 QuickSight evaluates queries, see [Order of evaluation in Amazon QuickSight](#).

The following calculation levels are supported for LAC-W functions:

- **PRE_FILTER** – Before applying filters from the analysis, QuickSight evaluates prefilter calculations. Then it applies any filters that are configured on these prefilter calculations.
- **PRE_AGG** – Before computing display-level aggregations, QuickSight 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 QuickSight

You can use the following functions and operators to create calculated fields.

Topics

- [Operators](#)
- [Functions by category](#)
- [Functions](#)
- [Aggregate functions](#)
- [Table calculation functions](#)

Operators

You can use the following operators in calculated fields. Amazon QuickSight 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 QuickSight 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 QuickSight 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: (<>) not equal

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 QuickSight, 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 QuickSight 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.

- [stdevlf](#) 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.
- [stdevplf](#) 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.
- [varlf](#) 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.
- [sumlf](#)) calculates the sum based on a conditional statement.

Conditional functions

The conditional functions for calculated fields in Amazon QuickSight 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.
- [nulllf](#) 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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight. 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 QuickSight.
- **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 QuickSight calculated field that returns a date value.

Example

```
calcFieldStartDate = addDateTime(10, "DD", startDate)
addWorkDays(calcFieldStartDate, numWorkDays)
```

- **Parameters** – Any QuickSight 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 QuickSight 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 QuickSight calculated field that returns a date value.

Example

```
someOtherIntegerCalcField = (num_days * 2) + 12
addWorkDays(initDate, someOtherIntegerCalcField)
```

- **Parameter** – Any QuickSight 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
```

Field wells

Group by	Value
employee_id	total_days_worked (Sum)
employment_end_date	bonus_per_day (Sum)
last_day_of_work	total_bonus (Sum)

netWorkDays addWorkDays

Total pro-rated bonus expenditure					
employee...	employment_end_date	last_day_of_work	total_days_worked	bonus_per_day	total_bonus
101	May 6, 2022	Oct 20, 2023	731	55	40,205
102	May 15, 2022	Oct 23, 2023	731	55	40,205
103	Aug 3, 2022	Oct 27, 2023	730	55	40,150
104	Jan 8, 2022	Oct 27, 2023	730	55	40,150
105	Jan 14, 2022	May 21, 2024	731	55	40,205
106	Sep 2, 2022	May 24, 2024	730	55	40,150
107	Aug 9, 2022	May 27, 2024	731	55	40,205
			5,114	385	281,270

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
```

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 QuickSight.
- **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](#), [stdevplf](#), [sumIf](#), [varIf](#), and [varplf](#).

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 QuickSight (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 `isNotNull(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',

/*<i>insert more regions here</i>*/

/*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 `f1_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(f1_date),["2015-05-14", "2015-05-15", "2015-05-16"])
```

Field wells

Group by

fl_date

flightInMidMay

Value

Add measures here

Sheet 1 ▾ +

Fl_date and Flightinmidmay

fl_date	flightInMidMay
May 12, 2015	0
May 13, 2015	0
May 15, 2015	1

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")
```

Field wells

Group by

dest_state_nm

StateCategory

Value

Add measures here

Sheet 1 +

Dest_state_nm and Statecategory

dest_state_nm	StateCategory
Alabama	Other US State
Alaska	Other US State
California	WestCoastUSState
Colorado	Other US State
Connecticut	Other US State
Florida	Other US State
Georgia	Other US State
Hawaii	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 QuickSight calculated field that returns a date value.
- Parameters: Any QuickSight `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)
```

Field wells

netWorkDays | addWorkDays | isWorkDay ▾ +

Filter employees with end date on week day and discern b/w start dates

employee...	employment_start_date	employment_end_date	is_start_date_work_day	is_end_date_work_day
101	Jan 1, 2021	May 6, 2022	1	1
103	Jan 9, 2021	Aug 3, 2022	0	1
105	Aug 3, 2021	Jan 14, 2022	1	1
106	Aug 7, 2021	Sep 2, 2022	0	1
107	Aug 9, 2021	Aug 9, 2022	1	1

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.

```
5
```

netWorkDays

netWorkDays returns the number of working days between the provided two date fields or even custom date values generated using other QuickSight 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 QuickSight calculated field that returns a date value.
- Parameters: Any QuickSight 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 QuickSight calculated field that returns a date value.
- Parameters: Any QuickSight `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)
```

Field wells

Group by

- employee_id
- employment_start_date
- employment_end_date
- salary_paid

Value

- days_worked (Sum)
- salary_per_day (Sum)

netWorkDays v +

Salary paid to each employee per day

employee_id	employment_start_date	employment_end_date	salary_paid	days_worked	salary_per_day
101	Jan 1, 2021	May 6, 2022	120,000	351	341.88
102	Jan 4, 2021	May 15, 2022	110,000	355	309.86
103	Jan 9, 2021	Aug 3, 2022	105,000	408	257.35
104	Jan 10, 2021	Jan 8, 2022	100,000	260	384.62
105	Aug 3, 2021	Jan 14, 2022	135,000	119	1,134.45
106	Aug 7, 2021	Sep 2, 2022	118,000	280	421.43
107	Aug 9, 2021	Aug 9, 2022	100,000	262	381.68
				2,035	3,231.27

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"])
```

Field wells

Group by

fl_date

flightNotInMidMay

Value

Add measures here

Sheet 1 +

FL_date and Flightnotinmidmay

fl_date	flightNotInMidMay
May 15, 2015	0
May 12, 2015	1
May 13, 2015	1

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")
```

Field wells

Group by

dest_state_nm
▼

StateCategory
▼

Value

Add measures here

Sheet 1 ▼ +

Dest_state_nm and Statecategory

dest_state_nm	StateCategory
Alabama	notWestCoastUSState
Alaska	notWestCoastUSState
California	WestCoastUSState
Colorado	notWestCoastUSState
Connecticut	notWestCoastUSState
Florida	notWestCoastUSState
Georgia	notWestCoastUSState
Hawaii	notWestCoastUSState

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.

Amazon QuickSight 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 data sets	All date format patterns specified in Supported date formats .
Direct query of Amazon Aurora, MariaDB, and MySQL databases	<ul style="list-style-type: none"> MM/dd/yyyy dd/MM/yyyy

Date source type	Supported date formats
	<ul style="list-style-type: none"> • 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 • MM-dd-yyyy HH:mm:ss • dd-MM-yyyy HH:mm:ss •

Date source type	Supported date formats
	<p>yyyy-MM-dd HH:mm:ss</p> <ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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
	<p data-bbox="886 212 1284 247">MM-dd-yyyy HH:mm:ss.SSS</p> <ul data-bbox="854 279 1317 695" style="list-style-type: none"><li data-bbox="854 279 1284 338">• dd-MM-yyyy HH:mm:ss.SSS<li data-bbox="854 369 1284 428">• yyyy-MM-dd HH:mm:ss.SSS<li data-bbox="854 459 1312 518">• MMM-dd-yyyy HH:mm:ss.SSS<li data-bbox="854 550 1312 609">• dd-MMM-yyyy HH:mm:ss.SSS<li data-bbox="854 640 1312 699">• yyyy-MMM-dd HH:mm:ss.SSS

Date source type	Supported date formats
Direct query of Microsoft SQL Server databases	<ul style="list-style-type: none"> • 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
	<p>yyyy-MM-dd HH:mm:ss</p> <ul style="list-style-type: none"> • 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
<p>Direct query of Amazon Redshift or PostgreSQL databases</p> <p>Also, datasets from any DBMS that are stored in QuickSight SPICE</p>	<ul style="list-style-type: none"> • 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
	<ul style="list-style-type: none"> 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 QuickSight 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. 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.

Field wells

Group by

- origin_city_name
- arr_delay
- newDelayToSomeCities

Value

Add measures here

Sheet 1 ▾ +

Origin_city_name, Arr_delay, and Newdelaytosomecities

origin_city_name	arr_delay	newDelayToSomeCities
New York, NY	-33	-66
New York, NY	-23	-46
New York, NY	-20	-40
New York, NY	-14	-28
Alexandria, LA	-11	-21
New York, NY	-9	-18
New York, NY	-1	-2
Albany, NY	-20	0
Albany, NY	-5	15
New York, NY	9	18
Albany, NY	14	34
New York, NY	20	40

toLower

toLower formats a string in all lowercase. toLower skips rows containing null values.

Syntax

```
toLower(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.

```
toLower('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 QuickSight.
- **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 QuickSight 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 QuickSight](#) .

The aggregate functions for calculated fields in QuickSight 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 [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 [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}])
```

Avg(sales), Sum of Avg(sales and [Country]) by Region, Country, and Product

Regi...	Country	Product	avg(sales)	avg(sales, [Country])
AMER	Argentina	Big Ol Database	3,299.95	275.11
AMER	Argentina	ChatBot Plugin	67.53	275.11
AMER	Argentina	ContactMatcher	232.22	275.11
AMER	Argentina	Data Smasher	113.73	275.11
AMER	Argentina	FinanceHub	209.86	275.11
AMER	Argentina	Marketing Suite	1,137.94	275.11
AMER	Argentina	Marketing Suite - Gold	518.79	275.11
AMER	Argentina	OneView	323.41	275.11
AMER	Argentina	SaaS Connector Pack	67.93	275.11
AMER	Argentina	SaaS Connector Pack - Gold	38.43	275.11
AMER	Argentina	Site Analytics	397.53	275.11
AMER	Argentina	Storage	10.82	275.11
AMER	Argentina	Support	73.15	275.11
APJ	Australia	Alchemy	1,479.97	162.94
APJ	Australia	Big Ol Database	751.26	162.94
APJ	Australia	ChatBot Plugin	25.14	162.94
APJ	Australia	ContactMatcher	113.16	162.94
API	Australia	Data Smasher	167.75	162.94

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 [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})
```

Field wells

Rows

Order Date (MONTH) ▾

Columns

Add dimensions here

Values Row Column

count({Sales}) (Custom) ▾

Sheet 1 ▾ +

Order Date	count({Sales})
Jan 2019	78
Feb 2019	47
Mar 2019	157
Apr 2019	135
May 2019	122
Jun 2019	135

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 [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}])
```

Count(sales), Sum of Count(sales and [Country]) by Region, Country, and Product

Regi...	Country	Product	count(Sales)	count(Sales, [Country])
AMER	Argentina	Big OL Database	3	130
AMER	Argentina	ChatBot Plugin	11	130
AMER	Argentina	ContactMatcher	17	130
AMER	Argentina	Data Smasher	9	130
AMER	Argentina	FinanceHub	13	130
AMER	Argentina	Marketing Suite	2	130
AMER	Argentina	Marketing Suite - Gold	9	130
AMER	Argentina	OneView	13	130
AMER	Argentina	SaaS Connector Pack	14	130
AMER	Argentina	SaaS Connector Pack - Gold	4	130
AMER	Argentina	Site Analytics	9	130
AMER	Argentina	Storage	5	130
AMER	Argentina	Support	21	130
APJ	Australia	Alchemy	4	492
APJ	Australia	Big OL Database	5	492
APJ	Australia	ChatBot Plugin	37	492
APJ	Australia	ContactMatcher	98	492

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 [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})
```

The screenshot shows the Amazon QuickSight interface. At the top, there are two field wells: 'Group by' with 'Region' selected and 'Value' with 'distinct_count({Order Date}) (Custom)' selected. Below this is a sheet titled 'Sheet 1' containing a table visualization titled 'Distinct_count({order Date}) by Region'. The table has two columns: 'Regi...' and 'distinct_count({Order Date})'. The data rows are: AMER (897), APJ (679), and EMEA (975).

Regi...	distinct_count({Order Date})
AMER	897
APJ	679
EMEA	975

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 [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_count(order Date), Sum of Distinct_count(order Date and [Country] by Region

Regi...	distinct_count(Order Date)	distinct_count(Order Date, [Country])
AMER	897	1,391
APJ	679	932
EMEA	975	1,885

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 [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})
```

Field wells

Group by

Region
▼

Value

Sales (Sum)
▼

max({Sales}) (Custom)
▼

min({Sales}) (Custom)
▼

median({Sales}) (Custom)
▼

Sheet 1 ▼ +

Max({sales}), Sum of Sales, Min({sales}), and Median({sales}) by Region

Region	Sales	max({Sales})	min({Sales})	median({Sales})
AMER	837,849.6	13,999.96	0.99	60.34
APJ	415,464.24	17,499.95	0.44	44.43
EMEA	1,043,887.02	22,638.48	0.85	53.42

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 [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])
```

Max(sales), Sum of Max(sales and [Country]) by Region

Regi...	max(Sales)	max(Sales, [Country])
AMER	13,999.96	51,730.23
APJ	17,499.95	47,417.48
EMEA	22,638.48	128,301.12

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 [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})
```

Field wells

Group by

Region

Value

Sales (Sum)

max({Sales}) (Custom)

min({Sales}) (Custom)

median({Sales}) (Custom)

Sheet 1 +

Max({sales}), Sum of Sales, Min({sales}), and Median({sales}) by Region

Region	Sales	max({Sales})	min({Sales})	median({Sales})
AMER	837,849.6	13,999.96	0.99	60.34
APJ	415,464.24	17,499.95	0.44	44.43
EMEA	1,043,887.02	22,638.48	0.85	53.42

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 [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])
```

Median(sales), Sum of Median(sales and [Country]) by Region

Regi...	median(Sales)	median(Sales, [Country])
AMER	60.34	489.27
APJ	44.43	641.82
EMEA	53.42	3,656.93

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 [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})
```

Field wells

Group by

Region ▼

Value

Sales (Sum) ▼

max({Sales}) (Custom) ▼

min({Sales}) (Custom) ▼

median({Sales}) (Custom) ▼

Sheet 1 +

Max({sales}), Sum of Sales, Min({sales}), and Median({sales}) by Region

Region	Sales	max({Sales})	min({Sales})	median({Sales})
AMER	837,849.6	13,999.96	0.99	60.34
APJ	415,464.24	17,499.95	0.44	44.43
EMEA	1,043,887.02	22,638.48	0.85	53.42

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 [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])
```

Min(sales), Sum of Min(sales and [Country]) by Region

Regi...	min(Sales)	min(Sales, [Country])
AMER	0.99	22.65
APJ	0.44	28.18
EMEA	0.85	1,716.48

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 QuickSight:

- [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 [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 [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])
```

Percentilecont(sales, 30), Sum of Percentilecont(sales, 30 and [Country]) by Region

Regi...	percentileCont(Sales, 30)	percentileCont(Sales, 30, [Country])
AMER	23.84	188.74
APJ	16.77	239.59
EMEA	22.32	2,480.17

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 [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 [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])
```


Percentile(sales, 30), Sum of Percentile(sales, 30 and [Country]) by Region

Regi...	percentile(Sales, 30)	percentile(Sales, 30, [Country])
AMER	23.84	187.64
APJ	16.74	237.73
EMEA	22.32	2,472.33

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. QuickSight 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"))
```

Field wells

Group by: payment_type

Value: PTDAverage (Custom)

Sheet 1

Ptdaverage by Payment_type

payment_ty...	PTDAverage
1	13.74

$(A+B+C)/3$

Average of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	14.33
Jun 29, 2021	1	C 13.93
Jun 28, 2021	1	B 13.44
Jun 27, 2021	1	A 13.85
Jun 26, 2021	1	14.81

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. QuickSight 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"))
```

Field wells

Group by

payment_type

Value

PTDCount (Custom)

Sheet 1



Ptdcount by Payment_type

payment_ty...	PTDCount
1	701,602

A+B+C

Count of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	240,926
Jun 29, 2021	1	C 240,644
Jun 28, 2021	1	B 230,231
Jun 27, 2021	1	A 230,727
Jun 26, 2021	1	209,895

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. QuickSight 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"))
```

Field wells

Group by: payment_type

Value: PTDMax (Custom)

Sheet 1 +

Ptdmax by Payment_type

payment_ty...	PTDMax
1	490

Max of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	500
Jun 29, 2021	1	400
Jun 28, 2021	1	320
Jun 27, 2021	1	490
Jun 26, 2021	1	500

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. QuickSight 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"))
```

Field wells

Group by: **payment_type**

Value: **PTDMedian (Custom)**

Sheet 1 +

Ptdmedian by Payment_type

payment_ty...	PTDMedian
1	10

Median of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	10.50
Jun 29, 2021	1	10.00
Jun 28, 2021	1	10.00
Jun 27, 2021	1	10.00
Jun 26, 2021	1	10.50

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. QuickSight 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"))
```

Field wells

Group by: payment_type

Value: PTDMin (Custom)

Sheet 1 +

Ptdmin by Payment_type

payment_ty...	PTDMin
1	-7

Min of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	-52
Jun 29, 2021	1	0
Jun 28, 2021	1	0
Jun 27, 2021	1	-7
Jun 26, 2021	1	0

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. QuickSight 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"))
```

Field wells

Group by
payment_type

Value
PTDPercentile (Custom)

Sheet 1 +

Ptdpercentile by Payment_type

payment_ty...	PTDPercentile
1	27

P90 of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	29.00
Jun 29, 2021	1	27.50
Jun 28, 2021	1	26.00
Jun 27, 2021	1	28.50
Jun 26, 2021	1	30.00

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. QuickSight 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"))
```

Field wells

Group by: payment_type

Value: PTDPercntileCont (Custom)

Sheet 1 +

Ptdpercentilecont by Payment_type

payment_ty...	PTDPercntileCont
1	27

P90 of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	29.00
Jun 29, 2021	1	27.50
Jun 28, 2021	1	26.00
Jun 27, 2021	1	28.50
Jun 26, 2021	1	30.00

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. QuickSight 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"))
```

Field wells

Group by: payment_type

Value: PTDStDev (Custom)

Sheet 1

Ptdstdev by Payment_type

payment_ty...	PTDStDev
1	11.78

Standard deviation of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	12.26
Jun 29, 2021	1	11.65
Jun 28, 2021	1	11.45
Jun 27, 2021	1	12.21
Jun 26, 2021	1	12.87

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. QuickSight 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"))
```

Field wells

Group by: payment_type

Value: PTDStDevP (Custom)

Sheet 1 +

Ptdstdevp by Payment_type

payment_ty...	PTDStDevP
1	11.78

Standard deviation - population of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	12.26
Jun 29, 2021	1	11.65
Jun 28, 2021	1	11.45
Jun 27, 2021	1	12.21
Jun 26, 2021	1	12.87

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. QuickSight 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"))
```

Field wells

Group by: payment_type

Value: PTDSum (Custom)

Sheet 1

Ptdsum by Payment_type

payment_ty...	PTDSum
1	9,642,352.38

A+B+C

Sum of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	3,452,526.24
Jun 29, 2021	1	C 3,352,298.42
Jun 28, 2021	1	B 3,095,150.7
Jun 27, 2021	1	A 3,194,903.26
Jun 26, 2021	1	3,108,392.02

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. QuickSight 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"))
```

Field wells

Group by

payment_type

Value

PTDVar (Custom)

Sheet 1



Ptdvar by Payment_type

payment_ty...	PTDVar
1	138.65

Variance of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	150.22
Jun 29, 2021	1	135.78
Jun 28, 2021	1	131.13
Jun 27, 2021	1	149.02
Jun 26, 2021	1	165.76

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. QuickSight 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"))
```

Field wells

Group by: payment_type

Value: PTDVarP (Custom)

Sheet 1 +

Ptdvarp by Payment_type

payment_ty...	PTDVarP
1	138.65

Variance - population of Fare_amount by Pickupdatetime and Payment_type

pickupDatetime	payment_type	fare_amount
Jun 30, 2021	1	150.22
Jun 29, 2021	1	135.78
Jun 28, 2021	1	131.13
Jun 27, 2021	1	149.02
Jun 26, 2021	1	165.76

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 [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 [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 [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 [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 [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 [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])
```

Sum(sales), Sum of Sum(sales and [Country]) by Region, Product, and Country

Regi...	Country	Product	sum(Sales)	sum(Sales, [Country])
AMER	Argentina	Big Ol Database	9,899.85	35,764.31
AMER	Argentina	ChatBot Plugin	742.8	35,764.31
AMER	Argentina	ContactMatcher	3,947.81	35,764.31
AMER	Argentina	Data Smasher	1,023.56	35,764.31
AMER	Argentina	FinanceHub	2,728.24	35,764.31
AMER	Argentina	Marketing Suite	2,275.88	35,764.31
AMER	Argentina	Marketing Suite - Gold	4,669.08	35,764.31
AMER	Argentina	OneView	4,204.36	35,764.31
AMER	Argentina	SaaS Connector Pack	950.97	35,764.31
AMER	Argentina	SaaS Connector Pack - Gold	153.7	35,764.31
AMER	Argentina	Site Analytics	3,577.75	35,764.31
AMER	Argentina	Storage	54.12	35,764.31
AMER	Argentina	Support	1,536.19	35,764.31
APJ	Australia	Alchemy	5,919.89	80,166.1
APJ	Australia	Big Ol Database	3,756.31	80,166.1
APJ	Australia	ChatBot Plugin	930.03	80,166.1

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 screenshot shows the 'Field wells' interface. On the left, the 'Group by' section has 'Segment' selected. On the right, the 'Value' section has 'Sales (Sum)' and 'sumif (Custom)' selected. Below the field wells, a table titled 'Sum of Sales and Sumif by Segment' is displayed. The table has three columns: 'Segment', 'Sales', and 'sumif'. The rows are: Enterprise (Sales: 429,653), SMB (Sales: 1,161,401, sumif: 1,161,401.35), and Strategic (Sales: 706,146).

Segment	Sales	sumif
Enterprise	429,653	
SMB	1,161,401	1,161,401.35
Strategic	706,146	

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 [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 [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 [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 [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 QuickSight.

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.

Field wells

Sum of Billed Amount and Difference by Customer Region		
Customer Region	difference	Billed Amount
APAC		a 8,390,654
EMEA	+b 2,647,510	=c 11,038,164
US	12,509,001	23,547,165

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 QuickSight](#).

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
(  
  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.

Field wells

Group by
Order Date

Value
Sales (Sum)
lag1 (Custom)

Sheet 1

Sum of Sales and Lag1 by Order Date

Order Date	Sales	lag1
Jan 4, 2020	16	
Jan 5, 2020	288	16
Jan 6, 2020	20	288
Jan 7, 2020	4,407	20
Jan 8, 2020	87	4,407
Jan 10, 2020	41	87
Jan 11, 2020	55	41
Jan 12, 2020	10	55

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.

Field wells

Group by

Order Date



Segment



Value

Sales (Sum)

lag2 (Custom)

Sheet 1

Sheet 2



Sum of Sales and Lag2 by Order Date and Segment

Order Date	Segment	Sales	lag2
Jan 4, 2020	SMB	16	
Jan 5, 2020	Enterprise	288	
Jan 6, 2020	SMB	20	16
Jan 7, 2020	Enterprise	4,375	288
Jan 7, 2020	SMB	19	20
Jan 7, 2020	Strategic	13	

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.

Field wells

The screenshot shows two field wells. The 'Group by' well has a dropdown menu with 'Customer Segment' selected. The 'Value' well has two items: 'Billed Amount (Sum)' and 'lead (Custom)'. Both wells are highlighted with a green border.

Lead and Sum of Billed Amount by Customer Segment

Customer Segment	Billed Amount	lead
Enterprise	14,643,518	3,857,503
SMB	3,857,503	24,474,962
Startup	24,474,962	

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.

Field wells

The screenshot shows two field wells at the top. The 'Group by' well has a dropdown menu with 'Customer Region' selected. The 'Value' well has two dropdown menus: 'Billed Amount (Sum)' and 'percentDifference (Custom)'. Below the wells is a table visualization titled 'Percentdifference and Sum of Billed Amount by Customer Region'. The table has three columns: 'Customer Region', 'Billed Amount', and 'percentDifference'. The rows are APAC, EMEA, and US. The APAC row has 'a' next to '8,390,654' and 'c' next to '-24.0%'. The EMEA row has 'b' next to '11,038,164' and '-53.1%'. The US row has '23,547,165' and the formula '(b-a) / b = c'.

Customer Region	Billed Amount	percentDifference
APAC	a 8,390,654	c -24.0%
EMEA	b 11,038,164	-53.1%
US	23,547,165	(b-a) / b = c

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.

Field wells

Group by

Customer Region

Service Line

Value

Billed Amount (Sum)

avgOver (Custom)

Sum of Billed Amount and Avgover by Customer Region and Service Line

Customer Region	Service Line	Billed Amount		avgOver
APAC	Billing	3,569,780		2,796,885
APAC	HR	3,441,106	=	2,796,885
APAC	Marketing	1,379,768		2,796,885
EMEA	Billing	4,446,586		3,679,388
EMEA	HR	4,316,700		3,679,388
EMEA	Marketing	2,274,878		3,679,388

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 Amazon QuickSight](#).

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 Amazon QuickSight](#).

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]  
)
```

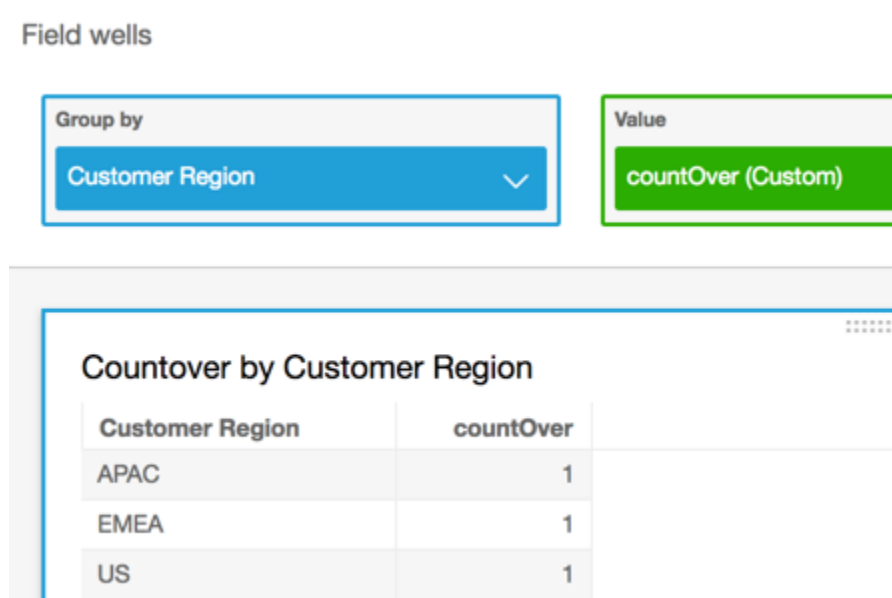
The following example shows the count of Billed Amount over Customer Region. The fields in the table calculation are in the field wells of the visual.

```
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.

Field wells



The screenshot displays the configuration for a calculated field in Amazon QuickSight. The 'Field wells' section shows the 'Group by' field set to 'Customer Region' and the 'Value' field set to 'countOver (Custom)'. Below this, a table titled 'Countover by Customer Region' displays the results of the calculation.

Customer Region	countOver
APAC	1
EMEA	1
US	1

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.

Field wells

Group by

Customer Segment ▼

Service Line ▼

Customer Region ▼

Value

countOver (Custom) ▼

.....

Countover by Customer Segment, Service Line, and Customer Region

Customer Segment	Service Line	Customer Region	countOver
Enterprise	Billing	APAC	9
Enterprise	Billing	EMEA	9
Enterprise	Billing	US	9
Enterprise	HR	APAC	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.

Field wells

Group by

Customer Segment ▼

Service Line ▼

Customer Region ▼

Value

countOver (Custom) ▼

.....

Countover by Customer Region, Service Line, and Customer Segment

Customer Segment	Service Line	Customer Region	countOver
Enterprise	Billing	APAC	1
Enterprise	Billing	EMEA	1
Enterprise	Billing	US	1
Enterprise	HR	APAC	1
Enterprise	HR	EMEA	1
Enterprise	HR	US	1

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 Amazon QuickSight](#).

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.

Field wells



Sum of Billed Amount and Maxover by Customer Region and Service Line

Customer Region	Service Line	Billed Amount	maxOver
APAC	Billing	3,569,780	3,569,780
APAC	HR	3,441,106	= 3,569,780
APAC	Marketing	1,379,768	3,569,780
EMEA	Billing	4,446,586	4,446,586
EMEA	HR	4,316,700	4,446,586
EMEA	Marketing	2,274,878	4,446,586

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 Amazon QuickSight](#).

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.

Field wells



Customer Region	Service Line	Billed Amount	minOver
APAC	Billing	3,569,780	1,379,768
APAC	HR	3,441,106	= 1,379,768
APAC	Marketing	1,379,768	1,379,768
EMEA	Billing	4,446,586	2,274,878
EMEA	HR	4,316,700	2,274,878
EMEA	Marketing	2,274,878	2,274,878

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 QuickSight:

- [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

```
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, 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 QuickSight](#) and [Using level-aware calculations in Amazon QuickSight](#).

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 QuickSight](#) and [Using level-aware calculations in Amazon QuickSight](#).

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.

Customer Region	Billed Amount	percOver50	percOver98
APAC	\$598,114	\$850,406	\$1,648,034
EMEA	\$850,406	\$850,406	\$1,648,034
US	\$1,648,034	\$850,406	\$1,648,034

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.

Field wells

The screenshot shows two field wells. The 'Group by' well is on the left and contains two items: 'Service Line' and 'Customer Region'. The 'Value' well is on the right and contains two items: 'percentOfTotal (Custom)' and 'Billed Amount (Sum)'. Both wells have a dropdown arrow on the right side of each item.

Percentoftotal and Sum of Billed Amount by Service Line and Customer Region

Service Line	Customer Region	percentOfTotal	Billed Amount
Billing	APAC	20.6%	3,569,779.71
Billing	EMEA	25.6%	4,446,586.13
Billing	US	53.8%	9,330,832.51
HR	APAC	20.0%	3,441,106.16
HR	EMEA	25.1%	4,316,700.48
HR	US	55.0%	9,464,168.33

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 defaults 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 screenshot shows the Amazon QuickSight interface. On the left, the 'Group by' dropdown is set to 'Order Date'. On the right, the 'Value' section contains two fields: 'Sales (Sum)' and 'PeriodOverPeriod (Custom)'. Below this, the calculated field 'PeriodOverPeriodDifference' is selected. The resulting visualization is a table titled 'Sum of Sales and Periodoverperiod by Order Date'.

Order Date	Sales	PeriodOverPeriod
Jan 4, 2020	16	
Jan 5, 2020	288	272
Jan 6, 2020	20	-269
Jan 7, 2020	4,407	4,388
Jan 8, 2020	87	-4,320

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)
```

Group by: Order Date (MONTH)

Value: Sales (Sum), PeriodOverPeriod2 (Custom)

Calculation bar: Lag, Sumif, PeriodOverPeriodDifference

Sum of Sales and Periodoverperiod2 by Order Date

Order ...	Sales	PeriodOverPeriod2
Jan 2020	13,946	
Feb 2020	4,811	
Mar 2020	55,691	41,745
Apr 2020	28,295	23,485
May 2020	23,648	-32,043

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)
```

Field wells

Group by

Order Date (MONTH) ▾

Value

MoMLastValue (Custom) ▾

Sales (Sum) ▾

PeriodOverPeriod | PTDOverTime / PTD Agg | **PeriodOverPeriod** ▾ +

Sum of Sales and Momlastvalue by Order Date

Order Date	MoMLastValue	Sales
Jan 2015		274,766.92
Feb 2015	274,766.92	326,101.47
Mar 2015	326,101.47	271,696.67
Apr 2015	271,696.67	389,831.95
May 2015	389,831.95	306,572.07
Jun 2015	306,572.07	355,368.8

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`, `HOURL`, `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)
```

Field wells
edit, replace, and remove datasets.

Group by
Order Date (MONTH) ✓

Value
MoMPercentDifference (Custom) ✓
Sales (Sum) ✓

PeriodOverPeriod | PTDOverTime / PTD Agg | **PeriodOverPeriod** ✓ +

Sum of Sales and Mompercentdifference by Order Date

Order Date	MoMPercentDifference	Sales
Jan 2015		274,766.92
Feb 2015	18.68%	326,101.47
Mar 2015	-16.68%	271,696.67
Apr 2015	43.48%	389,831.95
May 2015	-21.36%	306,572.07
Jun 2015	15.92%	355,368.8

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)
```


Field wells

Group by	Value
pickupDatetime (WEEK) ▼	fare_amount (Sum) ▼
	PTDOverTimeAvg (Custom) ▼

Sheet 1 ▼ +

Ptdovertimeavg and Sum of Fare_amount by Pickupdatetime

pickupDatetime	fare_amount	PTDOverTimeAvg	
Jun 27, 2021	D 20276346.25	31143423.01	$(A+B+C+D)/4$
Jun 20, 2021	C 34294039.05	34765781.93	$(A+B+C)/3$
Jun 13, 2021	B 34311811.14	35001653.37	$(A+B)/2$
Jun 6, 2021	A 35691495.60	35691495.60	$A/1$
May 30, 2021	33927942.61	35770454.41	
May 23, 2021	32781949.16	36231082.36	
May 16, 2021	38225816.10	37380793.42	

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)
```

Field wells

Group by

- pickupDatetime (WEEK) ▼
- vendorid ▼

Value

- PTDOverTimeCount (Custom) ▼

Sheet 1 ▼ +

Ptdovertimecount by Pickupdatetime and Vendorid

pickupDatetime	vendorid	PTDOverTimeCount
Jun 27, 2021	1	4
Jun 27, 2021	2	4
Jun 20, 2021	1	3
Jun 20, 2021	2	3
Jun 13, 2021	1	2
Jun 13, 2021	2	2

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)
```

Field wells

Group by

pickupDatetime (WEEK) ▾

Value

fare_amount (Max) ▾

PTDOverTimeMax (Custom) ▾

Sheet 1 ▾ +

Max of Fare_amount and Ptdovertimemax by Pickupdatetime

pickupDatetime	fare_amount	PTDOverTimeMax
Jun 27, 2021	628544.74	628,544.74
Jun 20, 2021	8007.00	187,440.96
Jun 13, 2021	8452.00	187,440.96
Jun 6, 2021	187440.96	187,440.96
May 30, 2021	133057.84	133,057.84
May 23, 2021	4886.00	8,007
May 16, 2021	1520.40	8,007

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)
```

Field wells

Group by

pickupDatetime (WEEK) ▾

Value

fare_amount (Min) ▾

PTDOverTimeMin (Custom) ▾

Sheet 1 ▾ +

Min of Fare_amount and Ptdovertimemin by Pickupdatetime

pickupDatetime	fare_amount	PTDOverTimeMin
Jun 27, 2021	-250.00	-450
Jun 20, 2021	-450.00	-450
Jun 13, 2021	-273.00	-287
Jun 6, 2021	-287.00	-287
May 30, 2021	-199.00	-410
May 23, 2021	-300.00	-410

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)
```

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)
```

Field wells

Group by

pickupDatetime (WEEK) ▼

Value

fare_amount (Sum) ▼

PTDOverTimeSum (Custom) ▼

Sheet 1 +

Sum of Fare_amount and Ptdovertimesum by Pickupdatetime

pickupDatetime	fare_amount	PTDOverTimeSum	
Jun 27, 2021	D 20,276,346.25	124,573,692.04	A+B+C+D
Jun 20, 2021	C 34,294,039.05	104,297,345.79	A+B+C
Jun 13, 2021	B 34,311,811.14	70,003,306.74	A+B
Jun 6, 2021	A 35,691,495.6	35,691,495.6	A
May 30, 2021	33,927,942.61	178,852,272.03	
May 23, 2021	32,781,949.16	144,924,329.42	
May 16, 2021	38,225,816.1	112,142,380.26	
May 9, 2021	36,938,239.42	73,916,564.16	

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 Amazon QuickSight](#).

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 Amazon QuickSight](#).

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 Amazon QuickSight](#).

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 Amazon QuickSight](#).

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 Amazon QuickSight](#).

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.

Field wells

The screenshot shows two field wells. The 'Group by' well contains two items: 'Customer Region' and 'Customer Segment'. The 'Value' well contains two items: 'Billed Amount (Sum)' and 'sumOver (Custom)'. Both wells have a dropdown arrow on the right side of each item.

Sum of Billed Amount and Sumover by Customer Region and Customer Segment

Customer Region	Customer Segment	Billed Amount	sumOver
APAC	Enterprise	2,035,949	8,390,654
APAC	SMB	660,097	= 8,390,654
APAC	Startup	5,694,609	8,390,654
EMEA	Enterprise	5,678,783	11,038,164
EMEA	SMB	1,341,834	11,038,164
EMEA	Startup	4,017,547	11,038,164

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 Amazon QuickSight](#).

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 Amazon QuickSight](#).

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.

Field wells

The screenshot shows two field wells. The 'Group by' well has a dropdown menu with 'Customer Region' selected. The 'Value' well has two items: 'rank (Custom)' and 'Billed Amount (Sum)', both with dropdown arrows.

Rank and Sum of Billed Amount by Customer Region

Customer Region	rank	Billed Amount
APAC	3	8,390,654
EMEA	2	11,038,164
US	1	23,547,165

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 Amazon QuickSight](#).

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.

The screenshot displays the configuration for a table calculation. In the 'Field wells' section, the 'Group by' well contains 'Customer Region'. The 'Value' well contains two fields: 'Billed Amount (Sum)' and 'Percentile (Custom)'. Below the field wells, a table visualization titled 'Sum of Billed Amount and Percentile by Customer Region' is shown. The table has three columns: 'Customer Region', 'Billed Amount', and 'Percentile'. The data rows are as follows:

Customer Region	Billed Amount	Percentile
APAC	8,390,654.34	66.6666
EMEA	11,038,164.3	33.3333
US	23,547,164.89	0

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.

Field wells Rows Date Columns Values # Billed Amount (Sum) = runningSum (Custom)

Sum of Billed Amount and Runningsum by Date

Date	Billed Amount	runningSum
Jan 2012	54,675.45	a 54,675.45
Feb 2012	+b 57,127.93	=c 111,803.38
Mar 2012	66,303.97	178,107.35
Apr 2012	66,694.23	244,801.58
May 2012	75,906.62	a 320,708.2
Jun 2012	+b 83,531.67	=c 404,239.87

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},  
    truncDate('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 Amazon QuickSight start page, choose **Datasets**.
2. 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 QuickSight to join objects from one or more data sources. By using Amazon QuickSight 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 QuickSight *logical tables*, where each logical table contains information about how to fetch data. When editing a dataset in QuickSight, 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 QuickSight: 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 QuickSight data source:
 - All of the logical tables in this dataset must refer to the same QuickSight data source. This doesn't apply if two separate QuickSight data sources refer to the same underlying database.

It must be the exact same QuickSight 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 QuickSight dataset that is a parent dataset:
 - The parent dataset must use direct query.
 - The parent dataset must refer to the same QuickSight 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?

Same-source datasets have no restrictions on data size.

With cross-source datasets QuickSight automatically detects the largest logical table in the dataset, and the total size of the other logical tables must sum up to less than 1 GB. The logical size of these tables is recomputed each time your dataset is refreshed into SPICE.

Note that the logical size as computed during a join is separate from the size logic used by SPICE. A SPICE dataset with a size of 1.5 GB could have a logical size of 1 GB.

Additionally, because cross-source datasets must use SPICE, they must adhere to the quotas and limits for SPICE. For more information, see [SPICE quotas for imported data](#).

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 QuickSight](#) 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 QuickSight, 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 ore more manually uploaded file, your dataset is automatically stored in SPICE.

If you choose **SPICE**, the data is ingested into QuickSight. 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.

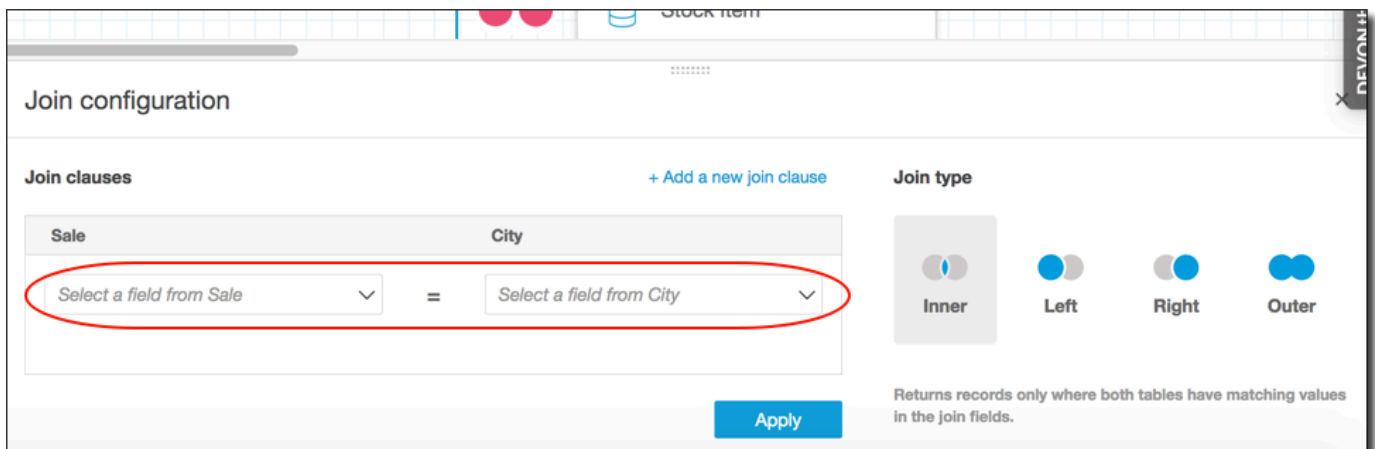
An icon with red dots appears to indicate that you need to configure this join. Two red dots appear for joins that aren't yet configured. To create joins, choose the first join configuration icon.




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.

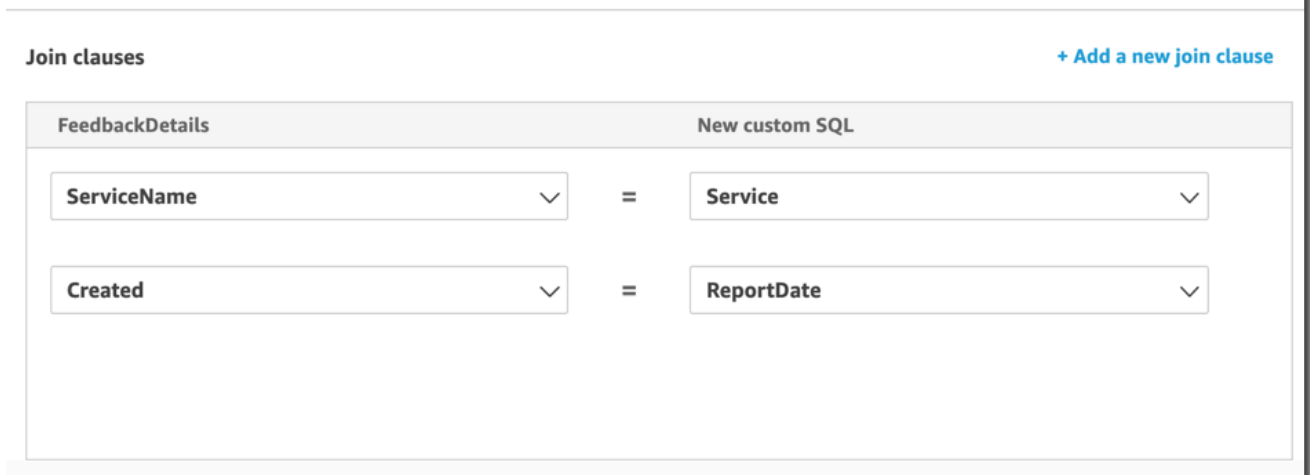


- In the **Join clauses** section, choose the join column for each table.




(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.

Join configuration



- 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 QuickSight](#).

Join types

Amazon QuickSight 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 QuickSight

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 QuickSight](#).

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 QuickSight 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 QuickSight](#)
- [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. From the QuickSight start page, choose **Datasets**.

The screenshot shows the Amazon QuickSight interface. At the top, there is a search bar with the text "2018-2019 Platfor..." and a prompt "Type a question about your data". Below the search bar, there is a navigation menu on the left with the following items: "Find analyses & more" (with a magnifying glass icon), "Favorites" (with a star icon), "Recent" (with a clock icon), "My folders" (with a folder icon), "Shared folders" (with a folder icon), "Dashboards" (with a bar chart icon), "Analyses" (with a bar chart icon and highlighted in light blue), "Datasets" (with a database icon and highlighted with a red box), and "Topics" (with a speech bubble icon). The main content area is titled "Analyses" and displays a grid of analysis cards. The top card is titled "Business Review analysis" and shows a value of "8.03% ↑" and "Updated 2 hours ago". The bottom card is titled "Sales Pipeline analysis" and shows a "SAMPLE" label. To the right of the main content area, there are partial views of other analysis cards.

2. Choose the dataset that you want, and then choose **Edit dataset**.
3. 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.

Augment with SageMaker

Search fields

Focus

All fields

Select All | None

- Customer
- Date
- Customer Region
State
- Segment-1
- Service Line
- Revenue Goal
- Billed Amount
- Cost
- Channel
- Distinct ID

Excluded fields No fields excluded

Filters No filters applied [Add filter](#)

Query mode [Refresh now](#)

Business Review

Date	Customer ID	Customer ...	Customer ...
<input checked="" type="checkbox"/> Date	<input type="checkbox"/> String	<input type="checkbox"/> String	<input checked="" type="checkbox"/> State
2012-01-01...	DXegKx8qH...	SMB10	APAC
2012-01-01...	DXegKx8qH...	SMB10	APAC
2012-01-01...	A28Dzrr5dn...	SMB64	APAC
2012-01-01...	A28Dzrr5dn...	SMB64	APAC
2012-01-01...	A28Dzrr5dn...	SMB64	APAC
2012-01-01...	mbaEj8eHB...	SMB55	APAC
2012-01-01...	mbaEj8eHB...	SMB55	APAC

Viewing filters in analyses

Use the following procedure to view filters in analyses.

To view a filter in an analysis

1. From the QuickSight start page, choose **Analyses**.

2018-2019 Platfor... Type

Find analyses & more

★ Favorites

🕒 Recent

📁 My folders

📁 Shared folders

📊 Dashboards

📊 Analyses

🗄️ Datasets

🗨️ Topics

Analyses

Analysis

8.03% ↑

Business Review analysis

Updated 2 hours ago

Analysis

web-and-social-analytics 1....

Updated 3 months ago

2. On the **Analyses** page, choose the analysis that you want to work with.
3. In the analysis, choose **Filter** at left.

Any filters applied to the analysis appear here.

You can see how a filter is scoped in an analysis. To do so, choose the filter in the pane and view the **Applied to** section at the top of the pane. 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. From the QuickSight start page, choose **Datasets**.

The screenshot shows the Amazon QuickSight interface. At the top, there is a search bar with the text "2018-2019 Platfor..." and a prompt "Type a question about your data". Below the search bar, there is a navigation sidebar on the left with the following options: "Find analyses & more" (with a magnifying glass icon), "Favorites" (with a star icon), "Recent" (with a clock icon), "My folders" (with a folder icon), "Shared folders" (with a folder icon), "Dashboards" (with a bar chart icon), "Analyses" (with a bar chart icon and highlighted in light blue), "Datasets" (with a database icon and highlighted with a red box), and "Topics" (with a speech bubble icon). The main content area is titled "Analyses" and displays a grid of analysis cards. The top card is titled "Business Review analysis" and shows a value of "8.03% ↑". The bottom card is titled "Sales Pipeline analysis" and has a "SAMPLE" label. Each card includes a star icon and a three-dot menu icon.

2. Choose the dataset that you want, and then choose **Edit dataset**.
3. On the data preparation page that opens, choose **Add filter** at lower left, and then choose a field that you want to filter.

The screenshot displays the Amazon QuickSight interface. At the top, there's a header with 'Web and Social Media Analytics' and '2018-2019 Platfor...'. Below this, the 'Fields' section is visible, containing 'All fields included', 'Add calculated field', 'Augment with SageMaker', and a search bar for fields. Underneath, there's a 'Focus' dropdown set to 'All fields' and a 'Select' section with 'All' and 'None' options. A list of fields is shown, including 'populated_event', 'Date', 'New visitors SEO', 'New visitors CPC', 'New visitors Social Media', and 'Return visitors'. Below the fields, there's an 'Excluded fields' section with 'No fields excluded'. At the bottom of the fields section, there's a 'Filters' section with 'No filters applied' and an 'Add filter' button highlighted with a red box. To the right, the 'Data' section shows a 'Tutor' dataset.

The filter is added to the **Filters** section at left.

4. 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 QuickSight](#).

5. When finished, choose **Apply**.

Note

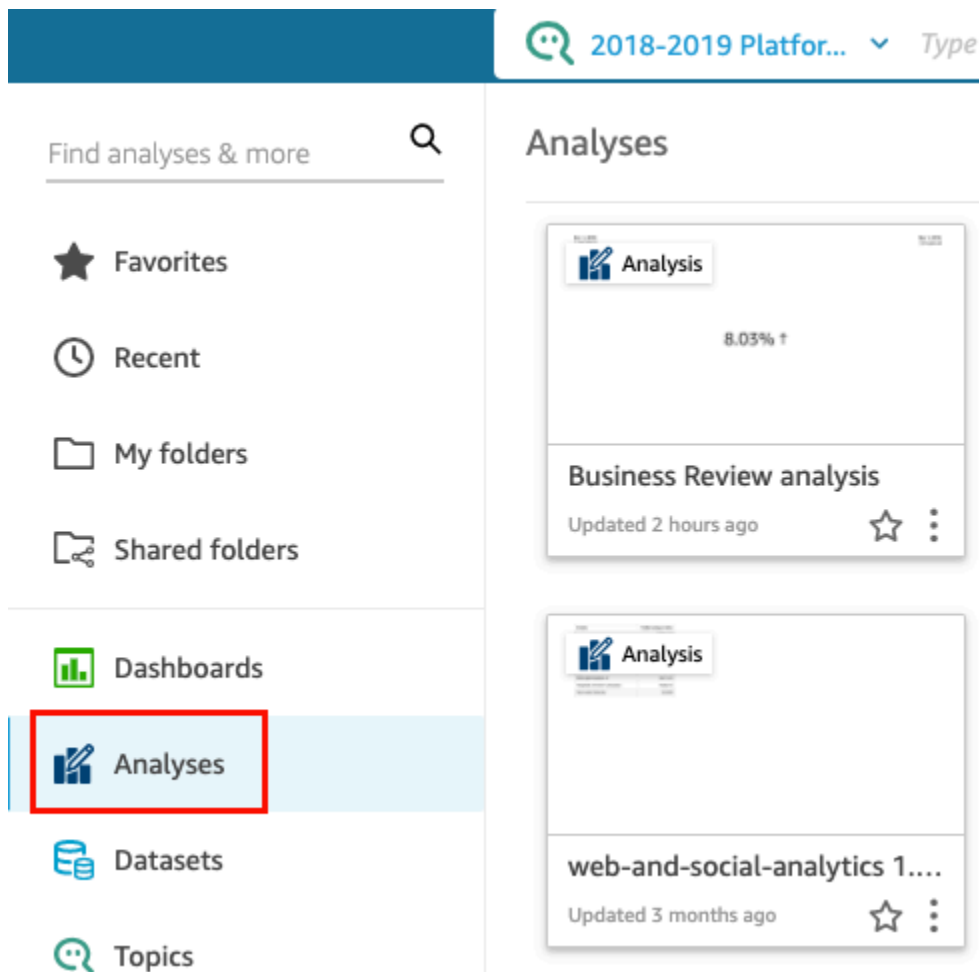
The data preview shows you the results of your combined filters only as they apply to the first 1,000 rows. If all of the first 1,000 rows are filtered out, then no rows show in the preview. This effect occurs even when rows after the first 1,000 aren't filtered out.

Adding filters in analyses

Use the following procedure to add filters to analyses.

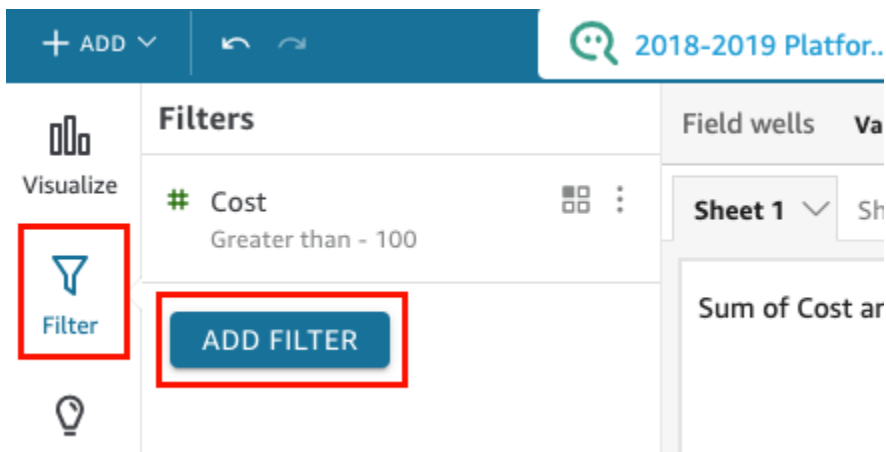
To add a filter to an analysis

1. From the QuickSight start page, choose **Analyses**.

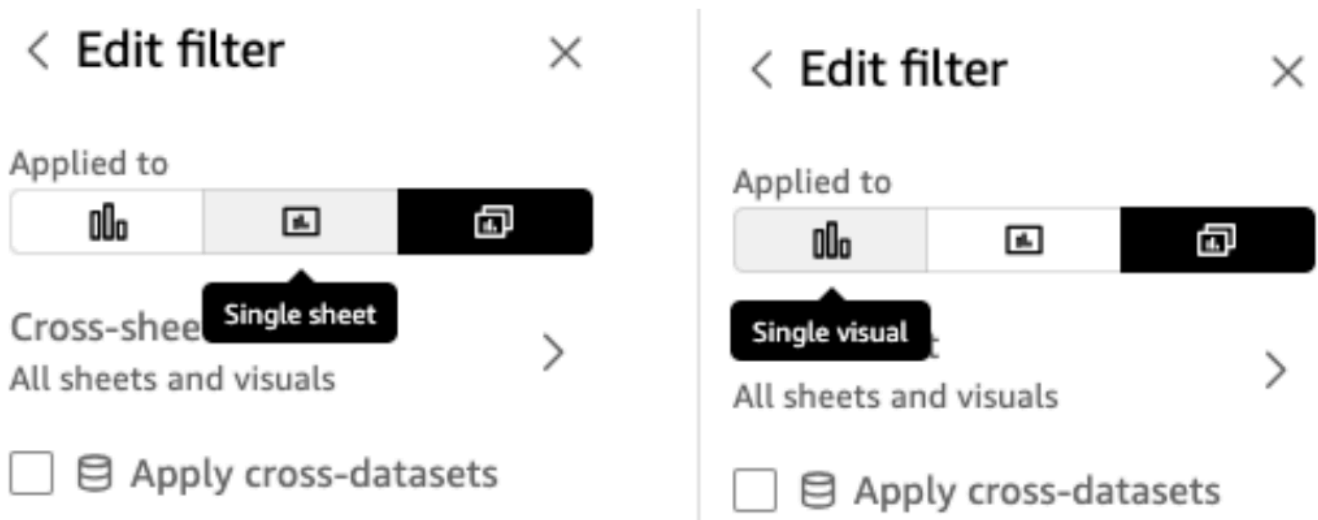


2. On the **Analyses** page, choose the analysis that you want to work with.

3. In the analysis, choose **Filter** at left, choose **Add filter**, and then choose a field that you want to filter.



4. 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**.
5. 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 QuickSight](#).

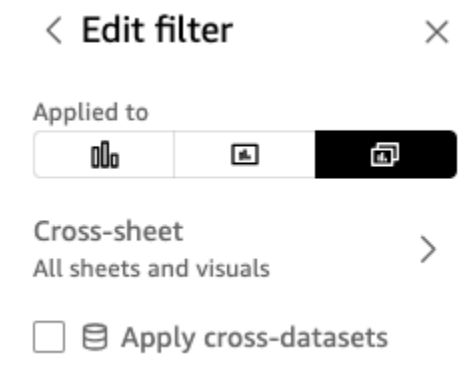
Cross-sheet filters and controls

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.

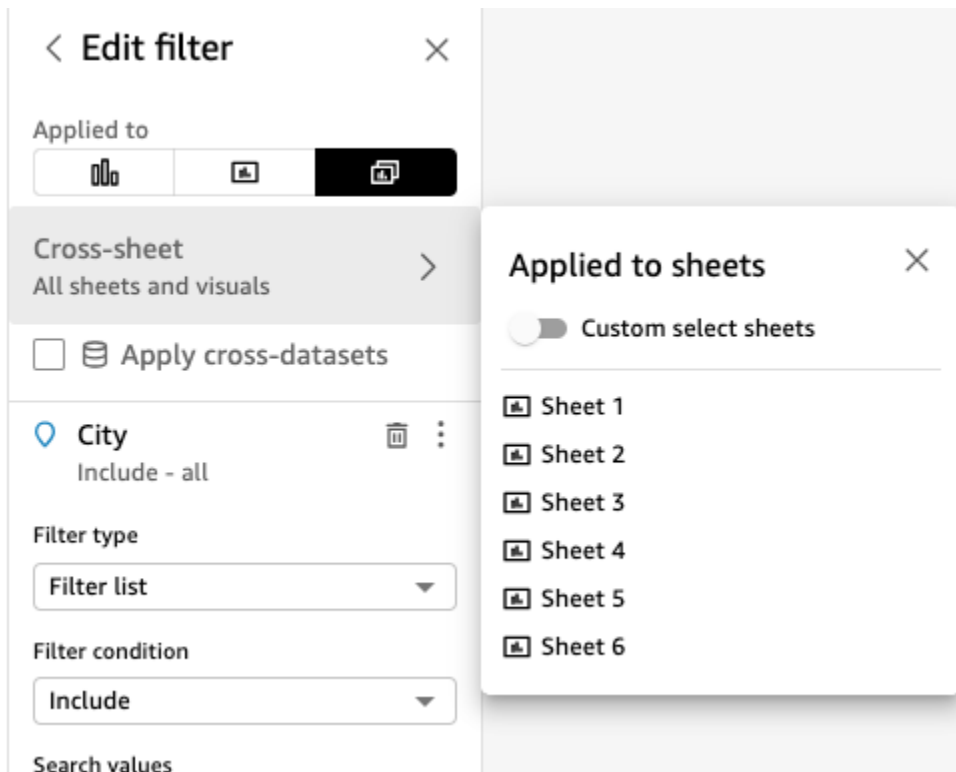
Filters

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 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.

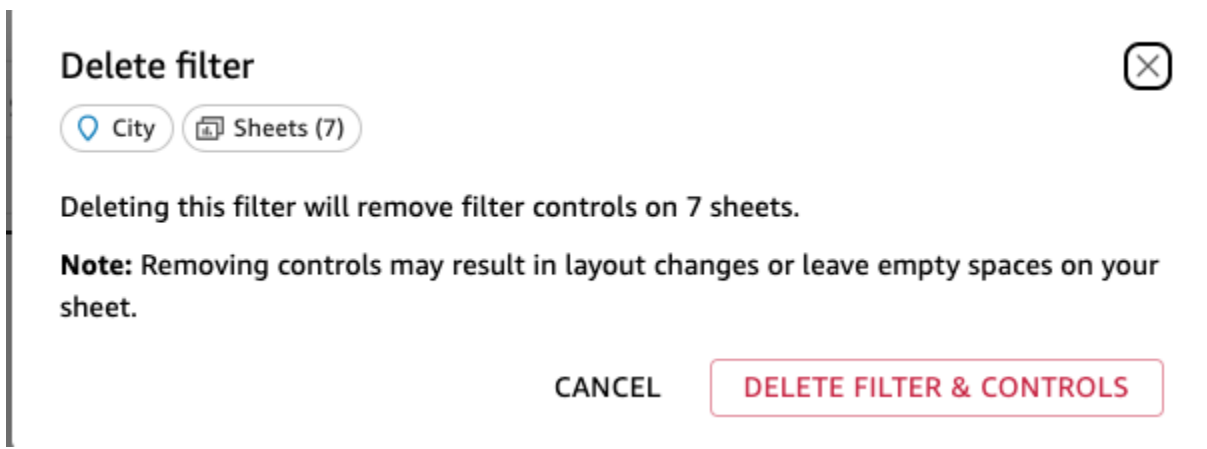
Removing a Cross-Sheet Filter

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. You will see the following modal when you choose **Delete**:

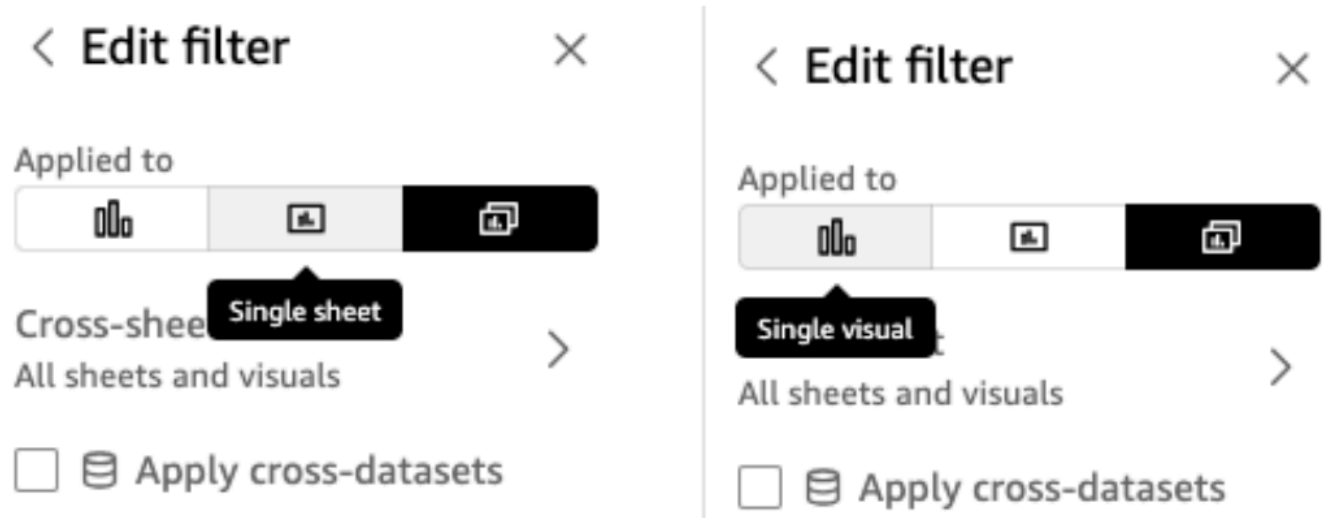


3. If you choose **Delete Filter and Controls**, the controls will be deleted from all pages. This may impact the layout of your analysis. Alternatively, you can remove these controls individually.

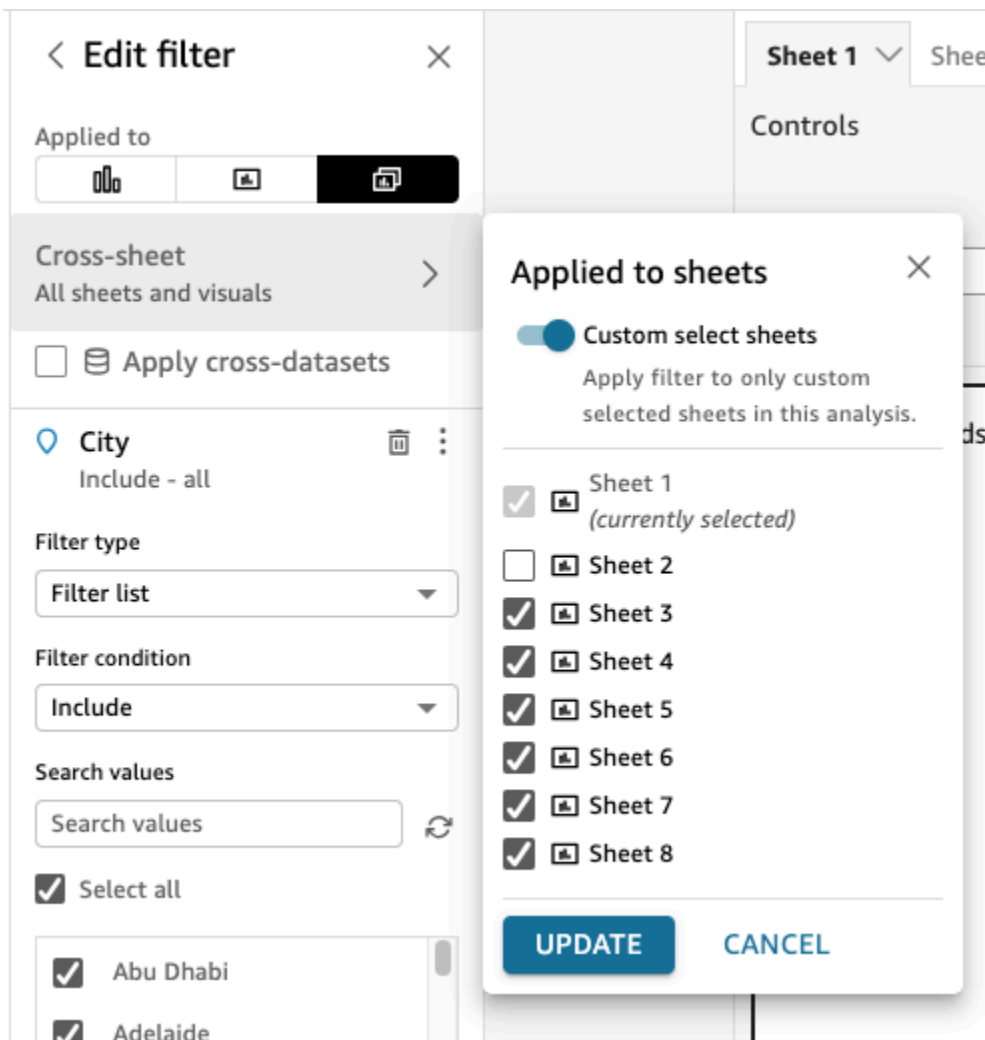
Downscoping

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 get the following 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.

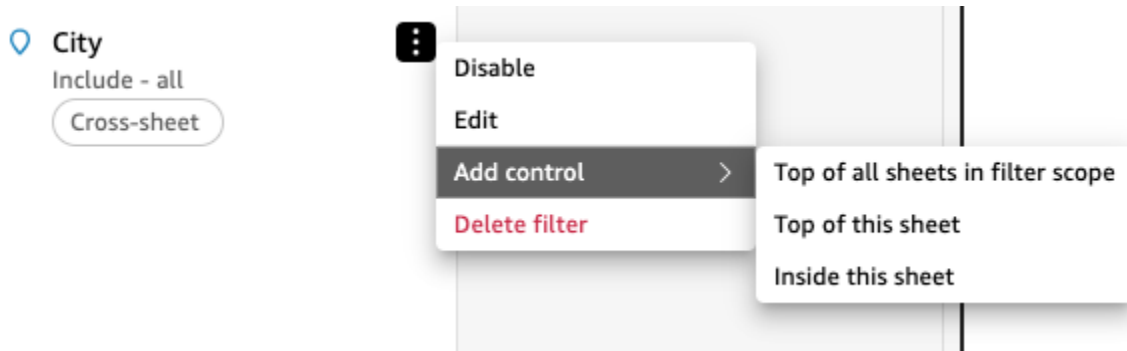
Controls

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](#).

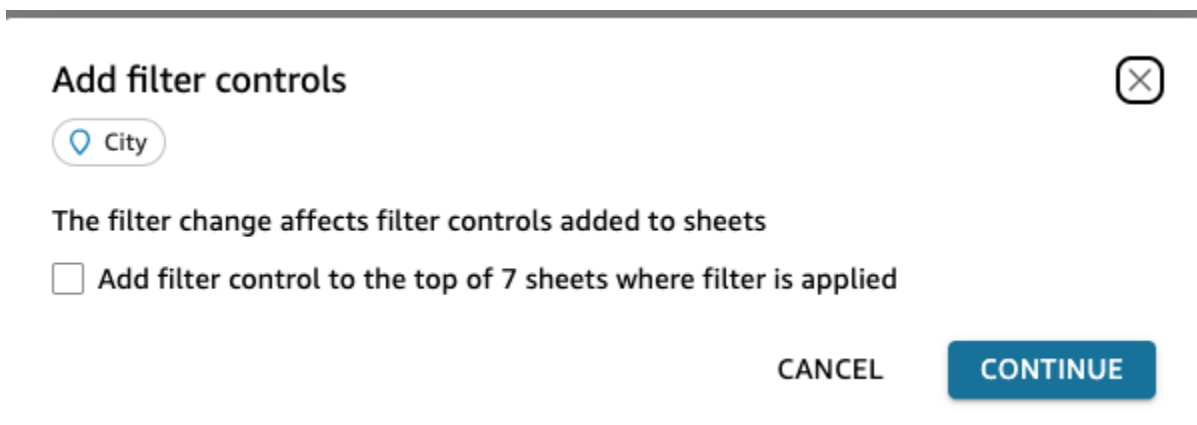


Increasing Scope of Existing Control

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 get the following 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

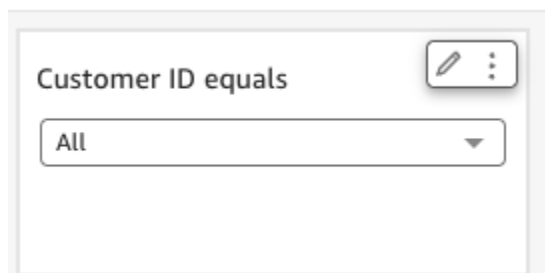
1. Go to the cross-sheet control and select the three-dot menu if the control is pinned to the top or the edit pencil icon if the control is on the sheet. You will be presented with the following options:

- **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**

On sheet



Top of sheet



2. Choose **Edit**. This brings up the **Format Control** pane on the right side of your analysis.

Format control



Cross-sheet filter

City

Cross-sheet settings

Sheets (8)

DISPLAY SETTINGS

Title

City equals



Show info icon on control

Info icon text

Add more information for readers

CONTROL OPTIONS

Dropdown - multiselect

Values

Filter

Specific values

Hide Select all option from the control values

Relevant value

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)** (as shown following:

Format control [X]

Cross-sheet filter
City

Cross-sheet settings ⓘ ^

Added to sheets [X]

- Sheet 1 [Top]
- Sheet 2 [Top]
- Sheet 3 [Sheet]
- Sheet 4 [Top]
- Sheet 5 [Top]
- Sheet 6 [Top]
- Sheet 7 [Top]
- Sheet 8 [Top]

DISPLAY SETTINGS

Title
City equals

Show info icon on control

Info icon text
Add more information for readers

CONTROL OPTIONS

Dropdown - multiselect [v]

Values

- Filter
- Specific values

Hide Select all option from the control values ⓘ

Relevant value [v]

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 QuickSight

You can create several different types of filters in Amazon QuickSight. 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

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 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 QuickSight](#).

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 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 QuickSight](#). 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 timestampz, 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 QuickSight 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

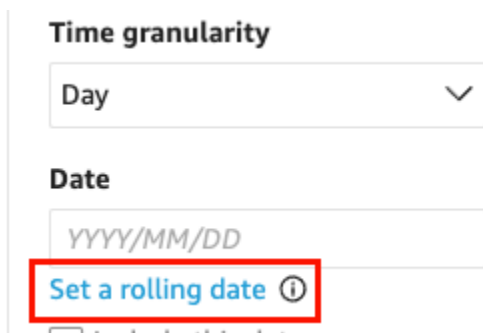
1. Create a new filter using a text field. For more information about creating filters, see [Adding filters](#).
2. In the **Filters** pane, choose the new filter to expand it.
3. For **Filter type**, choose **Date & time range**.
4. For **Condition**, choose a comparison type: **Between**, **After**, **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**.



Time granularity

Day

Date

YYYY/MM/DD

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 QuickSight](#).

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 QuickSight](#).

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 QuickSight](#).

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, Amazon QuickSight 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 QuickSight](#).

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 Amazon QuickSight, 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 QuickSight](#), [Using a control with a parameter in Amazon QuickSight](#), 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 [???](#).

Topics

- [Adding filter controls](#)
- [Pinning filter controls to the top of a sheet](#)
- [Customizing filter controls](#)

Adding filter controls


Use the following procedure to add a filter control.

To add a filter control

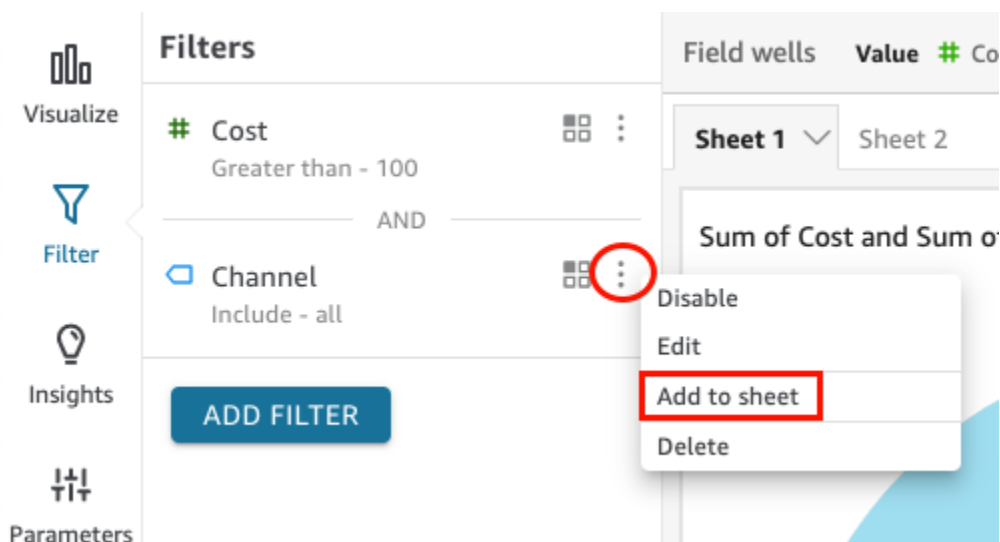
1. From the QuickSight start page, choose **Analyses**, and then choose the analysis that you want to work with.

The screenshot shows the Amazon QuickSight interface. At the top, there is a blue header bar with a search icon, the text "2018-2019 Platfor...", a dropdown arrow, and the word "Type". Below the header, the left sidebar contains a search bar labeled "Find analyses & more" with a magnifying glass icon. Underneath are several menu items: "Favorites" (star icon), "Recent" (clock icon), "My folders" (folder icon), "Shared folders" (folder with hand icon), "Dashboards" (bar chart icon), "Analyses" (bar chart with pencil icon, highlighted with a red box and a light blue background), "Datasets" (cylinder icon), and "Topics" (speech bubble icon). The main content area is titled "Analyses" and displays two analysis cards. The first card shows a bar chart with a value of "8.03% ↑" and is titled "Business Review analysis", updated "2 hours ago". The second card shows a bar chart and is titled "web-and-social-analytics 1...", updated "3 months ago".

2. In the analysis, choose **Filter** at left.



3. If you don't already have some filters available, create one. For more information about creating filters, see [Adding filters](#).
4. 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 Amazon QuickSight](#).
 - **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 Amazon QuickSight](#).

- **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.

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.

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.

- **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. (Optional) 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:

- 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).

To create a cascading control, do the following:

- a. Choose **Show relevant values only**. Note that this option might not be available for all filter control types.
 - b. In the **Show relevant values only** pane that opens, choose one or more controls from the available list.
 - c. Choose a field to match the value to.
 - d. Choose **Update**.
3. When finished, choose **Apply**.

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.

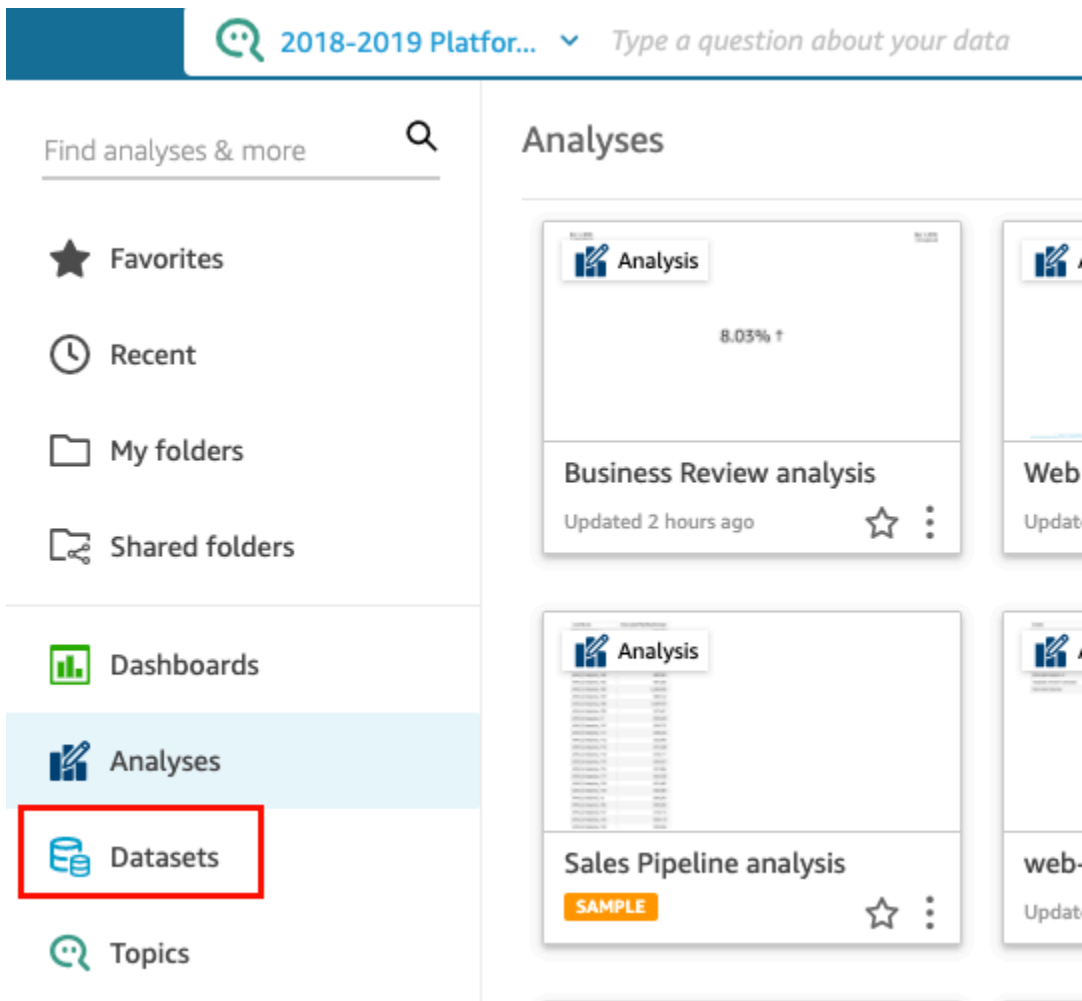
Use the following procedures to learn how to edit filters.

Editing filters in datasets

Use the following procedure to edit filters in datasets.

To edit a filter in a dataset

1. From the QuickSight start page, choose **Datasets**.



2. Choose the dataset that you want, and then choose **Edit dataset**.
3. On the data preparation page that opens, choose **Filters** at lower left.

Augment with SageMaker

Search fields

Focus

All fields

Select All | None

- Customer
- Date
- Customer Region
State
- Segment-1
- Service Line
- Revenue Goal
- Billed Amount
- Cost
- Channel
- Distinct ID

Excluded fields No fields excluded

Filters No filters applied [Add filter](#)

Query mode [Refresh now](#)

Business Review

Date	Customer ID	Customer ...	Customer ..
<input checked="" type="checkbox"/> Date	<input type="checkbox"/> String	<input type="checkbox"/> String	<input type="checkbox"/> State
2012-01-01...	DXegKx8qH...	SMB10	APAC
2012-01-01...	DXegKx8qH...	SMB10	APAC
2012-01-01...	A28Dzrr5dn...	SMB64	APAC
2012-01-01...	A28Dzrr5dn...	SMB64	APAC
2012-01-01...	A28Dzrr5dn...	SMB64	APAC
2012-01-01...	mbaEj8eHB...	SMB55	APAC
2012-01-01...	mbaEj8eHB...	SMB55	APAC
2012-01-01...	A28Dzrr5dn...	SMB64	APAC

- Choose the filter that you want to edit.
- 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

- From the QuickSight start page, choose **Analyses**.

2018-2019 Platfor... Type

Find analyses & more

- Favorites
- Recent
- My folders
- Shared folders

Dashboards

Analyses

Datasets

Topics

Analyses

Analysis

8.03% ↑

Business Review analysis

Updated 2 hours ago

Analysis

web-and-social-analytics 1....

Updated 3 months ago

- On the **Analyses** page, choose the analysis that you want to work with.
- In the analysis, choose **Filter** at left.

The screenshot shows the Amazon QuickSight interface. At the top, there is a header with a '+ ADD' button, a search bar containing '2018-2019 Platfor...', and a prompt 'Type a question at'. Below the header, the interface is divided into several sections:

- Visualize:** A section with a bar chart icon and the text 'Visualize'.
- Dataset:** A dropdown menu showing 'SPICE Business Review' with a '100%' indicator.
- Fields list:** A search bar labeled 'Search fields' and a list of fields:
 - Customer
 - Billed Amount
 - Channel
 - Cost (highlighted in green)
 - Customer Region
 - Date (highlighted in blue)
 - Distinct ID
 - Revenue Goal
- Filter:** A filter icon (a funnel) is highlighted with a red box.
- Insights:** A lightbulb icon and the text 'Insights'.
- Parameters:** An icon with four arrows and the text 'Parameters'.
- Actions:** A cursor icon and the text 'Actions'.

On the right side, there is a 'Field wells' section with 'Value # Cost (Sum)'. Below that, there are tabs for 'Sheet 1' and 'Sheet 2'. The main visualization area shows a partial view of a chart with the title 'Sum of Cost and Sum of Revenue'.

4. Choose the filter that you want to edit.
5. 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.

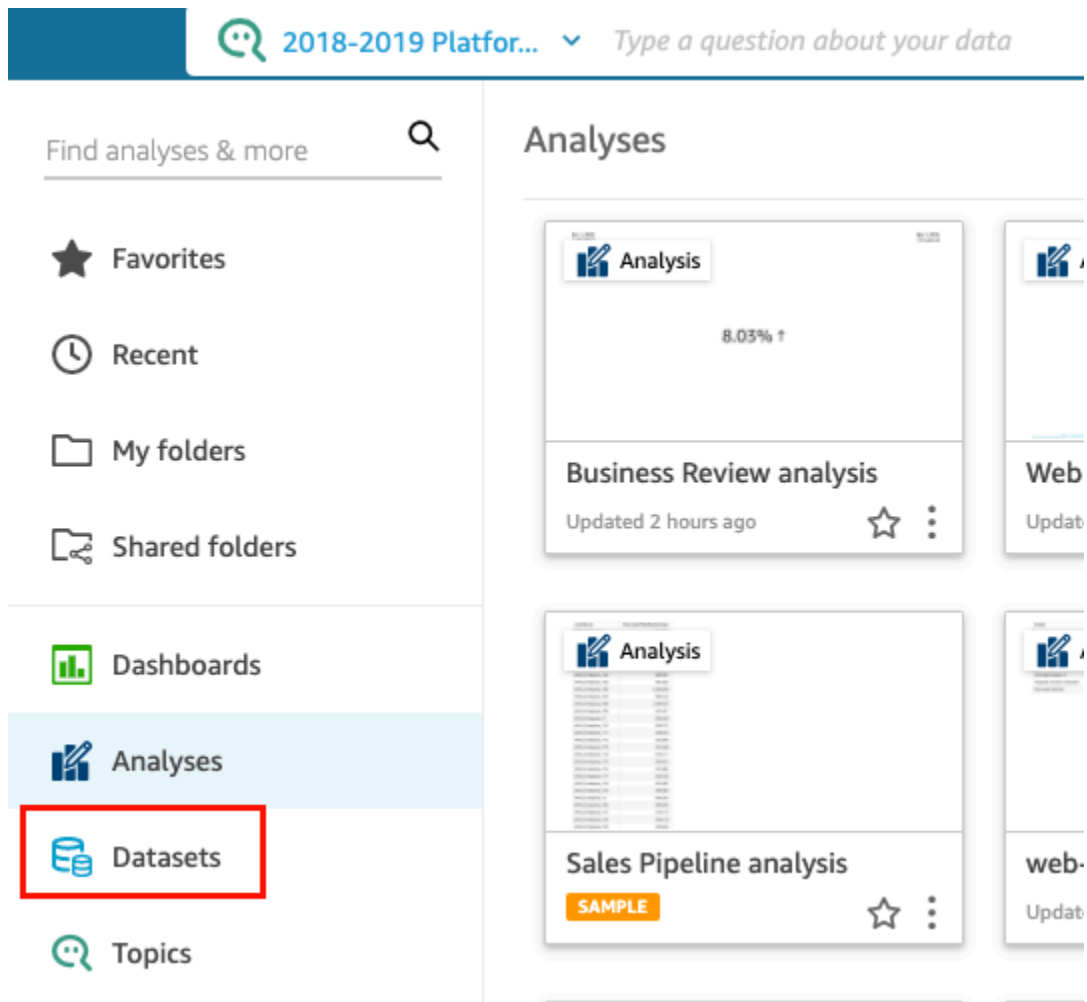
Use the following procedures to learn how to enable or disable filters.

Disabling filters in datasets

Use the following procedure to disable filters in datasets.

To disable a filter in a dataset

1. From the QuickSight start page, choose **Datasets**.



The screenshot shows the Amazon QuickSight interface. At the top, there is a search bar with the text "2018-2019 Platfor..." and a dropdown arrow, followed by the prompt "Type a question about your data". Below this is a navigation sidebar on the left with the following items: "Find analyses & more" (with a search icon), "Favorites", "Recent", "My folders", "Shared folders", "Dashboards", "Analyses" (highlighted in light blue), "Datasets" (highlighted with a red box), and "Topics". The main content area is titled "Analyses" and displays a grid of analysis cards. The top card is titled "Business Review analysis" and shows a value of "8.03% ↑" and "Updated 2 hours ago". The bottom card is titled "Sales Pipeline analysis" and shows a "SAMPLE" label and "Updated 2 hours ago".

2. Choose the dataset that you want, and then choose **Edit dataset**.
3. On the data preparation page that opens, choose **Filters** at lower left.

The screenshot shows the Amazon QuickSight interface. On the left, there is a sidebar with a search bar for fields, a 'Focus' dropdown set to 'All fields', and a list of filters. The 'Filters' section is highlighted with a red box and shows 'No filters applied' with an 'Add filter' button. Below the filters, there is a 'Query mode' section with a 'Refresh now' button. On the right, a data table is displayed with columns: Date, Customer ID, Customer ... (SMB), and Customer ... (State). The table contains several rows of data, with the last row highlighted in blue. At the top right, a 'Business Review' dropdown menu is visible.

- 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

- From the QuickSight start page, choose **Analyses**.

2018-2019 Platfor... Type

Find analyses & more

- Favorites
- Recent
- My folders
- Shared folders

Dashboards

Analyses

Datasets

Topics

Analyses

Analysis

8.03% ↑

Business Review analysis

Updated 2 hours ago

Analysis

web-and-social-analytics 1....

Updated 3 months ago

- On the **Analyses** page, choose the analysis that you want to work with.
- In the analysis, choose **Filter** at left.

The screenshot shows the Amazon QuickSight interface. At the top, there is a header with a '+ ADD' button, a search bar containing '2018-2019 Platfor...', and a prompt 'Type a question at'. Below the header, the interface is divided into several sections:

- Visualize:** A section with a bar chart icon and the text 'Visualize'.
- Dataset:** A section showing 'SPICE Business Review' with a dropdown arrow and '100%'.
- Fields list:** A section with a search bar 'Search fields' and a list of fields:
 - Customer
 - Billed Amount
 - Channel
 - Cost (highlighted in green)
 - Customer Region
 - Date (highlighted in blue)
 - Distinct ID
 - Revenue Goal
- Field wells:** A section with 'Value # Cost (Sum)' and 'Sheet 1' selected.
- Visualization:** A section showing a partial view of a chart titled 'Sum of Cost and Sum of Revenue'.

The 'Filter' icon in the left sidebar is highlighted with a red box.

- 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

- From the QuickSight start page, choose **Datasets**.

The screenshot shows the Amazon QuickSight user interface. At the top, there is a header with a search icon, the text "2018-2019 Platfor...", and a prompt "Type a question about your data". Below the header is a navigation sidebar on the left with the following items: "Find analyses & more" (with a search icon), "Favorites", "Recent", "My folders", "Shared folders", "Dashboards", "Analyses" (highlighted in light blue), "Datasets" (highlighted with a red box), and "Topics". The main content area is titled "Analyses" and displays a grid of analysis cards. The top card is titled "Business Review analysis" and shows a value of "8.03% ↑". The bottom card is titled "Sales Pipeline analysis" and has a "SAMPLE" label. Each card includes a star icon and a three-dot menu icon.

2. Choose the dataset that you want, and then choose **Edit dataset**.
3. On the data preparation page that opens, choose **Filters** at lower left.

The screenshot shows the Amazon QuickSight interface. On the left sidebar, there is a search bar for fields, a 'Focus' dropdown set to 'All fields', and a list of fields to select. The 'Filters' section is highlighted with a red box. The main view displays a 'Business Review' visualization and a data table.

Field Selection List:

- Customer
- Date
- Customer Region (State)
- Segment-1
- Service Line
- Revenue Goal
- Billed Amount
- Cost
- Channel
- Distinct ID

Excluded fields: No fields excluded

Filters: No filters applied [Add filter](#)

Query mode: [Refresh now](#)

Dataset Table:

Date	Customer ID	Customer ...	Customer ..
2012-01-01...	DXegKx8qH...	SMB10	APAC
2012-01-01...	DXegKx8qH...	SMB10	APAC
2012-01-01...	A28Dzrr5dn...	SMB64	APAC
2012-01-01...	A28Dzrr5dn...	SMB64	APAC
2012-01-01...	A28Dzrr5dn...	SMB64	APAC
2012-01-01...	mbaEj8eHB...	SMB55	APAC
2012-01-01...	mbaEj8eHB...	SMB55	APAC

- 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

- From the QuickSight start page, choose **Analyses**.

2018-2019 Platfor... Type

Find analyses & more

- Favorites
- Recent
- My folders
- Shared folders

Dashboards

Analyses

Datasets

Topics

Analyses

Analysis

8.03% ↑

Business Review analysis

Updated 2 hours ago

Analysis

web-and-social-analytics 1....

Updated 3 months ago

- On the **Analyses** page, choose the analysis that you want to work with.
- In the analysis, choose **Filter** at left.

The screenshot shows the Amazon QuickSight interface. On the left sidebar, the 'Filter' icon is highlighted with a red box. The main area displays a dataset named 'Business Review' with a 'Fields list' containing 'Cost' and 'Date'. The 'Field wells' section shows 'Value # Cost (Sum)' and a visualization titled 'Sum of Cost and Sum of Revenue'.

4. 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
- **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 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 QuickSight can display them on a map. Amazon QuickSight 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

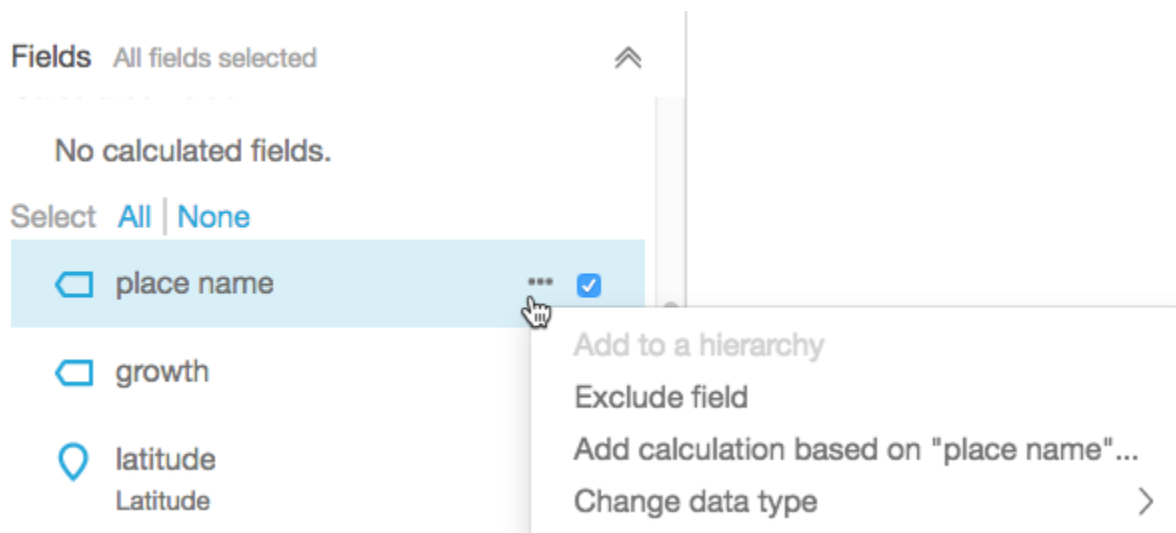
Geospatial charts in Amazon QuickSight aren't currently supported in some Amazon Web Services Regions, including in China. We are working on adding support for more Regions.

Use the following procedures to add geospatial data types and hierarchies to your dataset.

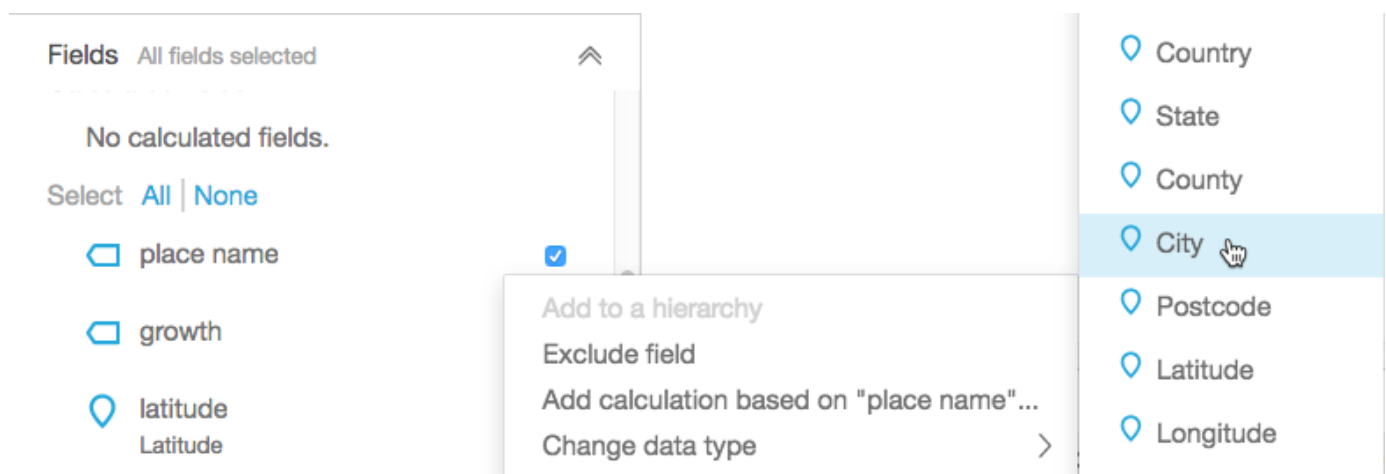
To add geospatial data types and hierarchies to your dataset

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.

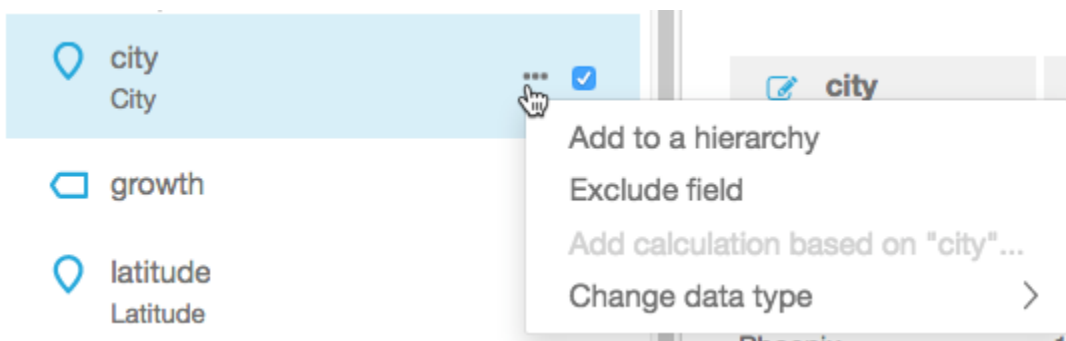


Also check the name of the data type, for example latitude or country.

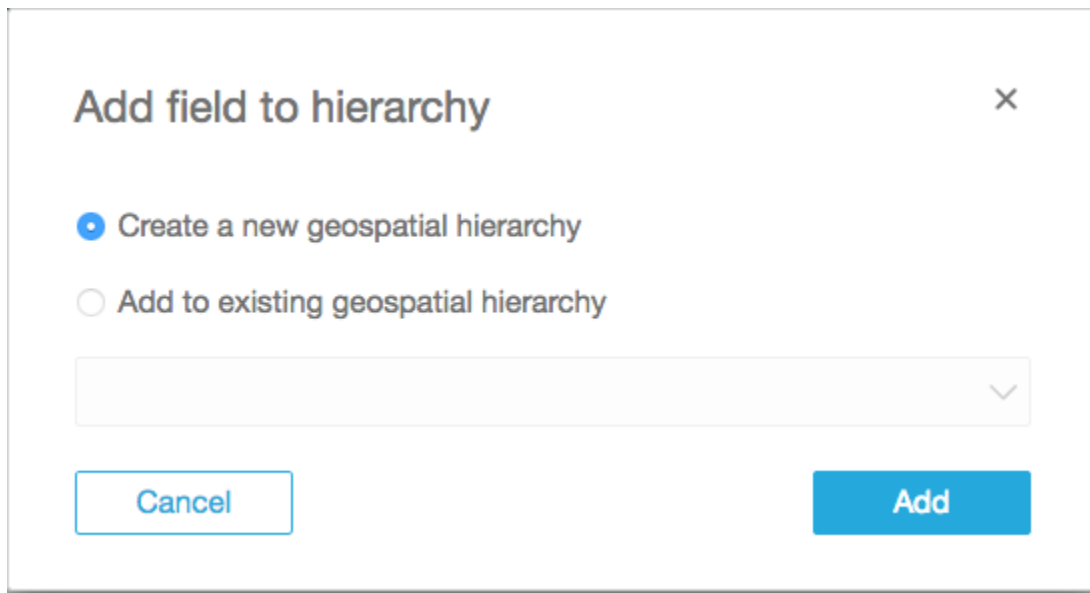
3. (Optional) You can set up a hierarchy or grouping for geographical components (state, city), or for latitude and longitude coordinates. For coordinates, you must add both latitude and longitude to the geospatial field wells.

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.



The screenshot shows a dialog box titled "Add field to hierarchy" with a close button (X) in the top right corner. It contains two radio button options: "Create a new geospatial hierarchy" (which is selected) and "Add to existing geospatial hierarchy". Below these options is a text input field with a downward arrow on the right side. At the bottom of the dialog, there are two buttons: "Cancel" on the left and "Add" on the right.

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 as follows. 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**.

Create hierarchy ×

Select fields that correspond to latitude and longitude

Name your latitude and longitude hierarchy

Latitude-Longitude

Field to use for latitude

Latitude

Field to use for longitude

Longitude

Cancel Update

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.

Create hierarchy ×

Select a country or a field to use for country

Name your hierarchy

state-city

This hierarchy is for a single country

United States ▾

This hierarchy is for multiple countries

▾

Cancel Update

For either hierarchy type, choose **Update** to continue.

- Continue by adding as many fields to the hierarchy as you need to.


Your geospatial groupings appear in the **Fields** pane.

Fields All fields selected

Select [All](#) | [None](#)

state-city

United States

 state
State

 city
City


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 QuickSight requirements for correctly processing geospatial data. If Amazon QuickSight 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 QuickSight 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 QuickSight team from within the Amazon QuickSight console.

Topics

- [Geocoding issues](#)
- [Issues with latitude and longitude](#)

Geocoding issues

Amazon QuickSight geocodes place names into latitude and longitude coordinates. It uses these coordinates to display place names on the map. Amazon QuickSight 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 QuickSight 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 QuickSight 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 QuickSight. 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 QuickSight supports latitude and longitude coordinates within specific ranges.

Coordinate	Valid range
Latitude	Between -90 and 90
Longitude	Between -180 to 180

Amazon QuickSight 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 QuickSight, 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

Using unsupported or custom dates

Amazon QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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.

Revenue ...	Distinct ID	Service Line	
APAC	0	HR	
APAC	0	Billing	
APAC	1	HR	
APAC	0	Billina	

Note

Integer dates that aren't Unix epoch datetimes don't work as is. For example, these formats are not supported as integers: MMdyy, MMddy, ddMMyy, ddMMyyy, 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 QuickSight 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**.

Integrating Amazon SageMaker models with Amazon QuickSight

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 QuickSight.

You can augment your Amazon QuickSight Enterprise edition data with Amazon SageMaker machine learning models. You can run inferences on data stored in SPICE imported from any data source supported by Amazon QuickSight. For a full list of supported data sources, see [Supported data sources](#).

Using Amazon QuickSight with SageMaker 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 Amazon QuickSight serverless approach, the process scales seamlessly, so you don't need to worry about inference or query capacity.

Amazon QuickSight supports SageMaker 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 Amazon QuickSight to provide predictions, the SageMaker 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. Amazon QuickSight doesn't support multiple values inside a single column.

Topics

- [How SageMaker integration works](#)
- [Costs incurred \(no additional costs with integration itself\)](#)
- [Usage guidelines](#)
- [Defining the schema file](#)
- [Adding a SageMaker model to your QuickSight dataset](#)
- [Build predictive models with SageMaker Canvas](#)

How SageMaker integration works

In general, the process works like this:

1. An Amazon QuickSight administrator adds permissions for Amazon QuickSight to access SageMaker. To do this, open **Security & Permissions** settings from the **Manage QuickSight** page. Go to **QuickSight access to Amazon services**, and add SageMaker.

When you add these permissions, Amazon QuickSight is added to an Amazon Identity and Access Management (IAM) role that provides access to list all the SageMaker models in your Amazon account. It also provides permissions to run SageMaker jobs that have names that are prefixed with `quicksight-auto-generated-`.

2. We recommend that you connect to an SageMaker model that has an inference pipeline, because it automatically performs data preprocessing. For more information, see [Deploy an Inference Pipeline](#) in the *SageMaker 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. 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 Amazon QuickSight, 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 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. You use these fields just as you use other fields, within the guidelines described in [Usage guidelines](#).

When you run SageMaker integration, Amazon QuickSight passes a request to SageMaker to run batch transform jobs with inference pipelines. Amazon QuickSight 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, Amazon QuickSight names all its SageMaker 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 batch transform to update the output fields with new data.

You can use the Amazon QuickSight SPICE ingestion API operations to control the data refresh process. For more information about using these API operations, see the [Amazon QuickSight 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, 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 if the dataset has inferred fields. Costs are incurred in the same Amazon account where your QuickSight subscription is.
- Your QuickSight subscription costs are as follows:
 - The cost of storing your data in the in-memory calculation engine in QuickSight (SPICE). If you are adding new data to SPICE, you might need to purchase enough SPICE capacity to accommodate it.
 - QuickSight 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 QuickSight, 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 QuickSight 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, the following usage guidelines apply to a pretrained model that you use with Amazon QuickSight:

- When you create the model, associate it with the Amazon Resource Name (ARN) for the appropriate IAM role. The IAM role for the SageMaker model needs to have access to the Amazon S3 bucket that Amazon QuickSight 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 QuickSight 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, not real-time data. You can't use an SageMaker endpoint.

Defining the schema file

Before you use an SageMaker model with Amazon QuickSight data, create the JSON schema file that contains the metadata that Amazon QuickSight needs to process the model. The Amazon QuickSight 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 model expects for the input data. The only supported value for this is "text/csv". QuickSight 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 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. QuickSight 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`, QuickSight drops rows that don't contain this value before calling SageMaker. Doing this helps avoid causing SageMaker to fail on missing required data.

output

A list of output columns that the SageMaker model produces. QuickSight 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 QuickSight. You can override the name specified here in QuickSight. 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 can provision to run the transform job. The list is provided to the QuickSight user to choose from. This list is limited to the types supported by SageMaker. For more information on supported types, see [TransformResources](#) in the *SageMaker Developer Guide*.

defaultInstanceType

(Optional) The instance type that is presented as the default option in the SageMaker wizard in QuickSight. Include this instance type in `instanceTypes`.

instanceCount

(Optional) The instance count defines how many of the selected instances for SageMaker 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 model to communicate with the person who is using this model in QuickSight. 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.

Adding a SageMaker model to your QuickSight dataset

Using the following procedure, you can add a pretrained SageMaker 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 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 QuickSight dataset with SageMaker

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 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 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. QuickSight 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 batch transform job to update the output fields with new data.

Build predictive models with SageMaker Canvas

QuickSight authors can export data into SageMaker Canvas to build ML models that can be sent back to QuickSight. Authors can use these ML models to augment their datasets with predictive analytics that can be used to build analyses and dashboards.

Prerequisites

- A QuickSight account that's integrated with IAM Identity Center. If your QuickSight account isn't integrated with IAM Identity Center, create a new QuickSight 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 QuickSight with IAM Identity Center, see [Configure your Amazon QuickSight account with IAM Identity Center](#).
 - To import assets from an existing QuickSight account to a new QuickSight account that's integrated with IAM Identity Center, see [Asset bundle operations](#).
- A new SageMaker domain that is integrated with IAM Identity Center. For more information about onboarding to SageMaker Domain with IAM Identity Center, see [Onboard to SageMaker Domain using IAM Identity Center](#).

Topics

- [Build a predictive model in SageMaker Canvas from Amazon QuickSight](#)
- [Create a dataset with a SageMaker Canvas model](#)

- [Considerations](#)

Build a predictive model in SageMaker Canvas from Amazon QuickSight

To build a predictive model in SageMaker Canvas

1. Log in to QuickSight 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 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 Canvas console.
5. In SageMaker Canvas, create a predictive model with the data that you exported from QuickSight. 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 Canvas, see [Build a model](#).
6. Send the predictive model back to QuickSight. For more information about sending a model from SageMaker Canvas to Amazon QuickSight, see [Send your model to Amazon QuickSight](#).

Create a dataset with a SageMaker Canvas model

After you create a predictive model in SageMaker Canvas and send it back to QuickSight, 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 QuickSight console, navigate to the **Datasets** page, and choose **Datasets**.
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** modal.
5. For **Model**, choose the model that you sent to QuickSight from SageMaker 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 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 Canvas, the data is imported into SPICE and a batch inference job begins in SageMaker. It can take up to 10 minutes for these processes to complete.

Considerations

The following limitations apply to the creation of SageMaker Canvas models with QuickSight data.

- The **Build a predictive model** option that is used to send data to SageMaker 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 QuickSight 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 QuickSight](#)
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 QuickSight 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 Amazon QuickSight](#)
7. Check the [SPICE](#) indicator to see if you have enough space to import the dataset. Amazon QuickSight 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 you 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 Amazon QuickSight](#)
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 datasets from new database data sources](#). 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 QuickSight if you choose to use joined tables. For more information about joining data using the Amazon QuickSight 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 Amazon QuickSight](#)

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.

Visualizing data in Amazon QuickSight

Following, you can find descriptions of how to create and customize Amazon QuickSight charts, arrange charts in a dashboard, and more.

Topics

- [Working with an analysis in Amazon QuickSight](#)
- [Adding sheets](#)
- [Working with interactive sheets in Amazon QuickSight](#)
- [Working with paginated reports in Amazon QuickSight](#)
- [Working with items on sheets in Amazon QuickSight analyses](#)
- [Using themes in Amazon QuickSight](#)
- [Accessing Amazon QuickSight using keyboard shortcuts](#)

Working with an analysis in Amazon QuickSight

In Amazon QuickSight, 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 QuickSight analysis.

Topics

- [Starting an analysis in Amazon QuickSight](#)
- [Adding a title and description to an analysis](#)
- [Renaming an analysis](#)
- [Duplicating analyses](#)
- [Viewing analysis details](#)
- [Customize date and time values of an analysis](#)
- [The analysis menu](#)
- [Saving changes to analyses](#)
- [Exporting data from Amazon QuickSight analyses](#)
- [Deleting an analysis](#)

Starting an analysis in Amazon QuickSight

In Amazon QuickSight, 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 QuickSight 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 QuickSight](#). For more information on sheets, see [Adding 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 QuickSight](#).

New sheet ×




Interactive sheet

Single page, interactive content

Layout

Optimize for viewing on



Paginated report

Multi-page, highly formatted document

Paper size

Portrait Landscape

CANCEL ADD

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)

The default paper size is US letter (8.5 x 11 in.)

- (Optional) Choose the orientation of the sheet. You can choose between **Portrait** or **Landscape**. The default option is portrait.

Before you can create Amazon QuickSight paginated reports, first get the paginated reporting add-on for your QuickSight account. For more information on getting the paginated reporting add-on, see [Get the QuickSight paginated reports add-on](#).

For more information on formatting paginated reports, see [Working with paginated reports in Amazon QuickSight](#).

New sheet

Interactive sheet
Single page, interactive content
Layout: Free-form
Optimize for viewing on: 1600px

Paginated report
Multi-page, highly formatted document
Paper size: US letter - 8.5 x 11 in
 Portrait Landscape

CANCEL ADD

6. Choose **Add**.
7. Create a visual. For more information about creating visuals, see [Adding visuals to Amazon QuickSight 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 **Add** in the application bar and then choose **Add title**.
2. For **Sheet title**, enter a title and press **Enter**. To remove a title, choose **Add** 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 QuickSight](#).

3. Choose **Add** in the application bar, and then choose **Add description**.



4. Enter a description in the **Description** box and press **Enter**. To remove a description, choose **Add** in the application bar and then choose **Remove description**. Or, to remove the description, you can select the description and then choose the x-shaped delete icon.



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 Amazon QuickSight. Use the following procedure to learn how.

To duplicate an analysis

1. From the QuickSight start page, 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 QuickSight start page and selecting **Analyses**.

Viewing analysis details

To view an analysis, locate the analysis on the **All analyses** tab of the Amazon QuickSight start page. Then choose the analysis.



Customize date and time values of an analysis

In Amazon QuickSight, 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

QuickSight 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.

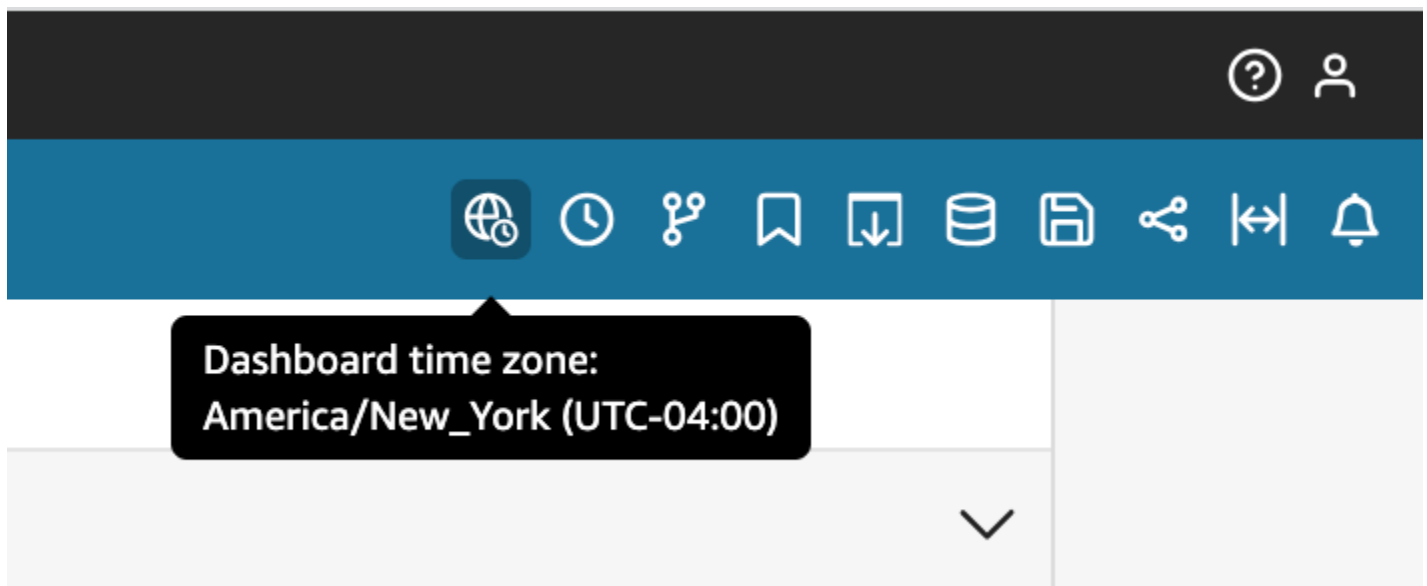
- **Datetime displays with a granularity that is lower than hour are converted to the selected time zone.** For example, if you set the timezone of an analysis to `America/New_York` (UTC-04:00), the datetime value `Dec . 1, 2020 12:00am` in UTC+00:00 is converted and displayed as `Nov . 30, 2020 7:00pm`. 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, QuickSight 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 QuickSight 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.
- QuickSight 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.
- QuickSight 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

QuickSight 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**.

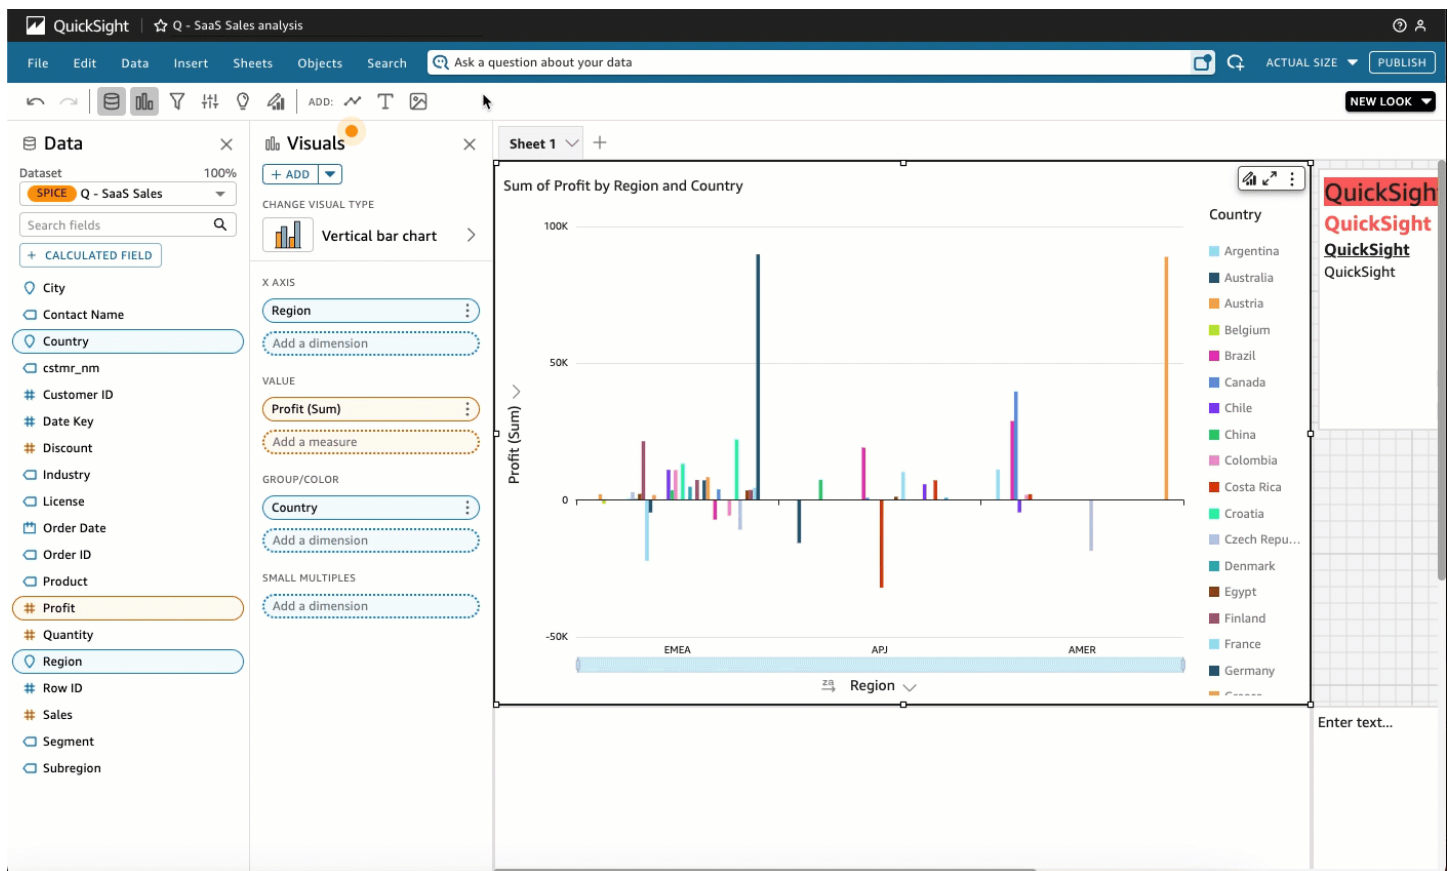
Considerations

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.
- QuickSight 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.
- QuickSight readers can't change the week start day of a dashboard.


The analysis menu

While working on an analysis, Amazon QuickSight provides menu options, as shown at the top of the following screenshot. 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 **Format visual** 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.

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

To change the Autosave mode for an analysis, choose **Autosave** at upper right in the application bar, 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 QuickSight gives you the option to do one of two things. You can either let Amazon QuickSight 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 Amazon QuickSight 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. See the following sections for more information.

Export data from an analysis to a CSV file


To export data from an analysis or dashboard to a comma-separated values (CSV) file, follow the procedure in [Exporting data from visuals](#).

Export data from an analysis to a PDF file

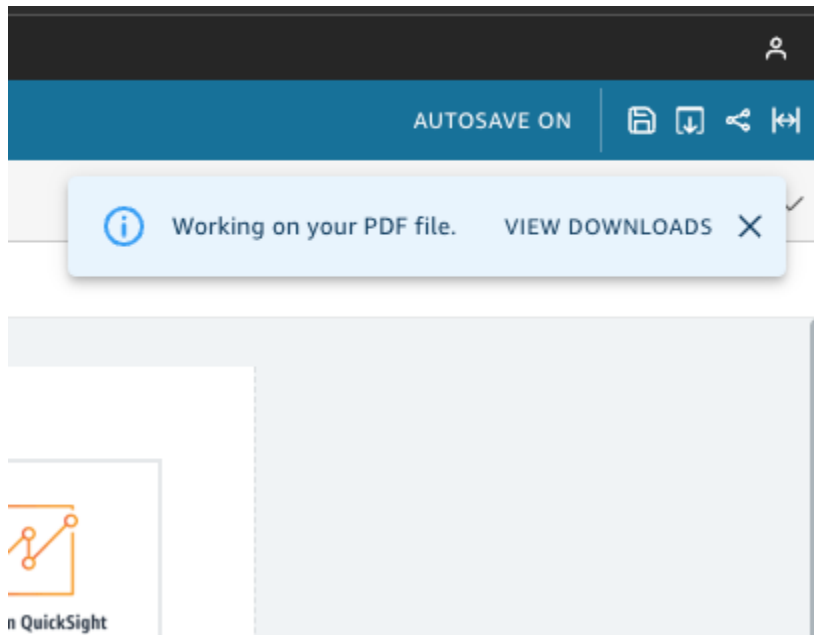
You can export content from a dashboard into a Portable Document Format file (PDF). Similar to a print-out, this format provides a snapshot of the current sheet as it appears on-screen at the time of download.

To export an analysis as a PDF

1. From the analysis that you want to export, choose the **Export** icon at the top right.
2. Choose **Generate PDF**.

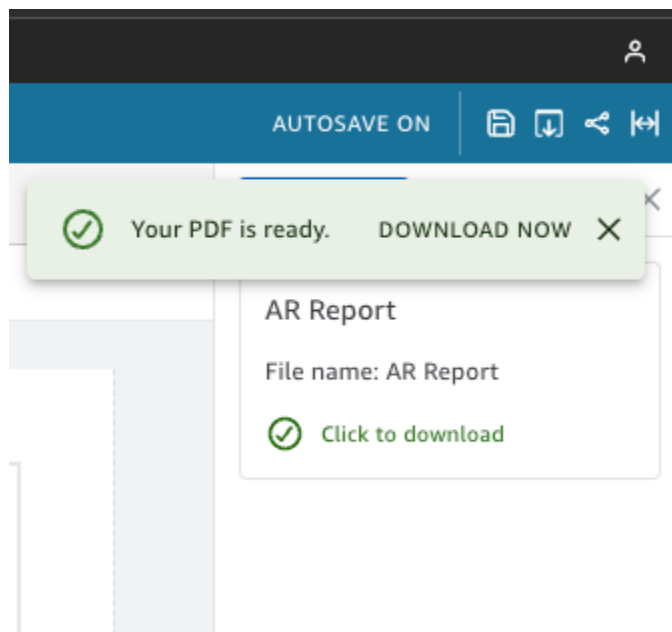


3. When you choose **Generate PDF**, QuickSight begins preparing the analysis for download. Choose **View downloads** in the blue pop-up to open the **Downloads** pane on the right.

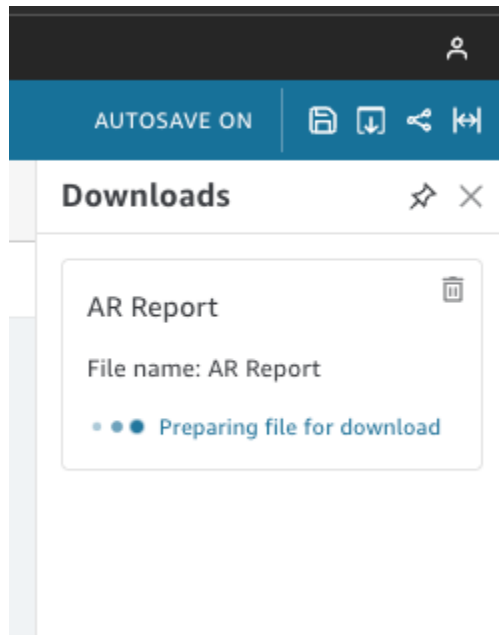


4. There are two ways to download your analysis:

- 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 analysis or report that is ready to download.



The process for exporting to a PDF works the same way in both dashboards and analyses.

You can also attach a PDF to dashboard email reports. For more information, see [Scheduling and sending reports by email](#).

Deleting an analysis

You can delete an analysis by using the **All analyses** tab of the Amazon QuickSight start page. Deleting an analysis doesn't affect any dashboards that are based on that analysis.

To remove an analysis, choose the details icon (:) on the analysis, and choose **Delete**. Confirm your choice by choosing **Delete** again. You can't undo this action.

Adding 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. QuickSight 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 [Arranging visuals in an interactive dashboard](#).

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 QuickSight, add the **Paginated reporting Add-on** to your Amazon QuickSight account. To get the **Paginated reporting Add-on** and start working with paginated reports, see [Working with paginated reports in Amazon QuickSight](#).

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.

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 QuickSight

An *Interactive Sheet* is a collection of data expressed in visuals that users can interact with when the sheet is published to a dashboard. QuickSight authors can add different 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 Amazon QuickSight](#).

For more information on formatting interactive sheets, see the following topics.

Topics

- [Arranging visuals in an interactive dashboard](#)

Arranging visuals in an interactive dashboard

You can customize the type of layout you want for visuals in your interactive dashboards. You can also create parameters and custom actions to add interactivity to dashboards.

See the following topics to learn more.

Topics

- [Customizing dashboard layouts in Amazon QuickSight](#)

- [Adding interactivity to dashboards in Amazon QuickSight](#)

Customizing dashboard layouts in Amazon QuickSight

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.

The screenshot displays the Amazon QuickSight interface. On the left is the 'Settings' panel, and on the right is the dashboard view.


Settings Panel:

- Visualize:** Filter, Insights, Parameters, Actions, Themes, Settings.
- Layout:**
 - Tiled**: Visuals snap to grid with standard spacing and alignment. Dashboards are displayed as designed, with options to fit to screen or view at actual size.
 - Free-form**: Visuals can be placed anywhere (including overlap) with precise coordinates. Dashboards are displayed as designed, with options to fit to screen or view at actual size.
 - Classic**: Visuals snap to grid with standard spacing and alignment. Dashboards will hide data or change formatting to fit smaller screen sizes.
- Optimize for viewing on:** 1600px (default)

Dashboard View:

- Field wells:** Sheet 1
- Sum of Mailing List Adds:** 9,889
- Sum of Free Sign Up and Sum of Mailing List Adds by Date:** A line chart showing trends from Jan 2015 to Dec 2016. The legend indicates 'Mailing list ...' (light blue) and 'Free sign up' (dark blue).

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.

The screenshot shows the Amazon QuickSight interface. On the left is the 'Settings' panel with a 'Dashboard layout' section. The 'Free-form' layout is selected, and the resolution is set to '1600px (default)'. The main area displays a dashboard with two visualizations: a horizontal bar chart titled 'Sum of Free Sign Up by Events' and a KPI card titled 'Sum of Mailing List Adds' showing the value '9,889'. The bar chart data is as follows:

Event	Sum of Free Sign Up
empty	~35M
2016 CPC campaign	~25M
New website promo	~5M
New mobile site promo	~3M
SEO optimization v2	~2M
Targeted outreach campaign	~1M
SEO optimization v1	~1M
2013 CPC campaign	~1M
Tech news features	~0.5M

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 QuickSight 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. From the Amazon QuickSight start page, choose **Analyses**, and then choose the analysis that you want to change.
2. On the analysis page, choose **Settings** at left.

Settings

Dashboard layout

Tiled
Visuals snap to grid with standard spacing and alignment. Dashboards are displayed as designed, with options to fit to screen or view at actual size.

Free-form
Visuals can be placed anywhere (including overlap) with precise coordinates. Dashboards are displayed as designed, with options to fit to screen or view at actual size.

On mobile devices, dashboard displays as published with no changes to the layout.

Once designed in this layout, switching to another dashboard layout may cause some visual elements to shift and will need adjusting.

Optimize for viewing on
1600px (default) ▾

Show page breaks
See break lines and page numbers for easier printing. This view is only visible when authoring.

Classic
Visuals snap to grid with standard spacing and alignment. Dashboards will hide data or change formatting to fit smaller screen sizes.

Optimize performance

Reload visuals each time I switch sheets

Update visuals manually
To manually refresh visuals and see your changes, click Update.

Apply

3. In the **Settings** pane, under **Dashboard layouts**, choose the layout you want.
4. 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 **Format visual** pane that opens at left, choose **Placement**.

Format visual ×

Title ⤴

KPI ⤴


Placement ⤵

X Y

Width Height

↔


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 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 and border can be customized in the **Style** pane of the **Format visual** menu.

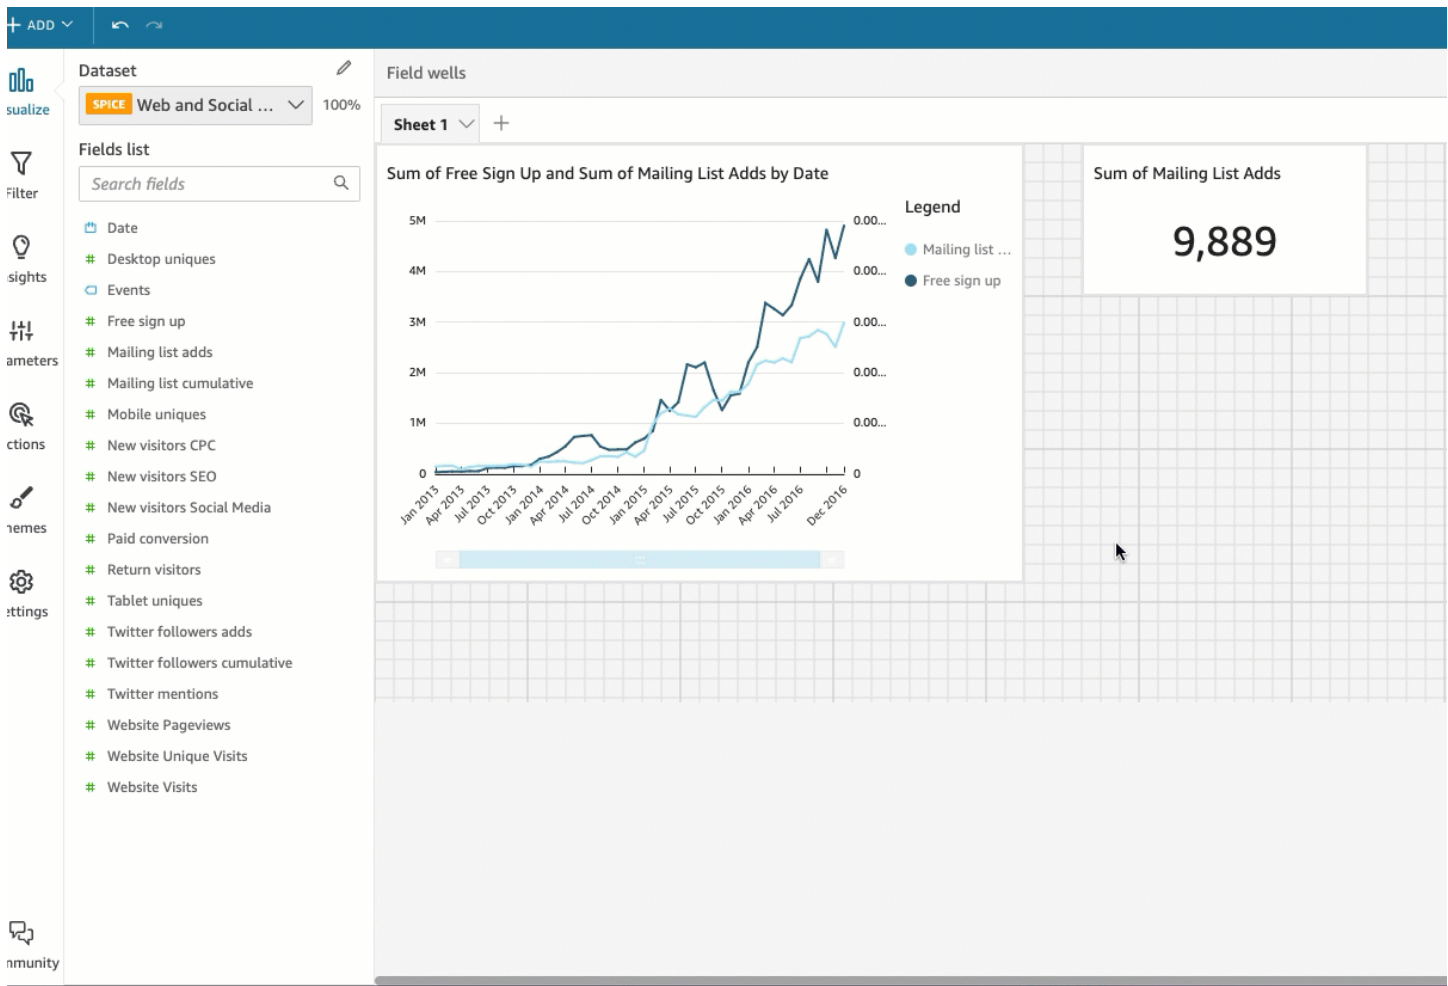
To change the color of a visual's background

1. Choose the visual that you want to change.
2. On the menu in the upper-right hand side of the visual, choose **Format visual**.

3. In the **Format visual** pane that appears on the left, choose **Style**.
4. Under the **Background color** switch, click the color box and choose **Custom color**.
5. Choose the color that you want.

You can also enter a color's hexadecimal code or change the color's opacity.

6. 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 **Format visual**.
3. In the **Format visual** pane that appears on the left, choose **Style**.

4. In the **Format visual** pane that appears on the left, choose **Reset**.

Changing the color of visual Borders

You can also customize the color of a visual's border.

To change the color of a visual's border

1. Choose the visual that you want to change.
2. On the menu in the upper-right hand side of the visual, choose **Format visual**.
3. In the **Format visual** pane that appears on the left, choose **Style**.
4. Under the **Border color** switch, click on the color box and choose **Custom color**.
5. Choose the color you want.

You can also enter a color's hexadecimal code or change the color's opacity.

6. Choose **Apply**.

Changing a visual's selection color

You can also customize the color that appears around a visual when it's selected.

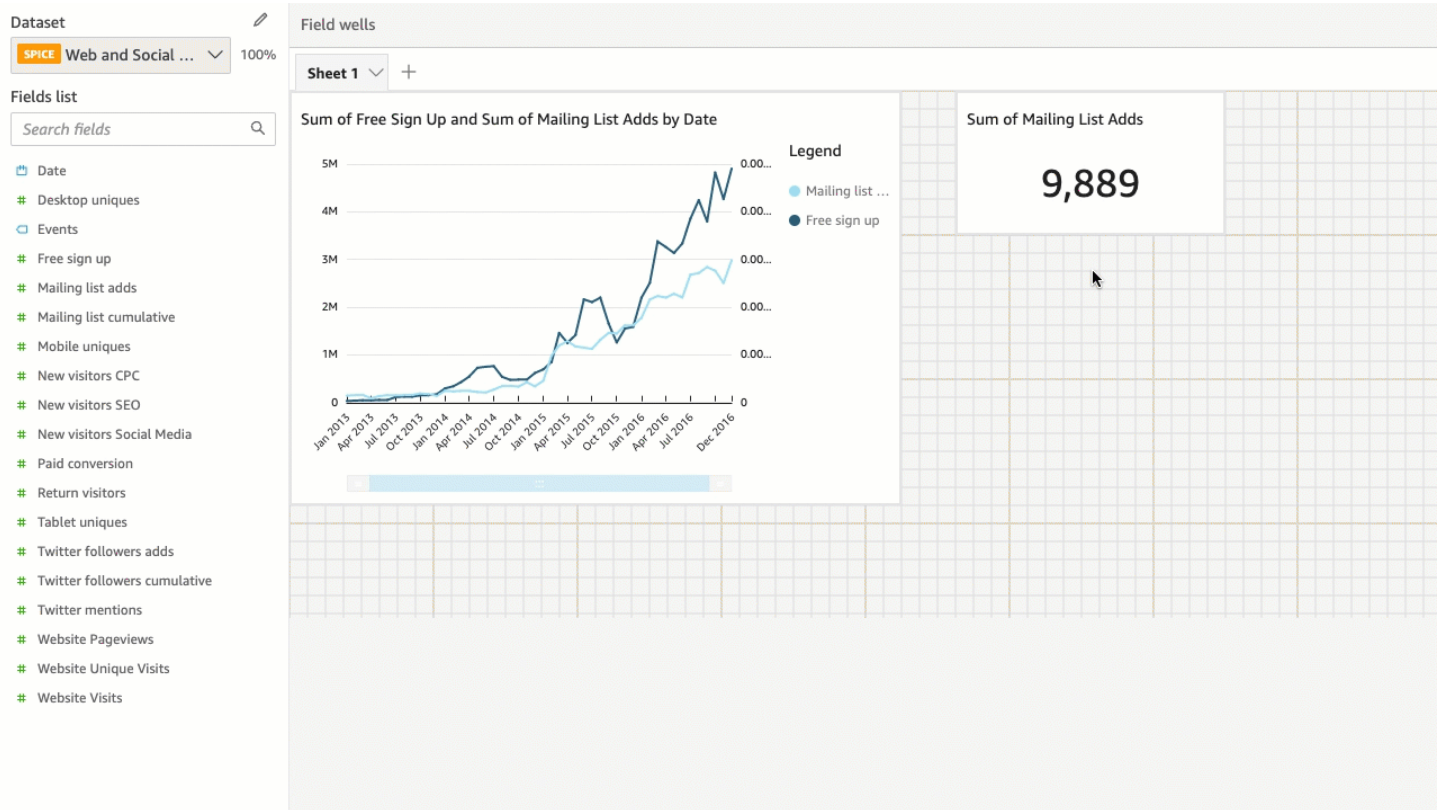
To change the color of a visual's selection border

1. Choose the visual that you want to change.
2. On the menu in the upper-right hand side of the visual, choose **Format visual**.
3. In the **Format visual** pane that appears on the left, choose **Style**.
4. Under the **Selection color** switch, click on the color box and choose **Custom color**.
5. Choose the color you want. (Optional)You can also enter a color's hex code or change the color's opacity.
6. Choose **Apply**.

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

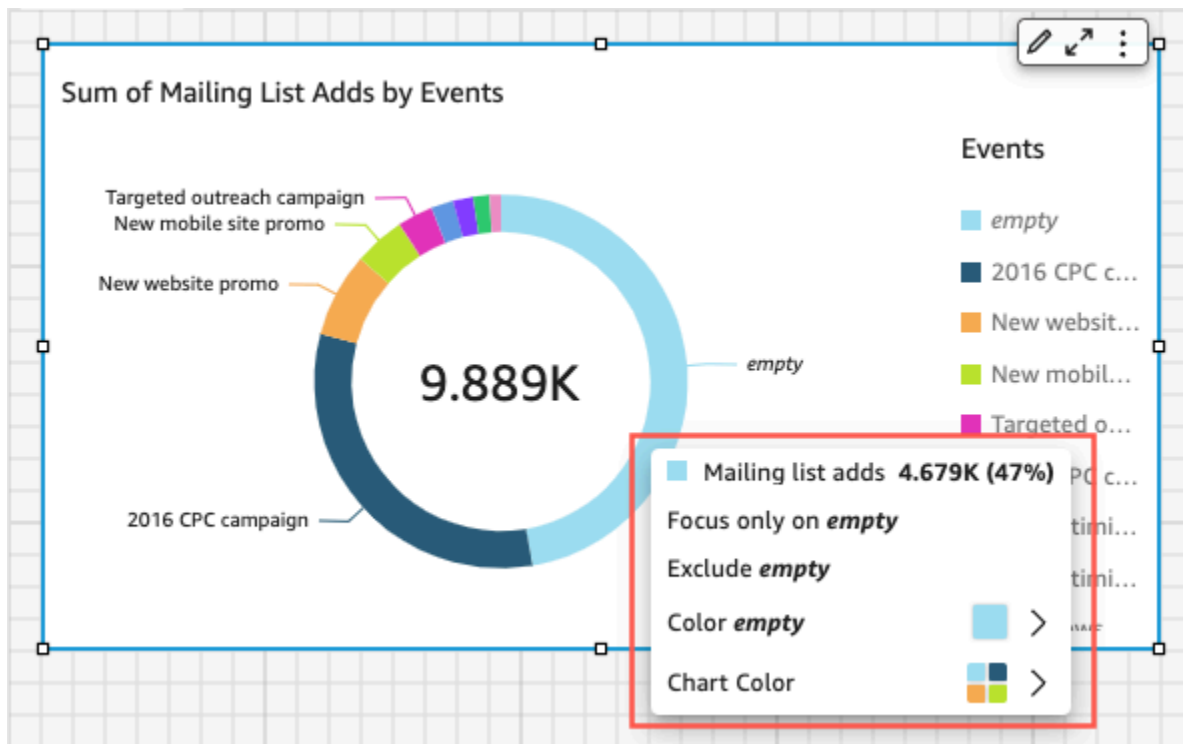
selection colors by clearing the **Show border**, **Show background**, and **Selection color** switches. You can also remove a visual's loading animation by clearing the **Show loading animation** switch.




Disabling visual menus

By using the **Interactions** panel of the **Format visual** menu, you can turn off the **Context** menu and **On-visual** menu from appearing on selected visuals. You can disable 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**.



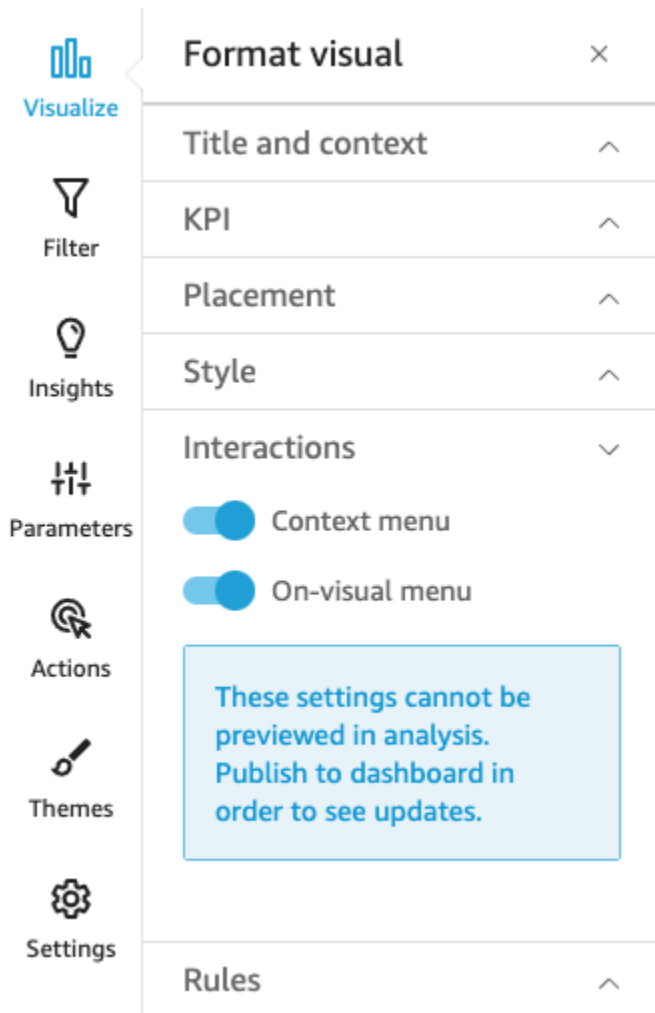
The **On-visual** menu appears on the top-right side of a visual. The **On-visual** menu is used to access the **Format visual** menu, **Maximize** the visual, access the **Menu options** panel, and review an **Anomaly insight**.



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**

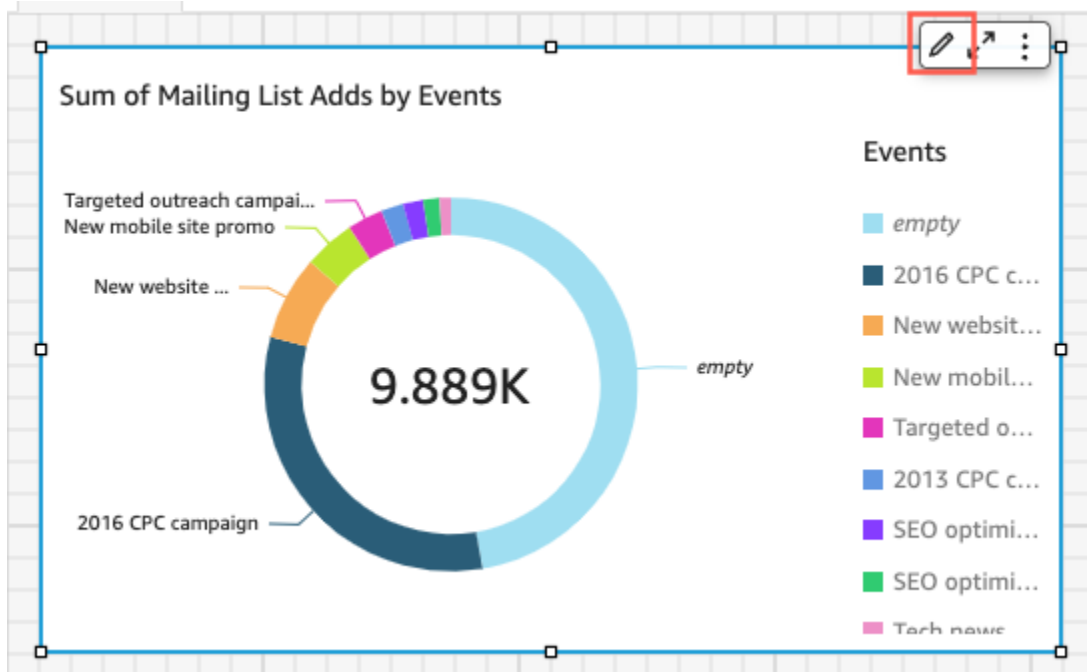
This feature is currently available with the **Free-form** layout. Conditional rules are used to hide or show visuals when specific conditions are met. This can be useful when you have multiple versions of the same visual overlapped with each other and want the dashboard viewer to see a version that 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 QuickSight](#).

Hiding a visual by default

In the **Rules** section of the **Format visual** menu, 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.



Format visual	×
Title	⤴
Donut chart	⤴
Group/Color	⤴
Legend	⤴
Data labels	⤴
Placement	⤴
Tooltip	⤴
Style	⤴
Dashboard interactions	⤴
Rules	⤵
Conditional rules	
Render this visual based on parameters.	
<input type="checkbox"/> Hide this visual by default	
Add Rule	

To hide a visual by default

1. From the QuickStart start page, 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 **Format visual**.
4. In the **Format visual** menu, choose **Rules**.
5. In the **Rules** pane, choose **Hide this visual by default**.

The screenshot shows the 'Format visual' menu for a donut chart titled 'Sum of Mailing List Adds by Events'. The chart displays a total of 9.889K adds, broken down by event categories. The 'Rules' section is expanded, showing a conditional rule to 'Hide this visual by default'.

Field wells: Group/Color (Events), Value (Mailing list adds (Sum))

Visual Title: Sum of Mailing List Adds by Events

Visual Value: 9.889K

Events Legend:

- empty
- 2016 CPC c...
- New websit...
- New mobil...
- Targeted o...
- 2013 CPC c...
- SEO optimi...
- SEO optimi...
- Tech newe...

Conditional rules:

Render this visual based on parameters.

Hide this visual by default

Show if toggle (i) ⋮

Equals 2

[Add Rule](#)

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** pane of the **Format visual** menu.

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 QuickSight](#).

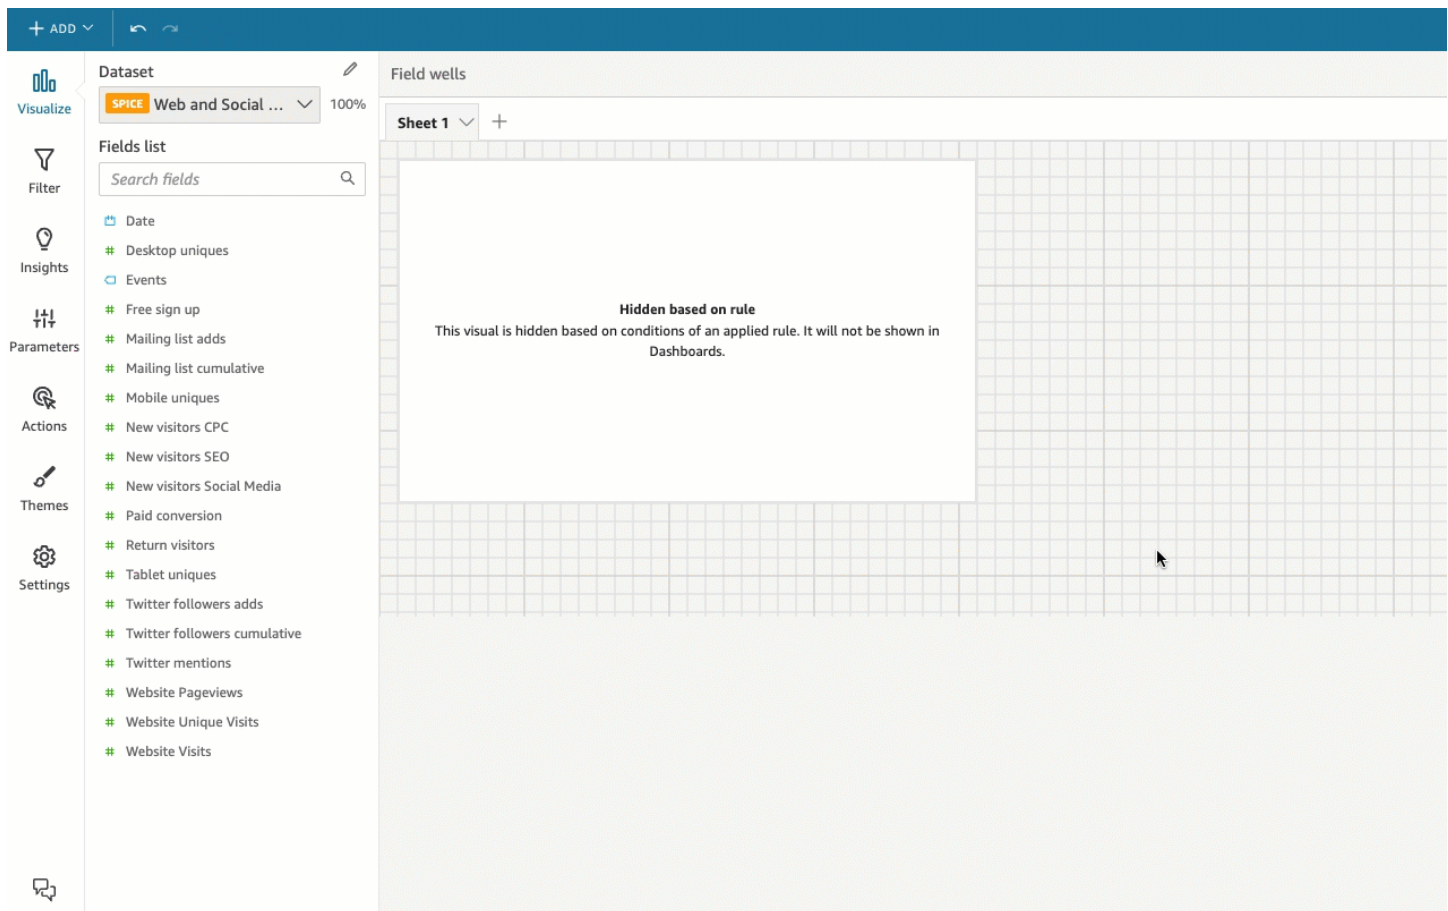
To set a conditional rule

1. From the Amazon QuickSight start page, 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 **Format visual**.
4. In the **Format visual** pane that appears on the left, choose **Rules**.
5. Choose the **Add** box
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 **Format visual**.
2. In the **Format visual** pane that appears on the left, 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 **Format visual**.
2. In the **Format visual** pane that appears on the left, 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 QuickStart start page, 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 QuickStart start page, 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 **Pin to top**.



Adding interactivity to dashboards in Amazon QuickSight

To add interactivity to your dashboards, you can create custom actions and parameters. See the following topics for more information.

Topics

- [Parameters in Amazon QuickSight](#)
- [Using custom actions for filtering and navigating](#)

Parameters in Amazon QuickSight

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 QuickSight](#)
- [Using a control with a parameter in Amazon QuickSight](#)
- [Creating parameter defaults in Amazon QuickSight](#)
- [Connecting to parameters in Amazon QuickSight](#)

Setting up parameters in Amazon QuickSight

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** pane from the left side of the screen.

3. Add a new parameter by choosing the plus sign (+) 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).

Default date

Rolling: Start of this day, 2022/06/03

 Fixed date Relative date Today Yesterday Filter condition

Range

Period

Start of

This

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 \(RLS\) 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 QuickSight

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 v 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 Amazon QuickSight](#).

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 QuickSight](#).

- **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 QuickSight

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 \(RLS\) 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 QuickSight](#).
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 QuickSight 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 QuickSight. 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](#) QuickSight 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/AIDAJ64EIEIOPX5CEIEI0"
  },
  {
    "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 QuickSight. To identify groups that you can add to the dataset, use one or more of the following QuickSight API operations or CLI commands:
 - [ListGroups](#) – Lists QuickSight 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 QuickSight group.
 - [ListUserGroups](#) – Lists the QuickSight groups that a QuickSight 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 QuickSight 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 QuickSight, 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 multivalued DDP

1. Create dataset columns with multivalued 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 QuickSight Enterprise edition.
2. For each field that displays a dynamic default for a multivalued 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, multivalued 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 multivalued 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 multivalued 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 multivalued 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 QuickSight, 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 QuickSight console, open the **Parameters** menu at left 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 QuickSight](#)

Amazon QuickSight 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 QuickSight

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 choose your next step from the shortcuts on this screen.

Parameter added ×

Connect your parameter:


Create a filter, using the combination of parameters, new control, and a filter.


Create a new control for a filter or a calculated field.


Use a parameter in a calculated field.


Create a URL action with parameters.

Close

 Filter


 Control

 Calculated field

 Custom actions

You can navigate to each of these options in another way, as follows:

- To create a filter, choose **Filter** to the left of the screen. 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 **Parameters** on the left. 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 QuickSight](#)
- [Using calculated fields with parameters in Amazon QuickSight](#)
- [Using custom actions with parameters in Amazon QuickSight](#)
- [Using parameters in a URL](#)
- [Using parameters in titles and descriptions in Amazon QuickSight](#)

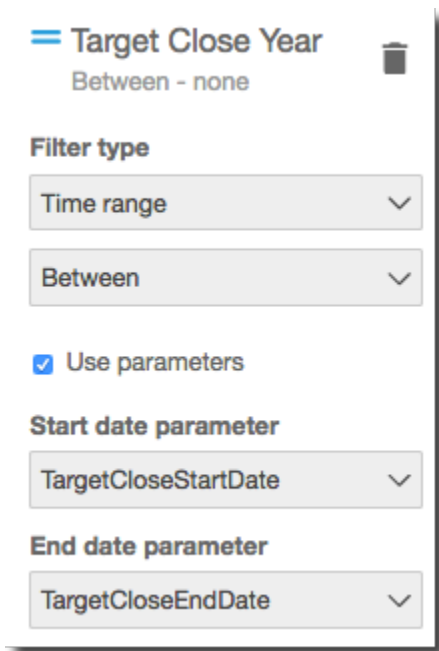
Using filters with parameters in Amazon QuickSight

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.

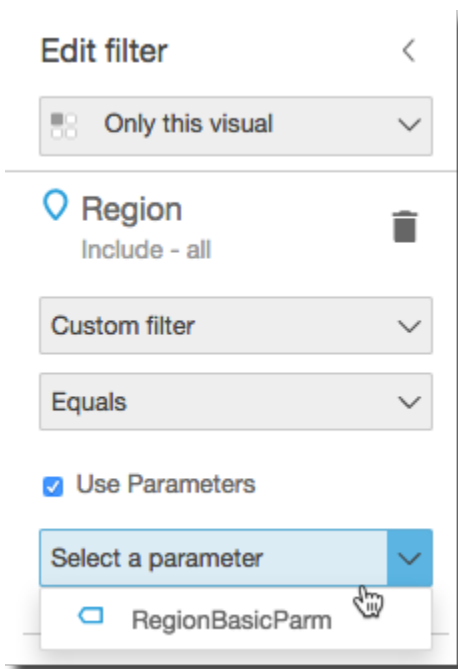
Before using a filter with a parameter, you should already know how to work with filters.

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 QuickSight

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.

The screenshot shows the 'New calculated field' dialog box. It features three columns: 'Function list' on the left with various mathematical and date functions; 'Field list' in the middle with fields like 'Active Opportunity', 'Closed Opportunity', 'Date', 'Forecasted Mont...', 'Latest Status Entry', and 'Lead Name'; and 'Parameter list' at the bottom with 'Year' and 'Region'. A red arrow points from the 'Parameter list' section towards the 'Formula' input field on the right. The 'Formula' field is currently empty and has a placeholder text 'Enter a formula...'. Above it is a 'Calculated field name' field with a placeholder 'Enter a field name...'. A blue 'Create' button is located at the bottom right of the dialog.

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 QuickSight

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 QuickSight](#). 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 QuickSight 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]".
 - 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](#).

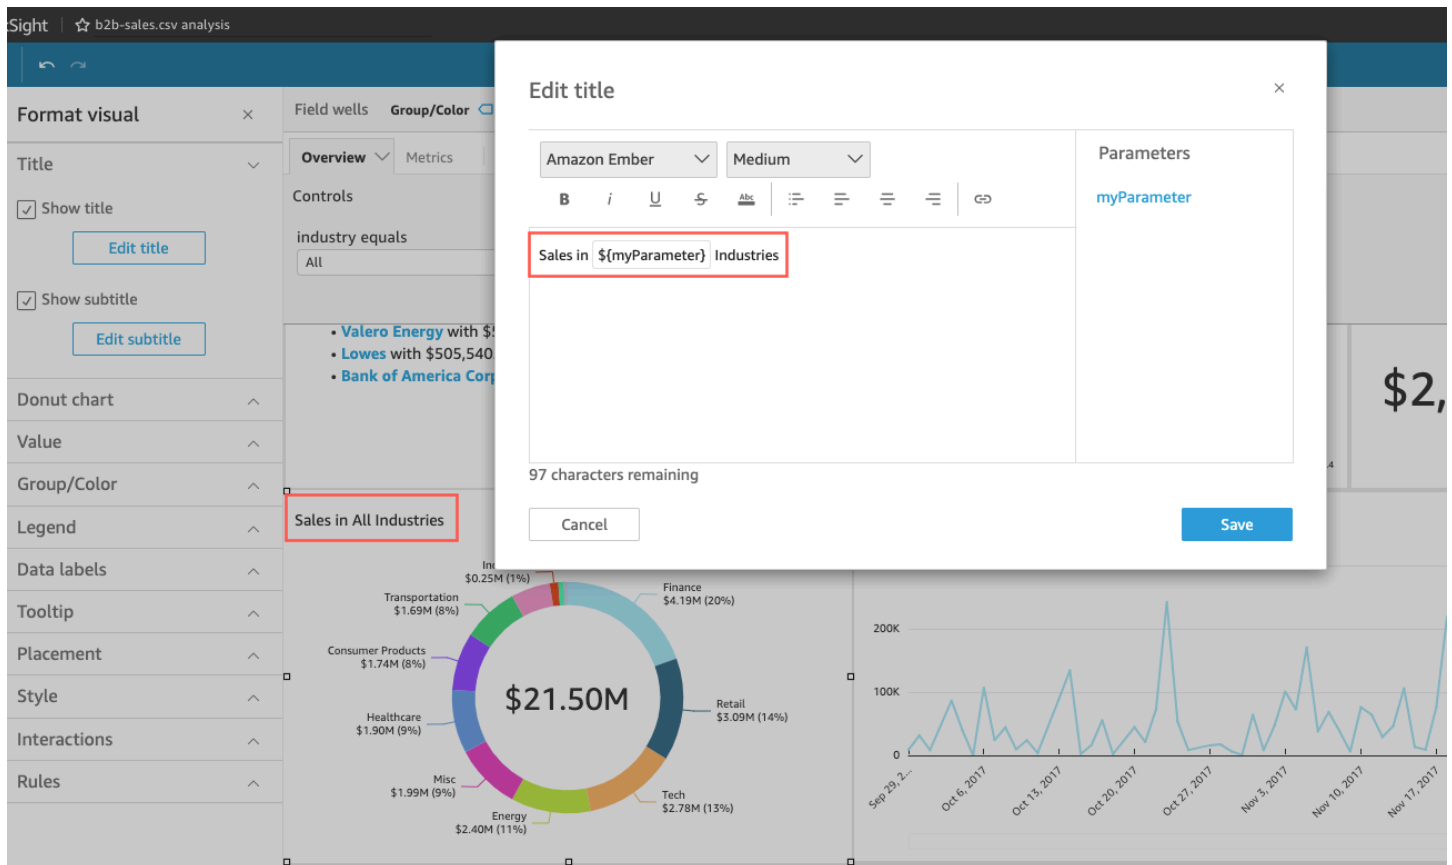
Using parameters in titles and descriptions in Amazon QuickSight

When you create parameters in Amazon QuickSight, 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 **Format visual** pane for the visual that you want to format.
2. In the **Format visual** 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 **Format visual** 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 QuickSight in QuickSight](#).


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 **Format visual** pane for the visual that you want to format.
2. In the **Format visual** 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 **Format visual** pane. In the chart, the parameter value is displayed in the axis title.

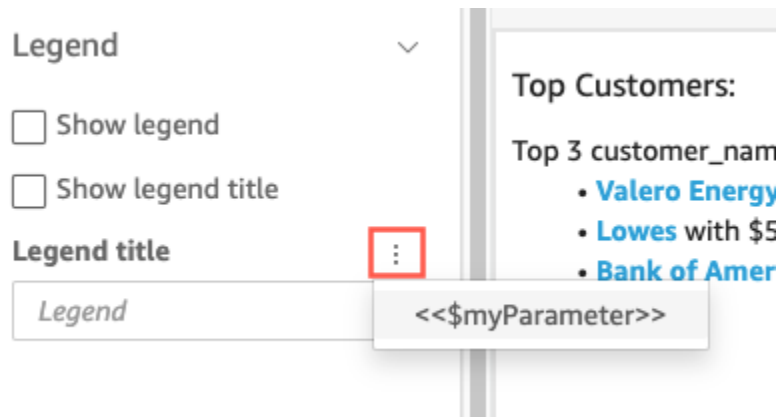
For more information about editing axis titles, see [Axes and grid lines](#).

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 **Format visual** pane for the visual that you want to format.
2. In the **Format visual** 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 **Format visual** pane. In the chart, the parameter value is displayed in the legend title.

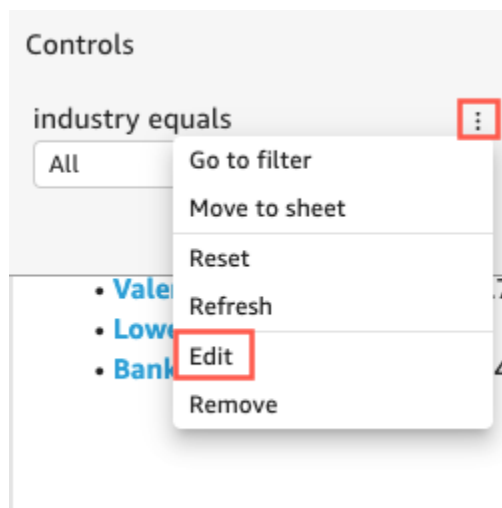
For more information about formatting legends, see [Legends on visual types in QuickSight](#).

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 **Add** 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 (QuickSight 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 QuickSight](#)
- [Repairing custom actions](#)
- [Understanding field mapping for custom actions in Amazon QuickSight](#)

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** at left.
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 QuickSight

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
	URL action

Action to perform	Type of action
Send an email (mailto)	

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



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 QuickSight 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



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** at left.

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



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



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  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 QuickSight 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**.
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



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 QuickSight

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.

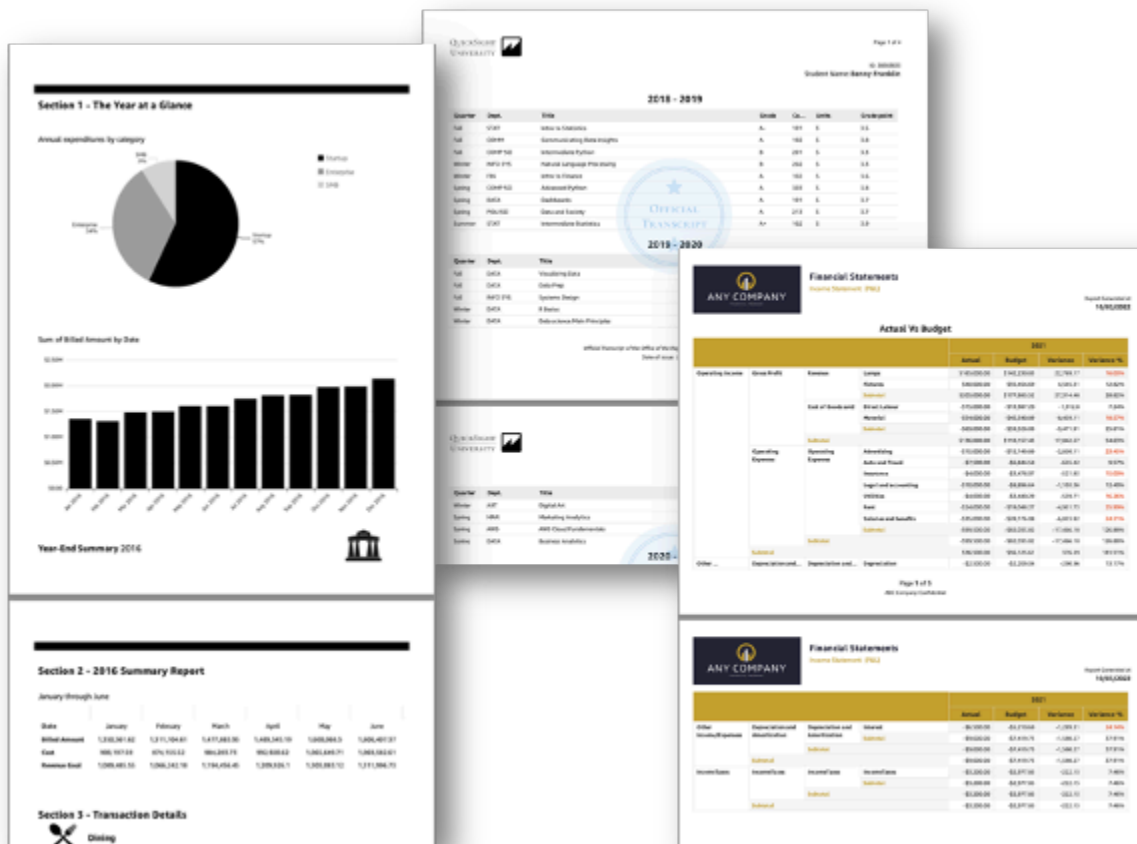
The new visuals won't respond to filter actions, because there is no field mapping to connect them. When you make changes, remember to save your custom actions again to redo the field mappings.

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 QuickSight

With Amazon QuickSight Paginated Reports, you can create, schedule, and share highly formatted multipage PDF reports. You can also schedule data exports as CSV files using QuickSight's existing web interface. This unifies historically separate systems for dashboards and reports.



Report creators can use QuickSight'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 QuickSight'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.

Topics

- [Getting started](#)
- [Creating reports from an analysis in Amazon QuickSight](#)
- [Formatting reports in Amazon QuickSight](#)
- [Consuming paginated reports in Amazon QuickSight](#)

Getting started

To get started working with Amazon QuickSight paginated reports, first get the paginated reporting add-on for your QuickSight account. Pricing for the add-on applies to your entire QuickSight account and isn't specific to a Region. After you subscribe to QuickSight reporting, authors can begin creating, scheduling, and sending paginated reports.

For more information about pricing for paginated reporting in Amazon QuickSight, see [Amazon QuickSight Pricing](#).

Get the QuickSight paginated reports add-on

Before you can work with paginated reports in Amazon QuickSight, you must add the **Paginated Reports add-on** to your QuickSight subscription.

To get the paginated reporting add-on in Amazon QuickSight

1. On the QuickSight start page, choose your user name at the upper right, and then choose **Manage QuickSight**.
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 QuickSight.

Unsubscribe from paginated reporting in QuickSight

You can unsubscribe from QuickSight paginated reporting at any time. Once you unsubscribe from paginated reporting, you will lose the ability to create and schedule paginated reports in QuickSight. 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 Amazon QuickSight

1. From any page in Amazon QuickSight, choose your user name in the upper right, and choose **Manage QuickSight**.
2. Choose **Your subscriptions** on the left.
3. On the **Your subscriptions** page, scroll down to the **QuickSight Paginated Reports add-on** section and choose **Manage**.
4. Scroll down to your chosen subscription plan and choose **Cancel subscription**.

Creating reports from an analysis in Amazon QuickSight

Creating reports from an analysis in Amazon QuickSight

Paginated reports are created at the sheet level of an analysis in Amazon QuickSight. 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.

To create a paginated report from a new analysis

1. On the QuickSight 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 **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 QuickSight

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 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 QuickSight

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 name of the sheet that you want to convert.

2. Choose **Duplicate**.



Formatting reports in Amazon QuickSight

Use this section to learn how to format a paginated report in Amazon QuickSight.

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](#)

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 QuickSight 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.

Adding, moving, and deleting sections

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 QuickSight start menu, 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**.

The screenshot displays a paginated report interface. The top section, labeled 'SECTION 4', contains a table titled 'Late Invoice Details'. The table has five columns: 'invoiceNumber', 'InvoiceDate', 'DueDate', 'DaysLate', and 'InvoiceAmount'. Below the table, a dashed line separates it from 'SECTION 5'. In Section 5, a menu is visible with two options: 'Remove page break' and 'Add section'. The 'Add section' option is highlighted with a red box.

invoiceNumber	InvoiceDate	DueDate	DaysLate	InvoiceAmount
9,632,048,192	Jul 9, 2012	Aug 8, 2012	15	\$128.28
857,712,918	Feb 22, 2013	Mar 24, 2013	11	\$93.39
9,390,786,866	Feb 24, 2013	Mar 26, 2013	11	\$74.62
915,652,542	Jan 5, 2012	Feb 4, 2012	9	\$78.29
246,081,324	Sep 6, 2012	Oct 6, 2012	7	\$92.53
106,360,977	Feb 19, 2012	Mar 20, 2012	6	\$93.48

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.

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.

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**.

Adding and deleting 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 QuickSight 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 the **ADD (+)** icon in the top left.
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.

Changing paper size, margins, and orientation

After you create a paginated report in Amazon QuickSight, you can change the report format, orientation, and margins from the **Settings** menu whenever you want.

To change the paper size of a paginated report

1. From the QuickSight start menu, choose **Analyses**, and then choose the analysis that contains the paginated report that you want to change.
2. Choose the **Settings** on the left.
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 QuickSight start menu, 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**.

Changing the page margins of a report

You can customize the margins of a paginated report in the report **Settings** menu. The 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.

To change the margins of a report

1. From the QuickSight start menu, choose **Analyses**, and then choose the analysis that contains the paginated report that you want to change.
2. Choose the **Settings** on the left.
3. Enter the margin values that you want your report to have, and then choose **Apply**.

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 right 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**.

The screenshot shows a report with two sections. SECTION 4, titled 'Late Invoice Details', contains a table with the following data:

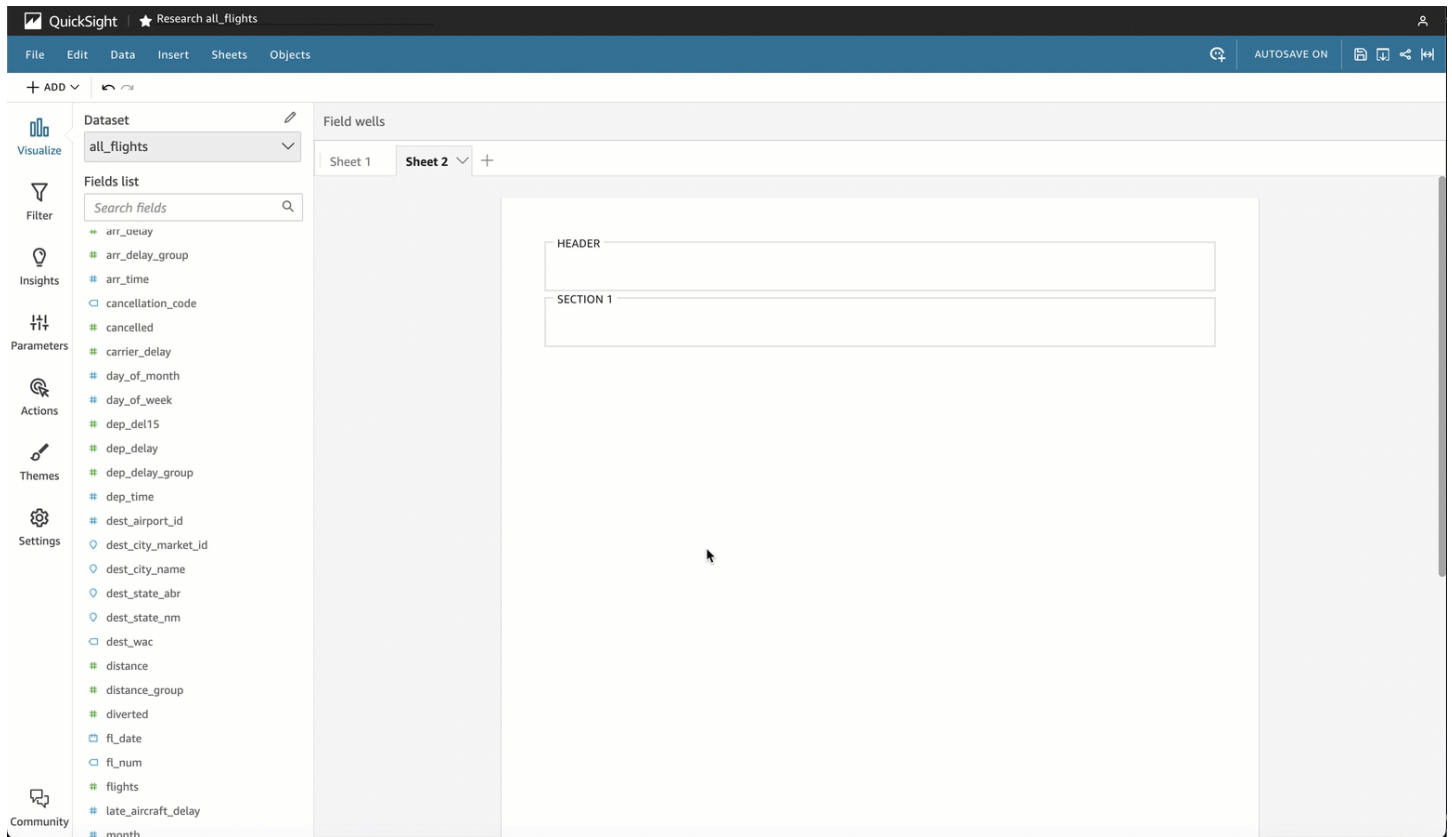
invoiceNumber	InvoiceDate	DueDate	DaysLate	InvoiceAmount
9,632,048,192	Jul 9, 2012	Aug 8, 2012	15	\$128.28
857,712,918	Feb 22, 2013	Mar 24, 2013	11	\$93.39
9,390,786,866	Feb 24, 2013	Mar 26, 2013	11	\$74.62
915,652,542	Jan 5, 2012	Feb 4, 2012	9	\$78.29
246,081,324	Sep 6, 2012	Oct 6, 2012	7	\$92.53
106,360,977	Feb 19, 2012	Mar 20, 2012	6	\$93.48

Below the table, SECTION 5 is partially visible with the text 'Disputed Am'. A context menu is open over SECTION 5, showing two options: 'Remove page break' (highlighted with a red box) and 'Add section'.

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 the upper left corner.
3. Choose **Add visual**.



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 QuickSight 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 QuickSight, see [Formatting in Amazon QuickSight](#).

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 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 **ADD (+)** icon in the upper left corner.
3. Choose **Add text box**.

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.

The screenshot displays the Amazon QuickSight interface. The top navigation bar includes 'QuickSight', 'Research all_flights', and a menu with 'File', 'Edit', 'Data', 'Insert', 'Sheets', and 'Objects'. Below the navigation bar is a toolbar with '+ ADD', undo, and redo icons. The left sidebar contains various tool categories: Visualize, Filter, Insights, Parameters, Actions, Themes, Settings, and Community. The main workspace shows a report with a donut chart titled 'Sum of Flights by Dest_city_name' and a text box. The donut chart is labeled '39.62M' and 'SHOWING TOP 20 IN DEST_CITY_NAME'. The text box is titled 'SECTION 1' and contains the text 'Dest City N...'. The chart and text box are both selected, indicated by a black border and a small toolbar in the top right corner of the text box.

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.

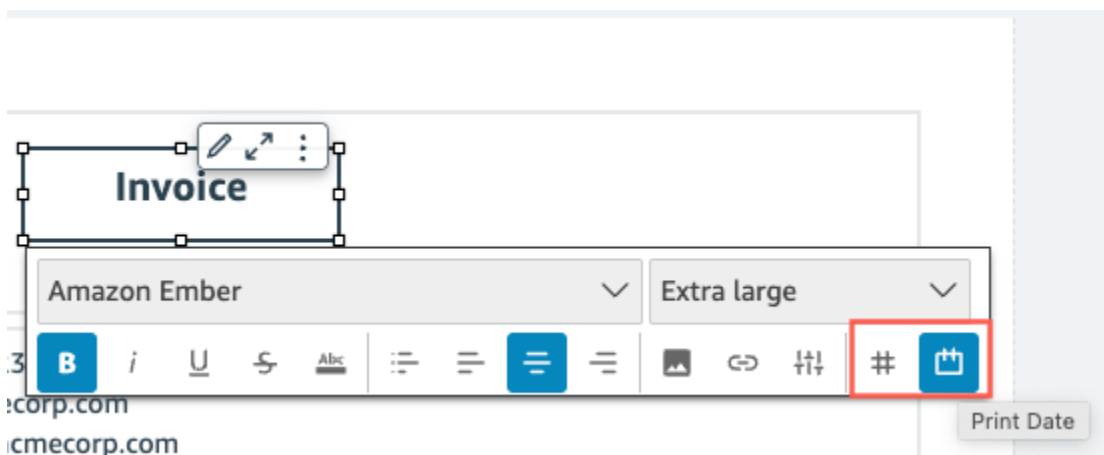
2. Choose the three-dot icon in the upper right corner of the visual to open the on-text box menu.
3. Choose **Delete**.

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.

Consuming paginated reports in Amazon QuickSight

When you publish and then send a scheduled paginated report, Amazon QuickSight 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 reports by email](#).



Schedule to generate content

Create a schedule and choose to generate a PDF of this sheet. Your content will appear here once available.

[ADD SCHEDULE](#)

[GENERATE PDF](#)

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 QuickSight saves a copy of the generated snapshot that is sent out for your reference. You can view these snapshots at any time in the QuickSight console.

To view a report's snapshot history

1. From the QuickSight start page, 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.

Working with items on sheets in Amazon QuickSight analyses

Use this section to learn how to work with visuals and other items as you author sheets in Amazon QuickSight

Topics

- [Adding visuals to Amazon QuickSight analyses](#)
- [Using Q Topics on sheets in Amazon QuickSight](#)
- [Visual types in Amazon QuickSight](#)
- [Formatting in Amazon QuickSight](#)
- [Customizing data presentation](#)

Adding visuals to Amazon QuickSight 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.

Amazon QuickSight supports up to 50 datasets in a single analysis, and up to 30 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 QuickSight 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 QuickSight 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 **Fields 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 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 QuickSight](#) 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 **Fields list** 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. On the Amazon QuickSight start page, choose the analysis that you want to add a visual to.
2. On the analysis page, choose the dataset that you want to use from the dataset list at the top of the **Fields list** pane. For more information, see [Adding a dataset to an analysis](#).
3. Choose **Add** on the application bar, and then choose **Add visual**.

A new, blank visual is created and receives focus.


4. Use one of the following options:
 - Choose the fields to use from the **Fields list** pane at left. If the Fields list isn't visible, choose **Visualize** to display it. Amazon QuickSight creates the visual, using the visual type it determines is most compatible with the data you selected.
 - Create a visual by choosing a visual type and then choosing fields 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.



Click anywhere on the field wells to open them.



- 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 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 QuickSight 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.

Duplicating Amazon QuickSight 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 QuickSight 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 QuickSight

Amazon QuickSight 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 Amazon QuickSight](#)

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 Amazon QuickSight 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


The following screenshot shows how to use the legend for focusing on, or excluding, a dimension.



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 QuickSight](#).


Searching for specific values in your data in Amazon QuickSight

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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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.
- QuickSight 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 QuickSight. 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 Amazon QuickSight

When you work in an Amazon QuickSight 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 QuickSight, open the analysis.
2. In the analysis, choose **Settings** at left.
3. In the **Settings** pane that opens, for **Refresh Options**, toggle on **Reload visuals each time I switch sheets**.
4. Choose **Apply**.

Deleting Amazon QuickSight 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 Q Topics on sheets in Amazon QuickSight

QuickSight 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 Amazon QuickSight Q topics in your analysis workspace, you activate the ML-powered automated data prep for Q, which speeds Natural Language (NL) topic creation. Automated data prep for Q 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 for Q 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 QuickSight Q topic is automatically selected, making it easier for them to query the dataset.

The following rules apply to working with Q 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 Q topic

1. Open the analysis that you want to use with automated data prep for Q.
2. On the top navigation bar, choose the Q topic icon.



3. Choose one of the following:
 - To activate a new Q 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 Q 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 Q topic for an analysis or dashboard, you are starting a process where automated data prep for Q 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 Q topic, the better prepared it becomes to answer your questions.

To learn more about interacting with Q in an analysis, see <https://docs.amazonaws.cn/quicksight/latest/user/quicksight-q-starting-from-sheets.html>.

Visual types in Amazon QuickSight

Amazon QuickSight 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 histograms](#)
- [Using KPIs](#)
- [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 QuickSight 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 QuickSight 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 QuickSight](#).

For dashboards and analyses, Amazon QuickSight 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.

The following image shows the **other** category on a bar chart.

Sum of Weighted Revenue by Opportunity Stage and Salesperson

SHOWING TOP 3 IN OPPORTUNITY STAGE AND TOP 5 IN SALESPERSON



The following image shows the **other** category on a pivot table.

Sum of Weighted Revenue by Salesperson and Opportunity...

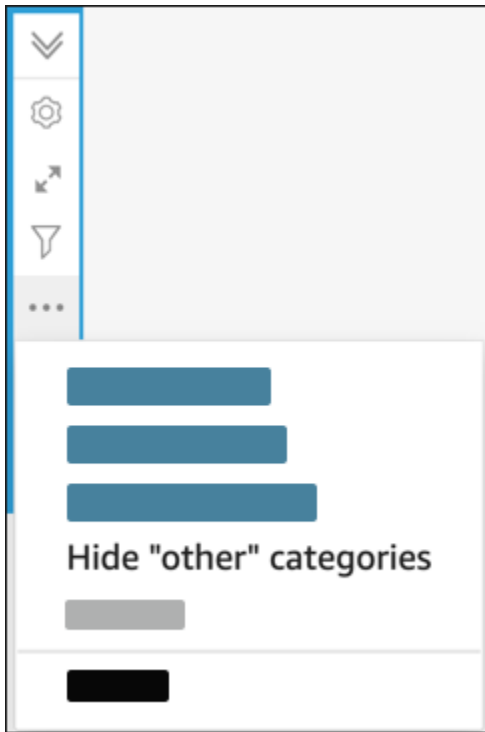
Salesperson	Opportunity Stage	Weighted Revenue
Anne Smith	Closed Won	26,768,347
	Contracting	24,283,888
	Prospect	18,952,686
David King	Closed Won	20,797,243
	Contracting	21,273,660
Other		297,131,502

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	2,500
Bar chart, vertical	X-axis – Number of data points displayed	2,500
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 QuickSight to choose the visual type for you. When you create a visual by choosing AutoGraph and then selecting fields, Amazon QuickSight uses the most appropriate visual type for the number and data types of the fields you select.

The icon for AutoGraph is as follows.



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 QuickSight 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. For example, the revenue goal per region, as shown in the following image.

Single Measure Bar Chart (horizontal)




- **Multi-measure** – A *multi-measure bar chart* shows values for multiple measure for a dimension. For example, the revenue goal and cost per region, as shown in the following image.

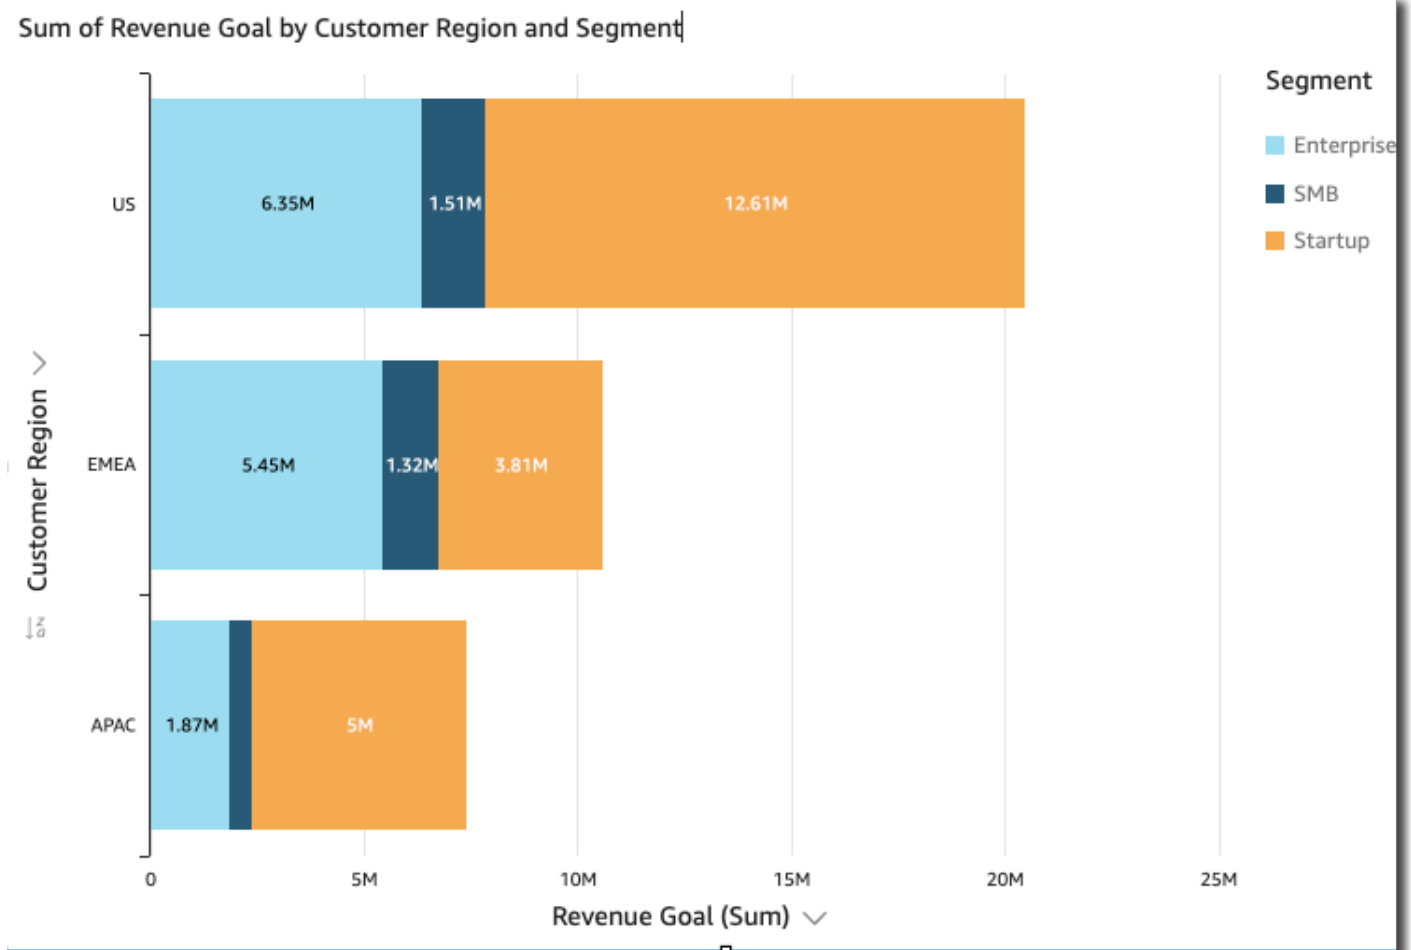
Multi-Measure Bar Chart (vertical)



- **Clustered** – A *clustered bar chart* shows values for a single measure for a dimension, grouped by another dimension. For example, the cost for each channel in a region, as shown in the following image.




- **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. For example, the revenue goal for each segment by region, as shown in the following image.



- **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. For example, the percent each segment costs per region, as shown in the following image.

Stacked 100% Bar Chart (horizontal)



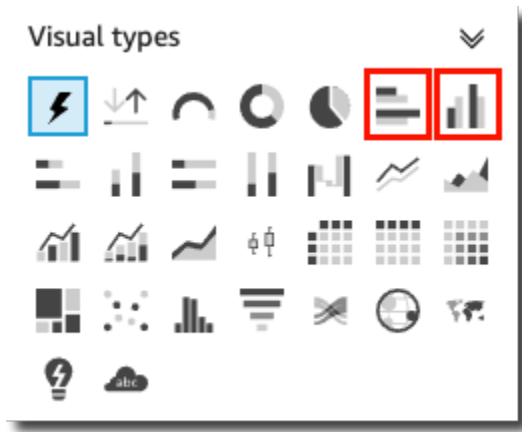
Bar charts show up to 2,500 data points on the axis for visuals that don't use group or color. For visuals that do use group or color, they show up to 50 data points on the axis and up to 50 data points for group or color. For more information about how Amazon QuickSight handles data that falls outside 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, as shown following.



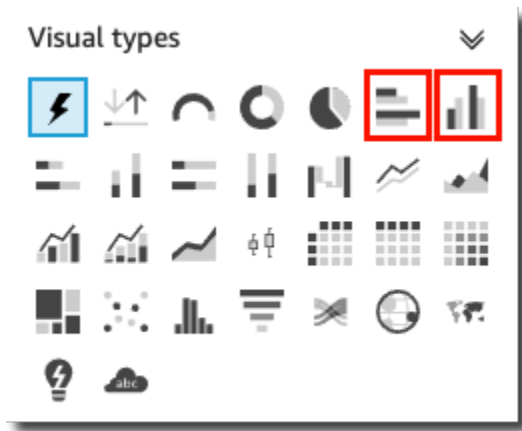
4. From the **Fields list** pane, drag a dimension to the **X-axis** or **Y-axis** field well.
5. From the **Fields list** pane, drag a measure to the **Value** field well.

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, as shown following.



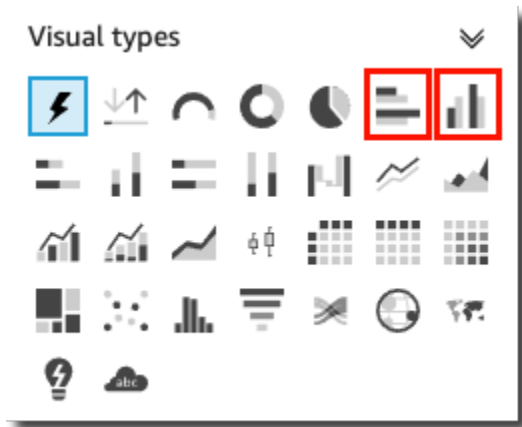
4. From the **Fields list** pane, drag a dimension to the **X-axis** or **Y-axis** field well.
5. From the **Fields list** pane, drag two or more measures to the **Value** field well.

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, as shown following.



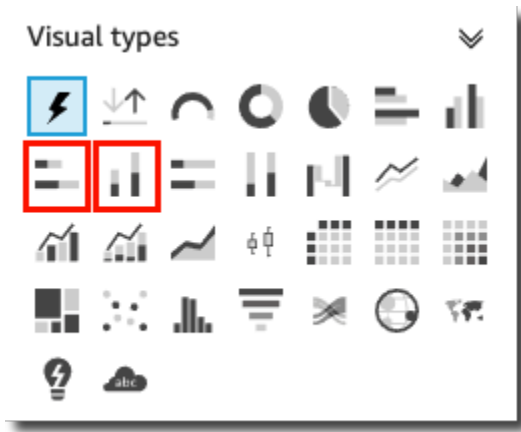
4. From the **Fields list** pane, drag a dimension to the **X-axis** or **Y-axis** field well.
5. From the **Fields list** pane, drag a measure to the **Value** field well.
6. From the **Fields list** pane, drag a dimension to the **Group/Color** field well.

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, as shown following.



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 **Format visual** pane that opens at left, choose **Data labels**.
 - c. Select **Show data labels**.

Labels for each measure value appear in the chart and the option to show totals appears in the pane.

- d. Choose **Show totals**.

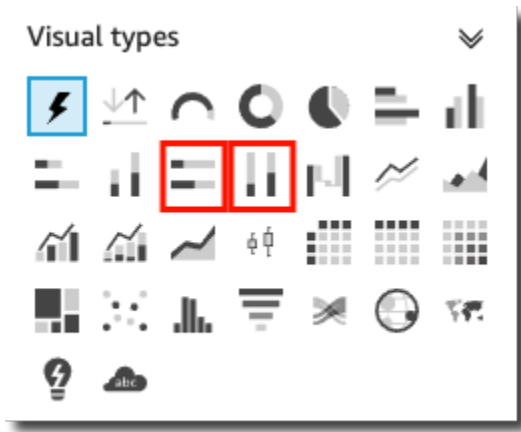
Totals appear for each bar in the chart.

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, as shown following.



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 QuickSight
Changing the title display	Yes		Titles and subtitles on visual types in QuickSight in QuickSight
Changing the axis range	Yes		Range and scale on visual types in QuickSight
Showing or hiding axis lines, grid lines, axis labels, and axis sort icons	Yes		Axes and grid lines on visual types in QuickSight

Feature	Supported?	Comments	For more information
Changing the visual colors	Yes		Colors in visual types in QuickSight
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 QuickSight
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 Amazon QuickSight
Showing data labels	Yes		Data labels on visual types in QuickSight
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.

The following screenshot shows a box plot.



The icon for a box plot is as follows.



To create a basic box plot visual

1. Sign in to Amazon QuickSight at <https://quicksight.aws.amazon.com/>.
2. Open Amazon QuickSight 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 Amazon QuickSight](#).
 - 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 Amazon QuickSight](#).

To understand the features supported by box plots, see [Analytics formatting per type in QuickSight](#). For customization options, see [Formatting in Amazon QuickSight](#).

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 QuickSight supports the following types of combo charts:

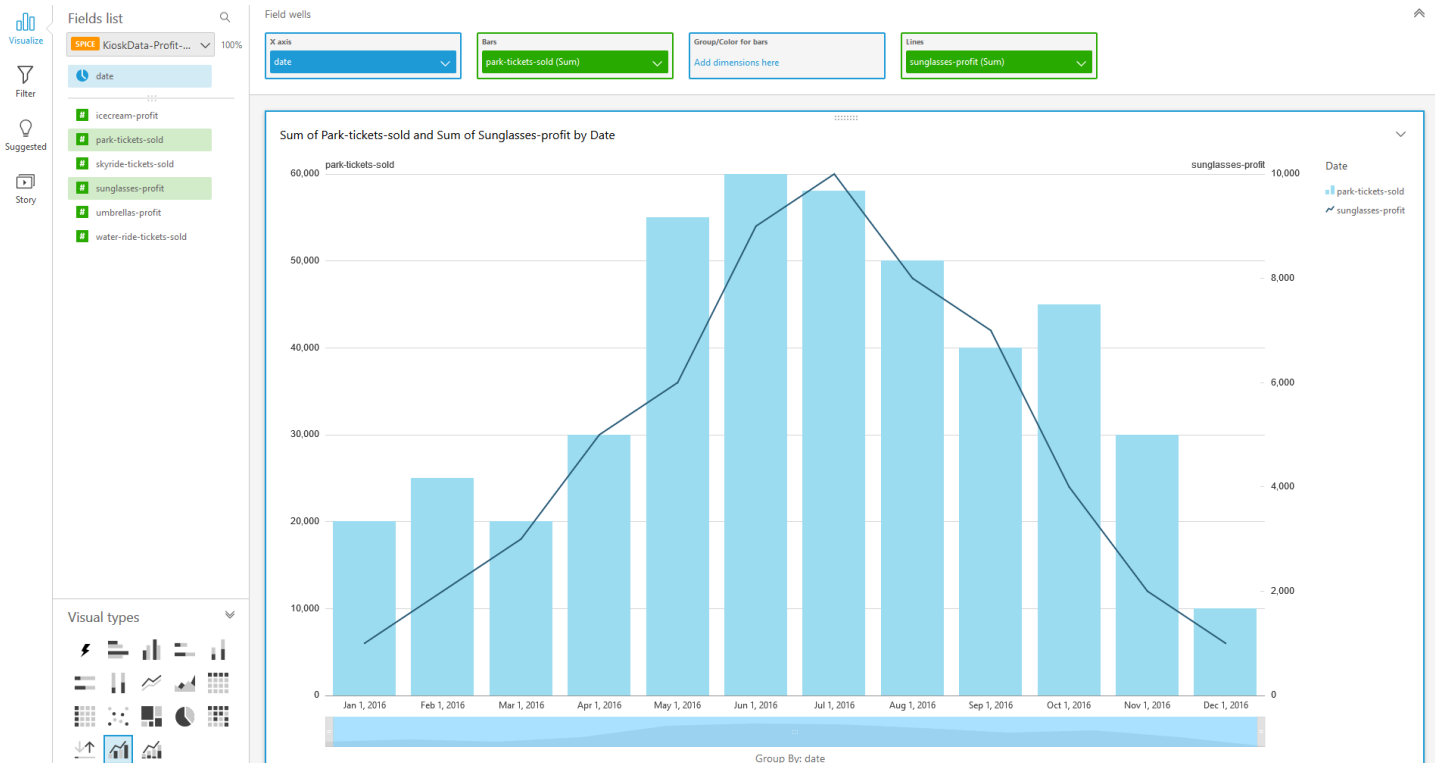
- **Clustered bar combo charts** – display sets of single-color bars where each set represents a parent dimension and each bar represents a child dimension. Use this chart to make it easy to determine values for each bar.
- **Stacked bar combo charts** – display multi-color bars where each bar represents 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 QuickSight handles data that falls outside display limits, see [Display limits](#).

The icons for combo charts are as follows.



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 QuickSight
Changing the title display	Yes		Titles and subtitles on visual types in QuickSight in QuickSight
Changing the axis range	Yes	You can set the range for the axis.	Range and scale on visual types in QuickSight
Showing or hiding axis lines,	Yes		Axes and grid lines on visual types in QuickSight


Feature	Supported?	Comments	For more information
grid lines, axis labels, and axis sort icons			
Changing the visual colors	Yes		Colors in visual types in QuickSight
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 QuickSight
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 Amazon QuickSight
Synchronizing y-axis	Yes	Synchronize the y-axes for both bars and lines into a single axis.	Range and scale on visual types in QuickSight

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 Amazon QuickSight](#).



Using custom visual content

You can embed webpages and online videos, forms, and images in your Amazon QuickSight 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, as shown in the following image.



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 QuickSight, 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. Blocking third party cookies results in images not loading in QuickSight.

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 **Format visual** 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 **Format visual** 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.

The screenshot shows the 'Format visual' pane on the left and the 'Field wells' pane on the right. The 'Format visual' pane has a 'Custom Content' section with a 'URL' field containing the embeddable URL: <https://docs.google.com/forms/d/e/1FAIpQLSf4KtE5m>. Below the URL field is an 'Apply' button and a checkbox labeled 'Show as image'. The 'Field wells' pane shows a dashboard with a visual titled 'Help us improve our content'. The form content is displayed within this visual, including the title, a request for feedback, a 'Was this dashboard helpful?' question with 'Yes' and 'No' radio buttons, a 'Comments' section with a text input field, and a 'Submit' button. To the right of the form, another visual titled 'YTD Opportunities by Stage and Segment' is visible, showing a horizontal bar chart with segments for Startup, SMB, and Enterprise.

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 **Format visual** 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 **Format visual** 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.

The following screenshot shows an example of a donut chart.



To learn how to use donut charts in Amazon QuickSight, 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 QuickSight handles data that falls outside display limits, see [Display limits](#).

The icon for a donut chart is as follows.



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 QuickSight
Changing the title display	Yes		Titles and subtitles on visual types in QuickSight in QuickSight
Changing the axis range	Not applicable		Range and scale on visual types in QuickSight
Changing the visual colors	Yes		Colors in visual types in QuickSight
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 QuickSight
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 Amazon QuickSight

Feature	Supported?	Comments	For more information
Choosing size	Yes	You can choose how thick the donut chart is: small, medium, and large.	Formatting in Amazon QuickSight
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 QuickSight

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 Amazon QuickSight](#).

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.

The icon for a funnel chart is as follows.



To create a basic funnel chart visual

1. Open Amazon QuickSight 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 Amazon QuickSight](#).
 - 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 Amazon QuickSight](#).


To understand the features supported by funnel charts, see [Analytics formatting per type in QuickSight](#). For customization options, see [Formatting in Amazon QuickSight](#).

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.

The following screenshot shows an example of a gauge chart.



To learn how to use gauge charts in Amazon QuickSight, 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)

The icon for a gauge chart is as follows.



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 QuickSight
Changing the title display	Yes		Titles and subtitles on visual types in QuickSight in QuickSight
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		

Feature	Supported?	Comments	For more information
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 QuickSight
Focusing on or excluding elements	No		
Sorting	No		Sorting visual data in Amazon QuickSight
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 QuickSight 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.

For example, the following heat map shows which products are most used by the customers in these countries, measured by a simple count.

Customers by Product and Country



Heat maps show up to 50 data points for rows and up to 50 data points for columns. For more information about how Amazon QuickSight handles data that falls outside display limits, see [Display limits](#).

The icon for a heat map is as follows.



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 QuickSight

Feature	Supported?	Comments	For more information
Changing the title display	Yes		Titles and subtitles on visual types in QuickSight in QuickSight
Changing the axis range	Not applicable		Range and scale on visual types in QuickSight
Changing the visual colors	No		Colors in visual types in QuickSight
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
Sorting	Yes	You can sort by the fields you choose for the columns and the values.	Sorting visual data in Amazon QuickSight
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 Amazon QuickSight
Conditional formatting	No		Conditional formatting on visual types in QuickSight

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 Amazon QuickSight](#).

Using histograms

Use a histogram chart in Amazon QuickSight to display the distribution of continuous numerical values in your data. Amazon QuickSight 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 QuickSight handles data that falls outside display limits, see [Display limits](#).

The icon for histograms is as follows.



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 QuickSight

Feature	Supported?	Comments	For more information
Changing the title display	Yes		Titles and subtitles on visual types in QuickSight in QuickSight
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 QuickSight
Changing the visual colors	Yes		Colors in visual types in QuickSight
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 **Format visual** 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 QuickSight 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 QuickSight](#).

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. The following histogram displays a left skewed distribution.




- *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. The following histogram displays a right skewed distribution.




- 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. The following histogram displays a normal distribution.



- Bimodal – A chart that has two peaks representing the most common values. The following histogram displays a bimodal distribution.



- **Multimodal** – A chart that has three or more peaks representing the most common values. The following histogram displays a multimodal distribution.



- **Uniform** – A chart that has no peaks or valleys, with a relatively equal distribution of data. The following histogram displays a uniform distribution.



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.

Histogram	Bar chart
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.
If a line is included in a histogram, it represents 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 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.

Sum of Population by Region

South
84,673,176

West
67,928,970

16,744,206 ↑



The icon for a KPI is as follows.



KPI features

To understand the features supported by the KPI visual type in Amazon QuickSight, use the following table.

Feature	Supported?	Comments	For more information
Changing the title display	Yes		Titles and subtitles on visual types in QuickSight in QuickSight

Feature	Supported?	Comments	For more information
Removing the title	Yes	You can choose not to display a title.	
Changing comparison method	Yes	By default, Amazon QuickSight automatically chooses a method. The settings are auto, difference, percent, and difference as percent.	
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.

- (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

- Navigate to the KPI visual that you want to change and choose **KPI layouts**.
- In the **KPI Layouts** pane, choose the KPI layout that you want to use.

The following image shows the **KPI Layouts** menu.



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.

Sum of Sales (Sum) by Industry and Ship_region




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.

Sum of Sales (Sum) by Industry and Ship_region




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.

Sum of Sales (Sum) by Industry and Ship_region




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.

Sum of Sales (Sum) by Industry and Ship_region



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](#).

The icons for line charts are as follows.





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 QuickSight
Changing the title display	Yes		Titles and subtitles on visual types in QuickSight in QuickSight
Changing the axis range	Yes	You can set the range for the Y axis.	Range and scale on visual types in QuickSight
Showing or hiding axis lines, grid lines, axis labels, and axis sort icons	Yes		Axes and grid lines on visual types in QuickSight
Adding a second Y-axis	Yes		Creating a dual-axis line chart
Changing the visual colors	Yes		Colors in visual types in QuickSight
Focusing on or excluding elements	Yes, with exceptions	You can focus on or exclude any line on the chart, except in the following cases:	Focusing on visual elements

Feature	Supported?	Comments	For more information
		<ul style="list-style-type: none"> 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. <p>In these cases, you can only focus on a line, not exclude it.</p>	Excluding visual elements
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 QuickSight
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 Amazon QuickSight

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 Amazon QuickSight](#).

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.

Sum of Free Sign Up and Sum of Mailing List Adds by Date



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.


Sum of Free Sign Up and Sum of Mailing List Adds by Date



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**.

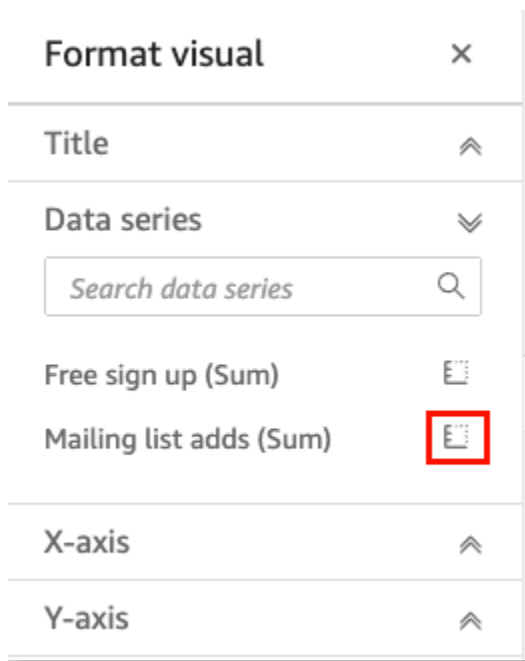


Or you can create a dual-axis line chart using the **Format Visual** pane:

- a. On the menu in the upper-right corner of the line chart, choose the **Format visual** icon.



- b. In the **Format visual** pane that opens at left, 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 **Format visual** 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 **Format visual** 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 Amazon QuickSight: 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 Amazon QuickSight 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. QuickSight 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 QuickSight, see the following.

Topics

- [Creating point maps](#)
- [Creating filled maps](#)
- [Interacting with maps](#)

Creating point maps

You can create point maps in Amazon QuickSight 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 QuickSight.

To create point maps in QuickSight, make sure that you have the following:

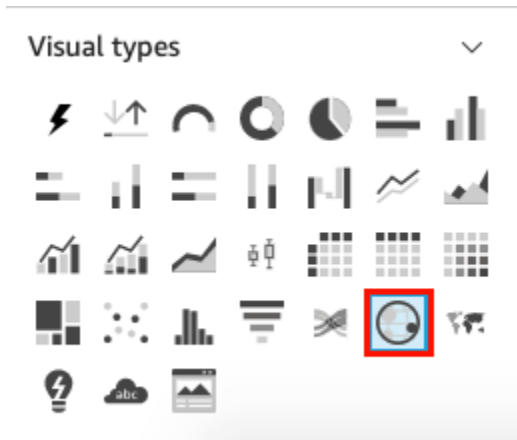
- One geospatial field (such as country, state or region, county or district, city, or ZIP code or postal code). Or you can use one latitude field and one longitude field.
- One numeric field (measure) for size.
- (Optional) A categorical field (dimension) for color.

For information on formatting geospatial maps, see [Map and geospatial chart formatting options](#).

Creating point maps

To create a point map

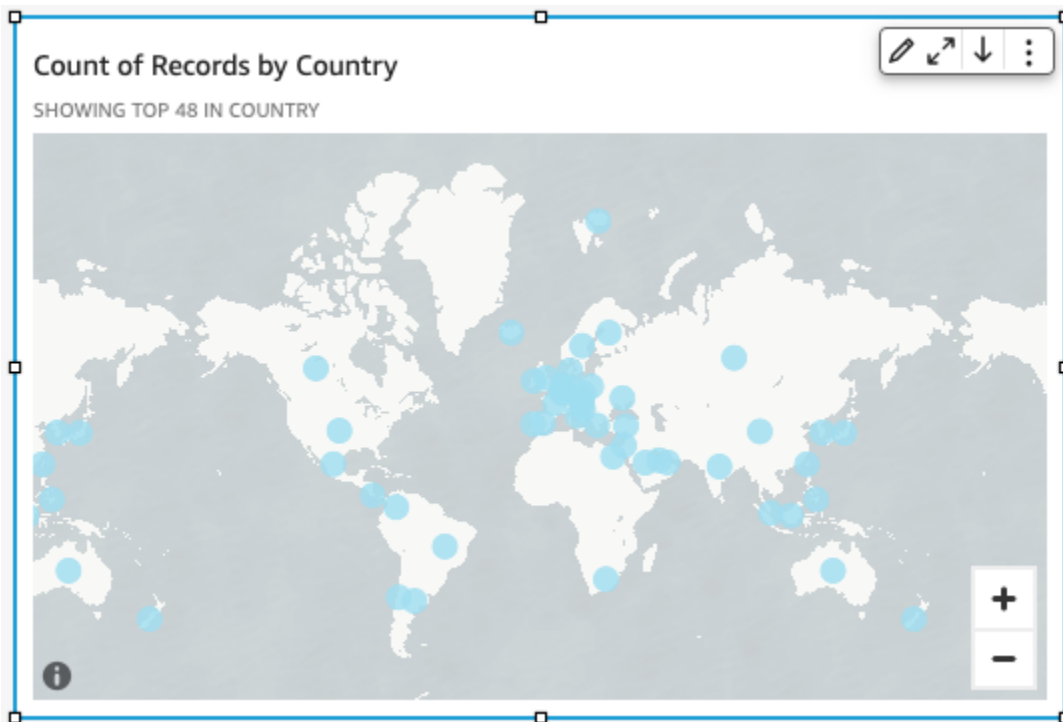
1. Add a new visual to your analysis. For more information about starting analyses, see [Starting an analysis in Amazon QuickSight](#). For more information about adding visuals to analyses, see [Adding a visual](#).
2. For **Visual type**, choose the **Points on map** icon. It looks like a globe with 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.

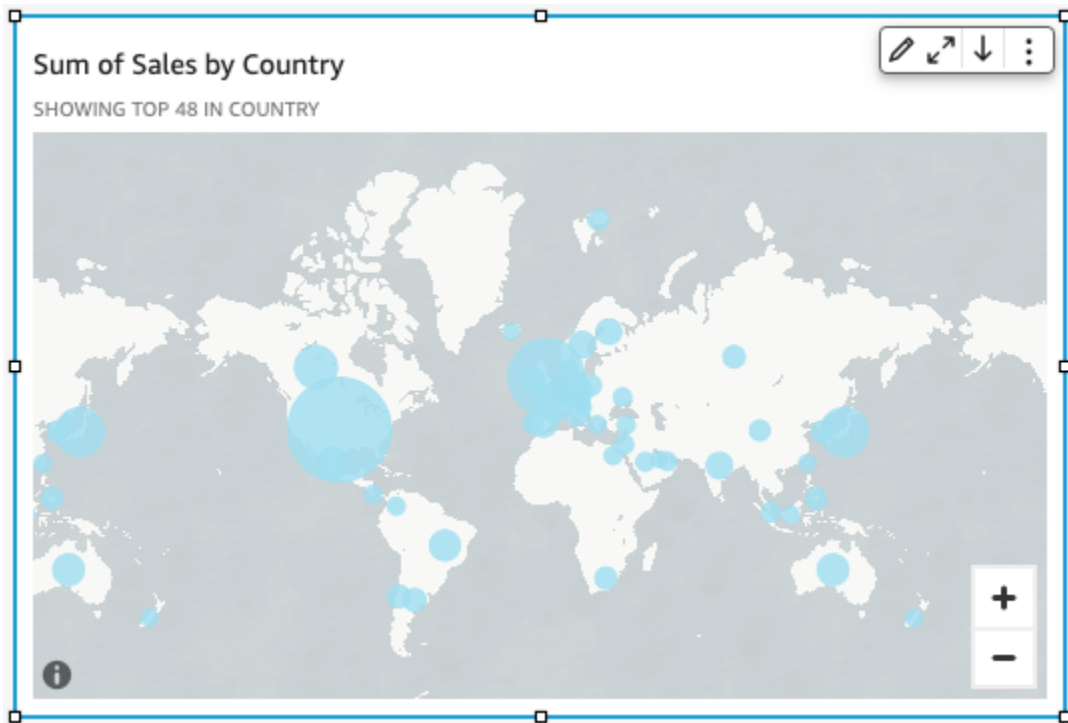
A point map appears with a point for each location in your data.

If the field is part of a geographic hierarchy, the hierarchy displays in the field well.



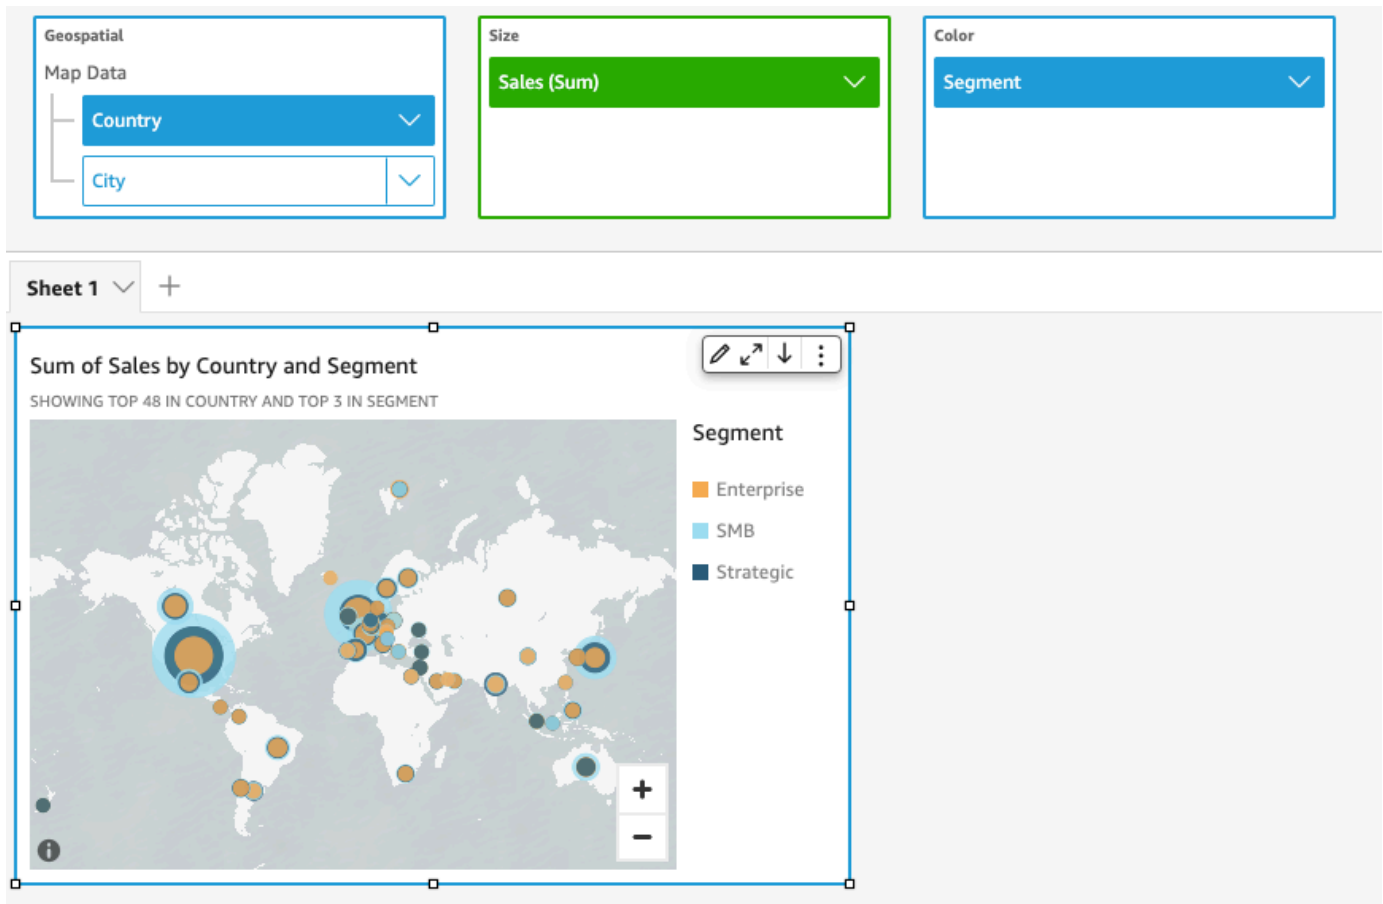
4. Drag a measure from the **Fields list** pane to the **Size** field well.

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 Amazon QuickSight 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 QuickSight.

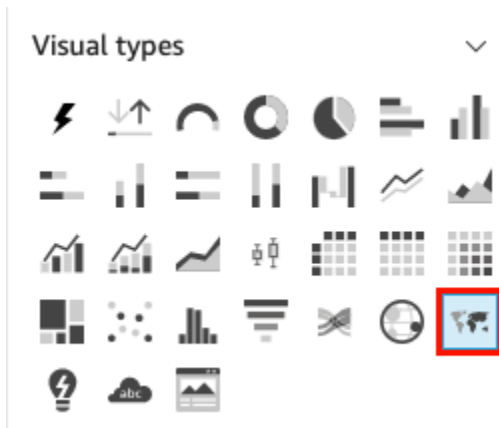
To create filled maps in QuickSight, make sure that you have the following:

- 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

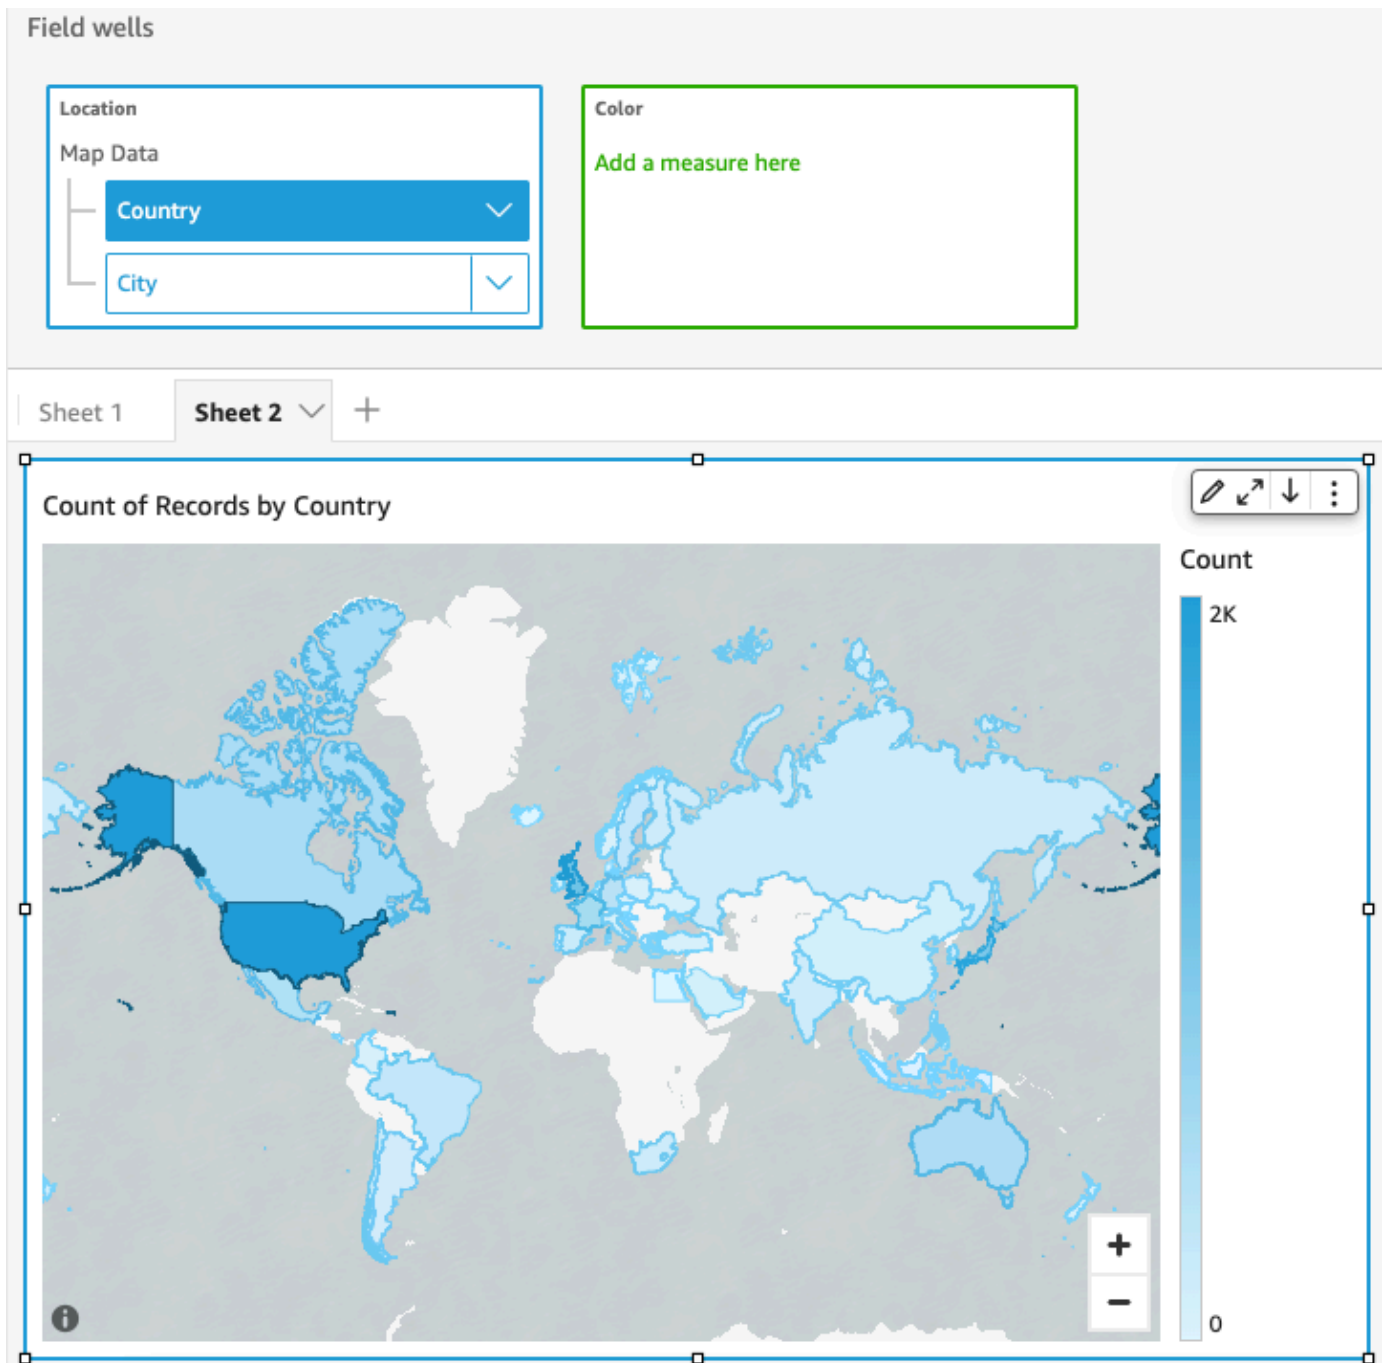
1. Add a new visual to your analysis. For more information about starting analyses, see [Starting an analysis in Amazon QuickSight](#). For more information about adding visuals to analyses, see [Adding a visual](#).
2. For **Visual type**, choose the **Filled map** icon.



3. Drag a geographic field from the **Fields list** pane to the **Location** field well, for example Country.

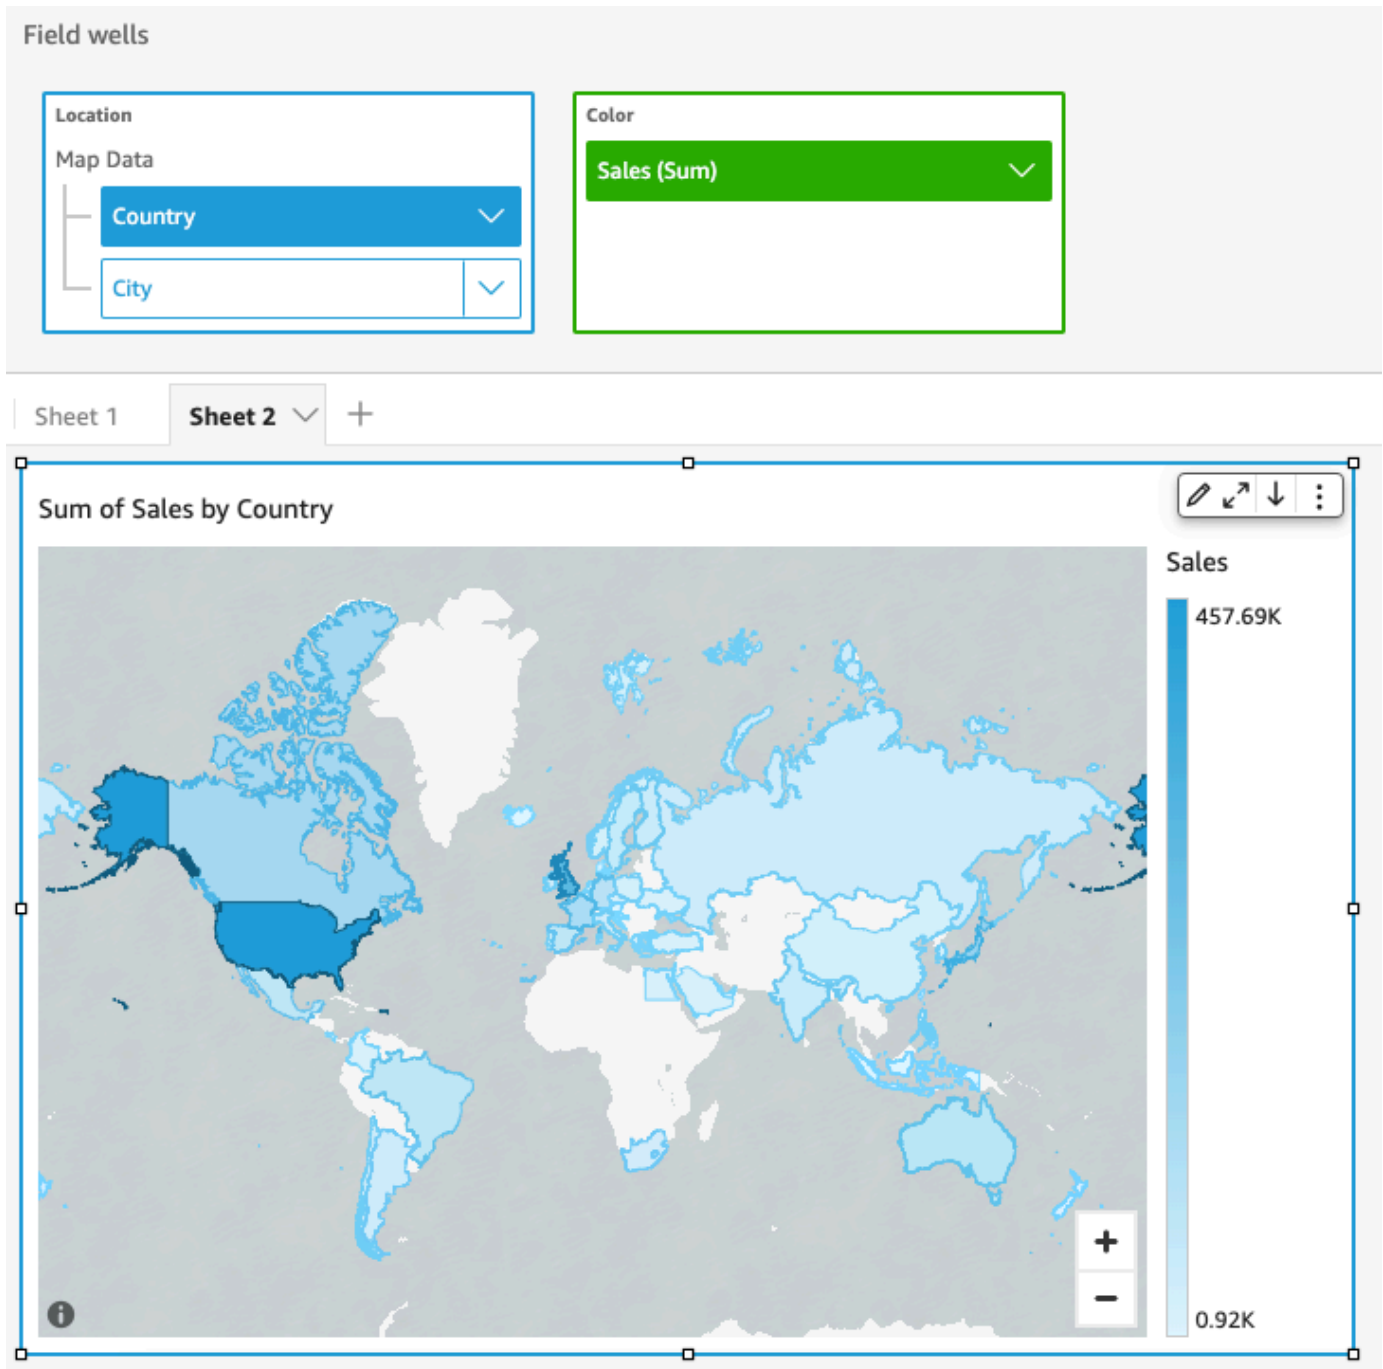
A filled map appears with each location in your data filled in by the number of times they appear in your dataset (the count).

If the field is part of a geographic hierarchy, the hierarchy displays in the field well.



- (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 Amazon QuickSight 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

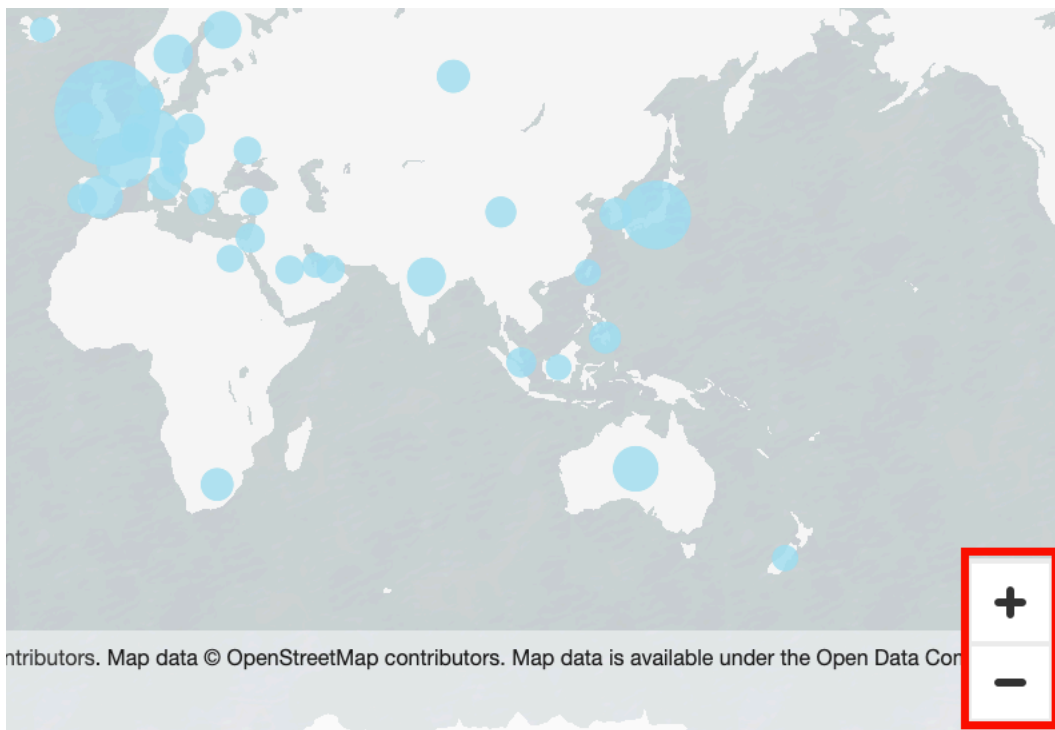
icons at bottom right of the map. Using this option, you can quickly zoom back to the underlying data.

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 QuickSight 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.

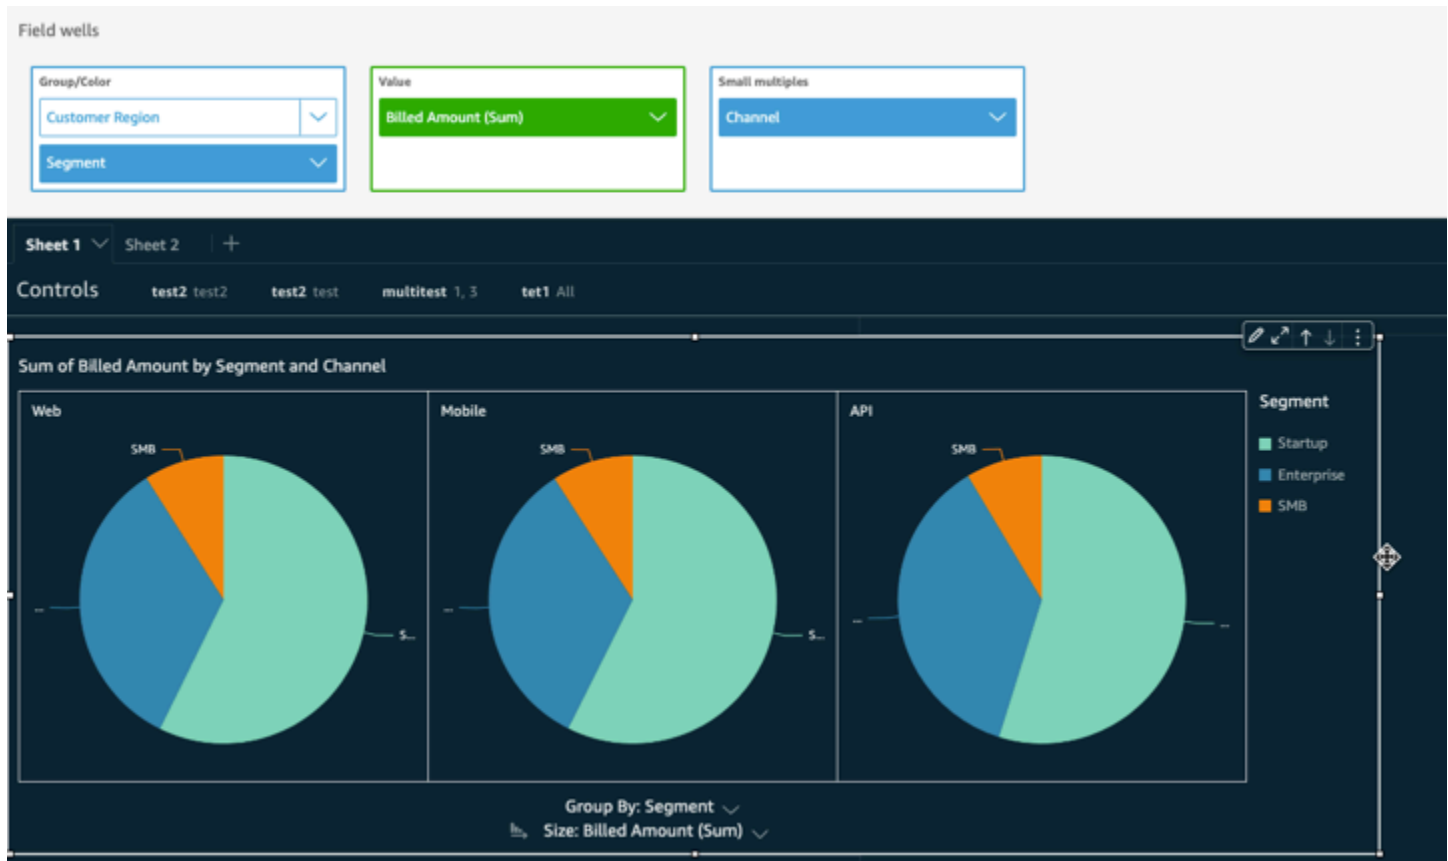
Small multiples aren't listed in the palette visualization icons. Instead, the option to create small multiples appears as a field well, in the visuals that support it.


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**
- Title options
 - **Panel title** (toggle)
 - Font size and color
 - Font weight
 - Text alignment
- **Border options**
 - Line thickness, style, and color
 - **Panel gutter** (toggle), with an option for **Spacing**
 - **Panel background** (toggle), with an option for **Background color**

The following screenshots show examples of small multiples.







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 QuickSight handles data that falls outside display limits, see [Display limits](#).

The icon for a pie chart is as follows.



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 QuickSight
Changing the title display	Yes		Titles and subtitles on visual types in QuickSight in QuickSight
Changing the axis range	Not applicable		Range and scale on visual types in QuickSight
Showing or hiding axis labels.	Yes		Axes and grid lines on visual types in QuickSight
Changing the visual colors	Yes		Colors in visual types in QuickSight
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
Sorting	Yes	You can sort on the field that you choose for the value or the group or color.	Sorting visual data in Amazon QuickSight

Feature	Supported?	Comments	For more information
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 Amazon QuickSight

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 Amazon QuickSight](#).

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 QuickSight 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 V-shaped 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 QuickSight](#).

For information on formatting pivot tables, see [Formatting in Amazon QuickSight](#).

For information on custom total aggregations for pivot tables, see [Custom total values](#).

The icon for a pivot table is as follows.



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 Amazon QuickSight](#)
- [Sorting pivot tables in Amazon QuickSight](#)
- [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 QuickSight
Changing the title display	Yes		Titles and subtitles on visual types in QuickSight in QuickSight
Changing the axis range	Not applicable		Range and scale on visual types in QuickSight

Feature	Supported?	Comments	For more information
Changing the visual colors	No		Colors in visual types in QuickSight
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 QuickSight Sorting pivot tables in Amazon QuickSight
Performing field aggregation	Yes	<p>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.</p> <p>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.</p>	Changing field aggregation

Feature	Supported?	Comments	For more information
Adding drill-downs	No		Adding drill-downs to visual data in Amazon QuickSight
Showing and hiding totals and subtotals	Yes	<p>You can show or hide totals and subtotals for rows and columns.</p> <p>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.</p>	
Exporting or copying data	Yes	<p>You can export all of the data to a CSV file.</p> <p>You can select and copy the content of the cells.</p>	Exporting data from visuals
Conditional formatting	Yes	You can add conditional formatting for values, subtotals and totals.	Conditional formatting on visual types in QuickSight

Creating a pivot table

Use the following procedure to create a pivot table.

To create a pivot table

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 pivot table icon.
4. From the **Fields list** pane, choose the fields that you want to include. Amazon QuickSight 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 QuickSight, 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 QuickSight](#).

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.

Field wells

Rows	Columns	Values
Customer Region	Customer Segment	Billed Amount (Sum)
Consumption Channel	Service Line	

Revenue by Service Line, Customer Region, Customer Segment, and Consumption Channel

Customer Region	Consumption C...	Customer Segment > Service Line					
		Enterprise			SMB		
		Marketing	HR	Billing	Marketing	HR	Billing
		Billed Amount	Billed Amount	Billed Amount	Billed Amount	Billed Amount	Billed Amount
APAC	API	104,350.34	132,225.58	127,772.08	17,566.44	44,803.39	36,607.5
	Mobile	191,448.63	247,966.85	238,773.51	40,337.72	104,994.76	86,602.13
	Web	282,733.31	370,020.33	340,658.32	57,737.54	150,060.36	121,386.85
EMEA	API	253,154.34	353,785.82	427,324.86	35,788.24	79,706.78	84,140.91
	Mobile	445,424.8	663,531.12	787,623.59	83,356.67	189,671.89	201,325.78
	Web	659,433.53	951,639.98	1,136,865.12	118,399.79	266,358.12	283,085.52
US	API	213,715.84	440,195.3	383,297.6	61,647.74	101,238.29	108,180.43
	Mobile	497,424.85	1,029,193.5	899,185.99	150,126.31	248,692.51	257,293.02
	Web	714,712.03	1,466,952.72	1,284,108.35	210,907.85	350,534.51	366,952.17

- Choose **Row** for a row format.

Field wells

Rows

- Customer Region
- Consumption Channel

Columns

- Customer Segment
- Service Line

Values

- Billed Amount (Sum)

Revenue by Service Line, Customer Region, Customer Segment, and Consumption Channel

Customer Segment > Service Line

Customer Region	Consumption C...	Billed Amount	Enterprise			SMB
			Marketing	HR	Billing	Marketing
APAC	API	Billed Amount	104,350.34	132,225.58	127,772.08	17,566.44
	Mobile	Billed Amount	191,448.63	247,966.85	238,773.51	40,337.72
	Web	Billed Amount	282,733.31	370,020.33	340,658.32	57,737.54
EMEA	API	Billed Amount	253,154.34	353,785.82	427,324.86	35,788.24
	Mobile	Billed Amount	445,424.8	663,531.12	787,623.59	83,356.67
	Web	Billed Amount	659,433.53	951,639.98	1,136,865.12	118,399.79
US	API	Billed Amount	213,715.84	440,195.3	383,297.6	61,647.74
	Mobile	Billed Amount	497,424.85	1,029,193.5	899,185.99	150,126.31
	Web	Billed Amount	714,712.03	1,466,952.72	1,284,108.35	210,907.85

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. For more information, see [Font and style on visual types in QuickSight](#).

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.

Revenue by Service Line, Customer Region, Customer Segment, and Consumption Channel

		Customer Segment > Service Line					
Customer Region	Consumption C...	Billed Amount	Enterprise			SMB	Startup
			Marketing	HR	Billing		
APAC	API	Billed Amount	104,350.34	132,225.58	127,772.08	98,977.33	854,403.2
	Mobile	Billed Amount	191,448.63	247,966.85	238,773.51	231,934.61	1,985,478.41
	Web	Billed Amount	282,733.31	370,020.33	340,658.32	329,184.75	2,854,727.09
EMEA	API	Billed Amount	253,154.34	353,785.82	427,324.86	199,635.93	599,882.97
	Mobile	Billed Amount	445,424.8	663,531.12	787,623.59	474,354.34	1,399,320.6
	Web	Billed Amount	659,433.53	951,639.98	1,136,865.12	667,843.43	2,018,343.87
US	API	Billed Amount	213,715.84	440,195.3	383,297.6	271,066.46	2,193,052.91
	Mobile	Billed Amount	497,424.85	1,029,193.5	899,185.99	656,111.84	5,188,415.35
	Web	Billed Amount	714,712.03	1,466,952.72	1,284,108.35	928,394.53	7,381,337.62

Showing and hiding pivot table columns in Amazon QuickSight

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

- In your analysis, select the pivot table visual that you want to work with.
- Choose the column or row header that you want to hide, and then choose **Hide**.

Field wells Rows Channel Segment Columns Customer Region Values # Cost (Sum)

Sheet 1 +

Sum of Cost by Channel, Customer Region, and Segment

Channel	Segment	Customer Region		
		US Cost	EMEA Cost	APAC Cost
API	Enterprise	697,877.37	756,282.96	285,249.97
	SMB	182,391.93	147,228.72	80,661.19
	Startup	1,478,997.31	441,312.46	701,144.36
Mobile	Enterprise	1,548,957.28	1,475,086.85	435,460.37
	SMB	419,653.32	369,625.7	149,825.23
	Startup	3,322,570.56	1,088,844.05	1,291,006.02
Web	Enterprise	2,198,689.04	1,823,214.96	640,056.09
	SMB	589,733.39	444,692.8	212,877.77
	Startup	4,698,355.09	1,339,746.34	1,860,356.93

Or you can choose the field in the **Rows** or **Columns** field wells, and then choose **Hide**.

Field wells

Rows

Channel

Segment

Columns

Customer Region

Values

Cost (Sum)

Sheet 1

Sum of Cost by Channel, Customer Region, and Segment

Channel	Segment	Customer Region		
		US	EMEA	APAC
		Cost	Cost	Cost
API	Enterprise	697,877.37	756,282.96	285,249.97
	SMB	182,391.93	147,228.72	80,661.19
	Startup	1,478,997.31	441,312.46	701,144.36
Mobile	Enterprise	1,548,957.28	1,475,086.85	435,460.37
	SMB	419,653.32	369,625.7	149,825.23
	Startup	3,322,570.56	1,088,844.05	1,291,006.02

To hide values in a pivot table

1. In your analysis, select the pivot table visual that you want to work with.
2. In the **Fields well**, choose a field in the **Values** field well, and then choose **Hide**.

The screenshot shows the Amazon QuickSight interface with the following configuration:

- Rows:** Channel, Segment
- Columns:** Customer Region
- Values:** Cost (Sum), Revenue Goal (Sum)

The pivot table visual is titled "Sum of Cost and Sum of Revenue Goal by Channel, Customer Region, and Segment". The data is summarized in the following table:

Channel	Segment	Customer Region						R
		US		EMEA		APAC		
		Cost	Revenue Goal	Cost	Revenue Goal	Cost	Revenue Goal	
API	Enterprise	697,877.37	1,053,807.5	756,282.96	1,125,773.55	285,249.97		
	SMB	182,391.93	269,606.06	147,228.72	226,997.4	80,661.19		
	Startup	1,478,997.31	1,984,731.88	441,312.46	670,274.87	701,144.36		
Mobile	Enterprise	1,548,957.28	2,186,232.78	1,475,086.85	1,808,674.71	435,460.37		
	SMB	419,653.32	497,849.23	369,625.7	518,028.28	149,825.23		
	Startup	3,322,570.56	4,761,599.22	1,088,844.05	1,334,073.8	1,291,006.02		
Web	Enterprise	2,198,689.04	3,113,705.38	1,823,214.96	2,515,864.3	640,056.09		
	SMB	589,733.39	739,488.07	444,692.8	576,226.05	212,877.77		
	Startup	4,698,355.09	5,865,657.05	1,339,746.34	1,803,841.21	1,860,356.93		

To show a hidden field in a pivot table

1. In your analysis, select the pivot table visual that you want to work with.
2. In the **Fields well**, choose the field in the **Rows**, **Columns**, **Values** field well, and then choose **Show**.

Field wells

Rows

Channel

Segment

Columns

Customer Region

Values

Cost (Su

Sheet 1 +

Sum of Cost by Channel, Customer Region, and Segment

Channel	Customer Region		
	US	EMEA	APAC
	Cost	Cost	Cost
API	697,877.37	756,282.96	285,249.97
	182,391.93	147,228.72	80,661.19
	1,478,997.31	441,312.46	701,144.36
Mobile	1,548,957.28	1,475,086.85	435,460.37
	419,653.32	369,625.7	149,825.23
	3,322,570.56	1,088,844.05	1,291,006.02

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. In the pivot table, choose any column or row header, and then choose **Show all hidden fields**.

Sheet 1 ▾ +

Sum of Cost by Channel, Customer Region, and Segment

Channel	Customer Region		
	US	EMEA	APAC
API	Cost	Cost	Cost
	697,877.37	756,282.96	285,249.97
	182,391.93	147,228.72	80,661.19
Mobile	1,478,997.31	441,312.46	701,144.36
	1,548,957.28	1,475,086.85	435,460.37
	419,653.32	369,625.7	149,825.23
Web	3,322,570.56	1,088,844.05	1,291,006.02
	2,198,689.04	1,823,214.96	640,056.09
	589,733.39	444,692.8	212,877.77
	4,698,355.09	1,339,746.34	1,860,356.93

Or you can choose any field in the **Fields well** and choose **Show all hidden fields**.

Field wells

Rows

Channel ▼

Segment 🔄 ▼

Columns

Customer Region 🔄 ▼

Values Row Column

Cost (Sum) ▼

Sheet 1 +

Sum of Cost by Channel, Customer Region, and Segment 🔗 ↻ 🔄 ⋮

Channel	Cost	Cost	Cost
API	697,877.37	756,282.96	285,249.97
	182,391.93	147,228.72	80,661.19
	1,478,997.31	441,312.46	701,144.36
Mobile	1,548,957.28	1,475,086.85	435,460.37
	419,653.32	369,625.7	149,825.23
	3,322,570.56	1,088,844.05	1,291,006.02
Web	2,198,689.04	1,823,214.96	640,056.09
	589,733.39	444,692.8	212,877.77
	4,698,355.09	1,339,746.34	1,860,356.93

Sorting pivot tables in Amazon QuickSight

In Amazon QuickSight, 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.

i 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.

Customer Region	Segment	Cost
APAC	SMB	443,364.19
	Enterprise	1,360,766.43
	Startup	3,852,507.31
EMEA	SMB	961,547.22
	Startup	2,869,902.85
	Enterprise	4,054,584.77
US	SMB	1,191,778.64
	Enterprise	4,445,523.69
	Startup	9,499,922.96

Customer Region	Segment	Cost
APAC	SMB	443,364.19
EMEA	SMB	961,547.22
US	SMB	1,191,778.64
APAC	Enterprise	1,360,766.43
EMEA	Startup	2,869,902.85
APAC	Startup	3,852,507.31
EMEA	Enterprise	4,054,584.77
US	Enterprise	4,445,523.69
US	Startup	9,499,922.96

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).



Customer Region	Channel	Revenue Goal	Cost
US	API	\$3,308,145	\$2,359,267
	Mobile	\$7,445,681	\$5,291,181
	Web	\$9,718,851	\$7,486,778
	Subtotal	\$20,472,677	\$15,137,225
EMEA	API	\$2,023,046	\$1,344,824
	Mobile	\$3,660,777	\$2,933,557
	Web	\$4,895,932	\$3,607,654
	Subtotal	\$10,579,754	\$7,886,035
APAC	API	\$1,392,361	\$1,067,056
	Mobile	\$2,606,978	\$1,876,292
	Web	\$3,400,757	\$2,713,291
	Subtotal	\$7,400,096	\$5,656,638

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.

Pivot Table: Segment sorted by Cost in ascending order

Customer Region	Segment	Cost
APAC	Enterprise	1,360,766.43
	SMB	443,364.19
	Startup	3,852,507.31
EMEA	Enterprise	4,054,584.77
	SMB	961,547.22
	Startup	2,869,902.85
US	Enterprise	4,445,523.69
	SMB	1,191,778.64
	Startup	9,499,922.96

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**.

Customer Region	Segment	Revenue Goal	Cost
APAC	Enterprise	1,873,330.74	1,360,766.43
	SMB	523,381.34	443,364.19
	Startup	5,003,383.53	3,852,507.31
EMEA	Enterprise	5,450,312.56	4,054,584.77
	SMB	1,321,251.73	961,547.22
	Startup	3,808,189.88	2,869,902.85
US	Enterprise	6,353,745.66	4,445,523.69
	SMB	1,506,943.36	1,191,778.64
	Startup	12,611,988.15	9,499,922.96

Sorting by value headers in a pivot table also works on subtotals.

Sorting tabular pivot tables using the field wells

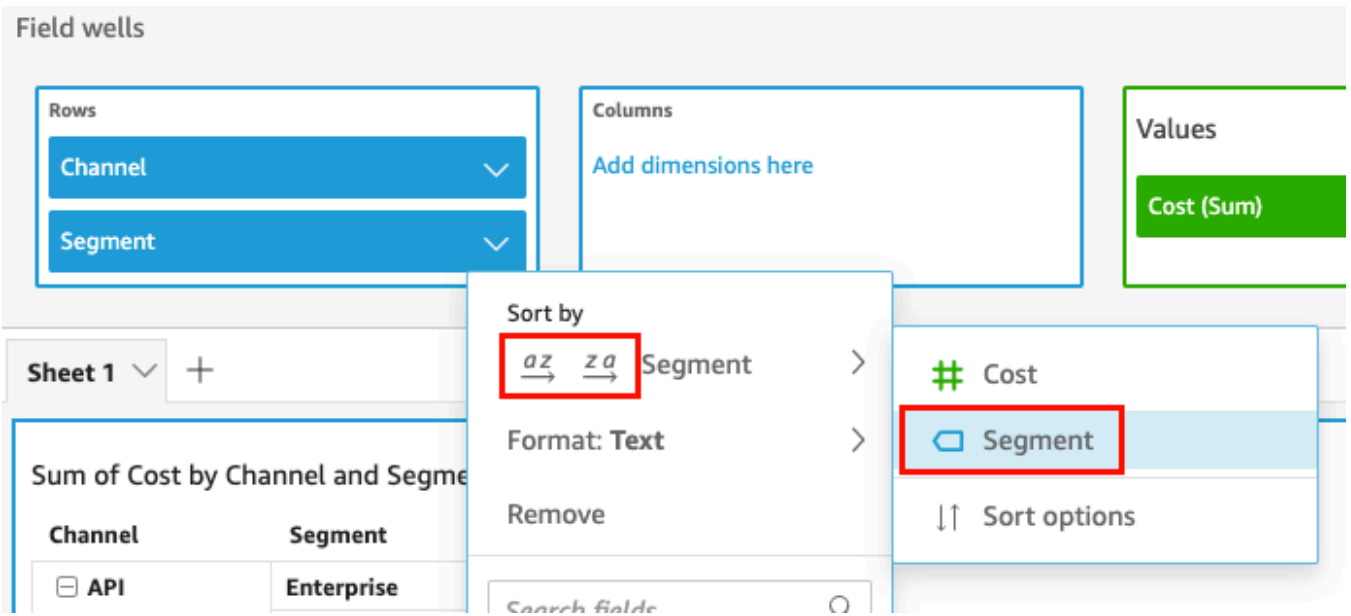
Use the following procedure to sort values in a tabular pivot table using the field wells.

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.

Region > Ship Mode

Rows	☐ Central			☐ East		
	Regular Air	Express Air	Delivery Truck	Regular Air	Express Air	De
	Sales	Sales	Sales	Sales	Sales	
Expand >						
Collapse >						
Product Category >	Sort by: Product Category >			☐ Product Category		
Product Sub-Category >	Sort order: Ascending >			# Sales		
Product Container >	393.38		276,138.02			
Product Name >	45,362.14	301.98		14,712.99		
☐ Medium Box	3,331.77	2,382.29		18,724.15	11,350.75	
☐ Office Furnishings						
☐ Jumbo Drum	176.33		298.51			
☐ Large Box	19,088.67	3,470.48		24,270.16	7,495.49	
☐ Medium Box	36,487.98	6,193.95		10,621.23	6,376.6	
☐ Small Box	28,754.9	4,993.22		17,615.19	1,601.23	
☐ Small Pack	20,554.86	1,329.09		12,245.54	2,983.67	
☐ Wrap Bag	12,150.68	2,057.71		10,627.04	479.78	
☐ Tables						
☐ Jumbo Box	95.97		230,167.66			

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

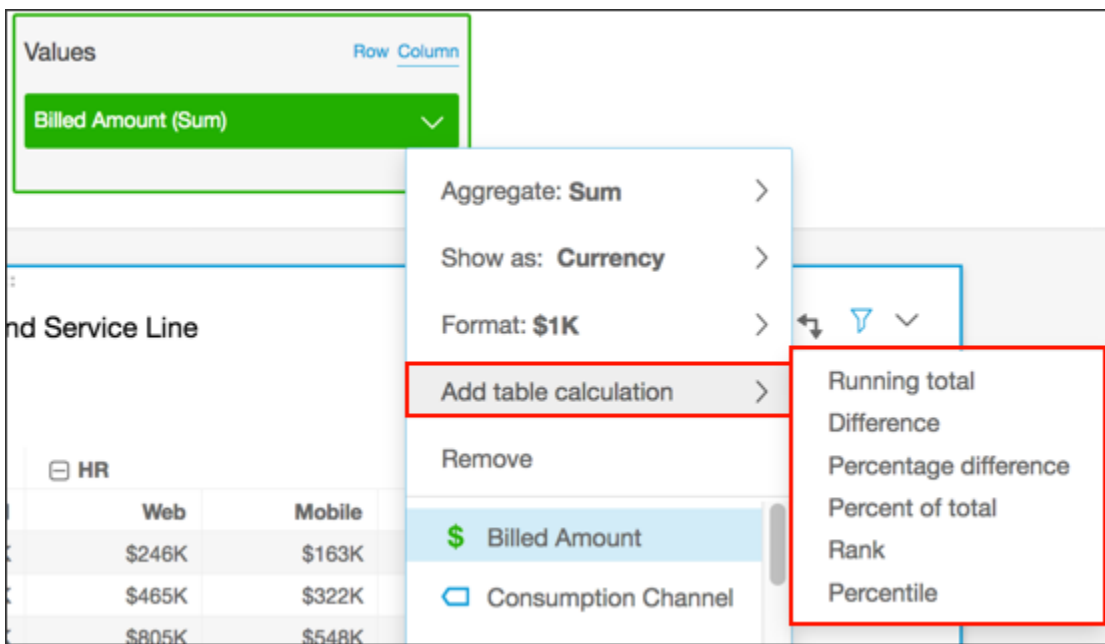
Use the following procedure to add a table calculation to a pivot table.

To add a table calculation to a pivot table

1. Expand the **Field wells** pane by choosing the field wells at the top 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

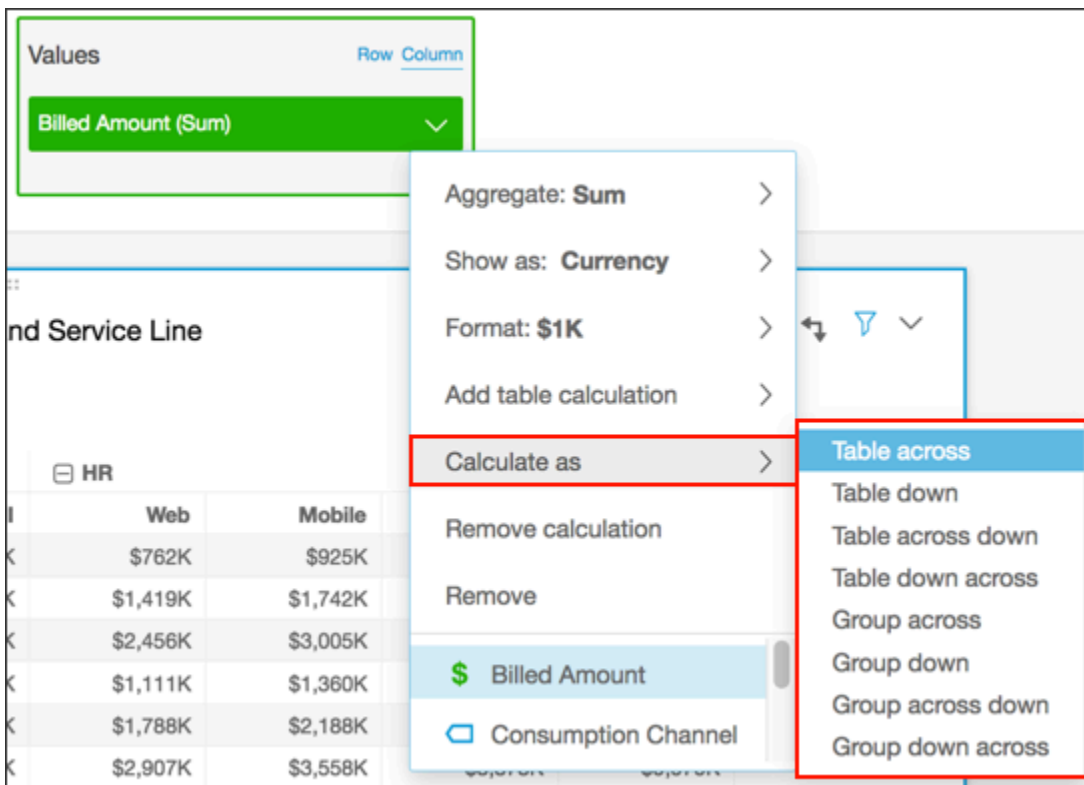
Use the following procedure to change the way a table calculation is applied to a pivot table.

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 at the top 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](#)

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. For example, suppose that you have the following data.

Sum of Billed Amount by Date, Customer Region, Consumption Channel, and Service Line

Service Line > Consumption Channel

Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014	\$257K	\$178K	\$82K	\$246K	\$163K	\$74K
	2015	\$471K	\$335K	\$149K	\$465K	\$322K	\$146K
	2016	\$819K	\$572K	\$260K	\$805K	\$548K	\$253K
EMEA	2014	\$376K	\$259K	\$122K	\$353K	\$249K	\$114K
	2015	\$589K	\$414K	\$209K	\$576K	\$401K	\$198K
	2016	\$968K	\$671K	\$336K	\$932K	\$651K	\$316K
US	2014	\$709K	\$496K	\$213K	\$729K	\$501K	\$217K
	2015	\$1,253K	\$872K	\$375K	\$1,267K	\$903K	\$375K
	2016	\$2,181K	\$1,538K	\$645K	\$2,192K	\$1,554K	\$654K

Applying the **Running total** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Sum of Billed Amount by Date, Customer Region, Consumption Channel, and Service Line

Service Line > Consumption Channel

Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014	\$257K	\$435K	\$516K	\$762K	\$925K	\$999K
	2015	\$471K	\$806K	\$954K	\$1,419K	\$1,742K	\$1,888K
	2016	\$819K	\$1,391K	\$1,651K	\$2,456K	\$3,005K	\$3,258K
EMEA	2014	\$376K	\$635K	\$757K	\$1,111K	\$1,360K	\$1,473K
	2015	\$589K	\$1,003K	\$1,212K	\$1,788K	\$2,188K	\$2,386K
	2016	\$968K	\$1,639K	\$1,975K	\$2,907K	\$3,558K	\$3,873K
US	2014	\$709K	\$1,205K	\$1,417K	\$2,146K	\$2,647K	\$2,864K
	2015	\$1,253K	\$2,126K	\$2,500K	\$3,767K	\$4,670K	\$5,044K
	2016	\$2,181K	\$3,720K	\$4,365K	\$6,557K	\$8,111K	\$8,765K

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. For example, suppose that you have the following data.

Sum of Billed Amount by Date, Customer Region, Consumption Channel, and Service Line

Service Line > Consumption Channel

Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014	\$257K	\$178K	\$82K	\$246K	\$163K	\$74K
	2015	\$471K	\$335K	\$149K	\$465K	\$322K	\$146K
	2016	\$819K	\$572K	\$260K	\$805K	\$548K	\$253K
EMEA	2014	\$376K	\$259K	\$122K	\$353K	\$249K	\$114K
	2015	\$589K	\$414K	\$209K	\$576K	\$401K	\$198K
	2016	\$968K	\$671K	\$336K	\$932K	\$651K	\$316K
US	2014	\$709K	\$496K	\$213K	\$729K	\$501K	\$217K
	2015	\$1,253K	\$872K	\$375K	\$1,267K	\$903K	\$375K
	2016	\$2,181K	\$1,538K	\$645K	\$2,192K	\$1,554K	\$654K

Applying the **Difference** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Sum of Billed Amount by Date, Customer Region, Consumption Channel, and Service Line

Service Line > Consumption Channel

Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014		-\$78K	-\$96K	\$164K	-\$83K	-\$90K
	2015		-\$136K	-\$186K	\$317K	-\$143K	-\$176K
	2016		-\$247K	-\$312K	\$545K	-\$257K	-\$295K
EMEA	2014		-\$117K	-\$137K	\$231K	-\$104K	-\$135K
	2015		-\$176K	-\$204K	\$366K	-\$175K	-\$203K
	2016		-\$297K	-\$335K	\$595K	-\$281K	-\$335K
US	2014		-\$213K	-\$284K	\$517K	-\$228K	-\$284K
	2015		-\$381K	-\$498K	\$892K	-\$364K	-\$528K
	2016		-\$643K	-\$893K	\$1,547K	-\$638K	-\$900K

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 $Cell1=(Cell1-null)/null$, $Cell2=(Cell2-Cell1)/Cell1$, $Cell3=(Cell3-Cell2)/Cell2$, and so on. Because $(Cell1-null)/null = null$, the Cell1 value is always empty. For example, take the following rows.

Sum of Billed Amount by Date, Customer Region, Consumption Channel, and Service Line

Service Line > Consumption Channel

Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014	\$257K	\$178K	\$82K	\$246K	\$163K	\$74K
	2015	\$471K	\$335K	\$149K	\$465K	\$322K	\$146K
	2016	\$819K	\$572K	\$260K	\$805K	\$548K	\$253K
EMEA	2014	\$376K	\$259K	\$122K	\$353K	\$249K	\$114K
	2015	\$589K	\$414K	\$209K	\$576K	\$401K	\$198K
	2016	\$968K	\$671K	\$336K	\$932K	\$651K	\$316K
US	2014	\$709K	\$496K	\$213K	\$729K	\$501K	\$217K
	2015	\$1,253K	\$872K	\$375K	\$1,267K	\$903K	\$375K
	2016	\$2,181K	\$1,538K	\$645K	\$2,192K	\$1,554K	\$654K

Applying the **Percentage Difference** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Sum of Billed Amount by Date, Customer Region, Consumption Channel, and Service Line

Service Line > Consumption Channel

Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014		-30.56%	-54.15%	200.96%	-33.64%	-54.88%
	2015		-28.91%	-55.61%	213.06%	-30.67%	-54.71%
	2016		-30.17%	-54.52%	209.77%	-31.94%	-53.80%
EMEA	2014		-31.03%	-52.93%	189.54%	-29.52%	-54.34%
	2015		-29.79%	-49.37%	174.91%	-30.38%	-50.64%
	2016		-30.68%	-49.86%	177.00%	-30.16%	-51.51%
US	2014		-29.99%	-57.15%	243.08%	-31.27%	-56.68%
	2015		-30.37%	-57.04%	237.94%	-28.71%	-58.50%
	2016		-29.48%	-58.06%	239.84%	-29.10%	-57.92%

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 $Cell1 = Cell1 / (\text{sum of all cells})$, $Cell2 = Cell2 / (\text{sum of all cells})$, and so on. For example, suppose that you have the following data.

Sum of Billed Amount by Date, Customer Region, Consumption Channel, and Service Line

Service Line > Consumption Channel

Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014	\$257K	\$178K	\$82K	\$246K	\$163K	\$74K
	2015	\$471K	\$335K	\$149K	\$465K	\$322K	\$146K
	2016	\$819K	\$572K	\$260K	\$805K	\$548K	\$253K
EMEA	2014	\$376K	\$259K	\$122K	\$353K	\$249K	\$114K
	2015	\$589K	\$414K	\$209K	\$576K	\$401K	\$198K
	2016	\$968K	\$671K	\$336K	\$932K	\$651K	\$316K
US	2014	\$709K	\$496K	\$213K	\$729K	\$501K	\$217K
	2015	\$1,253K	\$872K	\$375K	\$1,267K	\$903K	\$375K
	2016	\$2,181K	\$1,538K	\$645K	\$2,192K	\$1,554K	\$654K

Applying the **Percent of Total** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Sum of Billed Amount by Date, Customer Region, Consumption Channel, and Service Line

Service Line > Consumption Channel

Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014	21.47%	14.90%	6.83%	20.56%	13.64%	6.15%
	2015	20.54%	14.60%	6.48%	20.29%	14.06%	6.37%
	2016	20.40%	14.24%	6.47%	20.06%	13.65%	6.30%
EMEA	2014	20.00%	13.79%	6.49%	18.79%	13.24%	6.04%
	2015	19.07%	13.39%	6.77%	18.63%	12.97%	6.40%
	2016	19.31%	13.38%	6.71%	18.59%	12.98%	6.29%
US	2014	19.61%	13.72%	5.88%	20.17%	13.86%	6.00%
	2015	19.38%	13.49%	5.79%	19.59%	13.96%	5.79%
	2016	19.38%	13.66%	5.73%	19.48%	13.81%	5.81%

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.

Customer Region	Service Line		
	Marketing	HR	Billing
APAC	\$1,357K	\$3,023K	\$3,122K
EMEA	\$2,247K	\$3,788K	\$3,944K
US	\$4,657K	\$8,392K	\$8,282K

Applying the **Rank** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Customer Region	Service Line		
	Marketing	HR	Billing
APAC	3	2	1
EMEA	3	2	1
US	3	1	2

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.

$$\text{percentile rank}(x) = 100 * B / N$$

Where:

B = number of scores below x

N = number of scores

For example, suppose that you have the following data.

Sum of Billed Amount by Date, Customer Region, Consumption Channel, and Service Line

Service Line > Consumption Channel

Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014	\$257K	\$178K	\$82K	\$246K	\$163K	\$74K
	2015	\$471K	\$335K	\$149K	\$465K	\$322K	\$146K
	2016	\$819K	\$572K	\$260K	\$805K	\$548K	\$253K
EMEA	2014	\$376K	\$259K	\$122K	\$353K	\$249K	\$114K
	2015	\$589K	\$414K	\$209K	\$576K	\$401K	\$198K
	2016	\$968K	\$671K	\$336K	\$932K	\$651K	\$316K
US	2014	\$709K	\$496K	\$213K	\$729K	\$501K	\$217K
	2015	\$1,253K	\$872K	\$375K	\$1,267K	\$903K	\$375K
	2016	\$2,181K	\$1,538K	\$645K	\$2,192K	\$1,554K	\$654K

Applying the **Percentile** function across the table rows, using **Table across** for **Calculate as**, gives you the following results.

Sum of Billed Amount by Customer Region, Date, Consumption Channel, and Service Line

Service Line > Consumption Channel

Customer Region	Date	Billing			HR		
		Web	Mobile	API	Web	Mobile	API
APAC	2014	88.89	66.67	33.33	77.78	55.56	22.22
	2015	88.89	66.67	33.33	77.78	55.56	22.22
	2016	88.89	66.67	33.33	77.78	55.56	11.11
EMEA	2014	88.89	66.67	22.22	77.78	55.56	11.11
	2015	88.89	66.67	22.22	77.78	55.56	11.11
	2016	88.89	66.67	22.22	77.78	55.56	11.11
US	2014	77.78	55.56	11.11	88.89	66.67	22.22
	2015	77.78	55.56	22.22	88.89	66.67	11.11
	2016	77.78	55.56	11.11	88.89	66.67	22.22

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.

Sum of Billed Amount by Customer Region, Date, and Consumption Channel ↺ 🔍 ⌵

Customer Region	Date	Consumption Channel		
		Web	Mobile	API
APAC	2014	\$599K	\$410K	\$186K
	2015	\$1,137K	\$795K	\$361K
	2016	\$1,997K	\$1,379K	\$638K
EMEA	2014	\$932K	\$646K	\$302K
	2015	\$1,511K	\$1,046K	\$531K
	2016	\$2,461K	\$1,702K	\$848K
US	2014	\$1,816K	\$1,257K	\$539K
	2015	\$3,230K	\$2,272K	\$963K
	2016	\$5,613K	\$3,971K	\$1,669K

Applying the **Running total** function using **Table across** gives you the following results, with row totals in the last column.

Sum of Billed Amount by Customer Region, Date, and Consumption Channel

Customer Region	Date	Consumption Channel		
		Web	Mobile	API
APAC	2014	\$599K	\$1,009K	\$1,195K
	2015	\$1,137K	\$1,932K	\$2,293K
	2016	\$1,997K	\$3,376K	\$4,014K
EMEA	2014	\$932K	\$1,578K	\$1,880K
	2015	\$1,511K	\$2,557K	\$3,088K
	2016	\$2,461K	\$4,163K	\$5,011K
US	2014	\$1,816K	\$3,074K	\$3,613K
	2015	\$3,230K	\$5,502K	\$6,465K
	2016	\$5,613K	\$9,584K	\$11,253K

Table down

Using **Table down** applies the calculation down the columns of the pivot table, regardless of any grouping. For example, take the following pivot table.

Sum of Billed Amount by Customer Region, Date, and Consumption Channel

Customer Region	Date	Consumption Channel		
		Web	Mobile	API
APAC	2014	\$599K	\$410K	\$186K
	2015	\$1,137K	\$795K	\$361K
	2016	\$1,997K	\$1,379K	\$638K
EMEA	2014	\$932K	\$646K	\$302K
	2015	\$1,511K	\$1,046K	\$531K
	2016	\$2,461K	\$1,702K	\$848K
US	2014	\$1,816K	\$1,257K	\$539K
	2015	\$3,230K	\$2,272K	\$963K
	2016	\$5,613K	\$3,971K	\$1,669K

Applying the **Running total** function using **Table down** gives you the following results, with column totals in the last row.

Sum of Billed Amount by Customer Region, Date, and Consumption Channel

Customer Region	Date	Consumption Channel		
		Web	Mobile	API
APAC	2014	\$599K	\$410K	\$186K
	2015	\$1,735K	\$1,205K	\$547K
	2016	\$3,733K	\$2,584K	\$1,185K
EMEA	2014	\$4,665K	\$3,230K	\$1,487K
	2015	\$6,176K	\$4,276K	\$2,017K
	2016	\$8,637K	\$5,978K	\$2,865K
US	2014	\$10,454K	\$7,235K	\$3,404K
	2015	\$13,684K	\$9,507K	\$4,367K
	2016	\$19,297K	\$13,478K	\$6,036K

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. For example, take the following pivot table.

Sum of Billed Amount by Customer Region, Date, and Consumption Channel

Customer Region	Date	Consumption Channel		
		Web	Mobile	API
APAC	2014	\$599K	\$410K	\$186K
	2015	\$1,137K	\$795K	\$361K
	2016	\$1,997K	\$1,379K	\$638K
EMEA	2014	\$932K	\$646K	\$302K
	2015	\$1,511K	\$1,046K	\$531K
	2016	\$2,461K	\$1,702K	\$848K
US	2014	\$1,816K	\$1,257K	\$539K
	2015	\$3,230K	\$2,272K	\$963K
	2016	\$5,613K	\$3,971K	\$1,669K

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.

Sum of Billed Amount by Customer Region, Date, and Consumption Channel				
Customer Region	Date	Consumption Channel		
		Web	Mobile	API
APAC	2014	\$599K	\$1,009K	\$1,195K
	2015	\$2,332K	\$3,127K	\$3,488K
	2016	\$5,485K	\$6,864K	\$7,501K
EMEA	2014	\$8,433K	\$9,079K	\$9,381K
	2015	\$10,893K	\$11,939K	\$12,469K
	2016	\$14,931K	\$16,633K	\$17,480K
US	2014	\$19,297K	\$20,554K	\$21,093K
	2015	\$24,323K	\$26,595K	\$27,558K
	2016	\$33,171K	\$37,142K	\$38,811K

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.

Sum of Billed Amount by Customer Region, Date, and Consumption Channel				
Customer Region	Date	Consumption Channel		
		Web	Mobile	API
APAC	2014	21	24	27
	2015	13	18	25
	2016	6	11	20
EMEA	2014	16	19	26
	2015	10	14	23
	2016	4	8	17
US	2014	7	12	22
	2015	3	5	15
	2016	1	2	9

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. For example, take the following pivot table.

Customer Region	Date	Consumption Channel		
		Web	Mobile	API
APAC	2014	\$599K	\$410K	\$186K
	2015	\$1,137K	\$795K	\$361K
	2016	\$1,997K	\$1,379K	\$638K
EMEA	2014	\$932K	\$646K	\$302K
	2015	\$1,511K	\$1,046K	\$531K
	2016	\$2,461K	\$1,702K	\$848K
US	2014	\$1,816K	\$1,257K	\$539K
	2015	\$3,230K	\$2,272K	\$963K
	2016	\$5,613K	\$3,971K	\$1,669K

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.

Sum of Billed Amount by Customer Region, Date, and Consumption Channel

Customer Region	Date	Consumption Channel		
		Web	Mobile	API
APAC	2014	\$599K	\$1,009K	\$1,195K
	2015	\$2,332K	\$3,127K	\$3,488K
	2016	\$5,485K	\$6,864K	\$7,501K
EMEA	2014	\$8,433K	\$9,079K	\$9,381K
	2015	\$10,893K	\$11,939K	\$12,469K
	2016	\$14,931K	\$16,633K	\$17,480K
US	2014	\$19,297K	\$20,554K	\$21,093K
	2015	\$24,323K	\$26,595K	\$27,558K
	2016	\$33,171K	\$37,142K	\$38,811K

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.

Sum of Billed Amount by Customer Region, Date, and Consumption Channel

Customer Region	Date	Consumption Channel		
		Web	Mobile	API
APAC	2014	21	24	27
	2015	13	18	25
	2016	6	11	20
EMEA	2014	16	19	26
	2015	10	14	23
	2016	4	8	17
US	2014	7	12	22
	2015	3	5	15
	2016	1	2	9

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.

		Consumption Channel		
Customer Region	Date	Web	Mobile	API
APAC	2014	\$599K	\$410K	\$186K
	2015	\$1,137K	\$795K	\$361K
	2016	\$1,997K	\$1,379K	\$638K
EMEA	2014	\$932K	\$646K	\$302K
	2015	\$1,511K	\$1,046K	\$531K
	2016	\$2,461K	\$1,702K	\$848K
US	2014	\$1,816K	\$1,257K	\$539K
	2015	\$3,230K	\$2,272K	\$963K
	2016	\$5,613K	\$3,971K	\$1,669K

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.

Sum of Billed Amount by Customer Region, Date, Service Line, and Consumption Channel

Service Line > Consumption Channel

Customer Region	Date	Billing		Marketing	
		API	Mobile	API	Mobile
APAC	2012	\$23K	\$74K		
	2013	\$45K	\$153K	\$3K	\$11K
	2014	\$82K	\$260K	\$31K	\$100K
EMEA	2012	\$22K	\$78K		
	2013	\$50K	\$172K	\$4K	\$15K
	2014	\$122K	\$381K	\$66K	\$204K
US	2012	\$43K	\$140K		
	2013	\$114K	\$382K	\$14K	\$48K
	2014	\$213K	\$709K	\$110K	\$370K

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).

Sum of Billed Amount by Customer Region, Date, and Consumption Channel

Customer Region	Date	Consumption Channel		
		Web	Mobile	API
APAC	2014	\$599K	\$410K	\$186K
	2015	\$1,137K	\$795K	\$361K
	2016	\$1,997K	\$1,379K	\$638K
EMEA	2014	\$932K	\$646K	\$302K
	2015	\$1,511K	\$1,046K	\$531K
	2016	\$2,461K	\$1,702K	\$848K
US	2014	\$1,816K	\$1,257K	\$539K
	2015	\$3,230K	\$2,272K	\$963K
	2016	\$5,613K	\$3,971K	\$1,669K

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.

Sum of Billed Amount by Customer Region, Date, Service Line, and Consumption Channel

Customer Region	Date	Service Line > Consumption Channel			
		Billing		Marketing	
		API	Mobile	API	Mobile
APAC	2012	\$23K	\$51K		
	2013	\$68K	\$159K	\$3K	\$8K
	2014	\$149K	\$337K	\$34K	\$77K
EMEA	2012	\$22K	\$56K		
	2013	\$72K	\$178K	\$4K	\$11K
	2014	\$194K	\$438K	\$70K	\$148K
US	2012	\$43K	\$97K		
	2013	\$157K	\$365K	\$14K	\$34K
	2014	\$370K	\$861K	\$124K	\$294K

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).

Sum of Billed Amount by Customer Region, Date, and Consumption Channel

Customer Region	Date	Consumption Channel		
		Web	Mobile	API
APAC	2014	\$599K	\$410K	\$186K
	2015	\$1,137K	\$795K	\$361K
	2016	\$1,997K	\$1,379K	\$638K
EMEA	2014	\$932K	\$646K	\$302K
	2015	\$1,511K	\$1,046K	\$531K
	2016	\$2,461K	\$1,702K	\$848K
US	2014	\$1,816K	\$1,257K	\$539K
	2015	\$3,230K	\$2,272K	\$963K
	2016	\$5,613K	\$3,971K	\$1,669K

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.

Sum of Billed Amount by Customer Region, Date, Service Line, and Consumption Channel

Customer Region	Date	Service Line > Consumption Channel			
		Billing		Marketing	
		API	Mobile	API	Mobile
APAC	2012	\$23K	\$201K		
	2013	\$68K	\$309K	\$3K	\$42K
	2014	\$149K	\$487K	\$34K	\$111K
EMEA	2012	\$22K	\$250K		
	2013	\$72K	\$372K	\$4K	\$81K
	2014	\$194K	\$632K	\$70K	\$219K
US	2012	\$43K	\$466K		
	2013	\$157K	\$734K	\$14K	\$157K
	2014	\$370K	\$1,231K	\$124K	\$418K

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.

Sum of Billed Amount by Customer Region, Date, Service Line, and Consumption Channel

Customer Region	Date	Service Line > Consumption Channel			
		Billing		Marketing	
		API	Mobile	API	Mobile
APAC	2012	6	4		
	2013	5	2	4	3
	2014	3	1	2	1
EMEA	2012	6	4		
	2013	5	2	4	3
	2014	3	1	2	1
US	2012	6	5		
	2013	4	2	4	3
	2014	3	1	2	1

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 QuickSight 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).

Sum of Billed Amount by Customer Region, Date, and Consumption Channel

Customer Region	Date	Consumption Channel		
		Web	Mobile	API
APAC	2014	\$599K	\$410K	\$186K
	2015	\$1,137K	\$795K	\$361K
	2016	\$1,997K	\$1,379K	\$638K
EMEA	2014	\$932K	\$646K	\$302K
	2015	\$1,511K	\$1,046K	\$531K
	2016	\$2,461K	\$1,702K	\$848K
US	2014	\$1,816K	\$1,257K	\$539K
	2015	\$3,230K	\$2,272K	\$963K
	2016	\$5,613K	\$3,971K	\$1,669K

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.

Sum of Billed Amount by Customer Region, Date, Service Line, and Consumption Channel

Customer Region	Date	Service Line > Consumption Channel			
		Billing		Marketing	
		API	Mobile	API	Mobile
APAC	2012	\$23K	\$201K		
	2013	\$68K	\$309K	\$3K	\$42K
	2014	\$149K	\$487K	\$34K	\$111K
EMEA	2012	\$22K	\$250K		
	2013	\$72K	\$372K	\$4K	\$81K
	2014	\$194K	\$632K	\$70K	\$219K
US	2012	\$43K	\$466K		
	2013	\$157K	\$734K	\$14K	\$157K
	2014	\$370K	\$1,231K	\$124K	\$418K

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.

Sum of Billed Amount by Customer Region, Date, Service Line, and Consumption Channel

Customer Region	Date	Service Line > Consumption Channel			
		Billing		Marketing	
		API	Mobile	API	Mobile
APAC	2012	6	4		
	2013	5	2	4	3
	2014	3	1	2	1
EMEA	2012	6	4		
	2013	5	2	4	3
	2014	3	1	2	1
US	2012	6	5		
	2013	4	2	4	3
	2014	3	1	2	1

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 QuickSight. 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 QuickSight, 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.

Employee satisfaction by department



To work with radar charts, choose the following radar chart icon.



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 QuickSight
Changing the title display	Yes		Titles and subtitles on visual types in QuickSight in QuickSight
Changing the axis range	Yes		Range and scale on visual types in QuickSight
Changing the visual colors	Yes		Colors in visual types in QuickSight
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 QuickSight
Performing field aggregation	Yes		Changing field aggregation
Adding drill-downs	Not supported		Adding drill-downs to visual data in Amazon QuickSight
Choosing size	Yes		Formatting in Amazon QuickSight
Showing totals	Not supported		Formatting in Amazon QuickSight

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 QuickSight, 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 QuickSight, 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

Dimension (Source)	Dimension (Destination)	Measure (Weight)
W	Z	300
X	Z	5
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 QuickSight. 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 QuickSight. 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
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 QuickSight 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 QuickSight to order the rows in the Step Number column. For more information about using custom SQL in QuickSight, 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

Customer ID	Time	Action	Step Number
1	9:08 am	Step 3	3
2	11:44 am	Step 1	1
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

Customer ID	Time	Action	Step Number	Next Step Number
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 QuickSight, 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. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells.

Sankey diagrams are made of a source dimension, a destination dimension, and a measure.

To create a Sankey diagram, drag a dimension to the **Source** field well, a dimension to the **Destination** field well, and a measure to the **Weight** field well.

Customizing the number of nodes

Use the following procedure to customize the number of nodes that appear in a Sankey diagram. QuickSight 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 **Format visual** pane that opens at left, 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.

QuickSight 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	

Feature	Supported?	For more information
Changing the title display	Yes	Titles and subtitles on visual types in QuickSight in QuickSight
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. For more information about how Amazon QuickSight handles data that falls outside display limits, see [Display limits](#).

The icon for a scatter plot is as.



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 QuickSight
Changing the title display	Yes		Titles and subtitles on visual types in QuickSight in QuickSight
Changing the axis range	Yes	You can set the range for both the X and Y axes.	Range and scale on visual types in QuickSight
Showing or hiding axis lines, grid lines, axis labels, and axis sort icons	Yes		Axes and grid lines on visual types in QuickSight
Changing the visual colors	Yes		Colors in visual types in QuickSight
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	Focusing on visual elements

Feature	Supported?	Comments	For more information
		a date field as a dimension . In that case, you can only focus on a bubble, not exclude it.	Excluding visual elements
Sorting	No		Sorting visual data in Amazon QuickSight
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 Amazon QuickSight

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 Amazon QuickSight](#).


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 QuickSight](#).

The icon for a table is as follows.



To create a table visual

1. Open Amazon QuickSight 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 Amazon QuickSight](#).

- To use an existing analysis, choose the analysis that you want to edit.
3. Choose **Add (+), 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 **Add (+), 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 **Format visual** 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 **Format visual** pane, choose **Table options**, and select **Wrap header text**.

To rearrange columns in a table chart

1. Open the analysis that holds your table and click anywhere on **Field wells** to expand the field wells.
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 QuickSight 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. 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.

Url options

- None
- Make URLs hyperlinks
- Show URLs as images

Open in

- New tab
- Same tab
- New window

Style as

- Link
- Icon
- Plain text
- Custom link

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 QuickSight 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 QuickSight 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**.

Sum of Population by Region

Region	population
#N/A	
Midwest	
Northeast	
South	
West	
	2

Aggregate: **Sum** >

Sort by  >

Show as: **Number** >

Format: **1,234.57** >

Total: **Default** >

Hide

Move < >

Freeze column

Freeze up to this column

Conditional formatting

Remove


Custom total values

QuickSight 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.

Using text boxes

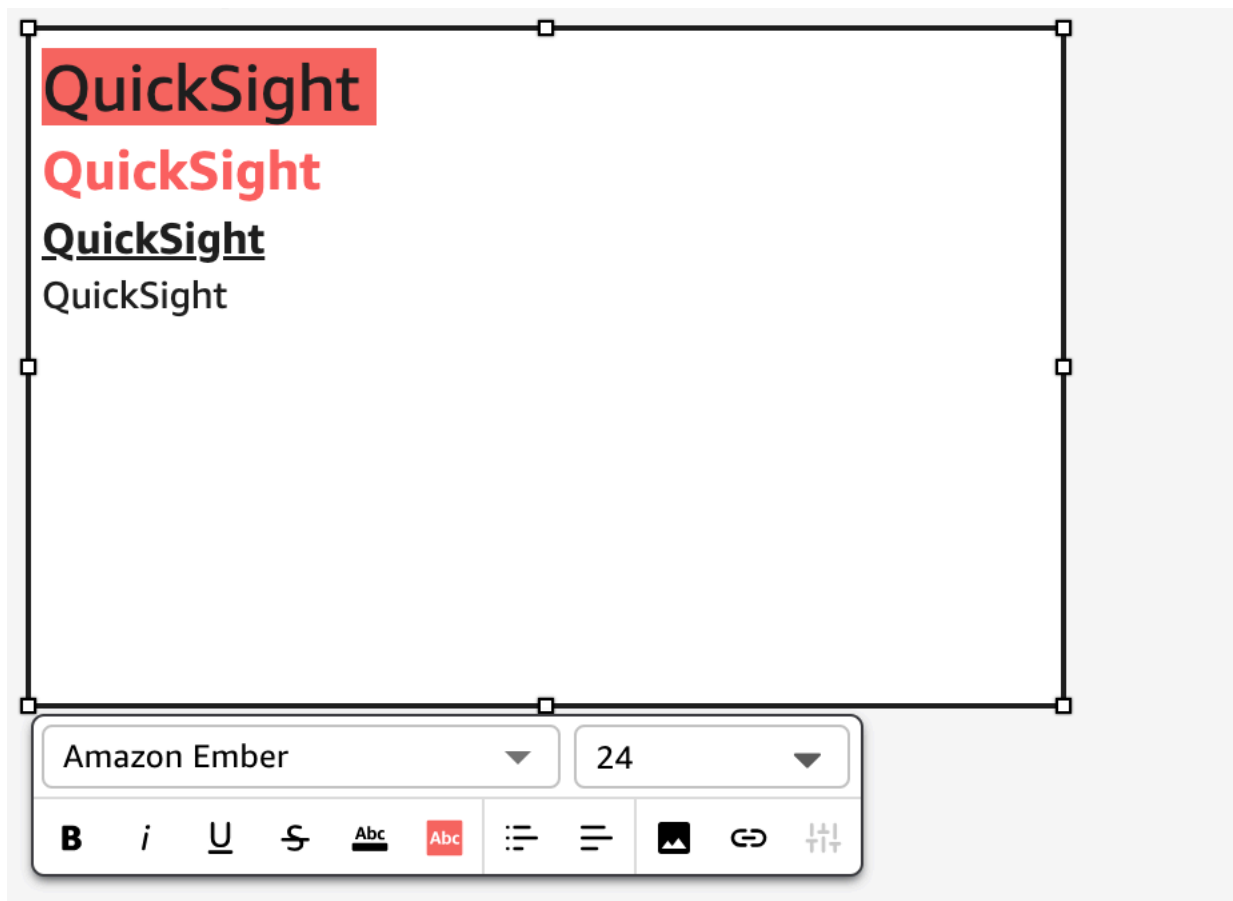
Add text to add context to sheets in an analysis by using a text box. Text can hold directions, descriptions, or even hyperlinks to external websites. The toolbar on the text box offers font settings so you can customize the font type, style, color, size, spacing, size in pixels, text highlights, and alignment. The text box itself has no format settings.

To add text to a new text box, simply select it and begin typing.

The icon for a text box is as follows.



The following screenshot shows an example of a text box.



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 QuickSight handles data that falls outside display limits, see [Display limits](#).

The icon for a tree map is as follows.



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 QuickSight
Changing the title display	Yes		Titles and subtitles on visual types in QuickSight in QuickSight
Changing the axis range	Not applicable		Range and scale on visual types in QuickSight
Changing the visual colors	No		Colors in visual types in QuickSight
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	Focusing on visual elements Excluding visual elements

Feature	Supported?	Comments	For more information
		, you can only focus on a rectangle, not exclude it.	
Sorting	No	Default sorting is in descending order by the measure in the Size column.	Sorting visual data in Amazon QuickSight
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 Amazon QuickSight

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.

- (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 Amazon QuickSight](#).

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 icon for a waterfall chart is as follows.



The following screenshot shows a waterfall chart.



To create a basic waterfall chart visual

1. Open Amazon QuickSight 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 Amazon QuickSight](#).
 - 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 Amazon QuickSight](#).

To understand the features supported by waterfall charts, see [Analytics formatting per type in QuickSight](#). For customization options, see [Formatting in Amazon QuickSight](#).

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 QuickSight handles data that falls outside display limits, see [Display limits](#).

The icon for a word cloud is as follows.



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 QuickSight
Changing the title display	Yes		Titles and subtitles on visual types in QuickSight in QuickSight
Changing the axis range	Not applicable		Range and scale on visual types in QuickSight
Changing the visual colors	Yes	To change the color, choose a word and then choose a color.	Colors in visual types in QuickSight
Focusing on or excluding elements	Yes		Focusing on visual elements Excluding visual elements
Sorting	Yes		Sorting visual data in Amazon QuickSight
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 Amazon QuickSight

Feature	Supported?	Comments	For more information
Using format options	Yes	You can choose to allow vertical words, emphasize 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).	Formatting in Amazon QuickSight
Showing totals	No		Formatting in Amazon QuickSight

Creating a word cloud

Use the following procedure to create a word cloud.

To create a word cloud

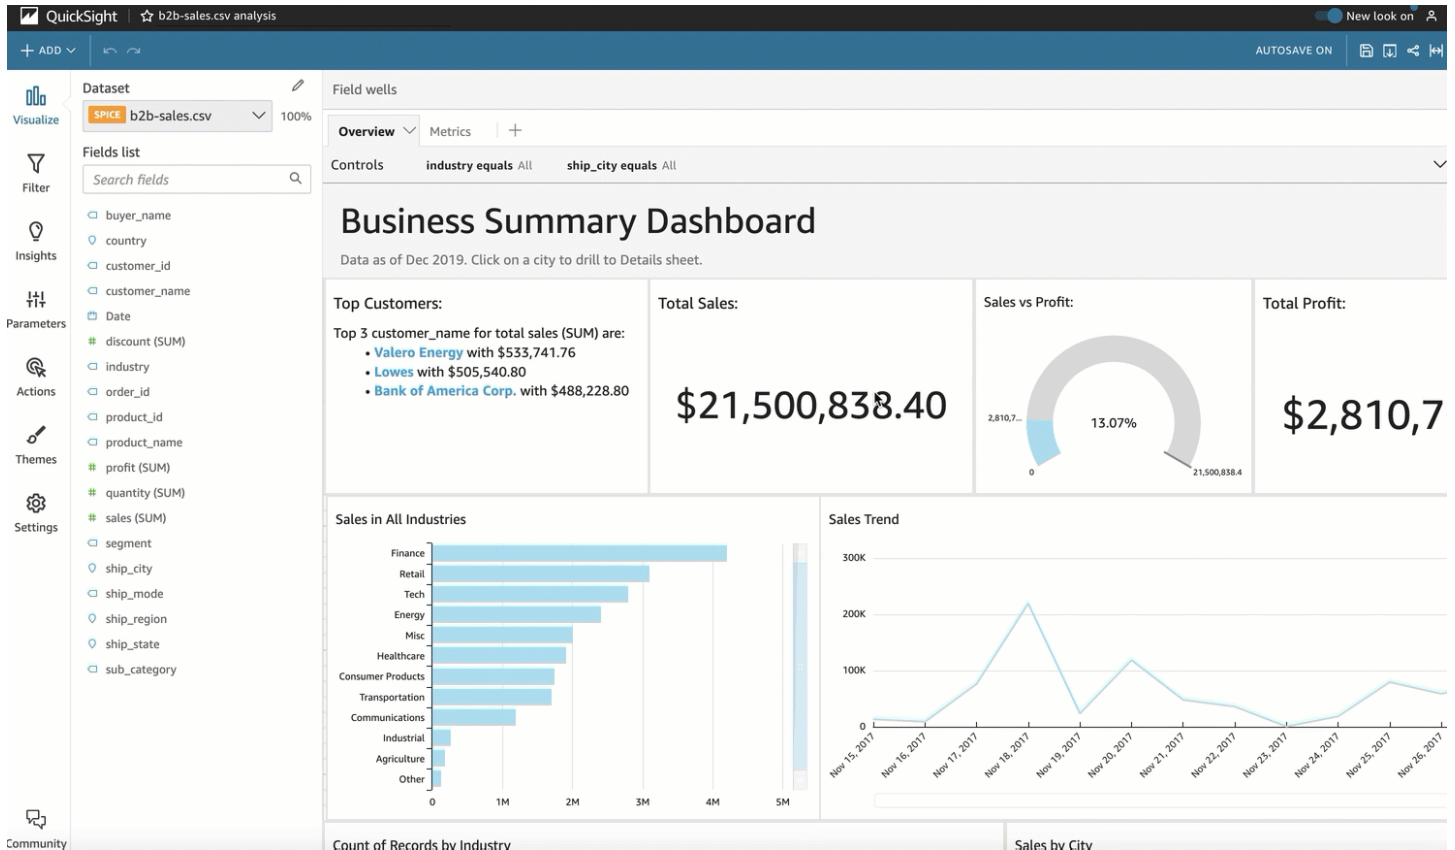
1. On the analysis page, choose **Visualize** on the tool bar.
2. Choose **Add** on the application bar, and then choose **Add visual**.
3. On the **Visual types** pane, choose the word cloud icon.
4. From the **Fields list** pane, drag the fields that you want to use to the appropriate field wells. 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 Amazon QuickSight](#).

Formatting in Amazon QuickSight

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 QuickSight](#).



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 QuickSight](#)
- [Table and pivot table formatting options in QuickSight](#)
- [Adding data bars to tables in QuickSight](#)

- [Map and geospatial chart formatting options in QuickSight](#)
- [Axes and grid lines on visual types in QuickSight](#)
- [Colors in visual types in QuickSight](#)
- [Working with field level coloring in Amazon QuickSight](#)
- [Conditional formatting on visual types in QuickSight](#)
- [Font and style on visual types in QuickSight](#)
- [KPI options](#)
- [Labels on visual types in QuickSight](#)
- [Formatting visual numeric data based on language settings in QuickSight](#)
- [Legends on visual types in QuickSight](#)
- [Line and marker styling on line charts in QuickSight](#)
- [Missing data on visual types in QuickSight](#)
- [Reference lines on visuals types in QuickSight](#)
- [Formatting radar charts in QuickSight](#)
- [Range and scale on visual types in QuickSight](#)
- [Small multiples axis options](#)
- [Titles and subtitles on visual types in QuickSight in QuickSight](#)
- [Tooltips on visual types in QuickSight](#)

Analytics formatting per type in QuickSight

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

KPIs don't support sorting.

- 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 Amazon QuickSight](#).

- 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 QuickSight

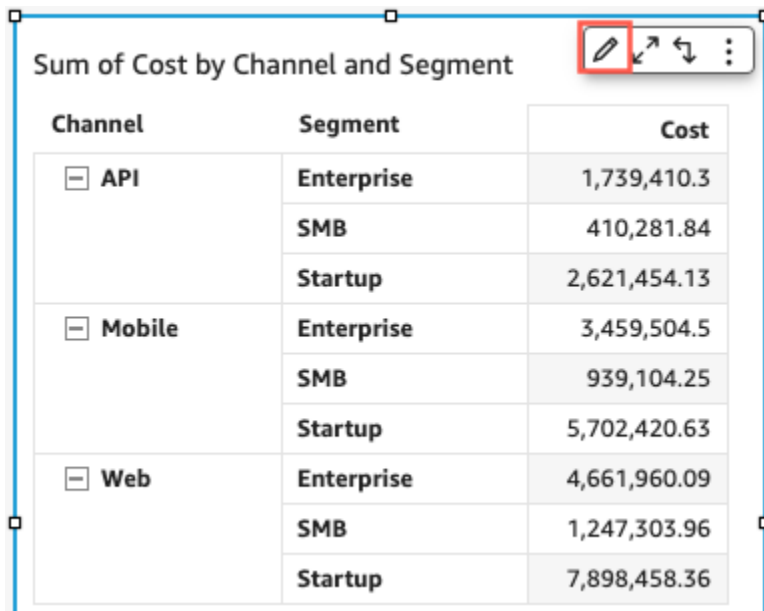
You can customize tables and pivot tables in Amazon QuickSight to meet your business needs. You can customize table headers, cells, and totals by specifying the color, size, wrap, and alignment of text in each. You can also specify the height of rows in a table, add borders and grid lines, and add custom 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 screenshot shows a table titled "Sum of Cost by Channel and Segment". The table has three columns: "Channel", "Segment", and "Cost". The data is grouped by Channel (API, Mobile, Web) and further subdivided by Segment (Enterprise, SMB, Startup). The "Format visual" icon (a pencil) is highlighted with a red box in the top right corner of the table's header area.

Channel	Segment	Cost
API	Enterprise	1,739,410.3
	SMB	410,281.84
	Startup	2,621,454.13
Mobile	Enterprise	3,459,504.5
	SMB	939,104.25
	Startup	5,702,420.63
Web	Enterprise	4,661,960.09
	SMB	1,247,303.96
	Startup	7,898,458.36

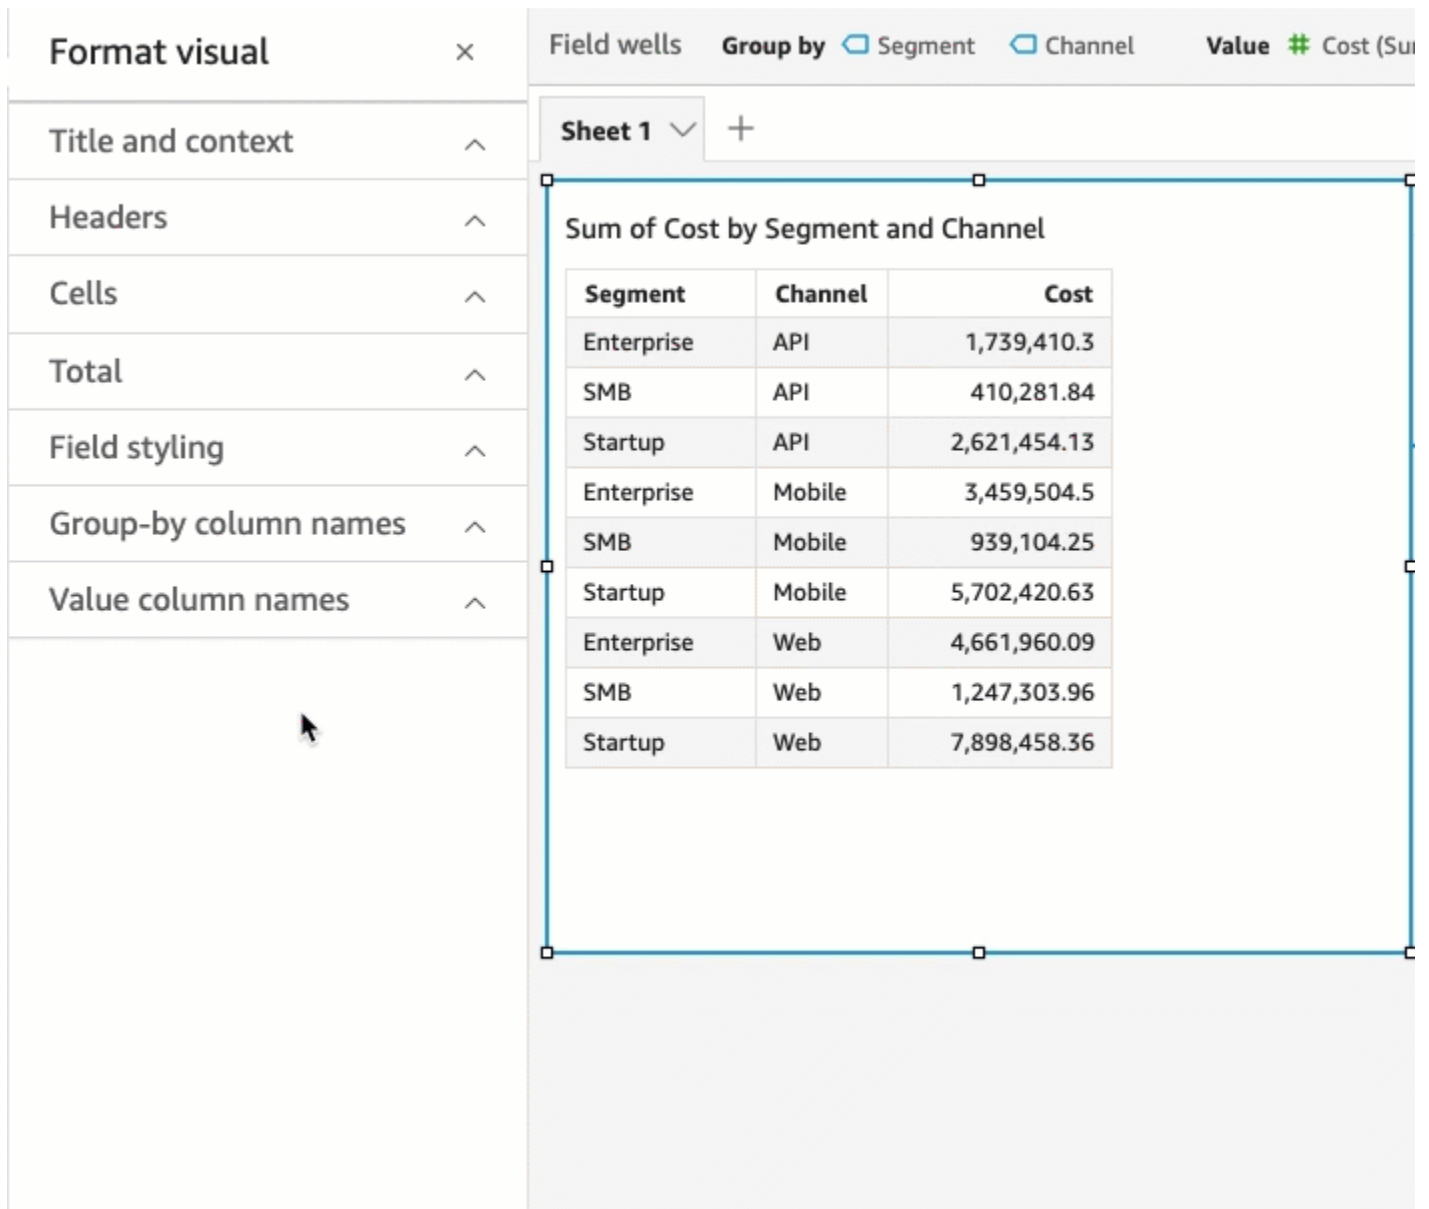
The **Format visual** pane opens at left.

Following, you can find descriptions for options for customizing each area of your table or pivot table in the **Format visual** pane.

Topics

- [Headers](#)
- [Cell formatting](#)
- [Totals and subtotals](#)
- [Row and column size in tables and pivot tables in QuickSight](#)
- [Customize pivot table data](#)

Headers



The screenshot shows the Amazon QuickSight interface. On the left is a navigation menu with the following items: 'Format visual' (closed), 'Title and context' (expanded), 'Headers' (expanded), 'Cells' (expanded), 'Total' (expanded), 'Field styling' (expanded), 'Group-by column names' (expanded), and 'Value column names' (expanded). The main area displays a pivot table visualization titled 'Sum of Cost by Segment and Channel'. The visualization is grouped by 'Segment' and 'Channel', with 'Cost' as the value field. The data is as follows:

Segment	Channel	Cost
Enterprise	API	1,739,410.3
SMB	API	410,281.84
Startup	API	2,621,454.13
Enterprise	Mobile	3,459,504.5
SMB	Mobile	939,104.25
Startup	Mobile	5,702,420.63
Enterprise	Web	4,661,960.09
SMB	Web	1,247,303.96
Startup	Web	7,898,458.36

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 **Format visual** 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 **Format visual** 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

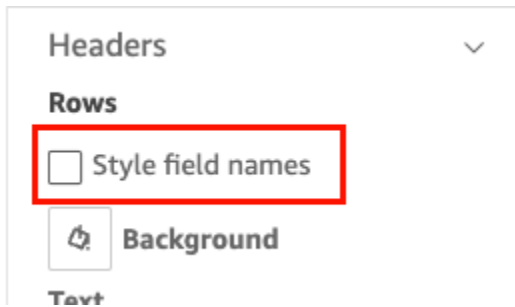
1. In the **Format visual** pane, choose **Headers**.
2. In the **Headers** section, do one or more of the following:
 - To wrap text in headers that are too long to fit, select **Wrap text**. Wrapping text in a header doesn't automatically increase the height of the header. Follow the previous procedure for increasing header height.
 - To customize the size of the text, choose a text size. You can choose between extra small and extra large text.
 - To change the font color, choose the **Abc** 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 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 **Format visual** 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, choose **Style field names** under **Rows**.



- To customize the header text size, choose a text size for **Text**. You can customize the text size for column headers in the **Columns** section, and row headers in the **Rows** section.
- To change the header font color, choose the **Abc** color icon, and then choose a color. You can choose a font color for column headers in the **Columns** section, and for row headers in the **Rows** section. You can choose one of the provided colors, reset the header text color to the default color, or create a custom color.
- 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 column field names, choose **Hide column field names**.

Columns

Hide column field names

Row height (pixels)

Auto



Background

Header background color

You can customize table headers' background color.

To customize the background color of table headers

1. In the **Format visual** 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 **Format visual** 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 **Format visual** 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 **Format visual** 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

The screenshot shows the 'Format visual' pane on the left, which is open to the 'Cells' section. The main visual area displays a table titled 'Sum of Cost by Segment and Channel'. The table has three columns: Segment, Channel, and Cost. The data is grouped by Segment and Channel. A 'Reading List' button is visible at the bottom of the pane.

Segment	Channel	Cost
Enterprise	API	1,739,410.3
SMB	API	410,281.84
Startup	API	2,621,454.13
Enterprise	Mobile	3,459,504.5
SMB	Mobile	939,104.25
Startup	Mobile	5,702,420.63
Enterprise	Web	4,661,960.09
SMB	Web	1,247,303.96
Startup	Web	7,898,458.36

Row height

You can customize table row height.

To customize the height of rows in a table or pivot table

1. In the **Format visual** 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 **Format visual** pane, choose **Cells**.

The **Cells** section expands to show options for customizing cells.

2. For **Text**, do one or more of the following:
 - 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 customize the size of the text, choose a text size. You can choose between extra small and extra large text.
 - To change the font color, choose the **Abc** color icon, and then choose a color. You can choose one of the provided colors, reset the cell text color to the default color, or create a custom color.
 - 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.

Headers ▼

Rows

Style field names



Background

Text

Medium ▼ Abc
≡ ≡ ≡
Automatic
↑ — ↓

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 **Format visual** 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 **Format visual** 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 QuickSight

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 **Format visual**, choose **Total**, and choose one of the options below.
 - In the **Rows** section, choose **Show totals** to show totals on the bottom row of the visual. Choose **Pin totals** to keep the totals visible as you scroll through the table.

- In the **Columns** section, choose **Show totals** to show totals on the last column of the visual.

The screenshot shows the Amazon QuickSight interface. On the left, the 'Dataset' is 'SPICE Business Review' at 100% zoom. Below it is the 'Fields list' with a search bar and various fields like 'Billed Amount', 'Channel', 'Cost', etc. The main area is titled 'Field wells' and shows a pivot table titled 'Revenues vs Goals'. The table has columns for 'Customer Region', 'Date', and 'Channel' (Web, Mobile, API). The rows are grouped by 'Customer Region' (APAC and EMEA) and 'Date' (2012-2016). The table shows revenue values for each combination of region and date across the three channels.

Customer Region	Date	Channel		
		Web Billed ...	Mobile Billed ...	API Billed ...
APAC	2012	\$141,419.15	\$100,566.79	\$42,777.10
	2013	\$303,371.82	\$210,949.34	\$90,286.49
	2014	\$598,770.06	\$410,354.34	\$186,005.91
	2015	\$1,136,687.48	\$794,905.82	\$360,931.11
	2016	\$1,997,075.29	\$1,378,825.72	\$637,727.92
EMEA	2012	\$163,883.27	\$114,805.38	\$45,919.83
	2013	\$365,314.83	\$261,725.61	\$107,455.64
	2014	\$932,084.13	\$645,923.78	\$301,930.42
	2015	\$1,511,398.73	\$1,046,015.37	\$530,805.86
	2016	\$2,461,444.97	\$1,701,784.31	\$847,672.17

2. To display subtotals, choose **Subtotal** and choose one of the options below.

- In the **Rows** section, choose **Show subtotals** to show subtotals on rows.
- In the **Columns** section, choose **Show subtotals** to show subtotals on columns.
- In the **Level** section, 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.

The screenshot shows the Amazon QuickSight interface. On the left, the 'Dataset' is 'SPICE Business Review' and the 'Fields list' includes 'Billed Amount', 'Channel', 'Cost', 'Customer ID', 'Customer Name', 'Customer Region', 'Date', 'Distinct ID', 'Revenue Goal', 'Segment', and 'Service Line'. The main area displays a table visualization titled 'Revenues vs Goals' with columns for 'Date', 'Channel', 'Service Line', and 'Billed ...'. The table data is as follows:

Date	Channel	Service Line	Billed ...
2012	API	Billing	\$88,118.04
		HR	\$84,106.02
	Mobile	Billing	\$203,872.24
		HR	\$204,240.45
	Web	Billing	\$294,528.47
		HR	\$292,196.51
2013	API	Billing	\$208,658.63
		HR	\$214,378.28
		Marketing	\$21,157.13
	Mobile	Billing	\$498,719.75
		HR	\$498,224.63
		Marketing	\$52,830.18
	Web	Billing	\$705,197.96
		HR	\$726,099.36
		Marketing	\$72,233.94
2014	API	Billing	\$416,224.30
		HR	\$404,326.33

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 **Format visual** 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 **Format visual** 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 **Format visual** pane, choose **Total** or **Subtotal**.
2. For **Text**, do one or more of the following.
 - To customize the size of the text, choose a text size. You can choose between extra small and extra large text.
 - To change the font color, choose the **Abc** color icon, and then choose a color. You can choose one of the provided colors, reset the cell text color to the default color, or create a custom color.

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 **Format visual** 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 **Format visual** 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 **Format visual** 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 QuickSight

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.

Industry	Product	Region		
		EMEA	APJ	AMER
		Sales	Sales	Sales
☐ Communications	Big Ol Database	5,178.58	287.91	3,617.66
	ChatBot Plugin	550.1	94.67	490.45
	ContactMatcher	14,814.21	1,984.05	10,844.09
	Data Smasher	7,911.58	1,864.1	3,418.41
	FinanceHub	7,811.35	1,203.52	7,866.2
	Marketing Suite	1,359.89	362.94	1,668.13
	Marketing Suite ...	4,816.37	2,561.22	3,795.39
	OneView	3,051.7	796.06	1,527.89
	SaaS Connector ...	2,769.74	1,328.2	1,785.16
	SaaS Connector ...	481.34	362.46	195.36
	Site Analytics	11,169.85	1,380.05	6,720.26
	Storage	98.31	25.25	102.18
	Support	4,414.38	1,265.26	987.94
☐ Consumer Products	Alchemy	5,379.87	20,459.89	13,999.96
	Big Ol Database	836.08	868.8	2,966.2

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.

Sum of Sales by Industry, Product, and Region

Industry	Product	Region		
		EMEA	APJ	AMER
		Sales	Sales	Sales
☐ Communications	Big Ol Database	5,178.58	287.91	3,617.66
	ChatBot Plugin	550.1	94.67	490.45
	ContactMatcher	14,814.21	1,984.05	10,844.09
	Data Smasher	7,911.58	1,864.1	3,418.41
	FinanceHub	7,811.35	1,203.52	7,866.2
	Marketing Suite	1,359.89	362.94	1,668.13
	Marketing Suite ...	4,816.37	2,561.22	3,795.39
	OneView	3,051.7	796.06	1,527.89
	SaaS Connector ...	2,769.74	1,328.2	1,785.16
	SaaS Connector ...	481.34	362.46	195.36
	Site Analytics	11,169.85	1,380.05	6,720.26
	Storage	98.31	25.25	102.18
	Support	4,414.38	1,265.26	987.94
☐ Consumer Products	Alchemy	5,379.87	20,459.89	13,999.96
	Big Ol Database	836.08	868.8	2,966.2

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.

Headers

Cells

Wrap text

Row height (pixels)

46

Column width (pixels)

100

Expand to view

Text

Medium

Automatic

Customize pivot table data

You can customize how QuickSight 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 QuickSight authors remove clutter from their pivot tables and provide an easier reader experience for QuickSight 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.
 - **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.

Pivot options



- Hide +/- buttons
- Hide single metric
- Hide collapsed columns

Expand to view

Adding data bars to tables in QuickSight

You can use data bars to add visual context to your table visuals in Amazon QuickSight. 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.

Sum of Sales (Sum) and Sum of Profit (Sum) by Industry

industry	sales (SUM)	profit (SUM)
Agriculture	166,710.24	-34,659.2
Communications	1,177,239.36	110,953.52
Consumer Products	1,738,051.36	411,494.66
Energy	2,399,247.52	344,129.9
Finance	4,193,206.24	508,255.05
Healthcare	1,904,002.08	195,355.16
Industrial	253,124.96	55,607.64
Misc	1,991,786.24	222,764.91
Other	115,406.88	7,703.52
Retail	3,089,122.08	452,895.46
Tech	2,781,517.28	347,760.8
Transportation	1,691,424.16	188,505.49

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 **Format visual** pane, open the **Visuals** dropdown list and choose **ADD DATA BARS**.

Visuals



Data bars
sales (SUM) (Sum)

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.

Data bars



Value field

Select field



Positive color



Negative color

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.

Data bars



Value field

Select field



Positive color

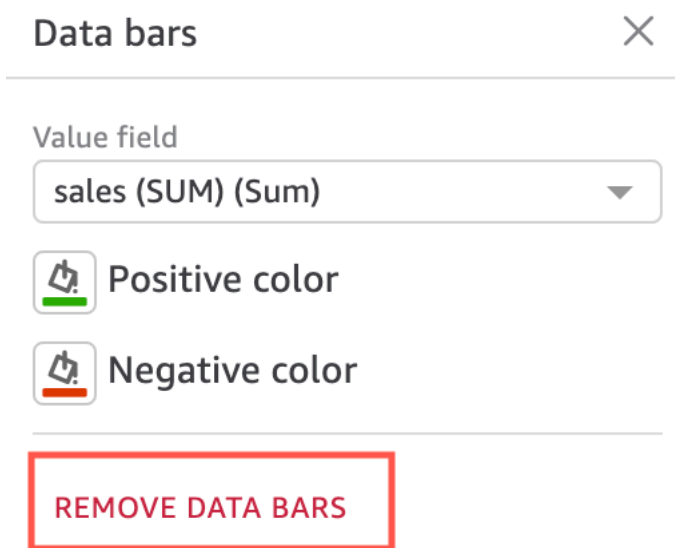


Negative color

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 **Format visual** 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 **Format visual** pane opens.
2. In the **Format visual** 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 QuickSight

In QuickSight, you can choose from multiple formatting options for your maps and geospatial charts. You can view formatting options by opening the **Format visual** pane from the on-visual menu located at the top right of the currently selected geospatial map.

Format visual	×
Title	^
Legend	^
Points	∨
Style	
<input checked="" type="radio"/> Basic points	
<input type="radio"/> Cluster points	
<input type="radio"/> Heatmap	
Base map	^
Tooltip	^

QuickSight authors and readers can also toggle the different formatting options of a geospatial map visual from the on visual menu.



Topics

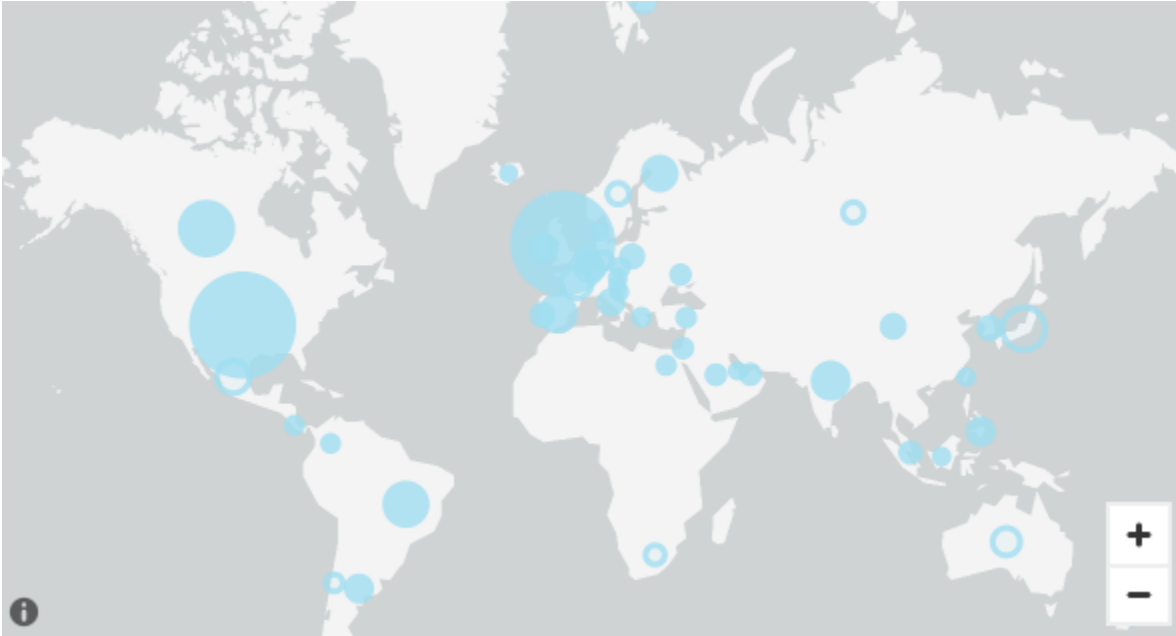
- [Base maps on geospatial maps in QuickSight](#)
- [Geospatial heatmaps in Amazon QuickSight](#)
- [Marker clustering on geospatial point maps in QuickSight](#)

Base maps on geospatial maps in QuickSight

When you create a map visual in Amazon QuickSight, 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 QuickSight, there are four options for base maps: light gray canvas, dark gray canvas, streets, and imagery. An example of each appears following:

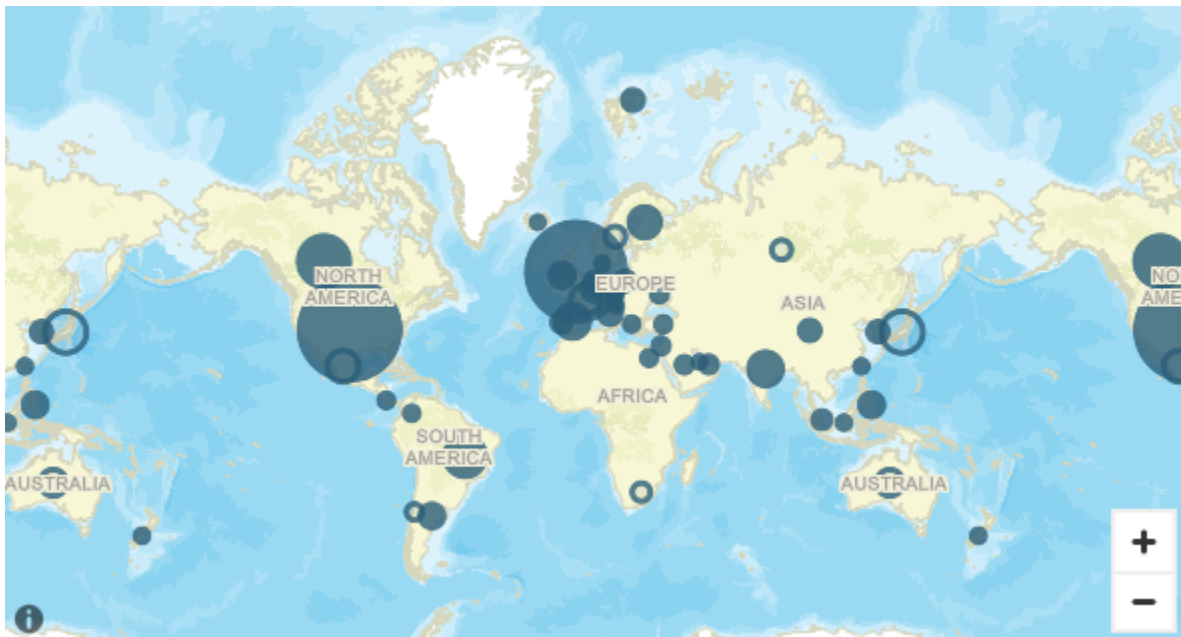
- Light gray canvas




- 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 **Format visual** pane that opens at left, choose the **Base map** section and then choose the base map that you want.

Geospatial heatmaps in Amazon QuickSight

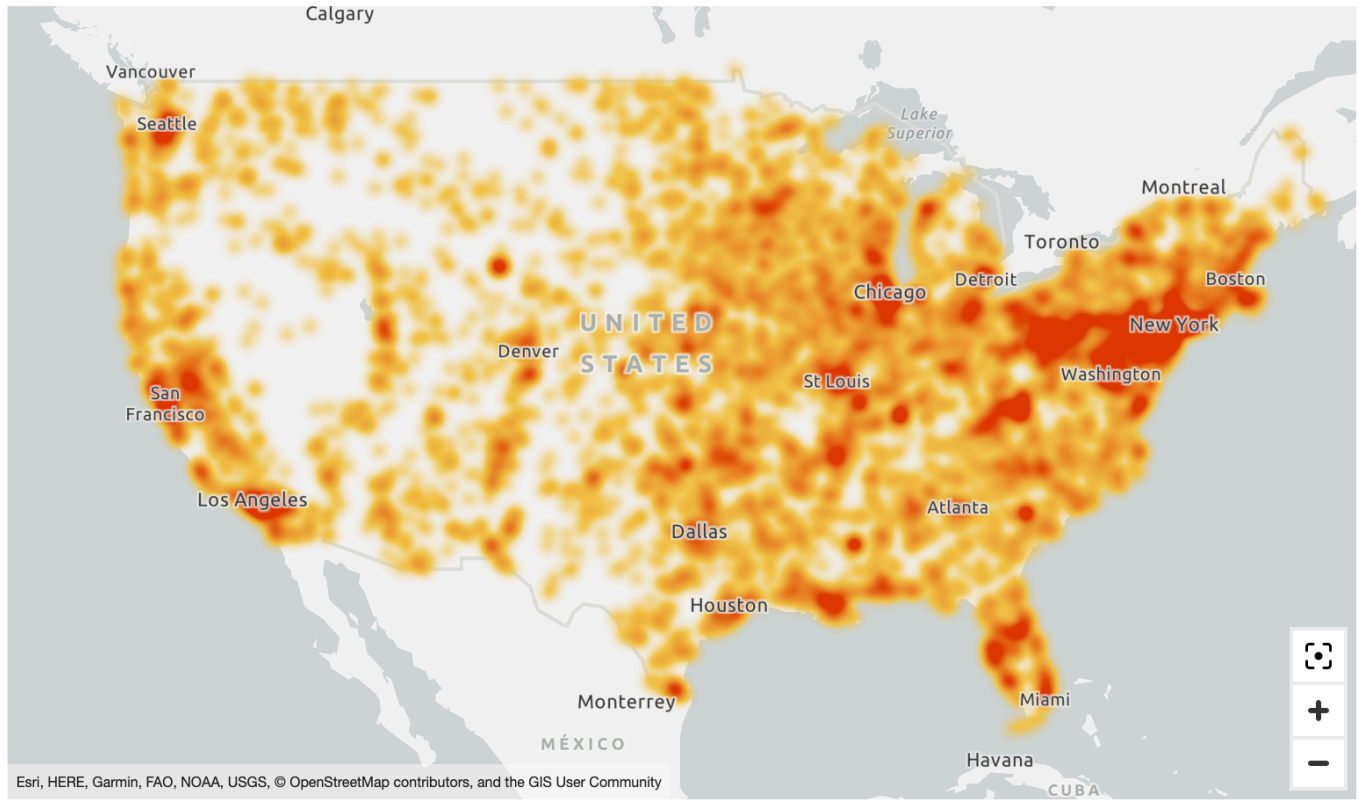
Use geospatial heatmaps to reveal patterns of marker concentration in your geospatial visuals. Heat maps display concentrations of data points using a colored overlay that highlights the intensity or concentration of the visual's markers.

Field wells

Sheet 1 ▾ +

Count of Records by City

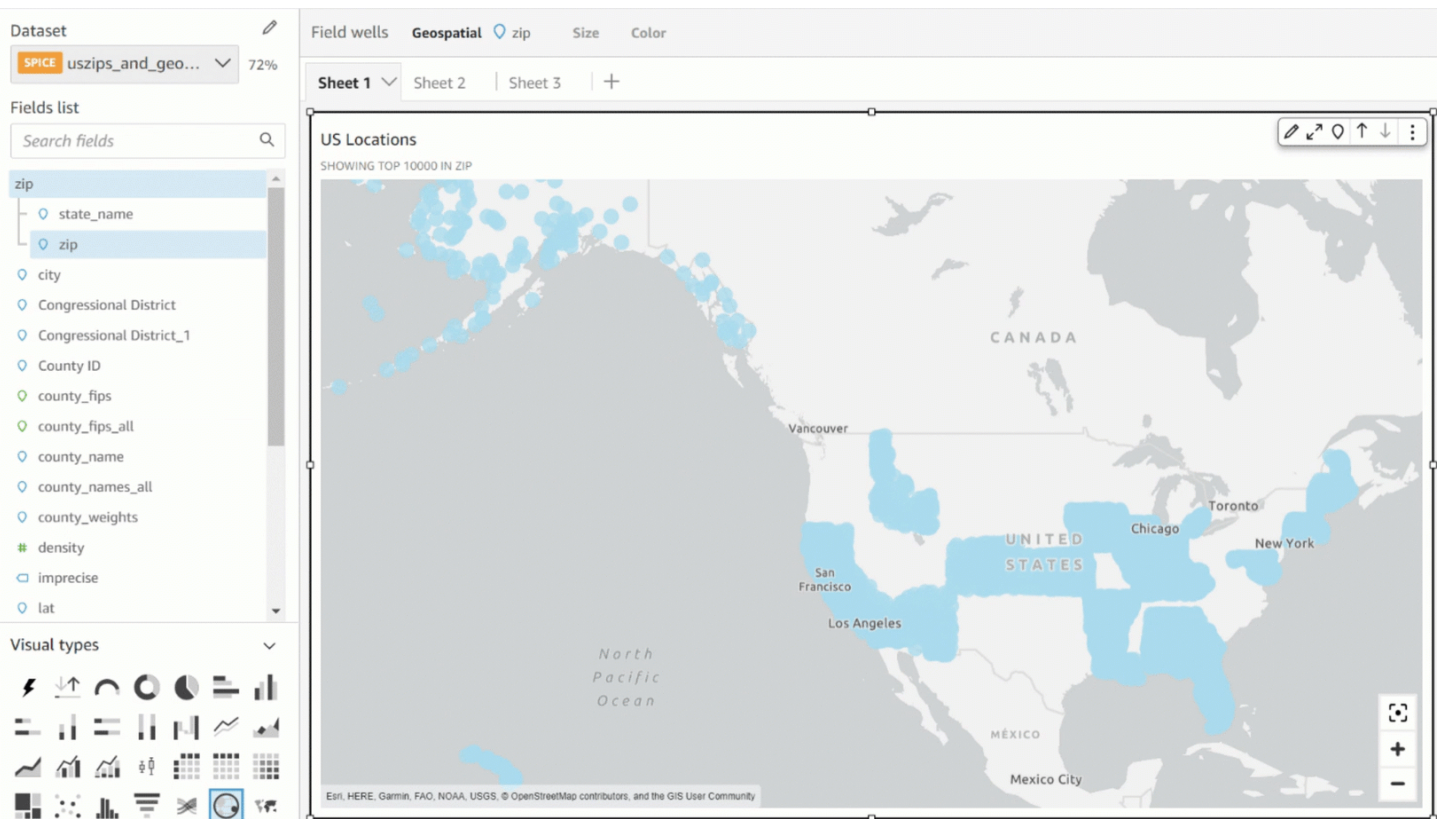
SHOWING TOP 10000 IN CITY

**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 **Format visual** 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 QuickSight

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 **Format visual** 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 QuickSight


When you create a chart in Amazon QuickSight, 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 **Format visual** pane opens at left.

To show or hide axis lines

1. In the **Format visual** 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 **Format visual** 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 show or hide the sort icon

1. In the **Format visual** 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 **Format visual** 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 **Format visual** 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 **Format visual** pane, choose the axis that you want to format.
2. For **Label size**, choose a size.

To change the label orientation

1. In the **Format visual** 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 **Format visual** 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 QuickSight

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

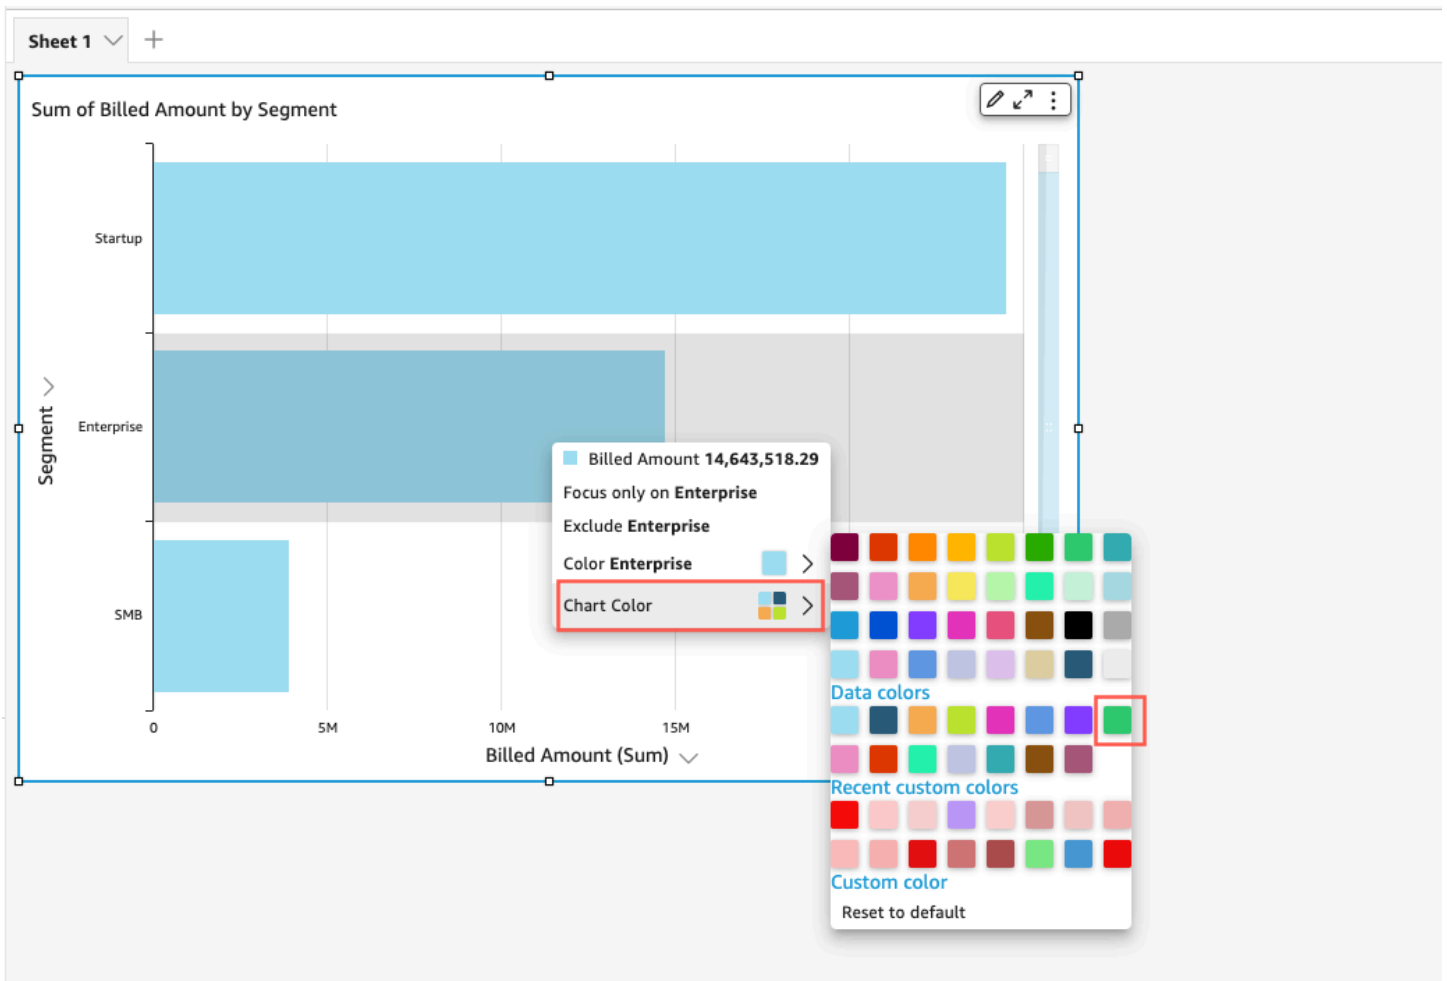
To change colors on bar charts, donut charts, gauge charts, line charts, and scatter plots, see [Changing colors on charts](#).

To change colors on heat maps and tree maps, see [Changing colors on heat maps and tree maps](#).

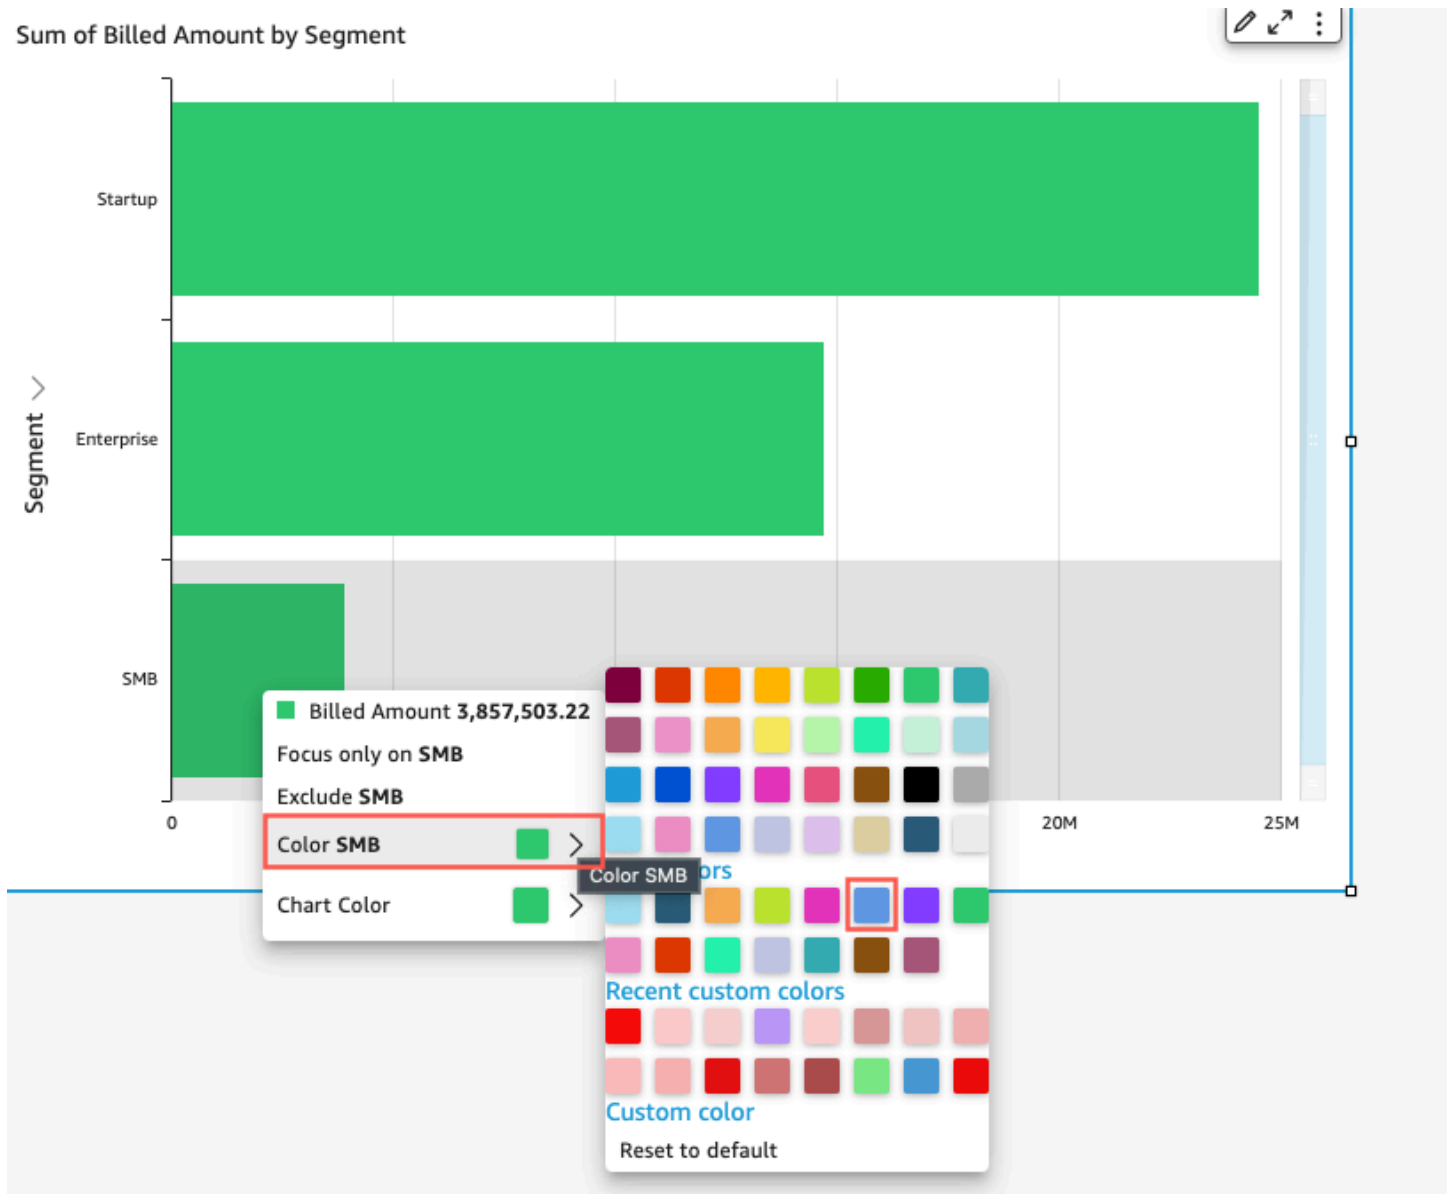
Changing colors on charts

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.

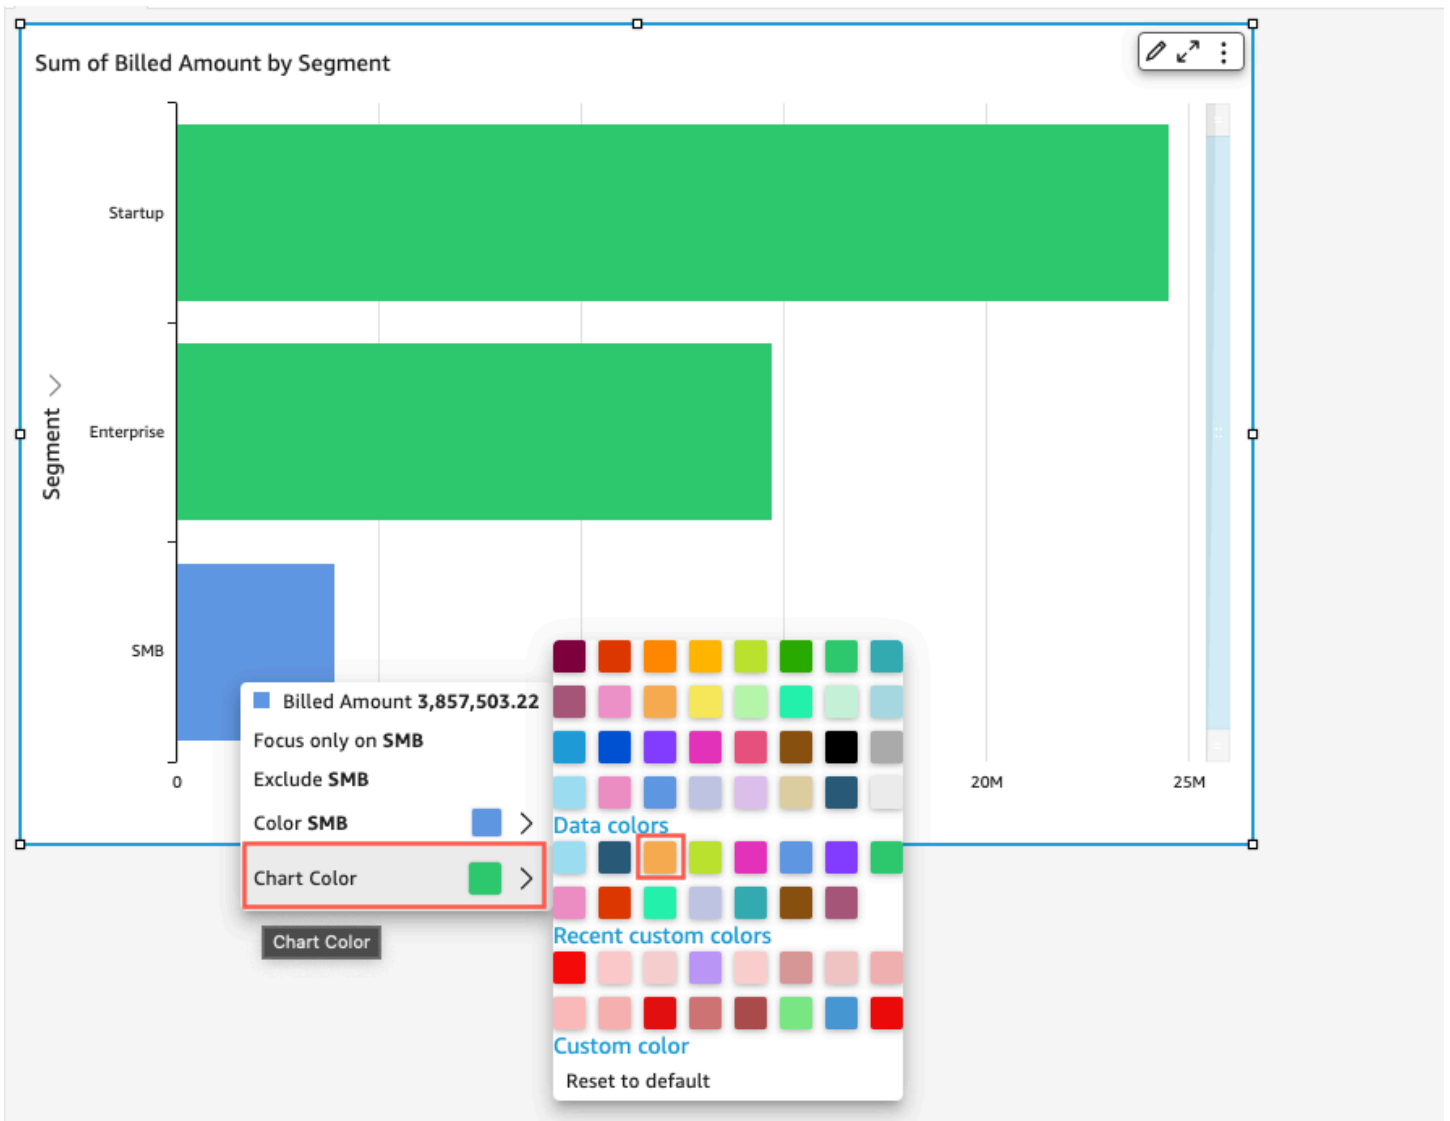
For example, suppose that you set the chart color to green.




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.



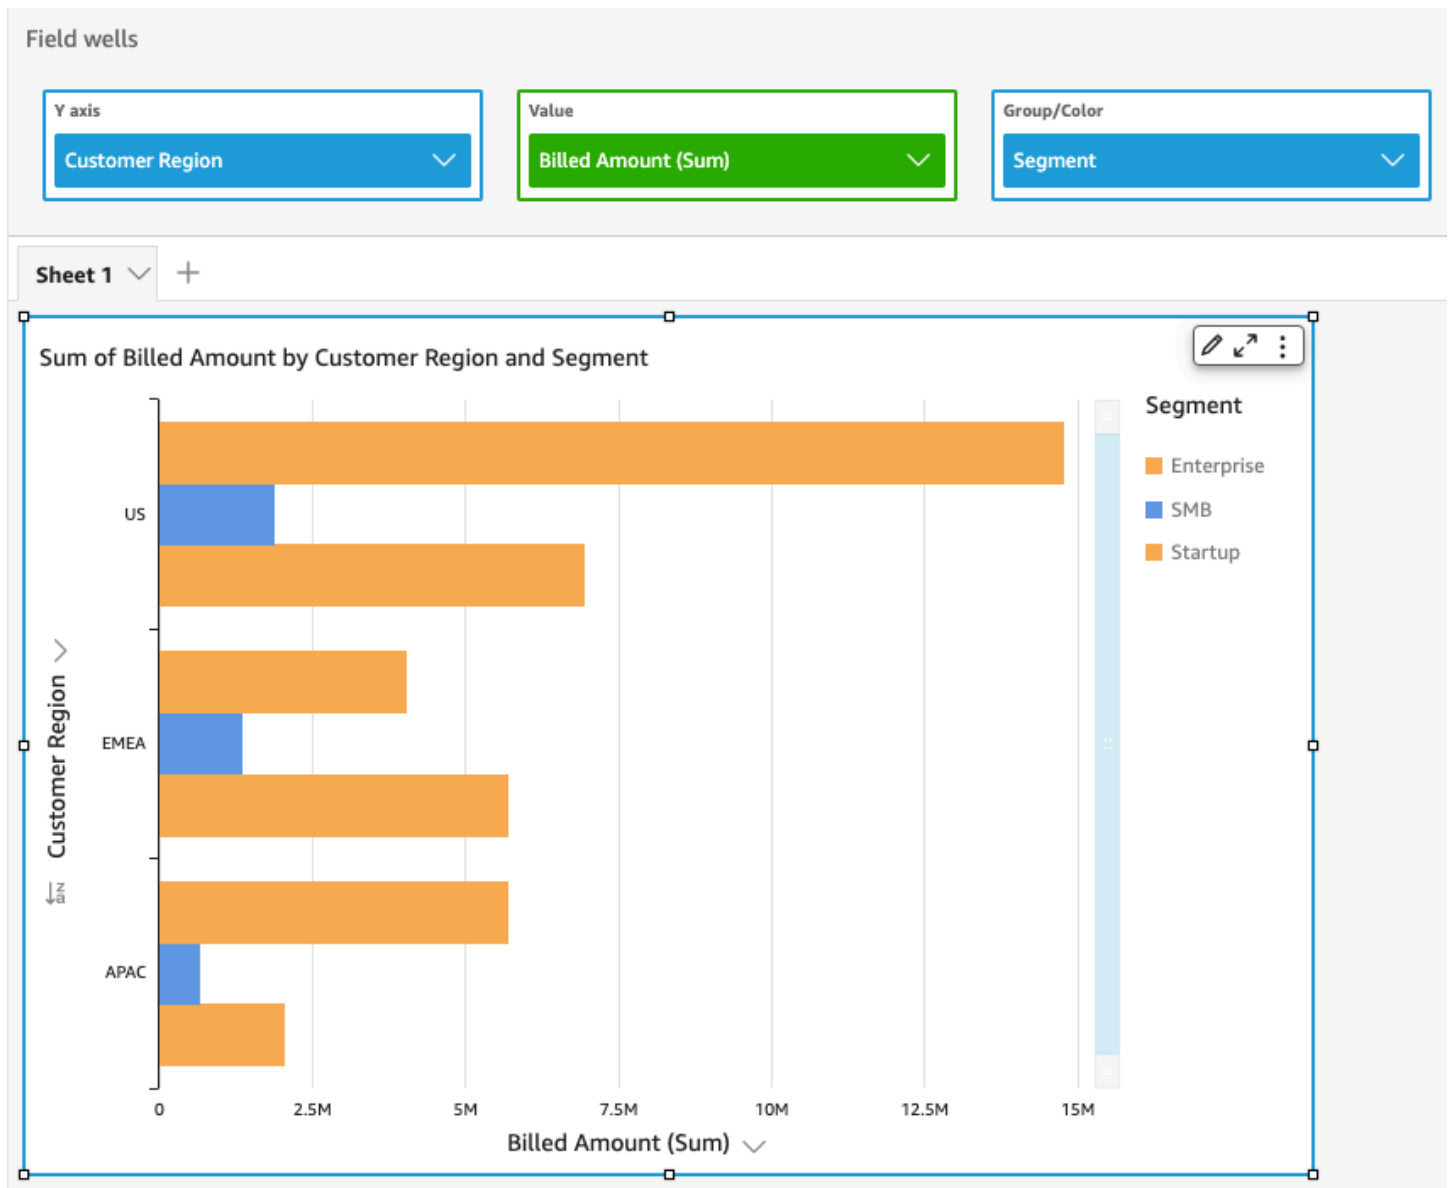
Looking at the result, you decide that you need more contrast between the green and blue bars, so you change the chart color to orange. If you are changing the chart color, it doesn't matter which bar you choose to open the context menu from.



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.

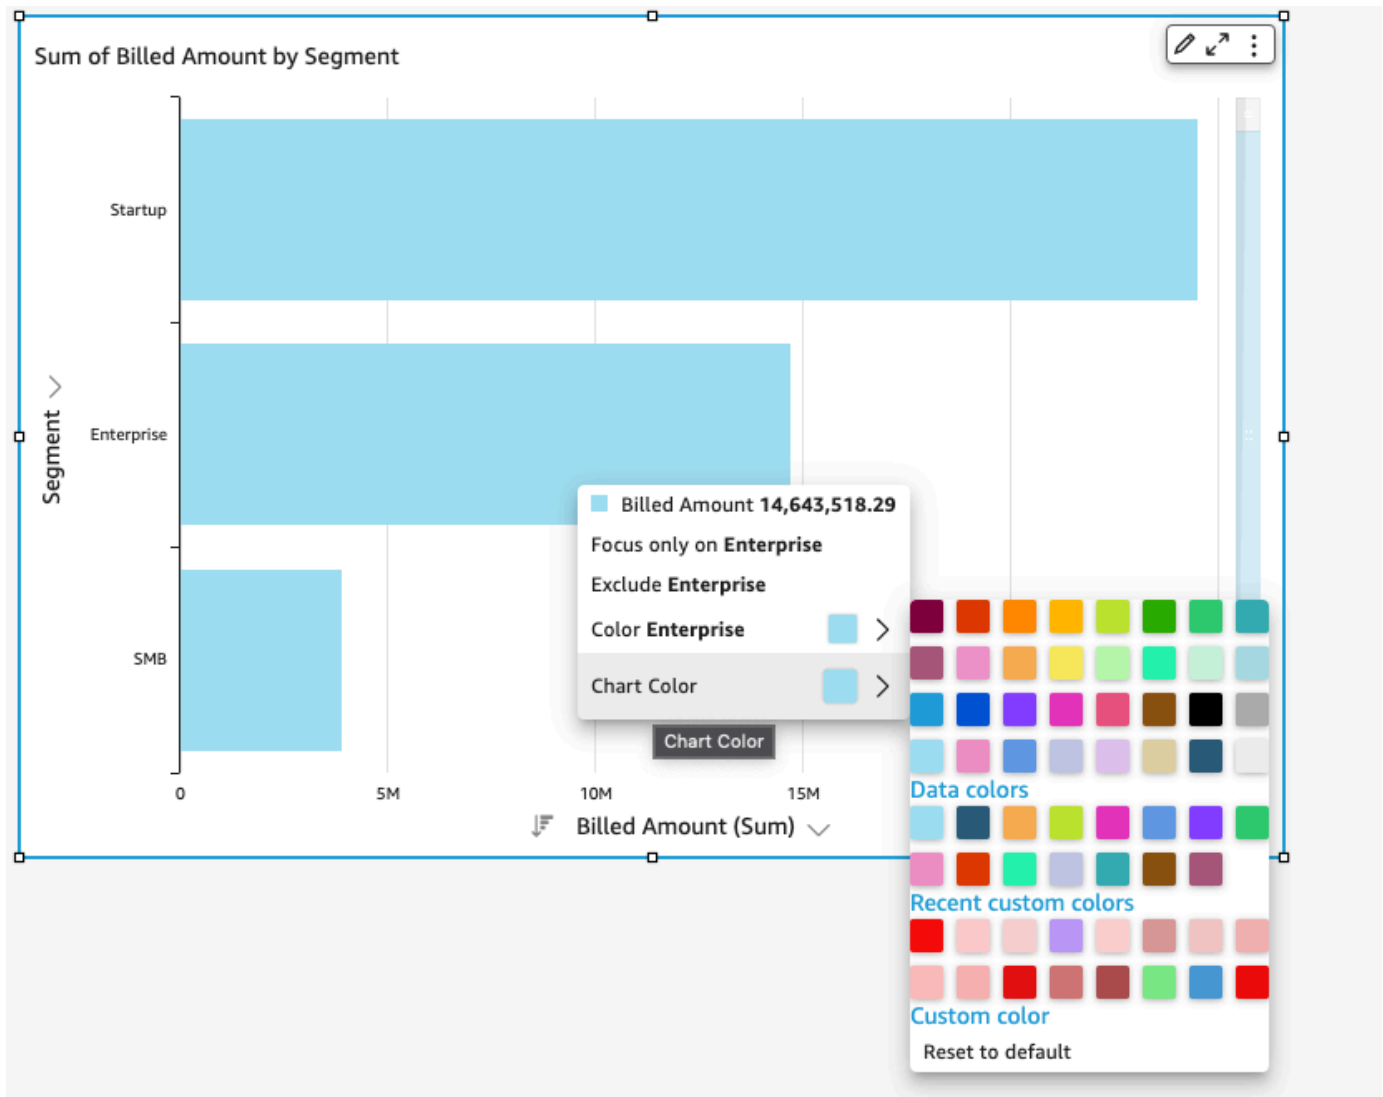
To change the colors for a visual

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.



- 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 **Format visual** 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.

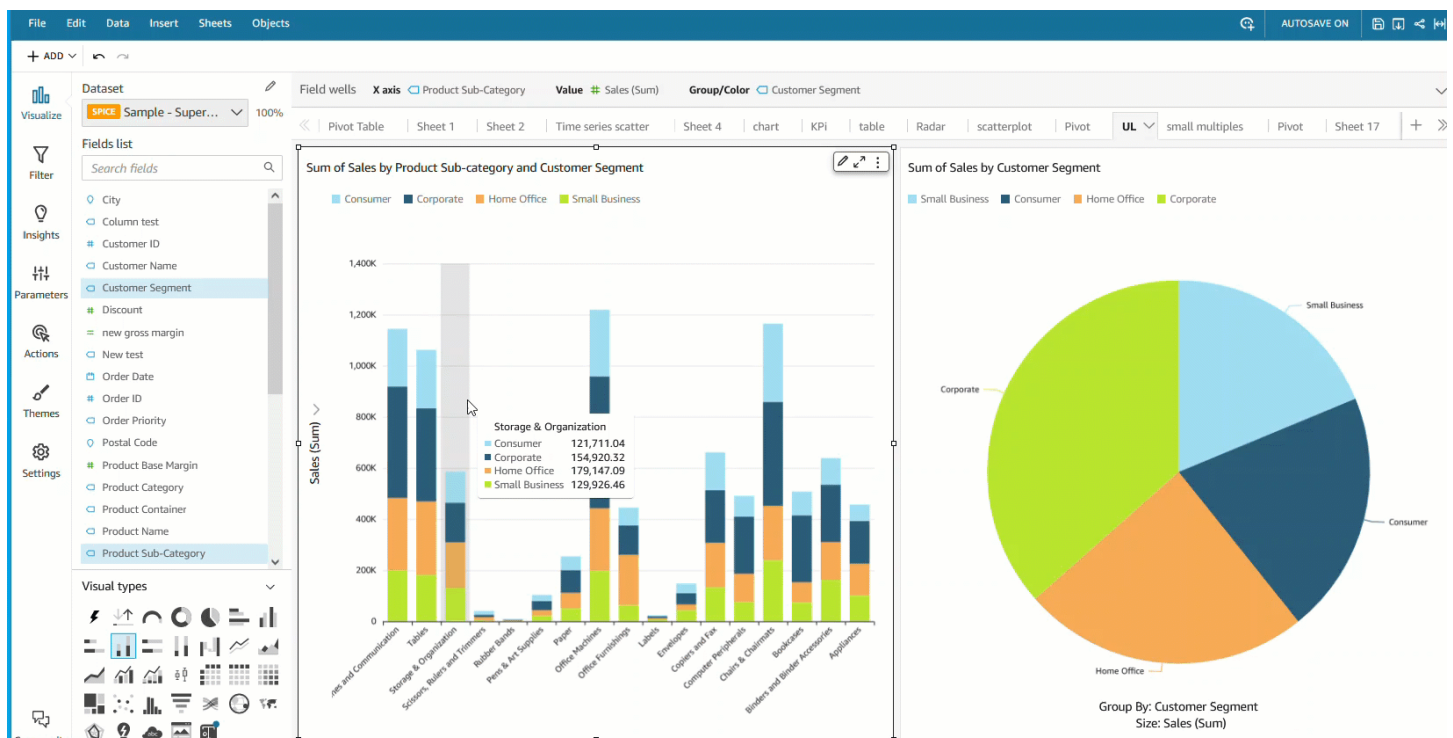
You can enter a number that defines the midpoint between the two main gradient colors. If you add a value, the middle color represents the number you entered. If you leave this blank, the middle color acts like the other colors in the gradient.

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 QuickSight

With field level coloring, you can assign specific colors to specific field values across all visuals in a QuickSight 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.

QuickSight 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 QuickSight

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**.

Font and style on visual types in QuickSight

You can choose from several options for styling, including fitting the table to your current view, hiding column field names, changing font sizes. You can also choose to hide the metric label when you use a single metric, to avoid seeing the same metric label repeated on the pivot table.

To customize the styling for a visualization

1. On the analysis page, choose the chart that you want to format.

2. Choose the menu on the visualization



and then choose **Format visual**



3. Choose **Styling**.

4. To prevent displaying a single metric label repeatedly, enable **Hide single metric**.

5. To hide labels for fields in the **Columns** field well, enable **Hide column field names**.

6. To hide collapse



and expand icons



enable **Hide +/- buttons**.

7. To expand the table to fill your current view, choose **Fit table to view**. You can't undo this action.

To shrink the table to fit your current view, you can adjust the width of each column. To do this, grab the right edge of the column, near the column title. Drag the edge in either direction.

Startup	SMB
Billed Amount	Billed Amount
102,832.76	47,500.14
357,995.33	44,803.39
393,575.11	36,607.5
242,010.98	40,337.72

8. Choose your preferred font size for each of the following chart types:
- For KPIs, choose the font sizes for the primary and secondary values.
 - For pivot tables and tables, choose the font sizes for table headers, cells, totals, and subtotals.
 - For other visuals, you can choose font sizes depending on which chart type you are using.

KPI options

You can customize KPIs in Amazon QuickSight 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 QuickSight, navigate to the KPI that you want to change and choose **Format visual** to open the **Format visual**. The icon for the **Format visual** menu is as follows.



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 QuickSight. 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 **Format visual** 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.

- (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.
- (Optional) Choose **Add tooltip** to add a tooltip to the KPI visual.

The following image shows the **Visual** section of the **Format visual** menu.

The screenshot displays the 'Format visual' menu for a KPI. The KPI visual shows 'Sum of Population by Region' with a primary value of 16,744,206 and an upward trend arrow. A secondary value of 84,673,176 is shown for 'South' and 67,928,970 for 'West'. The 'Visual' section of the menu is highlighted with a red box, showing 'Area sparkline' selected and 'Tooltip' unchecked.

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

- 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.

The following images shows the **KPI** menu.

The image shows a configuration menu for a KPI. At the top, the word "KPI" is displayed with an upward-pointing chevron icon. Below this, the "Primary value" section is visible. It includes a "Font" dropdown menu currently set to "Auto" and a color selection icon labeled "Abc". Underneath, the "Primary value displayed" section has two radio button options: "Actual" (which is selected) and "Comparison". A horizontal line separates this from the "Secondary value" section, which is currently checked with a square box. This section also includes a "Font" dropdown menu set to "Extra large" and a color selection icon labeled "Abc".


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 **Format visual** menu.

To change the color of positive and negative values

1. In the **Format visual** menu, 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. The following image shows the **Conditional formatting** menu of a KPI.




Labels on visual types in QuickSight

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 **Format visual**.

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 **Format Visual** panel, which you open by choosing the **V**-shaped icon at top right.

3. On the **Format Visual** pane, enable or disable **Show title**. This option removes the axis title.
4. Close the **Format Visual** pane by choosing the **X** icon in the upper-right corner of the pane.

Data labels on visual types in QuickSight

To customize data labels on a visual, you can use the **Format Visual** pane to show data labels, and then use the settings to configure them. Data label customization is supported on bar, line, combo, scatter, 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
formatting

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

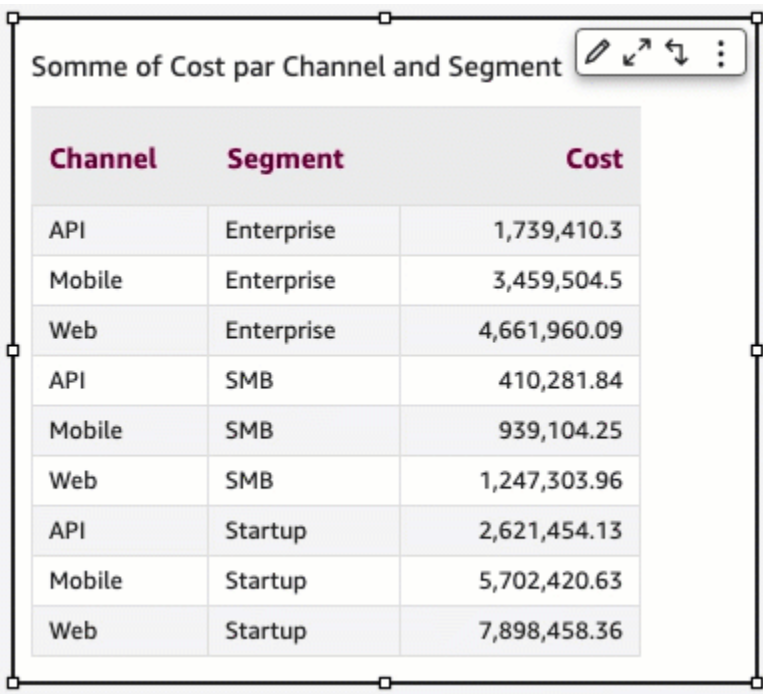
1. On the analysis page, choose the visual that you want to format.
2. Choose the on-visual menu from the down icon at the upper-right corner of the visual, and then choose **Format visual**.
3. On the **Format Visual** pane, choose **Data Labels**.
4. Enable **Show data labels** to show and customize 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. Close the **Format Visual** pane by choosing the X icon in the upper-right corner of the pane.

Formatting visual numeric data based on language settings in QuickSight


In Amazon QuickSight, you can choose how your numeric data values appear in visuals so that they align with the regional language that you have chosen.



The image shows a screenshot of a table in Amazon QuickSight. The table is titled "Somme of Cost par Channel and Segment" and has three columns: "Channel", "Segment", and "Cost". The data is organized into three groups based on the Segment: Enterprise, SMB, and Startup. Each group contains three rows corresponding to the Channel: API, Mobile, and Web. The "Cost" column displays numerical values with varying decimal places, indicating different currency or unit formats. A toolbar with icons for edit, zoom, and refresh is visible in the top right corner of the table.

Channel	Segment	Cost
API	Enterprise	1,739,410.3
Mobile	Enterprise	3,459,504.5
Web	Enterprise	4,661,960.09
API	SMB	410,281.84
Mobile	SMB	939,104.25
Web	SMB	1,247,303.96
API	Startup	2,621,454.13
Mobile	Startup	5,702,420.63
Web	Startup	7,898,458.36

As a QuickSight author, you can choose the language formatting that best fits your audience. Amazon QuickSight configures numeric data languages at the analysis level based on the language that you have chosen to view QuickSight in. You can change the format of numbers, currencies, and dates. You can change your QuickSight language settings in the **Language** dropdown list of the QuickSight **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 **Visualize** 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**.

The screenshot shows the Amazon QuickSight interface. On the left, the 'Dataset' is 'SPICE Business Review' at 100% zoom. The 'Fields list' contains various fields, with 'Cost' highlighted. On the right, the 'Field wells' section shows a visualization titled 'Sum of Cost by Channel and Segment'. A context menu is open over the 'Cost' field, displaying options: 'Add to visual', 'Add filter for this field', 'Show as: Number', 'Format: 1,234.57', and 'Convert to dimension'. The 'Format: 1,234.57' option is selected, and a sub-menu is open showing different numeric formats: '1,234.57', '1234.57', '1,235', and '1235'. A 'More formatting options...' link is also present in the sub-menu.

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.

Apply language format

Reset to defaults

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**.

Sum of Cost by Channel and Segment

Channel	Segment	Cost
API	Enterprise	\$1,739,4
Mobile	Enterprise	\$3,459,50
Web	Enterprise	\$4,661,90
API	SMB	\$410,28
Mobile	SMB	\$939,10
Web	SMB	\$1,247,30
API	Startup	\$2,621,49
Mobile	Startup	\$5,702,42
Web	Startup	\$7,898,49

Aggregate: **Sum**

Sort by

Show as: **Currency**


Format: \$1,234.57

- \$1,234.57 Dollar
- £1,234.57 Pound
- €1,234.57 Euro
- ¥1,234.57 Yen
- ₩1,234.57 Won
- kr1,234.57 Krone
- NT\$1,234.57 TWD
- ₹1,234.57 Rupee
- More formatting options...

Search fields

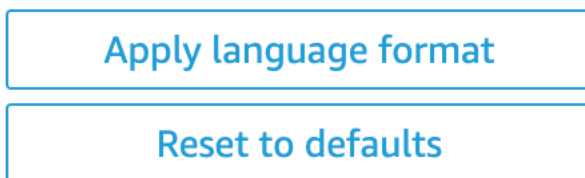
Billed Amount

- 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 QuickSight

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.

To display or hide a visual legend

1. Sign in to Amazon QuickSight at <https://quicksight.aws.amazon.com/>.
2. On the analysis page, choose the visual that you want to format.

3. At the visual's right, for **Menu options**



choose **Hide legend** or **Show legend** to hide or show the visual legend.

When shown, the legend displays the values in alphabetical order.

To customize a visual legend

1. At the visual's right, choose **Format Visual**



2. In the **Format Visual** pane, expand the **Legend** section. In this section, you can format the legend title and position.

3. Choose the **X** icon at upper right to close the **Format Visual** pane.

Line and marker styling on line charts in QuickSight

In QuickSight 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**.

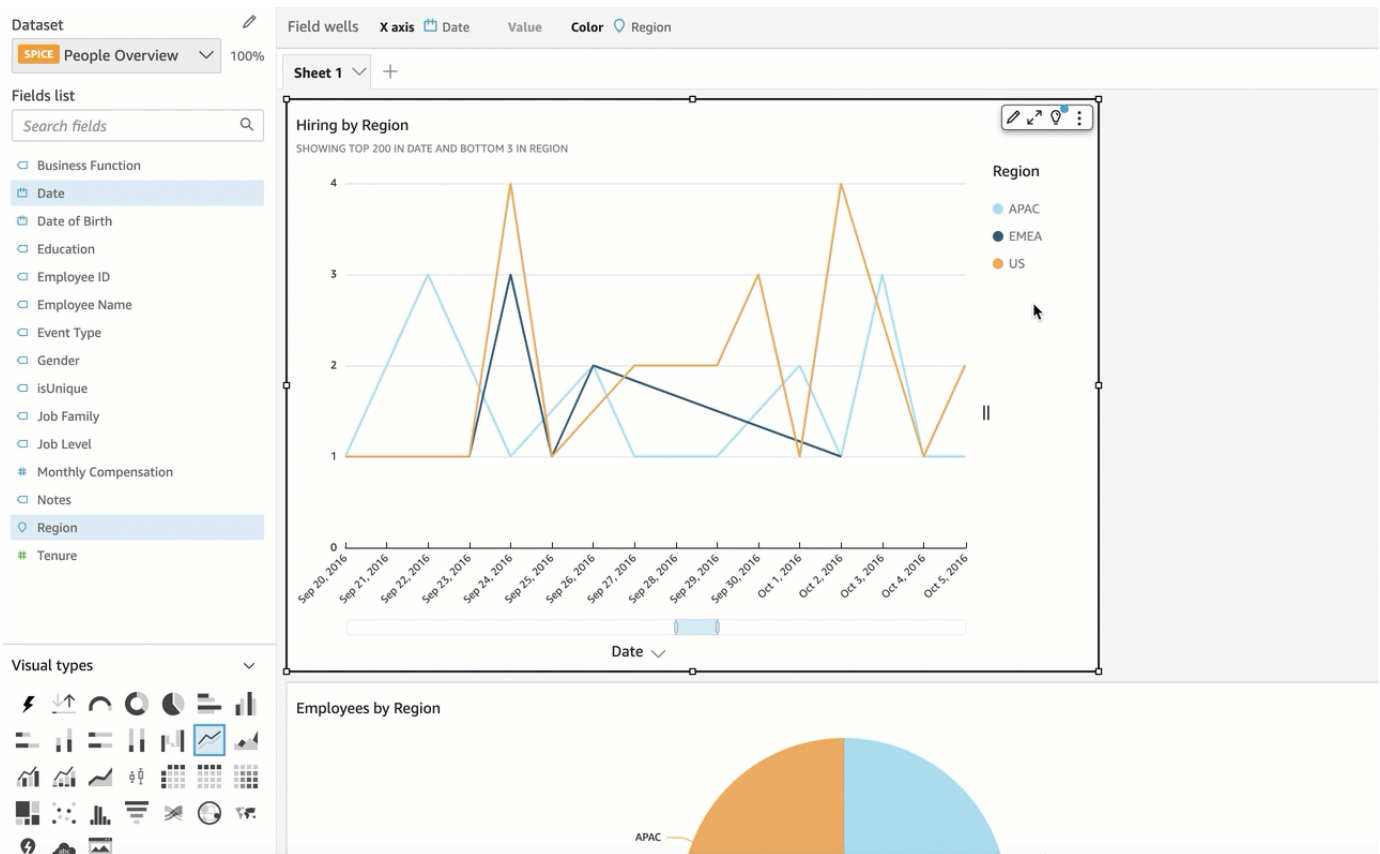
Missing data on visual types in QuickSight

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.
- **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 QuickSight

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.



Amazon QuickSight 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. Choose your visualization so that it's highlighted and its menu appears. To open the formatting options, choose the cog icon



icon on the visualization menu



2. Open the **Reference lines** section.
3. Add or edit a reference line. To add a reference line, choose **Add new line**. To edit a reference line, choose the line to edit.
4. With the reference line settings open in the **Format** pane at left, you can modify its properties:

- **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.
5. Choose **Done** to save your selections.

To list existing reference lines

1. Choose your visualization so that it's highlighted and its menu appears. To open the formatting options, choose the cog icon



icon on the visualization menu



2. Open the **Reference lines** section.

Existing reference lines appear in a list, followed by an **Add new line** button.

To disable a reference line

1. Choose your visualization so that it's highlighted and its menu appears. To open the formatting options, choose the cog icon



icon on the visualization menu



2. Open the **Reference lines** section.

3. Choose **Disable** from the context menu (...) for the reference line that you want to disable.

To delete a reference line

1. Choose your visualization so that it's highlighted and its menu appears. To open the formatting options, choose the cog icon



icon on the visualization menu



2. Open the **Reference lines** section.

3. Choose **Delete** from the context menu (...) for the reference line that you want to delete.

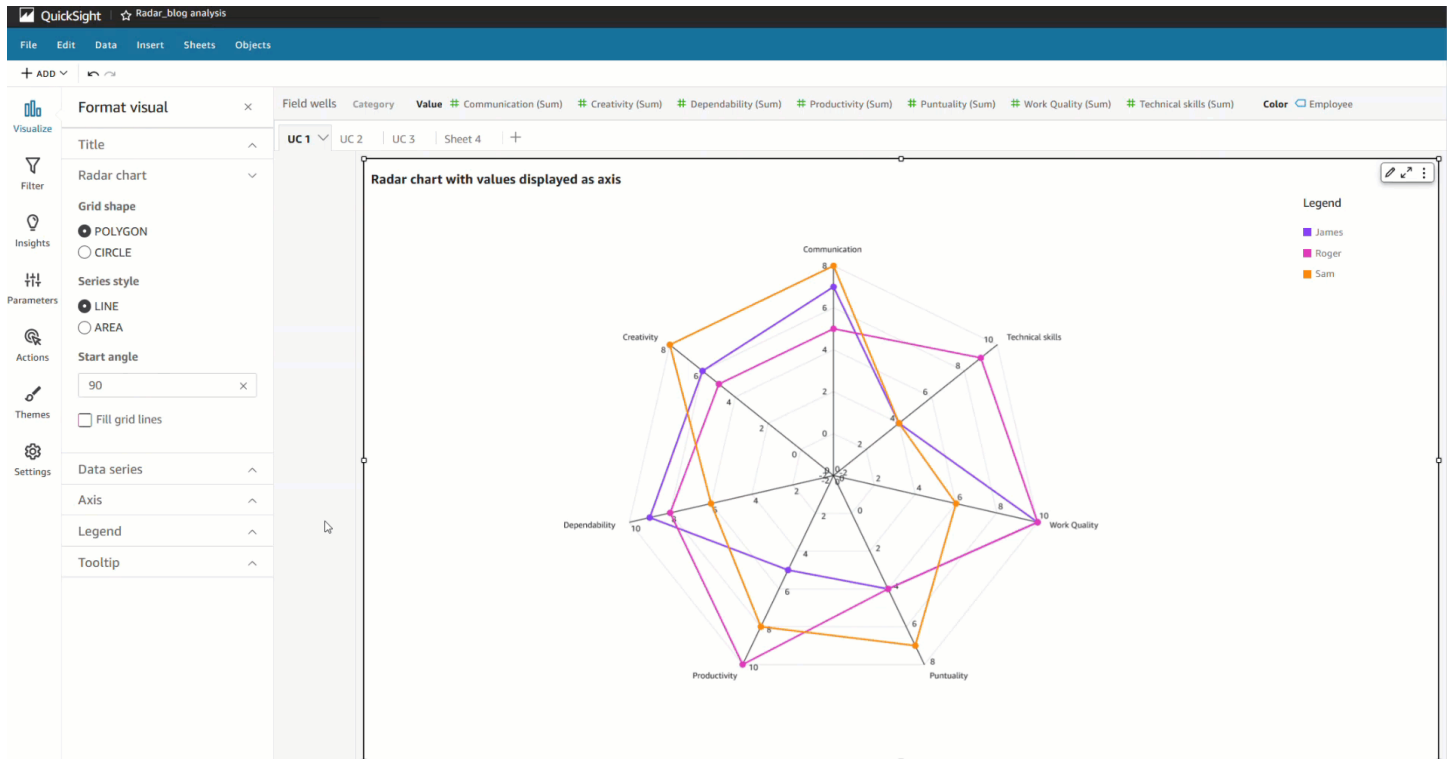
Formatting radar charts in QuickSight

You can customize radar charts in Amazon QuickSight 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 **Format visual** 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**.



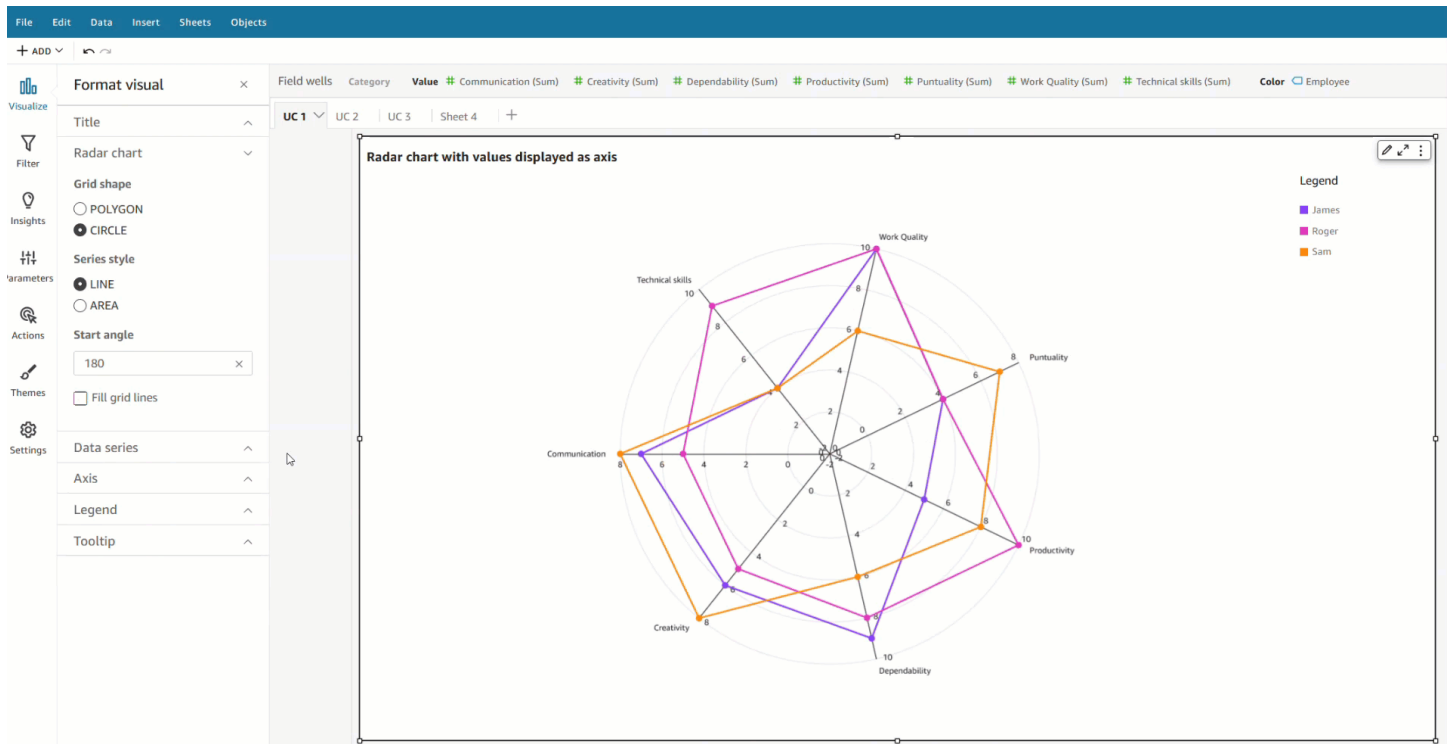
To choose the start angle of a radar chart

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 **Format visual** 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 **Format visual** pane on the left, open the **Radar chart** 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 **Format visual** 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 QuickSight

To change the scale of the values shown on the visual, you can use the **Format Visual** 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 **Format Visual** 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. QuickSight 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 **Format Visual** 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 **Format visual** 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 **Format visual** 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 QuickSight in QuickSight


In Amazon QuickSight, you can format visual titles and subtitles to meet your business needs. QuickSight offers rich text formatting for titles and subtitles, and the ability to add hyperlinks and parameters in titles. You can edit titles in the Format visual pane, or by double-clicking on a title or subtitle in the visual.

Showing or hiding visual titles or subtitles

Use the following procedure to hide or display the title or subtitle for a visual. The visual title is shown by default. After subtitles are created, they're also shown by default.


1. Sign in to Amazon QuickSight at <https://quicksight.aws.amazon.com/>.
2. On the analysis page, choose the visual that you want to format.
3. At the visual's right, choose the **Format visual** icon.

Sum of Cost by Channel and Segment



4. In the **Format Visual** pane that opens at left, choose the **Title** tab, and then choose from the following settings:
 - Select or clear the **Show title** check box.
 - Select or clear the **Show subtitle** check box.

Editing visual titles or subtitles



Use the following procedure to edit the title or subtitle for a visual.

1. Sign in to Amazon QuickSight at <https://quicksight.aws.amazon.com/>.
2. On the analysis page, choose the visual that you want to format and double-click on the title or subtitle in the visual.

Or, if your titles or subtitles are hidden, you can do the following:

- a. At the visual's right, choose the **Format visual** icon.
 - b. In the **Format Visual** pane that opens at left, choose the **Title** tab.
 - c. Choose **Edit title** or **Edit subtitle**.
3. In the **Edit title** or **Edit subtitle** page that opens, highlight the text that you want to edit, and then choose from the following options:

- 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 QuickSight](#).
- 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 QuickSight, 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**.

4. When you're finished editing, choose **Save**.

Tooltips on visual types in QuickSight

When you hover your cursor over any graphical element in an Amazon QuickSight 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. Use the following procedures to learn how.




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.

Sum of Cost by Channel and Segment



3. In the **Format visual** pane that opens at left, choose **Tooltip**.
4. For **Type**, choose **Detailed tooltip**.

A new set of options appear, as shown following.

Tooltip ⌵

Show tooltip

Type

Detailed tooltip ⌵

Display options

Use primary value as title

Show aggregations

Fields

📅 Date ⋮

📁 Opportunity Stage ⋮

== Count ⋮

[Add field](#)

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, QuickSight 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 QuickSight, 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

1. Select the field menu for the field that you want to customize (the three dots) and choose **Edit**.
2. In the **Edit tooltip field** page that opens, for **Label**, enter the label that you want to appear in the tooltip.
3. 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

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 **Format visual** pane that opens at left, choose **Tooltip**.
4. 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 Amazon QuickSight 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 QuickSight](#)
- [Sorting visual data in Amazon QuickSight](#)

Changing fields used by a visual in Amazon QuickSight

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 QuickSight](#) 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 QuickSight](#).

Use the following topics to learn more about adding, removing, and modifying fields on a visual.

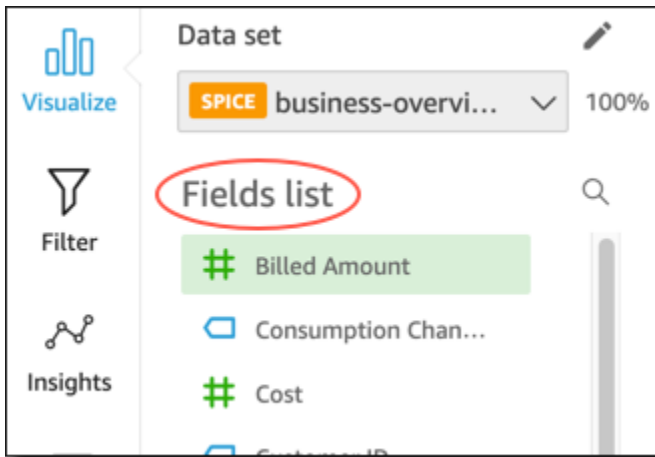
Topics

- [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](#)

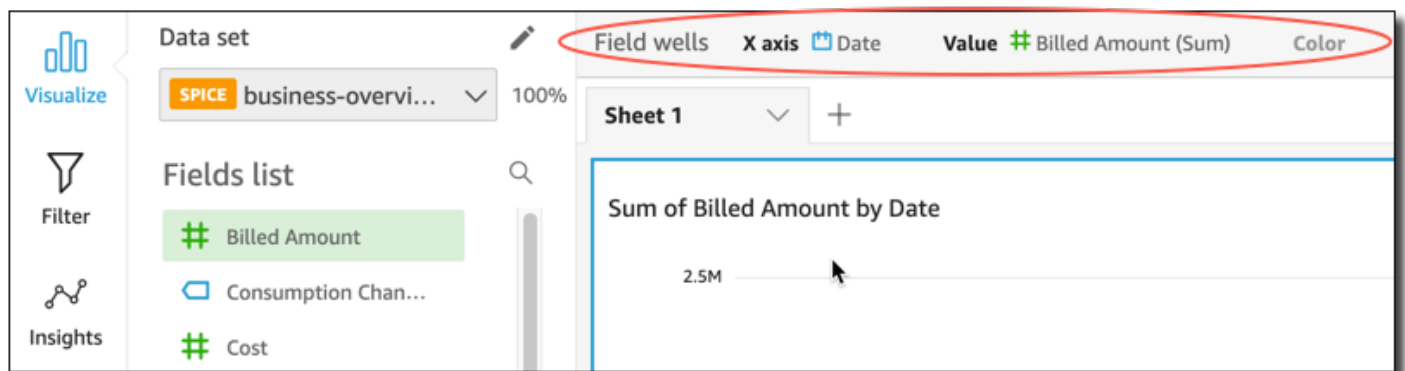
Using visual field controls

You can edit the fields used by a visual by using the following user interface (UI) controls:


- The **Fields list** pane.



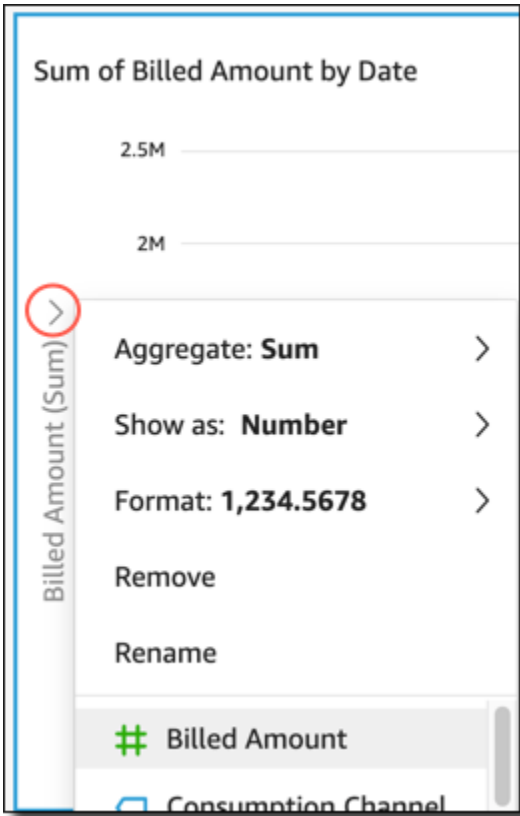
- The field wells. The following screenshot shows the field wells in the default closed state.



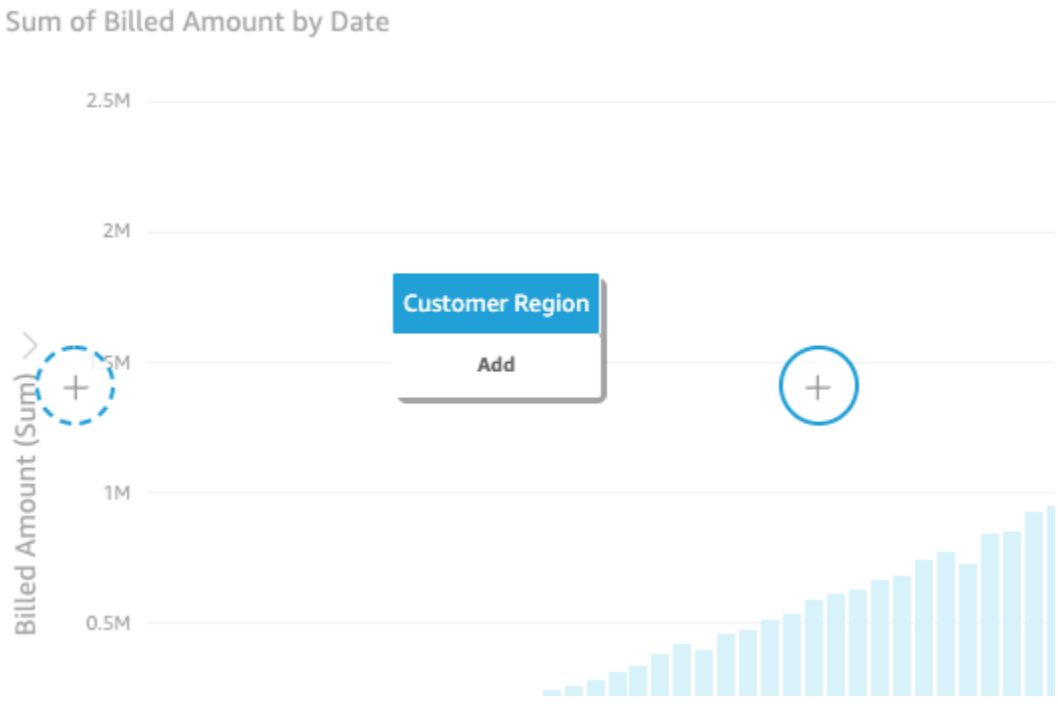
Click anywhere on the **Field wells** to open the field wells. The following screenshot shows the field wells in the open state.



- The on-visual editors.



- The drop targets on the visual.



You can use these controls as follows:


- You can 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.
- You can 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.
- You can 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.

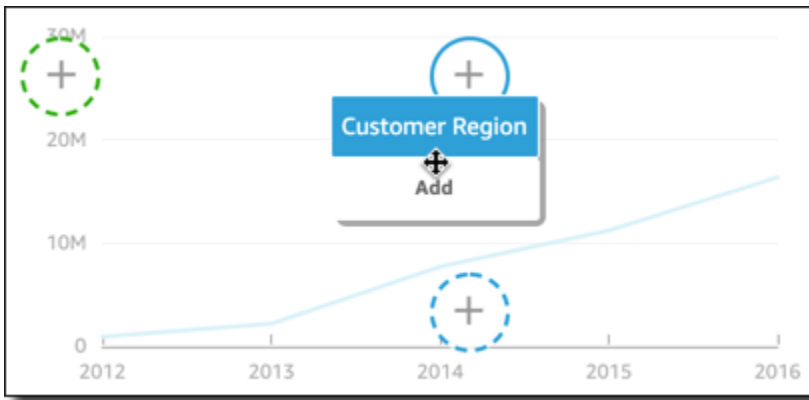
Dragging fields to drop targets or field wells

When you drag a field to either a drop target or field well, Amazon QuickSight provides you with information about whether the target element expects a measure or a dimension. Amazon QuickSight 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 **Axistitle** 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.

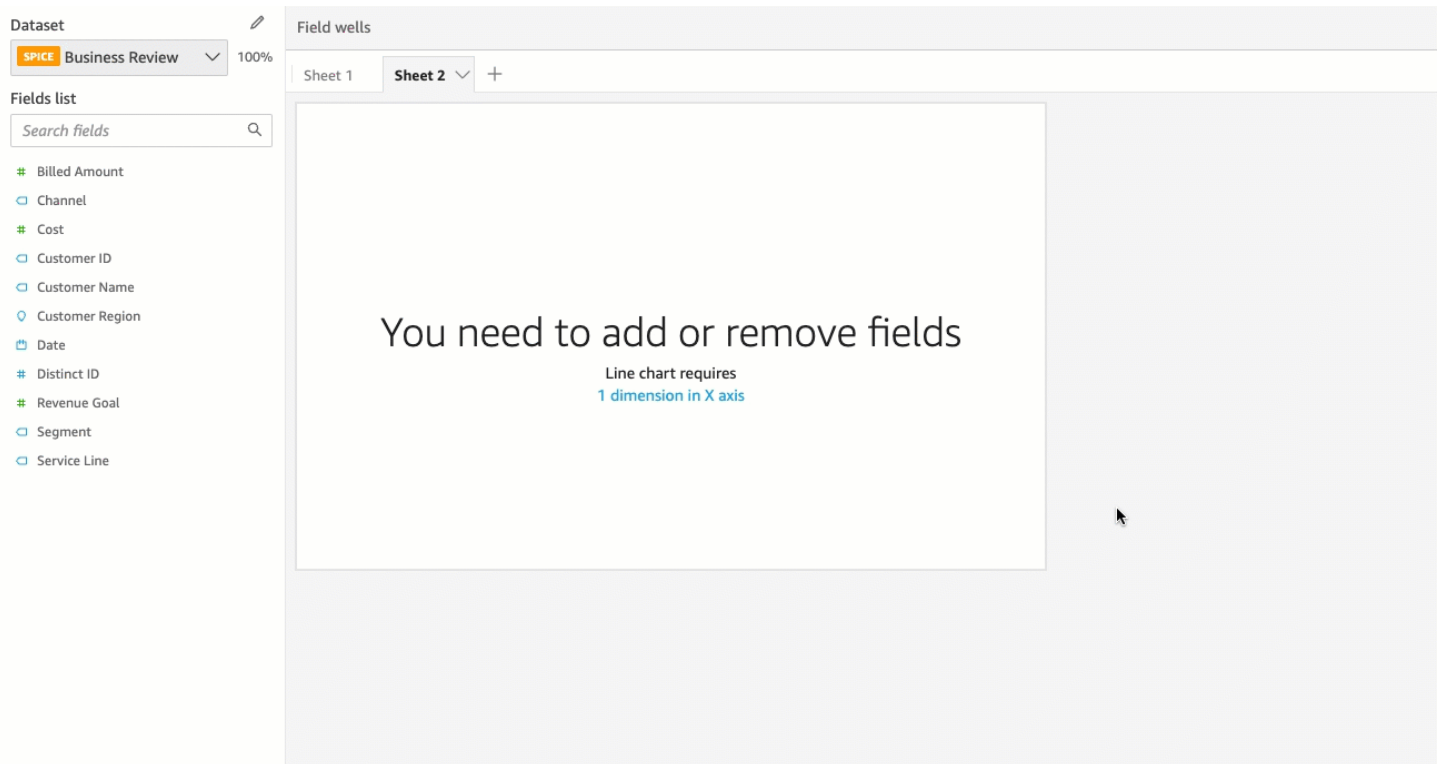
Adding a field by selecting it in the fields list pane

You can also let Amazon QuickSight map the field to the most appropriate visual element. To do so, choose the field in the **Fields list** pane. Amazon QuickSight 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 QuickSight 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

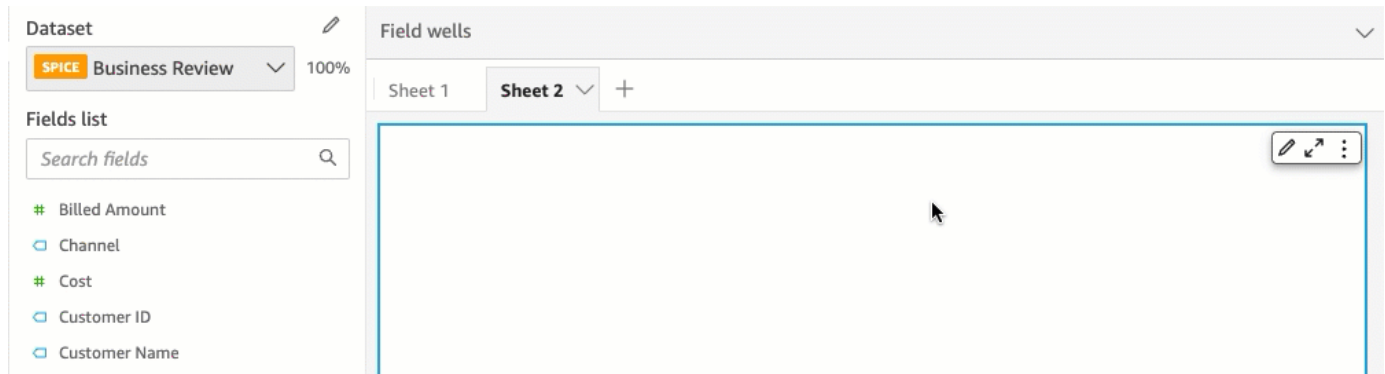
To add a field to a visual by using a drop target, first choose a field in the **Fields list** pane. Then drag the field to your chosen drop target on the visual, making sure the drop indicator shows that the field is being added.



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. Click anywhere on the **Field wells** to expand them.



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 **Format Visual** 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

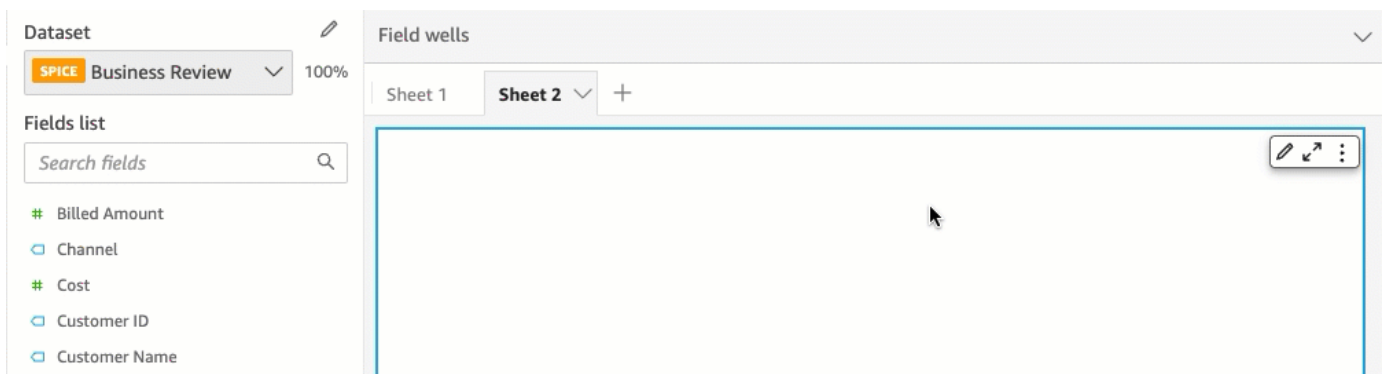
To modify the mapping of a field to a visual element by using a drop target, choose a field in the **Fields list** pane. Then drag the field to a drop target on the visual, making sure that the drop indicator shows that the field is being replaced.

Changing a field mapping by using a field well

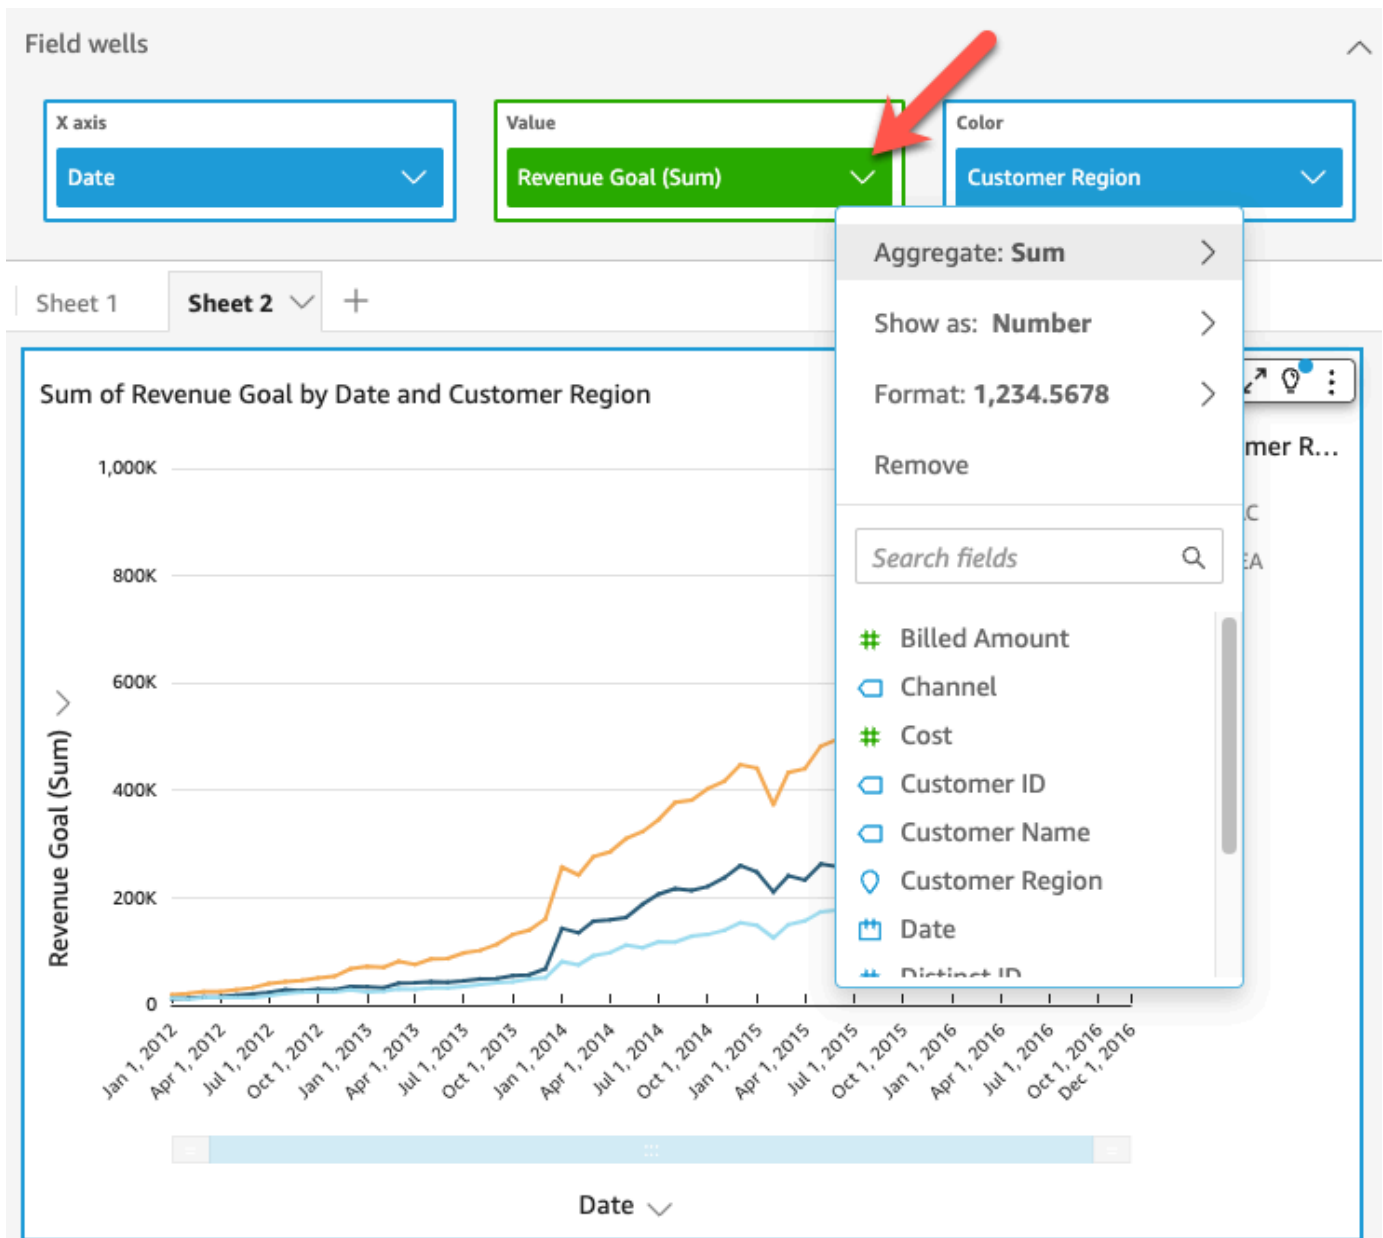
Use the following procedure to modify the mapping of a field to a visual element.

To modify the mapping of a field by using a field well

1. Click anywhere on the **Field wells** to expand them.



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 QuickSight:

- 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.

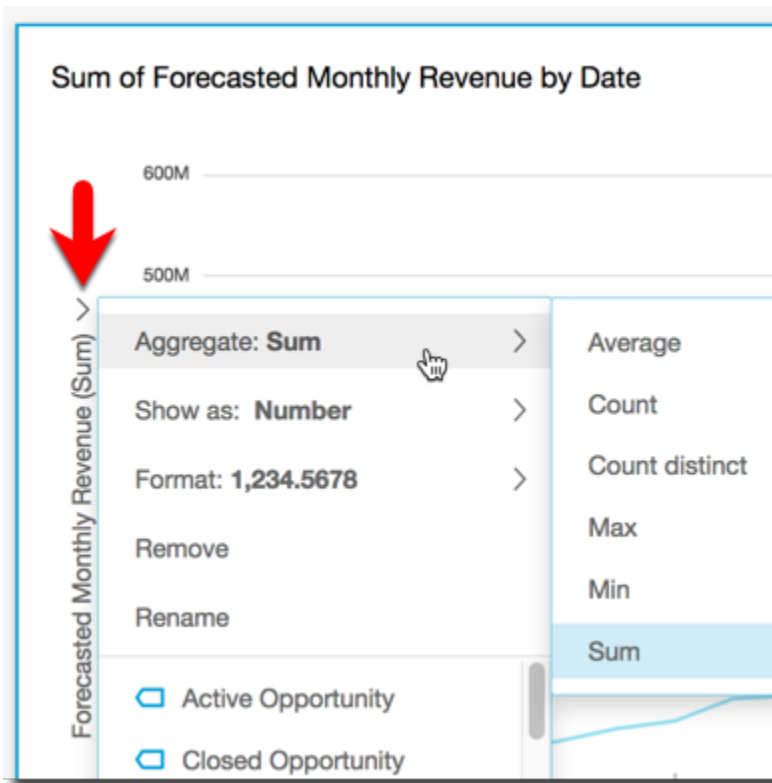
The visual elements that support aggregated fields varies by visual type.

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.

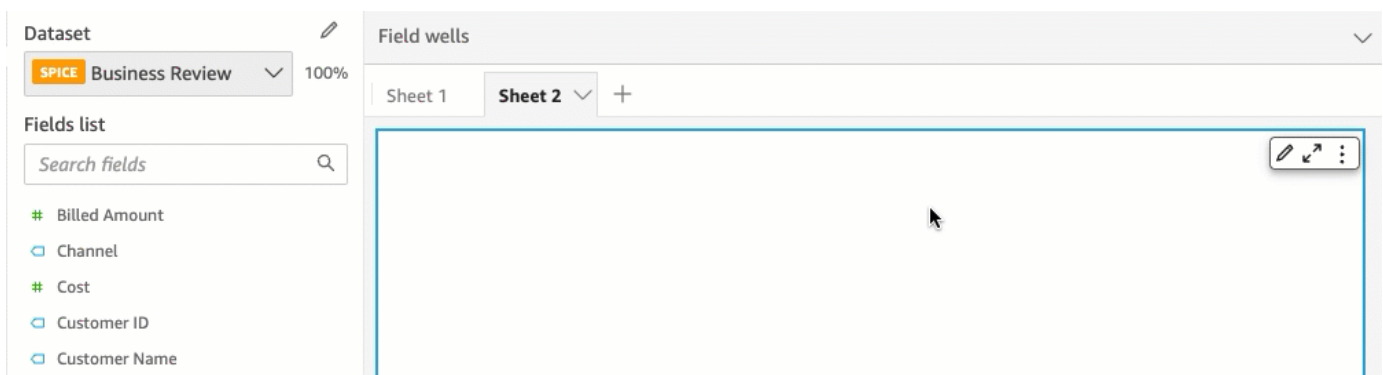


Changing or adding aggregation to a field by using a field well

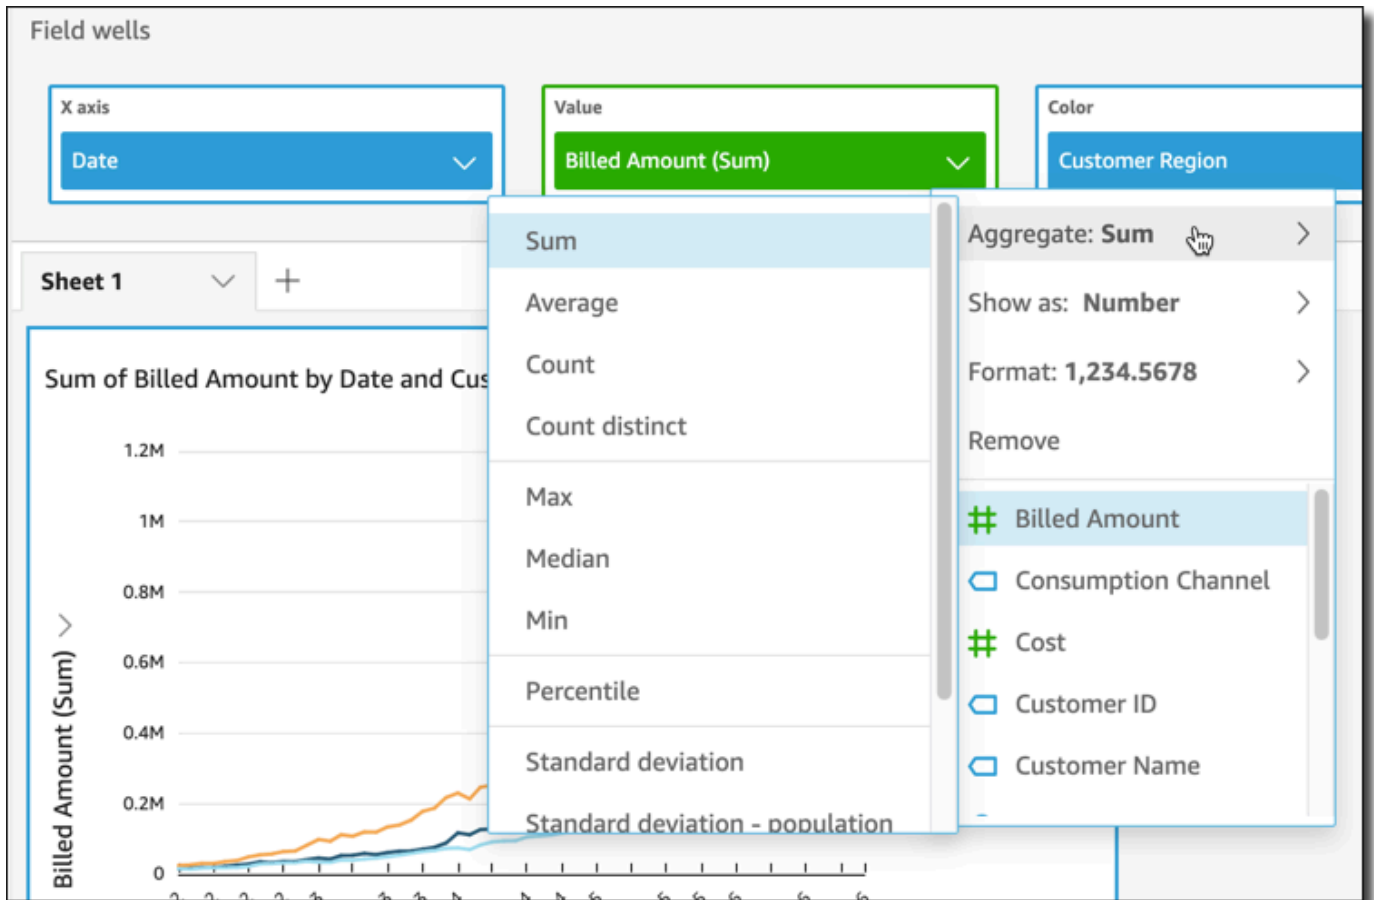
Use the following procedure to add aggregation to a field for a pivot table visual.

To add aggregation to a field for a pivot table visual

1. Click anywhere on the **Field wells** to expand them.



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. Click anywhere on the **Field wells** to expand them.



2. Choose the field well containing the date field, and then choose **Aggregate**. Choose the date granularity that you want to use.

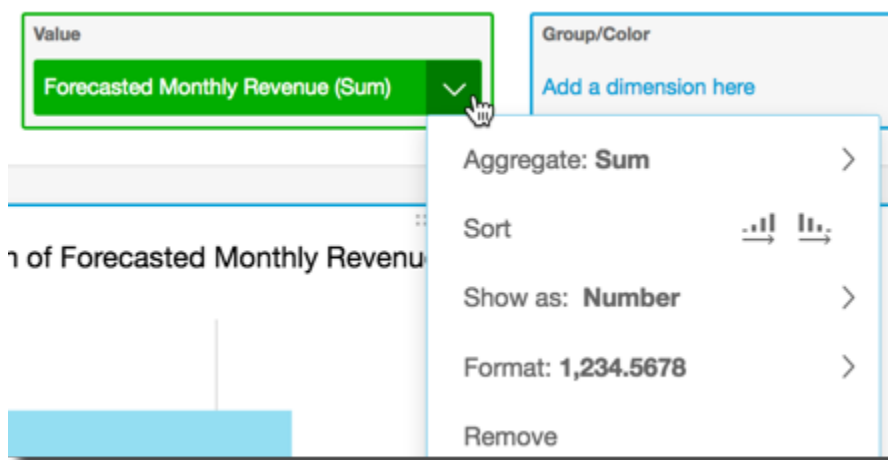


Customizing a field format

Use the following procedure to customize the appearance of fields in an analysis.

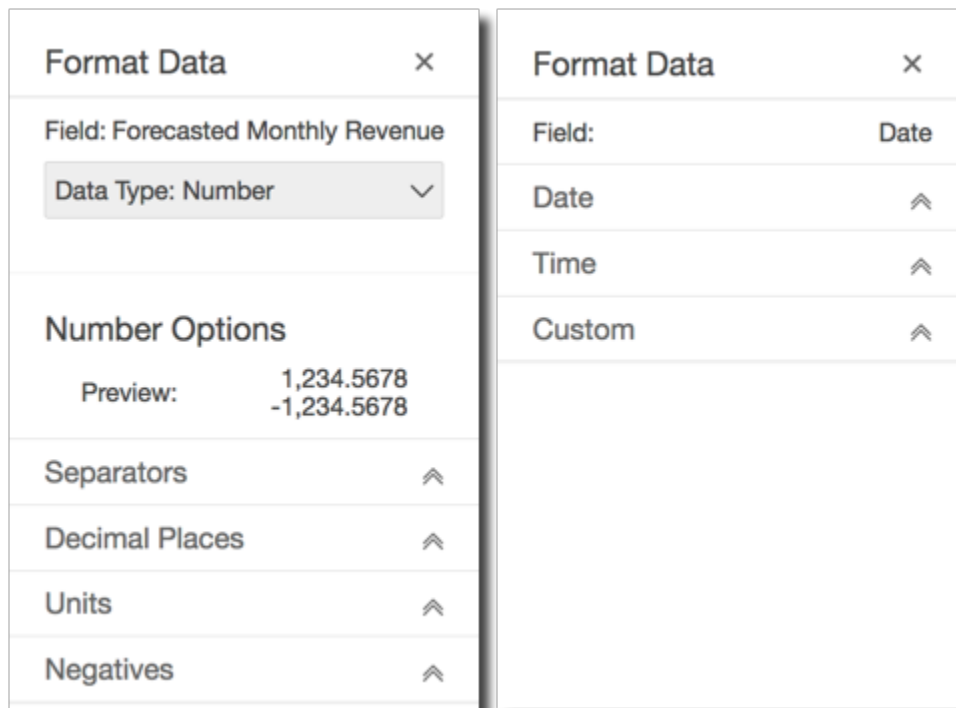
To customize the appearance of fields in an analysis

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 following screenshots show the **Format Data** pane.



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 QuickSight. 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:mm, 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.

- 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

Time zone related display tokens (Z and z) from the Moment.js library are not supported in QuickSight.

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.

- (Optional) Expand the **Custom** section and use **Preview** to verify your specified format.
- 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 Amazon QuickSight

In Amazon QuickSight, 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 QuickSight.

Example	Description	Token
0-6		d

Example	Description	Token
	Numeric representation of a particular day of the week. 0 is Sunday and 6 is Saturday.	
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
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

Example	Description	Token
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
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

Example	Description	Token
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
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

Example	Description	Token
AM or PM	AM/PM	A
1632184215	Unix timestamp.	X
1632184215000	Millisecond Unix timestamp.	x

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 QuickSight preset tokens.

Example	Token
8:30 PM	LT
8:30:25 PM	LTS
August 2 1985	LL
Aug 2 1985	ll
August 2 1985 08:30 PM	LLL
Aug 2 1985 08:30 PM	lll

Example	Token
Thursday, August 2 1985 08:30 PM	LLLL
Thu, Aug 2 1985 08:30 PM	1111

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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight

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 QuickSight](#).

Pivot tables behave differently than tables when sorting values. For more information about sorting pivot tables, see [Sorting pivot tables in Amazon QuickSight](#).

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.

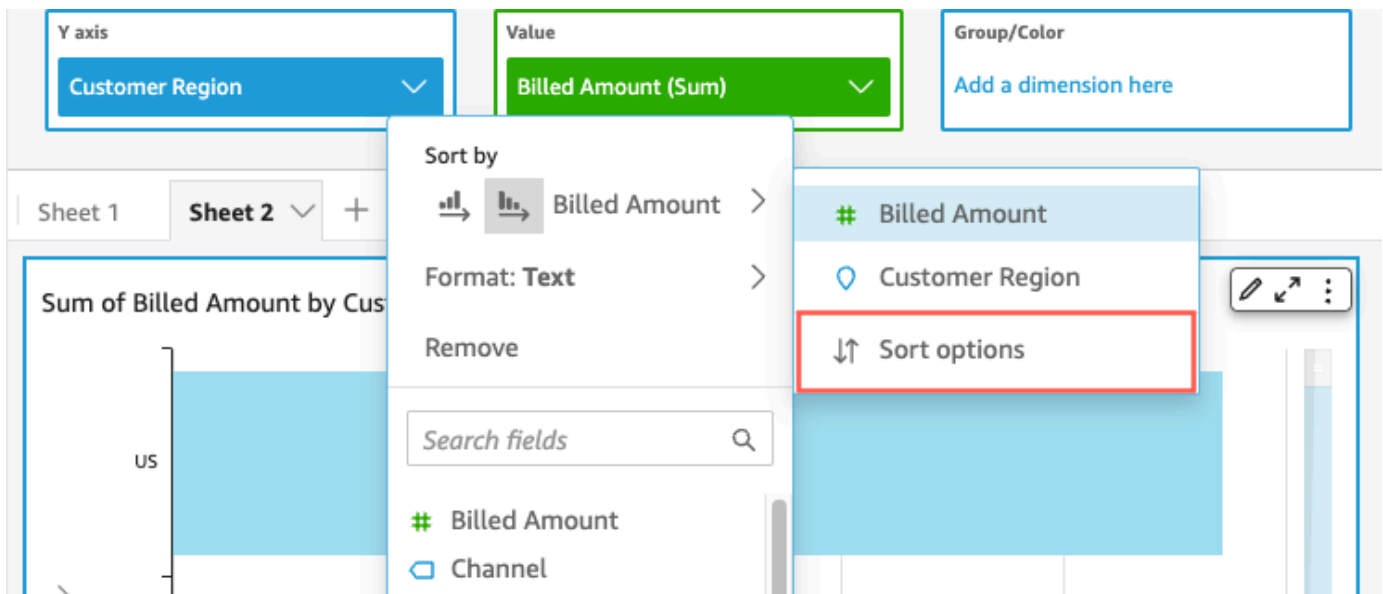


To sort by using an off-visual metric

- Open the analysis with the visual that you want to sort and click anywhere on **Field wells** to expand the field wells.




2. Choose a field well that supports sorting, then choose **Sort by**, **Sort options**.



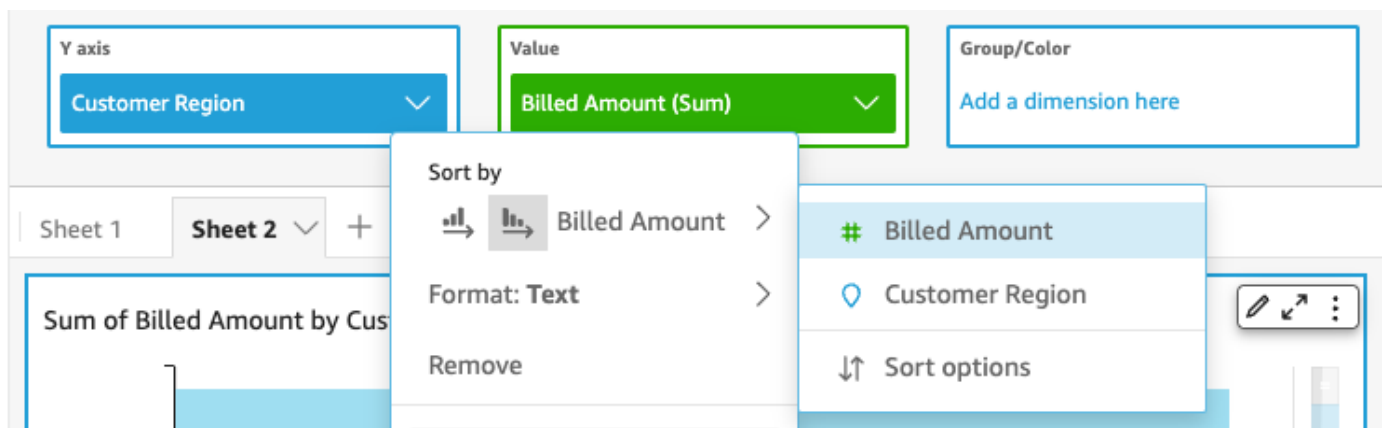
3. On the **Sort options** pane, sort by specific fields, choose an aggregation, or sort ascending or descending, or do a combination of these.
Amazon QuickSight Sort options pane to sort by off-visual fields.
4. Choose **Apply** to save your changes. Or choose **Clear** to start over or **Cancel** to go back.

To sort by using a field well

1. Open the analysis with the visual you want to sort and click anywhere on **Field wells** to expand the field wells.



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 QuickSight

In Amazon QuickSight, a *theme* is a collection of settings that you can apply to multiple analyses and dashboards. Amazon QuickSight 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 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	<p>This setting is used as an interactive hint for the following:</p> <ul style="list-style-type: none"> • Buttons • Borders around the selected visual • Loading indicators • Narration customizations

Group	Setting	What the setting changes
		<ul style="list-style-type: none">• 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.
Main	Font	The font to use for all of the text. You can choose from a variety of fonts supported by Amazon QuickSight.
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.

Group	Setting	What the setting changes
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.
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.

Group	Setting	What the setting changes
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 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.

2. Choose **Themes** on the left. The theme panel opens.
3. 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 QuickSight.
4. 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 QuickSight 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.
5. 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 an analysis, or create a new one. Choose **Themes** at left.

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.

2. Choose **Main**. The color picker used in each of these settings is the standard one used throughout Amazon QuickSight.

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.

3. 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.

4. 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.

5. 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.

6. 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.

7. 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 (...).
8. 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 QuickSight using keyboard shortcuts

You can use the following keyboard shortcuts to navigate a Amazon QuickSight 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.

Visualize Dataset: **SPICE b2b-sales.csv** 100%

Fields list

Search fields

- buyer_name
- country
- customer_id
- customer_name
- Date
- discount (SUM)
- industry
- order_id
- product_id
- product_name
- profit (SUM)
- quantity (SUM)
- sales (SUM)
- segment
- ship_city
- ship_mode
- ship_region
- ship state

Field wells

Overview Metrics +

Controls industry equals All ship_city equals All

Business Summary Dashboard

Data as of Dec 2019. Click on a city to drill to Details sheet.

Top Customers: Top 3 customer_name for total sales (SUM) are: <ul style="list-style-type: none">Valero Energy with \$533,741.76Lowes with \$505,540.80Bank of America Corp. with \$488,228.80	Total Sales: <h2>\$21,500,838.40</h2>
---	---

Sales in All Industries

Industry	Sales
Finance	~300K
Retail	~250K
Tech	~200K
Energy	~150K
Misc	~100K

Sales Trend

Gaining insights with machine learning (ML) in Amazon QuickSight

Amazon QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight provides three major features:

- **ML-powered anomaly detection** – Amazon QuickSight 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 QuickSight 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 QuickSight 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 QuickSight, 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 QuickSight is like providing a personal data analyst to all of your users.

Topics

- [Understanding the ML algorithm used by Amazon QuickSight](#)
- [Dataset requirements for using ML insights with Amazon QuickSight](#)
- [Working with insights in Amazon QuickSight](#)
- [Creating autonarratives with Amazon QuickSight](#)
- [Detecting outliers with ML-powered anomaly detection](#)
- [Forecasting and creating what-if scenarios with Amazon QuickSight](#)

Understanding the ML algorithm used by Amazon QuickSight

You don't need any technical experience in machine learning to use the ML-powered features in Amazon QuickSight. 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 QuickSight 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 QuickSight.

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 RCF is and what it does](#)
- [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 QuickSight 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 QuickSight 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 RCF is and what it does

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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight

To begin using the machine learning capabilities of Amazon QuickSight, you need to connect to or import your data. You can use an existing Amazon QuickSight 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 QuickSight

In Amazon QuickSight, you can add ready-to-use analytical computations to your analysis as widgets. You can work with insights in two ways:

- Suggested insights

Amazon QuickSight 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 QuickSight](#).

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.

Visualize

Filter

Insights

Parameters

Actions

Themes

Settings

Suggested insights

SALES IN <<\$3BA9A2ED-17EF-4A91-A1B4-180E1BAA35B3>> INDUSTRIES

TOP 3 INDUSTRIES

Top 3 industrys for total sales (SUM) are:

- Finance** with 4,193,206.24
- Retail** with 3,089,122.08
- Tech** with 2,781,517.28

BOTTOM 3 INDUSTRIES

Bottom 3 industrys for total sales (SUM) are:

- Other** with 115,406.88
- Agriculture** with 166,710.24
- Industrial** with 253,124.96

TOTAL AGGREGATION

Total sales (sum) is 21,500,838.4.

COUNT OF RECORDS BY INDUSTRY

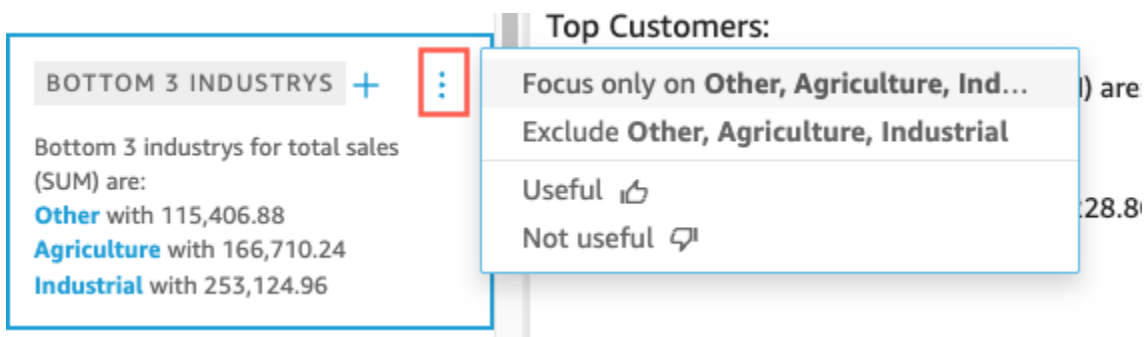
TOP 3 INDUSTRIES

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.

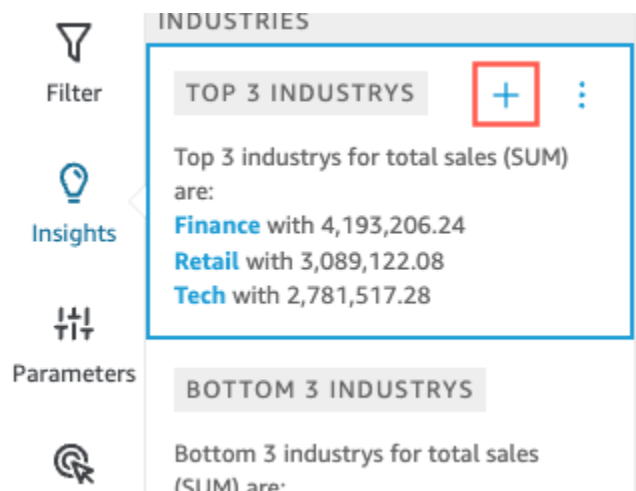
The insights that appear are controlled by the data type of the fields you choose to include in your visual. This list is generated each time you change your visual. If you make changes, check **Insights** to see what is new. To get a specific insight, see [Adding custom insights to your analysis](#).

- (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.
- Add a suggested insight to your analysis by choosing the plus sign (+) near the insight title.



- (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 QuickSight](#).

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](#).

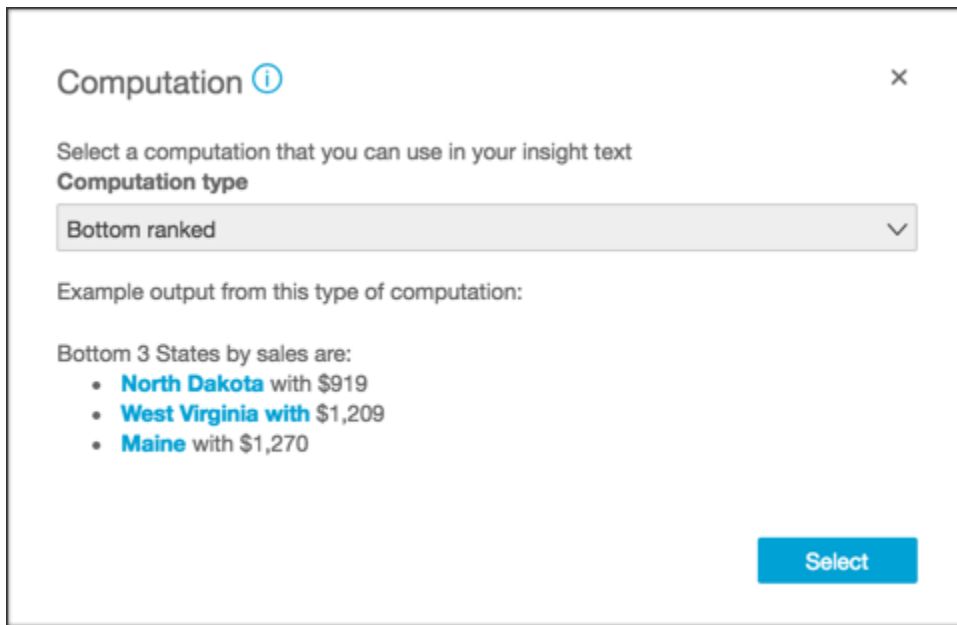
6. (Optional) To remove the insight from your analysis, choose the v-shaped on-visual menu at the top right of the visual. Then choose **Delete**.

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. The following screen appears.



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 QuickSight](#).

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 QuickSight

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 QuickSight 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](#)
- [Working with the expression editor screen and menus](#)
- [Adding URLs](#)
- [Walkthrough: Use the narrative expression editor](#)
- [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](#).

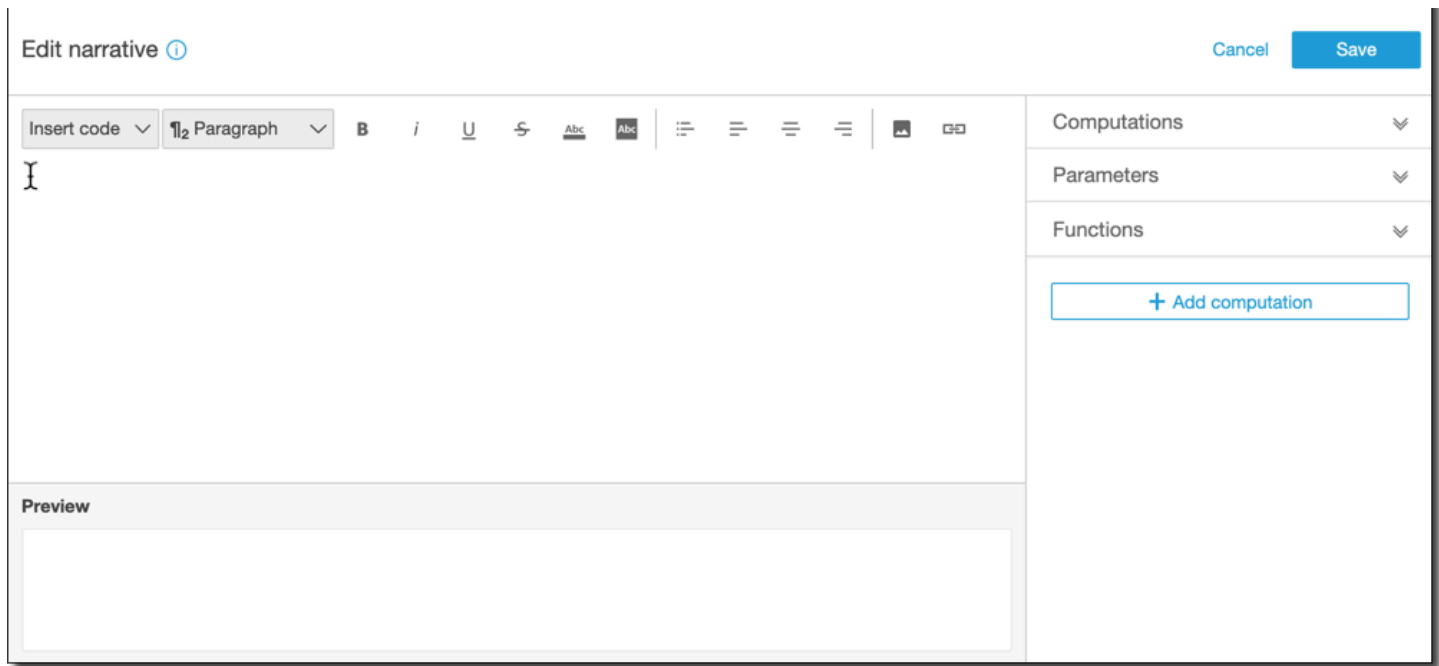
- **Bottom ranked** – For example, the bottom three states by sales revenue. Requires that you have at least one dimension in the **Categories** field well.
- **Bottom movers** – For example, the bottom three products sold, by sales revenue. Requires that you have at least one dimension in the **Time** 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.

Working with the expression editor screen and menus

The following screenshot shows a new blank narrative. In this image, the browser window is smaller than usual, so you can see the icons on the menu bar. You can maximize the browser to make the editor as large as your screen.



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 [Walkthrough: 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 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 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](#).

Walkthrough: 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.

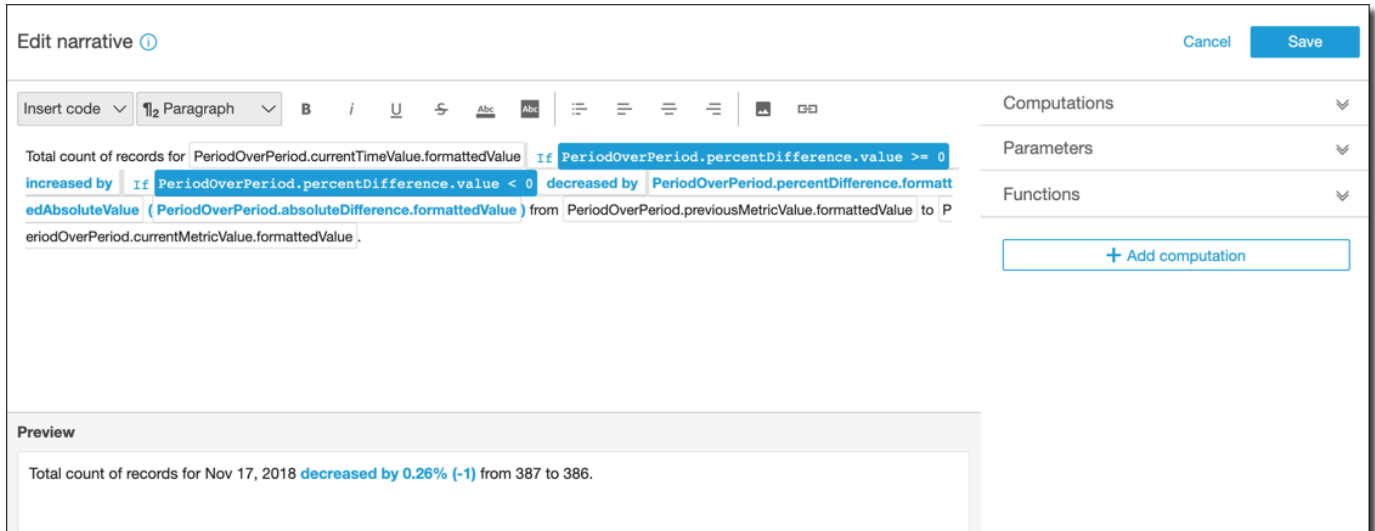
1. Begin with an existing analysis. Add a **period over period** insight to it. The easiest way to do this is to choose the + icon, then **Add insight**, then choose a type of insight from the list. To learn what type of computational insights you can add as autonarratives, see [Insights that include 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**. The following screen appears, filling the entire browser window except for the Amazon QuickSight menu.



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 QuickSight 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**.

- 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**.
- In the workspace, add the following text after the final period: **Compounded growth rate for the last**, then add a space.
- 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**.

- 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.

Edit narrative ⓘ Cancel Save

Insert code ▾ | Paragraph ▾ | B | i | U | | | | | |

Total count of records for `PeriodOverPeriod.currentTimeValue.formattedValue` **increased by** `If PeriodOverPeriod.percentDifference.value >= 0` **decreased by** `PeriodOverPeriod.percentDifference.formattedAbsoluteValue (PeriodOverPeriod.absoluteDifference.formattedValue)` from `PeriodOverPeriod.previousMetricValue.formattedValue` to `PeriodOverPeriod.currentMetricValue.formattedValue`. Compounded growth rate for the last `GrowthRate.timePeriods` days is `GrowthRate.compoundedGrowthRate.formattedValue`.

Preview
 Total count of records for Nov 17, 2018 **decreased by 0.26% (-1)** from 387 to 386. Compounded growth rate for the last 4 days is -0.06%.

GrowthRate
 > timeField
 > metricField
 > previousMetricValue
 > previousTimeValue
 > currentMetricValue
 > currentTimeValue
 timePeriods
 compoundedGrowthRate
 value
 formattedValue
 formattedAbsoluteValue
 > absoluteDifference

Parameters ▾
Functions ▾


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.

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

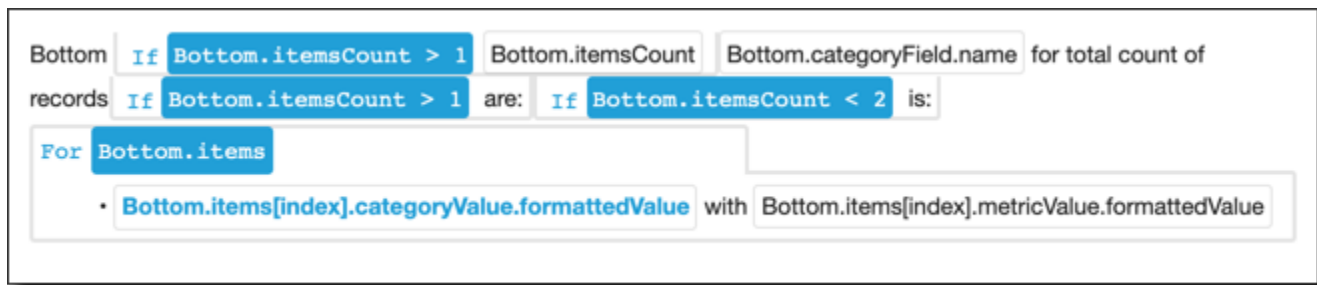
These are the same output parameters as the ones that are returned by the top ranked computation.

- **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 QuickSight](#).

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, QuickSight shows the year to date sales of the latest month and the percentage difference compared with the previous month.



- (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, QuickSight shows the year to date sales of the latest month and the percentage difference compared with the previous month.



- (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

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 QuickSight 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 QuickSight notifies you in your visuals if it detects that you can analyze an anomaly or do some forecasting on your data.

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. To learn more about usage-based pricing, see [Amazon QuickSight Pricing](#).

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 QuickSight 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.

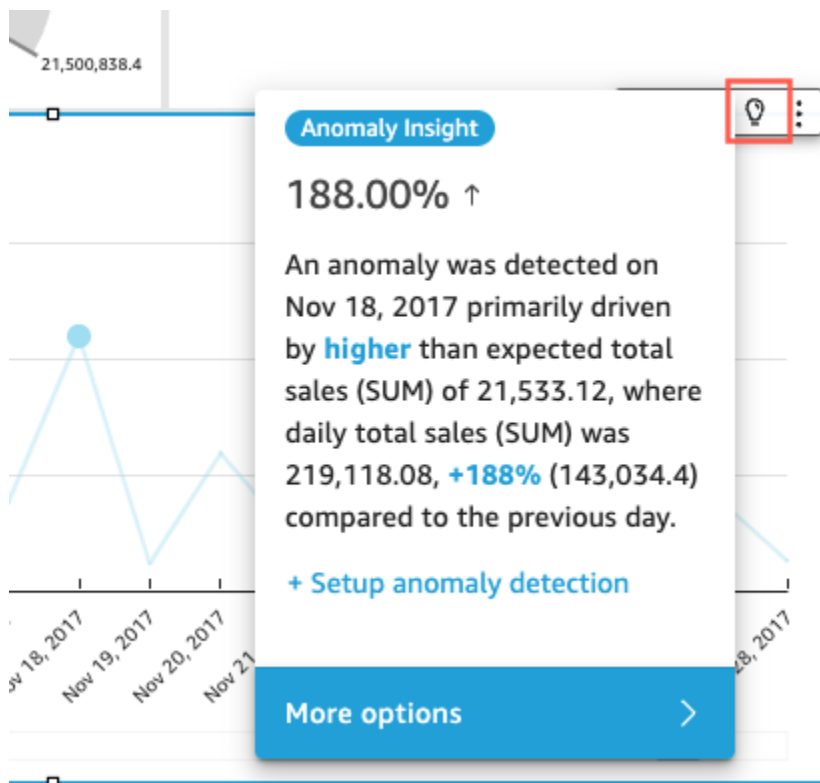
Topics

- [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 QuickSight 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, as shown in the following screenshot.



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 QuickSight 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 QuickSight 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, on the top menu, choose **Add**, then **Add insight**. 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), QuickSight 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), QuickSight 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), QuickSight 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, QuickSight 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 QuickSight 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 QuickSight 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 QuickSight generates. Amazon QuickSight 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 QuickSight to analyze the key drivers when an outlier (anomaly) is detected.

For example, Amazon QuickSight 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.

Appears on the Page	Status
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.
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 QuickSight 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 QuickSight to analyze the key drivers when an outlier (anomaly) is detected. For example, Amazon QuickSight 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**.

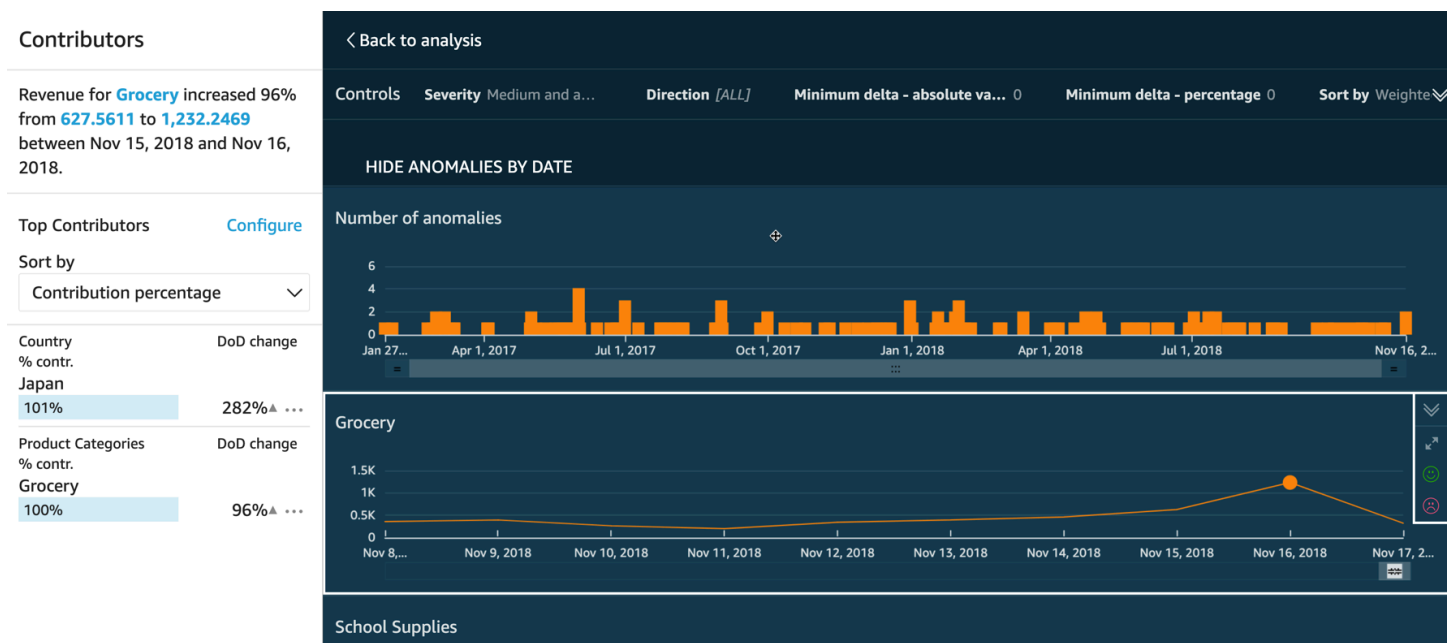
If you want to change the fields you're using as key drivers, change the fields that are enabled in this list. If you disable all of them, QuickSight won't perform any contribution analysis in this insight.

- 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.

Exploring contributors (key drivers)

If your anomaly insight is set up to detect key drivers, QuickSight runs the contribution analysis to determine which categories (dimensions) are influencing the outliers. The **Contributors** section appears on the left.

Contributors

Revenue for **Grocery** increased 96% from **627.5611** to **1,232.2469** between Nov 15, 2018 and Nov 16, 2018.

Top Contributors [Configure](#)

Sort by

Contribution percentage ▼

Country % contr.	DoD change
Japan 101%	282%▲ ...

Product Categories % contr.	DoD change
Grocery 100%	96%▲ ...

[← Back to analysis](#)

Controls Severity Medium

Anomalies

HIDE ANOMALIES BY D

Number of anomalies

Grocery

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 QuickSight 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 screenshot shows the Amazon QuickSight interface. At the top, there is a dark blue navigation bar with a '< Back to analysis' link on the left and a 'Controls' button circled in red. To the right of the 'Controls' button are several filters: 'Severity Medium and a...', 'Direction [ALL]', 'Minimum delta - absolute va... 0', 'Minimum delta - percentage 0', and 'Sort by Weighte...'. Below this bar, a red arrow points down to an expanded 'Controls' panel. This panel has a title 'Controls' and a double-up arrow icon on the right. It contains several settings: 'Severity' set to 'Medium and above', 'Direction' set to '[ALL]', 'Minimum delta - absolute value' set to '0', 'Minimum delta - percentage' set to '0', 'Sort by' set to 'Weighted difference from ...', and 'Product Categories' set to '[ALL]'.


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 QuickSight

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 QuickSight is designed to handle complex real-world scenarios. Amazon QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight](#).

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**.



QuickSight 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.



- 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 QuickSight 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.

Sharing and subscribing to data in Amazon QuickSight

A *dashboard* is a read-only snapshot of an analysis that you can share with other Amazon QuickSight 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 [Working with embedded analytics](#).

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

- [Sharing Amazon QuickSight analyses](#)
- [Publishing dashboards](#)
- [Sharing Amazon QuickSight dashboards](#)
- [Sharing your view of a Amazon QuickSight dashboard](#)
- [Scheduling and sending reports by email](#)
- [Subscribing to email reports in Amazon QuickSight](#)
- [Working with threshold alerts in Amazon QuickSight](#)
- [Printing a dashboard or analysis](#)
- [Exporting Amazon QuickSight analyses or dashboards as PDFs](#)
- [Error codes for failed PDF export jobs](#)
- [Organizing assets into folders for Amazon QuickSight](#)

Sharing Amazon QuickSight 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 Amazon QuickSight 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. On the analysis page, choose **Share** on the application bar, and then choose **Share analysis**.

You can only share analyses with users or groups who are in your Amazon QuickSight account.

2. 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.

Edit sharing for this analysis by choosing **Manage analysis access**.

The **Share with users and groups in your account** screen appears, as shown following. On this screen, you can edit permissions and add more users or groups.



Share with users and groups in your account ✕

Co-owners can access the analysis and the data set it contains, including retrieving new data.
Co-owners can't access connection information, like credentials.

3. 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.

Share with users and groups in your account ✕

Co-owners can access the analysis and the data set it contains, including retrieving new data.
Co-owners can't access connection information, like credentials.

Name	Email	Permission
 <i>[Redacted]</i>	<i>[Redacted]</i>	COOWNER ⌵ 

4. 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. On the analysis page, choose **Share** on the application bar, and then choose **Share analysis**.
2. Choose **Manage analysis access**.
3. 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. On the analysis page, choose **Share** on the application bar, and then choose **Share analysis**.
2. Choose **Manage analysis access**.
3. Locate the user or group whose access you want to revoke, and then choose the trash-can icon next to the user or group.
4. Choose **Confirm**.

Publishing dashboards

When you publish an analysis, that analysis becomes a dashboard that can be shared and interacted with by users of your Amazon QuickSight 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 QuickSight. 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 reports by email](#) .

For more information about viewing a report's snapshots, see [Consuming paginated reports in Amazon QuickSight](#).

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 **Share** on the application bar, and then choose **Publish dashboard**.
2. Do one of the following:
 - To create a new dashboard, choose **Publish new dashboard as**, 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 choose **Rename**. Enter a new name to rename the existing dashboard. When you rename a dashboard, it also saves any changes you made to the analysis.
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. You can make changes to this by selecting or de-selecting sheets from the dropdown list.



- (Optional) To allow dashboard readers to share data stories, choose **Allow sharing data stories**. For more information about data stories, see [???](#).
- (Optional) Open the **Advanced publish options**. 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.

6. Choose **Publish dashboard**.

If you renamed the existing dashboard, the top of the screen refreshes to show the new name.

Publish a dashboard ×

Publish new dashboard as

New dashboard 1

Replace an existing dashboard

Test Dashboard for CSV export ▼

ALL SHEETS SELECTED ▼

Enable topic for analysis ⓘ

[Learn more about Q topic](#)

Advanced publish options ^

Dashboard options

- Expand on-sheet controls by default
- Enable ad hoc filtering ⓘ
- Enable on-hover tooltip
- Enable auto-refresh for visualizations ⓘ
- Enable dashboard access request ⓘ

Visual options

- Enable maximize visual option
- Enable visual menu
 - Enable download options
 - Enable export of hidden fields on supported visuals

Data point options

- Enable drill up/down

Publish dashboard

7. (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 **Share** from the application bar.
- To share the dashboard, follow the procedure in [Sharing Amazon QuickSight 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 QuickSight 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 Amazon QuickSight at <https://quicksight.aws.amazon.com/> and choose **Dashboards** from the start page.
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 QuickSight dashboards](#)
- [Sharing Amazon QuickSight analyses](#)

Deleting an Amazon QuickSight dashboard

When you delete a Amazon QuickSight 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 QuickSight start page, 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.



Are you sure you want to delete this dashboard? ×

The item will be permanently deleted from the account and will disappear from all folders.

Cancel

Delete

Sharing Amazon QuickSight dashboards

By default, dashboards in Amazon QuickSight 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 QuickSight account. You can also choose to share the dashboard with everyone in your QuickSight account and make the dashboard visible on the QuickSight 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 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 QuickSight account. Or you share them with anyone on the internet. You can share dashboards and visuals by using the QuickSight console or the QuickSight 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:

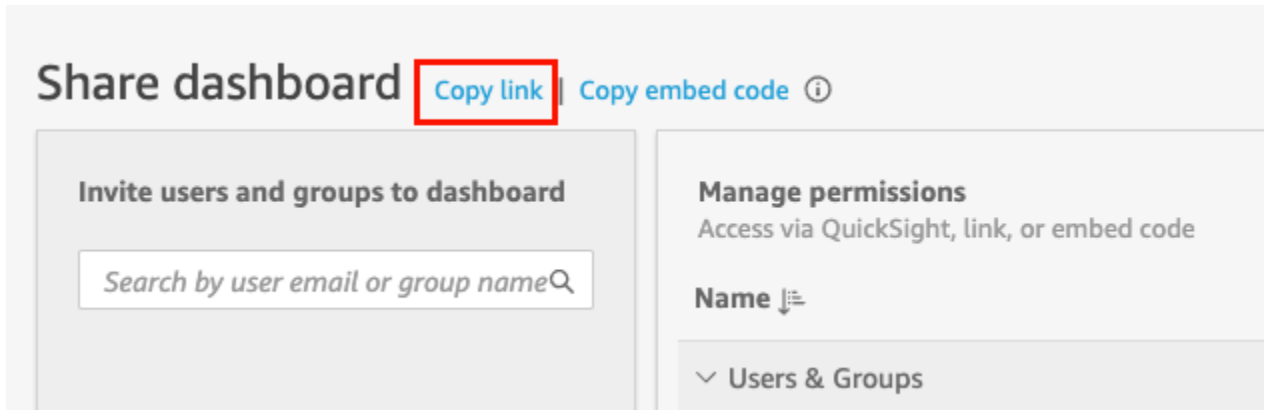
- [Granting individual Amazon QuickSight users and groups access to a dashboard in Amazon QuickSight](#)
- [Granting everyone in your Amazon QuickSight account access to a dashboard](#)
- [Granting anyone on the internet access to an Amazon QuickSight dashboard](#)
- [Granting everyone in your Amazon QuickSight account access to a dashboard with the QuickSight API](#)
- [Granting anyone on the internet access to an Amazon QuickSight dashboard using the QuickSight API](#)

Sharing a link to the 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 QuickSight account) who have access to this dashboard can access it by using the link. If they are accessing QuickSight for the first time, they will be asked to sign in with their email address or QuickSight user name and password for the account. After they sign in, they will have access to the dashboard.

Viewing who has access to a 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.

Revoking access to a 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.

Granting individual Amazon QuickSight users and groups access to a dashboard in Amazon QuickSight

Use the following procedure to grant access to a dashboard.

To grant users or groups access to a dashboard

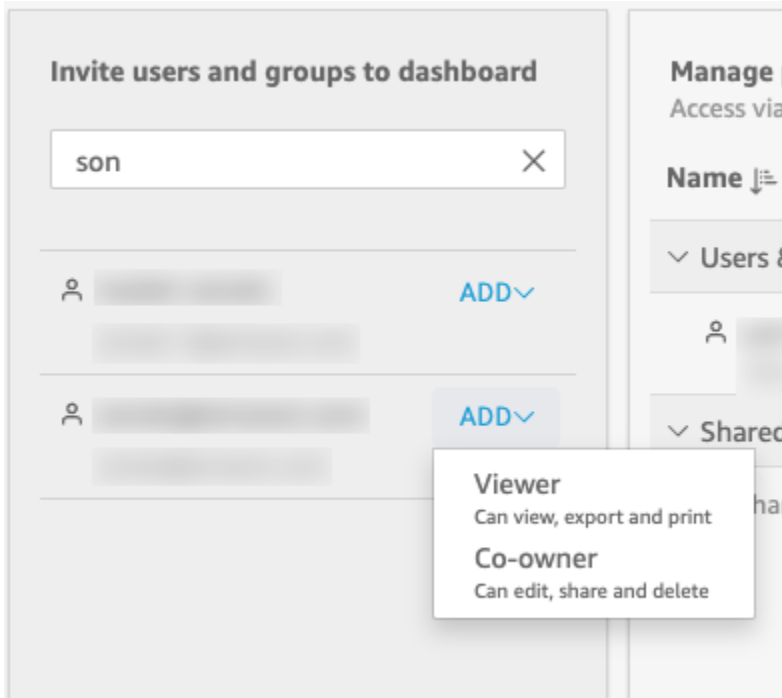
1. Open the published dashboard and choose **Share** at upper right. Then choose **Share dashboard**.



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 QuickSight role. The available permissions for each role are as follows:


- **Readers** – QuickSight 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** – QuickSight 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** – QuickSight 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.



To change the permission level for a user, choose the permission level menu in the **Permissions** column and select a permission.



- 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 QuickSight account access to a dashboard

Alternatively, you can share your Amazon QuickSight 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.

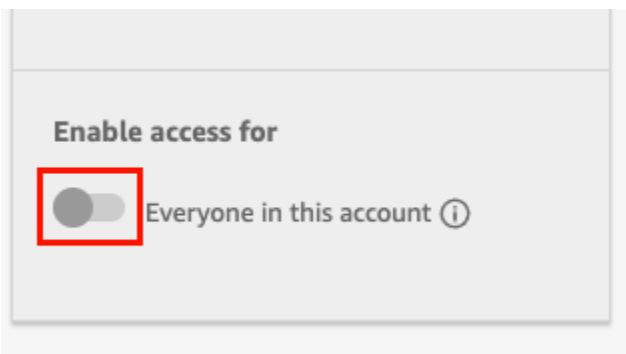
Sharing the dashboard with everyone in your account doesn't affect email reports. For example, suppose that you choose to share the dashboard with everyone in your account. Suppose also that you choose **Send email report to all users with access to dashboard** when setting up an email report for the 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**.

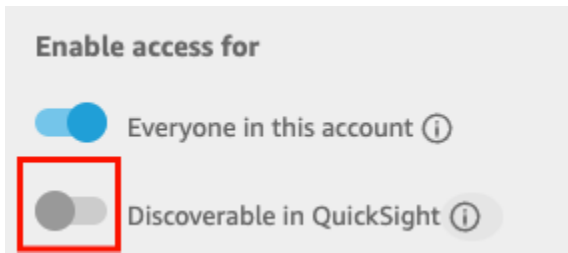


2. In the **Share dashboard** page that opens, for **Enable access for** at bottom left, toggle on **Everyone in this account**. Accounts that sign in with an Active Directory can't access the **Everyone in this account** switch. Accounts that use Active Directory can enable this setting with an `UpdateDashboardPermissions` API call. For more information on `UpdateDashboardPermissions`, see [UpdateDashboardPermissions](#) in the *Amazon QuickSight API Reference*.



3. (Optional) Toggle on **Discoverable in QuickSight**.

When you share a dashboard with everyone in the account, owners can also choose to make the dashboard discoverable in QuickSight. 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 QuickSight dashboard

Applies to: Enterprise Edition

You can also share your Amazon QuickSight dashboard with anyone on the internet from the **Share** menu in the QuickSight 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 QuickSight account, when you share the dashboard link or embed the dashboard.

Use the following procedure to grant anyone on the internet access to dashboard when you share it.

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. For more information about session capacity pricing, see <https://www.amazonaws.cn/quicksight/pricing/>.
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.

```
{
  "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`.

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

For more information on using IAM with QuickSight, see [Using Amazon QuickSight 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 QuickSight account.
 1. From the Amazon QuickSight start page, choose your user icon at the upper right of your browser window, and then choose **Manage QuickSight**.
 2. In the page that opens, choose **Security and permissions** at left.
 3. Scroll down and, in the **Public access to dashboards** section, choose **Manage**.

Public access to dashboards

Allow dashboard authors to share with unregistered users

Manage

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 QuickSight, 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 QuickSight 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 QuickSight console. Additionally, an eye icon appears on the dashboard on the QuickSight 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 to the dashboard](#). For more information about embedding dashboards for anyone on the internet, see [Turning on public access to visuals and dashboards 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:

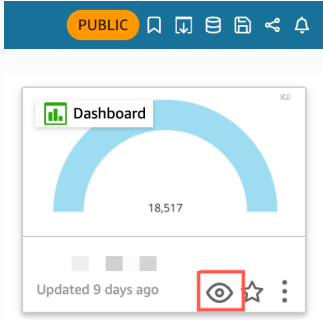

1. From the Amazon QuickSight start page, choose the analysis that is tied to the dashboard that you want to update and make your desired changes. You must be the owner or a co-owner of the analysis.
2. In the analysis, choose the **Share** icon at upper-right, and then choose **Publish dashboard**.
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**.

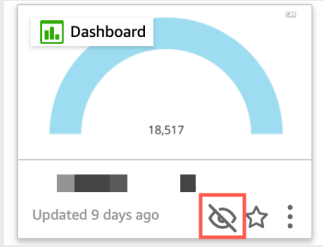
Once you choose **Publish dashboard**, your public dashboard is updated to reflect the new changes.

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	<p>An orange badge appears on the dashboard and an eye icon appears on the dashboard in the Dashboards page.</p>  <p>The screenshot shows a dashboard card with a blue gauge chart. At the top left of the card is a yellow 'PUBLIC' badge. At the bottom right of the card, there is an eye icon with a slash through it, which is highlighted with a red box. Other icons for bookmark, refresh, and share are also visible.</p>
Off	On	No	<p>A grey badge appears on the dashboard and an eye icon with 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.</p>  <p>The screenshot shows a dashboard card with a blue gauge chart. At the top left of the card is a grey 'PUBLIC' badge. At the bottom right of the card, there is a greyed-out eye icon with a slash through it. Other icons for bookmark, refresh, and share are also visible.</p>

Account-level public setting	Dashboard-level public setting	Public access	Visual indicators
			

To turn off public sharing for a single dashboard

1. In QuickSight, 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 QuickSight user account

1. From the Amazon QuickSight start page, choose your user icon at upper right of your browser window, and then choose **Manage QuickSight**.
2. In the page that opens, choose **Security and permissions** at left.
3. Scroll down and, in the **Public access to dashboards** section, choose **Manage**.
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 QuickSight account access to a dashboard with the QuickSight API

Intended audience: Amazon QuickSight developers

Alternatively, you can grant everyone in your account access to the dashboard with the QuickSight 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:AWS DIRECTORY REGION: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 QuickSight API Reference*.

When all users in a QuickSight 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 QuickSight dashboard using the QuickSight API

Alternatively, you can grant anyone on the internet access to the dashboard with the Amazon QuickSight 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 QuickSight account access to a dashboard with the QuickSight 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 QuickSight 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 your view of a Amazon QuickSight 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.

The link to your view of the dashboard is available to anyone who you send it to for up to three months.

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 QuickSight can access the link.

Scheduling and sending reports by email

Important

Amazon QuickSight in the China Beijing region (BJS) uses an internal email service (Amazon SES) in the China Ningxia region (ZHY) to send emails to QuickSight users. Customer data that's included in scheduled reports, alerts, and other features are passed by email from BJS to ZHY before it reaches QuickSight users.

The following QuickSight features send emails through Amazon SES in the ZHY 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 QuickSight subscription. To receive email reports, the users or group members must meet the following conditions:

- They are part of your Amazon QuickSight subscription.
- You already shared the dashboard with them.
- They have completed sign-up process to activate their subscription as Amazon QuickSight readers, authors, or admins.
- Amazon QuickSight can't send scheduled emails to more than 5,000 members.

Amazon QuickSight 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.

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 QuickSight](#).

You can create up to five schedules for each dashboard.

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.

Configuring email report settings for a dashboard in Amazon QuickSight Enterprise edition

Applies to: Enterprise Edition

In Amazon QuickSight 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 QuickSight](#).

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:

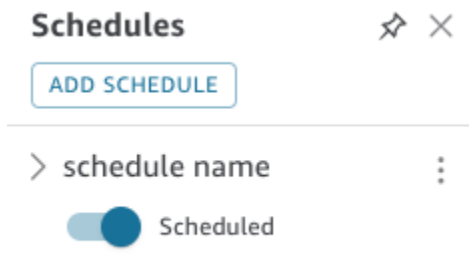
- (Recommended) Choose recipients from the provided access list to specify and maintain a list of people who you want to email reports to. You can use the search box to find people by email or group name.
- To send reports to all of the dashboard's subscribers, choose **Send email report to all users 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 QuickSight 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 QuickSight 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 QuickSight.




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 QuickSight Enterprise edition and that you have shared the dashboard with intended recipients.

To create or change an email report

1. Open Amazon QuickSight 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.

< New Schedule

Schedule name

100 characters remaining

Description (optional)

200 characters remaining

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.

Content ^

PDF
Paginated document

Sheet *

Sheet 1 ▼

CSV New
Data from tables

▼ Sheet 1

Pivot Table 1

▼ Sheet 2

Pivot Table 2

Table 1

> Sheet 3

2 / 5

CANCEL

DONE

Excel New
Data from tables

> Sheet 1

▼ Sheet 2

Pivot Table 2

Table 1

▼ Sheet 3

Table 1

2 / 5

CANCEL

DONE

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)**.

- For **Start date**, choose the start date and runtime that you want to send the first report on.
- For **Timezone**, choose the time zone from the dropdown.

The screenshot shows a configuration panel for a report. The 'Dates' section is expanded, revealing the following settings:

- Repeat:** Monthly
- On:** First
- Day(s):** Wednesday
- Start date:** 2023/06/15
- Begin run at:** 11:00 AM
- Time zone:** America/Los_Angeles

Below the 'Dates' section is the 'Email' section, which is currently collapsed. At the bottom of the panel are two buttons: 'SAVE' and 'SEND TEST'.

- In the **Email** tab, for **E-mail subject line**, enter a custom subject line, or leave it blank to use the report title.
- Enter the email addresses of the QuickSight 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.
- For **Email header**, enter the header that you want the email report to show.
- (Optional) For **E-mail body text**, leave it blank or enter a custom message to display at the beginning of the email.

Email ^

Email subject line (optional)

Email recipients


Send to all users with access


Email header text


Email body text (optional)

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.

Include sheet in email body

 PDF Sheet 1
<input checked="" type="radio"/> Download link Requires login. Expires after 1 year.
<input type="radio"/> File attachment 10 MB for all attachments

 CSVs 1 selected
<input checked="" type="radio"/> Download link Requires login. Expires after 1 year.
<input type="radio"/> File attachment 10 MB for all attachments

 Excel 1 workbook with 2 sheets
<input checked="" type="radio"/> Download link Requires login. Expires after 1 year.
<input type="radio"/> File attachment 10 MB for all attachments

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.



Subscribing to email reports in Amazon QuickSight

In Enterprise edition, Amazon QuickSight authors can set up subscriptions to a dashboard in report form. For more information, see [Scheduling and sending reports by email](#). QuickSight 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 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 QuickSight

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 QuickSight 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 QuickSight 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.

For more information on KPI, Gauge, Table, or Pivot table visuals, see [Visual types in Amazon QuickSight](#).

Note

You can't create alerts for visuals in an embedded dashboard or from the QuickSight 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.

Creating Alerts


Use the following procedure to create threshold alerts for KPI or Gauge visuals in a dashboard.

To create an alert

1. Open QuickSight and navigate to the dashboard that you want.

For more information about viewing dashboards as a dashboard subscriber in QuickSight, see [Interacting with Amazon QuickSight 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](#).

Category Purchases and Details			
Category	Segment	Ship Mode	Profit
[-] Auto	[-] Catalog	First Class	\$134.95
		Same Day	\$187.28
		Second Class	\$289.73
		Standard Class	\$1,392.72
	[-] In-Store	First Class	\$177.99
		Same Day	\$24.35
		Second Class	\$206.81
		Standard Class	\$659.72
	[-] Online	First Class	\$437.72
		Same Day	\$191.83
		Second Class	\$802.66
		Standard Class	\$2,022.52
[-] Baby Supplies	[-] Catalog	First Class	\$54.02
		Same Day	\$11.32
		Second Class	\$34.56

Create Alert

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.

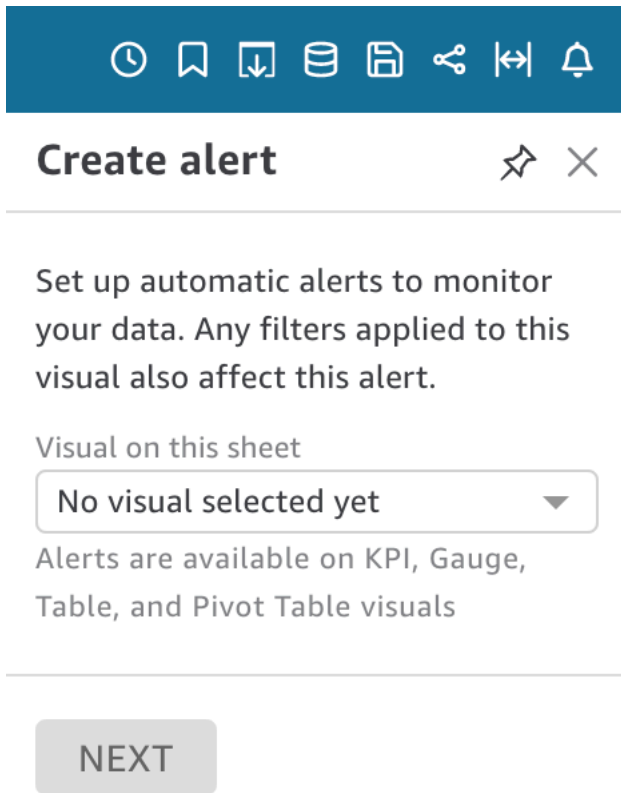
You can choose the following conditions.

- **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 QuickSight account.



Create alert ✨ ✕

Set up automatic alerts to monitor your data. Any filters applied to this visual also affect this alert.

Visual on this sheet

No visual selected yet ▼

Alerts are available on KPI, Gauge, Table, and Pivot Table visuals

NEXT

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 QuickSight, 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 QuickSight, 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.

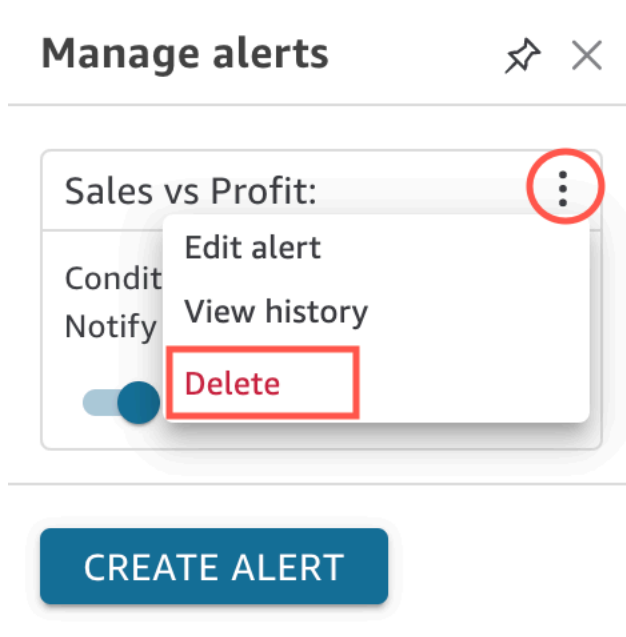
To turn on or turn off an existing alert

1. Open QuickSight, 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 QuickSight, 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, QuickSight 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, QuickSight 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, QuickSight attempts to check your data for threshold breaches again. After four failures, QuickSight turns off the alert and notifies you that the alert is turned off. If the alert can be checked again, QuickSight 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, QuickSight 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 QuickSight, choose **Datasets** in the navigation bar at left.
2. Choose the dataset name 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.

Alert Permissions

If you're an administrator, you can control who in your organization can set threshold alerts in QuickSight by creating a custom permissions policy. To set custom permissions in QuickSight, choose your user name at the upper-right corner of any QuickSight page, choose **Manage QuickSight**, and then choose **Manage permissions**.

Printing a dashboard or analysis

You can print a dashboard or an analysis in Amazon QuickSight.

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 QuickSight 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 Amazon QuickSight 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, QuickSight 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 reports by email](#).

Error codes for failed PDF export jobs

When you generate PDF reports in Amazon QuickSight, you may encounter instances where your request to generate a PDF report fails. There are many reasons why a failure might occur. QuickSight 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 QuickSight returns when a PDF export job fails.

Error code	Guidance
INVALID_DATAPREP_SYNTAX	Check the syntax for your calculated fields, and try again.
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.
SPIICE_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.

Error code	Guidance
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.
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 QuickSight

Applies to: Enterprise Edition

In Amazon QuickSight Enterprise edition, your team members can create personal and shared folders to add hierarchical structure to QuickSight 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 QuickSight:

- 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 QuickSight administrator.
 - **Shared restricted folders** are a type of shared folder in QuickSight 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.

Restricted folders can only be created with the QuickSight `CreateFolder` API operation.

All shared folders are visible to people who have access to them.

Topics

- [Overview of QuickSight folders](#)
- [Permissions overview for shared folders](#)
- [Create and manage permissions for shared folders](#)
- [Considerations](#)
- [Creating scaled folders with the Amazon CLI](#)

Overview of QuickSight folders

In Amazon QuickSight, 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)

QuickSight 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 overview for shared folders](#).

The following table summarizes the actions that a QuickSight 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

Action	Owner	Contributor	Viewer
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
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

*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 QuickSight 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
Modify assets in the folder.	Yes	No
Delete assets in 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

Action	Contributor	Viewer
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

The owner role is not supported for restricted shared folders.

Permissions overview for shared folders

Shared folders have three permission levels. To set folder-level permissions for a user or group, see [Create and manage permissions for 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:

- QuickSight 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 QuickSight 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.

Create and manage permissions for shared folders

Shared folders (unrestricted)

To create a shared folder and to share the folder with one or more groups in the QuickSight 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 permissions of a shared folder.

1. From the left navigation, choose **Shared folders** and 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.

```
aws quicksight update-folder-permissions \  
--aws-account-id AWSACCOUNTID \  
--region us-east-1 \  
--folder-id example-folder-name \  
--grant-permissions Principal=arn:aws-cn:quicksight::us-east-1::AWSACCOUNTID:user/  
default/:username,Actions=quicksight:CreateFolder,quicksight:DescribeFolder, \  
quicksight:CreateFolderMembership,quicksight>DeleteFolderMembership,quicksight:DescribeFolderPe  
\
```

Considerations

The following limitations apply to folders in Amazon QuickSight:

- You can't share folders in your Amazon account with people in other Amazon Web Services accounts.
- For people who have QuickSight 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 QuickSight CLI.
- Subfolders are not supported for restricted shared folders.
- Data sources that are located in restricted shared folders must be created with the `CreateDataSource` API operation.

For Amazon QuickSight quotas, the Service Quotas console provides the most accurate and up-to-date information. You can do the following in the Service Quotas console:

- [View the default Amazon QuickSight quotas for each Amazon Web Services Region](#)
- [Request quota increases for adjustable quotas](#)

Creating scaled folders with the Amazon CLI

You can use the Amazon QuickSight CLI 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 QuickSight assets with thousands of users.

Scaled folders can only be created with the QuickSight CLI. 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 QuickSight CLI or the QuickSight console.

Each Amazon QuickSight 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 QuickSight CLI.

Prerequisites

Before you begin, verify that you have an Amazon Identity and Access Management role that grants the CLI user access to call the QuickSight 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.

```
{
  "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 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"
```

Monitoring data in Amazon QuickSight

Amazon QuickSight sends metrics to Amazon CloudWatch that you can use to observe and respond to the availability and performance of your QuickSight environment in near real time. Currently, you can monitor metrics for QuickSight dashboards, visuals, and dataset ingestions to provide your readers with a consistent, high-performing, and uninterrupted experience on Amazon QuickSight.

For more information about using Amazon CloudWatch, see the [Amazon CloudWatch User Guide](#).

Following, you can find information on how to access CloudWatch and a list of all current supported metrics.

Accessing Amazon QuickSight metrics in Amazon CloudWatch

Use the following procedure to access Amazon QuickSight metrics in Amazon CloudWatch.

To access QuickSight metrics in CloudWatch

1. Sign in to the Amazon Web Services account that's associated with your QuickSight 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**.

Metrics

The Amazon/QuickSight namespace includes the following metrics for monitoring traffic and latency of your Amazon QuickSight dashboards and ingestions.

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	DashboardId	Count

Metric	Description	Dimension	Unit
	<p>been viewed. This number includes all access patterns such as web, mobile, and embedded.</p> <p>The most useful statistic for this metric is SUM, which represents the total number of dashboard views during a set period of time.</p>		
DashboardViewLoadTime	<p>The amount of time that it takes for a QuickSight dashboard to load. The measurement begins when a user reaches the QuickSight dashboard and ends when all of the dashboard's visuals finish rendering.</p> <p>The most useful statistic for this metric is AVERAGE, which represents the average load time of a QuickSight dashboard during a set period of time.</p>	DashboardId	Millisecond

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	<p>The number of failed ingestions.</p> <p>The most useful statistic for this metric is SUM, which represents the total number of failed ingestions during a set period of time.</p>	DatasetId	Count
IngestionInvocationCount	<p>The number of ingestions that have been initiated. This includes scheduled and manual ingestions that are initiated through the console and the Amazon QuickSight API operations.</p> <p>The most useful statistic for this metric is SUM, which represents the total number of ingestions initiated during a set period of time.</p>	DatasetId	Count

Metric	Description	Dimension	Unit
IngestionLatency	<p>The time period between the initiation of an ingestion to the completion of the ingestion.</p> <p>The most useful statistic for this metric is AVERAGE, which represents the average runtime of ingestions during a set period of time.</p>	DatasetId	Second
IngestionRowCount	<p>The number of successful row ingestions.</p> <p>The most useful statistic for this metric is SUM, which represents the total amount of data ingested during a set period of time.</p>	DatasetId	Count

Per-visual metrics

The following metrics track the load times and error counts of individual visuals on a QuickSight dashboard. You can find these metrics under the Amazon/QuickSight/Visual Metrics group in CloudWatch.

Metric	Description	Dimension	Unit
VisualLoadTime	<p>The time that it takes for a QuickSight 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 QuickSight service, and then back to client.</p> <p>The most useful statistic for this metric is AVERAGE, which represents the average load time of a visual during a set period of time.</p>	<ul style="list-style-type: none"> • DashboardId • SheetId • VisualId 	Millisecond
VisualLoadErrorCount	<p>The number of times that a QuickSight 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.</p> <p>The most useful statistic for this metric is SUM, which</p>	<ul style="list-style-type: none"> • DashboardId • SheetId • VisualId 	Count

Metric	Description	Dimension	Unit
	represents the total number of failed visual loads during a set period.		

Aggregate metrics

The Amazon/QuickSight namespace includes the following aggregate metrics for monitoring traffic and latency of your Amazon QuickSight dashboards and ingestions.

Aggregate dashboard metrics

The following metrics track view counts and load times of all dashboards in a QuickSight account and region. You can find these metrics under the Amazon/QuickSight/Aggregate Metrics group in CloudWatch.

Metric	Description	Unit
DashboardViewCount	<p>The number of times that all QuickSight dashboards have been viewed across the entire QuickSight account in the region. This number is an aggregate that includes all access patterns such as web, mobile, and embedded.</p> <p>The most useful statistic for this metric is SUM, which represents the total number of QuickSight dashboard views during a set period of time.</p>	Count

Metric	Description	Unit
DashboardViewLoadTime	<p>The amount of time that it takes for all QuickSight dashboards to load. The measurement begins when a user navigates to the QuickSight dashboard and ends when all of the dashboard's visuals finish rendering.</p> <p>The most useful statistic for this metric is AVERAGE, which represents the average load time of all QuickSight dashboard during a set period of time.</p>	Millisecond

Aggregate ingestion metrics

The following metrics track all ingestions associated with a QuickSight account and Amazon Web Services Region. You can find these metrics under the Amazon/QuickSight/Aggregate Metrics group in CloudWatch.

Metric	Description	Unit
IngestionErrorCount	<p>The number of failed ingestions.</p> <p>The most useful statistic for this metric is SUM, which represents the total number of failed ingestion during a set period.</p>	Count

Metric	Description	Unit
IngestionInvocationCount	<p>The number of ingestions that have been initiated . This includes scheduled and manual ingestions that are initiated through the console and the QuickSight API operations.</p> <p>The most useful statistic for this metric is SUM, which represents the total number of ingestions initiated during a set period of time.</p>	Count
IngestionLatency	<p>The time period between the initiation of an ingestion to the completion of the ingestion.</p> <p>The most useful statistic for this metric is AVERAGE, which represents the average runtime of ingestions during a set period of time.</p>	Second
IngestionRowCount	<p>The number of successful row ingestions.</p> <p>The most useful statistic for this metric is SUM, which represents the total amount of data ingested during a set period of time.</p>	Count

Aggregate visual metrics

The following metrics track load times and error counts of all visuals on a dashboard and in a QuickSight account in a Region. You can find these metrics under the Amazon/QuickSight/Aggregate Metrics group for CloudWatch.

Metric	Description	Unit
VisualLoadTime	<p>The time that it takes for all QuickSight 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 QuickSight service, and then back to the client.</p> <p>The most useful statistic for this metric is AVERAGE, which represents the average load time of all visuals during a set period of time.</p>	Millisecond
VisualLoadErrorCount	<p>The number of times that all QuickSight visuals that belong to the QuickSight account fail to complete a data query for an initial paint.</p> <p>The most useful statistic for this metric is SUM, which represents the total number of failed visuals during a set period.</p>	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 QuickSight 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 in consumed or purchased capacity is made.	MegaBytes

Dimensions

Following is a list of Amazon QuickSight metric dimensions that appear in Amazon CloudWatch.

Dimension	Description
DashboardId	The public ID of a QuickSight dashboard. You can use the <code>ListDashboards</code> API operation to

Dimension	Description
	<p>see a list of every dashboard in your Amazon QuickSight account. For more information, see ListDashboards in the <i>Amazon QuickSight API Reference</i>.</p>
DatasetId	<p>The public ID of a QuickSight dataset.</p> <p>You can use the <code>ListDataSets</code> API operation to see a list of every dataset in your Amazon QuickSight account. For more information, see ListDataSets in the <i>Amazon QuickSight API Reference</i>.</p>
SheetId	<p>The public ID of a QuickSight sheet.</p>
VisualId	<p>The public ID of a QuickSight visual.</p>

Graphing metrics by using the Amazon CloudWatch console

You can also use the Amazon CloudWatch console to graph metric data generated by Amazon QuickSight. For more information, see [Graphing metrics](#) in the *Amazon CloudWatch User Guide*.

Creating alarms by using the Amazon CloudWatch console

You can create a Amazon CloudWatch alarm that monitors CloudWatch metrics for your Amazon QuickSight 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*.

Developing with Amazon QuickSight

We provide API operations for Amazon QuickSight, and also software development kits (SDKs) for Amazon that enable you to access Amazon QuickSight 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 QuickSight 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 QuickSight 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 QuickSight

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)

- [QuickSight](#)

Amazon SDK for Ruby

- [Aws::QuickSight](#)

Terminology and concepts

This section provides a list of terms for development in Amazon QuickSight.

Anonymous QuickSight User: – A temporary Amazon QuickSight 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 QuickSight 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 QuickSight. Amazon approves callers through their invoker's identity. This is done to avoid having to explicitly add multiple accounts belonging to the same Amazon QuickSight 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 QuickSight 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 QuickSight API operations.

QuickSight dashboard: – An entity which identifies QuickSight reports, created from analyses or templates. QuickSight 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.

QuickSight 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.

QuickSight user: – This is an Amazon QuickSight 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 QuickSight.

Using the Amazon QuickSight developer portal

The [QuickSight dev portal](#) helps you learn by example how to use the QuickSight 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 [QuickSight 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 QuickSight icon, top left.

Developing applications with the Amazon QuickSight 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 QuickSight API Reference](#).

Before you can call the Amazon QuickSight 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 QuickSight Amazon Resource Name (ARN) identifier.

For more information, see the following:

- [IAM policy examples for Amazon QuickSight](#)
- [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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight API operation. This example returns a list of Amazon QuickSight 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 QuickSight.

1. To get started, create a Java project in your IDE.
2. Import the Amazon QuickSight SDK into your new project, for example:
`AWSQuickSightJavaClient-1.11.x.jar`
3. Once your IDE indexes the Amazon QuickSight 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 QuickSight 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 QuickSight 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();  
}
```

- Now, we can use the above client to list all the users in our Amazon QuickSight account.

Note

You have to provide the Amazon account ID that you used to subscribe to Amazon QuickSight. 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());  
    });
```

- 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 QuickSight interface. Alternatively, find it within the `com.amazonaws.services.quicksight` package in the Amazon QuickSight JavaClient JAR file.

JavaScript (Node.js) SDK

Use the following procedure to interact with Amazon QuickSight using Node.js.

- Set up your node environment using the following commands:
 - `npm install aws-sdk`
 - `npm install aws4`
 - `npm install request`
 - `npm install url`
- 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 QuickSight operations along with their URL request parameters, followed by a list of Amazon QuickSight 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 QuickSight.

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.

- 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.
- 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 QuickSight 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 QuickSight 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 QuickSight events integration

With Amazon EventBridge, you can respond automatically to events in Amazon QuickSight 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 QuickSight events to Amazon CloudWatch](#)
- [Creating rules to send Amazon QuickSight events to Amazon Lambda](#)

Supported events

QuickSight currently supports the following events.

Asset type	Action	Event detail type	Event detail
Dashboard	Create	QuickSight Dashboard Creation Successful	<pre>{ "dashboardId": "6fdb328-ebbd-457f-aa02-9780173afc83", "versionNumber": 1 }</pre>
Dashboard	Create	QuickSight Dashboard Creation Failed	<pre>{ "dashboardId": "6fdb328-ebbd-457f-aa02-9780173afc83", "versionNumber": 1, "errors": [{ "Type": "PARAMETER_NOT_FOUND", "Message": "Missing property abc" }, { "Type": "DATA_SET_NOT_FOUND", </pre>

Asset type	Action	Event detail type	Event detail
			<pre> "Message" : "Cannot find dataset with id abc" }] }</pre>
Dashboard	Update	QuickSight Dashboard Update Successful	<pre>{ "dashboar dId": "6fdb328 -ebbd-457f- aa02-9780173afc8 3", "versionN umber": 1 }</pre>

Asset type	Action	Event detail type	Event detail
Dashboard	Update	QuickSight Dashboard Update Failed	<pre>{ "dashboardId": "6fdb328-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	Publish	QuickSight Dashboard Published Version Updated	<pre>{ "dashboardId": "6fdb328-ebbd-457f-aa02-9780173afc83", "versionNumber": 2 }</pre>

Asset type	Action	Event detail type	Event detail
Dashboard	Delete	QuickSight Dashboard Deleted	<pre>{ "dashboardId": "6fdbbc328-ebbd-457f-aa02-9780173afc83" }</pre>
Analysis	Create	QuickSight Analysis Creation Successful	<pre>{ "analysisId": "e5f37119-e24c-4874-901a-af9032b729b5" }</pre>

Asset type	Action	Event detail type	Event detail
Analysis	Create	QuickSight Analysis Creation 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>
Analysis	Delete	QuickSight Analysis Deleted	<pre>{ "analysis Id": "e5f37119- e24c-4874-901a- af9032b729b5" }</pre>

Asset type	Action	Event detail type	Event detail
VPC connection	Create	QuickSight VPC Connection Creation Successful	<pre>{ "vpcConnectionId": "53d34238-57e7-488d-b99a-a0037d275a4e", "availabilityStatus": "CREATION_SUCCESSFUL" }</pre>
VPC connection	Create	QuickSight VPC Connection Creation Failed	<pre>{ "vpcConnectionId": "53d34238-57e7-488d-b99a-a0037d275a4e", "availabilityStatus": "CREATION_FAILED" }</pre>
VPC connection	Update	QuickSight VPC Connection Update Successful	<pre>{ "vpcConnectionId": "53d34238-57e7-488d-b99a-a0037d275a4e", "availabilityStatus": "UPDATE_SUCCESSFUL" }</pre>

Asset type	Action	Event detail type	Event detail
VPC connection	Update	QuickSight VPC Connection Update Failed	<pre>{ "vpcConnectionId": "53d34238-57e7-488d-b99a-a0037d275a4e", "availabilityStatus": "UPDATE_FAILED" }</pre>
VPC connection	Delete	QuickSight VPC Connection Deletion Successful	<pre>{ "vpcConnectionId": "53d34238-57e7-488d-b99a-a0037d275a4e", "availabilityStatus": "DELETED" }</pre>
VPC connection	Delete	QuickSight VPC Connection Deletion Failed	<pre>{ "vpcConnectionId": "53d34238-57e7-488d-b99a-a0037d275a4e", "availabilityStatus": "DELETION_FAILED" }</pre>

Asset type	Action	Event detail type	Event detail
Folder	Create	QuickSight Folder Created	<pre>{ "folderId": "77e307e8-b41b-472a-90e8-fe3f471537be", "parentFolderArn": "arn:aws-quicksight:us-east-1:123456789012:folder/098765432134" }</pre>
Folder	Update	QuickSight Folder Updated	<pre>{ "folderId": "77e307e8-b41b-472a-90e8-fe3f471537be" }</pre>
Folder	Delete	QuickSight Folder Deleted	<pre>{ "folderId": "77e307e8-b41b-472a-90e8-fe3f471537be" }</pre>

Asset type	Action	Event detail type	Event detail
Folder	Membership update	QuickSight Folder Membership Updated	<pre>{ "folderId": "77e307e8-b41b-472a-90e8-fe3f471537be", "membersAdded": ["arn:aws-cn:quicksight:us-east-1:123456789012:analysis/e5f37119-e24c-4874-901a-af9032b729b5"], "membersRemoved": [] }</pre>
Dataset	Create	QuickSight Dataset Created	<pre>{ "datasetId": "a6553a81-f97e-4ffa-a860-baea63196efa" }</pre>
Dataset	Update	QuickSight Dataset Updated	<pre>{ "datasetId": "a6553a81-f97e-4ffa-a860-baea63196efa" }</pre>

Asset type	Action	Event detail type	Event detail
Dataset	Delete	QuickSight Dataset Deleted	<pre>{ "datasetId": "a6553a81-f97e-4ffa-a860-baea63196efa" }</pre>
DataSource	Create	QuickSight DataSource Creation Successful	<pre>{ "dataSourceId": "230caa6e-dc87-406b-91fb-037f29c32824" }</pre>
DataSource	Create	QuickSight 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>

Asset type	Action	Event detail type	Event detail
DataSource	Update	QuickSight DataSource Update Successful	<pre>{ "datasourceId": "230caa6e-dc87-406b-91fb-037f29c32824" }</pre>
DataSource	Update	QuickSight 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	Delete	QuickSight DataSource Deleted	<pre>{ "datasourceId": "230caa6e-dc87-406b-91fb-037f29c32824" }</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 QuickSight events to Amazon CloudWatch

You can write simple rules to indicate which Amazon QuickSight events interest you and which automated actions to take when an event matches a rule. For example, you can configure Amazon QuickSight to send events to Amazon CloudWatch whenever a Amazon QuickSight 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**.

- In the **Creation method** section, choose **Custom pattern (JSON editor)**.
- 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 QuickSight. 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"
  }
}
```

- For **Targets**, choose **Amazon service > CloudWatch log group**.
- Choose from an existing log group or create a new one by entering a new log group name.
- Optionally, you can add another target for this rule.
- In **Configure tags**, choose **Next**.
- 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 QuickSight events to Amazon Lambda

In this tutorial, you create an Amazon Lambda function that logs the asset events in the Amazon QuickSight 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 QuickSight.

Step 1: Create an 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 QuickSight 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 QuickSight. 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 QuickSight 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 QuickSight event in JSON format, see [Overview of events for Amazon QuickSight](#).

Working with embedded analytics

Important

Amazon QuickSight 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 QuickSight 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 QuickSight developers

With Amazon QuickSight 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 QuickSight 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.

QuickSight supports embedding for these elements:

- QuickSight console (full authoring experience for registered users)
- QuickSight dashboards and visuals (for registered users, anonymous users, public end users)
- QuickSight Q search bar (for registered users and anonymous users)

With an embedded QuickSight console, you embed the full QuickSight experience. Doing this makes it possible to use QuickSight 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 QuickSight console need to be registered as QuickSight 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 QuickSight-supported authentication methods.

With an embedded QuickSight 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:

- QuickSight users authenticated in your Amazon Web Services account by any method supported by QuickSight.
- Unauthenticated visitors to a website or application – This option requires session packs with [capacity pricing](#).
- 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 QuickSight 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 QuickSight Enterprise edition in the Amazon Web Services account where you plan to use embedding.

QuickSight embedding is available in all supported Amazon Web Services Regions.

Topics

- [Embedding overview](#)
- [Customizing embedded assets](#)
- [1-click embedding and public embedding](#)
- [Embedding with the QuickSight APIs](#)

Embedding overview

Applies to: Enterprise Edition

To embed analytics, you can run the Amazon QuickSight embedding API to generate the embed code. Alternatively for dashboards, you can copy an embed code when you share the dashboard in QuickSight. Each option is described below.

1-click enterprise embedding

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 QuickSight 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 visuals and dashboards for registered users with a 1-click embed code](#).

Embedding with the QuickSight APIs

Embedding with the QuickSight API is best for when you want to embed the QuickSight 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 QuickSight APIs](#).

Customizing embedded assets

You can use Amazon QuickSight embedded analytics to embed custom QuickSight assets into your application that are tailored to meet your business needs. For embedded dashboards and visuals, QuickSight authors can add filters and drill downs that readers can access as they navigate the dashboard or visual. Amazon QuickSight developers can also use the QuickSight SDKs to build tighter integrations between their SaaS applications and their QuickSight embedded assets to add datapoint callback actions to visuals on a dashboard at runtime.

For more information about the Amazon QuickSight SDKs, see the `amazon-quicksight-embedding-sdk` on [GitHub](#) or [NPM](#).

Following, you can find descriptions of how to use the QuickSight SDKs to customize your QuickSight embedded analytics.

Topics

- [Add embedded callback actions at runtime in Amazon QuickSight](#)
- [Filtering data at runtime for embedded dashboards and visuals](#)
- [Customize the look and feel of embedded dashboards and visuals](#)

Add embedded callback actions at runtime in Amazon QuickSight

Use embedded datapoint callback actions to build tighter integrations between your software as a service (SaaS) application and your Amazon QuickSight embedded dashboards and visuals. Developers can register datapoints to be called back with the QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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 embedded dashboards and visuals

You can use filter methods in the Amazon QuickSight embedding SDK to leverage the power of QuickSight filters within your software as a service (SaaS) application at runtime. Runtime filters allow business owners to integrate their application with their embedded QuickSight 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 filters on an embedded dashboard or visual from their application with the QuickSight Embedding SDK. Create QuickSight 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 QuickSight Embedding SDK, see the [amazon-quicksight-embedding-sdk](#) on GitHub.

Prerequisites

Before you can get started, make sure that you are using the QuickSight 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 QuickSight Embedding SDK. For more information about using the QuickSight Embedding SDK, see [\[wrong link in Quip\]](#).

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

Method	Description
	ResponseMessage that indicates whether the addition was successful is returned.
<code>updateFilterGroups(filterGroups: FilterGroup[])</code>	Updates the FilterGroups on the embedded experience that contains the same FilterGroupId as the FilterGroup that is supplied in the parameter. A ResponseMessage that indicates whether the update was successful is returned.
<code>removeFilterGroups(filterGroupsOrIds: FilterGroup[] string[])</code>	Removes the supplied FilterGroups from the dashboard and returns a ResponseMessage 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 embedded dashboards and visuals

You can use the Amazon QuickSight Embedding SDK (version 2.5.0 and higher) to make changes to the theming of your embedded QuickSight dashboards and visuals at runtime. Runtime theming makes it easier to integrate your Software as a service (SaaS) application with your Amazon QuickSight embedded assets. Runtime theming allows you to synchronize the theme of your embedded content with the themes of the parent application that your QuickSight 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 QuickSight](#). For more information about using the QuickSight 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 QuickSight 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 QuickSight, make a `UpdateThemePermissions` API call or use the **Share** icon next to the theme in the QuickSight 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 QuickSight theme. Following is an example of custom theme ARN.

```
arn:aws-cn:quicksight:region:account-id:theme/theme-id
```

QuickSight 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 QuickSight account, make a call to the [ListThemes](#) API operation. For information on preset QuickSight themes, see [Setting a default theme for Amazon QuickSight analyses](#).
- 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>	<p>Replaces the active theme of a dashboard or visual with another theme. If applied, the theme override is removed.</p> <p>An error is returned if you don't have access to the theme or if the theme doesn't exist.</p>
<code>setThemeOverride(themeOverride: ThemeConfiguration)</code>	<p>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.</p> <p>An error is returned if the <code>ThemeConfiguration</code> that you supply is invalid.</p>

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 QuickSight. 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);
```

1-click embedding and public embedding

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 QuickSight.

You can embed a visual or dashboard in your internal application for your registered users. Or you can turn on public sharing in the QuickSight 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 visuals and dashboards for registered users with a 1-click embed code](#)
- [Turning on public access to visuals and dashboards with a 1-click embed code](#)

Embedding 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 QuickSight account. You do so using the embed code that you get when you share the dashboard or from the **Embed visual** menu in QuickSight. You don't have to run the QuickSight embedding API to generate the embed code. You can copy the embed code from QuickSight and paste it in your internal application's HTML code.

When users and groups (or all users on your QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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**.

Embed visual

Embed code 

```
<iframe width="600"
height="400" src="https://us-
east-
1.quicksight.aws.amazon.com/e
mbed/...."></iframe>
```



Visual accessible to all users with whom this dashboard is shared.

IDs for developers

[Copy all IDs](#)

Dashboard

[COPY](#)

Sheet

[COPY](#)

Visual

[COPY](#)

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 QuickSight 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 QuickSight 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 QuickSight 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.

Turning on public access to visuals and dashboards 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 QuickSight. You can also turn on public sharing by using the QuickSight 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 QuickSight 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 QuickSight account.

Step 3: Get the embed code

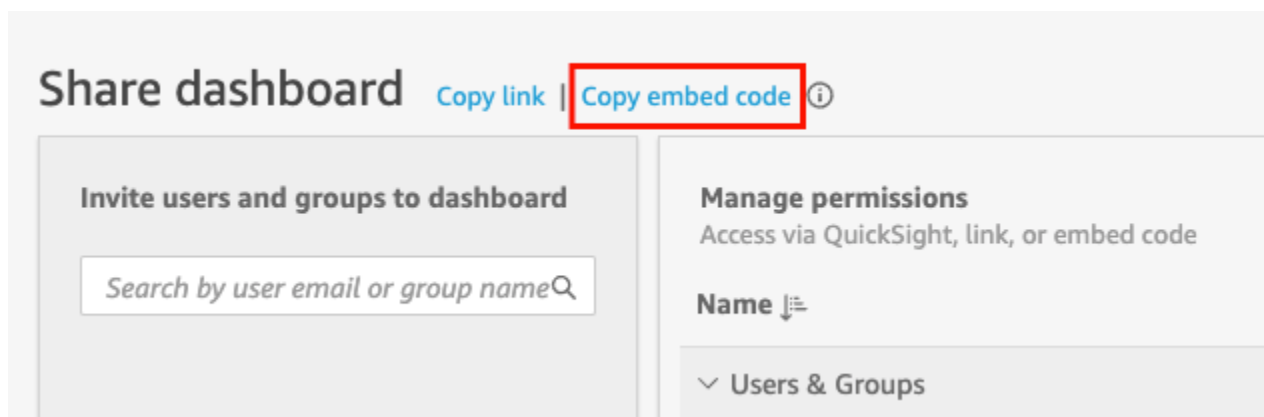
Use the following procedure to get the visual or dashboard embed code.

To get the dashboard embed code

1. Open the published dashboard in QuickSight 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 QuickSight 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 QuickSight 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**.

Embed visual

Embed code 

```
<iframe width="600"
height="400" src="https://us-
east-
1.quicksight.aws.amazon.com/e
mbed/...."></iframe>
```



Visual accessible to all users with whom this dashboard is shared.

IDs for developers

[Copy all IDs](#)

Dashboard

[COPY](#)

Sheet

[COPY](#)

Visual

[COPY](#)

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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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 QuickSight APIs

Applies to: Enterprise Edition

Intended audience: Amazon QuickSight developers

There are only a few steps involved in the actual process of embedding analytics using the QuickSight 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 QuickSight 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 QuickSight folder. If you prefer to use the QuickSight API, use the `DescribeDashboardPermissions` and `UpdateDashboardPermissions` API operations. For more information, see [DescribeDashboardPermissions](#) or [UpdateDashboardPermissions](#) in the *Amazon QuickSight 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 QuickSight administrator must explicitly enable domains where you plan to embed your QuickSight analytics. You can do this by using the **Manage QuickSight, 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 QuickSight administrators. You can also add subdomains as part of a domain. For more information, see [Allow listing domains at runtime with the QuickSight 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 QuickSight API operations.

After all the prerequisites are complete, embedding QuickSight 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 QuickSight using the generated URL. To simplify this process, you can use the Amazon QuickSight Embedding SDK, available on [NPMJS](#) and [GitHub](#). This customized JavaScript SDK is designed to help you efficiently integrate QuickSight 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 dashboards with the QuickSight API](#)
- [Embedding visuals with the QuickSight APIs](#)
- [Embedding the full functionality of the Amazon QuickSight console for registered users](#)
- [Embed the Amazon Q in QuickSight Generative Q&A experience](#)
- [Embedding the Amazon QuickSight Q search bar \(Classic\)](#)
- [Embedding analytics using the GetDashboardEmbedURL and GetSessionEmbedURL API operations](#)

Embedding dashboards with the QuickSight API

Use the following topics to learn about embedding dashboards with the Amazon QuickSight API.

Topics

- [Embedding QuickSight data dashboards for anonymous \(unregistered\) users](#)
- [Embedding QuickSight data dashboards for registered users](#)

Embedding QuickSight data dashboards for anonymous (unregistered) users

Important

Amazon QuickSight 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 QuickSight 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 QuickSight developers

In the following sections, you can find detailed information about how to set up embedded Amazon QuickSight 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 QuickSight 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 QuickSight 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 QuickSight** 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 QuickSight 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.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GenerateEmbedUrlForAnonymousUser"
      ],
      "Resource": [
        "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:namespace/{{namespace}}",
        "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:dashboard/{{dashboardId-1}}",
        "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:dashboard/{{dashboardId-2}}"
      ],
      "Condition": {
        "ForAllValues:StringEquals": {
          "quicksight:AllowedEmbeddingDomains": [
            "https://my.static.domain1.com",
            "https://*.my.static.domain2.com"
          ]
        }
      }
    }
  ]
}
```

```
    }  
  }  
}
```

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.

```
{  
  "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 QuickSight 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 QuickSight, 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");
        }

        @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
          "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::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_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,Visua
```

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 QuickSight developers

In the following section, you can find out how you can use the [QuickSight 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 [QuickSight Embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), QuickSight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, QuickSight provides a responsive layout that is modified as your window size changes. By using the QuickSight 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 Amazon 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 QuickSight 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,
    emitSizeChangeEventOnSheetChange: 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
          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>
</html>
```

```
</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 QuickSight Embedding SDK to load the embedded dashboard on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon QuickSight Embedding SDK](#) from GitHub. This repository is maintained by a group of QuickSight 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 QuickSight data dashboards for registered users

Important

Amazon QuickSight 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 QuickSight 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 QuickSight developers

In the following sections, you can find detailed information about how to set up embedded Amazon QuickSight dashboards for registered users of Amazon QuickSight.

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 QuickSight 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 QuickSight** 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 QuickSight API Reference*.

The following sample policy provides these permissions.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GenerateEmbedUrlForRegisteredUser"
      ],
      "Resource":
        "arn:partition:quicksight:region:accountId:user/namespace/userName",
      "Condition": {
        "ForAllValues:StringEquals": {
          "quicksight:AllowedEmbeddingDomains": [
            "https://my.static.domain1.com",
            "https://*.my.static.domain2.com"
          ]
        }
      }
    }
  ]
}
```

Additionally, if you are creating first-time users who will be Amazon QuickSight 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 QuickSight readers.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
```

```

    "Action": "quicksight:RegisterUser",
    "Resource": "*",
    "Effect": "Allow"
  },
  {
    "Effect": "Allow",
    "Action": [
      "quicksight:GenerateEmbedUrlForRegisteredUser"
    ],
    "Resource": [
      "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:namespace/
{{namespace}}",
      "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:dashboard/
{{dashboardId-1}}",
      "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:dashboard/
{{dashboardId-2}}"
    ],
    "Condition": {
      "ForAllValues:StringEquals": {
        "quicksight:AllowedEmbeddingDomains": [
          "https://my.static.domain1.com",
          "https://*.my.static.domain2.com"
        ]
      }
    }
  }
]
}

```

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 QuickSight. The following example shows a sample trust policy.

```

{
  "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 QuickSight 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 QuickSight, 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 QuickSight. 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: AWS 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 QuickSight from multiple locations.

The role session ID also becomes the user name in QuickSight. You can use this pattern to provision your users in QuickSight 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 QuickSight API operations, see the [QuickSight 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 QuickSight. For Microsoft AD users, you can use `DescribeUser` to get the user ARN.

The first time a user accesses QuickSight, 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 QuickSight, 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 QuickSight 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 [QuickSight Embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), QuickSight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, QuickSight provides a responsive layout that is modified as your window size changes. By using the Amazon QuickSight 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 Amazon 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 QuickSight 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 QuickSight Embedding SDK to load the embedded dashboard on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon QuickSight Embedding SDK](#) from GitHub. This repository is maintained by a group of QuickSight 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 visuals with the QuickSight APIs

You can embed individual visuals that are a part of a published dashboard in your application with the Amazon QuickSight API.

Topics

- [Embedding QuickSight visuals for anonymous \(unregistered\) users](#)
- [Embedding QuickSight visuals for registered users](#)

Embedding QuickSight visuals for anonymous (unregistered) users

Applies to: Enterprise Edition

Intended audience: Amazon QuickSight developers

In the following sections, you can find detailed information about how to set up embedded Amazon QuickSight 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 QuickSight 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 QuickSight 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 QuickSight** 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 QuickSight 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 visual, the error `UnsupportedPricingPlanException` is returned.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GenerateEmbedUrlForAnonymousUser"
      ],
      "Resource": [
        "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:namespace/
        {{namespace}}",
        "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:dashboard/
        {{dashboardId-1}}",
        "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:dashboard/
        {{dashboardId-2}}"
      ],
      "Condition": {
        "ForAllValues:StringEquals": {
          "quicksight:AllowedEmbeddingDomains": [
            "https://my.static.domain1.com",
            "https://*.my.static.domain2.com"
          ]
        }
      }
    }
  ]
}
```

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.

```
{
  "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 QuickSight 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 QuickSight, 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,Visua
```

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 QuickSight developers

In the following section, you can find out how you can use the [QuickSight 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 QuickSight 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 QuickSight [Embedding SDK](#) or by adding this URL into an iframe. If you set a fixed height and width number (in pixels), QuickSight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, QuickSight provides a responsive layout that is modified as your window size changes. By using the QuickSight 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 Amazon 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 QuickSight 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 occured 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 QuickSight Embedding SDK to load the embedded visual on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon QuickSight Embedding SDK](#) from GitHub. This repository is maintained by a group of QuickSight 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 QuickSight visuals for registered users

Applies to: Enterprise Edition

Intended audience: Amazon QuickSight developers

In the following sections, you can find detailed information about how to set up embedded Amazon QuickSight visuals for registered users of Amazon QuickSight.

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 QuickSight 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 QuickSight** 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 QuickSight API Reference*.

The following sample policy provides these permissions.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GenerateEmbedUrlForRegisteredUser"
      ],
      "Resource":
"arn:partition:quicksight:region:accountId:user/namespace/userName",
      "Condition": {
        "ForAllValues:StringEquals": {
          "quicksight:AllowedEmbeddingDomains": [
            "https://my.static.domain1.com",
            "https://*.my.static.domain2.com"
          ]
        }
      }
    }
  ]
}
```

```

    }
  }
]
}

```

Additionally, if you are creating first-time users who will be Amazon QuickSight 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 QuickSight readers.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "quicksight:RegisterUser",
      "Resource": "*",
      "Effect": "Allow"
    },
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GenerateEmbedUrlForRegisteredUser"
      ],
      "Resource": [
        "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:namespace/
        {{namespace}}",
        "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:dashboard/
        {{dashboardId-1}}",
        "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:dashboard/
        {{dashboardId-2}}"
      ],
      "Condition": {
        "ForAllValues:StringEquals": {
          "quicksight:AllowedEmbeddingDomains": [
            "https://my.static.domain1.com",
            "https://*.my.static.domain2.com"
          ]
        }
      }
    }
  ]
}

```

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 QuickSight. The following example shows a sample trust policy.

```
{
  "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 QuickSight user and get the embeddable visual URL on your application server. If you plan to embed visuals for IAM or QuickSight identity types, share the visual with the QuickSight users.

When a QuickSight user accesses your app, the app assumes the IAM role on the QuickSight user's behalf. Then it adds the user to QuickSight, if that QuickSight 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 QuickSight. 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 QuickSight 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
        "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: AWS 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 QuickSight from multiple locations.

The role session ID also becomes the user name in QuickSight. You can use this pattern to provision your users in QuickSight 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 QuickSight API operations, see the [QuickSight 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 QuickSight. For Microsoft AD users, you can use `DescribeUser` to get the user ARN.

The first time a user accesses QuickSight, 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 QuickSight, 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,Visua
```

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 QuickSight 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 QuickSight 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 [QuickSight Embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), QuickSight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, QuickSight provides a responsive layout that is modified as your window size changes. By using the Amazon QuickSight 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 Amazon 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 QuickSight 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;
              }
            }
          },
        };
      };
    </script>
  </head>
</html>
```

```
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 occured 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 QuickSight Embedding SDK to load the embedded visual on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon QuickSight Embedding SDK](#) from GitHub. This repository is maintained by a group of QuickSight 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 full functionality of the Amazon QuickSight console for registered users

Important

Amazon QuickSight 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 QuickSight 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 QuickSight developers

With Enterprise edition, in addition to providing read-only dashboards you can also provide the Amazon QuickSight 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 QuickSight through an embedded console need to belong to the author or admin security cohort. Readers don't have enough access to use the QuickSight 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 QuickSight Console](#).
- For information about using namespaces to isolate multitenancy users, groups, and QuickSight assets, see [QuickSight Namespaces](#).
- For information about adding your own branding to an embedded QuickSight console, see [Using Themes in QuickSight](#) and the [QuickSight Theme API Operations](#).

In the following sections, you can find detailed information about how to set up embedded Amazon QuickSight 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](#)

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 QuickSight assumes a role that gives them Amazon QuickSight access and permissions to the console session. To make this possible, create an IAM role in your AWS 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 QuickSight 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 QuickSight** 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.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "quicksight:RegisterUser",
      "Resource": "*",
      "Effect": "Allow"
    },
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GenerateEmbedUrlForRegisteredUser"
      ],
      "Resource": [
        "arn:partition:quicksight:region:accountId:user/namespace/userName"
      ],
      "Condition": {
        "ForAllValues:StringEquals": {
          "quicksight:AllowedEmbeddingDomains": [
            "https://my.static.domain1.com",
            "https://*.my.static.domain2.com"
          ]
        }
      }
    }
  ]
}
```


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.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GenerateEmbedUrlForRegisteredUser"
      ],
      "Resource": [
        "arn:partition:quicksight:region:accountId:user/namespace/userName"
      ],
      "Condition": {
        "ForAllValues:StringEquals": {
          "quicksight:AllowedEmbeddingDomains": [
            "https://my.static.domain1.com",
            "https://*.my.static.domain2.com"
          ]
        }
      }
    }
  ]
}
```

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 QuickSight. The following example shows a sample trust policy.

```
{
  "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 QuickSight, 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 QuickSight. 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.RegisteredUserQuickSightConsoleEmbeddingConfigurati

/**
 * 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
                "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: AWS 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 QuickSight, 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 QuickSight from multiple locations.

The role session ID also becomes the user name in QuickSight. You can use this pattern to provision your users in QuickSight 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 QuickSight API operations, see the [QuickSight 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 QuickSight. For Microsoft AD users, you can use `DescribeUser` to get the user ARN.

The first time a user accesses QuickSight, 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 \  
  --group-name TeamA1 \  
  --user-name jhnd
```

```
--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 QuickSight, and who has access to the QuickSight 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 QuickSight 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 QuickSight [Embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), QuickSight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, QuickSight provides a responsive layout that is modified as your window size changes. By using the Amazon QuickSight 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 Amazon 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 QuickSight 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 QuickSight 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 QuickSight Embedding SDK](#) from GitHub. This repository is maintained by a group of QuickSight 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
```

Embed the Amazon Q in QuickSight Generative Q&A experience

Intended audience: Amazon QuickSight 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

- [Embed the Amazon Q in QuickSight Generative Q&A experience for registered users](#)
- [Embedding the Generative Q&A experience for anonymous \(unregistered\) users](#)

Embed the Amazon Q in QuickSight 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 QuickSight.

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 QuickSight 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 QuickSight** 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 QuickSight API Reference*.

The following sample policy provides these permissions.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GenerateEmbedUrlForRegisteredUser"
      ],
      "Resource":
        "arn:partition:quicksight:region:accountId:user/namespace/userName",
      "Condition": {
        "ForAllValues:StringEquals": {
          "quicksight:AllowedEmbeddingDomains": [
            "https://my.static.domain1.com",
            "https://*.my.static.domain2.com"
          ]
        }
      }
    }
  ]
}
```

```

    }
  }
}

```

Also, if you're creating first-time users who will be Amazon QuickSight 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 QuickSight readers.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "quicksight:RegisterUser",
      "Resource": "*",
      "Effect": "Allow"
    },
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GenerateEmbedUrlForRegisteredUser"
      ],
      "Resource": [
        "arn:partition:quicksight:region:accountId:user/namespace/userName"
      ],
      "Condition": {
        "ForAllValues:StringEquals": {
          "quicksight:AllowedEmbeddingDomains": [
            "https://my.static.domain1.com",
            "https://*.my.static.domain2.com"
          ]
        }
      }
    }
  ]
}

```

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 QuickSight.

The following example shows a sample trust policy.

```
{
  "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 QuickSight 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 QuickSight, 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 QuickSight. 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.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 QuickSight from multiple locations.

The role session ID also becomes the user name in QuickSight. You can use this pattern to provision your users in QuickSight 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 QuickSight API operations, see the [QuickSight 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 QuickSight. For Microsoft AD users, you can use `DescribeUser` to get the user Amazon Resource Name (ARN).

The first time a user accesses QuickSight, 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 QuickSight, 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 QuickSight 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.

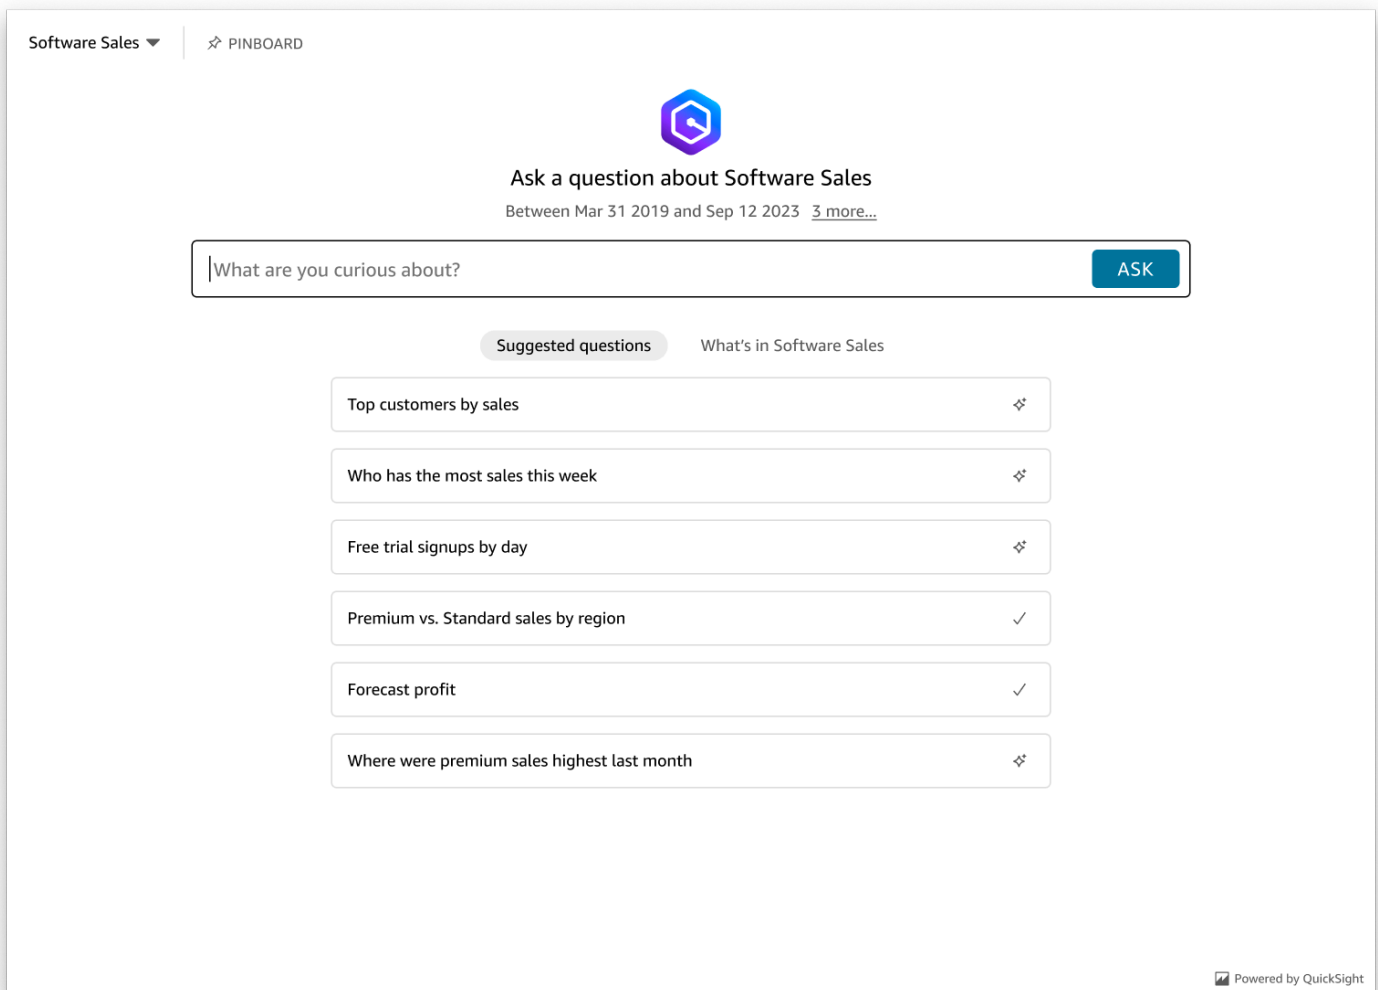
The following shows an example response from `generate-embed-url-for-registered-user`.

```
//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 [QuickSight embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), QuickSight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, QuickSight 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 QuickSight 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 and embedding](#).

You can use the QuickSight 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 QuickSight users have in the 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 QuickSight 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.

When you use the QuickSight Embedding SDK, the Generative Q&A experience on your page is dynamically resized based on the state. By using the QuickSight 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;
              }
            }
          },
        };
      };
    </script>
  </head>
  <body>
    <div id="experience-container">
      <div id="quicksight-embedding">
        <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;
                  }
                }
              },
            };

            embedGenerativeQnA(embeddingContext, frameOptions);
          };

          embedGenerativeQnA();
        </script>
      </div>
    </div>
  </body>
</html>
```



```

        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 QuickSight 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 QuickSight embedding SDK](#) from GitHub. This repository is maintained by a group of QuickSight 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 [QuickSight embedding SDK](#).

Embedding the Generative Q&A experience for anonymous (unregistered) users

Intended audience: Amazon QuickSight 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 QuickSight 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 QuickSight** 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 QuickSight API Reference*.

The following sample policy provides these permissions.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GenerateEmbedUrlForAnonymousUser"
      ],
      "Resource": [
        "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:namespace/
        {{namespace}}",
        "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:dashboard/
        {{dashboardId-1}}",
        "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:dashboard/
        {{dashboardId-2}}"
      ],
      "Condition": {
        "ForAllValues:StringEquals": {
          "quicksight:AllowedEmbeddingDomains": [
            "https://my.static.domain1.com",
            "https://*.my.static.domain2.com"
          ]
        }
      }
    }
  ]
}
```

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.

```
{
  "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 QuickSight, 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 QuickSight 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 QuickSight 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 [QuickSight embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), QuickSight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, QuickSight 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 QuickSight 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 and embedding](#).

You can use the QuickSight 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 QuickSight users have in the 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'`.

Ask about Software Sales



What are you curious about?

ASK

After you configure the `panelType` property, use the QuickSight 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.

When you use the QuickSight Embedding SDK, the Generative Q&A experience on your page is dynamically resized based on the state. With the QuickSight 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 QuickSight 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 QuickSight embedding SDK](#) from GitHub. This repository is maintained by a group of QuickSight 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 [QuickSight embedding SDK](#).

Embedding the Amazon QuickSight Q search bar (Classic)

Intended audience: Amazon QuickSight developers

Note

The embedded QuickSight Q search bar provides the classic QuickSight Q&A experience. QuickSight 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 [Embed the Amazon Q in QuickSight Generative Q&A experience](#).

Use the following topics to learn about embedding the Q search bar with the QuickSight APIs.

Topics

- [Embedding the Amazon QuickSight Q search bar for registered users](#)
- [Embedding the Amazon QuickSight Q search bar for anonymous \(unregistered\) users](#)

Embedding the Amazon QuickSight Q search bar for registered users

Applies to: Enterprise Edition

Intended audience: Amazon QuickSight developers

Note

The embedded QuickSight Q search bar provides the classic QuickSight Q&A experience. QuickSight 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 [Embed the Amazon Q in QuickSight Generative Q&A experience](#).

In the following sections, you can find detailed information about how to set up an embedded Amazon QuickSight Q search bar for registered users of QuickSight.

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 QuickSight Q search bar embedding functionalities](#)

Step 1: Set up permissions

Note

The embedded QuickSight Q search bar provides the classic QuickSight Q&A experience. QuickSight 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 [Embed the Amazon Q in QuickSight 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 QuickSight 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 QuickSight** 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 QuickSight API Reference*.

The following sample policy provides these permissions.

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

```

    "Resource":
      "arn:partition:quicksight:region:accountId:user/namespace/userName",
    "Condition": {
      "ForAllValues:StringEquals": {
        "quicksight:AllowedEmbeddingDomains": [
          "https://my.static.domain1.com",
          "https://*.my.static.domain2.com"
        ]
      }
    }
  }
]
}

```

Also, if you're creating first-time users who will be Amazon QuickSight 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 QuickSight readers.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Action": "quicksight:RegisterUser",
      "Resource": "*",
      "Effect": "Allow"
    },
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GenerateEmbedUrlForRegisteredUser"
      ],
      "Resource": [
        "arn:partition:quicksight:region:accountId:user/namespace/userName"
      ],
      "Condition": {
        "ForAllValues:StringEquals": {
          "quicksight:AllowedEmbeddingDomains": [
            "https://my.static.domain1.com",
            "https://*.my.static.domain2.com"
          ]
        }
      }
    }
  ]
}

```

```
    }
  }
]
}
```

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 QuickSight.

The following example shows a sample trust policy.

```
{
  "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 QuickSight Q search bar provides the classic QuickSight Q&A experience. QuickSight 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 [Embed the Amazon Q in QuickSight 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 QuickSight 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 QuickSight, 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 QuickSight. 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
                        "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: AWS 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 QuickSight from multiple locations.

The role session ID also becomes the user name in QuickSight. You can use this pattern to provision your users in QuickSight 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 QuickSight API operations, see the [QuickSight 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 QuickSight. For Microsoft AD users, you can use `DescribeUser` to get the user Amazon Resource Name (ARN).

The first time a user accesses QuickSight, 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 QuickSight, 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 QuickSight Q search bar provides the classic QuickSight Q&A experience. QuickSight 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 [Embed the Amazon Q in QuickSight 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 QuickSight 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 [QuickSight embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), QuickSight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, QuickSight 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 QuickSight 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 QuickSight Embedding SDK, the Q search bar on your page is dynamically resized based on the state. By using the QuickSight 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 QuickSight Embedding SDK to load the embedded dashboard on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon QuickSight embedding SDK](#) from GitHub. This repository is maintained by a group of QuickSight 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 QuickSight Q search bar embedding functionalities

Note

The embedded QuickSight Q search bar provides the classic QuickSight Q&A experience. QuickSight 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 [Embed the Amazon Q in QuickSight 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 [QuickSight embedding SDK](#).

Embedding the Amazon QuickSight Q search bar for anonymous (unregistered) users

Intended audience: Amazon QuickSight developers

Note

The embedded QuickSight Q search bar provides the classic QuickSight Q&A experience. QuickSight 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 [Embed the Amazon Q in QuickSight Generative Q&A experience](#).

In the following sections, you can find detailed information about how to set up an embedded Amazon QuickSight 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 QuickSight Q search bar embedding functionalities](#)

Step 1: Set up permissions

Note

The embedded QuickSight Q search bar provides the classic QuickSight Q&A experience. QuickSight 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 [Embed the Amazon Q in QuickSight 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 QuickSight 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 QuickSight** 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 QuickSight API Reference*.

The following sample policy provides these permissions.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "quicksight:GenerateEmbedUrlForAnonymousUser"
      ],
      "Resource": [
        "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:namespace/
        {{namespace}}",
        "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:dashboard/
        {{dashboardId-1}}",
        "arn:{{partition}}:quicksight:{{region}}:{{accountId}}:dashboard/
        {{dashboardId-2}}"
      ],
      "Condition": {
        "ForAllValues:StringEquals": {
          "quicksight:AllowedEmbeddingDomains": [
            "https://my.static.domain1.com",
            "https://*.my.static.domain2.com"
          ]
        }
      }
    }
  ]
}
```

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.

```
{
  "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 QuickSight Q search bar provides the classic QuickSight Q&A experience. QuickSight 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 [Embed the Amazon Q in QuickSight 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 QuickSight, 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
                "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': '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);
  });
```

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 = "U4zJMVZ2n2stZf1c80u3iKySEb3BEV6f"
            };
            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 QuickSight 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 QuickSight Q search bar provides the classic QuickSight Q&A experience. QuickSight 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 [Embed the Amazon Q in QuickSight 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 QuickSight 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 [QuickSight embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), QuickSight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, QuickSight 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 QuickSight 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 QuickSight Embedding SDK, the Q search bar on your page is dynamically resized based on the state. By using the QuickSight 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;
              }
            }
          },
        };
      };
    </script>
  </head>
</html>
```

```

        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
          }
        }
      }
    </script>
  </head>
</html>
```

```
        };
        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 QuickSight 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 QuickSight embedding SDK](#) from GitHub. This repository is maintained by a group of QuickSight 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 QuickSight Q search bar embedding functionalities

Note

The embedded QuickSight Q search bar provides the classic QuickSight Q&A experience. QuickSight 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 [Embed the Amazon Q in QuickSight 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 [QuickSight embedding SDK](#).

Embedding analytics using the `GetDashboardEmbedURL` and `GetSessionEmbedURL` API operations

Applies to: Enterprise Edition

Intended audience: Amazon QuickSight developers

The following API operations for embedding Amazon QuickSight dashboards and the QuickSight 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 overview](#)

- The [GetDashboardEmbedUrl](#) API operation embeds interactive dashboards.
- The [GetSessionEmbedUrl](#) API operation embeds the QuickSight console.

Topics

- [Embedding dashboards for everyone using GetDashboardEmbedURL \(old API\)](#)
- [Embedding dashboards for registered users using GetDashboardEmbedUrl \(old API\)](#)
- [Embedding the QuickSight console using GetSessionEmbedUrl \(old API\)](#)

Embedding dashboards for everyone using GetDashboardEmbedURL (old API)

Important

Amazon QuickSight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the QuickSight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding overview](#).

Applies to: Enterprise Edition

Intended audience: Amazon QuickSight developers

In the following sections, you can find detailed information on how to set up embedded Amazon QuickSight 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 QuickSight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the QuickSight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding overview](#).

Applies to: Enterprise Edition

Intended audience: Amazon QuickSight 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 QuickSight 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.

```
{
  "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.

```
{  
  "Version": "2012-10-17",  
  "Statement": {  
    "Effect": "Allow",  
    "Action": "sts:AssumeRole",  
    "Resource": "arn:aws-cn:iam::11112222333:role/QuickSightEmbeddingAnonymousPolicy"  
  }  
}
```

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 QuickSight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the QuickSight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding overview](#).

Applies to: Enterprise Edition

Intended audience: Amazon QuickSight 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 QuickSight, 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 \  
  --undo-redo-disabled true \  
  --reset-disabled true \  
  --user-arn arn:aws-cn:quicksight:us-east-1:111122223333:user/  
default/QuickSightEmbeddingAnonymousPolicy/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 QuickSight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the QuickSight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding overview](#).

Applies to: Enterprise Edition

Intended audience: Amazon QuickSight developers

In the following section, you can find out how you can use the [QuickSight 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 QuickSight [Embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), QuickSight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, QuickSight provides a responsive layout that is modified as your window size changes. By using the QuickSight 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 QuickSight Embedding SDK to load the embedded dashboard on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon QuickSight embedding SDK](#) from GitHub. This repository is maintained by a group of QuickSight 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 QuickSight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the QuickSight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding overview](#).

In the following sections, you can find detailed information on how to set up embedded Amazon QuickSight 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 QuickSight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the QuickSight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding overview](#).

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 QuickSight 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`.

```
{
  "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 QuickSight readers.

```
{
  "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.

```
{
  "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 QuickSight. The following example shows a role called `embedding_quicksight_dashboard_role`, which has the sample policy preceding as its resource.

```
{
  "Version": "2012-10-17",
  "Statement": {
    "Effect": "Allow",
    "Action": "sts:AssumeRole",
    "Resource": "arn:aws-
cn:iam::11112222333:role/embedding_quicksight_dashboard_role"
  }
}
```

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 QuickSight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the QuickSight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding overview](#).

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 QuickSight, 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 QuickSight. 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
                    key here
                    return new BasicAWSCredentials("access-
key", "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 ROLEARN
        final String roleArn, // IAM USER ROLE TO USE FOR EMBEDDING
        final String sessionName, // SESSION NAME FOR THE ROLEARN 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
  ) {
    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 QuickSight from multiple locations.

The role session ID also becomes the user name in QuickSight. You can use this pattern to provision your users in QuickSight 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 QuickSight API operations, see the [QuickSight 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 QuickSight. For Microsoft AD users, you can use `DescribeUser` to get the user ARN.

The first time a user accesses QuickSight, 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 QuickSight, 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 QuickSight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the QuickSight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding overview](#).

In the following section, you can find out how you can use the [Amazon QuickSight 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 [QuickSight embedding SDK](#) or by adding this URL into an iframe. If you set a fixed height and width number (in pixels), QuickSight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, QuickSight provides a responsive layout that is modified as your window size changes. By using the Amazon QuickSight 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 QuickSight Embedding SDK to load the embedded dashboard on your website using JavaScript. To get your copy, do one of the following:

- Download the [Amazon QuickSight embedding SDK](#) from GitHub. This repository is maintained by a group of QuickSight 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 QuickSight console using `GetSessionEmbedUrl` (old API)

Important

Amazon QuickSight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the QuickSight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding overview](#).

Applies to: Enterprise Edition

Intended audience: Amazon QuickSight developers

In the following sections, you can find detailed information on how to provide the Amazon QuickSight 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 QuickSight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the QuickSight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding overview](#).

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 QuickSight assumes a role that gives them Amazon QuickSight access and permissions to the console session. To make this possible, create an IAM role in your AWS 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 QuickSight 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`.

```
{
  "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.

```

{
  "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.

```

{
  "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 QuickSight. The following example shows a role called `embedding_quicksight_console_session_role`, which has the sample policy preceding as its resource.

```
{
  "Version": "2012-10-17",
  "Statement": {
    "Effect": "Allow",
    "Action": "sts:AssumeRole",
    "Resource": "arn:aws-
cn:iam::11112222333:role/embedding_quicksight_console_session_role"
  }
}
```

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 QuickSight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the QuickSight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding overview](#).

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 QuickSight, 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 QuickSight. 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");
                }

                @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 QuickSight, 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 QuickSight from multiple locations.

The role session ID also becomes the user name in QuickSight. You can use this pattern to provision your users in QuickSight 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 QuickSight API operations, see the [QuickSight 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 QuickSight. For Microsoft AD users, you can use `DescribeUser` to get the user ARN.

The first time a user accesses QuickSight, 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 \  
  --group-name TeamA1 \  
  --user-name john.doe@example.com
```

```
--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 QuickSight, and who has access to the QuickSight 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/embeddingsession
```

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 QuickSight has new APIs for embedding analytics: `GenerateEmbedUrlForAnonymousUser` and `GenerateEmbedUrlForRegisteredUser`. You can still use the `GetDashboardEmbedUrl` and `GetSessionEmbedUrl` APIs to embed dashboards and the QuickSight console, but they do not contain the latest embedding capabilities. For the latest up-to-date embedding experience, see [Embedding overview](#).

In the following section, you can find out how you can use the [Amazon QuickSight 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 QuickSight [Embedding SDK](#) or by adding this URL into an `iframe`. If you set a fixed height and width number (in pixels), QuickSight uses those and doesn't change your visual as your window resizes. If you set a relative percent height and width, QuickSight provides a responsive layout that is modified as your window size changes. By using the Amazon QuickSight 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 QuickSight 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 QuickSight embedding SDK](#) from GitHub. This repository is maintained by a group of QuickSight 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
```

Troubleshooting Amazon QuickSight

Use this information to help you diagnose and fix common issues that you can encounter when using Amazon QuickSight.

Note

Need more help? You can visit the Amazon QuickSight [User Community](#) or the [Amazon forums](#). See also the [Amazon QuickSight Resource Library](#).

Topics

- [Resolving Amazon QuickSight issues and error messages](#)
- [Connectivity issues when using Amazon Athena with Amazon QuickSight](#)
- [Data source connectivity issues for Amazon QuickSight](#)
- [Login issues with Amazon QuickSight](#)
- [Visual issues with Amazon QuickSight](#)

Resolving Amazon QuickSight 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 QuickSight

Following, you can find information about troubleshooting issues that you might encounter when using Amazon Athena with Amazon QuickSight.

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 QuickSight. For additional troubleshooting information, see [Troubleshooting](#) in the *Athena User Guide*.

Topics

- [Column not found when using Athena with Amazon QuickSight](#)
- [Invalid data when using Athena with Amazon QuickSight](#)
- [Query timeout when using Athena with Amazon QuickSight](#)
- [Staging bucket no longer exists when using Athena with Amazon QuickSight](#)
- [Table incompatible when using Amazon Glue with Athena in Amazon QuickSight](#)
- [Table not found when using Athena with Amazon QuickSight](#)
- [Workgroup and output errors when using Athena with Amazon QuickSight](#)

Column not found when using Athena with Amazon QuickSight

You can receive a "column not found" error if the columns in an analysis are missing from the Athena data source.

In Amazon QuickSight, 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 QuickSight 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 QuickSight

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 QuickSight

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 QuickSight 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 QuickSight

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 QuickSight 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 QuickSight

If you are getting errors when using Amazon Glue tables in Athena with Amazon QuickSight, 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 QuickSight 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_datebase_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.q1.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.q1.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.q1.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 QuickSight

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 QuickSight 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 QuickSight 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 Amazon QuickSight

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 QuickSight permissions to access the S3 output location, the Amazon QuickSight 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 QuickSight

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 QuickSight 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 QuickSight to Athena, check the following settings:

- Amazon resource permissions inside of Amazon QuickSight
- 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 QuickSight](#).

Make sure that you authorized Amazon QuickSight to use Athena

Intended audience: Amazon QuickSight administrators

Use the following procedure to make sure that you successfully authorized Amazon QuickSight to use Athena. Permissions to Amazon resources apply to all Amazon QuickSight users.

To perform this action, you must be an Amazon QuickSight 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 QuickSight to access Athena

1. Choose your profile name (upper right). Choose **Manage QuickSight**, and then choose **Security & permissions**.
2. Under **QuickSight access to Amazon Web Services**, 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 QuickSight.

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. 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 QuickSight access to Amazon Web 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 Amazon QuickSight, see [IAM policy examples for Amazon QuickSight](#).

Make sure that the IAM user has read/write access to your S3 location

Intended audience: Amazon QuickSight administrators

To access Athena data from Amazon QuickSight, 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 QuickSight 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, chose **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 QuickSight to access the S3 account. It's not enough that you, the user, are authorized. Amazon QuickSight must be authorized separately.

To authorize Amazon QuickSight 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 QuickSight, choose your profile name (upper right). Choose **Manage QuickSight**, and then choose **Security & permissions**.
3. 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 QuickSight. 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 QuickSight 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, Amazon QuickSight 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 QuickSight and choose **Datasets**.
2. Choose **New dataset**.
3. On the **Create a dataset** page, scroll down to the **FROM EXISTING DATASOURCES** section, and then choose the Adobe Analytics data source that you want to update.
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 QuickSight home screen, choose **Manage data**.
2. Choose **New dataset**.
3. Scroll to **FROM 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 datasets from new database data sources](#).

You can also refer to the Amazon RDS documentation on troubleshooting connections, [Cannot Connect to Amazon RDS DB Instance](#).

Login issues with Amazon QuickSight

Use the following section to help you troubleshoot login and access issues with the Amazon QuickSight console.

Topics

- [Insufficient permissions when using Athena with Amazon QuickSight](#)
- [Amazon QuickSight isn't working in my browser](#)
- [How do I delete my Amazon QuickSight account?](#)
- [Individuals in my organization get an "External Login is Unauthorized" message when they try to access Amazon QuickSight](#)
- [My email sign-in stopped working](#)

Insufficient permissions when using Athena with Amazon QuickSight

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 QuickSight can access the Amazon S3 buckets used by Athena:
 - a. To do this, choose your profile name (upper right). Choose **Manage QuickSight**, and then choose **Security & permissions**.
 - b. 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 QuickSight.

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. 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 QuickSight access to Amazon Web 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 QuickSight 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 QuickSight Role ARN> --operations Decrypt
```

The Amazon Resource Name (ARN) for the Amazon QuickSight 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, QuickSight 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 QuickSight uses:

```
arn:aws-cn:iam::AWS-ACCOUNT-ID:role/service-role/aws-quicksight-service-role-v0
```

Amazon QuickSight isn't working in my browser

If you can't view Amazon QuickSight correctly in your Google Chrome browser, take the following steps to fix the problem.

To view Amazon QuickSight 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 QuickSight account?

In some cases, you might need to delete your Amazon QuickSight account even when you can't access Amazon QuickSight 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/): `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 QuickSight users, and Amazon QuickSight administrators. If you have further difficulty, contact support.

Individuals in my organization get an "External Login is Unauthorized" message when they try to access Amazon QuickSight

Intended audience: Amazon QuickSight administrators

When an individual in your organization is federating into Amazon QuickSight using **AssumeRoleWithWebIdentity**, QuickSight 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 QuickSight and get the following unexpected error message.

The external login used for federation is unauthorized for the QuickSight 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 QuickSight, the single sign-on (IAM Identity Center) setup might use the `CognitoIdentityCredentials` API operation to assume the QuickSight role. This method maps all users in the Amazon Cognito identity pool to a single QuickSight user and isn't supported by Amazon QuickSight.

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 QuickSight role trust policy is set up like the following example.


```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Federated": "cognito-identity.amazonaws.com"
      },
      "Action": "sts:AssumeRoleWithWebIdentity",
      "Condition": {
        "StringEquals": {
          "cognito-identity.amazonaws.com:aud": "us-west-2:cognito-pool-id"
        },
        "ForAnyValue:StringLike": {
          "cognito-identity.amazonaws.com:amr": "unauthenticated"
        }
      }
    }
  ]
}
```

This setup allows a temporary Amazon Cognito user to assume a role session mapped to a unique QuickSight user. Because unauthenticated identities are temporary, they aren't supported by QuickSight.

We recommend that you don't use this setup, which setup isn't supported by Amazon QuickSight. For Amazon QuickSight, 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 Amazon QuickSight 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 QuickSight user, the session name might correspond to the same QuickSight role-based user.

We recommend that you remap the QuickSight 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 QuickSight

Mapping multiple Amazon Cognito user pools in different Amazon Web Services accounts to one identity pool and QuickSight isn't supported by Amazon QuickSight.

How can I fix it?

You can use QuickSight 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 QuickSight 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-AROAAAAAAAAAAAAAAAAA: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 QuickSight 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 QuickSight 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 Amazon QuickSight

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 QuickSight](#)

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 QuickSight](#).

- If you are having trouble adding a visual to an analysis, try the following:
 - Check that you aren't trying to add more objects than the quota allows. Amazon QuickSight 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, verify that the network's firewall settings permit traffic from *.aws.amazon.com, amazonaws.com, https://mobileanalytics.*.amazonaws.com, and cloudfront.net.
 - Add exceptions to your ad blocker for *.aws.amazon.com, amazonaws.com, https://mobileanalytics.*.amazonaws.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 QuickSight 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 QuickSight

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 Amazon QuickSight. For example, the value 1.59964E+11, which is actually 159964032802, formats as 159964000000 in QuickSight. 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 QuickSight.

Administration for Amazon QuickSight

Use the following section to learn about Amazon QuickSight administrative tasks. This section contains information about controlling access, managing accounts, and choosing Amazon Web Services Regions.

Topics

- [Different editions of Amazon QuickSight](#)
- [Amazon Web Services Regions, websites, IP address ranges, and endpoints](#)
- [Supported browsers](#)
- [Managing Amazon QuickSight](#)
- [Supporting multitenancy with isolated namespaces](#)
- [Customizing the QuickSight console](#)
- [Tracking Amazon account cost and usage data with Billing and Cost Management and Amazon QuickSight](#)

Different editions of Amazon QuickSight

Amazon QuickSight offers Standard and Enterprise editions. To learn more about the differences in availability, user management, permissions, and security between the two versions, see the following topic.

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 Amazon QuickSight users and administrators.

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

Availability of editions

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

The capacity region in which you start your Amazon QuickSight 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.

You can start a new Amazon QuickSight subscription using Standard edition, choosing any default capacity region. You can then upgrade it to Enterprise edition at any time.

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 QuickSight.

User management between editions

User management is different between the Amazon QuickSight 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).

User management for standard edition

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

When you create a QuickSight 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 QuickSight. You perform all management of users by adding, changing, and deleting accounts in Amazon QuickSight.

User management for enterprise edition

In Enterprise edition, 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 QuickSight as 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 QuickSight as users.

⚠ Important

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

With Active Directory, Amazon QuickSight Administrators and users aren't automatically notified of their access to Amazon QuickSight. 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 QuickSight. When you add a QuickSight 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 QuickSight.

To remove a user's access to QuickSight, 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 QuickSight.

Permissions for the different editions

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

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

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

- Any IAM user who is an Amazon QuickSight administrator
- The IAM user or Amazon root account that created the Amazon QuickSight account

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

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

The following table summarizes the admin actions that you can perform in QuickSight based on the access type that you choose.

Admin action	IAM permissions	QuickSight administrator (non-IAM)
Manage assets	Yes	
Security & permissions	Yes	
Manage VPC connections	Yes	
KMS keys	Yes	
Account settings	Yes	
Account customization		Yes
Manage users		Yes
Your subscriptions		Yes
Mobile settings		Yes
Domains and embedding		Yes
SPICE capacity		Yes

Amazon Web Services Regions, websites, IP address ranges, and endpoints

Amazon cloud-computing resources are housed in highly available facilities in different areas of the world (for example, North America, Europe, and Asia). These facilities are each part of an Amazon Web Services Region. For more information about Amazon Web Services Regions and Availability Zones (AZs), see [Global infrastructure](#).

Amazon QuickSight is currently supported in the following Amazon Web Services Regions. The following list provides websites, IP address ranges, and endpoints for Amazon QuickSight in each Amazon Web Services Region.

The IP addresses listed following are the ranges where QuickSight traffic originates from when making outbound connections to databases. They are not the IP address ranges that you use to connect to the QuickSight website or service API. For more information about authorizing QuickSight, see [Authorizing connections to Amazon data stores](#).

- US East (Ohio) (us-east-2)
 - **Website for user access** – `https://us-east-2.quicksight.aws.amazon.com`
 - **Service API endpoint** – `quicksight.us-east-2.amazonaws.com`
 - **IP address range for data source connectivity** – `52.15.247.160/27`
- US East (N. Virginia) (us-east-1)
 - **Website for user access** – `https://us-east-1.quicksight.aws.amazon.com`
 - **Service API endpoint** – `quicksight.us-east-1.amazonaws.com`
 - **IP address range for data source connectivity** – `52.23.63.224/27`
- US West (Oregon) (us-west-2)
 - **Website for user access** – `https://us-west-2.quicksight.aws.amazon.com`
 - **Service API endpoint** – `quicksight.us-west-2.amazonaws.com`
 - **IP address range for data source connectivity** – `54.70.204.128/27`
- Asia Pacific (Mumbai) (ap-south-1)
 - **Website for user access** – `https://ap-south-1.quicksight.aws.amazon.com`
 - **Service API endpoint** – `quicksight.ap-south-1.amazonaws.com`
 - **IP address range for data source connectivity** – `52.66.193.64/27`
- Asia Pacific (Seoul) (ap-northeast-2)
 - **Website for user access** – `https://ap-northeast-2.quicksight.aws.amazon.com`
 - **Service API endpoint** – `quicksight.ap-northeast-2.amazonaws.com`
 - **IP address range for data source connectivity** – `13.124.145.32/27`
- Asia Pacific (Singapore) (ap-southeast-1)
 - **Website for user access** – `https://ap-southeast-1.quicksight.aws.amazon.com`
 - **Service API endpoint** – `quicksight.ap-southeast-1.amazonaws.com`
 - **IP address range for data source connectivity** – `13.229.254.0/27`

- Asia Pacific (Sydney) (ap-southeast-2)
 - **Website for user access** – <https://ap-southeast-2.quicksight.aws.amazon.com>
 - **Service API endpoint** – quicksight.ap-southeast-2.amazonaws.com
 - **IP address range for data source connectivity** – 54.153.249.96/27
- Asia Pacific (Tokyo) (ap-northeast-1)
 - **Website for user access** – <https://ap-northeast-1.quicksight.aws.amazon.com>
 - **Service API endpoint** – quicksight.ap-northeast-1.amazonaws.com
 - **IP address range for data source connectivity** – 13.113.244.32/27
- Canada (Central) (ca-central-1)
 - **Website for user access** – <https://ca-central-1.quicksight.aws.amazon.com>
 - **Service API endpoint** – quicksight.ca-central-1.amazonaws.com
 - **IP address range for data source connectivity** – 15.223.73.0/27
- China (Beijing) (cn-north-1)
 - **Website for user access** – <https://cn-north-1.quicksight.amazonaws.cn>
 - **Service API endpoint** quicksight.cn-north-1.amazonaws.com.cn
 - **IP address range for data source connectivity** – 71.136.65.64/27
- Europe (Frankfurt) (eu-central-1)
 - **Website for user access** – <https://eu-central-1.quicksight.aws.amazon.com>
 - **Service API endpoint** – quicksight.eu-central-1.amazonaws.com
 - **IP address range for data source connectivity** – 35.158.127.192/27
- Europe (Ireland) (eu-west-1)
 - **Website for user access** – <https://eu-west-1.quicksight.aws.amazon.com>
 - **Service API endpoint** – quicksight.eu-west-1.amazonaws.com
 - **IP address range for data source connectivity** – 52.210.255.224/27
- Europe (London) (eu-west-2)
 - **Website for user access** – <https://eu-west-2.quicksight.aws.amazon.com>
 - **Service API endpoint** – quicksight.eu-west-2.amazonaws.com
 - **IP address range for data source connectivity** – 35.177.218.0/27
- ~~Europe (Paris) (eu-west-3)~~
 - **Website for user access** – <https://eu-west-3.quicksight.aws.amazon.com>

- **Service API endpoint** – `quicksight.eu-west-3.amazonaws.com`
- **IP address range for data source connectivity** – `13.38.202.0/27`
- Europe (Stockholm) (eu-north-1)
 - **Website for user access** – `https://eu-north-1.quicksight.aws.amazon.com`
 - **Service API endpoint** – `quicksight.eu-north-1.amazonaws.com`
 - **IP address range for data source connectivity** – `13.53.191.64/27`
- South America (São Paulo) (sa-east-1)
 - **Website for user access** – `https://sa-east-1.quicksight.aws.amazon.com`
 - **Service API endpoint** – `quicksight.sa-east-1.amazonaws.com`
 - **IP address range for data source connectivity** – `18.230.46.192/27`
- Amazon GovCloud (US-West) (gov-west-1)
 - **Website for user access** – `quicksight.us-gov-west-1.amazonaws.com`
 - **Service API endpoint** – `quicksight.us-gov-west-1.amazonaws.com`
 - **IP address range for data source connectivity** – `160.1.180.32/27`

Supported browsers

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

Note

Amazon QuickSight ended support for Microsoft Internet Explorer 11 on July 31, 2022. We can no longer ensure that the features and web pages of Amazon QuickSight 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	10 or later	Open Safari. On the menu, choose Safari , and then choose About Safari . The

Browser	Version	Check your version
		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.

Managing Amazon QuickSight

If you are a QuickSight administrator, the account you use to sign in to QuickSight is in the ADMIN QuickSight 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 QuickSight.

QuickSight asset management

Use this section to manage all of the assets in your Amazon QuickSight account in one unified view.

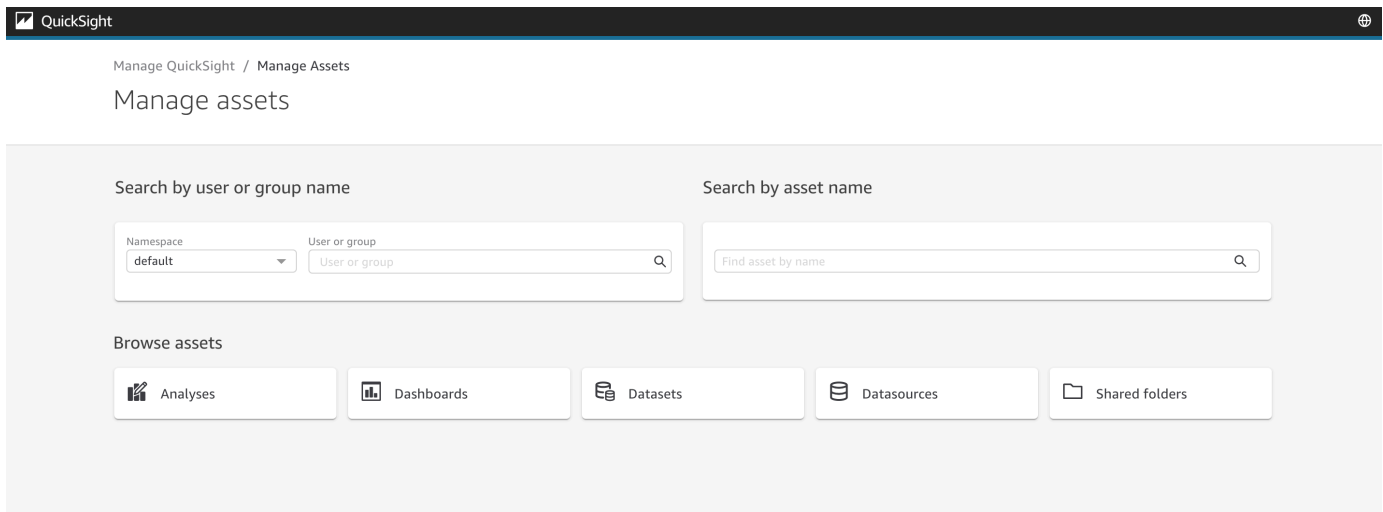
Here are some common reasons to use the asset manager:

- **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 QuickSight assets

1. Choose the profile icon, then **Manage QuickSight**.
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.

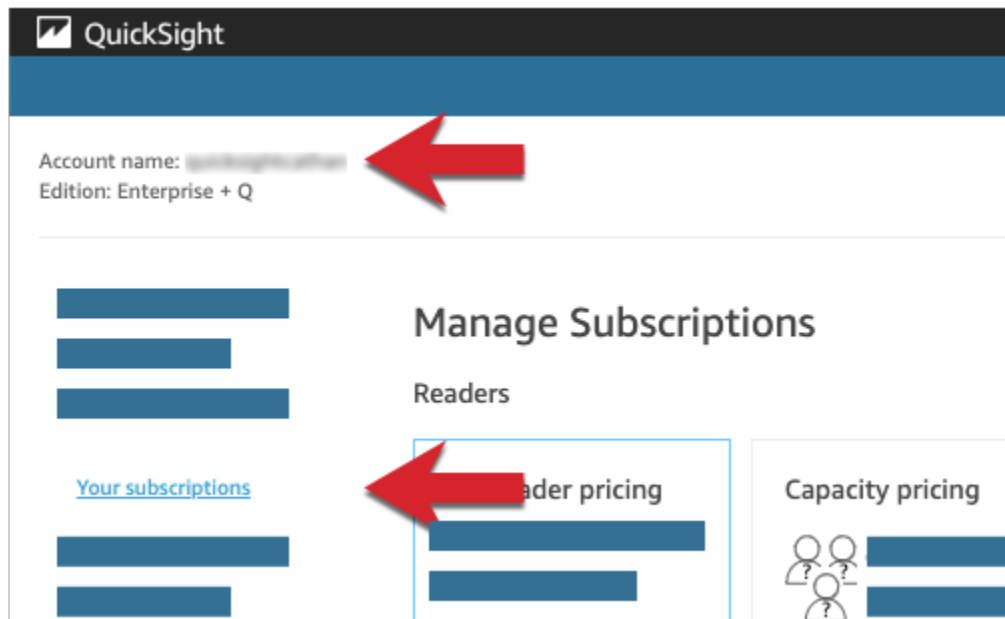
- Change the type you are browsing for by selecting a different asset type from the **Asset type** list.
- Use the vertical dot menu at right to perform an action on the asset in that row.
- Use the **Share** button to share all of the selected assets. A popup window displays sharing options to apply 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.

Amazon QuickSight subscriptions

Use this section to change your Amazon QuickSight subscriptions.

To open your account subscriptions

1. Choose the profile icon, and then select **Your subscriptions**.
2. Click on **Your subscriptions**.
3. Subscription details display on the screen.



Use the procedures in the following sections to manage subscription settings.

Manage Subscriptions

Per reader pricing

Capacity pricing

Get monthly subscription Get annual subscription

QuickSight Q add-on

Manage regions

Authors

Add more authors

Topics

- [Managing subscriptions](#)
- [Upgrading your Amazon QuickSight subscription from Standard edition to Enterprise edition](#)

Managing subscriptions

You can purchase standard user subscriptions to get discounted pricing on Amazon QuickSight. When you invite additional users to Amazon QuickSight, 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.

When you purchase an annual subscription, you pay for a QuickSight 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. For more information about pricing, see [Amazon QuickSight](#).

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.

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 QuickSight**.



2. Choose **Your Subscriptions**.
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.

Purchasing subscriptions

Use the following procedure to purchase subscriptions.

To purchase subscriptions

1. Choose your user name on the application bar and then choose **Manage QuickSight**.




2. Choose **Your Subscriptions**.
3. Choose **Purchase subscription**.
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 subscription**.

Editing subscriptions

Use the following procedure to edit subscriptions.

To edit subscriptions

1. Choose your user name on the application bar and then choose **Manage QuickSight**.



2. Choose **Your Subscriptions**.
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.

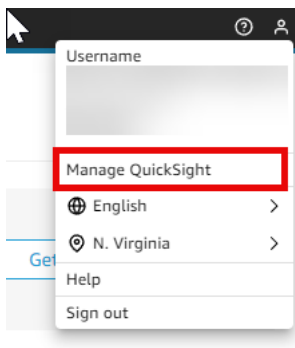
5. Choose whether the subscriptions autorenew.
6. Choose **Save changes**.

Deleting subscriptions

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 QuickSight**.



2. Choose **Your Subscriptions**.
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 QuickSight, 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 QuickSight.

Upgrading your Amazon QuickSight subscription from Standard edition to Enterprise edition

You can upgrade from Amazon QuickSight Standard edition to Amazon QuickSight Enterprise edition. In Enterprise edition, Amazon QuickSight 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 QuickSight 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 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 QuickSight editions](#).

When you upgrade your account, your administrators and authors are billed at the Amazon QuickSight Enterprise edition rates. For up-to-date information on rates, see [Pricing](#). 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 QuickSight, 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 QuickSight 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 QuickSight, with security permissions to subscribe. The person performing the upgrade is usually an Amazon administrator who is also an Amazon QuickSight 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**.

The following screen appears. For the latest prices, see [Amazon QuickSight pricing](#).



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 QuickSight 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 QuickSight.

When you upgrade to Enterprise edition, your admin and author users retain their roles.

Managing SPICE memory capacity

SPICE (Super-fast, Parallel, In-memory Calculation Engine) is the robust in-memory engine that QuickSight 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 Amazon QuickSight](#).

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 QuickSight in a single Amazon Web Services Region. The other Amazon Web Services Regions have no SPICE capacity unless you choose to purchase some.

QuickSight 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 QuickSight subscription and can be used in multiple Amazon Web Services Regions. For information about additional SPICE pricing, see [Amazon QuickSight pricing](#).

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 QuickSight using the `http://quicksight.aws.amazon.com` URL, QuickSight automatically selects your account and Amazon Web Services Region. You can't view your Amazon Web Services account from QuickSight. We recommend using a different method to open QuickSight when you want to work with SPICE capacity.
- If you open QuickSight from the Amazon Web Services Management Console, QuickSight 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 QuickSight. The Amazon Web Services Management Console and the QuickSight 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 QuickSight.

- If you use the Amazon Command Line Interface (Amazon CLI) to run QuickSight commands, make sure to provide the relevant Amazon Web Services account for each QuickSight 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 QuickSight administrator to view or manage SPICE capacity.

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 QuickSight by first entering "**quicksight**" into the **Find Services** search box. When the word QuickSight appears following the search box, choose it from the list.
3. In QuickSight, open the profile menu by choosing your profile icon at top right. The Amazon name of the Amazon Web Services Region that QuickSight 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 QuickSight 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 QuickSight. 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 Web Services Region](#).
2. Open the administration page by choosing **Manage QuickSight** 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 QuickSight:
 - Purchased SPICE capacity – This is the additional SPICE capacity.
 - SPICE capacity bundled with QuickSight – 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 QuickSight:
 - 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.

Purchasing SPICE capacity in an Amazon Web Services Region

To purchase more SPICE capacity (console)

1. Open QuickSight. 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 Web Services Region](#).
2. Open the administration page by choosing **Manage QuickSight** 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 QuickSight.
6. To confirm your choice, choose **Purchase SPICE capacity**. To exit without making any changes, choose **Cancel**.

Turning SPICE auto capacity purchasing on

Turn on SPICE auto capacity purchasing to allow Amazon QuickSight to automatically manage your QuickSight account's SPICE capacity. When you turn auto capacity purchasing on, QuickSight evaluates how much capacity is needed based on your account's usage. As your account uses more SPICE storage, 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. QuickSight 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 QuickSight accounts that are created in the console have auto capacity purchasing turned on in the region that their capacity is located. To turn on auto capacity purchasing for other regions, QuickSight account admins can manually turn on auto capacity from the **SPICE capacity** management page.

By default, all new QuickSight accounts that were created with the QuickSight API and all existing QuickSight accounts have auto capacity purchasing turned off. To turn on auto capacity

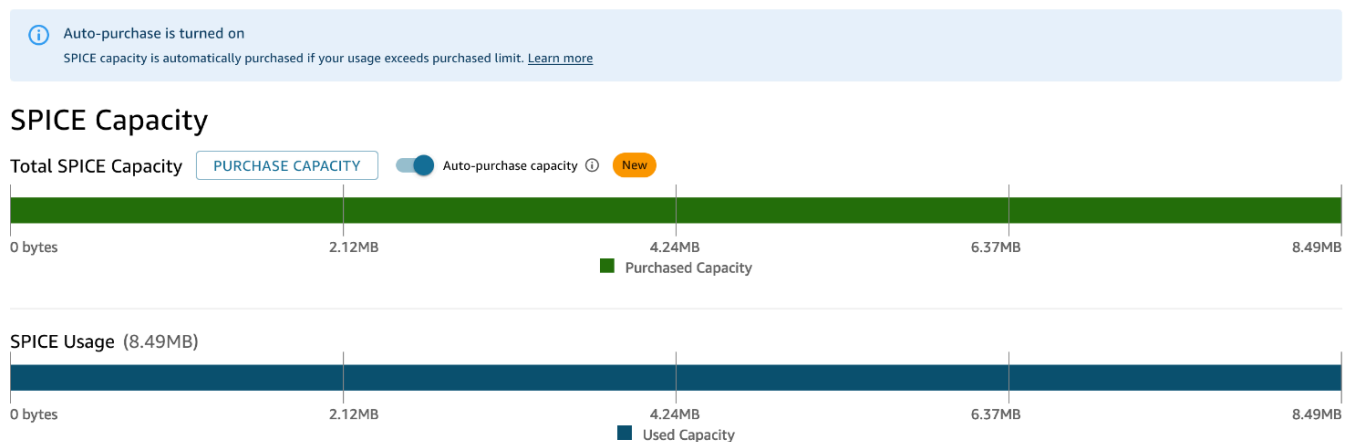
purchasing, QuickSight account admins can manually turn on auto capacity from the **SPICE capacity** management page.

To turn SPICE capacity purchasing on or off

1. From any page in the QuickSight console, choose your profile name, and then choose **Manage QuickSight**.

The **Manage QuickSight** menu is only available to QuickSight admins. If you are unable to access this menu, contact your QuickSight account admin for assistance.

2. Choose **SPICE capacity**.
3. On the **SPICE Capacity** page that opens, toggle the **Auto-purchase capacity** on, as shown in the image below.



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.

QuickSight 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, QuickSight automatically matches your account's capacity to your current usage. After QuickSight 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 QuickSight. 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 Web Services Region](#).
2. Open the administration page by choosing **Manage QuickSight** 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 QuickSight, choose **Release all**.
 - To release some of gigabytes of SPICE capacity from the Amazon Web Services Region that is currently selected in QuickSight, enter the number of gigabytes to release.
6. To confirm your choice, choose **Release SPICE capacity**. To exit without making any changes, choose **Cancel**.

Account settings

Use this section to change the account-wide settings in Amazon QuickSight.

To open your account settings

1. Choose the profile icon, and then select **Manage QuickSight**.
2. Click on **Account settings**.

Topics

- [Changing your notification email](#)
- [Deleting your Amazon QuickSight subscription and closing the account](#)

Changing your notification email

You can change the notification email address for access requests and service notifications.

Use the following procedure to change your Amazon QuickSight notification email and to enable or disable IAM user access requests.

To change your notification email and enable or disable IAM user access requests

1. Choose the profile icon, and then select **Manage QuickSight**.
2. Click on **Account settings**.
3. Under **Notification email address**, enter the email address you want to use.

Choose whether to send IAM user access requests to the same email address. Toggle **Enable IAM user access requests to this account** for this setting.

Deleting your Amazon QuickSight subscription and closing the account

The act of deleting Amazon QuickSight from your is immediate and final. Deletion removes every QuickSight 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 QuickSight account from the **Manage QuickSight** menu or by using the API. To prevent someone from deleting a QuickSight user account accidentally or maliciously, QuickSight 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 QuickSight 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 QuickSight 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 QuickSight account.
- You understand that your Amazon Web Services account is not deleted when you terminate your Amazon QuickSight 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).

To terminate your QuickSight account without the QuickSight UI

1. Sign in to Amazon where you want to remove Amazon QuickSight.
2. [Use this direct link to open the Amazon QuickSight **Account termination** screen.](#)

This approach works no matter which Amazon Web Services Regions you use.

To terminate your account by using the QuickSight UI

1. Choose your profile on the application bar, and then choose **Manage QuickSight**.
2. Use one of the following methods to open the **Account termination** screen.
 - Use this [direct link](#) to the screen.
 - Choose **Account settings, Manage**.



Account name: quicksightcathan
Edition: Enterprise

Manage users

Manage groups

Manage assets

Your subscriptions

SPICE capacity

[Account settings](#)

Security & permissions

Manage VPC connections

Mobile settings

Domains and Embedding

Account customization

Single sign-on (SSO)

KMS keys

Account settings

Notification email address

This will be where access requests and service notifications will be sent.

Enable IAM user access requests to this account.

Account termination

Manage account termination protection or delete this account

Manage

3. On the **Account termination** page, confirm that you are viewing the correct QuickSight account by checking the name listed for account name.

Account termination

QuickSight account name



quicksightcahan

Account termination protection ⓘ

Account termination protection is an extra safe-guard to help prevent accidental deletion of accounts. Turn account termination protection off to delete your account.



Account termination protection is on.



Delete account

Deleting this account can't be undone and will permanently delete all users, dashboards, analyses, along with other related data.

Type "confirm" to delete this account

confirm



Cancel

Delete account

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.

Permissions and access to account termination

You need the following special permissions to terminate a QuickSight account. Without these permissions, you won't be able to terminate a QuickSight user account. Contact your account administrator for help.

- You're a QuickSight administrator and have an Admin role in QuickSight.
- You need permissions to run the following (except if you're the root admin user (IAM) who added QuickSight)
 - `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.

Managing domains and embedding

Applies to: Enterprise Edition

Intended audience: Amazon QuickSight administrators

In Amazon QuickSight Enterprise edition, you can embed QuickSight 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 Amazon QuickSight subscription. This requirement protects your data by preventing unapproved domains from hosting embedded dashboards. To embed a QuickSight dashboard, visual, console, or Q search bar to a web page or app, add approved domains to a static allow list in the QuickSight console. Alternatively, add them at runtime with the QuickSight 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 QuickSight API](#)

Allow listing static domains

You can add static domains to your allow list through the QuickSight 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 QuickSight**. You must be an Amazon QuickSight 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 QuickSight 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 QuickSight** 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 QuickSight 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 [Working with embedded analytics](#).

Supporting multitenancy with isolated namespaces

Amazon QuickSight Enterprise edition supports multitenancy through namespaces. A QuickSight *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 QuickSight 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 [Amazon security in Amazon QuickSight](#).
- 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 QuickSight account with IAM Identity Center. For more information on integrating your QuickSight account with IAM Identity Center, see [Configure your Amazon QuickSight account with IAM Identity Center](#).

- 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 [Amazon QuickSight 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](#).

If you don't have an existing Amazon Web Services account or you need to sign up for QuickSight, read the following guidelines, then follow the applicable instructions in [Signing up for an Amazon QuickSight 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 QuickSight [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 QuickSight](#). 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 QuickSight user with administrator privileges. For more information, see [Upgrading your Amazon QuickSight subscription from Standard edition to Enterprise edition](#).

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 QuickSight 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

QuickSight assets and resources, including datasets, data sources, dashboards, analyses, and so on, exist outside of any namespace. They are visible only to users who have resource permissions granted to them.

To implement namespaces, you use the following QuickSight API operations:

- [CreateNamespace](#)
- [DescribeNamespace](#)
- [ListNamespaces](#)
- [DeleteNamespace](#)

📘 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:

- QuickSight 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`
- QuickSight 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 QuickSight 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 QuickSight administrators

To migrate existing users in one namespace to a different namespace

1. Identify the users that you want to transfer to a different namespace by using the QuickSight user and group API operations. For more information, see [API operations for controlling access](#) in the [Amazon QuickSight 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 QuickSight 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 QuickSight 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 QuickSight.

3. To identify the asset and resource permissions that the new namespace users need, use the QuickSight 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 [Amazon QuickSight 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 use `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 [Amazon QuickSight 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 QuickSight 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 [Amazon QuickSight 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.

Customizing the QuickSight console

Using Amazon QuickSight, you can create a customized experience for people using either the Amazon Web Services Management Console or QuickSight consoles embedded in your application.

Currently, different options for customizing QuickSight are available separately in the console and the QuickSight API. Following, you can find information about the available options.

The following customization options are currently available:

- You can customize the welcome content QuickSight 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 QuickSight home page.
- You can create and specify a default theme.
- You can create and set defaults for sharing dashboards via email with email templates.

⚠ Important

All customizations apply only to the Amazon Web Services Region that you are using in the API or that is selected in the QuickSight console.

To check your Region setting, you can use one of the following procedures.

To check your Amazon Web Services Region on the QuickSight 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

- [Customizing QuickSight welcome content](#)
- [Customizing email report templates](#)
- [Setting a default theme for Amazon QuickSight analyses](#)

Customizing QuickSight welcome content**To customize QuickSight welcome content**

1. In Amazon QuickSight, choose your profile icon at upper right to open the menu.

2. Choose **Manage QuickSight** to open the administration page.
3. On the navigation pane, choose **Account customization** to open the customization options.
4. Select the **Show introductory videos** check box to show the default tutorial videos and the introductory animation. Clear the check box to hide QuickSight videos and the intro animation for all users in your current region.
5. Select the **Create sample datasets and analyses** check box to accept sample datasets and analyses for new users. Doing this also applies to existing users who open QuickSight in a new Amazon Web Services Region. Clear the check box if you want to decline sample datasets and analyses. You can also provide your own versions of these to your users.
6. Choose **Update**.

Any changes to customizations take about 10 minutes to appear. They apply only to your current Amazon Web Services Region.

Customizing email report templates

Intended audience: System administrators and Amazon QuickSight administrators

In Amazon QuickSight, you can customize how dashboard email reports appear and behave for account users. You can customize the sender display name (who the email says it's from), the logo, and the footer that displays in the email. You can also customize where the dashboard opens when recipients click on it in the email report. You can even include a user-friendly name in place of your custom email address, such as Sales, for the email to be from.

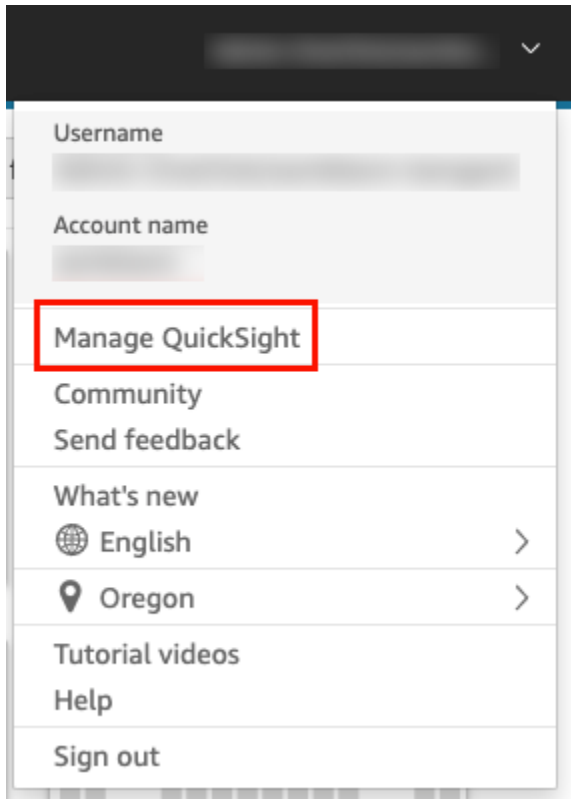
Before you can create a custom email report template, make sure you have the appropriate IAM permissions to create and update account customizations, including creating and updating custom email templates. If you plan to use a custom email address to send reports from, such as a company email address, make sure you also have permissions to obtain SES identity attributes. For more information about the permissions you need, and to see an example IAM policy, see [IAM identity-based policies for Amazon QuickSight: customizing email report templates](#).

Important

All customizations apply only to the Amazon Web Services Region and account that's selected in the QuickSight console.

To create a custom email report template

1. In QuickSight, choose your user name at upper right, and then choose **Manage QuickSight**.



2. In the toolbar at left, choose **Account customization**.
3. On the **Account customization** page that opens, under **Email report template**, choose **Update**.

Manage users

Your subscriptions

SPICE capacity

Account settings

Security & permissions

Manage VPC connections

Mobile settings

Domains and Embedding

[Account customization](#)

Single sign-on (SSO)

Account customization

Email report template

Customize email reports for sharing dashboards via email

[Update](#)

Welcome content

Show introductory videos

Create sample datasets and analyses

Customization changes take about 10 minutes and apply o

[Update](#)

The **Customize email template** page opens with the following options for customizing dashboard reports.

- Customizing the "Sent from" display name
- Customizing the logo to display in the email
- Customizing where the linked dashboard opens
- Customizing text in the email footer

The following sections describe each of these options. To create a custom email template that fits your business needs, follow the steps in each section. When you're finished, choose **Save**.

Customizing the "sent from" display name

You can customize who email reports are from by selecting a "Sent from" display name. This name displays in the "Sent from" line in email reports. You can choose to send emails from and display the QuickSight email address (the default) or a custom email address that you've verified within Amazon Simple Email Service under the same Amazon account.

Before you can send email reports from a custom email address, make sure that you set up a verified SES identity for the email address. An SES *identity* is an email address or domain that your

email is sent from. For more information, see [Verifying an email address](#) in the *Amazon Simple Email Service Developer Guide*.

To select a "sent from" display name

- For **Select "Sent from" display name configuration**, choose one of the following options:
 - **Custom email address within the Simple Email Service (SES) Amazon account #111222333**. – When you choose this option, you can enter an email address to display in the "Sent from" line in email reports.

If you choose this option, QuickSight sends email reports for the email address that you provide.

To use a custom email address, you first confirm that the email address is a verified SES identity. Then you create a custom policy for that identity using the provided authorization policy code in SES, and then verify the authorization policy in QuickSight. You can also provide a user friendly display name (optional) for the email. For more information, see following.

- **QuickSight: no-reply@quicksight.aws** – When you select this option, the QuickSight no-reply email address displays in the "Sent from" line in email reports.

Configuring a custom email address for email reports

Use the following procedure to set up a custom email address within your SES account to use for QuickSight email reports.

To set up a custom email address in your SES account for QuickSight email reports

1. For **STEP 1**, enter a verified SES email address in the text box, and then choose **Verify**.

If the email address has been verified in your SES account, a verification message appears. If it hasn't, go to your SES account and verify the email address. For more information, see [Verifying an email address](#) in the *Amazon Simple Email Service Developer Guide*.

2. For **STEP 2**, choose **Copy authorization policy**, and then do the following:
 - a. Choose **Go to SES**.

- b. Sign in to your SES account and create a custom policy for the email address that you verified in the previous step. You can paste the authorization policy code that you copied from QuickSight in the SES policy editor.

For more information about creating identity policies in SES, see [Creating a custom policy](#) in the *Amazon Simple Email Service Developer Guide*.

3. For **STEP 3**, choose **Verify Authorization** to verify that the SES identity has authorized QuickSight to send emails on its behalf.

If it's verified, a verification message appears.

4. (Optional) For **STEP 4**, enter a user-friendly name to display in the "Sent from" line in email reports, and then choose **Save**.

Customizing the logo in the email

You can choose to display the QuickSight logo in email reports or a custom logo, such as your company logo. You can also choose not to display a logo.

To choose the type of logo to display in the email

- For **Select logo type**, choose one of the following options:
 - **Custom logo** – When you choose this option, you can upload an image to display a custom logo in your email reports. You can customize the logo header background color as well.

The image you upload must be a JPG, JPEG, or PNG file, and can have a maximum size of 1 MB. When uploaded, the logo automatically resizes to a max height of 32 pixels.

- To upload an image for the custom logo, choose **Upload an image** and choose an image file.

Make sure that the image that you upload is a .jpg, .jpeg, or .png file. The file can have a maximum size of 1 MB. When uploaded, the logo automatically resizes to a maximum height of 32 pixels.

- To change the header background color, choose the color picker icon, or enter a hexadecimal number in the space provided.
- **QuickSight logo** – This is the default logo that displays in email reports unless specified otherwise.

- **No logo** – When you choose this option, no logo is displayed in the email.

Customizing where the dashboard opens from the email

Email reports link to dashboards that recipients can view in QuickSight or your application. You can choose where the dashboard opens, or you can hide the dashboard link in the email.

To select where the dashboard opens

- For **Select where the dashboard opens**, choose one of the following options:
 - **Open in custom application** – When you choose this option, users are redirected to your application when they click on the link to the dashboard in the email report.
 - To open the dashboard in your application, enter the URL for your application. You can use parameters in the URL. Any parameters that you add are replaced at runtime with the appropriate information. The following parameters are supported: <<\$accountId>>, <<\$dashboardId>>, and <<\$awsRegion>>.

For example, let's say that you enter the following URL with parameters:

```
https://www.example.com/analytics?account-id=<<$accountId>>&dashboard-id=<<$dashboardId>>&region=<<$awsRegion>>.
```

When the email report is sent to subscribers, QuickSight replaces the parameters with the appropriate values at runtime. The URL in the dashboard report email might be similar to the following:

```
https://www.example.com/analytics?account-id=111222333&dashboard-id=28ab58b4-8b53-441c-b52b-bc475f620d7f&region=us-west-2.
```

- To enter a custom call to action for the dashboard link in the email, enter text for **Enter custom call to action text**.
- **Open in quicksight.aws.com** – When you choose this option, users are redirected to QuickSight when they click on the link to the dashboard in the email report.
- **Hide dashboard link in email** – When you choose this option, a link to view the dashboard isn't shown.

Customizing the email footer

To customize the text in the email footer

- For **Select footer type**, choose one of the following options:
 - **Custom footer** – When you choose this option, you can enter a custom footer of up to 500 characters.
 - **QuickSight footer** – When you choose this option, the following default QuickSight footer is used.

Sent by <dashboardowner@email.com> from Amazon account <accountname>

A link to unsubscribe is also included.

- **No footer** – When you choose this option, no footer message appears in the email.

Amazon CloudTrail logs

When you or someone in your account sets up an email template, the following snippet is added to the CloudTrail log as part of the eventName DescribeAccountCustomization and DescribeEmailCustomizationTemplate, and the eventCategory Management.

```
DescribeAccountCustomization
{
  "eventSource": "quicksight.amazonaws.com",
  "eventName": "DescribeAccountCustomization",
  "requestParameters": {
    "awsAccountId": "111222333",
    "resolved": false
  },
  "responseElements": null,
  "eventCategory": "Management"
}

DescribeEmailCustomizationTemplate
{
  "eventSource": "quicksight.amazonaws.com",
  "eventName": "DescribeEmailCustomizationTemplate",
  "requestParameters": {
    "awsAccountId": "111222333",
    "emailCustomizationTemplateId": "TemplateId"
  }
}
```

```

},
"responseElements": null,
"eventCategory": "Management"
}

```

When the template is saved, the following snippets are added as part of the eventName for CreateAccountCustomization and CreateEmailCustomizationTemplate.

```

CreateAccountCustomization
{
  "eventSource": "quicksight.amazonaws.com",
  "eventName": "CreateAccountCustomization",
  "requestParameters": {
    "accountCustomization": {
      "defaultEmailCustomizationTemplate": "arn:aws-cn:quicksight:us-
west-2:111222333:email-customization-template/template-id"
    },
    "awsAccountId": "111222333"
  },
  "responseElements": {
    "status": 201,
    "arn": "arn:aws-cn:quicksight:us-west-2:111222333:customization/account/111222333",
    "awsAccountId": "111222333",
    "accountCustomization": {
      "defaultEmailCustomizationTemplate": "arn:aws-cn:quicksight:us-
west-2:111222333:email-customization-template/template-id"
    },
    "requestId": "6b6f2ce8-584b-47cb-9f56-4273ab7061a6"
  },
  "eventCategory": "Management"
}

```

```

CreateEmailCustomizationTemplate
{
  "eventSource": "quicksight.amazonaws.com",
  "eventName": "CreateEmailCustomizationTemplate",
  "requestParameters": {
    "fromEmailAddressCurrentOption": "DEFAULT",
    "description": "",
    "awsAccountId": "111222333",
    "emailCustomizationTemplateId": "template-id",
    "name": "Email Customization Template",
    "dashboardLinkCurrentOption": "DEFAULT",

```

```

    "footerCurrentOption": "DEFAULT",
    "logoCurrentOption": "DEFAULT"
  },
  "responseElements": {
    "emailCustomizationTemplateId": "template-id",
    "status": 200,
    "requestId": "17dea6c9-7811-4ee2-9c79-00c4d376a2c2",
    "arn": "arn:aws-cn:quicksight:us-west-2:111222333:email-customization-template/template-id"
  },
  "eventCategory": "Management"
}

```

When the template is saved, the following snippets are added as part of the eventName for UpdateAccountCustomization and UpdateEmailCustomizationTemplate.

```

UpdateAccountCustomization
{
  "eventSource": "quicksight.amazonaws.com",
  "eventName": "UpdateAccountCustomization",
  "requestParameters": {
    "accountCustomization": {
      "defaultEmailCustomizationTemplate": "arn:aws-cn:quicksight:us-west-2:111222333:email-customization-template/template-id"
    },
    "awsAccountId": "111222333"
  },
  "responseElements": {
    "status": 200,
    "arn": "arn:aws-cn:quicksight:us-west-2:111222333:customization/account/111222333",
    "awsAccountId": "111222333",
    "accountCustomization": {
      "defaultEmailCustomizationTemplate": "arn:aws-cn:quicksight:us-west-2:111222333:email-customization-template/template-id"
    },
    "requestId": "6b6f2ce8-584b-47cb-9f56-4273ab7061a6"
  },
  "eventCategory": "Management"
}

```

```

UpdateEmailCustomizationTemplate
{
  "eventSource": "quicksight.amazonaws.com",

```

```
"eventName": "UpdateEmailCustomizationTemplate",
"requestParameters": {
  "fromEmailAddressCurrentOption": "DEFAULT",
  "description": "",
  "awsAccountId": "1112223333",
  "emailCustomizationTemplateId": "template-id",
  "name": "Email Customization Template",
  "dashboardLinkCurrentOption": "DEFAULT",
  "footerCurrentOption": "DEFAULT",
  "logoCurrentOption": "DEFAULT"
},
"responseElements": {
  "emailCustomizationTemplateId": "template-id",
  "status": 200,
  "requestId": "17dea6c9-7811-4ee2-9c79-00c4d376a2c2",
  "arn": "arn:aws-cn:quicksight:us-west-2:1112223333:email-customization-template/
template-id"
},
"eventCategory": "Management"
}
```

Setting a default theme for Amazon QuickSight analyses

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 QuickSight ARNs, see [ARN formats](#) in the *Amazon QuickSight 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

Tracking Amazon account cost and usage data with Billing and Cost Management and Amazon QuickSight

Applies to: Enterprise Edition

With Billing and Cost Management, you can visualize your Amazon account's billing and cost management data with a pre-built cost and usage dashboard that's powered by Amazon QuickSight. For more information about creating a cost and usage dashboard, see [Creating a cost and usage dashboard](#) in the *Amazon Billing User Guide*

Amazon security in Amazon QuickSight

Amazon QuickSight 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 Amazon QuickSight, 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 QuickSight. The following topics show you how to configure Amazon QuickSight 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 QuickSight resources.

Amazon QuickSight 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 QuickSight can also support FedRAMP, HIPAA, PCI DSS, ISO, and SOC compliance to help you meet industry-specific or regulatory requirements.

Data protection in Amazon QuickSight

The Amazon [shared responsibility model](#) applies to data protection in Amazon QuickSight. 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 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.
- Use Amazon encryption solutions, along with all default security controls within Amazon Web 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 140-2 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 \(FIPS\) 140-2](#).

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 QuickSight or other Amazon Web 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 QuickSight](#)
- [Inter-network traffic privacy in Amazon QuickSight](#)
- [Accessing data sources](#)

Data encryption in Amazon QuickSight

Amazon QuickSight uses the following data encryption features:

- Encryption at rest
- Encryption in transit
- Key management

You can find more detail on these topics in the following sections.

Topics

- [Encryption at rest](#)
- [Encryption in transit](#)
- [Key management](#)

Encryption at rest

Amazon QuickSight securely stores your Amazon QuickSight metadata. This includes the following:

- Amazon QuickSight user data, including Amazon QuickSight user names, email addresses, and passwords. Amazon QuickSight 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
- Names of your uploaded files, data source names, and data set names.
- Statistics that Amazon QuickSight uses to populate machine learning (ML) insights

Amazon QuickSight securely stores your Amazon QuickSight 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, Query result cache.

When you delete a user, all of that user's metadata is permanently deleted. If you don't transfer that user's Amazon QuickSight objects to another user, all of the deleted user's Amazon QuickSight objects (data sources, datasets, analyses, and so on) are also deleted. When you unsubscribe from Amazon QuickSight, all metadata and any data you have in SPICE is completely and permanently deleted.

Encryption in transit

Amazon QuickSight 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 QuickSight secures all encrypted transfers by using Secure Sockets Layer (SSL).

Key management

Using Amazon managed keys in QuickSight

All non-customer managed keys associated with Amazon QuickSight 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](#).

Using customer-managed keys from Amazon KMS with SPICE datasets in Amazon QuickSight

QuickSight enables you to encrypt your SPICE datasets using 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 datasets in QuickSight SPICE 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 QuickSight SPICE dataset. A QuickSight administrator can then use a CMK to encrypt SPICE datasets and control access.

The following rules apply to using CMKs with SPICE datasets:

- Amazon QuickSight 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.
- The key that is currently the default CMK is automatically used to encrypt new SPICE datasets.
- Some features always use QuickSight's default encryption instead of applying SPICE CMK settings:
 - Amazon S3 analytics dashboard
 - Augmenting data with Amazon SageMaker
 - Direct file uploads
 - Exporting data with the following methods:
 - Exporting visual data to a .csv, .xlsx, or .pdf file
 - Reporting data in a .csv, .xlsx, or .pdf file
 - ML-powered anomaly detection
 - QuickSight Q

Note

If you use Amazon Key Management Service with Amazon QuickSight, 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 QuickSight.

Add a CMK to your account

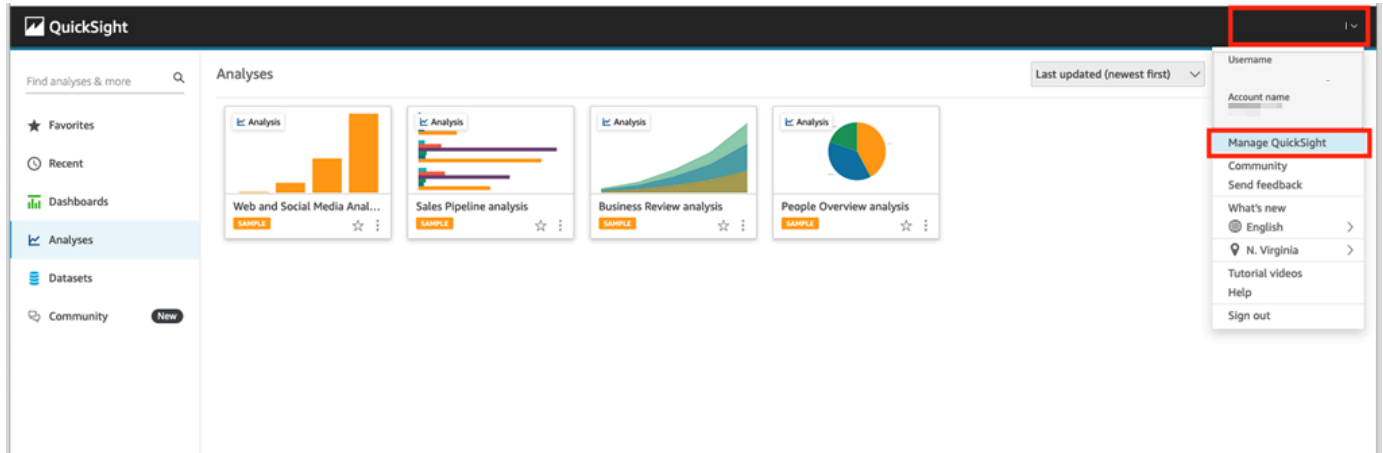
Before you begin, make sure that you have an IAM role that grants the admin user access to the Amazon QuickSight admin key management console. For more information on the required permissions, see [IAM identity-based policies for Amazon QuickSight: using the admin key management console](#).

You can add keys that already exist in Amazon KMS to your QuickSight account, so that you can encrypt your SPICE datasets. Keys that you add only affect new datasets created in SPICE. If you have an existing SPICE dataset that you want to encrypt, perform a full refresh on the dataset to encrypt it with the default CMK.


To learn more about how you can create a key to use in QuickSight, see the [Amazon Key Management Service Developer Guide](#).

To add a new CMK to your QuickSight account.

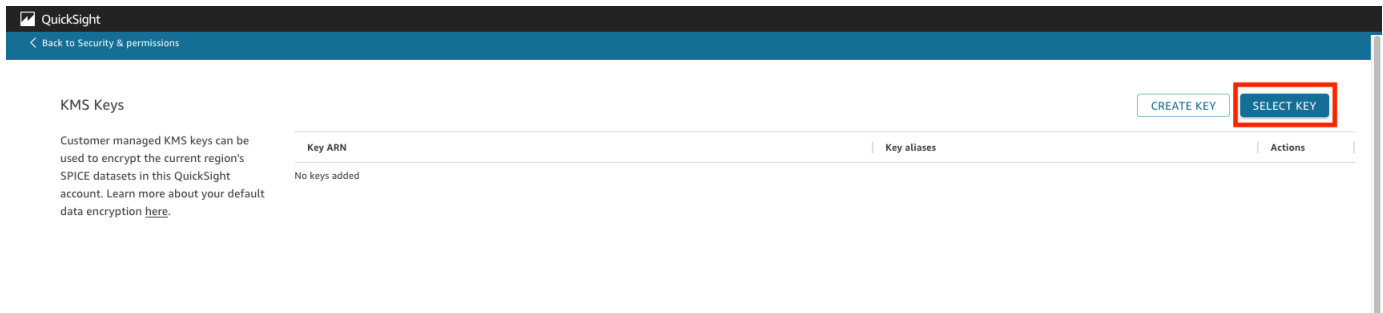
1. On the QuickSight start page, choose **Manage QuickSight**, and then choose **KMS keys**.



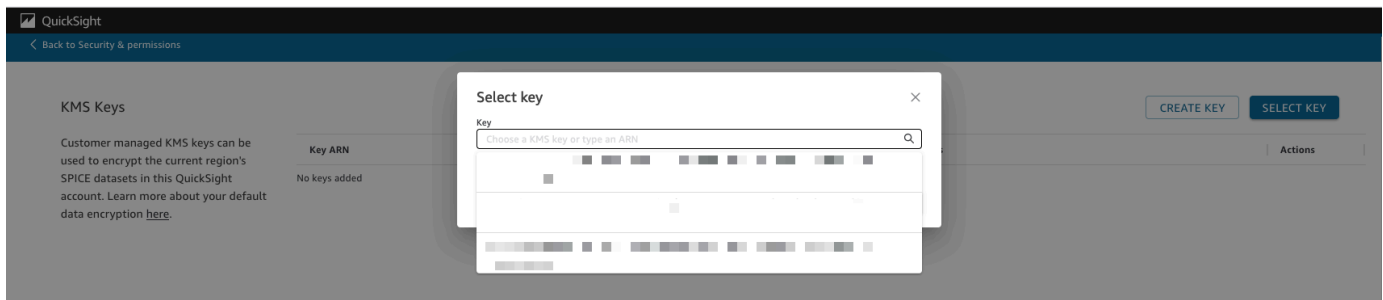
2. On the **KMS keys** page, choose **Manage**. The **KMS keys** dashboard opens.



3. On the **KMS Keys** dashboard, choose **Select key**.



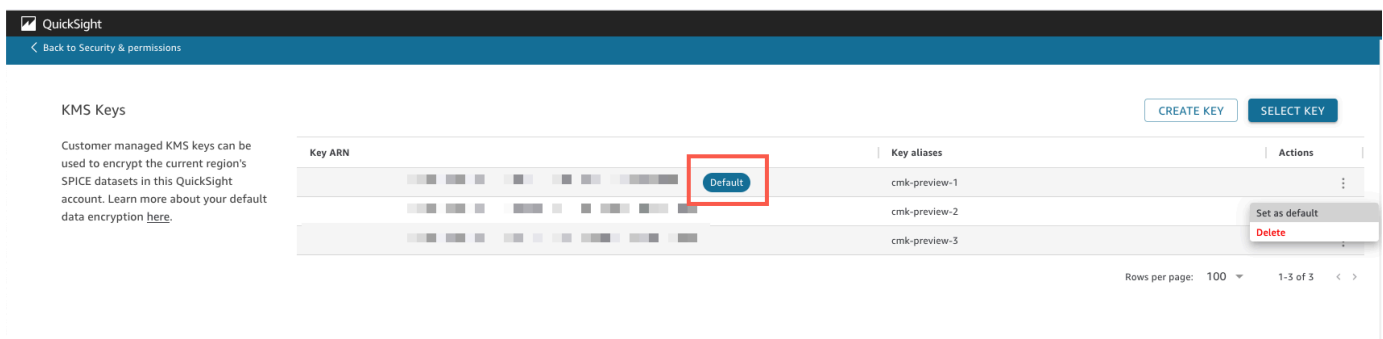
- On the **Select key** pop-up box, choose **Key** to open the list. Then, select the key that you want to add.



If your key isn't in the list, you can manually enter the key's ARN.

- (Optional) Select the **Use as default encryption key for all new SPICE datasets in this QuickSight account** to set the selected key as your default key. A blue badge appears next to the default key to indicate its status.

When you choose a default key, all new SPICE datasets that are created in the Region that hosts your QuickSight account are encrypted with the default key.



- (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.

Note

To use a specific key for an existing dataset, switch the account default key to the new key, then run a full refresh on the SPICE dataset.

Verify the key used by a SPICE dataset

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 a SPICE dataset is encrypted by, you can find this information in CloudTrail.

Verify the CMK that's currently used by a SPICE dataset

1. Navigate to your CloudTrail log. For more information, see [Logging operations with Amazon 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 QuickSight 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.

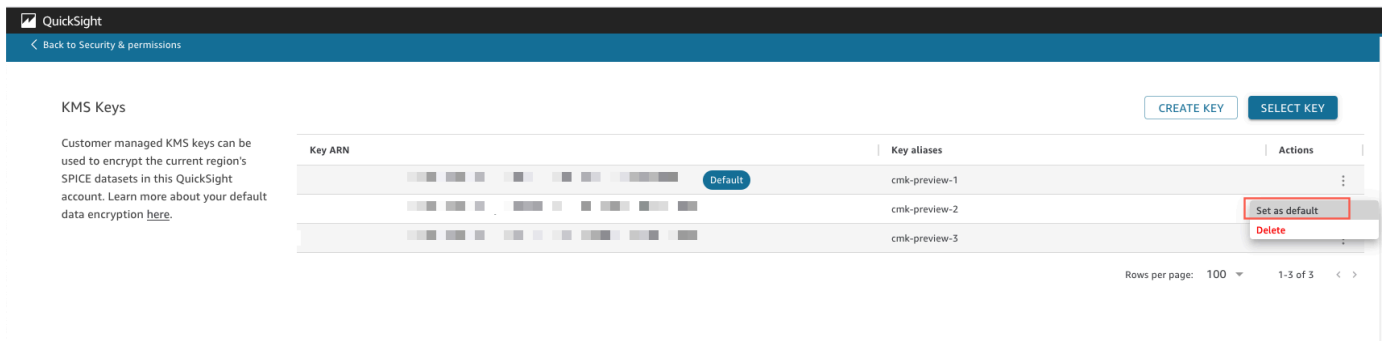
RetireGrant – If latest CloudTrail event of the SPICE dataset is RetireGrant, there is no key ID and the SPICE dataset is no longer CMK encrypted.

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 datasets created in SPICE are encrypted on the new key. The new default key changes how new SPICE datasets are encrypted. However, existing datasets continue to use the previous default key until the dataset is fully refreshed. To encrypt the dataset with a new default key, perform a full refresh on the dataset.

To change the default key to an existing key

1. On the QuickSight start page, choose **Manage QuickSight**, and then choose **KMS keys**.
2. On the **KMS keys** page, 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**.



The selected key is now your default key.

Removing CMK encryption on your QuickSight account

You can remove the default key to disable SPICE dataset encryption in your QuickSight account. Removing the key prevents new datasets from encrypting on a CMK.

To remove CMK encryption for new SPICE datasets

1. On the QuickSight start page, choose **Manage QuickSight**, 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, QuickSight stops encrypting new SPICE datasets. Any existing encrypted datasets stay encrypted until a full refresh occurs.

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-encrypted dataset

You can revoke access to your CMK-encrypted SPICE datasets. When you revoke access to a key that is used to encrypt a dataset, access to the dataset 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 QuickSight KMS policy in IAM.

Use the following procedure to revoke access to your CMK-encrypted datasets 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 SPICE dataset to become inaccessible.

Recovering an encrypted SPICE dataset

To recover a SPICE dataset while its access is revoked

1. Restore access to the CMK. Usually, this is enough to recover the dataset.
2. Test the SPICE dataset to see if you can see the data.
3. (Optional) If the data is not fully recovered, even after you restored its access to the CMK, perform a full refresh on the dataset.

Inter-network traffic privacy in Amazon QuickSight

To use Amazon QuickSight, users need access to the internet. They also need access to a compatible browser or a mobile device with the Amazon QuickSight mobile app installed. They don't need access to the data sources they want to analyze. This access is handled inside Amazon QuickSight. User connections to Amazon QuickSight are protected through the use of SSL. So that users can access Amazon QuickSight, 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 QuickSight accesses data sources by using connection information supplied by the data source owner in Amazon QuickSight. Connections are protected both between Amazon QuickSight and on-premises applications and between Amazon QuickSight and other Amazon resources within the same Amazon Web Services Region. For connections to any source, the data source must allow connections from Amazon QuickSight.

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 QuickSight 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 QuickSight is a logical entity within a VPC that allows connectivity only to Amazon QuickSight. The VPC routes requests to Amazon QuickSight and routes responses back to the VPC. For more information, see the following:

- [VPC endpoints](#) in the *Amazon VPC User Guide*
- [Connecting to a VPC with Amazon QuickSight](#)

Accessing data sources

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators and Amazon QuickSight administrators

Use this section to help you configure access to resources in other Amazon services.

We recommend that you use SSL to secure Amazon QuickSight connections to your data sources. To use SSL, you must have a certificate signed by a recognized certificate authority (CA). Amazon QuickSight doesn't accept certificates that are self-signed or issued from a nonpublic CA. For more information, see [QuickSight SSL and CA certificates](#).

Topics

- [Required permissions](#)
- [Allowing autodiscovery of Amazon resources](#)
- [Authorizing connections to Amazon data stores](#)
- [Accessing Amazon resources](#)
- [Exploring your Amazon data in Amazon QuickSight](#)

Required permissions

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

When you connect to a data source that requires a user name, the user name must have SELECT permissions on some system tables. These permissions allow Amazon QuickSight 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 dba_segments all_segments user_segments
PostgreSQL	pg_stats

Instance type	Tables
	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](#).

Allowing autodiscovery of Amazon resources

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Each Amazon service that you access from Amazon QuickSight needs to allow traffic from QuickSight. Instead of opening each service console separately to add permissions, a QuickSight 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 Amazon QuickSight account, Amazon QuickSight 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 QuickSight to use if you want Amazon QuickSight to autodiscover your Amazon resources.

You can have Amazon QuickSight 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 QuickSight 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 QuickSight to Amazon RDS DB instances](#). In this topic, you can find information on creating a security group to allow connections from Amazon QuickSight 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 QuickSight to Amazon Redshift clusters](#). In this topic, you can find information on creating a security group to allow connections from Amazon QuickSight servers.
- For an Amazon RDS DB instance or Amazon Redshift cluster that is in a nondefault VPC, see [Authorizing connections from Amazon QuickSight to Amazon RDS DB instances](#) or [Authorizing connections from Amazon QuickSight to Amazon Redshift clusters](#). In these topics, you can find information on first creating a security group to allow connections from Amazon QuickSight 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 QuickSight Region's public IP address.

Enabling autodiscovery is the easiest way to make this data available in Amazon QuickSight. You can still manually create data connections whether or not you enable autodiscovery.

Authorizing connections to Amazon data stores

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

For Amazon QuickSight to access your Amazon resources, you must create security groups for them that authorize connections from the IP address ranges used by Amazon QuickSight 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 QuickSight connections.

Topics

- [Authorizing connections from Amazon QuickSight to Amazon RDS DB instances](#)
- [Authorizing connections from Amazon QuickSight to Amazon Redshift clusters](#)
- [Authorizing connections from Amazon QuickSight to Amazon EC2 instances](#)
- [Authorizing connections through Amazon Lake Formation](#)
- [Authorizing connections to Amazon OpenSearch Service](#)
- [Authorizing connections to Amazon Athena](#)

Authorizing connections from Amazon QuickSight to Amazon RDS DB instances

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

For Amazon QuickSight 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 Amazon QuickSight servers in that Amazon Web Services Region. To learn more about authorizing Amazon QuickSight 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 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 QuickSight 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 datasets from new database data sources](#).

Manually enabling access to an Amazon RDS instance in a VPC

Use the following procedure to enable Amazon QuickSight 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 QuickSight) or that has Internet Gateways attached, see [Connecting to a VPC with Amazon QuickSight](#).


To enable Amazon QuickSight 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 QuickSight.

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 QuickSight in supported Amazon Web Services Regions, see [Amazon Web Services Regions, websites, IP address ranges, and endpoints](#).

 **Note**

If you have activated Amazon QuickSight in multiple Amazon Web Services Regions, you can create inbound rules for each Amazon QuickSight endpoint CIDR. Doing this allows Amazon QuickSight to have access to the Amazon RDS DB instance from any Amazon Region defined in the inbound rules.

Anyone who uses Amazon QuickSight in multiple Amazon Web Services Regions is treated as a single user. In other words, even if you are using Amazon QuickSight in every Amazon Web Services Region, both your Amazon QuickSight 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 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.

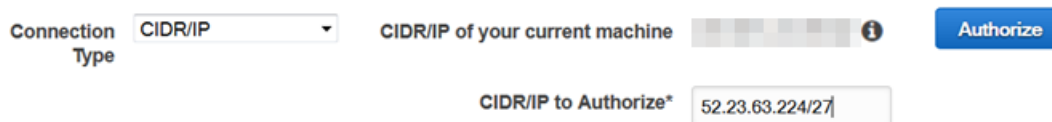
To use the console to access an Amazon RDS DB instance that is not in a VPC

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.



- For **Connection Type**, choose **CIDR/IP**.
- For **CIDR/IP to Authorize**, enter the appropriate CIDR address block. For more information on the IP address ranges for Amazon QuickSight in supported Amazon Web Services Regions, see [Amazon Web Services Regions, websites, IP address ranges, and endpoints](#).



- Choose **Authorize**.
- 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**.
- 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.
- Choose **Continue**, and then choose **Modify DB Instance**.

Authorizing connections from Amazon QuickSight 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 QuickSight 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 QuickSight asset to the user. QuickSight 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 QuickSight 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 QuickSight Q
- CSV, Excel, and PDF exports
- Anomaly detection

For Amazon QuickSight 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 QuickSight servers in that Amazon Web Services Region. To learn more about authorizing Amazon QuickSight connections, see [Manually enabling access to an Amazon Redshift cluster in a VPC](#).

Enabling connection from Amazon QuickSight 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 datasets from new database data sources](#).

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 QuickSight 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 QuickSight 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 QuickSight. 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 QuickSight.

Trusted identity propagation data sources are supported only in Direct Query datasets. SPICE datasets do not currently support 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 QuickSight accounts that are integrated with IAM Identity Center. For more information, see [Configure your Amazon QuickSight 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 QuickSight account that you want to use. The cluster must also be configured with the same organization instance in IAM Identity Center that your QuickSight account is configured to. For more information about configuring a Amazon Redshift cluster, see [Integrating IAM Identity Center](#).

Enabling trusted identity propagation in QuickSight

To configure QuickSight to connect to Amazon Redshift data sources with trusted identity propagation, configure Amazon Redshift OAuth scopes to your QuickSight account.

To add a scope that allows QuickSight to authorize identity propagation to Amazon Redshift, specify the Amazon Web Services account ID of the QuickSight 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 QuickSight 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, QuickSight authorizes users from any Amazon Redshift cluster that share the same IAM Identity Center instance. The example below configures QuickSight 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 QuickSight account.

```
aws quicksight delete-identity-propagation-config --aws-account-id "AWSACCOUNTID" --service "REDSHIFT" --authorized-targets "arn:aws-cn:sso::XXXXXXXXXXXX:application/ssoins-XXXXXXXXXXXXapl-XXXXXXXXXXXX" "arn:aws-cn:sso::XXXXXXXXXXXX:application/ssoins-XXXXXXXXXXXX/apl-XXXXXXXXXXXX"
```

The following example lists all OAuth scopes that are currently on a QuickSight 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 QuickSight. 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 QuickSight access to an Amazon Redshift cluster in a VPC.

To enable Amazon QuickSight 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 QuickSight.
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 QuickSight, navigate to the **Manage QuickSight** 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 QuickSight, 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 QuickSight. If the instance ID doesn't auto populate, choose the VPC that you created from the dropdown list.
- Enter the database credentials. If your QuickSight 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 QuickSight 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 QuickSight 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 QuickSight access to Amazon S3 or Athena. Amazon QuickSight 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 QuickSight to Amazon EC2 instances

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators


For Amazon QuickSight 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 Amazon QuickSight 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 Amazon QuickSight 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 datasets from new database data sources](#).

To enable Amazon QuickSight 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 QuickSight. 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

QuickSight in supported Amazon Regions, see [Amazon Web Services Regions, websites, IP address ranges, and endpoints](#).

Note

If you have activated Amazon QuickSight in multiple Amazon Web Services Regions, you can create inbound rules for each Amazon QuickSight endpoint CIDR. Doing this allows Amazon QuickSight to have access to the Amazon RDS DB instance from any Amazon Web Services Region defined in the inbound rules.

An Amazon QuickSight user or administrator who uses Amazon QuickSight in multiple Amazon Regions is treated as a single user. In other words, even if you are using Amazon QuickSight in every Amazon Web Services Region, both your Amazon QuickSight account and your users are global.

Create Security Group [X]

Security group name ⓘ Amazon-QuickSight-access

Description ⓘ Amazon-QuickSight-access

VPC ⓘ vpc-b423b7d1 (172.31.0.0/16) *
* denotes default VPC

Security group rules:

Inbound Outbound

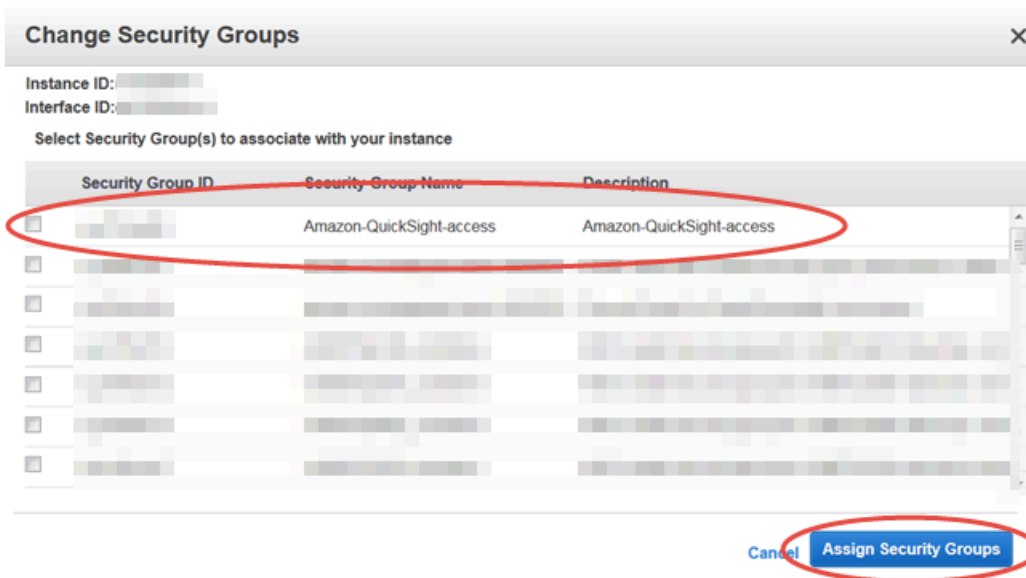
Type ⓘ	Protocol ⓘ	Port Range ⓘ	Source ⓘ
Custom TCP Rule ▾	TCP	1433	Custom IP ▾ 52.23.63.224/27 [X]

Add Rule

Cancel Create

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 QuickSight. 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 QuickSight to access databases and tables by name or through SQL queries. Amazon QuickSight 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*.

Enabling connection from Lake Formation

Before you begin using this solution with Amazon QuickSight, 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 QuickSight 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 Amazon QuickSight users and groups that need to access the data in Lake Formation. These users should be Amazon QuickSight authors or administrators.

To find Amazon QuickSight user and group ARNs

1. Use the Amazon CLI to find user ARNs for Amazon QuickSight 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/AIDAJVCZOVSR3DESMJ7TA
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 QuickSight 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 QuickSight 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 QuickSight console. For more information, see [Creating and managing groups in Amazon QuickSight](#).

To avoid using the Amazon CLI, you can construct the ARNs for each group manually.

Enabling connection from Amazon QuickSight

To work with Lake Formation and Athena, make sure that you have Amazon resource permissions configured in Amazon QuickSight:

- 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 Amazon QuickSight, 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 QuickSight dataset, there are a few tasks for the QuickSight 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 QuickSight administrator enables authorized connections from QuickSight to OpenSearch Service. This process is required for every Amazon service that you connect to from QuickSight. 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 QuickSight to OpenSearch Service

1. Within Amazon QuickSight, choose **Administrator** and **Manage QuickSight**.
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.

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 QuickSight 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 VPC data source](#).

If a connection isn't already defined in QuickSight 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 QuickSight to a VPC and add a connection from QuickSight to the VPC, see [Connecting to a VPC with Amazon QuickSight](#).

Using OpenSearch permissions

After you configure QuickSight 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 QuickSight.
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 QuickSight. 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 QuickSight

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/fields/get*`
 - `indices:data/read/search*`
 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 QuickSight. 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 QuickSight 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 QuickSight 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 QuickSight:

- Athena stores query results from QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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, QuickSight 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 QuickSight 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 QuickSight, 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 QuickSight 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 QuickSight 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 QuickSight.

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 QuickSight 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 QuickSight administrator. This error means need to recheck your connection authorization settings to find the discrepancy. .

5. After you can connect successfully, you or your QuickSight authors can create data sources connections and share them with other QuickSight authors. The authors can then create multiple datasets from the connections, to use in QuickSight dashboards.

For troubleshooting information on Athena, see [Connectivity issues when using Amazon Athena with Amazon QuickSight](#).

Accessing Amazon resources

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators and Amazon QuickSight administrators

You can control the Amazon resources that Amazon QuickSight 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.

Use the following sections to help you configure your Amazon resources to work with Amazon QuickSight.

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 QuickSight. For more information, see [Passing IAM roles to Amazon QuickSight](#)

To enable or disable the Amazon services that Amazon QuickSight can access

1. Sign in to Amazon QuickSight at <https://quicksight.aws.amazon.com/>.
2. At the upper right, choose your user name, and then choose **Manage QuickSight**.
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 QuickSight 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 QuickSight 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 QuickSight for Amazon resource access was modified outside of Amazon QuickSight, so you can no longer edit this policy to provide Amazon resource permission to Amazon QuickSight. 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 QuickSight 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 default resource access to Amazon services](#)
- [Setting granular access to Amazon services through IAM](#)
- [Using Amazon Secrets Manager secrets instead of database credentials in Amazon QuickSight](#)

Setting default resource access to Amazon services

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon QuickSight administrators

In Enterprise edition, you can configure specific permissions for the Amazon services that an Amazon QuickSight user can access. If no such configuration occurs, Amazon QuickSight 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 QuickSight 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 **Default resource access**, choose **Change**.
5. Choose one of the following:

- Allow access to all Amazon data and resources.
- Deny access to all Amazon data and resources.

Setting granular access to Amazon services through IAM

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon QuickSight administrators

In Enterprise edition, Amazon QuickSight provides a way for you to set up detailed access to resources in Amazon services. Like every other Amazon service, Amazon QuickSight 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 Amazon QuickSight, see [Identity and access management in Amazon QuickSight](#).

To assign an IAM policy to a user or group

1. Sign in to Amazon QuickSight 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 QuickSight user or group. To determine permissions, Amazon QuickSight 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 QuickSight 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 Amazon QuickSight

Intended audience: Amazon QuickSight Administrators and Amazon QuickSight 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](#).

Amazon QuickSight administrators can grant QuickSight 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 QuickSight API.

QuickSight 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 Amazon QuickSight, 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 QuickSight.

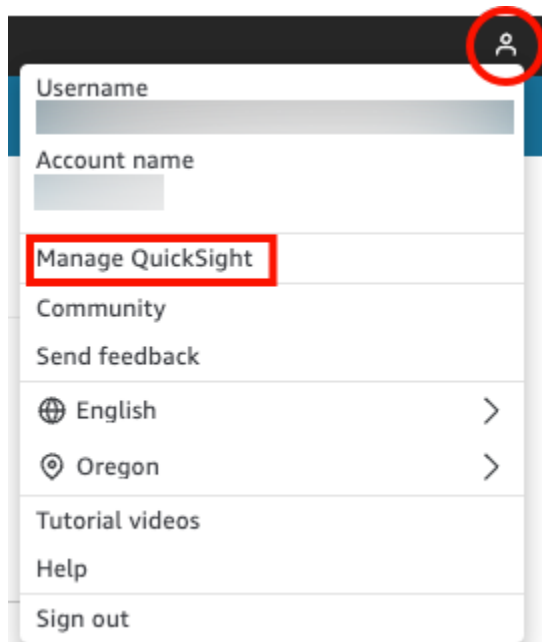
Use the following procedures described in the following sections to integrate Secrets Manager with Amazon QuickSight.

Granting QuickSight access to Secrets Manager and selected secrets

If you're an administrator and you have secrets in Secrets Manager, you can grant Amazon QuickSight read-only access to selected secrets.

To grant QuickSight access to Secrets Manager and selected secrets

1. In QuickSight, choose your user icon on the upper right, and then choose **Manage QuickSight**.



2. Choose **Security & permissions** on the left.
3. Choose **Manage** in **QuickSight access to Amazon resources**.

Manage users

Manage groups

Your subscriptions

SPICE capacity

Account settings

Security & permissions

Manage VPC connections

Mobile settings

Domains and Embedding

Account customization

Single sign-on (SSO)

Security & permissions

QuickSight can control access to AWS resources for the entire account in addition to individual users and groups

QuickSight access to AWS services

By configuring access to AWS services, QuickSight can access the data in those services. Access by users and groups can

IAM role in use

Quicksight-managed role (default)

Access granted to 4 services

 Amazon Redshift

 Amazon RDS

 IAM

 Amazon Athena

Manage

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 QuickSight read-only access to.

Secrets in your QuickSight 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**.

QuickSight 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:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "secretsmanager:GetSecretValue"
      ],
      "Resource": [
        "arn:aws-cn:secretsmanager:region:accountId:secret:secret_name"
      ]
    }
  ]
}
```



```
    ]  
  }  
]  
}
```

When QuickSight users create analyses from or view dashboards that use a data source with secrets, QuickSight 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 QuickSight 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, QuickSight 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 QuickSight 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 QuickSight API

After the QuickSight administrator has granted QuickSight 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 QuickSight. 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 QuickSight 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 QuickSight. They can also view and update the data source permissions. Instead of a QuickSight 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, QuickSight 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 QuickSight 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 QuickSight 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

Quicksight requires the following JSON format to access your secret:

```
{
  "username": "username",
  "password": "password"
}
```

The username and password fields are required for QuickSight to access secrets. All other fields are optional and are ignored by QuickSight.

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*.

Modifying the secret

To modify a secret, you use Secrets Manager. After you make changes to a secret, the updates become available the next time QuickSight requests access to the secret.

Exploring your Amazon data in Amazon QuickSight

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Use this section to learn how to explore Amazon data in Amazon QuickSight using the Amazon Web Services Management Console. Using the **Explore in QuickSight** shortcut, you can access a customizable dashboard template showing your data. Just as with any Amazon QuickSight dashboard, this dashboard can be refreshed on a schedule, published, and shared with other users in your organization.

Topics

- [Exploring Amazon S3 analytics data](#)

Exploring Amazon S3 analytics data

Amazon QuickSight contains a dashboard designed to provide insight into your Amazon S3 analytics data. To use this feature, you must first enable S3 analytics storage class analysis for your S3 buckets. For more on enabling storage class analysis in S3, see [Amazon S3 analytics – Storage class analysis](#) in the *Amazon S3 Developer Guide*.

After you have enabled storage class analysis, you can use Amazon QuickSight to explore your S3 analytics data.

To explore S3 analytics data in Amazon QuickSight

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 QuickSight account, you're prompted to create one before you can use the dashboard.

The screenshot shows the Amazon S3 console interface. At the top, there are tabs for 'Objects', 'Properties', 'Permissions', and 'Management'. Below these are buttons for 'Lifecycle', 'Analytics', 'Metrics', and 'Inventory'. The 'Analytics' tab is selected. On the left, there is a search bar and a list of filters, including 'anotherS3A' and 'test'. The main content area is titled 'Storage class analysis' and features a blue button labeled 'Explore in QuickSight' which is circled in red. Below this, there is a status message: 'Analyzing your data (started on 5/22/2017) We are analyzing your storage usage and will share the observed infrequent access patterns.' A table shows 'Standard' storage with a 'Total size as of 6/21/2017' of '2.4 MB' and 'Data retrieved for 6/21/2017' of '0 B'. At the bottom, there is a chart titled 'How much of my storage did I retrieve?' with a legend for 'Storage' (blue) and 'Data retrieved' (purple). The chart shows a blue area representing storage and a purple area representing data retrieved, with a peak in the purple area around the middle of the x-axis.

When you choose the option to explore in Amazon QuickSight, your 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 S3 analytics.

To learn more about using S3 analytics and storage class analysis, see [Amazon S3 analytics – Storage class analysis](#) in the *Amazon S3 Developer Guide*.

Identity and access management in Amazon QuickSight

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators and Amazon QuickSight administrators

Amazon QuickSight 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 Amazon QuickSight Standard edition, you can manage users entirely within Amazon QuickSight. 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 QuickSight:

- [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 \(IAM Identity Center\)](#) (Standard and Enterprise edition)
- [Multifactor authentication \(MFA\)](#) (Standard and Enterprise edition)

Using service control policies to restrict Amazon QuickSight 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 QuickSight. You can restrict the edition of Amazon QuickSight 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 Amazon QuickSight 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](#)

Restricting the Amazon QuickSight edition

To restrict the edition of Amazon QuickSight 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>	QuickSight Standard Edition
	<code>enterprise</code>	QuickSight Enterprise Edition

Restricting user management options

To restrict the user management options that individuals in your organization can use to sign up for Amazon QuickSight, 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 QuickSight-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 Directory Service for Microsoft Active Directory
	<code>iam_identity_center</code>	Users managed in a QuickSight account that is integrated with IAM Identity Center.

Example SCP

The following example for Amazon QuickSight shows a service control policy that denies signing up for a QuickSight Standard Edition and turns off the ability to sign up using QuickSight or Active Directory credentials. This policy uses the `quicksight:subscribe` action, in addition to the condition keys previously described. For a list of QuickSight-specific keys for use in IAM permission policies, see [Actions, resources, and condition keys for Amazon QuickSight](#) 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 QuickSight Enterprise Edition. Additionally, they can sign up only by using the **IAM Identity Center enabled application** option. If they try to sign up for QuickSight 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 QuickSight.

Using Amazon Identity and Access Management (IAM)

Following, you can find an introduction to the concepts involved in using Amazon Identity and Access Management (IAM). This section also covers how to use IAM with Amazon QuickSight.

Amazon Identity and Access Management (IAM) is an Amazon Web 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 QuickSight resources. IAM is an Amazon Web Service that you can use with no additional charge.

Topics

- [Introduction to IAM concepts](#)
- [Using Amazon QuickSight with IAM](#)
- [IAM policy examples for Amazon QuickSight](#)
- [Provisioning users for Amazon QuickSight](#)
- [Troubleshooting Amazon QuickSight 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 QuickSight resources. IAM is an Amazon service that you can use with no additional charge.

IAM is used with Amazon QuickSight 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 QuickSight.

- If you want your Amazon QuickSight 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 QuickSight.
- If you want to create specialized access for specific groups of QuickSight users or to specific resources, you can use IAM policies to accomplish this.

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 QuickSight.

Service user – In some cases, you might use Amazon QuickSight to interact with data, analyses, and dashboards through the Amazon QuickSight 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 QuickSight.

Amazon QuickSight administrator – If you're in charge of Amazon QuickSight resources at your company, you probably have full access to Amazon QuickSight. It's your job to determine which Amazon QuickSight features and resources your team members should access. If you have specialized requirements that you can't solve by using the Amazon QuickSight admin panel, then you can work with your administrator to create permissions policies for your Amazon QuickSight 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 QuickSight, see [Using Amazon QuickSight with IAM](#).

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

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 [Signing Amazon 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 [Using multi-factor authentication \(MFA\) in Amazon](#) in the *IAM User Guide*.

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 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 [When to create an IAM user \(instead of a role\)](#) 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. You can temporarily assume

an IAM role in the Amazon Web Services Management Console by [switching roles](#). 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 [Using IAM roles](#) 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 [Creating a role for a third-party Identity Provider](#) 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, 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 [How IAM roles differ from resource-based policies](#) in the *IAM User Guide*.
- **Cross-service access** – Some Amazon Web Services use features in other Amazon Web 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 Service, combined with the requesting Amazon Web 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 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 [Creating a role to delegate permissions to an Amazon Web 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 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 [Using an IAM role to grant permissions to applications running on Amazon EC2 instances](#) in the *IAM User Guide*.

To learn whether to use IAM roles or IAM users, see [When to create an IAM role \(instead of a user\)](#) 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 [Creating IAM 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 [Choosing 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.

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 [How SCPs work](#) 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 Amazon QuickSight with IAM

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Before you use IAM to manage access to Amazon QuickSight, you should understand what IAM features are available to use with Amazon QuickSight. To get a high-level view of how Amazon

QuickSight and other Amazon services work with IAM, see [Amazon Services That Work with IAM](#) in the *IAM User Guide*.

Topics

- [Amazon QuickSight Policies \(identity-based\)](#)
- [Amazon QuickSight policies \(resource-based\)](#)
- [Authorization based on Amazon QuickSight tags](#)
- [Amazon QuickSight IAM roles](#)
- [Passing IAM roles to Amazon QuickSight](#)

Amazon QuickSight 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 QuickSight 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 QuickSight account. Amazon root and administrator credentials already have all of the required permissions for managing Amazon QuickSight 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 QuickSight. The policy must include the appropriate statements for the Amazon QuickSight administrative tasks you need to perform, as described in the following sections.

Important

Be aware of the following when working with Amazon QuickSight and IAM policies:

- Avoid directly modifying a policy that was created by Amazon QuickSight. When you modify it yourself, Amazon QuickSight 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 QuickSight account, see [Actions Defined by Amazon QuickSight](#) in the *IAM User Guide*.
- In some cases, you might have an Amazon QuickSight 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 QuickSight account, then recreate it. For more information, see [Deleting your Amazon QuickSight subscription and closing the account](#).

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

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

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 QuickSight 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 QuickSight identity-based policies, see [IAM identity-based policies for Amazon QuickSight](#).

Amazon QuickSight policies (resource-based)

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

Authorization based on Amazon QuickSight tags

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

Amazon QuickSight 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 QuickSight actions.

Amazon QuickSight doesn't support the following role features:

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

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

Passing IAM roles to Amazon QuickSight

Applies to: Enterprise Edition

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

Prerequisites

For your users to pass IAM roles to QuickSight, 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 QuickSight to assume the role.** Use the following example to create a trust policy for the role. The following example trust policy allows the Amazon QuickSight 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 QuickSight administrators the permission to update resource-level permissions in QuickSight. For more information about resource types defined by QuickSight, see [Actions, resources, and condition keys for Amazon QuickSight](#) in the *IAM User Guide*.
 - `iam:PassRole` – This grants users permission to pass roles to QuickSight. 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 QuickSight. 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 Amazon QuickSight.

```
{
  "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 QuickSight, see [IAM policy examples for Amazon QuickSight](#).

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

After your administrator completes the prerequisites, your IAM users can pass IAM roles to QuickSight. They do so by choosing an IAM role when they [sign up for QuickSight](#), or by [switching to an IAM role](#) on their QuickSight Security and Permissions page. To learn how to switch to an existing IAM role in QuickSight, see the following section.

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 QuickSight permission to perform specific actions. You can then

attach the policy to the IAM roles that you later pass to QuickSight. 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 QuickSight in Athena, see [AWSQuicksightAthenaAccess Managed Policy](#) in the *Amazon Athena User Guide*. IAM users can access this role in QuickSight using the following ARN: `arn:aws-cn:iam::aws:policy/service-role/AWSQuicksightAthenaAccess`.

The following is an example of a permissions policy for QuickSight 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 QuickSight to an Amazon S3 bucket in another account, see [How do I set up cross-account access from Amazon QuickSight 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 Amazon QuickSight

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

Use the following procedure to learn how to pass IAM roles in QuickSight.

To use an existing IAM role in QuickSight

1. In QuickSight, choose your account name in the navigation bar at top right and choose **Manage QuickSight**.
2. On the **Manage QuickSight** page that opens, choose **Security & Permissions** in the menu at left.
3. In the **Security & Permissions** page that opens, under **QuickSight 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 Amazon QuickSight

This section provides examples of IAM policies that you can use with Amazon QuickSight.

IAM identity-based policies for Amazon QuickSight

This section shows examples of identity-based policies to use with Amazon QuickSight.

Topics

- [IAM identity-based policies for QuickSight IAM console administration](#)
- [IAM identity-based policies for Amazon QuickSight: dashboards](#)
- [IAM identity-based policies for Amazon QuickSight: namespaces](#)
- [IAM identity-based policies for Amazon QuickSight: custom permissions](#)
- [IAM identity-based policies for Amazon QuickSight: customizing email report templates](#)
- [IAM identity-based policies for Amazon QuickSight: creating users](#)
- [IAM identity-based policies for Amazon QuickSight: creating and managing groups](#)
- [IAM identity-based policies for Amazon QuickSight: All access for Standard edition](#)
- [IAM identity-based policies for Amazon QuickSight: All access for Enterprise edition with IAM Identity Center](#)
- [IAM identity-based policies for Amazon QuickSight: all access for Enterprise edition with Active Directory](#)
- [IAM identity-based policies for Amazon QuickSight: active directory groups](#)
- [IAM identity-based policies for Amazon QuickSight: using the admin asset management console](#)
- [IAM identity-based policies for Amazon QuickSight: using the admin key management console](#)
- [Amazon resources Amazon QuickSight: scoping policies in Enterprise edition](#)

IAM identity-based policies for QuickSight IAM console administration

The following example shows the IAM permissions needed for QuickSight IAM console administration actions.

```
{
  "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 Amazon QuickSight: 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 Amazon QuickSight: namespaces

The following examples show IAM policies that allow a QuickSight 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 Amazon QuickSight: custom permissions

The following example shows an IAM policy that allows a QuickSight 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 Amazon QuickSight: customizing email report templates

The following example shows a policy that allows viewing, updating, and creating email report templates in QuickSight, as well as obtaining verification attributes for an Amazon Simple Email Service identity. This policy allows a QuickSight 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 Amazon QuickSight: creating users

The following example shows a policy that allows creating Amazon QuickSight 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 Amazon QuickSight: creating and managing groups

The following example shows a policy that allows QuickSight 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 Amazon QuickSight: All access for Standard edition

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

```

{
  "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:Subscribe"
      ],
      "Resource": "*"
    },
    {
      "Effect": "Deny",
      "Action":
        "quicksight:Unsubscribe",
      "Resource": "*"
    }
  ]
}

```

IAM identity-based policies for Amazon QuickSight: All access for Enterprise edition with IAM Identity Center

The following example for Amazon QuickSight Enterprise edition shows a policy that allows subscribing, creating users, and managing Active Directory in a QuickSight account that is integrated with IAM Identity Center. This example explicitly denies permission to unsubscribe from Amazon QuickSight.

```
{
  "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:DescribeGroup",
        "sso:SearchGroups",
        "sso:GetProfile",
        "sso:CreateApplicationAssignment",
```



```

        "sso:DeleteApplicationAssignment",
        "sso:ListInstances",
        "sso:DescribeRegisteredRegions",
        "organizations:DescribeOrganization"
    ],
    "Resource": [
        "*"
    ]
}
]
}

```

IAM identity-based policies for Amazon QuickSight: all access for Enterprise edition with Active Directory

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

```

{
  "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 Amazon QuickSight: active directory groups

The following example shows an IAM policy that allows Active Directory group management for an Amazon QuickSight 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 Amazon QuickSight: 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 Amazon QuickSight: 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": [
        "kms:CreateGrant",
        "kms:ListAliases",
        "kms:ListKeys",
        "quicksight:ListCustomerManagedKeys",
        "quicksight:ListKMSKeysForUser",

```

```
        "quicksight:RegisterCustomerManagedKey"  
        "quicksight:RemoveCustomerManagedKey",  
    ],  
    "Resource": "*" ]  
}
```

Amazon resources Amazon QuickSight: scoping policies in Enterprise edition

The following example for Amazon QuickSight 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 QuickSight

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators and Amazon QuickSight administrators

Self-provisioning an Amazon QuickSight administrator

Amazon QuickSight administrators are users who can also manage Amazon QuickSight features such as account settings and accounts. They can also purchase additional Amazon QuickSight user subscriptions, purchase [SPICE](#) capacity, and cancel the subscription to Amazon QuickSight 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 QuickSight. 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 QuickSight. To set up self-provisioning, give these users permission to use the `quicksight:CreateAdmin` action.

Granting permissions with IAM actions only affects the specified user's ability to create a QuickSight account for their specified role. After a user has created and logged into their account, you use a separate set of permissions within QuickSight to manage QuickSight-specific features. For more information, see [Customizing access to the Amazon QuickSight console](#).

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

To make a user the Amazon QuickSight administrator

1. Create the Amazon user:
 - Use IAM to create the user that you want to be the administrator of Amazon QuickSight. 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 QuickSight.

Self-provisioning an Amazon QuickSight author

Amazon QuickSight authors can create data sources, data sets, analyses, and dashboards. They can share analyses and dashboards with other Amazon QuickSight users in your Amazon QuickSight account. However, they don't have access to the **Manage Amazon QuickSight** menu. They can't change account settings, manage accounts, purchase additional Amazon QuickSight user subscriptions or [SPICE](#) capacity, or cancel the subscription to Amazon QuickSight for your Amazon Web Services account.

You can use an Amazon user or group policy to give users the ability to create an Amazon QuickSight author account for themselves. Their accounts become active and billable the first time

they open Amazon QuickSight. To set up self-provisioning, you need to give them permission to use the `quicksight:CreateUser` action.

Self-provisioning an Amazon QuickSight read-only user

Amazon QuickSight 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 QuickSight readers can't create data sources, data sets, 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.

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

Reader accounts become active and billable the first time they open Amazon QuickSight. 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. For more information about user and capacity pricing, see [Amazon QuickSight Pricing](#).

Troubleshooting Amazon QuickSight 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 QuickSight and IAM.

Topics

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

I am not authorized to perform an action in Amazon QuickSight

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 QuickSight.

Some Amazon Web 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 QuickSight. 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 QuickSight 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 QuickSight supports these features, see [Using Amazon QuickSight 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 [How IAM roles differ from resource-based policies](#) in the *IAM User Guide*.

Identity management in Amazon QuickSight

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators and Amazon QuickSight administrators

Amazon QuickSight is an Amazon IAM Identity Center enabled application. We recommend that you integrate new QuickSight subscriptions with for identity management.

Use IAM Identity Center to configure an external identity provider or to create users in the IAM Identity Center identity store. When QuickSight is configured with IAM Identity Center, users

and groups created in IAM Identity Center are used to share assets and to assign users to roles in QuickSight.

IAM Identity Center is supported in QuickSight Enterprise Edition subscriptions. For more information about IAM Identity Center, see [Amazon IAM Identity Center](#).

Following is a list of other supported identity configurations in Amazon QuickSight:

- IAM (Standard and Enterprise editions)
- IAM federation (Standard and Enterprise editions)
- Amazon Directory Service for Microsoft Active Directory (Enterprise Edition only)

Configure your Amazon QuickSight 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 QuickSight and IAM Identity Center so that you can sign up for a new QuickSight 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 QuickSight. Identity methods can't be changed after your account is created.

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

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

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

Admin action	IAM permissions	QuickSight admin role permissions
Manage assets	Yes	
Security & permissions	Yes	
Manage VPC connections	Yes	
KMS keys	Yes	
Account settings	Yes	
Account customization		Yes
Manage users		Yes
Your subscriptions		Yes
Mobile settings		Yes
Domains and embedding		Yes
SPICE capacity		Yes

Role level custom permissions are supported for accounts that are integrated with IAM Identity Center. User level custom permissions are not supported for QuickSight accounts that are integrated with IAM Identity Center. For more information about customizing access to the QuickSight console, see [Customizing access to the Amazon QuickSight console](#).

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

Considerations

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

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

Using external identity federation and single sign-on with Amazon QuickSight

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Note

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

Amazon QuickSight 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 Amazon QuickSight.

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 QuickSight 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 QuickSight 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 QuickSight using your existing identity provider.

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*

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](#)

Topics

- [IAM federation](#)
- [Using Active Directory with Amazon QuickSight Enterprise edition](#)
- [Directory integration with Amazon QuickSight Enterprise edition](#)
- [Using multi-factor authentication \(MFA\) with Amazon QuickSight](#)

IAM federation

Note

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

Topics

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

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 QuickSight.

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 QuickSight. After QuickSight checks that they are authorized, your users can access QuickSight.

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

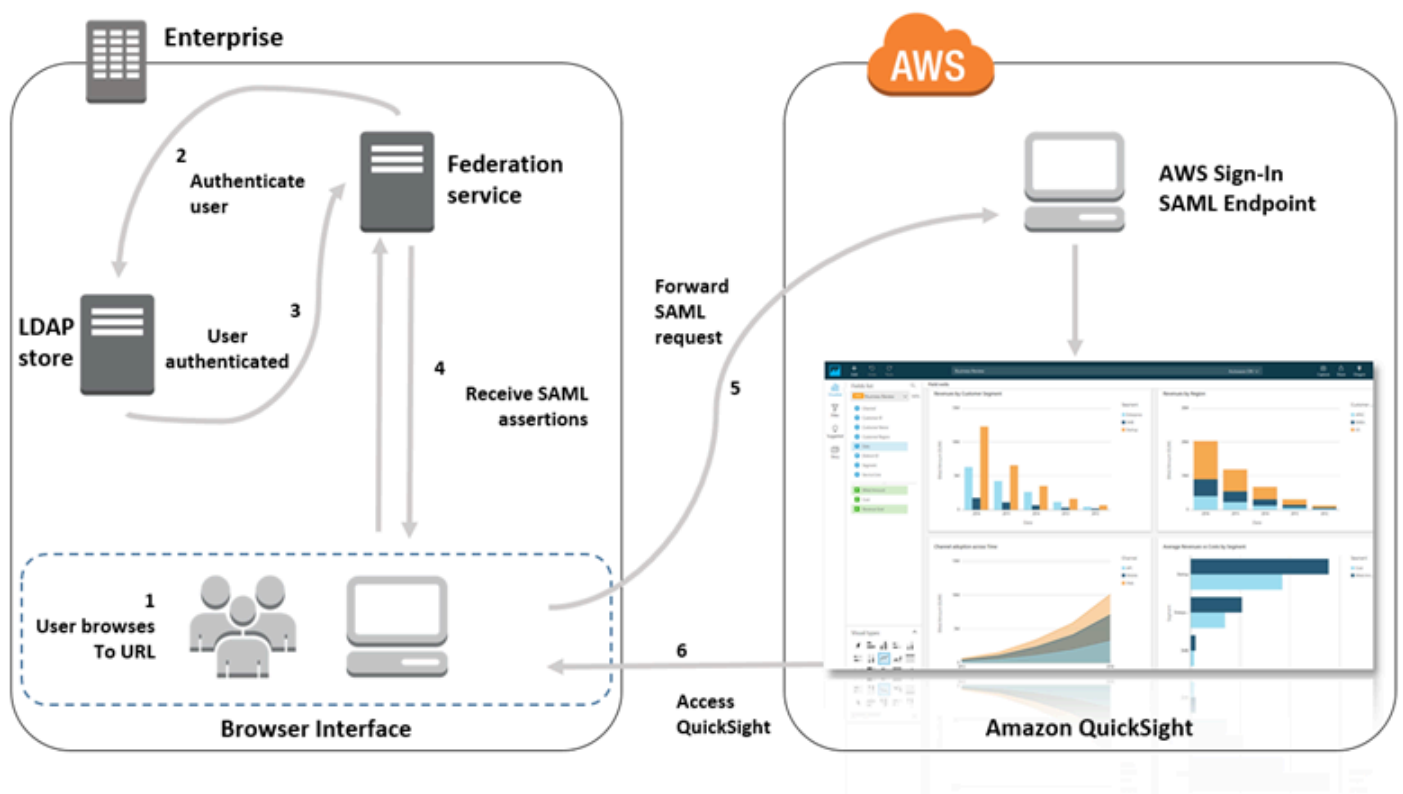
1. The user browses to `https://applications.example.com` and signs on to the IdP. At this point, the user isn't signed in to the service provider.
2. The federation service and the IdP authenticate the user:
 - a. The federation service requests authentication from the organization's identity store.
 - 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 QuickSight:

- 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 QuickSight service.
4. Amazon QuickSight accepts the authentication token from Amazon and presents QuickSight 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 QuickSight application portal, without ever having to supply any Amazon credentials.

In the following diagram, you can find an authentication flow between Amazon QuickSight and a third-party identity provider (IdP). In this example, the administrator has set up a sign-in page to access Amazon QuickSight, 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.



Setting up IdP federation using IAM and QuickSight

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Note

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

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 QuickSight. 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 Amazon QuickSight. For more information, see [Step 5: Configure the relay state of your federation](#).
- Create an Amazon QuickSight account and note the name to use when you configure your IAM policy and IdP. For more information on creating an Amazon QuickSight account, see [Signing up for an Amazon QuickSight 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 QuickSight, 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 QuickSight.

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 QuickSight. 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 QuickSight by attaching an inline policy to the role. The following sample policy document provides access to Amazon QuickSight. This policy allows the user access to Amazon QuickSight 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 QuickSight and also the ability to create Amazon QuickSight 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 QuickSight, 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 QuickSight relay state URL. After successful authentication by Amazon, the user is directed to Amazon QuickSight, defined as the relay state in the SAML authentication response.

The relay state URL for Amazon QuickSight is as follows.

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

Initiating sign-on from Amazon QuickSight

Applies to: Enterprise Edition

Intended audience: System administrators

Note

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

In this scenario, your user initiates the sign-on process from an Amazon QuickSight 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 QuickSight. QuickSight sends an authentication request to the IdP. After the user is authenticated, QuickSight opens.

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

1. The user opens QuickSight. At this point, the user isn't signed in to the IdP.
2. The user attempts to sign in to QuickSight.
3. QuickSight 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 QuickSight service.
5. Amazon QuickSight accepts the authentication token from Amazon and presents QuickSight to the user.

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

Setting up service provider–initiated federation with Amazon QuickSight Enterprise edition

Applies to: Enterprise Edition

Intended audience: System administrators

Note

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

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 QuickSight Enterprise Edition. For QuickSight-initiated IAM federation to work, you need to authorize QuickSight to send the authentication request to your IdP. A QuickSight administrator can configure this by adding the following information provided by the IdP:

- The IdP URL – QuickSight 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 Amazon QuickSight URL that you provide.

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

QuickSight supports connecting to one IdP per Amazon Web Services account. The configuration page in QuickSight 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, QuickSight provides a parameter (`enable-sso=0`) to temporarily turn off QuickSight initiated IAM federation, in case you need to disable it temporarily.

To set up QuickSight 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 QuickSight. To test this setup, check if you can share a dashboard with another person in your company's domain.
2. Open QuickSight, and choose **Manage QuickSight** from your profile menu at upper right.

To perform this procedure, you need to be a QuickSight administrator. If you aren't, you can't see **Manage QuickSight** 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 QuickSight, for example `https://quicksight.aws.amazon.com/sn/start`.
 - To test signing in with QuickSight 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 QuickSight, and choose **Manage QuickSight** 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 QuickSight.

To disable service provider initiated IAM federation

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

Tutorial: Amazon QuickSight and IAM identity federation

Applies to: Enterprise Edition and Standard Edition

Intended audience: Amazon QuickSight Administrators and Amazon QuickSight developers

Note

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

In the following tutorial, you can find a walkthrough for setting up the IdP Okta as a federation service for Amazon QuickSight. 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 QuickSight-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 - QuickSight**.
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 QuickSight. You can do one of the following:

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

- Provision QuickSight users in advance by using the [QuickSight API](#). By choosing this option, you can provision users and add them to groups at the same time. For more information, see [Creating and managing groups in Amazon QuickSight](#).

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 QuickSight user, use [register-user](#). To add a QuickSight group, use [create-group](#). To add users to the QuickSight 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 QuickSight users automatically, rather than using the QuickSight 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 QuickSight administrator can later upgrade them from the QuickSight profile menu, **Manage QuickSight, 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 - QuickSight**—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 QuickSight 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 - QuickSight**.
3. Choose the application icon to launch **Amazon Account Federation - QuickSight**.

You can now manage identities using Okta and use federated access with Amazon QuickSight.

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

(Optional) To set up QuickSight to send authentication requests to Okta

1. Open QuickSight, and choose **Manage QuickSight** 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 QuickSight, for example `https://quicksight.aws.amazon.com/sn/start`.
 - To test signing in with QuickSight 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`, QuickSight doesn't send the authentication request, and you sign in to QuickSight as before.
6. For **Status**, choose **ON**.
7. Choose **Save** to keep your settings.

You can create a deep link to a QuickSight 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 QuickSight 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 QuickSight 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 QuickSight 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.

Configuring email syncing for federated users in Amazon QuickSight

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon QuickSight administrators

Note

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

In Amazon QuickSight Enterprise edition, as an administrator you can restrict new users from using personal email addresses when provisioning through their identity provider (IdP) directly to QuickSight. QuickSight 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 QuickSight account through your IdP.

Note

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

When you configure email syncing for federated users in QuickSight, users who log in to your QuickSight 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.

QuickSight 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 QuickSight.

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 QuickSight. 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 QuickSight

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 QuickSight as described in the following procedure.

To turn on email syncing for federated users

1. From any page in QuickSight, choose your username at top right, and then choose **Manage QuickSight**.
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, QuickSight 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, QuickSight asks users to input their email address manually when provisioning new users to your account. They can use any email addresses that they want.

Using Active Directory with Amazon QuickSight Enterprise edition

Applies to: Enterprise Edition

Intended audience: System administrators

Note

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

Amazon QuickSight 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 Amazon QuickSight, 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 QuickSight, 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 Amazon QuickSight, 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 Amazon QuickSight?](#) 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 Amazon QuickSight by creating at least three groups for users:

- **Amazon QuickSight admins** – Admins can change account settings, manage accounts. Admins can also purchase additional Amazon QuickSight user subscriptions or [SPICE](#) capacity, or cancel the subscription to Amazon QuickSight for your Amazon Web Services account.

- **Amazon QuickSight authors** – Amazon QuickSight authors can create data sources, datasets, analyses, and dashboards. They can share analyses and dashboards with other Amazon QuickSight users.
- **Amazon QuickSight readers** – Readers can view and interact with dashboards that were created by someone else.

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

Directory integration with Amazon QuickSight Enterprise edition

Applies to: Enterprise Edition

Intended audience: System administrators

Note

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

Amazon QuickSight 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 Amazon QuickSight. In this example, users authenticate using their Active Directory login credentials. To make access to Amazon QuickSight transparent to your users, use IAM federation in this scenario.

Using multi-factor authentication (MFA) with Amazon QuickSight

Applies to: Enterprise Edition and Standard Edition

Intended audience: System administrators

Note

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

There are several ways that you can use multi-factor authentication (MFA) with Amazon QuickSight. 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

Managing user access inside Amazon QuickSight

Intended audience: System administrators and Amazon QuickSight administrators

QuickSight administrators manage user access in QuickSight. User access management in QuickSight is determined by your QuickSight account identity configuration. For accounts that use IAM Identity Center or Active Directory, groups are assigned to QuickSight roles. Groups can be assigned the Admin, Author, Reader, Admin Pro, Author Pro, or Reader Pro roles. For more information about Pro roles in see [???](#). For more information about integrating your QuickSight account with IAM Identity Center, see [Managing access for IAM Identity Center users](#).

QuickSight accounts that use QuickSight and IAM users create users directly in QuickSight. These users and their roles are managed at the user level. For more details, see [Managing access for QuickSight and IAM users](#).

Topics

- [Managing access for IAM Identity Center users](#)
- [Managing access for QuickSight and IAM users](#)

Managing access for IAM Identity Center users

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon QuickSight 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 QuickSight accounts that use Active Directory.

To manage QuickSight users, you must have administrative privileges in Amazon QuickSight and also the appropriate Amazon permissions. For more information about the necessary Amazon permissions, see [IAM policy examples for Amazon QuickSight](#). If you are using directory groups, you need to be a network administrator.

Each Amazon QuickSight Enterprise edition account can have an unlimited number of users. User names that contain a semicolon (;) aren't supported.

Use the following procedures to add, view, and deactivate Amazon QuickSight users.

Important

You can't remap Amazon QuickSight 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 QuickSight. 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 QuickSight assets and settings before the migration.

Adding users

With IAM Identity Center, add users to QuickSight by associating their IAM Identity Center group to an Admin, Admin Pro, Author, Author Pro, Reader, or Reader Pro role in QuickSight. All users in the selected groups are authorized to sign in to Amazon QuickSight.

For more information about Pro roles in see [???](#).

To see which groups are integrated with your Amazon QuickSight 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 Amazon QuickSight.

1. Open the [QuickSight console](#).
2. Choose **Manage QuickSight**, 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 QuickSight.

Deactivating user accounts

Deactivating a QuickSight group or user account removes that group or user's access to Amazon QuickSight resources, like analyses or data sets. IAM Identity Center or Active Directory users that are removed from a group that grants them access to QuickSight lose access to QuickSight. These users appear in the **Inactive users** list in QuickSight 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 QuickSight user's account, put the user into a group with access to Amazon QuickSight. Doing this restores their access to Amazon QuickSight and to any existing resources that are still associated with that user.

Note

With IAM Identity Center integrated into your QuickSight 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 QuickSight role. If a user is in multiple groups that are mapped to different QuickSight role types, the user is able to access QuickSight 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 QuickSight.

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

QuickSight. You can also perform this task by adding a new group to a role in QuickSight. To do this, you need both administrative privileges in Amazon QuickSight 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 QuickSight role. If a user belongs to multiple groups that are mapped to different role types, the user is able to access QuickSight with the role that offers the broadest level of access.

When you make changes to users or groups in Amazon QuickSight, 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 QuickSight, the user no longer exists in QuickSight. You do not need to delete the user in the QuickSight application. The deleted user will appear in the **Inactive users** list in QuickSight until the first day of the following month. After that date passes, the user is automatically removed from the list.

Managing access for QuickSight and IAM users

Inviting users to access Amazon QuickSight

Applies to: Enterprise Edition and Standard Edition

Intended audience: Amazon QuickSight administrators

Use the following procedure to invite a user to access Amazon QuickSight.

1. Choose your user name on the application bar and then choose **Manage QuickSight**.
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 Amazon QuickSight. 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.

The image below shows the **Invite users to this account** table.

Invite users to this account ×

Enter email addresses separated by commas (,)

Username	Email	Role	IAM user ⓘ	
reader+pro@amazon.com	reader+pro@amazon.com	Reader Pro ▾	No ▾	🗑️

Reader Pro
Author Pro
Admin Pro
Reader
Author
Admin

Close Invite

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 QuickSight 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 QuickSight Admin and utilize applicable QuickSight Generative BI capabilities. For more information about Pro roles in QuickSight, see [???](#).

There are some differences in the administrative tasks that IAM users and Amazon QuickSight administrators can perform. These differences occur because some administrative tasks require permissions in Amazon, which Amazon QuickSight–only users lack. The differences are these:

- QuickSight 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 QuickSight permissions to Amazon resources, upgrade to Enterprise edition, and unsubscribe from Amazon QuickSight.

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 QuickSight resources. For more information about what statements are required, see [IAM policy examples for Amazon QuickSight](#).

- **AUTHOR roles:**

- **AUTHOR**– The user is able to author analyses and dashboards in Amazon QuickSight but not perform any administrative tasks in QuickSight.
- **AUTHOR PRO**– The user is able to perform all actions of a QuickSight Author and utilize applicable QuickSight Generative BI capabilities. For more information about Pro roles in QuickSight, see [???](#).

- **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 QuickSight Reader and utilize applicable QuickSight Generative BI capabilities. For more information about Pro roles in QuickSight, see [???](#).

7. For **IAM User**, verify that it says **Yes** for accounts that are associated with IAM users, and **No** for those that are Amazon QuickSight-only.
8. (Optional) To delete a user, choose the delete icon at the end of the relevant row.
9. Choose **Invite**.

Resend an invitation to a user

The sign-up URL in the invitation email expires after 7 days. To resend an invitation to someone, use the following procedure.

1. Choose your user name on the application bar and then choose **Manage QuickSight**.
2. Choose **Manage Users**.
3. Find the entry for the person you want to re-invite, and choose **Resend invitation** for that user.
4. Choose **Confirm**.

Viewing Amazon QuickSight account details

Intended audience: Amazon QuickSight administrators

You can view Amazon QuickSight accounts on the **Manage Users** page. To view a QuickSight user account, use the following procedure.

1. Choose your user name on the application bar and then choose **Manage QuickSight**.
2. Choose **Manage Users** to view details about people who are QuickSight 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 QuickSight 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 QuickSight user account

Intended audience: Amazon QuickSight administrators

Accounts can be deleted by either an Amazon administrator or an Amazon QuickSight administrator. Deleting a QuickSight user account works the same in both the Standard and Enterprise editions of Amazon QuickSight.

Deleting a QuickSight user account removes or transfers their resources. In Enterprise edition, the network administrator can temporarily deactivate a QuickSight user account by removing it from the network group that has access to Amazon QuickSight. If a user is deleted, but not deactivated, that user can still access Amazon QuickSight as a new user. For more information about deactivating an Enterprise account, see [Deactivating user accounts](#).

Use the following procedure to delete a QuickSight user account.

1. Choose your user name on the application bar and then choose **Manage QuickSight**.
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**.

Delete user ×

What would you like to do with the resources that will be orphaned by deleting this user?

This includes analyses, dashboards, data sets and data sources that are shared with others but owned solely by this user.

Transfer ownership of all orphaned resources to a different user in this account

Delete all orphaned resources

Cancel 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**.

Confirm user account deletion ×

Delete user and transfer ownership of all orphaned resources to a different user (author and admin only) in this account

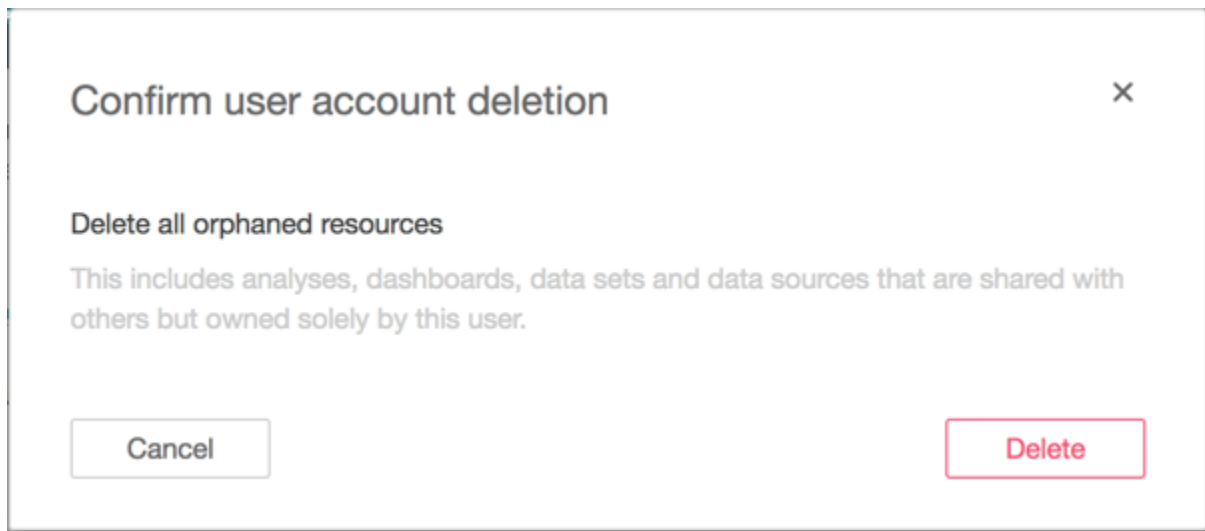
This includes analyses, dashboards, data sets and data sources that are shared with others but owned solely by this user.

Select user to receive all resources

catmi (catmi+qsdocs+66231+isengard@amazon.com)

Cancel 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 QuickSight

Intended audience: System administrators

Applies to: Enterprise Edition

Note

If you're using IAM Identity Center or Active Directory, you can't create and manage groups in Amazon QuickSight. Instead, you manage the assignment of your identity provider's groups to roles in QuickSight.

Admins with IAM credentials who have access to the Amazon QuickSight 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 QuickSight assets with all at once. You can create and manage groups using the QuickSight 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](#).

Creating and managing groups using the Amazon QuickSight console

Use the following procedures to create and manage groups in the Amazon QuickSight console.

To create a user group in the QuickSight console:

1. On the Amazon QuickSight start page, choose **Manage QuickSight**, and then choose **Manage groups**.

[Manage users](#)

Manage groups ←

Your subscriptions

SPICE capacity

Account settings

Security & permissions

Manage VPC connections

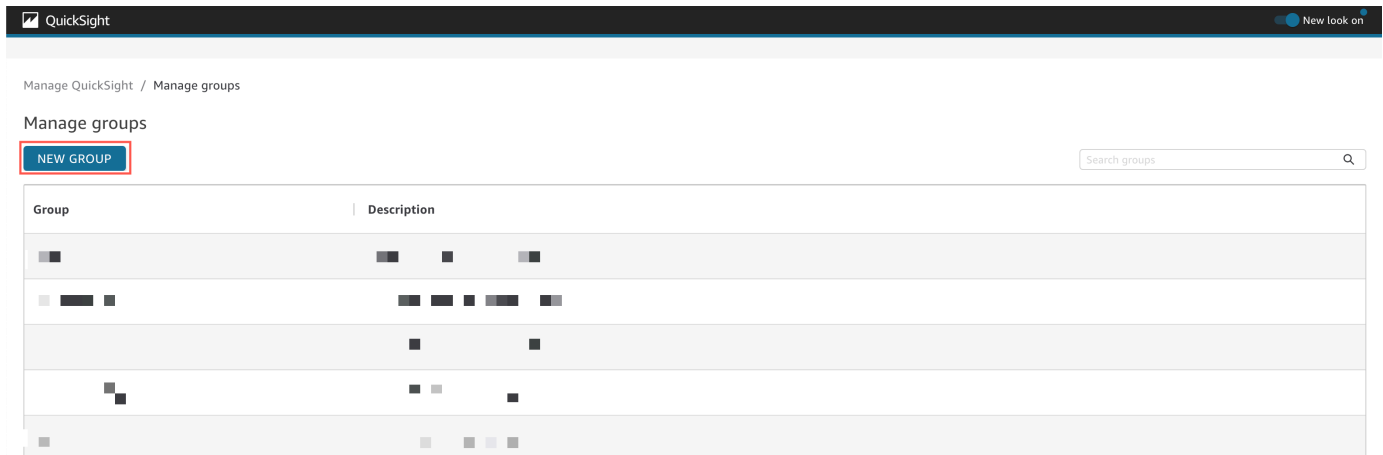
Mobile settings

Domains and Embedding

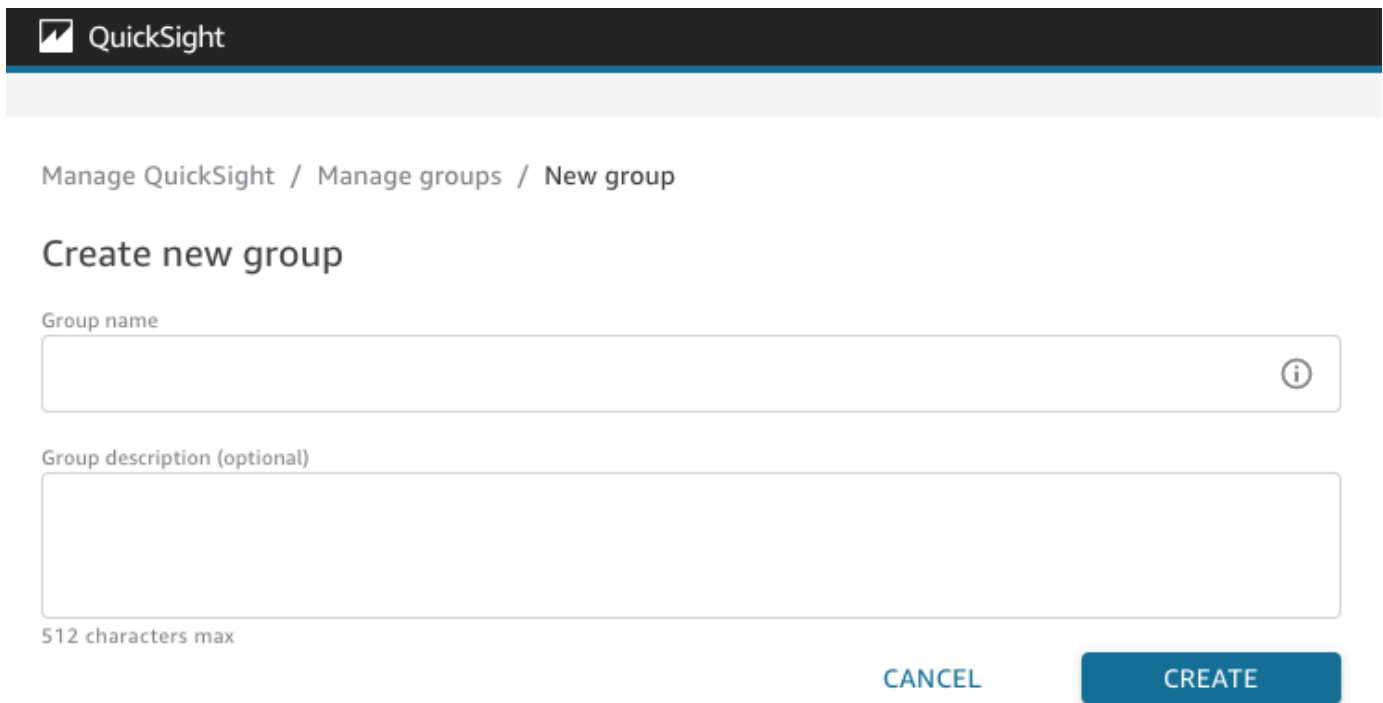
Account customization

Single sign-on (SSO)

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.


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 QuickSight start page, choose **Manage QuickSight**, 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.

Manage QuickSight / Manage groups / Marketing-East

Marketing-East


Regional report for eastern market 

3. In the **Edit description** box that appears, enter the new description and choose **Save**.

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 QuickSight account, see [Managing user access inside Amazon QuickSight](#).

To add a user to a group


1. On the Amazon QuickSight start page, choose **Manage QuickSight**, 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 QuickSight start page, choose **Manage QuickSight**, 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.



You can't delete a group from the QuickSight console, but you can delete a group with the Amazon CLI. For more information on deleting a QuickSight group with the Amazon CLI, see [Deleting groups from Amazon QuickSight](#).

Creating and managing groups using 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 Amazon QuickSight 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 QuickSight 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 QuickSight 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
```

Deleting groups from Amazon QuickSight

You can delete a group from the Amazon CLI. Use the following procedure to delete a Amazon QuickSight user group.

To delete a group in Amazon QuickSight

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 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 QuickSight API reference](#).

Turning on Internet Protocol (IP) and VPC endpoint restrictions in Amazon QuickSight

You can limit access to your organization's Amazon QuickSight 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 Amazon QuickSight 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 QuickSight account only from the VPC that is used for Amazon Direct Connect.

For more information about setting up VPC endpoints in QuickSight, see [Amazon QuickSight 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 Amazon QuickSight console pages can access the IP and VPC endpoint restrictions table.

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 Amazon QuickSight 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.

When you add, change, or delete a rule, a yellow box appears at the top of the table. This box tracks unsaved changes.

To apply changes to the IP and VPC endpoint restrictions table, choose **Save changes** in the box. The changes don't apply to the rules table until you save them. After you choose **Save changes**, it can take up to 10 minutes for a change to take effect.

To add an IP or VPC endpoint rule

1. On the Amazon QuickSight 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.

To update an existing IP or VPC endpoint rule

1. On the Amazon QuickSight 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.

To delete an IP rule

1. On the Amazon QuickSight 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 Amazon QuickSight 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 Amazon QuickSight 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](#).

Customizing access to the Amazon QuickSight console

Applies to: Enterprise Edition

Intended audience: Administrators and Amazon QuickSight developers

In Enterprise edition, you can restrict the functionality that people can access in Amazon QuickSight. Amazon QuickSight custom permissions are applied through IAM policies. You can configure custom permissions for roles (admin, author, reader) for all identity types in QuickSight. You can also apply user level custom permissions to Amazon Identity and Access Management users. *User level custom permissions* override a role's existing default or custom role level permissions for the specified user.

The following limitations apply to user level 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 permissions, you need to be a QuickSight administrator with permissions to use "quicksight:CustomPermissions".

IAM policies and QuickSight permissions are not the same thing. A user can be granted access permissions and assigned a role with an IAM policy, but the IAM policy doesn't control what that user can do within QuickSight. QuickSight assets have their own sets of permissions that are used to customize QuickSight-specific features. These permissions are handled at the resource level outside of IAM.

You can create custom permissions profiles to restrict access to any combination of the following operations.

Asset	Customizable permissions
Data sources and datasets	Create or update data source
	Create or update dataset
	Share dataset
Dashboards and analyses	Add or run anomaly detection
	Create or update theme
	Export to CSV or Excel
	Share
Folders	Create shared folder
	Rename shared folder
Reports	Create
	Update
	Subscribe to email report

Items that are added to shared folders are shared regardless of the asset's custom permissions. This applies to dashboards, analyses, datasets and data sources.

Use the following procedure to create a custom permissions profile in QuickSight.

To create a custom permissions profile

1. From any page in the QuickSight console, choose **Manage QuickSight** at the top right corner.

Only QuickSight administrators have access to the **Manage QuickSight** menu option. If you don't have access to the **Manage QuickSight** menu, contact your QuickSight administrator for assistance.

2. Choose **Security & permissions**.
3. Under **Manage permissions**, choose **Manage**.
4. Choose one of the following options.
 - To edit or view an existing custom permissions profile, choose the ellipsis (three dots) next to the profile that you want, and then choose **View/Edit**.
 - To create a new custom permissions profile, choose **Create**.
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.

Apply a custom permissions profile to a QuickSight role with the QuickSight API

After you create a custom permissions profile, use the QuickSight API 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 QuickSight 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 an IAM user with the QuickSight API

The following example adds custom permissions to a new 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 QuickSight, 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 QuickSight, see [Managing user access inside Amazon QuickSight](#).

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.

Incident response, logging, and monitoring in Amazon QuickSight

Intended audience: System administrators and Amazon QuickSight administrators

Amazon QuickSight is integrated with Amazon CloudTrail, which creates a record of calls from the Amazon QuickSight console and code calls to Amazon QuickSight API operations. For more information, see [Logging operations with Amazon CloudTrail](#).

Amazon QuickSight 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 QuickSight service status can be viewed on the [Service Health Dashboard](#).

Topics

- [Logging operations with Amazon CloudTrail](#)

Logging operations with Amazon CloudTrail

Intended audience: System administrators

Amazon QuickSight is integrated with Amazon CloudTrail. This service provides a record of actions taken by a user, role, or an Amazon service in Amazon QuickSight. CloudTrail captures all API calls for Amazon QuickSight as events. The calls captured include some calls from the Amazon QuickSight console and all code calls to Amazon QuickSight API operations. If you create a trail, you can enable continuous delivery of CloudTrail events to an Amazon S3 bucket, including events for Amazon QuickSight. 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 QuickSight, the IP address from which the request was made, who made the request, when it was made, and additional details.

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

- [Amazon QuickSight Information in CloudTrail](#)
- [Tracking non-API events by using CloudTrail logs](#)
- [Example: Amazon QuickSight log file entries](#)

Amazon QuickSight Information in CloudTrail

Intended audience: System administrators

CloudTrail is enabled on your Amazon account when you create the account. When supported event activity occurs in Amazon QuickSight, 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 QuickSight, 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 QuickSight 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 QuickSight 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 QuickSight)
- **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 QuickSight. 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

Data set

- **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 QuickSight 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"
]
}
}
```

Compliance validation for Amazon QuickSight

Third-party auditors assess the security and compliance of Amazon QuickSight 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 QuickSight 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 QuickSight

Amazon QuickSight 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 QuickSight 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 QuickSight

Intended audience: Amazon QuickSight administrators

Amazon QuickSight is delivered as a web application, hosted on dedicated Amazon EC2 hosts, separate from Amazon virtual private clouds (VPCs). Instead of deploying QuickSight on your own hosts, you access the QuickSight service through Regional public endpoints. QuickSight 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 QuickSight 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 Web Services Regions, websites, IP address ranges, and endpoints](#)
- [Connecting to a VPC with Amazon QuickSight](#)

As a managed service, Amazon QuickSight 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 QuickSight 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 QuickSight does support resource-based access policies, which can include restrictions based on the source IP address. You can also use QuickSight policies to control access from specific Amazon Virtual Private Cloud (Amazon VPC) endpoints or specific VPCs. Effectively, this isolates network access to a given

QuickSight resource from only the specific VPC within the Amazon network. For more information on using QuickSight in a VPC, see [Connecting to a VPC with Amazon QuickSight](#).

Topics

- [Network and database configuration requirements](#)
- [Connecting to a VPC with Amazon QuickSight](#)

Network and database configuration requirements

To serve as data sources, databases need to be configured so that Amazon QuickSight 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](#).

Network configuration requirements

Intended audience: System administrators

For you to use your database server from QuickSight, your server must be accessible from the internet. It must also allow inbound traffic from QuickSight servers.

If the database is on Amazon and in the same Amazon Web Services Region as your QuickSight account, you can auto-discover the instance to make connecting to it easier. To do this, you must grant QuickSight permissions to access it. For more information, see [Accessing data sources](#).

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 QuickSight servers to your Amazon cluster or instance. For further details on how to do this, choose the appropriate topic following:

- [Authorizing connections from Amazon QuickSight to Amazon RDS DB instances](#)
- [Authorizing connections from Amazon QuickSight to Amazon Redshift clusters](#)
- [Authorizing connections from Amazon QuickSight 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 QuickSight servers. These ACLs must do the following:
 - Allow inbound traffic from the appropriate QuickSight 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 QuickSight IP address range and all ports.

For more information about QuickSight IP address ranges, see [IP address ranges for QuickSight](#) following.

For more information about configuring VPC ACLs, see [Network ACLs](#).

- Security group rules that allow traffic between the cluster or instance and QuickSight servers. For further details on how to create appropriate security group rules, see [Authorizing connections to Amazon data stores](#).

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 QuickSight. For more information, see [Connecting to a VPC with Amazon QuickSight](#).

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 QuickSight servers. For further details on how to do this, choose the appropriate topic following:

- [Authorizing connections from Amazon QuickSight to Amazon RDS DB instances](#)
- [Authorizing connections from Amazon QuickSight to Amazon Redshift clusters](#)
- [Authorizing connections from Amazon QuickSight 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). QuickSight doesn't accept certificates that are self-signed or issued from a nonpublic CA. For more information, see [QuickSight 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 QuickSight IP address range. For more information about QuickSight IP address ranges, see [IP address ranges for QuickSight](#). For any other steps that you need to take to enable internet connectivity, see your operating system documentation.

QuickSight 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
- AddTrust Class 1 CA Root
- AddTrust External CA Root
- AddTrust Qualified CA Root
- AffirmTrust Commercial
- QuoVadis Root CA 2
- QuoVadis Root CA 3
- QuoVadis Root Certification Authority
- SecureTrust CA

- AffirmTrust Networking
- AffirmTrust Premium
- AffirmTrust Premium ECC
- America Online Root Certification Authority 1
- America Online Root Certification Authority 2
- Baltimore CyberTrust Code Signing Root
- Baltimore CyberTrust Root
- Bypass Class 2 Root CA
- Bypass 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
- Sonera Class1 CA
- Sonera Class2 CA
- Starfield Root Certificate Authority - G2
- Starfield Services Root Certificate Authority - G2
- 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

- 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
- GeoTrust Primary Certification Authority - G2
- GeoTrust Primary Certification Authority - G3
- GeoTrust Universal CA
- Global Chambersign Root - 2008
- GlobalSign
- GlobalSign Root CA
- 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
- VeriSign Universal Root Certification Authority
- XRamp Global Certification Authority

- Go Daddy Root Certificate Authority - G2
- GTE CyberTrust Global Root
- KEYNECTIS ROOT CA

IP address ranges for QuickSight

For more information on the IP address ranges for QuickSight in supported Regions, see [Amazon Web Services Regions, websites, IP address ranges, and endpoints](#).

Database configuration requirements for self-administered instances

Intended audience: System administrators and Amazon QuickSight administrators

For a database to be accessible to QuickSight, 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.

Connecting to a VPC with Amazon QuickSight

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon QuickSight administrators

Amazon QuickSight 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 QuickSight, you're adding elastic network interfaces in your VPC. These network interfaces allow QuickSight 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 QuickSight in the same way that they apply to traffic between other instances in your VPC.

When you register a VPC connection with QuickSight, 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 operations with Amazon CloudTrail](#). You can even create a dashboard to help you analyze your CloudTrail logs. By combining QuickSight 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 QuickSight, because QuickSight 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](#).

If you're a system administrator – we recommend that you focus on [Setting up a VPC to use with Amazon QuickSight](#) and [Finding information to connect to a VPC](#). The sections after that deal with setting up the connections in QuickSight and testing them.

If you're a QuickSight administrator – if you have the information that you need to configure a VPC connection in the QuickSight console, focus on [Configuring the VPC connection in Amazon QuickSight](#) and [Testing the connection to your VPC data source](#).

Topics

- [VPC terminology](#)
- [Supported VPC data sources](#)
- [Setting up a VPC to use with Amazon QuickSight](#)
- [Finding information to connect to a VPC](#)
- [Configuring the VPC connection in Amazon QuickSight](#)
- [Testing the connection to your VPC data source](#)

VPC terminology

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

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 QuickSight is called a *QuickSight network interface*. Each network interface in a VPC connection is configured based on the subnet it's attached to. You can view your QuickSight 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 QuickSight 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 QuickSight, you specify the security group ID to use.
- An optional description. We recommend that you add the word **QuickSight** to the description for QuickSight 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 QuickSight 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 QuickSight VPC connections.

Supported VPC data sources

The following data sources can connect to QuickSight 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 QuickSight, 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 QuickSight.

4. You registered the VPC with QuickSight by creating or using a VPC connection with the QuickSight console.

Setting up a VPC to use with Amazon QuickSight

Applies to: Enterprise Edition

Intended audience: System administrators

To set up a VPC to use with Amazon QuickSight 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 QuickSight. 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 QuickSight, 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 QuickSight. To set up a successful VPC connection, make sure that the following components work together to allow network traffic to pass between QuickSight and your data source:

- The Amazon VPC service
- The subnets that your data source is using
- The QuickSight 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 QuickSight network interface's security group).
 - Security group attached to the QuickSight 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 QuickSight VPC connection. The network interface for QuickSight that is automatically created during setup is called the *QuickSight 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](#)
- [QuickSight 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 QuickSight 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 QuickSight network interface. QuickSight determines which subnet to route traffic through on the backend. If the availability zone that the subnet is attached to experiences an outage, QuickSight 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 QuickSight 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 QuickSight. 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 Amazon QuickSight to successfully connect to an instance in your VPC, configure your security group rules to allow traffic between the QuickSight 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 QuickSight is connecting to
- From one of the following options:
 - The security group ID that's associated with QuickSight network interface (recommended)
 - or
 - The private IP address of the QuickSight network interface

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

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 QuickSight 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 QuickSight network interface security group doesn't automatically allow return traffic. Because of this, adding an egress rule to the QuickSight network interface security group doesn't work. To make it work for the QuickSight 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 QuickSight 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 QuickSight to connect to any instance in the VPC, you can configure the QuickSight 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 QuickSight 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 QuickSight, 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 QuickSight. 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 QuickSight 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 QuickSight network interface should have outbound rules that allow traffic to each of the database instances in your VPC that you want QuickSight to connect to. To restrict QuickSight 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 RDS for MySQL

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

QuickSight Network interface security group: inbound rule

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

QuickSight Network interface security group: outbound rule

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

RDS MySQL: inbound rule

Type	MYSQL/Aurora
Protocol	TCP
Port Range	3306

Source	sg-ENI3333333
Description	QuickSight to RDS MySQL

VPC connection rules: Amazon Redshift

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

QuickSight network interface security group: inbound rule

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

QuickSight network interface security group: outbound rule

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

Amazon Redshift: inbound rule

Type	Amazon Redshift
------	-----------------

Protocol	TCP
Port Range	5439
Source	sg-ENI3333333
Description	QuickSight–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 QuickSight. 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"
```

QuickSight elastic network interface

The *QuickSight elastic network interface* is a logical networking component in a VPC that represents a virtual network card. QuickSight 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 QuickSight data source you create. The QuickSight network interface alone doesn't give QuickSight direct access to your databases. The VPC connection works only for the QuickSight data sources that are configured to use it.

When you use the QuickSight data source to query a database or other instance within your VPC, all the network traffic from QuickSight originates from this QuickSight network interface. Because the QuickSight network interface exists inside your VPC, traffic originating from it can reach destinations within your VPC by using their private IP addresses. Each QuickSight 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.

QuickSight 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 QuickSight, 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 QuickSight 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 Amazon QuickSight. For each of these, note the database's private IP, security group, and subnets. QuickSight 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 QuickSight 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 [Setting up a VPC to use with Amazon QuickSight](#).

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 Amazon QuickSight in the same Amazon Web Services Region with the VPC.

You can't use QuickSight 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 QuickSight, you can see your current Amazon Web Services Region displayed at the upper right of the QuickSight home screen. You can change the Amazon Web Services Region you're using in QuickSight by changing the Region at the upper right of the QuickSight home screen. All the people who plan to use the data in the VPC must be using the same Amazon Web Services Region in QuickSight.

Note

The Amazon Web Services Region that displays in the QuickSight console doesn't have to match your Amazon CLI configuration. Take care not to mistake your current QuickSight 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 QuickSight 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 [Setting up a VPC to use with Amazon QuickSight](#). 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. QuickSight creates its QuickSight elastic network interface (QuickSight network interface) for the subnets that you choose. The QuickSight 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 the VPC connection in Amazon QuickSight

Applies to: Enterprise Edition

Intended audience: System administrators and Amazon QuickSight administrators

With QuickSight Enterprise Edition, account admins can configure a secure, private VPC connection to a QuickSight account from the QuickSight console or from the QuickSight CLI. Read the following walkthroughs on how to create, edit, and delete a VPC connection from a QuickSight account.

Topics

- [Configuring the VPC connection in the QuickSight console](#)
- [Configuring the VPC connection with the QuickSight CLI](#)

Configuring the VPC connection in the QuickSight console

To create a secure private connection to the Amazon VPC service from the Amazon QuickSight console, use the following procedure.

Prerequisites

- Sign in to QuickSight as a QuickSight admin to set up a VPC connection in QuickSight. To verify that you're a QuickSight administrator, choose your profile icon in the upper-right. If your profile menu contains the option **Manage QuickSight**, then you're a QuickSight 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.

```
{
  "Version": "2012-10-17",
  "Statement": [{
    "Effect": "Allow",
    "Action": [
      "iam:PassRole"
    ],
    "Resource": "arn:aws-cn:iam::account-id: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 QuickSight 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:

```
{
  "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 QuickSight 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 QuickSight access to the role:

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

- Subnet IDs – The IDs of the subnets that the QuickSight 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 Amazon QuickSight Enterprise edition

1. In QuickSight, choose your profile icon in the upper-right, then choose **Manage QuickSight**.

Only QuickSight administrators can view the **Manage QuickSight** option. If you don't see this option on your profile menu, you're not an administrator. In this case, contact your QuickSight 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 QuickSight 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 QuickSight 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 QuickSight. 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**.

QuickSight

Manage QuickSight / Manage VPC connections / Add VPC Connection

Add VPC Connection

Securely connect your data to QuickSight using a Virtual Private Cloud (VPC) connection. [Learn more](#)

AWS console links

- [VPC](#)
- [Subnet](#)
- [Security group](#)
- [DNS resolvers](#)
- [IAM console](#)

VPC connection name

Configuration name in Quicksight

VPC ID

This can not be changed later.

Execution role

Subnets (Select at least two)

Availability Zone	Subnet ID
No rows	

Security Group IDs

DNS resolver endpoints (optional)

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 QuickSight 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 QuickSight 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 QuickSight VPC Connection** pop-up that appears, confirm that you want to delete the connection, and then choose **Delete**.

VPC connection name	VPC ID	VPC connection ARN	Security Group IDs	DNS resolvers	Status	Actions
		arn:aws:quicksight:us-west-			AVAILABLE	Edit
		arn:aws:quicksight:us-west-			AVAILABLE	Delete

Configuring the VPC connection with the QuickSight CLI

To create a secure private connection to the Amazon VPC service from Amazon QuickSight with the QuickSight 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 QuickSight 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:

```
{
  "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 QuickSight 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 QuickSight access to the role:

```

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

```

- Subnet IDs – The IDs of the subnets that the QuickSight 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 QuickSight 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 QuickSight 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 QuickSight administrators and authors

To test whether you can connect to your data source through an existing Amazon QuickSight 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 QuickSight 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 QuickSight](#).

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 QuickSight administrator to

verify that the VPC connection is correct in QuickSight. 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 QuickSight 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 QuickSight administrator to verify the settings in the QuickSight VPC connection.
 - Contact your Amazon administrator to verify that the VPC is correctly configured for use with QuickSight.
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.

Best practices for security in Amazon QuickSight

Amazon QuickSight 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 QuickSight, allow access to HTTPS and WebSockets Secure (wss://) protocol. To allow Amazon QuickSight to reach a database that is on a non-Amazon server, change that server's firewall configuration to accept traffic from the applicable Amazon QuickSight IP address range.

SSL – Use SSL to connect to your databases, especially if you are using public networks. Using SSL with Amazon QuickSight requires the use of certificates signed by a publicly-recognized certificate authority (CA).

Enhanced security – Use Amazon QuickSight 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 QuickSight 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 QuickSight

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*.

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 QuickSight. 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 QuickSight also attaches the policy to a service role that allows Amazon QuickSight 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 QuickSight.

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 QuickSight.

- `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 QuickSight.
- `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 QuickSight.
- `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 QuickSight.

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 QuickSight. `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 QuickSight also attaches this policy to a service role that allows Amazon QuickSight to perform actions on your behalf.

This policy grants read-only permissions that allow access to OpenSearch resources from Amazon QuickSight.

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 QuickSight.
- `es` – Allows principals to use `es:ListDomainNames` to list your OpenSearch domains. This is required to initiate access of Amazon OpenSearch Service from QuickSight.

- `es` – Allows principals to use `es:DescribeElasticsearchDomain` and `es:DescribeDomain` to search your OpenSearch domains. This is required to use Amazon OpenSearch Service from QuickSight.
- `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 QuickSight.

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 resources from Amazon QuickSight.

You can attach `AWSQuickSightSageMakerPolicy` to your IAM entities. Amazon QuickSight also attaches this policy to a service role that allows Amazon QuickSight to perform actions on your behalf.

This policy grants read-only permissions that allow access to Amazon SageMaker resources from Amazon QuickSight.

To view the `AWSQuickSightSageMakerPolicy`, see [AWSQuickSightSageMakerPolicy](#) in the [Amazon Managed Policy reference](#).

Permissions details

This policy includes the following permissions:

- `sagemaker` – Allows principals to use `sagemaker:DescribeModel` to access predictive models in `.`. This is required to support automatic loading of predictive model schemas shared from `.`
- `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 default buckets. This is required to load models shared from Amazon SageMaker Canvas to the default Amazon SageMaker 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 QuickSight to Amazon SageMaker Canvas to build predictive models.

- `s3` – Allows principals to use `s3:ListBucket` to allow Amazon QuickSight to validate an existing Amazon SageMaker Canvas bucket in Amazon S3. This is required to allow the export of data from Amazon QuickSight to Amazon SageMaker Canvas to build predictive models.
- `s3` – Allows principals to use `s3:GetObject` on all Amazon QuickSight– owned Amazon S3 buckets that start with the prefix `arn:aws-cn:s3:::quicksight-m1`. This is required to allow Amazon QuickSight to access the predictions that are generated by Amazon SageMaker Canvas. The generated predictions can be appended to a Amazon QuickSight dataset.
- `sagemaker` – Allows principals to use `sagemaker:CreateTransformJob`, `sagemaker:DescribeTransformJob`, and `sagemaker:StopTransformJob` to perform SageMaker transform jobs on your behalf. This is required for Amazon QuickSight to request predictions from SageMaker models that can be appended to a Amazon QuickSight dataset.
- `sagemaker` – Allows principals to use `sagemaker:ListModelJobs` to list your SageMaker models. This is required to allow generated SageMaker models to appear in Amazon QuickSight.

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 QuickSight 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 QuickSight assets and their corresponding permissions.
- `quicksight` – Allows principals to use `quicksight:ListTagsForResource` to fetch tags of QuickSight 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 QuickSight 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 QuickSight assets and their permissions.
- `quicksight` – Allows principals to use `quicksight:Create*` and `quicksight:Update*` to make changes to the QuickSight 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 QuickSight 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 QuickSight updates to Amazon managed policies

View details about updates to Amazon managed policies for Amazon QuickSight since this service began tracking these changes. For automatic alerts about changes to this page, subscribe to the RSS feed on the [Amazon QuickSight Document History](#) page.

Change	Description	Date
AWSQuickSightAssetBundleExportPolicy – New policy	Amazon QuickSight added new permissions to simplify Asset bundle export operations.	March 27, 2024
AWSQuickSightAssetBundleImportPolicy – New policy	Amazon QuickSight added new permissions to simplify Asset bundle import operations.	March 27, 2024
AWSQuickSightSageMakerPolicy – Update to an existing policy	Amazon QuickSight added new permissions to allow integration with Amazon SageMaker Canvas.	July 25, 2023
AWSQuickSightElasticsearchPolicy – Update to an existing policy	Amazon QuickSight added new permissions to provide access to Amazon OpenSearch Service resources.	September 08, 2021
AWSQuickSightOpenSearchPolicy – New policy	Amazon QuickSight added a new policy to allow access to Amazon OpenSearch Service resources from Amazon QuickSight.	September 08, 2021
Amazon QuickSight started tracking changes	Amazon QuickSight started tracking changes for its Amazon managed policies.	August 2, 2021

Amazon Glossary

For the latest Amazon terminology, see the [Amazon glossary](#) in the *Amazon Web Services Glossary Reference*.

Document history for the Amazon QuickSight User Guide

This page describes changes to the *Amazon QuickSight 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 Amazon QuickSight, see [What's New with Analytics](#). To find out about changes to the Amazon QuickSight service, sign up for the [Amazon QuickSight newsletter](#).

New Amazon QuickSight 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 QuickSight User Guide* release since March 4, 2019.

Change	Description	Date
Embed the Generative Q&A experience	You can now set up an embedded that uses enhanced NLQ capabilities powered by LLMs. This 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 QuickSight Generative Q&A experience..	April 30, 2024
Amazon Q Business in Amazon QuickSight: New Pro roles	Amazon Q Business integrates with Amazon QuickSight to give QuickSight users access to a suite of new	April 30, 2024

generative BI capabilities. With Amazon Q in QuickSight, 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](#).

[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 QuickSight .	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, websites, 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 QuickSight, 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 QuickSight](#).

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 Canvas integration](#)

QuickSight authors can now export data into SageMaker 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 Canvas](#).

October 6, 2023

Managed policy update	The AWSQuickSightSageMakerPolicy is updated to reflect the Amazon QuickSight–SageMaker 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 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 datapoint 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 datasources 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 colocated 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 Amazon 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 Amazon 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 Amazon 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 Amazon QuickSight Q](#).

June 2, 2022

[What's New in Amazon QuickSight](#)

To find out what's new in Amazon 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 Amazon QuickSight .	January 27, 2022

[Sharing dashboards](#)

You now have the option to grant everyone in your Amazon 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 Amazon 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	Amazon 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	Amazon 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 Amazon 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 Amazon QuickSight. For more information, see Using service control policies to restrict Amazon 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](#)

Amazon 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](#)

Amazon 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](#)

Amazon 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](#)

Amazon QuickSight now supports sorting values in pivot 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, Amazon 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, Amazon 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 Amazon 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, Amazon 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](#)

Amazon 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](#)

Amazon 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 [QuickSight 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](#)

Amazon QuickSight now supports connecting to Amazon OpenSearch. For more information, see [Using Amazon OpenSearch with Amazon QuickSight](#).

November 25, 2020

[Format metrics on pie charts](#)

Amazon 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 Amazon 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](#)

Amazon 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](#)

Amazon 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 Amazon 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	Amazon 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	Amazon 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	Amazon 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	Amazon 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	Amazon 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 Amazon QuickSight enterprise edition .	October 30, 2020
New chart type: Funnel charts	Amazon 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	Amazon 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 Amazon QuickSight - Amazon GovCloud (US) .	October 28, 2020

New data source: Oracle	Amazon QuickSight now supports connecting to Oracle. For more information, see Supported data sources .	October 23, 2020
Filter across datasets	In Amazon 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	Amazon 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	Amazon 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	Amazon 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	Amazon 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	Amazon QuickSight now supports duplicating sheets. For more information, see Working with multiple sheets in an Amazon QuickSight analysis .	September 14, 2020
Customize labels of totals and subtotals	In Amazon 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 Amazon 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	Amazon 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	Amazon QuickSight now supports using reference lines in bar, line, and combo charts. For more information, see Reference lines .	September 3, 2020
New table calculations	Amazon 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	Amazon QuickSight now supports dynamic default values for multi-valued parameters. For more information, see Creating parameter defaults in Amazon QuickSight .	August 25, 2020
Last active status for users	Amazon 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](#)

Amazon 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 Amazon 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](#)

Amazon 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 Amazon QuickSight assets into folders](#).

August 5, 2020

[Embed the Amazon QuickSight console](#)

In Amazon 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 Amazon QuickSight console](#).

July 23, 2020

[Datasets with 2,000 columns](#)

Amazon QuickSight now supports datasets that have up to 2,000 columns. For more information, see [Data source limits](#).

July 23, 2020

[Customized permissions](#)

In Amazon QuickSight Enterprise edition, you can now create custom roles to restrict a person from using specific functionality in the Amazon 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 Amazon QuickSight console](#).

July 23, 2020

[Customize the console](#)

In Amazon 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 Amazon QuickSight console](#).

July 23, 2020

[Amazon 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

[Larger SPICE datasets](#)

Amazon 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 Amazon 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](#)

Amazon 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 Amazon 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 Amazon QuickSight. For more information, see [Using themes in Amazon QuickSight](#).

July 8, 2020

Amazon Lake Formation	If you use Amazon 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	Amazon 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 Amazon QuickSight. For more information, see Adding conditional formatting to visuals .	June 12, 2020
Five new languages	Amazon QuickSight now supports Danish, Dutch, Finnish, Norwegian, and Swedish. These languages expand upon the existing 15 languages already available in Amazon 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](#)

Amazon 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 integration now generally available](#)

The integration of SageMaker with Amazon 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 together makes 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 Amazon 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 [Amazon QuickSight VPC connections](#).

May 18, 2020

[Use logarithmic scale with Amazon QuickSight](#)

You can now format visuals using log scale. For more information, see [Changing the visual scale](#).

May 1, 2020

[Severity levels for Amazon 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

[Improvements to visual controls in Amazon 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

[Amazon 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

[Amazon QuickSight activates legends by adding menus for quick feature access](#)

Amazon 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

[Amazon 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 Amazon QuickSight ML-powered anomaly detection](#)

Now you can control the minimum delta value that Amazon 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 Amazon QuickSight narrative expression editor](#)

Amazon 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 integration now available in more Amazon Web Services Regions](#)

The public preview for integrating SageMaker with Amazon 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

[Private VPC support for presto in Amazon QuickSight](#)

Now Presto can use a VPC connection. For more information, see [Supported data sources for VPC](#).

December 10, 2019

[New mathematical functions in Amazon 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 Amazon QuickSight](#)

Now you can use `rank`, `denseRank`, and `percentileRank` for level-aware aggregations. For more information, see [Using level-aware aggregations](#).

December 10, 2019

[Amazon 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

[Integrate SageMaker ML models into your Amazon QuickSight analyses and dashboards](#)

In this public preview, Amazon QuickSight launches the ability to integrate machine learning (ML) models created and trained in SageMaker. 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 Amazon 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

[New Amazon 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 QuickSight 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

[Amazon QuickSight supports the now function in SPICE](#)

For more information, see [now](#).

November 22, 2019

[Amazon 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

[Amazon 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 Amazon QuickSight](#).

November 22, 2019

[Amazon 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 (<code>toString</code> and <code>parseDecimal</code>) and two new date functions (<code>parseDate</code> and <code>formatDate</code>). 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	Amazon 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](#)

Amazon 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 dashboard](#).

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 Amazon 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

[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 Amazon 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

[Amazon QuickSight adds custom colors for charts.](#)

With the Amazon 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

[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 <i>table</i> 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 QuickSight User Guide* before March 4, 2019.

Change	Description	Date changed
New features	<p>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.</p> <p>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 reports by email.</p> <p>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.</p> <p>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.</p>	March 4, 2019

Change	Description	Date changed
New features	<p>There are new table calculations available:</p> <ul style="list-style-type: none">• percentileOver• runningAvg• runningCount• runningMax• runningMin	February 7, 2019
New features	<p>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:</p> <ul style="list-style-type: none">• Scheduling and sending reports by email• Using a control with a parameter in Amazon QuickSight• Adding date filters	January 22, 2019

Change	Description	Date changed
New features	<p>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> .</p> <p>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.</p> <p>For more information, see the following:</p> <ul style="list-style-type: none"> • Aggregate functions • Using pivot tables 	January 10, 2019
New feature	<p>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.</p> <p>For more information, see Joining data.</p>	January 3, 2019
New SDK	<p>You can embed dashboards and manage users or groups by using the Amazon QuickSight SDKs. For more information, see Developing with Amazon QuickSight.</p>	November 27, 2018
New features	<p>You can use groups with row-level security (RLS), and you can add cascading controls for parameters. For more information, see the following:</p> <ul style="list-style-type: none"> • Using row-level security (RLS) 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	<p>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.</p> <p>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.</p>	November 1, 2018
New feature	Use JSON native data types with <code>parseJson</code> . For more information see, parseJson .	October 30, 2018
New features	<p>Use date functions to find out what quarter a date is in:</p> <ul style="list-style-type: none"> • addDateTime • Extract • truncDate <p>Add parameters to URLs. For more information see, Using parameters in a URL.</p>	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	<p>You can schedule emailed reports, and add data labels to your visuals.</p> <p>For more information on these features, see the following sections:</p> <ul style="list-style-type: none">• Scheduling and sending reports by email• Subscribing to email reports in Amazon QuickSight• Data labels on visual types in QuickSight	August 15, 2018
New features	<p>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.</p> <p>For more information on these features, see the following sections:</p> <ul style="list-style-type: none">• Table calculations• Changing date field granularity	August 8, 2018

Change	Description	Date changed
New features	<p>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.</p> <p>For more information on these features, see the following sections:</p> <ul style="list-style-type: none">• Replacing datasets• Labels on visual types in QuickSight• Customizing a field format• Visualizing data in Amazon QuickSight	June 21, 2018

Change	Description	Date changed
New features and a new Amazon Web Services Region	<p>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.</p> <p>In both editions, Amazon QuickSight supports parameters with on-sheet controls, dashboard co-ownership, custom URL actions, and 25-GB SPICE datasets.</p> <p>Also, Amazon QuickSight is available in Asia Pacific (Tokyo).</p> <p>For more information on these features, see the following sections:</p> <ul style="list-style-type: none">• 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 Amazon QuickSight dashboards• Connecting to a VPC with Amazon QuickSight• Refreshing a dataset on a schedule• Parameters in Amazon QuickSight• 	May 30, 2018

Change	Description	Date changed
	<p>Reader experience: Exploring interactive dashboards in Amazon QuickSight (newly updated for read-only users)</p> <ul style="list-style-type: none"> • Using custom actions for filtering and navigating • Data source quotas • Amazon Web Services Regions, websites, IP address ranges, and endpoints 	
New feature	<p>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:</p> <ul style="list-style-type: none"> • Sorting visual data in Amazon QuickSight • Duplicating datasets • Totals and subtotals • Creating a basic SQL query 	May 25, 2018
New features	<p>You can use Amazon QuickSight to connect to Software as a Service (SaaS) providers. For more information, see Software as a service (SaaS) data.</p> <p>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.</p>	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 (RLS) 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 Amazon 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 Amazon 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 Amazon 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	<ul style="list-style-type: none"> 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 datasets. 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>Amazon QuickSight User Guide</i> .	November 15, 2016

Attributions

Following, you can find out about attributions and third-party licenses for material used in Amazon QuickSight 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 zeměměřičský 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, Licensor waives 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 Licensor 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 Licensor You must, to the extent practicable, remove from the Collective Work any reference to such Licensor or the Original Author, as requested. If You create a Derivative Work, upon notice from any Licensor You must, to the extent practicable, remove from the Derivative Work any reference to such Licensor 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 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 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, MERCHANTABILITY, 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, MERCHANTABILITY, 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 Licensor 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 Licensor 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 Licensor 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 Licensor – Licensed Material. Every recipient of the Licensed Material automatically receives an offer from the Licensor 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.