

SAP Guides

General SAP Guides





Table of Contents

Home	1
Overview	2
	2
Amazon overview	2
Amazon services	3
Amazon global infrastructure	7
Amazon security and compliance	8
Amazon resource provisioning and management	9
SAP on Amazon overview	10
SAP software and licenses on Amazon	10
SAP support on Amazon	11
Deploying SAP systems on Amazon	12
Partner services	13
SAP on Amazon planning	14
SAP notes	14
SAP on Amazon architectures	15
Choosing an Amazon Region and Availability Zone	17
Network and connectivity	18
Following security best practices	8
Amazon EC2 instance types for SAP	20
Operating systems	21
Databases	24
SAP installation media	26
SAProuter and SAP Solution Manager	26
Document revisions	27
Amazon EC2 instance types	28
Instance type availability	28
SAP NetWeaver supported instances	29
Current generation Amazon EC2 instances for SAP NetWeaver	29
Previous generation Amazon EC2 instances for SAP NetWeaver	53
SAP HANA certified and non-certified instances	55
Current generation certified Amazon EC2 instances for SAP HANA	55
Previous generation certified Amazon EC2 instances for SAP HANA	61
Non-certified Amazon EC2 instances for SAP HANA	62

SAP Business One certified instances, version for SAP HANA	64
Document history	65
Amazon Data Provider	67
	67
Introduction	67
Pricing	68
Technical requirements	69
Amazon VPC network topologies	69
Amazon VPC Endpoints	71
IAM Roles	72
DataProvider 4.3	78
Installing DataProvider 4.3	79
Updating to DataProvider 4.3	99
Uninstalling older versions	104
Troubleshooting	105
Troubleshooting on Linux	106
Troubleshooting on Windows	110
Customizing the DataProvider	116
Syntax Rules for configuration files	117
User-configurable EC2 instance types	117
User-configurable EBS volume types	120
Verification of monitoring	121
Checking Metrics with SAPOSCOL	121
Checking Metrics with SAP CCMS	122
Example of captured metrics	125
Version history	131
Cost estimation	136
Amazon Region	136
Compute	136
Storage	137
Amazon EBS	137
Amazon EFS	138
Amazon FSx for Windows File Server	138
Amazon FSx for NetApp ONTAP	138
Amazon S3	139
Network	139

Amazon VPC	140
Amazon Site-to-Site VPN	141
Amazon Direct Connect	141
Elastic Load Balancing	141
Data transfer pricing	142
Automation	144
Backup, restore, and recovery	144
Amazon Backint Agent for SAP HANA	144
Amazon Elastic Disaster Recovery	145
Amazon EBS snapshots	145
Migration	145
Migration Hub Orchestrator	145
Amazon DataSync	146
Monitoring	146
Amazon Data Provider	146
Amazon CloudWatch Application Insights for SAP HANA	146
Amazon CloudWatch	147
Amazon CloudTrail	147
VPC Flow Logs	147
Operating System licenses	147
Red Hat	148
SUSE	148
Windows	149
Oracle Enterprise Linux	
Amazon Marketplace	149
Amazon Support	150
Architecture guidance	151
Overview	
Prerequisites	
Specialized knowledge	
Recommended reading	152
Introduction	
SAP NetWeaver architecture single points of failure	
High availability and disaster recovery	
On premises vs. cloud deployment patterns	
Architecture guidelines and decisions	157

	Regions and Availability Zones	. 157
	Amazon accounts	. 159
	Compute	160
	Networking	. 164
	Storage	. 167
	Monitoring and audit	. 171
	Architecture patterns	171
	Failure scenarios	. 172
	Patterns	. 173
	Summary	196
	Microsoft SQL	198
	Patterns	. 198
	Comparison matrix	198
	Single Region patterns	. 199
	Multi-Region patterns	. 202
Di	saster recovery with Elastic Disaster Recovery	209
	Scenarios	210
	References	210
	SLAs and licenses	. 211
	Recovery time objective (RTO)	211
	Recovery point objective (RPO)	. 212
	Recovery consistency objective (RCO)	. 212
	SAP licenses	213
	Network, storage, and compute	213
	Network	214
	Storage	. 216
	Compute	218
	Scenarios	220
	Amazon In-Region disaster recovery	. 220
	Amazon Cross-Region disaster recovery	. 222
	Outside of Amazon to Amazon disaster recovery	224
	Shared storage resiliency	226
	Amazon In-Region disaster recovery	. 227
	Amazon Cross-Region disaster recovery	. 227
	Outside of Amazon to Amazon disaster recovery	228
	Implementation	228

SAP application layer	. 229
SAP database layer	229
RISE with SAP	233
Connectivity	234
On-premises to Amazon connectivity	235
Connecting to RISE with SAP VPC using Amazon VPN	235
Connecting to RISE with SAP VPC using Amazon Direct Connect	. 235
Connectivity between Amazon Web Services accounts	236
Connecting to RISE with SAP using your Amazon Web Services account	239
SAP Business Technology Platform (BTP) with RISE with SAP on Amazon	. 240
Security	242
Single Sign-On – SAP Business Technology Platform (BTP) and Amazon IAM Identity	
Center	242
Use Amazon services to enhance security in RISE with SAP	243
Reliability	245
Extensions	247
Performance	247
Data lake	248
Application integration	. 249
Document management	250
Development and extension	251
User front-end	251

General SAP Guides

This section covers the following guides.

- Overview
- Amazon EC2 instance types
- Estimation
- Amazon Data Provider
- Architecture guidance
- Disaster recovery with Amazon Elastic Disaster Recovery
- RISE with SAP on Amazon Cloud

Additional SAP on Amazon documentation

- SAP HANA on Amazon
- SAP NetWeaver on Amazon
- Databases for SAP applications on Amazon
- Amazon Launch Wizard for SAP
- Amazon Systems Manager for SAP
- Amazon SDK for SAP ABAP
- SAP BusinessObjects on Amazon
- Amazon Migration Hub Orchestrator

SAP on Amazon Overview and Planning

SAP specialists, Amazon Web Services

Last updated: January 2023

This guide provides overview and planning information for SAP customers and partners who are considering implementing or migrating SAP environments or systems to the Amazon Web Services Cloud.

This guide is intended for users who have previous experience installing, migrating, and operating SAP environments and systems on traditional on-premises infrastructure. It consists of three main sections:

- An overview of the Amazon Cloud and Amazon services, for readers who are new to the cloud.
- An overview of SAP on Amazon, including software and licenses, support options, and partner services.
- Technical considerations that will help you plan and get the most out of your SAP environment on Amazon.



Note

To access the SAP notes referenced in this guide, you must have an SAP One Support Launchpad user account. For more information, see the SAP Support website.

About this Guide

This guide is part of a content series that provides detailed information about hosting, configuring, and using SAP technologies in the Amazon Cloud. For the other guides in the series, ranging from overviews to advanced topics, see https://aws.amazon.com/sap/docs/.

Amazon Overview

Amazon offers a broad set of global, cloud-based services, including compute, storage, networking, Internet of Things (IoT), and many others. These services help organizations move faster, lower IT costs, and support scalability. Amazon is trusted by the largest enterprises and popular start-ups to

power a wide variety of workloads, such as web and mobile applications, game development, data processing and warehousing, storage, and archiving.

Amazon Services

Amazon provides over 200 cloud services that you can use in combinations tailored to your business or organizational needs. For information about all Amazon services, see the <u>Amazon Web</u> Services Cloud Platform documentation.

This section introduces the Amazon services that are most relevant for the deployment and operation of SAP solutions. The following list provides a high-level description of each service and its use for SAP systems. To view features, pricing, and documentation for an individual service, follow the *details*link after the description.

Area	Service	Description	SAP uses
Compute	Amazon Elastic Compute Cloud (Amazon EC2)	Secure, resizable compute capacity in the cloud. (details)	Virtual and bare metal servers for the installation and operation of SAP systems.
Storage	Amazon Elastic Block Store (Amazon EBS)	Persistent block storage volumes for use with EC2 instances. (details)	File systems for SAP software (e.g., /usr/sap), SAP database log and data files, and SAP local backups.
	Amazon Simple Storage Service (Amazon S3)	Object storage service that offers an extremely durable, highly available, and infinitely scalable data storage infrastru cture. (details)	Highly durable storage for file backups, database backups, archiving data, data lakes, and more.
	Amazon Elastic File System (Amazon EFS)	Simple, scalable, elastic file system	Shared file system for SAP applicati

Area	Service	Description	SAP uses
		for Linux-based workloads for use with Amazon Cloud services and on- premises resources. (details)	on servers (e.g., / sapmnt).
	Amazon FSx for Windows File Server	Fully managed, highly durable, and available native Microsoft Windows file system. (details)	Shared file system for SAP applicati on servers (e.g., / sapmnt).
	Amazon FSx for NetApp ONTAP	Fully managed, highly reliable, scalable, high- performing file storage built on NetApp ONTAP file system(details)	Shared file system for SAP applicati on servers (e.g., / sapmnt).
Networking	Amazon Virtual Private Cloud (Amazon VPC)	Logically isolated section of the Amazon Cloud where you can launch Amazon resources in a virtual network that you define. (details)	Network for SAP resources. You can control the level of isolation of your EC2 instance from other networks, instances, and onpremises network resources, such as those in production and non-production environments.

Area	Service	Description	SAP uses
	Amazon Site-to-Site VPN	Enables you to securely connect your on-premises network or branch office site to your VPC. (details)	Network connectivity between on-premis es systems/users and SAP systems on Amazon.
	Amazon Direct Connect	Lets you establish private network connectivity between Amazon and your data center, office, or co-location environment. (details)	Private network connectivity between on-premises systems/ users and the SAP system or environme nt on Amazon.
	Amazon Route 53	Highly available and scalable cloud Domain Name System (DNS) web service. (details)	Name and address resolution for SAP systems running on Amazon.
	Amazon Time Sync	Highly accurate and reliable time reference that is natively accessible from EC2 instances. (Linux Windows)	Time synchronization for your SAP systems on EC2 instances.
Management and operation tools	Amazon Web Services Management Console	Simple web interface to provision and manage Amazon resources. (details)	Provisioning and management of Amazon resources for your SAP environme nt on Amazon.

Area	Service	Description	SAP uses
	Amazon Command Line Interface (Amazon CLI)	Command-line tool set to provision and manage Amazon resources. (<u>details</u>)	Creation of scripts to automate the provisioning and management of Amazon resources for your SAP environme nt on Amazon.
	Amazon CloudForm ation	An easy way to create a collection of related Amazon resources and provision them in an orderly and predictable fashion. (details)	Automated provision ing of Amazon resources for new SAP landscapes, disaster recovery environments, and other use cases.
	Amazon CloudWatch	Monitoring for Amazon Cloud resources and the applications you run on Amazon: collect and track metrics, collect and monitor log files, and set alarms. (details)	Monitoring SAP systems running on Amazon using Amazon CloudWatch Application Insights.
	Amazon CloudTrail	Records activity made on your account and delivers log files to your S3 bucket. (details)	Audit capabilit ies within your Amazon Web Services account, such as use of the Amazon EC2 API.

Area	Service	Description	SAP uses
	Amazon Launch Wizard for SAP	Amazon Launch Wizard for SAP is a service that guides you through the sizing, configuration, and deployment of SAP applications on Amazon. (details)	Setup and configura tion of resources required for your SAP deployment.
	Amazon Backint Agent for SAP HANA	SAP certified solution to backup and restore SAP HANA database to and from Amazon S3. (details)	Backup solution to store SAP HANA database backups to Amazon S3.
Security, identity, and compliance	Amazon Identity and Access Management (IAM)	Manages access to Amazon services and resources securely. Using IAM, you can create and manage Amazon users and groups, and use permissions to allow and deny their access to Amazon resources. (details)	Fine-grained access control using a least privileged security model to access specific Amazon services and actions; e.g., to allow SAP BASIS resources to launch, to stop and start EC2 instances without terminating them.

Amazon Global Infrastructure

The Amazon Cloud infrastructure is built around Regions and Availability Zones. An Amazon Region is a physical location that provides multiple, physically separated and isolated Availability Zones. Each Availability Zone consists of one or more data centers that are connected with low-latency, high-throughput, and highly redundant networking. These Availability Zones offer an

easier and more effective way to design and operate your applications and databases, making them more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For a list of the available Amazon Regions and to learn more about the Amazon global infrastructure, see Global Infrastructure on the Amazon website.

Amazon Security and Compliance

Security

At Amazon, security is our top priority. As an Amazon customer, you will benefit from a data center and network architecture built to meet the requirements of the most security-sensitive organizations. Security in the cloud is much like security in your on-premises data centers—only without the costs of maintaining facilities and hardware. In the cloud, you don't have to manage physical servers or storage devices. Instead, you use software-based security tools to monitor and protect the flow of information into and out of your cloud resources.

As an Amazon customer you inherit all the best practices of Amazon policies, architecture, and operational processes built to satisfy the requirements of our most security-sensitive customers, and get the flexibility and agility you need in security controls.

The Amazon Cloud enables a shared responsibility model. While Amazon manages security **of** the cloud, you are responsible for security **in** the cloud. This means that you retain control of the security you choose to implement to protect your own data, platform, applications, systems, and networks no differently than you would in an on-site data center.

To learn more about Amazon security, see Amazon Cloud Security on the Amazon website.

Compliance

Amazon provides robust controls to help maintain security and data protection in the cloud. As systems are built on top of Amazon Cloud infrastructure, compliance responsibilities will be shared. By tying together governance-focused, audit-friendly service features with applicable compliance or audit standards, Amazon compliance enablers build on traditional programs and help you operate in an Amazon security control environment.

The IT infrastructure that Amazon provides to its customers is designed and managed in alignment with best security practices and a variety of IT security standards. The following is a partial list of assurance programs with which Amazon complies:

- SOC 1/ISAE 3402, SOC 2, SOC 3
- FISMA, FIPS, DIACAP, and FedRAMP
- PCI DSS Level 1
- ISO 9001, ISO 27001, ISO 27017, ISO 27701, ISO 27018

For more information, see Amazon Compliance Programs.

Amazon Provisioning and Management

The provisioning and management of Amazon services and resources use a self-service model managed by the customer or a partner. For an overview of the tools available for provisioning and management, see the management tools in the Amazon Services section.

Figure 1 shows the services managed by Amazon and the services managed by the customer or partner for SAP.

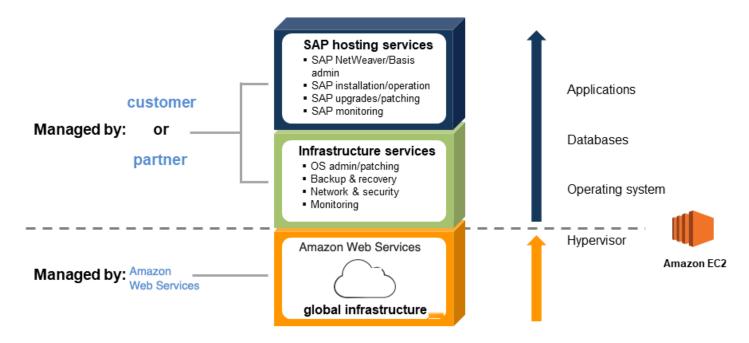


Figure 1: Managed services for SAP on Amazon

SAP on Amazon Overview

Amazon has been working with SAP since 2011 to help customers deploy and migrate their SAP applications to Amazon, and SAP supports running the vast majority of available SAP applications on Amazon.

SAP Software and Licenses on Amazon

This section describes the options available for SAP software and licenses on Amazon.

Bring Your Own Software and License

The majority of SAP solutions that can be run on Amazon use a bring-your-own-software and bring-your-own-license (BYOL) model. Running SAP systems on Amazon doesn't require special or new SAP licenses. If you're an existing SAP customer, you can use your existing SAP licenses when running SAP on Amazon. You are responsible for obtaining a valid SAP license, and you must ensure that you are in compliance with the SAP licensing policies. Amazon does not provide or sell SAP licenses.

Amazon Marketplace

<u>Amazon Marketplace</u> is a digital catalog with thousands of software listings from independent software vendors that makes it easy to find, test, buy, and deploy software that runs on Amazon. To view SAP-related offerings available in Amazon Marketplace, follow this link: <u>SAP in Amazon Marketplace</u>.

SAP Trial and Developer Licenses

The <u>SAP Cloud Appliance Library</u> provides access to an online repository of the latest preconfigured SAP solutions. You can quickly deploy these solutions on Amazon by using a launch wizard that automates deployment. Some of the solutions available in the SAP Cloud Appliance Library are provided with free trial or developer edition licenses.

SAP Hardware Key Generation

SAP hardware key generation on EC2 instances uses a specific process that is dependent on the SAP kernel patch level. If a hardware key is generated before patching the SAP kernel to the proper level, and the kernel is updated at a later time, the hardware key may change, making the installed license invalid. For details on how the SAP hardware ID is generated on EC2 instances and the

SAP on Amazon overview 10

required SAP kernel patch levels see the following SAP notes (SAP One Support Launchpad access required):

- SAP Note 2327159 SAP NetWeaver License Behavior in Virtual and CLoud Environments
- SAP Note 1178686 Linux: Alternative method to generate a SAP hardware key
- SAP Note 2327159 SAP NW License Behavior in Virtual and Cloud Environments
- SAP Note 1697114 Determination of hardware ID in Amazon clouds
- SAP Note 2113263 Additional public key for Amazon Hardware ID
- SAP Note 2823805 Additional public keys for Amazon Hardware ID
- SAP Note 2319387 Adjustment of the license check for Amazon China

SAP Support on Amazon

Amazon and SAP have worked together closely to ensure that you receive the same level of support via the same support channels, whether you're running your SAP systems on Amazon or on premises.

SAP Solutions Supported on Amazon

The majority of SAP solutions that run on traditional on-premises infrastructure are fully supported by SAP on Amazon. For the complete list of SAP solutions supported on Amazon, see <u>SAP Note</u> 1656099 and the other notes referenced within that note.

SAP Support on Amazon

To ensure full support of your SAP on Amazon environment from SAP and Amazon, you must follow the guidelines and requirements in <u>SAP Note 1656250</u>. Here are the primary requirements you must follow to ensure support of your SAP on Amazon environment:

- Enable detailed monitoring for Amazon CloudWatch on each EC2 instance to ensure that the
 required Amazon metrics are provided in one-minute intervals. For additional information on
 Amazon CloudWatch, see https://aws.amazon.com/cloudwatch.
- Install, configure, and run the <u>Amazon Data Provider for SAP</u> on each EC2 instance. The
 Amazon Data Provider collects the required performance and configuration data from a variety
 of sources, including the Amazon EC2 API, Amazon EC2 instance metadata, and Amazon
 CloudWatch, and shares it with SAP applications, to help monitor and improve the performance
 of business transations.

SAP support on Amazon 11

 Any Amazon Web Services account that you use for running SAP systems must have an <u>Amazon</u> support plan for either Business Support or Enterprise Support.

Deploying SAP Systems on Amazon

The section describes different options available for provisioning Amazon infrastructure and installing SAP systems on Amazon.

Manual Deployment

The majority of SAP solutions supported on Amazon can be installed by manually provisioning the required Amazon infrastructure resources and then following the relevant SAP installation document on Amazon.

Automated Deployment

Amazon Launch Wizard for SAP is a service that guides you through the sizing, configuration, and deployment of SAP applications on Amazon. Amazon Launch Wizard reduces the time it takes to deploy SAP applications on Amazon. You input your application requirements, including SAP HANA settings, SAP landscape settings, and deployment details on the service console, and Launch Wizard identifies the appropriate Amazon resources to deploy and run your SAP application.

For more information, see How Amazon Launch Wizard for SAP works.

Prebuilt Images

Some SAP solutions are available on Amazon as a prebuilt system image that contains a preinstalled and preconfigured SAP system. A prebuilt SAP system image enables you to rapidly provision a new SAP system without spending the time and effort required by a traditional manual SAP installation.

Prebuilt SAP system images are available from the following sources:

- Amazon Marketplace
- SAP Cloud Appliance Library

SAP solution	Deployment option(s)
SAP Business Suite (ERP, CRM, etc.)	Manual SAP CAL

SAP solution	Deployment option(s)
SAP NetWeaver	Manual Amazon Launch Wizard for SAP SAP CAL
SAP S/4HANA	Manual Amazon Launch Wizard for SAP <u>SAP</u> <u>CAL</u>
SAP BW/4HANA	Manual SAP CAL
SAP HANA	Manual Amazon Launch Wizard for SAP SAP CAL
SAP BusinessObjects BI	Manual Amazon Marketplace SAP CAL
SAP Commerce (Hybris)	Manual SAP CAL
SAP Business One, version for SAP HANA	Manual SAP CAL
SAP Business One, version for Microsoft SQL Server	Manual

Getting Assistance from APN Partners

There are Amazon Partner Networks (APN) partners who are experienced in deploying and operating SAP solutions, and can help you with your SAP workloads on Amazon. For additional information see the following section.

Partner Services for SAP on Amazon

The <u>Amazon Partner Network (APN)</u> is a community of companies that offer a wide range of services and products on Amazon. APN SAP partners can provide SAP-specific services to help you fully maximize the benefits of running SAP solutions on Amazon.

Types of Partner Services and Solutions for SAP on Amazon

• Cloud assessment services – Advisory services to help you develop an efficient and effective plan for your cloud adoption journey. Typical services include financial/TCO (total cost of ownership), technical, security and compliance, and licensing.

Partner services 13

• **Proof-of-concept services** – Services to help you test SAP on Amazon; for example: SAP ERP/ECC migration to SAP HANA or SAP S/4HANA, SAP Business Warehouse (BW) migration to SAP HANA or SAP BW/4HANA, SAP OS/DB migrations, new SAP solution implementation.

- **Migration services** Services to migrate existing SAP environments or systems to Amazon; for example: all-on-Amazon SAP migrations (PRD/QAS/DEV), hybrid SAP migrations (QAS/DEV), single SAP system (e.g., SAP BW) migrations.
- Managed services Managed services for SAP environments on Amazon, including: Amazon Web Services account and resource administration, OS administration/patching, backup and recovery, SAP Basis and SAP NetWeaver.
- Packaged solutions Bundled software and service offerings from SAP Partners that combine SAP software, licenses, implementation, and managed services on Amazon, such as SAP S/4HANA, SAP BusinessObjects BI, and many others.
- **ISV software solutions** Partner software solutions for the migration, integration, and operation of SAP solutions on Amazon; for example: system migration, high availability, backup and recovery, data replication, automatic scaling, disaster recovery.

How to Find Partner Solutions for SAP on Amazon

The **Amazon SAP Partner Solutions** provides a centralized place to search, discover, and connect with trusted APN partners who offer solutions and services to help your business achieve faster time to value and maximize the benefits of running SAP solutions on Amazon. For more information, see Amazon SAP Competency Partners.

SAP on Amazon Planning

If you are an experienced SAP Basis or SAP NetWeaver administrator, there are a number of Amazon-specific considerations relating to compute configurations, storage, security, management, and monitoring that will help you get the most out of your SAP environment on Amazon. This section provides guidelines for achieving optimal performance, availability, and reliability, and lower total cost of ownership (TCO) while running SAP solutions on Amazon.

SAP Notes

Before migrating or implementing an SAP environment on Amazon, you should read and follow the relevant SAP notes. Start from <u>SAP Note 1656099</u> for general information and follow the links to other relevant SAP notes (SAP One Support Launchpad access required).

SAP on Amazon planning 14

SAP on Amazon Architectures

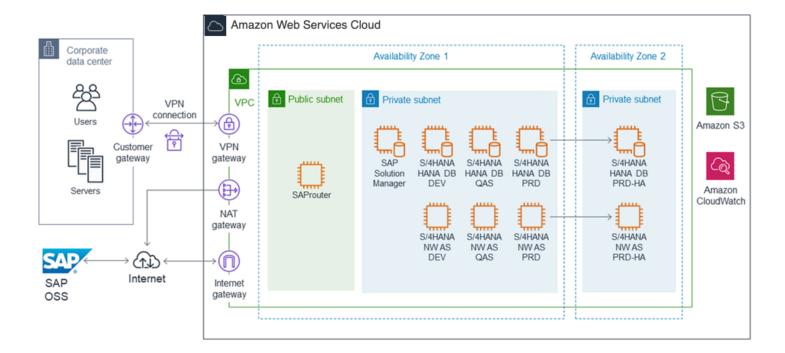
This section describes the two primary architectural patterns for SAP on Amazon: all systems on Amazon and hybrid.

All-on-Amazon Architecture

With the SAP All-on-Amazon architecture, all systems and components of your SAP environment are hosted on Amazon. Example scenarios of such an architecture include:

- Implementation of a complete, new SAP environment on Amazon
- Migration of a complete, existing SAP environment to Amazon

Figure 3 depicts an SAP all-on-Amazon architecture. The SAP environment running on Amazon is integrated with on-premises systems and users via a VPN connection or a dedicated network connection via Amazon Direct Connect. SAProuter is deployed in a public subnet and assigned a public IP address that is reachable from the internet to enable integration with the SAP OSS network via a secure network communications (SNC) connection. A <u>network address translation</u> (NAT) gateway enables instances in the private subnet to connect to the internet or other Amazon services, but prevents instances from receiving inbound traffic that is initiated by someone on the internet. For additional information, see the <u>Configuring Network and Connectivity</u> section.



SAP on Amazon architectures 15

Figure 3: SAP all-on-Amazon architecture

Hybrid Amazon Architecture

With an SAP hybrid Amazon architecture, some SAP systems and components are hosted on your on-premises infrastructure and others are hosted on the Amazon infrastructure. Example scenarios of such an architecture include:

- Running SAP test, trial, training, proof-of-concept (PoC), and similar systems on Amazon
- Running non-production SAP landscapes (for example, DEV and QAS) on Amazon, integrated with an SAP production landscape running on premises
- Implementing a new SAP application on Amazon and integrating it with an existing SAP onpremises environment

Figure 4 depicts an SAP hybrid Amazon architecture with SAP DEV and QAS landscapes and SAP test, training, and PoC systems running on Amazon. These systems are integrated with SAP systems and users on the corporate network. Connectivity between the VPC and the corporate network is provided with either a VPN connection or an Amazon Direct Connect connection. The existing SAProuter and SAP Solution Manager running on the corporate network are used to manage the SAP systems running within the VPC.

SAP on Amazon architectures 16

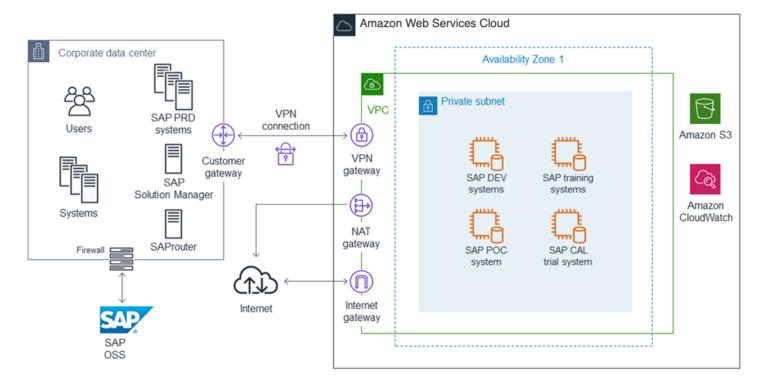


Figure 4: SAP hybrid Amazon architecture

Choosing an Amazon Region and Availability Zone

See the <u>Amazon Global Infrastructure</u> section of this guide for information about Amazon Regions and Availability Zones.

Choosing a Region

When choosing the Amazon Region to deploy your SAP environment in, consider the following factors:

- Proximity to your on-premises data center(s), systems, and end users to minimize network latency.
- · Data residency and regulatory compliance requirements.
- Availability of the Amazon products and services you plan to use in the region. For a detailed list
 of Amazon products and services by region, see the <u>Region Table</u> on the Amazon website.
- Availability of the EC2 instance types you plan to use in the region. To view Amazon Region availability for a specific instance type, see the Amazon EC2 Instance Types for SAP webpage.

Choosing an Availability Zone

No special considerations are required when choosing an Availability Zone for your SAP deployment on Amazon. All SAP applications (SAP ERP, CRM, SRM, and so on) and systems (SAP database system, SAP Central Services system, and SAP application servers) should be deployed in the same Availability Zone. If high availability (HA) is a requirement, use multiple Availability Zones. For more information, see Architecture guidance for availability and reliability of SAP on Amazon.

Network and Connectivity

Amazon VPC

Amazon VPC enables you to define a virtual network in your own, logically isolated area within the Amazon Cloud. You can launch your Amazon resources, such as instances, into your VPC. Your VPC closely resembles a traditional network that you might operate in your own data center, with the benefits of using the Amazon scalable infrastructure. You can configure your VPC; you can select its IP address range, create subnets, and configure route tables, network gateways, and security settings. You can connect instances in your VPC to the internet. You can connect your VPC to your own corporate data center, and make the Amazon Cloud an extension of your data center. To protect the resources in each subnet, you can use multiple layers of security, including security groups and network access control lists. For more information, see the Amazon VPC User Guide.

For detailed instructions for setting up and configuring a VPC, and connectivity between your network and VPC, see the Amazon VPC documentation.

Network Connectivity Options

Multiple options are available to provide network connectivity between your on-premises users and systems with your SAP systems running on Amazon, including a direct internet connection, hardware VPN, and private network connection.

Private Network Connection

<u>Amazon Direct Connect</u> makes it easy to establish a dedicated network connection from your premises to Amazon. Using Amazon Direct Connect, you can establish private connectivity between Amazon and your data center, office, or co-location environment. In many cases, this can reduce your network costs, increase bandwidth throughput, and provide a more consistent network experience than internet-based connections. For additional information, see the <u>Amazon Direct Connect User Guide</u>.

Network and connectivity 18

Use cases: Recommended for customers who require greater bandwidth and lower latency than possible with a hardware VPN.

For more information, see Amazon Virtual Private Cloud Cloud Connectivity Options.

Direct Internet Connection

The quickest and simplest way to connect to your SAP systems running on Amazon involves using a VPC with a single public subnet and an internet gateway to enable communication over the internet. For additional information, see Scenario 1: VPC with a Public Subnet Only in the Amazon VPC User Guide.

Use cases: Most suitable for SAP demo, training, and test type systems that do not contain sensitive data.

Site-to-Site / Hardware VPN

Amazon Site-to-Site VPN extends your data center or branch office to the cloud via Internet Protocol security (IPsec) tunnels, and supports connecting to both virtual private gateways and Amazon Transit Gateway. You can optionally run Border Gateway Protocol (BGP) over the IPsec tunnel for a highly available solution. For additional information, see Adding a Hardware Virtual Private Gateway to your VPC in the Amazon VPC User Guide.

Use cases: Recommended for any SAP environments on Amazon that require integration with onpremises users and systems.

Client VPN

Amazon Client VPN provides a fully-managed VPN solution that can be accessed from anywhere with an internet connection and an OpenVPN-compatible client. It is elastic, automatically scales to meet your demand, and enables your users to connect to both Amazon and on-premises networks. Amazon Client VPN seamlessly integrates with your existing Amazon infrastructure, including Amazon VPC and Amazon Directory Service, so you don't have to change your network topology.

Use cases: Provides quick and easy connectivity to your remote workforce and business partners.

Following Security Best Practices

In order to provide end-to-end security and end-to-end privacy, Amazon builds services in accordance with security best practices, provides appropriate security features in those services, and documents how to use those features. In addition, Amazon customers must use those features and best practices to architect an appropriately secure application environment. Enabling

customers to ensure the confidentiality, integrity, and availability of their data is of the utmost importance to Amazon, as is maintaining trust and confidence.

Shared Responsibility Environment

There is a shared responsibility model between you as the customer and Amazon. Amazon operates, manages, and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the services operate. In turn, you assume responsibility and management of the guest operating system (including updates and security patches), other associated application software, Amazon VPC setup and configuration, as well as the configuration of the Amazon-provided security group firewall. For additional information on Amazon security, visit the Amazon Cloud Security page and review the various Security Resources available there.

Amazon VPC

The foundation for security of an SAP environment on Amazon is the use of Amazon VPC for providing the overall isolation. Amazon VPC includes security details that you must set up to enable proper access and restrictions for your resources. Amazon VPC provides features that you can use to help increase and monitor the security for your VPC:

- **Security groups** act as a firewall for associated EC2 instances, controlling both inbound and outbound traffic at the instance level.
- **Network access control lists (ACLs)** act as a firewall for associated subnets, controlling both inbound and outbound traffic at the subnet level.
- Route tables consist of a set of rules, called routes, that determine where network traffic is directed. Each subnet in your VPC must be associated with a route table; the table controls the routing for the subnet.
- Flow logs capture information about the IP traffic going to and from network interfaces in your VPC.

For detailed documentation about how to set up and manage security within a VPC, see the <u>Security</u> section of the *Amazon VPC User Guide*.

EC2 Instance Types for SAP

Amazon EC2 provides a wide selection of instance types optimized to fit different use cases. Instance types comprise varying combinations of CPU, memory, storage, and networking capacity

and give you the flexibility to choose the appropriate mix of resources for your applications. Each instance type includes one or more instance sizes, allowing you to scale your resources to the requirements of your target workload.

SAP systems deployed on Amazon that will require support from SAP must be run on an EC2 instance type that has been certified with SAP. This section describes where you can find details about the EC2 instance types that have been certified with SAP and additional information for specific SAP solutions.

SAP NetWeaver-based Solutions

SAP solutions based on the SAP NetWeaver platform and that use <u>SAP Application Performance</u> <u>Standard (SAPS)</u> for sizing must be run on a specific subset of EC2 instance types in order to receive support from SAP Support. For details, see:

- SAP Note 1656099
- Amazon EC2 Types for SAP

SAP HANA

The SAP HANA platform and SAP solutions that run on top of an SAP HANA database—for example, SAP Suite on HANA, SAP S/4HANA, SAP Business Warehouse (BW) on HANA, SAP BW/4HANA— require specific EC2 instance types that have been certified for SAP HANA. For more information, see Amazon EC2 instance types for SAP on Amazon.

SAP Business One, version for SAP HANA

For information about the EC2 instance types that are certified for SAP Business One, version for SAP HANA, see:

- SAP Note 2058870
- SAP Business One on Amazon

Operating Systems

Supported Operating Systems

EC2 instances run on 64-bit virtual processors based on the Intel x86 instruction set. The following 64-bit operating systems and versions are available and supported for SAP solutions on Amazon.

Operating systems 21

- SUSE Linux Enterprise Server (SLES)
- SUSE Linux Enterprise Server for SAP Applications (SLES for SAP)
- Red Hat Enterprise Linux (RHEL)
- Red Hat Enterprise Linux for SAP Solutions (RHEL for SAP)
- Microsoft Windows Server
- Oracle Enterprise Linux

For additional information regarding SAP-supported operating systems on Amazon, see <u>SAP Note</u> 1656250.

SLES for SAP and RHEL for SAP

SUSE and Red Hat offer SAP-specific versions of their operating systems that provide the following benefits:

- · Configuration and tuning for SAP
- Extended release support
- High availability extension for SAP
- Dedicated support channel



Because of these benefits, we strongly recommend using SLES for SAP or RHEL for SAP with High Availability (HA) and Update Services (US) for your SAP on Amazon deployments.

To learn more about SUSE's and Red Hat's operating system versions for SAP, see the following information on the SLES and Red Hat websites.

SLES for SAP

- General information
- SUSE on Amazon for SAP Applications

RHEL for SAP

Operating systems 22

- General information
- Red Hat in the Public Cloud
- Red Hat Cloud Access
- How to Locate Red Hat Cloud Access Gold Images on Amazon EC2
- What is the Difference between Red Hat Cloud Access and Red Hat Enterprise Linux On-Demand Subscriptions in the public cloud?

Operating System Licenses

These operating system licensing options are available for SAP systems on Amazon:

- On-demand The operating system software and license are bundled in an Amazon Machine Images (AMI). The fee for the operating system license is included in the On-Demand Instance hourly fee or Reserved Instance fee for the instance type.
- **Bring Your Own License/Subscription (BYOL)** Bring your existing operating system license or subscription to the Amazon Cloud.
- Amazon Marketplace Purchase operating system licenses and subscriptions from Amazon Marketplace.

The following table lists the licensing options available for each operating system and version. To learn more about each option, follow the link in the table.

Operating system	License/subscription options
SLES	On-demand BYOL
SLES for SAP	Amazon Marketplace BYOL
RHEL	On-demand BYOL
RHEL for SAP with HA and US	Amazon Marketplace BYOL
Windows	On-demand BYOL
Oracle Linux	BYOL

Operating systems 23

Databases

Supported Databases

All the database platforms and versions supported by SAP for an on-premises infrastructure are also supported by SAP on Amazon. For details about the databases supported with specific SAP solutions on Amazon, see SAP Note 1656099.

Database Installation and Administration

Customer-Managed Database on Amazon EC2

The majority of SAP solutions use a customer-managed model on Amazon EC2. Installation, configuration, administration, and backup and recovery of the database are done by either the customer or a partner.

The following SAP solutions use a self-managed database model on Amazon EC2:

- SAP Business Suite and SAP NetWeaver-based applications
- SAP HANA
- SAP S/4HANA
- SAP BW/4HANA
- SAP BusinessObjects BI
- SAP Business One

Amazon RDS

Amazon Relational Database Service (Amazon RDS) is a managed service that makes it easy to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity, while managing time-consuming database administration tasks, freeing you up to focus on your applications and business. Amazon RDS is currently supported for the following SAP solutions:

- SAP BusinessObjects BI
- SAP Commerce (previously known as SAP Hybris Commerce)

Databases 24

Amazon Aurora

<u>Amazon Aurora (Aurora)</u> is a MySQL and PostgreSQL-compatible relational database built for the cloud. It combines the performance and availability of traditional enterprise databases with the simplicity and cost-effectiveness of open source databases. Aurora MySQL is currently supported for the following SAP solution:

SAP Commerce (previously known as SAP Hybris Commerce)

Database Licenses

These database licensing options are available for SAP systems on Amazon:

- On-demand The database software and license are bundled in an Amazon Machine Image (AMI). The fee for the database license is included in the On-Demand Instance hourly fee or Reserved Instance fee for the instance type.
- Bring Your Own License (BYOL) Bring your existing database licenses to the Amazon Cloud.
- Amazon Marketplace Purchase database software and licenses from Amazon Marketplace.

The following table lists the licensing options available on Amazon for each database. For additional information, follow the links in the *Licensing options* column.

Database	Licensing options
SAP HANA	BYOL
SAP Adaptive Server Enterprise (ASE) (SAP ASE)	BYOL
Microsoft SQL Server	BYOL*
IBM DB2	BYOL
Oracle	BYOL
Amazon Aurora	<u>On-demand</u>

Databases 25

* SQL Server runtime licenses purchased from SAP require either Microsoft Software Assurance or Amazon EC2 Dedicated Hosts to bring these licenses to Amazon. For additional information, see:

- SAP Note 2139358 Effect of changes in licensing terms of SQL Server
- Microsoft Licensing on Amazon

SAP Installation Media

The majority of SAP solutions on Amazon use a bring-your-own-software model. There are two primary options for copying SAP installation media to Amazon:

- Download from the SAP Software Download Center to Amazon EC2. From your EC2 instance, connect to the SAP Software Download Center and download the required installation media. This option will most likely be the fastest method for getting SAP installation media to Amazon, because EC2 instances have very fast connections to the internet. You can create a dedicated Amazon EBS volume to store installation media, and then attach the volume to different instances as needed. You can also create a snapshot of the Amazon EBS volume and create multiple volumes that you can attach to multiple instances in parallel.
- Copy from your network to Amazon EC2. If you already have the required SAP installation
 media downloaded to a location on your network, you can copy the media from your network
 directly to an EC2 instance.

SAProuter and SAP Solution Manager

The following sections describe options for SAProuter and SAP Solution Manager when running SAP solutions on Amazon.

For SAP All-on-Amazon Architecture

When setting up an SAP environment on Amazon, you will need to set up an SAP Solution Manager system and SAProuter with a connection to the SAP support network, as you would with any infrastructure. See the all-on-Amazon architecture diagram (Figure 3) for an illustration.

When setting up the SAProuter and SAP support network connection, follow these guidelines:

• Launch the instance that the SAProuter software is installed on into a public subnet of the VPC and assign it an Elastic IP address.

SAP installation media 26

• Create a specific security group for the SAProuter instance with the necessary rules to allow the required inbound and outbound access to the SAP support network.

• Use the Secure Network Communications (SNC) type of internet connection. For more information, see https://service.sap.com/internetconnection.

For SAP Hybrid Amazon Architecture

When using Amazon as an extension of your IT infrastructure, you can use your existing SAP Solution Manager system and SAProuter that are running in your data center to manage SAP systems running on Amazon within a VPC. See the hybrid architecture diagram (Figure 4) for additional information.

Document Revisions

Date	Change	Location
January, 2023	Update	Changes throughout guide
May, 2019	Update	Changes throughout guide
August, 2018	Initial publication	_

Document revisions 27

Amazon EC2 instance types for SAP on Amazon

Amazon Elastic Compute Cloud (Amazon EC2) offers a wide selection of instance types optimized to fit different use cases. The varying combinations of CPU, memory, storage, and networking capacity provide flexibility in selection of resources for your applications. Each instance type includes one or more instance sizes, enabling you to scale your resources according to the requirements of your workload. For more information, see Amazon EC2 instance types.

Amazon has worked closely with SAP to test and certify Amazon EC2 instance types for SAP on Amazon solutions. For more information, see SAP Note 1656099 - SAP Applications on Amazon: Supported DB/OS and Amazon EC2 products (requires SAP Portal access) and SAP Certified and Supported SAP HANA Hardware Directory.

Topics

- Instance type availability
- SAP NetWeaver supported instances
- SAP HANA certified and non-certified instances
- SAP Business One certified instances, version for SAP HANA
- Document history

Instance type availability

The availability of an Amazon EC2 instance type is based on your selection of Region. For more information about available instance types in your Region, see Amazon EC2 On-Demand pricing.



Note

Certain Amazon EC2 instance families, such as X1, X2idn, X2iedn, and High Memory may not be available across all Availability Zones within a Region. You must confirm while planning that the instance types required for your SAP workloads are available in your target Availability Zone.

You can also determine the availability of an instance type in a Region and its Availability Zone by using the describe-instance-type-offerings command and specifying the Region and the instance type.

Instance type availability

```
aws ec2 describe-instance-type-offerings --region us-east-2 --location-type availability-zone --filters Name=instance-type, Values=r5.8xlarge
```

```
{
    "InstanceTypeOfferings": [
        {
            "InstanceType": "r5.8xlarge",
            "LocationType": "availability-zone",
            "Location": "us-east-2a"
        },
        {
            "InstanceType": "r5.8xlarge",
            "LocationType": "availability-zone",
            "Location": "us-east-2b"
        },
        {
            "InstanceType": "r5.8xlarge",
            "LocationType": "availability-zone",
            "Location": "us-east-2c"
        }
    ]
}
```

SAP NetWeaver supported instances

Topics

- Current generation Amazon EC2 instances for SAP NetWeaver
- Previous generation Amazon EC2 instances for SAP NetWeaver

Current generation Amazon EC2 instances for SAP NetWeaver

General Purpose

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
m5	m5.metal	96	384	142,000	25	19,000	80,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
m5	m5.24xlar ge	96	384	135,230	25	19,000	80,000
m5	m5.16xlar ge	64	256	90,153	20	13,600	60,000
m5	m5.12xlar ge	48	192	67,615	10	9,500	40,000
m5	m5.8xlarg e	32	128	45,077	10	6,800	30,000
m5	m5.4xlarg e	16	64	22,538	High	4,750	18,750
m5	m5.2xlarg e	8	32	11,269	High	Up to 4,750	18,750
m5	m5.xlarge	4	16	5,535	High	Up to 4,750	18,750
m5	m5.large	2	8	2,817	High	Up to 4,750	18,750
m6i	m6i.32xla rge	128	512	198,080	50	40,000	160,000
m6i	m6i.24xla rge	96	384	148,560	37.5	30,000	120,000
m6i	m6i.16xla rge	64	256	99,040	25	20,600	80,000
m6i	m6i.12xla rge	48	192	74,280	18.75	15,000	60,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
m6i	m6i.8xlar ge	32	128	49,520	12.5	10,000	40,000
m6i	m6i.4xlar ge	16	64	24,760	Up to 12.5	Up to 10,000	40,000
m6i	m6i.2xlar ge	8	32	12,380	Up to 12.5	Up to 10,000	40,000
m6i	m6i.xlarg e	4	16	6,190	Up to 12.5	Up to 10,000	40,000
m6i	m6i.large	2	8	3,095	Up to 12.5	Up to 10,000	40,000
m6id	m6id.32xl arge	128	512	198,080	50	40,000	160,000
m6id	m6id.24xl arge	96	384	148,560	37.5	30,000	120,000
m6id	m6id.16xl arge	64	256	99,040	25	20,600	80,000
m6id	m6id.12xl arge	48	192	74,280	18.75	15,000	60,000
m6id	m6id.8xla rge	32	128	49,520	12.5	10,000	40,000
m6id	m6id.4xla rge	16	64	24,760	Up to 12.5	Up to 10,000	40,000
m6id	m6id.2xla rge	8	32	12,380	Up to 12.5	Up to 10,000	40,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
m6id	m6id.xlar ge	4	16	6,190	Up to 12.5	Up to 10,000	40,000
m6id	m6id.larg e	2	8	3,095	Up to 12.5	Up to 10,000	40,000
m6in	m6in.32xl arge	128	512	198,080	200	100,000	400,000
m6in	m6in.24xl arge	96	384	148,560	150	75,000	300,000
m6in	m6in.16xl arge	64	256	99,040	100	50,000	200,000
m6in	m6in.12xl arge	48	192	74,280	75	37,500	150,000
m6in	m6in.8xla rge	32	128	49,520	50	25,000	100,000
m6in	m6in.4xla rge	16	64	24,760	Up to 50	Up to 25,000	100,000
m6in	m6in.2xla rge	8	32	12,380	Up to 40	Up to 25,000	100,000
m6in	m6in.xlar ge	4	16	6,190	Up to 30	Up to 25,000	100,000
m6in	m6in.larg e	2	8	3,095	Up to 25	Up to 25,000	100,000
m6idn	m6idn.32x large	128	512	198,080	200	100,000	400,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
m6idn	m6idn.24x large	96	384	148,560	150	75,000	300,000
m6idn	m6idn.16x large	64	256	99,040	100	50,000	200,000
m6idn	m6idn.12x large	48	192	74,280	75	37,500	150,000
m6idn	m6idn.8xl arge	32	128	49,520	50	25,000	100,000
m6idn	m6idn.4xl arge	16	64	24,760	Up to 50	Up to 25,000	100,000
m6idn	m6idn.2xl arge	8	32	12,380	Up to 40	Up to 25,000	100,000
m6idn	m6idn.xla rge	4	16	6,190	Up to 30	Up to 25,000	100,000
m6idn	m6idn.lar ge	2	8	3,095	Up to 25	Up to 25,000	100,000
m6a	m6a.48xla rge	192	768	290,220	50	40,000	240,000
m6a	m6a.32xla rge	128	512	193,480	50	40,000	160,000
m6a	m6a.24xla rge	96	384	145,110	37.5	30,000	120,000
m6a	m6a.16xla rge	64	256	96,740	25	20,000	80,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
m6a	m6a.12xla rge	48	192	72,555	18.75	15,000	60,000
m6a	m6a.8xlar ge	32	128	48,370	12.5	10,000	40,000
m6a	m6a.4xlar ge	16	64	24,185	Up to 12.5	Up to 10,000	40,000
m6a	m6a.2xlar ge	8	32	12,093	Up to 12.5	Up to 10,000	40,000
m6a	m6a.xlarg e	4	16	6,046	Up to 12.5	Up to 10,000	40,000
m6a	m6a.large	2	8	3,023	Up to 12.5	Up to 10,000	40,000
m7a	m7a.48xla rge	192	768	429,720	50	40,000	240,000
m7a	m7a.32xla rge	128	512	286,464	50	40,000	160,000
m7a	m7a.24xla rge	96	384	214,848	37.5	30,000	120,000
m7a	m7a.16xla rge	64	256	143,232	25	20,000	80,000
m7a	m7a.12xla rge	48	192	107,424	18.75	15,000	60,000
m7a	m7a.8xlar ge	32	128	71,616	12.5	10,000	40,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
m7a	m7a.4xlar ge	16	64	35,808	Up to 12.5	Up to 10,000	Up to 40,000
m7a	m7a.2xlar ge	8	32	17,904	Up to 12.5	Up to 10,000	Up to 40,000
m7a	m7a.xlarg e	4	16	8,952	Up to 12.5	Up to 10,000	Up to 40,000
m7a	m7a.large	2	8	4,476	Up to 12.5	Up to 10,000	Up to 40,000
m7a	m7a.mediu	1	4	2,238	Up to 12.5	Up to 10,000	Up to 40,000
m7i	m7i.48xla rge*	192	768	306,202	50	40,000	240,000
m7i	m7i.24xla rge	96	384	167,470	37.5	30,000	120,000
m7i	m7i.16xla rge	64	256	123,300	25	20,000	80,000
m7i	m7i.12xla rge	48	192	101,200	18.75	15,000	60,000
m7i	m7i.8xlar ge	32	128	67,480	12.5	10,000	40,000
m7i	m7i.4xlar ge	16	64	33,740	Up to 12.5	Up to 10,000	Up to 40,000
m7i	m7i.2xlar ge	8	32	16,870	Up to 12.5	Up to 10,000	Up to 40,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
m7i	m7i.xlarg e	4	16	8,435	Up to 12.5	Up to 10,000	Up to 40,000
m7i	m7i.large	2	8	4,218	Up to 12.5	Up to 10,000	Up to 40,000

Compute Optimized

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
c5	c5.18xlar ge	72	144	95,400	25	19,000	80,000
c5	c5.9xlarg e	36	72	47,700	10	9,500	40,000
c5	c5.4xlarg e	16	32	21,200	Up to 10	4,750	20,000
c5	c5.2xlarg e	8	16	10,600	Up to 10	Up to 4,750	20,000
c5	c5.xlarge	4	8	5,300	Up to 10	Up to 4,750	20,000
c5	c5.large	2	4	2,650	Up to 10	Up to 4,750	20,000
c5n	c5n.18xla rge	72	192	95,400	100	19,000	80,000
c5n	c5n.9xlar ge	36	96	47,700	50	9,500	40,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
c5n	c5n.4xlar ge	16	42	21,200	Up to 25	4,750	20,000
c5n	c5n.2xlar ge	8	21	10,600	Up to 25	Up to 4,750	20,000
c5n	c5n.xlarg e	4	11	5,300	Up to 25	Up to 4,750	20,000
c5n	c5n.large	2	5	2,650	Up to 25	Up to 4,750	20,000
c5a	c5a.24xla rge	96	192	131,830	20	9,500	40,000
c5a	c5a.16xla rge	64	128	87,887	20	6,300	26,700
c5a	c5a.12xla rge	48	96	65,915	12	4,750	20,000
c5a	c5a.8xlar ge	32	64	43,943	10	3,170	13,300
c5a	c5a.4xlar ge	16	32	21,973	Up to 10	Up to 3,170	13,300
c5a	c5a.2xlar ge	8	16	10,986	Up to 10	Up to 3,170	13,300
c5a	c5a.xlarg e	4	8	5,493	Up to 10	Up to 3,170	13,300
c5a	c5a.large	2	4	2,746	Up to 10	Up to 3,170	13,300

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
c6i	c6i.32xla rge	128	256	188,780	50	40,000	160,000
c6i	c6i.24xla rge	96	192	141,585	37.5	30,000	120,000
c6i	c6i.16xla rge	64	128	94,390	25	20,600	80,000
c6i	c6i.12xla rge	48	96	70,793	18.75	15,000	60,000
c6i	c6i.8xlar ge	32	64	47,195	12.5	10,000	40,000
c6i	c6i.4xlar ge	16	32	23,598	Up to 12.5	Up to 10,000	40,000
c6i	c6i.2xlar ge	8	16	11,799	Up to 12.5	Up to 10,000	40,000
c6i	c6i.xlarg e	4	8	5,899	Up to 12.5	Up to 10,000	40,000
c6i	c6i.large	2	4	2,950	Up to 12.5	Up to 10,000	40,000
c6id	c6id.32xl arge	128	256	188,780	50	40,000	160,000
c6id	c6id.24xl arge	96	192	141,585	37.5	30,000	120,000
c6id	c6id.16xl arge	64	128	94,390	25	20,600	80,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
c6id	c6id.12xl arge	48	96	70,793	18.75	15,000	60,000
c6id	c6id.8xla rge	32	64	47,195	12.5	10,000	40,000
c6id	c6id.4xla rge	16	32	23,598	Up to 12.5	Up to 10,000	40,000
c6id	c6id.2xla rge	8	16	11,799	Up to 12.5	Up to 10,000	40,000
c6id	c6id.xlar ge	4	8	5,899	Up to 12.5	Up to 10,000	40,000
c6id	c6id.larg e	2	4	2,950	Up to 12.5	Up to 10,000	40,000
c6a	c6a.48xla rge	192	384	274,900	50	40,000	240,000
c6a	c6a.32xla rge	128	256	183,267	50	40,000	160,000
c6a	c6a.24xla rge	96	192	137,450	37.5	30,000	120,000
c6a	c6a.16xla rge	64	128	91,633	25	20,000	80,000
c6a	c6a.12xla rge	48	96	68,725	18.75	15,000	60,000
c6a	c6a.8xlar ge	32	64	45,817	12.5	10,000	40,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
c6a	c6a.4xlar ge	16	32	22,908	Up to 12.5	Up to 10,000	40,000
c6a	c6a.2xlar ge	8	16	11,454	Up to 12.5	Up to 10,000	40,000
c6a	c6a.xlarg e	4	8	5,727	Up to 12.5	Up to 10,000	40,000
c6a	c6a.large	2	4	2,864	Up to 12.5	Up to 10,000	40,000

Memory Optimized

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
r5	r5.metal	96	768	143,230	25	19,000	80,000
r5	r5.24xlar ge	96	768	138,770	25	19,000	80,000
r5	r5.16xlar ge	64	512	92,513	20	13,600	60,000
r5	r5.12xlar ge	48	384	69,385	10	9,500	40,000
r5	r5.8xlarg e	32	256	46,257	10	6,800	30,000
r5	r5.4xlarg e	16	128	23,128	Up to 10	4,750	18,750

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
r5	r5.2xlarg e	8	64	11,564	Up to 10	Up to 4,750	18,750
r5	r5.xlarge	4	32	5,782	Up to 10	Up to 4,750	18,750
r5	r5.large	2	16	2,891	Up to 10	Up to 4,750	18,750
r5n	r5n.metal	96	768	143,230	100	19,000	80,000
r5n	r5n.24xla rge	96	768	138,770	100	19,000	80,000
r5n	r5n.16xla rge	64	512	92,513	75	13,600	60,000
r5n	r5n.12xla rge	48	384	69,385	50	9,500	40,000
r5n	r5n.8xlar ge	32	256	46,257	25	6,800	30,000
r5n	r5n.4xlar ge	16	128	23,128	Up to 25	4,750	18,750
r5n	r5n.2xlar ge	8	64	11,564	Up to 25	Up to 4,750	18,750
r5n	r5n.xlarg e	4	32	5,782	Up to 25	Up to 4,750	18,750
r5n	r5n.large	2	16	2,891	Up to 25	Up to 4,750	18,750
r5b	r5b.metal	96	768	143,230	100	60,000	260,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
r5b	r5b.24xla rge	96	768	138,770	100	60,000	260,000
r5b	r5b.16xla rge	64	512	92,513	75	40,000	173,333
r5b	r5b.12xla rge	48	384	69,385	50	30,000	130,000
r5b	r5b.8xlar ge	32	256	46,257	25	20,000	86,667
r5b	r5b.4xlar ge	16	128	23,128	Up to 10	10,000	43,333
r5b	r5b.2xlar ge	8	64	11,564	Up to 10	Up to 10,000	43,333
r5b	r5b.xlarg e	4	32	5,782	Up to 10	Up to 10,000	43,333
r5b	r5b.large	2	16	2,891	Up to 10	Up to 10,000	43,333
r6i	r6i.metal	128	1,024	203,600	50	40,000	160,000
r6i	r6i.32xla rge	128	1,024	196,050	50	40,000	160,000
r6i	r6i.24xla rge	96	768	147,038	37.5	30,000	120,000
r6i	r6i.16xla rge	64	512	98,025	25	20,000	80,000
r6i	r6i.12xla rge	48	384	73,519	18.75	15,000	60,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
r6i	r6i.8xlar ge	32	256	49,013	12.5	10,000	40,000
r6i	r6i.4xlar ge	16	128	24,506	Up to 12.5	Up to 10,000	40,000
r6i	r6i.2xlar ge	8	64	12,253	Up to 12.5	Up to 10,000	40,000
r6i	r6i.xlarg e	4	32	6,127	Up to 12.5	Up to 10,000	40,000
r6i	r6i.large	2	16	3,063	Up to 12.5	Up to 10,000	40,000
r6id	r6id.meta l	128	1,024	203,600	50	40,000	160,000
r6id	r6id.32xl arge	128	1,024	196,050	50	40,000	160,000
r6id	r6id.24xl arge	96	768	147,038	37.5	30,000	120,000
r6id	r6id.16xl arge	64	512	98,025	25	20,000	80,000
r6id	r6id.12xl arge	48	384	73,519	18.75	15,000	60,000
r6id	r6id.8xla rge	32	256	49,013	12.5	10,000	40,000
r6id	r6id.4xla rge	16	128	24,506	Up to 12.5	Up to 10,000	40,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
r6id	r6id.2xla rge	8	64	12,253	Up to 12.5	Up to 10,000	40,000
r6id	r6id.xlar ge	4	32	6,127	Up to 12.5	Up to 10,000	40,000
r6id	r6id.larg e	2	16	3,063	Up to 12.5	Up to 10,000	40,000
r6in	r6in.32xl arge	128	1,024	196,050	200	100,000	400,000
r6in	r6in.24xl arge	96	768	147,038	37.5	75,000	300,000
r6in	r6in.16xl arge	64	512	98,025	25	50,000	200,000
r6in	r6in.12xl arge	48	384	73,519	18.75	37,500	150,000
r6in	r6in.8xla rge	32	256	49,013	12.5	25,000	100,000
r6in	r6in.4xla rge	16	128	24,506	Up to 25	Up to 25,000	100,000
r6in	r6in.2xla rge	8	64	12,253	Up to 25	Up to 25,000	100,000
r6in	r6in.xlar ge	4	32	6,127	Up to 25	Up to 25,000	100,000
r6in	r6in.larg e	2	16	3,063	Up to 25	Up to 25,000	100,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
r6idn	r6idn.32x large	128	1,024	196,050	200	100,000	400,000
r6idn	r6idn.24x large	96	768	147,038	150	75,000	300,000
r6idn	r6idn.16x large	64	512	98,025	100	50,000	200,000
r6idn	r6idn.12x large	48	384	73,519	75	37,500	150,000
r6idn	r6idn.8xl arge	32	256	49,013	50	25,000	100,000
r6idn	r6idn.4xl arge	16	128	24,506	Up to 50	Up to 25,000	100,000
r6idn	r6idn.2xl arge	8	64	12,253	Up to 40	Up to 25,000	100,000
r6idn	r6idn.xla rge	4	32	6,127	Up to 30	Up to 25,000	100,000
r6idn	r6idn.lar ge	2	16	3,063	Up to 25	Up to 25,000	100,000
r6a	r6a.48xla rge	192	1,536	286,450	50	40,000	240,000
r6a	r6a.32xla rge	128	1,024	190,967	50	40,000	160,000
r6a	r6a.24xla rge	96	768	143,225	37.5	30,000	120,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
r6a	r6a.16xla rge	64	512	95,483	25	20,000	80,000
r6a	r6a.12xla rge	48	384	71,613	18.75	15,000	60,000
r6a	r6a.8xlar ge	32	256	47,742	12.5	10,000	40,000
r6a	r6a.4xlar ge	16	128	23,871	Up to 12.5	Up to 10,000	40,000
r6a	r6a.2xlar ge	8	64	11,935	Up to 12.5	Up to 10,000	40,000
r6a	r6a.xlarg e	4	32	5,968	Up to 12.5	Up to 10,000	40,000
r6a	r6a.large	2	16	2,984	Up to 12.5	Up to 10,000	40,000
r7a	r7a.48xla rge	192	1536	432,980	50	40,000	240,000
r7a	r7a.32xla rge	128	1024	288,640	50	40,000	160,000
r7a	r7a.24xla rge	96	768	216,480	37.5	30,000	120,000
r7a	r7a.16xla rge	64	512	144,320	25	20,000	80,000
r7a	r7a.12xla rge	48	384	108,240	18.75	15,000	60,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
r7a	r7a.8xlar ge	32	256	72,160	12.5	10,000	40,000
r7a	r7a.4xlar ge	16	128	36,080	Up to 12.5	Up to 10,000	Up to 40,000
r7a	r7a.2xlar ge	8	64	18,040	Up to 12.5	Up to 10,000	Up to 40,000
r7a	r7a.xlarg e	4	32	9,020	Up to 12.5	Up to 10,000	Up to 40,000
r7a	r7a.large	2	16	4,510	Up to 12.5	Up to 10,000	Up to 40,000
r7a	r7a.mediu m	1	8	2,255	Up to 12.5	Up to 10,000	Up to 40,000
r7i	r7i.48xla rge*	192	1536	296,200	50	40,000	240,000
r7i	r7i.24xla rge	96	768	158,250	37.5	30,000	120,000
r7i	r7i.16xla rge	64	512	105,500	25	20,000	80,000
r7i	r7i.12xla rge	48	384	99,720	18.75	15,000	60,000
r7i	r7i.8xlar ge	32	256	66,480	12.5	10,000	40,000
r7i	r7i.4xlar ge	16	128	33,240	Up to 12.5	Up to 10,000	40,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
r7i	r7i.2xlar ge	8	64	16,620	Up to 12.5	Up to 10,000	40,000
r7i	r7i.xlarg e	4	32	8,310	Up to 12.5	Up to 10,000	40,000
r7i	r7i.large	2	16	4,155	Up to 12.5	Up to 10,000	40,000
x1	x1.32xlar ge	128	1,952	131,500	25	14,000	80,000
x1	x1.16xlar ge	64	976	65,750	10	7,000	40,000
x1e	x1e.32xla rge	128	3,904	131,500	25	14,000	80,000
x1e	x1e.16xla rge	64	1,952	65,750	10	7,000	40,000
x1e	x1e.8xlar ge	32	976	32,875	Up to 10	3,500	20,000
x1e	x1e.4xlar ge	16	488	16,437	Up to 10	1,750	10,000
x1e	x1e.2xlar ge	8	244	8,219	Up to 10	1,000	7,400
x1e	x1e.xlarg e	4	122	4,109	Up to 10	500	3,700
x2iezn	x2iezn.12 xlarge	48	1,536	85,020	100	19,000	80,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
x2iezn	x2iezn.8x large	32	1,024	56,680	75	12,000	55,000
x2iezn	x2iezn.6x large	24	768	42,510	50	9,500	40,000
x2iezn	x2iezn.4x large	16	512	28,340	Up to 25	4,175	20,000
x2iezn	x2iezn.2x large	8	256	14,170	Up to 25	3,170	13,333
x2idn	x2idn.32x large	128	2,048	196,050	100	80,000	300,000
x2idn	x2idn.24x large	96	1,536	147,038	75	60,000	200,000
x2idn	x2idn.16x large	64	1,024	98,025	50	40,000	150,000
x2iedn	x2iedn.32 xlarge	128	4,096	189,000	100	80,000	300,000
x2iedn	x2iedn.24 xlarge	96	3,072	141,750	75	60,000	200,000
x2iedn	x2iedn.16 xlarge	64	2,048	94,500	50	40,000	150,000
x2iedn	x2iedn.8x large	32	1,024	47,250	25	20,000	75,000
x2iedn	x2iedn.4x large	16	512	23,625	Up to 25	Up to 20,000	50,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
x2iedn	x2iedn.2x large	8	256	11,813	Up to 25	Up to 20,000	25,000
x2iedn	x2iedn.xl arge	4	128	5,906	Up to 25	Up to 20,000	12,500
High Memory	u-24tb1.m etal	448	24,576	517,480	100	38,000	160,000
High Memory	u-24tb1.1 12xlarge	448	24,576	508,720	100	38,000	160,000
High Memory	u-18tb1.m etal	448	18,432	534,130	100	38,000	160,000
High Memory	u-18tb1.1 12xlarge	448	18,432	520,330	100	38,000	160,000
High Memory	u-12tb1.m etal	448	12,288	480,600	100	38,000	160,000
High Memory	u-12tb1.1 12xlarge	448	12,288	475,500	100	38,000	160,000
High Memory	u-9tb1.me tal	448	9,216	480,600	100	38,000	160,000
High Memory	u-9tb1.11 2xlarge	448	9,216	475,500	100	38,000	160,000
High Memory	u-6tb1.me tal	448	6,144	480,600	100	38,000	160,000
High Memory	u-6tb1.11 2xlarge	448	6,144	475,500	100	38,000	160,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
High Memory	u-6tb1.56 xlarge	224	6,144	380,770	100	38,000	160,000
High Memory	u-3tb1.56 xlarge	224	3,072	237,750	50	19,000	80,000

Storage Optimized

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
i3en	i3en.meta l	96	768	143,230	100	19,000	80,000
i3en	i3en.24xl arge	96	768	138,770	100	19,000	80,000
i3en	i3en.12xl arge	48	384	69,385	50	9,500	40,000
i3en	i3en.6xla rge	24	192	34,693	25	4,750	20,000
i3en	i3en.3xla rge	12	96	17,346	Up to 25	Up to 4,750	20,000
i3en	i3en.2xla rge	8	64	11,564	Up to 25	Up to 4,750	20,000
i3en	i3en.xlar ge	4	32	5,782	Up to 25	Up to 4,750	20,000
i4i	i4i.metal	128	1,024	203,600	75	40,000	160,000

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	Max IOPS
i4i	i4i.32xla rge	128	1,024	196,050	75	40,000	160,000
i4i	i4i.24xla rge	96	768	147,038	56.25	30,000	120,000
i4i	i4i.16xla rge	64	512	98,025	37.5	20,000	80,000
i4i	i4i.12xla rge	48	384	73,519	28.125	15,000	60,000
i4i	i4i.8xlar ge	32	256	49,013	18.75	10,000	40,000
i4i	i4i.4xlar ge	16	128	24,506	Up to 25	Up to 10,000	40,000
i4i	i4i.2xlar ge	8	64	12,253	Up to 12.5	Up to 10,000	40,000
i4i	i4i.xlarge	4	32	6,127	Up to 10	Up to 10,000	40,000
i4i	i4i.large	2	16	3,063	Up to 10	Up to 10,000	40,000

^{*}m7i.48xlarge and r7i.48xlarge is supported on Windows 2016 and above, SLES 15 SP3 and above, and RHEL 8.6 and above.

Previous generation Amazon EC2 instances for SAP NetWeaver

General Purpose

Instance family	Instance type	vCPU	Memory (GiB)	SAPS
m2	m2.4xlarge	8	68.4	7,400
m2	m2.2xlarge	4	32.2	3,700
m4	m4.16xlarge	64	256	75,770
m4	m4.10xlarge	40	160	47,320
m4	m4.4xlarge	16	64	18,928
m4	m4.2xlarge	8	32	9,464
m4	m4.xlarge	4	16	4,732
m4	m4.large	2	8	2,366
N/A	cr1.8xlarge	32	244	30,430
N/A	cc2.8xlarge	32	60.5	90,330

Compute Optimized

Instance family	Instance type	vCPU	Memory (GiB)	SAPS
c 3	c3.8xlarge	32	60	31,830
c3	c3.4xlarge	16	30	15,915
c 3	c3.2xlarge	8	15	7,980
c3	c3.xlarge	4	7	3,990
c3	c3.large	2	3.75	1,995

Instance family	Instance type	vCPU	Memory (GiB)	SAPS
c4	c4.8xlarge	36	60	37,950
c4	c4.4xlarge	16	30	19,030
c4	c4.2xlarge	8	15	9,515
c4	c4.xlarge	4	7.5	4,758
c4	c4.large	2	3.75	2,379

Memory Optimized

Instance family	Instance type	vCPU	Memory (GiB)	SAPS
r3	r3.8xlarge	32	244	31,920
r3	r3.4xlarge	16	122	15,960
r3	r3.2xlarge	8	61	7,980
r3	r3.xlarge	4	30.5	3,990
r3	r3.large	2	15	1,995
r4	r4.16xlarge	64	488	76,400
r4	r4.8xlarge	32	244	38,200
r4	r4.4xlarge	16	122	19,100
r4	r4.2xlarge	8	61	9,550
r4	r4.xlarge	4	30.5	4,775
r4	r4.large	2	15.25	2,387

General SAP Guides SAP Guides SAP Guides

Previous generation Amazon EC2 instances for SAP NetWeaver are fully supported and these instance types retain the same features and functionality. We recommend using the current generation Amazon EC2 instance for new SAP NetWeaver implementations or migrations.

SAP HANA certified and non-certified instances

Topics

- Current generation certified Amazon EC2 instances for SAP HANA
- Previous generation certified Amazon EC2 instances for SAP HANA
- Non-certified Amazon EC2 instances for SAP HANA

Current generation certified Amazon EC2 instances for SAP HANA

Topics

- SAP HANA OLTP and OLAP Scale-up
- SAP HANA OLTP and OLAP Scale-out

SAP HANA OLTP and OLAP Scale-up

Instand family	Instand type	vCPU	Memoi (GiB)	SAPS		Storag (Mbps)		SAP HANA OLTP sizing	SAP HANA OLAP prod*	SAP HANA OLAP sizing	FSx for ONTAP
r5	r5.met	96	768	143,23	25	19,000	✓	Standa	✓	Standa	X
r5	r5.24xl ge	96	768	138,77	25	19,000	✓	Standa	✓	Standa	X
r5	r5.16xl ge	64	512	92,513	20	13,600	✓	Standa	✓	Standa	X
r5	r5.12xl ge	48	384	69,385	10	9,500	✓	Standa	✓	Standa	X

Instand family	Instand type	vCPU	Memoi (GiB)	SAPS		Storag (Mbps)		SAP HANA OLTP sizing	SAP HANA OLAP prod*	SAP HANA OLAP sizing	FSx for ONTAP
r5	r5.8xla e	32	256	46,257	10	6,800	✓	Standa	✓	Standa	X
r5b	r5b.me	96	768	143,23	100	60,000	✓	Standa	✓	Standa	x
r5b	r5b.24: rge	96	768	138,77	100	60,000	✓	Standa	✓	Standa	X
r5b	r5b.16: rge	64	512	92,513	75	40,000	✓	Standa	✓	Standa	X
r5b	r5b.12: rge	48	384	69,385	50	30,000	✓	Standa	✓	Standa	X
r5b	r5b.8xl ge	32	256	46,257	25	20,000	✓	Standa	✓	Standa	X
r6i	r6i.32x rge	128	1,024	196,05	50	40,000	✓	Standa	✓	Standa	✓
r6i	r6i.24x rge	96	768	147,03	37.5	30,000	✓	Standa	✓	Standa	✓
r6i	r6i.16x rge	64	512	98,025	25	20,000	✓	Standa	✓	Standa	✓
r6i	r6i.12x rge	48	384	73,519	18.75	15,000	✓	Standa	✓	Standa	✓
r6i	r6i.8xla ge	32	256	49,013	12.5	10,000	✓	Standa	X	N/A	X
r7i	r7i.48x rge	192	1536	296,20	50	40,000	✓	Standa	✓	Standa	✓

Instand family	Instand type	vCPU	Memoi (GiB)	SAPS		Storag (Mbps)		SAP HANA OLTP sizing	SAP HANA OLAP prod*	SAP HANA OLAP sizing	FSx for ONTAP
r7i	r7i.24x rge	96	768	158,25	37.5	30,000	✓	Standa	✓	Standa	✓
r7i	r7i.16x rge	64	512	105,50	25	20,000	✓	Standa	✓	Standa	✓
r7i	r7i.12x rge	48	384	99,720	18.75	15,000	✓	Standa	✓	Standa	✓
r7i	r7i.8xla ge	32	256	66,480	12.5	10,000	✓	Standa	✓	Standa	✓
x1	x1.32xl ge	128	1,952	131,50	25	14,000	✓	Standa	✓	Standa	X
x1	x1.16xl ge	64	976	65,750	10	7,000	✓	Standa	✓	Standa	X
x1e	x1e.32: rge	128	3,904	131,50	25	14,000	✓	Standa	✓	Worklo	X
x2idn	x2idn.3 large	128	2,048	196,05	100	80,000	✓	Standa	✓	Standa	✓
x2idn	x2idn.2 large	96	1,536	147,03	75	60,000	✓	Standa	✓	Worklo	✓
x2idn	x2idn.1 large	64	1,024	98,025	50	40,000	✓	Standa	✓	Worklo	✓
x2iedn	x2iedn. xlarge	128	4,096	189,00	100	80,000	✓	Standa	✓	Worklo	✓

Instand family	Instand type	vCPU	Memoi (GiB)	SAPS		Storag (Mbps)		SAP HANA OLTP sizing	SAP HANA OLAP prod*	SAP HANA OLAP sizing	FSx for ONTAP
x2iedn	x2iedn. xlarge	96	3,072	141,75	75	60,000	✓	Standa	✓	Worklo	✓
High Memor	u-24tb etal	448	24,576	517,48	100	38,000	✓	Worklo	X	N/A	X
High Memor	u-24tb 12xlarç	448	24,576	508,72	100	38,000	✓	Worklo	X	N/A	✓
High Memor	u-18tb etal	448	18,432	534,13	100	38,000	✓	Worklo	✓	Worklo	X
High Memor	u-18tb 12xlarç	448	18,432	520,33	100	38,000	✓	Worklo	✓	Worklo	✓
High Memor	u-12tb etal	448	12,288	480,60	100	38,000	✓	Standa	✓	Worklo	X
High Memor	u-12tb 12xlarç	448	12,288	475,50	100	38,000	✓	Standa	✓	Worklo	✓
High Memor	u-9tb1 tal	448	9,216	480,60	100	38,000	✓	Standa	✓	Worklo	X
High Memor	u-9tb1 2xlarge	448	9,216	475,50	100	38,000	✓	Standa	✓	Worklo	✓
High Memor	u-6tb1 tal	448	6,144	480,60	100	38,000	✓	Standa	✓	Standa	X
High Memor	u-6tb1 2xlarge	448	6,144	475,50	100	38,000	✓	Standa	✓	Standa	✓

Instand family	Instand type	vCPU	Memoi (GiB)	SAPS		Storag (Mbps)		SAP HANA OLTP sizing	OLAP	SAP HANA OLAP sizing	FSx for ONTAP
High Memor	u-6tb1 xlarge	224	6,144	380,77	100	38,000	✓	Standa	✓	Worklo	✓
High Memor	u-3tb1 xlarge	224	3,072	237,75	50	19,000	✓	Standa	✓	Worklo	✓

SAP HANA OLTP and OLAP Scale-out

Instar family		vCPU	Memo (GiB)	SAPS		Stora (Mbp:		SAP HANA OLTP sizing	OLTP Max. nodes	SAP HANA OLAP prod	SAP HANA OLAP sizing		FSx for ONTAP
r5	r5.24) ge	96	768	138,7	25	19,00	X	N/A	N/A	✓	Stand	16	X
r6i	r6i.32 rge	128	1,024	196,0	50	40,00	X	N/A	N/A	✓	Stand	16	✓
r6i	r6i.24 rge	96	768	147,0	37.5	30,00	X	N/A	N/A	✓	Stand	16	✓
x1	x1.32: ge	128	1,952	131,5	25	14,00	X	N/A	N/A	✓	Stand	25	X
x1	x1.16: ge	64	976	65,75	10	7,000	X	N/A	N/A	✓	Stand	7	X
x1e	x1e.32 rge	128	3,904	131,5	25	14,00	X	N/A	N/A	✓	Workl	25	X
x2idn	x2idn large	64	1,024	98,02	50	40,00	X	N/A	N/A	✓	Stand	16	✓

Instar famil		vCPU	Memo (GiB)	SAPS		Stora (Mbp:				HANA		OLAP Max. nodes	FSx for ONTAR
x2idn	x2idn large	96	1,536	147,0	75	60,00	X	N/A	N/A	✓	Workl	16	✓
x2idn	x2idn large	128	2,048	196,0	100	80,00	X	N/A	N/A	✓	Workl	16	✓
x2ied	x2iedı xlarge		3,072	141,7	75	60,00	X	N/A	N/A	✓	Workl	16	✓
x2ied	x2iedı xlarge		4,096	189,0	100	80,00	X	N/A	N/A	✓	Workl	16	✓
High Memo	u-12t etal	448	12,28	480,6	100	38,00	✓	Stand	4	X	N/A	N/A	X
_	u-12t 12xla	448	12,28	475,5	100	38,00	✓	Stand	4	✓	Workl	16	✓
	u-9tb 2xlarç		9,216	475,5	100	38,00	✓	Stand	4	✓	Workl	16	✓
High Memo	u-6tb tal	448	6,144	480,6	100	38,00	✓	Stand	4	✓	Stand	16	X
_	u-6tb 2xlarç		6,144	475,5	100	38,00	✓	Stand	4	✓	Stand	16	✓
_	u-6tb xlarge		6,144	380,7	100	38,00	✓	Stand	4	✓	Workl	16	✓



Note

All the current and previous generation Amazon EC2 instance types for SAP HANA can be used for running non-production workloads. For more information, see SAP Note 2271345.

Previous generation certified Amazon EC2 instances for SAP HANA

Topics

- SAP HANA OLTP and OLAP Scale-up
- SAP HANA OLTP and OLAP Scale-out

SAP HANA OLTP and OLAP Scale-up

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	SAP HANA OLTP prod*	SAP HANA OLTP sizing	SAP HANA OLAP prod*	SAP HANA OLAP sizing
r3	r3.8xlarg e	32	244	31,920	✓	Standard	✓	Standard
r3	r3.4xlarg e	16	122	15,960	X	Standard	X	Standard
r3	r3.2xlarg e	8	61	7,980	X	Standard	x	Standard
r4	r4.16xlar ge	64	488	76,400	25	14,000	✓	Standard
r4	r4.8xlarg e	32	244	38,200	10	7,000	✓	Standard

SAP HANA OLTP and OLAP Scale-out

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	SAP HANA OLTP prod*	SAP HANA OLTP sizing	SAP HANA OLAP prod*	SAP HANA OLAP sizing
r3	r3.8xlarg e	32	244	31,920	X	N/A	✓	Standard

Previous generation Amazon EC2 instances for SAP HANA are fully supported and these instance types retain the same features and functionality. We recommend using the current generation Amazon EC2 instance for new SAP HANA implementations or migrations.

Non-certified Amazon EC2 instances for SAP HANA

Amazon EC2 instances listed in the following table are not certified for production usage. You can use them for running non-production workloads. For more information, see <u>SAP Note 2271345 – Cost-Optimized SAP HANA Hardware for Non-Production Usage</u> (SAP portal access required).

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	FSx for ONTAP
r4	r4.4xlarg e	16	122	19,100	Up to 10	3,500	X
r4	r4.2xlarg e	8	61	9,550	Up to 10	1,700	X
r5	r5.4xlarg e	16	128	23,128	Up to 10	4,750	X
r5	r5.2xlarg e	8	64	11,564	Up to 10	Up to 4,750	X
r5b	r5b.4xlar ge	16	128	23,128	Up to 10	10,000	X

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Network (Gbps)	Storage (Mbps)	FSx for ONTAP
r5b	r5b.2xlar ge	8	64	11,564	Up to 10	Up to 10,000	X
r6i	r6i.4xlar ge	16	128	24,506	Up to 12.5	Up to 10,000	X
r6i	r6i.2xlar ge	8	64	12,253	Up to 12.5	Up to 10,000	X
x1e	x1e.4xlar ge	16	488	16,437	Up to 10	1,750	X
x1e	x1e.2xlar ge	8	244	8,219	Up to 10	1,000	X
x1e	x1e.xlarg e	4	122	4,109	Up to 10	500	X
x2iedn	x2iedn.16 xlarge	64	2,048	94,500	50	40,000	X
x2iedn	x2iedn.8x large	32	1,024	47,250	25	20,000	X
x2iedn	x2iedn.4x large	16	512	23,625	Up to 25	Up to 20,000	X
x2iedn	x2iedn.2x large	8	256	11,813	Up to 25	Up tp 20,000	X
x2iedn	x2iedn.xl arge	4	128	5,906	Up to 25	Up to 20,000	X

SAP Business One certified instances, version for SAP HANA

General Purpose

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Maximum concurrent users
m4	m4.16xlarge	64	256	75,770	100
m4	m4.10xlarge	40	160	47,320	50
r3	r3.8xlarge	32	244	31,920	50
c3	c3.8xlarge	32	60	31,830	25

Memory Optimized

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Maximum concurrent users
r5	r5.24xlarge	96	768	138,770	250
r5	r5.12xlarge	48	384	69,385	150
r5	r5.4xlarge	16	128	23,128	50
r5	r5.2xlarge	8	64	11,564	25
r6i	r6i.2xlarge	8	64	12,253	25
r6i	r6i.4xlarge	16	128	24,506	50
r6i	r6i.8xlarge	32	256	49,013	100
r7i	r7i.2xlarge	8	64	16,620	25
r7i	r7i.4xlarge	16	128	33,240	50

Instance family	Instance type	vCPU	Memory (GiB)	SAPS	Maximum concurrent users
r7i	r7i.8xlarge	32	256	66,480	100
r7i	r7i.12xlarge	48	384	99,720	150
x1	x1.16xlarge	64	976	65,750	200

Document history

Change	Date
Added r7i to SAP Business One	January 2024
Added m7i and r7i to SAP NetWeaver	October 2023
Added r7a to SAP NetWeaver	September 2023
Added m7a to SAP NetWeaver	August 2023
Added r6i to SAP Business One	February 2023
Added r6id, r6in, r6idn, m6id, m6in, and m6idn to SAP NetWeaver	January 2023
Added FSx for ONTAP support for u-18tb1.1 12 and u-24tb1.112	January 2023
Added u-24tb1 and u-18tb1 to SAP HANA and SAP NetWeaver	October 2022
Added i4i to SAP NetWeaver	September 2022
Added FSx for ONTAP support for SAP HANA	September 2022
Added r6a to SAP NetWeaver	July 2022

Document history 65

Change	Date
Added c6i, c6id, and c6a to SAP NetWeaver	July 2022
Added x2idn and r6i to SAP HANA OLAP scale- out	June 2022
Initial publication	April 2022

Document history 66

Amazon Data Provider for SAP

Amazon Data Provider for SAP is a tool that collects performance-related data from Amazon services. It makes this data available to SAP applications to help monitor and improve the performance of business transactions. SAP requires customers to install the agent as described in SAP note 1656250 (login credentials required).

The Amazon Data Provider for SAP uses operating system, network, and storage data that is most relevant to the operation of the SAP infrastructure. Its data sources include Amazon Elastic Compute Cloud (Amazon EC2) and Amazon CloudWatch. This guide provides installation, configuration, and troubleshooting information for the Amazon Data Provider for SAP on both Linux and Windows.

About this guide

This guide is part of a content series that provides detailed information about hosting, configuring, and using SAP technologies in the Amazon Cloud. For the other guides in the series, ranging from overviews to advanced topics, see the SAP on Amazon Technical Documentation home page.

Introduction

Many organizations of all sizes are choosing to host key SAP systems in the Amazon Web Services Cloud. With Amazon, you can quickly provision an SAP environment. Additionally, the elastic nature of the Amazon Cloud enables you to scale computing resources up and down as needed. As a result, your business can dedicate more resources (both people and funds) to innovation.

Many SAP systems operate daily business transactions and are critical to business functions. As an SAP customer, you need the ability to track and troubleshoot the performance of these transactions. The Amazon Data Provider for SAP is a tool that collects key performance data on an Amazon EC2 instance that SAP applications can use to monitor transactions built by SAP. The data is collected from a variety of sources within your Amazon Cloud operating environment, including Amazon EC2 and Amazon CloudWatch. This data includes information about the operating system, network, and storage that is relevant to your SAP infrastructure. Data from the Amazon Data Provider for SAP is read by the SAP Operating System Collector (SAPOSCOL) and the SAP CIM Provider.

Figure 1 provides a high-level illustration of the Amazon Data Provider for SAP, its data sources, and its outputs.

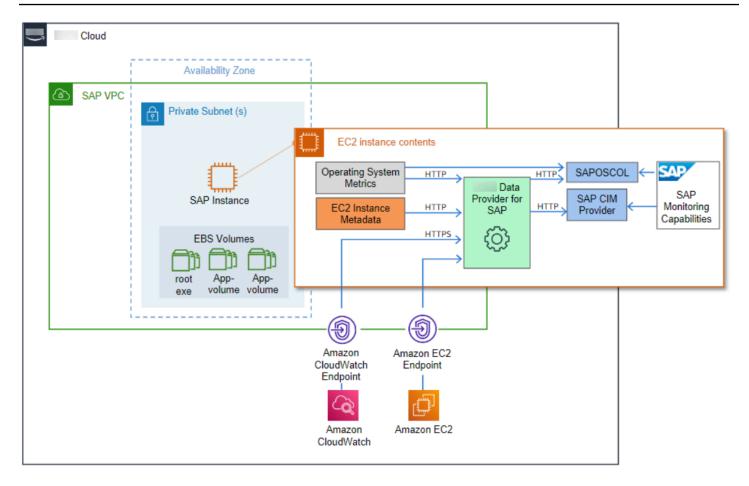


Figure 1: Data sources for the Amazon Data Provider for SAP

The purpose of this guide is to help you:

- Understand the technical requirements and components necessary to install and operate the Amazon Data Provider for SAP.
- Install the Amazon Data Provider for SAP.
- Understand the update process for the Amazon Data Provider for SAP.
- Troubleshoot installation issues.

Pricing

The DataProvider agent is provided free of charge. However, there are indirect costs associated with running the agent due to SAP requiring monitoring data to be delivered at specific intervals. This causes the DataProvider to do frequent **GetMetric** calls to Amazon CloudWatch and the Amazon EC2 API to retrieve the metric data. The expected costs for these calls ranges

Pricing 68

approximately from \$20.00 to \$40.00 per month per system and will vary based on how many disks are attached to the Amazon EC2 instance.

Example: Costs per month for using the DataProvider agent in the US East (N. Virginia) Region.

Fixed:

• Running the 2 required Amazon VPC endpoints (monitoring, Amazon EC2) is approximately **\$14.00 + \$0.01** per processed GB of data.



Note

These endpoints only need to be created once and are shared by the entire landscape. If you are already using these endpoints, you do not need to create them again.

Per System:

 You should expect around 70,000 API calls a day per instance (with 6 disks attached. At \$0.01 per 1,000 calls, it is approximately \$21.00 per month. The API call number increases or decreases based on the number of disks that are attached.

Technical Requirements

Before creating an SAP instance, ensure that the following technical requirements are met.

Amazon VPC Network Topologies

You need to deploy SAP systems that receive information from the Amazon Data Provider for SAP within an Amazon Virtual Private Cloud (Amazon VPC). You can use one of the following network topologies to enable routing to internet-based endpoints:

 The first topology configures routes and traffic directly to the Amazon Cloud through a NAT gateway within an Amazon VPC (see Figure 2). For more information about internet gateways, see the Amazon documentation.

Technical requirements

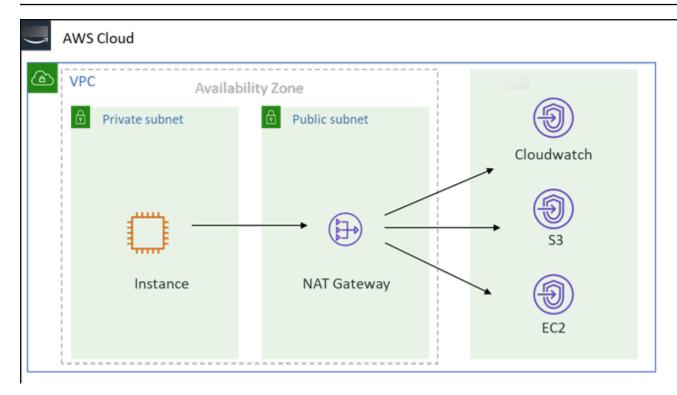


Figure 2: Connection to the Amazon Cloud via an internet gateway

• A second topology routes traffic from the Amazon VPC, through your organization's on-premises data center, and back to Amazon Cloud (see Figure 3). For more information about this topology, see the What is Amazon Site-to-Site VPN?

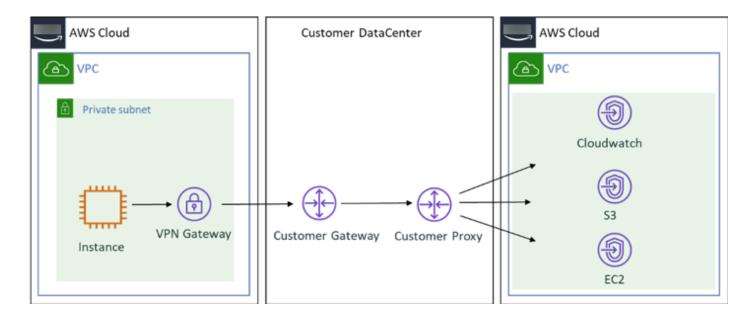


Figure 3: Connection to the Amazon Web Services Cloud via an on-premises data center

Amazon VPC Endpoints

Create endpoints for the following services that the DataProvider uses:

- Monitoring
- Amazon EC2

To create data endpoints in the Amazon console, use the following procedure for each of the two endpoints:

- 1. Sign in to the Amazon VPC console, navigate to **Endpoints**, and select **Create Endpoint**.
- 2. On the next screen, search for the service name, then select the appropriate VPC and route table, and select **Create Endpoint**.

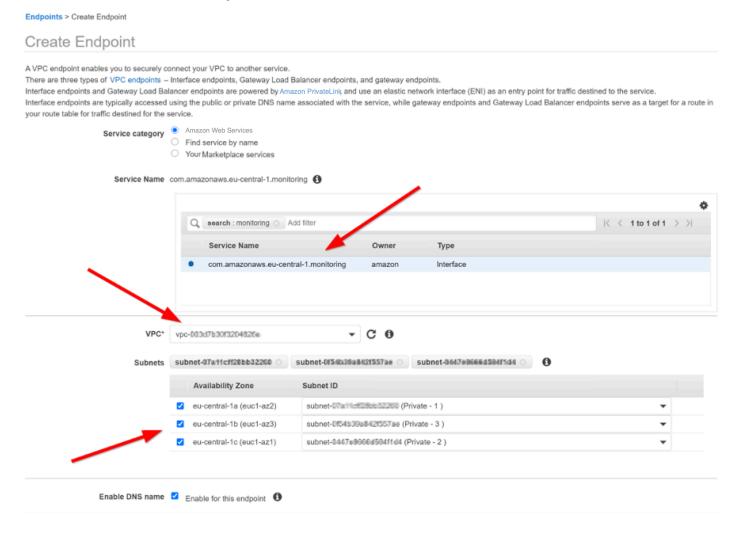


Figure 5: Completing VPC endpoint creation

Amazon VPC Endpoints 71

3. After creating all three endpoints you should see them in your list of endpoints as shown below:



Figure 6: VPC endpoints created

IAM Roles

You need to grant the Amazon Data Provider for SAP read-only access to the Amazon CloudWatch, Amazon Simple Storage Service (Amazon S3), and Amazon EC2 services so that you can use their APIs. You can do this by creating an Amazon Identity and Access Management (IAM) role for your Amazon EC2 instance and attaching a permissions policy.

Use the following procedure to create an IAM role and grant permissions to your Amazon EC2 instance:

- 1. Sign in to the Amazon Web Services Management Console and open the IAM console.
- 2. In the navigation pane, select **Roles**, and select **Create role**.
- 3. Choose the **Amazon service** role type, and select **EC2**.

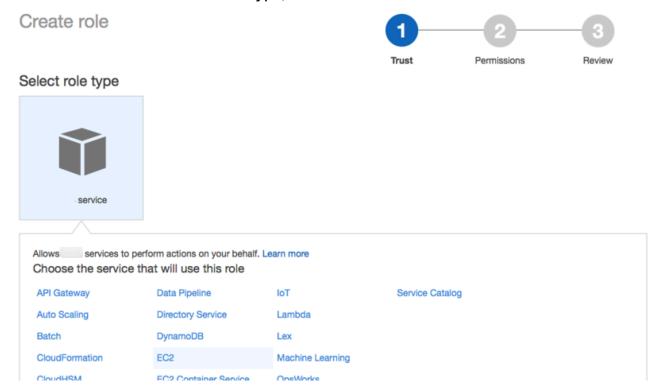


Figure 7: Creating a role for Amazon EC2

4. Select EC2 as the use case, and select Next Permissions.

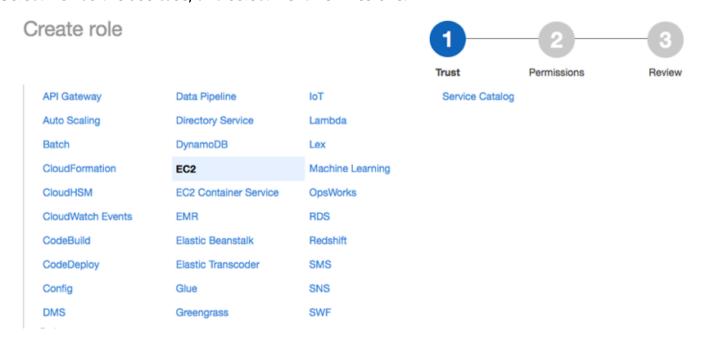


Figure 8: Select Amazon EC2 as your use case

5. Select Create Policy, and select JSON.

Create policy

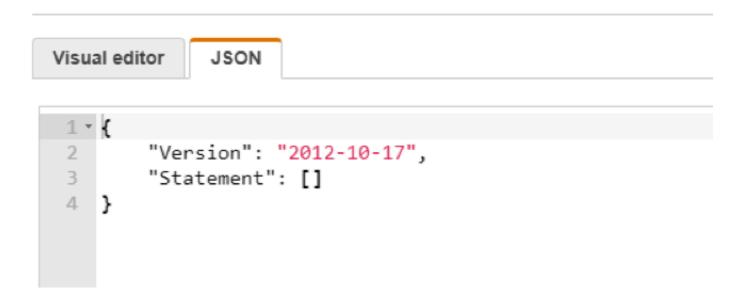


Figure 9: Create your own policy

6. Copy and paste the following policy into the input field, replace all existing text, and select **Review Policy.**



Note

If your Amazon EC2 instances are running in Beijing and Ningxia, you must update the **Resource line** with the correct region.

See the following example policies based on your Amazon Web Services Region.

Amazon Web Services Regions (except Amazon GovCloud (US-East), Amazon GovCloud (US-West), Beijing and Ningxia)

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Allow",
            "Action": [
                "EC2:DescribeInstances",
                "cloudwatch:GetMetricStatistics",
                "EC2:DescribeVolumes"
            ],
            "Resource": "*"
        },
        {
         "Sid": "VisualEditor1",
         "Effect": "Allow",
         "Action": "s3:GetObject",
         "Resource": [
          "arn:aws:s3:::aws-sap-dataprovider-<us-east-1>/config.properties"
            ]
        }
    ]
}
```

Beijing and Ningxia

```
"Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Allow",
            "Action": [
                "EC2:DescribeInstances",
                "cloudwatch:GetMetricStatistics",
                "EC2:DescribeVolumes"
            ],
            "Resource": "*"
        },
        {
         "Sid": "VisualEditor1",
         "Effect": "Allow",
         "Action": "s3:GetObject",
         "Resource": [
          "arn:aws-cn:s3:::aws-sap-dataprovider-<cn-north-1>/config.properties"
        }
    ]
}
```

Amazon GovCloud (US-East) and Amazon GovCloud (US-West)

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Allow",
            "Action": [
                "EC2:DescribeInstances",
                "cloudwatch:GetMetricStatistics",
                "EC2:DescribeVolumes"
            ],
            "Resource": "*"
        },
        {
         "Sid": "VisualEditor1",
         "Effect": "Allow",
         "Action": "s3:GetObject",
         "Resource": [
```

7. Provide a Name and Description for the role, and select Create Policy.

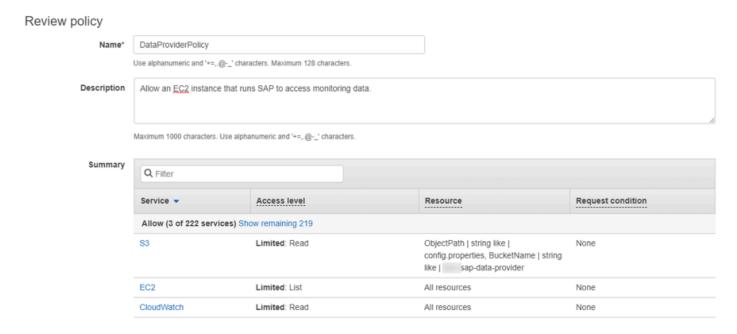


Figure 10: Provide name and description

8. Select **Create Policy**. The IAM console confirms the new policy with a message similar to the following.

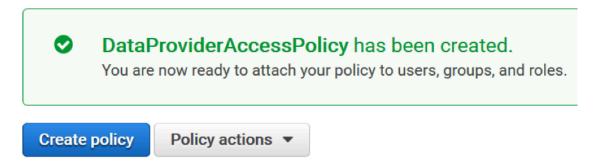


Figure 11: Policy confirmation

9. Navigate to the **Create Role** page, refresh the screen, search for the newly created role, and select the policy.



Figure 12: Policy confirmation

10Select Next:Tags.

11Add any tags if needed, otherwise select Next:Review.

12Provide a name for the Role and select Create Role.



Figure 13: Role confirmation

DataProvider 4.3

If you are new to Amazon Data Provider for SAP, see the section called "Installing DataProvider 4.3".

If you need to update or uninstall DataProvider 4.3, see the section called "Updating to DataProvider 4.3".

If you have an older version on your system, see the section called "Uninstalling older versions".



Important

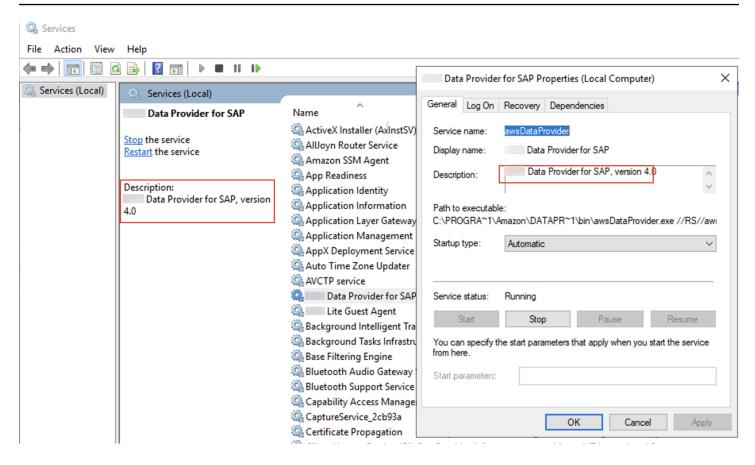
All the previous versions (v1, v2, v3) of the DataProvider have been deprecated and will no longer receive updates. For new DataProvider installations, you must install DataProvider 4.3 using an SSM distributor.

Run the following command to check the current version of DataProvider on your system.

```
rpm -qa | grep aws-sap
```

To check the current version of DataProvider on Windows, go to Services(Local), select Amazon Data Provider for SAP, and open Properties. You can see the current version in the Description field.

DataProvider 4.3



Installing DataProvider 4.3

The Amazon Data Provider for SAP runs as a service that automatically starts at boot and collects, aggregates, and exposes metrics to the SAP host agent. Metrics are sourced from a variety of providers that pull metrics from the relevant areas of the platform. The Amazon Data Provider for SAP is designed to continue operating, regardless of whether its providers have connectivity or permissions to access the Amazon service metrics they are requesting. Providers that cannot reach the metrics they are harvesting return blank values.

For example, if your Amazon EC2 instance does not have an IAM role associated with it that grants explicit access to the Amazon CloudWatch **GetMetricStatistics** API, the CloudWatch provider will be unable to perform the **GetMetricStatistics** action on the Amazon EC2 instance and will return blank values.

The provider needs to be installed on each SAP production system in order to be eligible for SAP support. You can only install one instance of the provider at a time on a system.

The Amazon Data Provider for SAP is designed to automatically update itself so that it can provide you with the most current metrics. When the Amazon Data Provider for SAP starts up, a built-

in update service retrieves the latest versions of its components and metric definitions from an Amazon managed Amazon S3 bucket. If the Amazon Data Provider for SAP cannot access the update service, it will continue to run as-is.

Installing with an SSM distributor – DataProvider 4.3 (Recommended)

The DataProvider 4.3 version enables you to install the package through a SSM distributor. Amazon recommends using this method for installation, you can install the DataProvider using a either a Linux or Windows platform.

Prerequisites for installing the DataProvider using a SSM distributor

SSM-Agent

You must have the ssm-agent installed on your instance before you can use the SSM distributor to install the DataProvider Agent. Use the following Amazon Systems Manager user guide to install the ssm-agent on your instances.

- RHEL: Manually installing SSM Agent on Red Hat Enterprise Linux instances
- SUSE: Manually install SSM Agent on SUSE Linux Enterprise Server instances
- Oracle: Manually installing SSM Agent on Oracle Linux instances
- Windows: Manually installing SSM Agent on Amazon EC2 instances for Windows Server

Java Runtime

The DataProvider is a java application that needs Java Runtime to be installed on the instance to run.

If your instance doesn't already have Java Runtime installed, you can use OpenJDK provided by Amazon Corretto to install the Java Runtime.

With DataProvider 4.3, the following Java Runtime versions are supported:

- Amazon Corretto 8 or OpenJDK 8
- Amazon Corretto 11 or OpenJDK 11
- Amazon Corretto 17 or OpenJDK 17

For more information on how to download and install JDK on your Amazon EC2 instance, see Amazon Corretto Documentation.

In the terminal, run the following commands to verify installation.

```
java -version
```

For example, expected output for Coretto-8.252.09.1 should be as follows:

```
openjdk version "1.8.0_252"OpenJDK Runtime Environment Corretto-8.252.09.1 (build 1.8.0_252-b09)OpenJDK 64-Bit Server VM Corretto-8.252.09.1 (build 25.252-b09, mixed mode)
```

GPG Key

If you are a SUSE user, you must download the DataProvider GPG key and import it before installation.

- GPG Key URL: GPG Key
- Sign in to your SUSE instance and run the following commands to import the key:

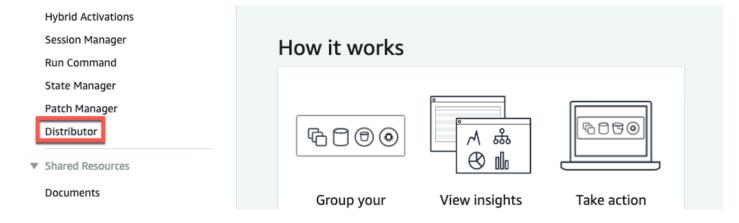
```
wget https//<url to GPG key>

rpm --import RPM-GPG-KEY-AWS
```

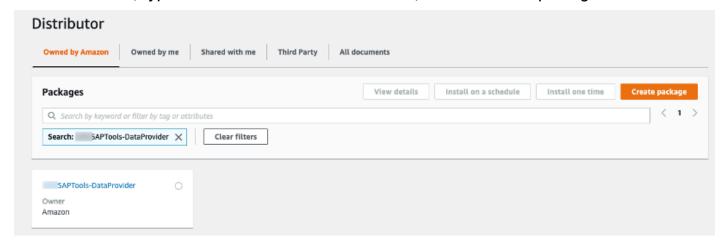
Installing the DataProvider Agent using an SSM distributor

Use the following procedure to install DataProvider 4.3.

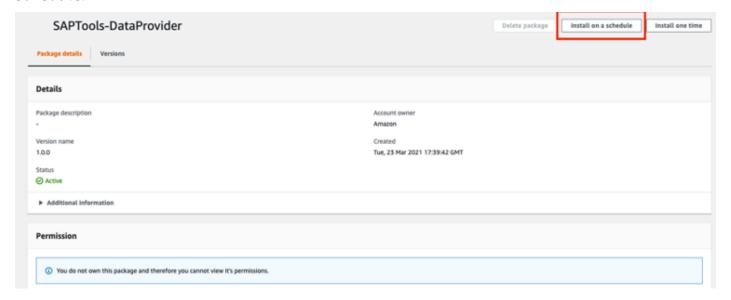
- 1. Open the Systems Manager console.
- 2. In the left navigation pane, under the Node Management section, choose **Distributor**.



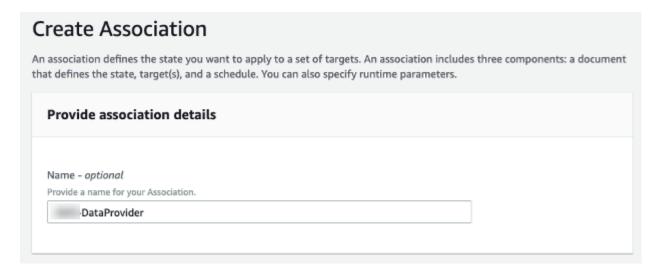
3. In the search bar, type AmazonSAPTools-DataProvider, and choose the package.



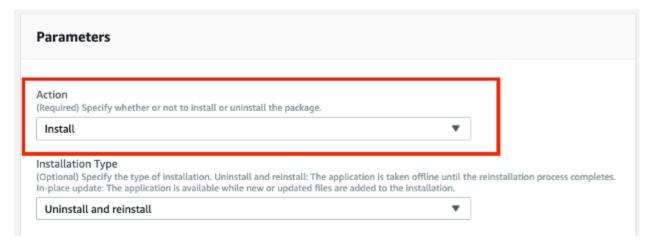
4. To receive auto updates for the DataProvider when there is a new release, choose **Install on a schedule**.



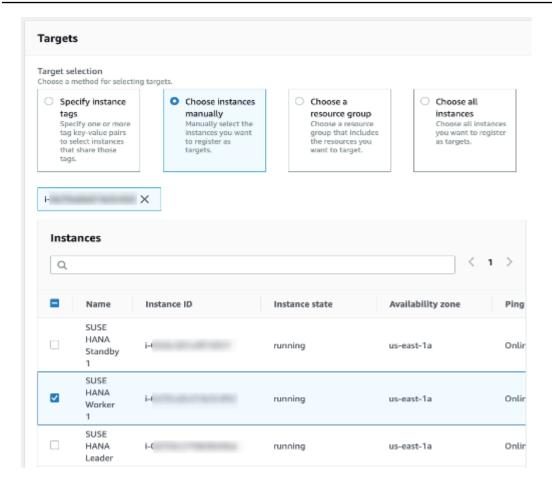
5. On the **Create Association** page, type a **Name** for your association.



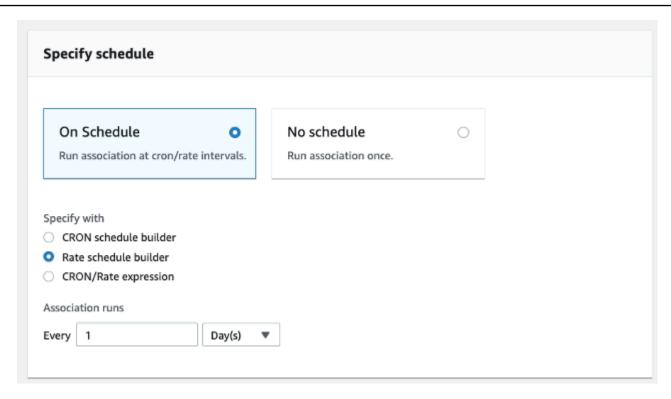
6. In the Parameters section, for Action, choose Install.



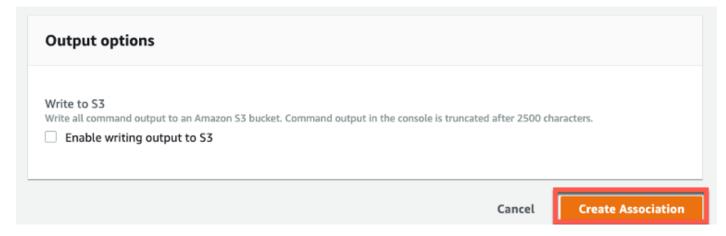
7. In the >Targets section, for Target selection, select Choose instances manually. Then, choose the instances where you want to install the DataProvider.



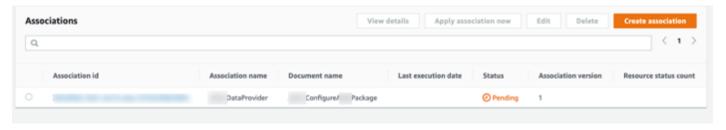
- 8. In the **Specify schedule** section, make the following selections:
 - Choose On Schedule
 - For Specify with, choose Rate schedule builder.
 - For Associate runs, choose 30 days. (Amazon recommends 30 days)



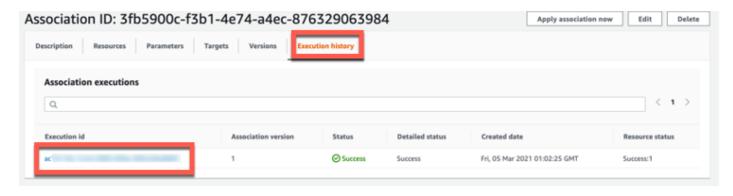
9. In the **Output options** section, choose **Create Association**.



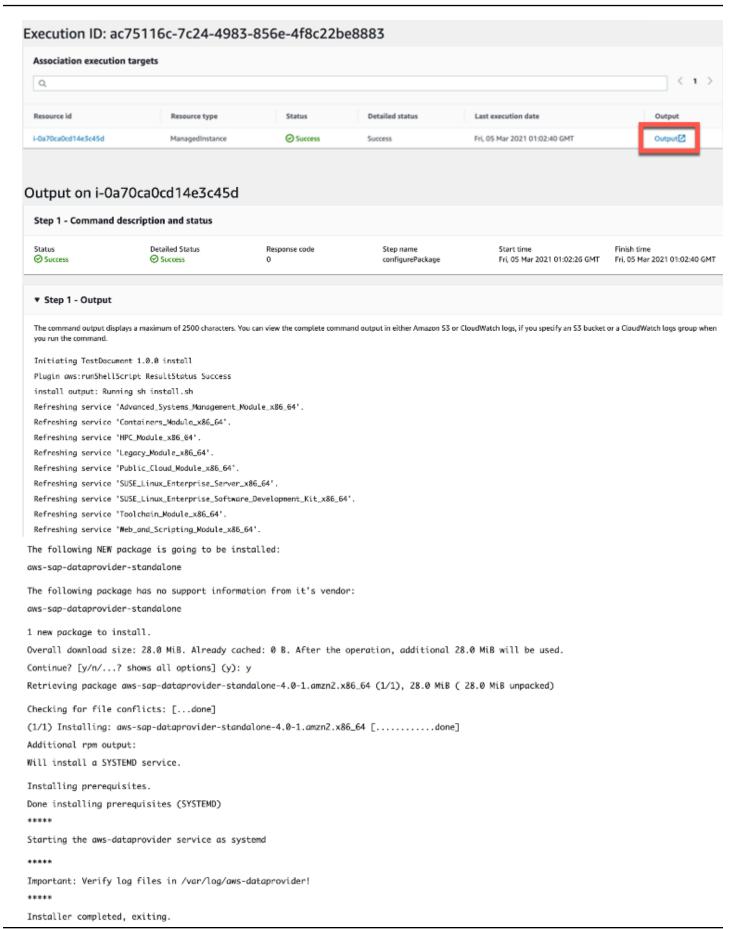
10Once the association is created, choose the **Association ID**.



11Choose the Execution History tab. Then, choose the Execution id.



12On the Execution ID page, choose Output to see the installation results.

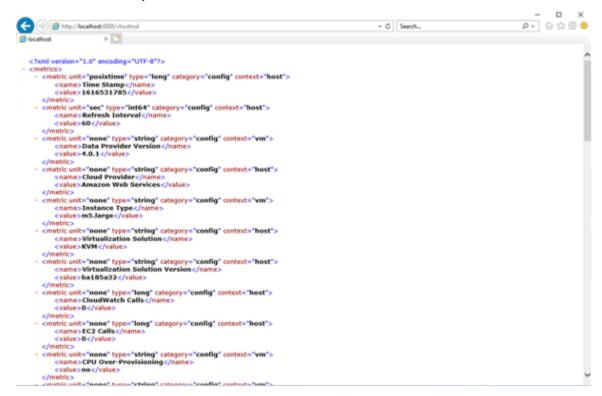


13Once the installation is completed, log in to the instance, and <u>call the endpoint</u> to enable the DataProvider to fetch metrics.

· Linux example

```
### A PRINT OF THE CONTROL OF THE PROPERTY OF
```

· Windows example



Installing with Downloadable Installer - DataProvider 4.3

If you choose to not use SSM to install the DataProvider 4.3, you can manually install the DataProvider using the following procedure.



Note

Before beginning the manual installation you must install the items listed in the Prerequisites section. You don't need to install the SSM-Agent. The downloadable DataProvider will not provide auto-updates, to get the latest versions, you must manually check for and download new releases manually.

Download the following file for your environment. By default the files will download in the useast-1 region, change the default region before downloading if you want the files to download in a different region.

- Red Hat https://aws-sap-dataprovider-us-east-1.s3.us-east-1.amazonaws.com/v4/installers/ linux/RHEL/aws-sap-dataprovider-rhel-standalone.x86_64.rpm
- SUSE https://aws-sap-dataprovider-us-east-1.s3.us-east-1.amazonaws.com/v4/installers/linux/ SUSE/aws-sap-dataprovider-sles-standalone.x86_64.rpm
- Oracle Linux https://aws-sap-dataprovider-us-east-1.s3.us-east-1.amazonaws.com/v4/ installers/linux/ORACLE/aws-sap-dataprovider-oel-standalone.x86_64.rpm
- Windows https://aws-sap-dataprovider-us-east-1.s3.us-east-1.amazonaws.com/v4/installers/ win/aws-data-provider-installer-win-x64-Standalone.exe
- **GPG Key** GPG Key: https://aws-sap-dataprovider-us-east-1.s3.us-east-1.amazonaws.com/ v4/installers/RPM-GPG-KEY-AWS

Installing on Linux

On Linux the data provider is delivered as an RPM package.

SUSE Linux Enterprise Server

To install the Amazon Data Provider for SAP on SUSE Linux Enterprise Server (SLES) download the following files:

Standard: aws-sap-dataprovider-sles.x86_64.rpm and GPG Key

• China: aws-sap-dataprovider-sles.x86_64.rpm and GPG Key

The files are identical but Amazon offers these two location options due to possible connectivity issues when working from China.

To install the data provider run the following commands:

```
wget https://<url to rpm package>
wget https://<url to GPG key>
rpm --import RPM-GPG-KEY-AWS
zypper install -y <rpm package>
```

Example:

```
wget https://aws-sap-dataprovider-us-east-1.s3.us-east-1.amazonaws.com/v4/installers/
linux/SUSE/aws-sap-dataprovider-sles-standalone.x86_64.rpm
wget https://aws-sap-dataprovider-us-east-1.s3.us-east-1.amazonaws.com/v4/installers/
RPM-GPG-KEY-AWS
rpm --import RPM-GPG-KEY-AWS
zypper install -y aws-sap-dataprovider-sles-standalone.x86_64.rpm
```

When the RPM package is installed, the agent starts as a daemon, as seen in the following image.

```
*****
Starting the aws-dataprovider service as systemd

*****
Important: Verify log files in /var/log/aws-dataprovider!

*****
Installer completed, exiting.
```

RPM package installation

Verify that the service is running by calling netstat -ant to determine if the listener is running on localhost port 8888.

```
-10-32-59-140:~ # netstat -ant
ctive Internet connections (servers and established)
roto Recv-Q Send-Q Local Address
                                              Foreign Address
                                                                       State
          0
                                              0.0.0.0:*
                                                                       LISTEN
ср
          0
                  0 0.0.0.0:8888
                                              0.0.0.0:*
ср
                                                                       LISTEN
сp
                                              0.0.0.0:*
                  0 12/.0.0.1:25
          0
                  0 0.0.0.0:111
                                             0.0.0.0:*
                                                                       LISTEN
сp
                 52 10.32.59.140:22
                                              69.120.21.212:49759
                                                                       ESTABLISHE
```

Verifying installation on Linux

You should also view the log files at /var/log/Amazon-dataprovider/messages to ensure the daemon has the appropriate connectivity and authorization to access the required metrics.

```
I 0001C Agent has started with version 4.2.0 : Tue Aug 25 18:13:00 UTC 2020
I 0000C Agent log level has been set to INFO
I 10001 Read in jar configuration; read on board configuration
I 10002 Is there a local, optional configuration file at /usr/local/ec2/aws-dataprovider/config.properties ?
I 10003 Read local, optional configuration file from /usr/local/ec2/aws-dataprovider/config.properties
I 03002 vhostmd agent is listening on localhost
I 0000D Agent is starting the vhostmd provider
I 08001 ** Running Diagnostics **
I 08006 Diagnostic : Amazon CloudWatch Connectivity & Access
I 08009 Diagnostic : Passed
I 08000 Diagnostic : Passed
I 08010 Diagnostic : Passed
I 08011 Diagnostic : Data Collector API
I 08011 Diagnostic : Passed
I 0800E ** Diagnostics Complete **
```

Verifying connectivity and authorization on Linux

At startup, the monitoring agent runs three sets of diagnostics:

- The Amazon connectivity diagnostic ensures network connectivity to Amazon S3 for obtaining automatic updates to the Amazon Data Provider for SAP.
- The second diagnostic tests for authorization to access CloudWatch. This authorization requires assigning an IAM role to the Amazon EC2 instance you are running on with an IAM policy that allows access to CloudWatch. For details, see IAM Roles, earlier in this guide.

• The third diagnostic tests for authorization to access Amazon EC2, which also requires an IAM role associated with the Amazon EC2 instance.

The Amazon Data Provider for SAP is designed to run with or without connectivity, but you can't obtain updates without connectivity. Amazon CloudWatch and Amazon EC2 will return blank values if you don't have the proper authorizations in place.

You can also call the Amazon Data Provider for SAP directly to view the metrics. Calling wget http://localhost:8888/vhostmd returns a file of metrics. You can look inside the file to see the metrics that were returned, as shown here.

```
ip-10-32-59-140:~ # wget http://localhost:8888/vhostmd
--2012-10-09 17:10:28-- http://localhost:8888/vhostmd
Resolving localhost... 127.0.0.1, ::1
Connecting to localhost | 127.0.0.1 | :8888... connected.
HTTP request sent, awaiting response... 200 OK
Length: unspecified [text/xml]
Saving to: 'whostmd'
    [ <=>
                                                                    ] 7,58
2012-10-09 17:10:31 (385 MB/s) - `vhostmd' saved [7589]
ip-10-32-59-140:~ # cat vhostmd
<?xml version="1.0" encoding="UTF-8"?>
  <metrics>
    <metric context="host" category="config" type="long" unit="posixtime"</pre>
     <name>Time Stamp</name>
     <value>1349802629284
    <metric context="host" category="config" type="int64" unit="sec">
      <name>Refresh Interval</name>
      <value>60</value>
    </metric>
    <metric context="host" category="config" type="string" unit="none">
      <name>Cloud Provider</name>
      <value>Amazon Web Services</value>
    </metric>
    <metric context="host" category="config" type="string" unit="none">
      <name>Instance Type</name>
      <value>m1.large</value>
    </metric>
```

Viewing metrics on Linux

The Amazon Data Provider for SAP now starts automatically each time the operating system starts. You can also manually stop and restart the Amazon Data Provider for SAP with the following command, which depends on your operating system version:

• SLES 11, Oracle Linux 6, and Red Hat Linux 6:

```
service aws-dataprovider [start|stop]
```

• SLES 12, SLES 15, Oracle Linux 7, Oracle Linux 8, Red Hat Linux 7, and Red Hat Linux 8.

```
systemctl [start|stop] aws-dataprovider
```

You can configure Amazon Data Provider to use proxies if you do not have transparent HTTP/ HTTPS access to the internet.

- 1. Stop Amazon Data Provider for SAP.
- 2. Enter proxy information in the file (as seen below) located at /usr/local/ec2/aws-dataprovider/proxy.properties.

```
# proxy.properties
# used to set web proxy settings for the AWS Data Provider for SAP
# Https is the only supported proxy method
# Blank values for everything means no proxy set
https.proxyHost=
https.proxyPort=
https.proxyDomain=
https.proxyUsername=
https.proxyPassword=
```

3. Start Amazon Data Provider for SAP.

Installing on Red Hat and Oracle Enterprise Linux

For Red Hat and Oracle Enterprise Linux, the installation steps are the same as described for SLES above but the RPM file and command to install the RPM package differs.

Red Hat

Default: aws-sap-dataprovider-rhel.x86_64.rpm

Oracle Enterprise Linux

Default: aws-sap-dataprovider-oel.x86_64.rpm

To install the data provider run the following commands:

```
wget https://<url to rpm package>
yum -y install <rpm package>
```

Example:

```
wget https://aws-sap-data-provider.s3.amazonaws.com/Installers/aws-sap-dataprovider-
rhel.x86_64.rpm
yum -y install aws-sap-dataprovider-rhel.x86_64.rpm
```

Installing on Windows

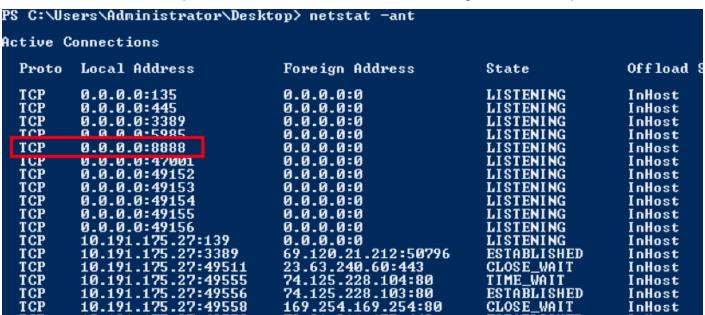
On Windows, the installer is delivered in the form of an NSIS (Nullsoft Scriptable Install System) executable.

- 1. Open a web browser and download the installer:
 - **Default:** aws-data-provider-installer-win-x64.exe
- 2. Run the downloaded exe file.
- 3. Verify the installation.
 - Once the installation is complete, you can see the file in C:\Program Files\Amazon
 \DataProvider directory.
 - The installation also creates and starts a Windows service called Amazon Data Provider for SAP.
 - Verify that the service is running by entering http://localhost:8888/vhostmd in a web browser. The page returns metrics from Amazon Data Provider for SAP if your installation is successful.
- 4. You can configure Amazon Data Provider to use proxies if you do not have transparent HTTP/ HTTPS access to the internet.
 - a. Stop Amazon Data Provider for SAP.

b. Enter proxy information in the file (as seen below) located at C:\Program Files\Amazon \DataProvider\proxy.properties.

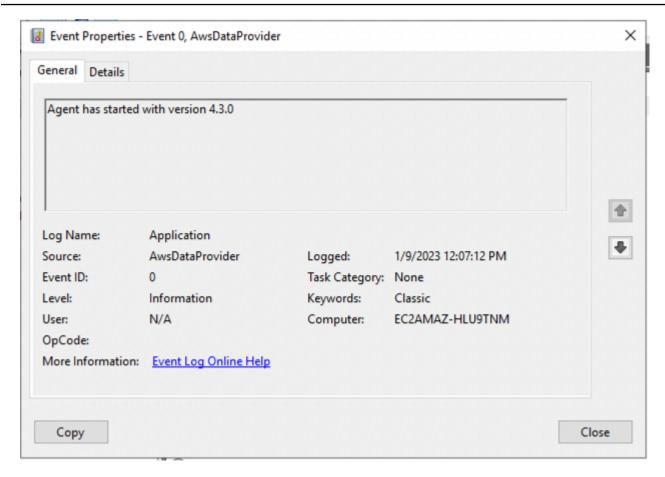
```
# proxy.properties
# used to set web proxy settings for the AWS Data Provider for SAP
# Https is the only supported proxy method
# Blank values for everything means no proxy set
https.proxyHost=
https.proxyPort=
https.proxyDomain=
https.proxyUsername=
https.proxyPassword=
```

- c. Start Amazon Data Provider for SAP.
- 5. Verify that the service is running by calling netstat -ant from a command window or from a Windows PowerShell script to determine if the listener is running on localhost port 8888.



Verifying the installation on Windows

6. Navigate to the Windows event log, and find the application log for startup events from the Amazon Data Provider for SAP. Check the diagnostics.



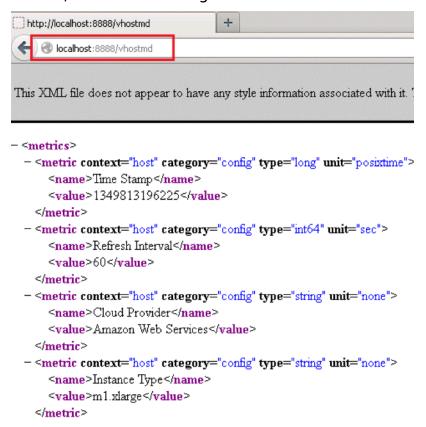
Checking diagnostics on Windows

At startup, the monitoring agent runs three sets of diagnostics:

- The Amazon connectivity diagnostic ensures network connectivity to Amazon S3 for obtaining automatic updates to the Amazon Data Provider for SAP.
- The second diagnostic tests for authorization to access CloudWatch, which requires assigning
 an IAM role to the EC2 instance you are running on with an IAM policy that allows access to
 CloudWatch. For details, see IAM Roles, earlier in this guide.
- The third diagnostic tests for authorization to access Amazon EC2, which also requires an IAM role associated with the Amazon EC2 instance.

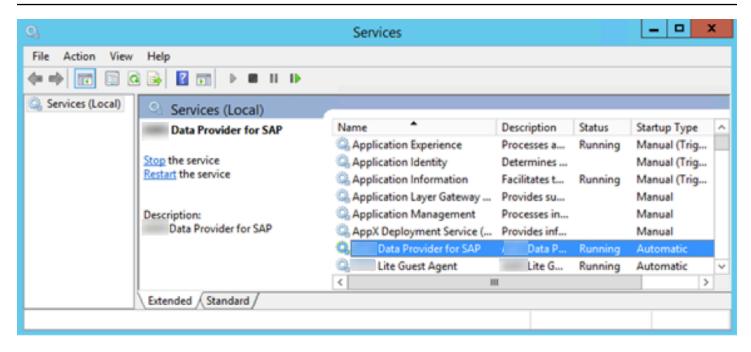
The Amazon Data Provider for SAP is designed to run with or without connectivity, but you can't obtain updates without connectivity. If you don't have the proper authorizations in place, Amazon CloudWatch and Amazon EC2 return blank values.

You can also call the Amazon Data Provider for SAP directly from your web browser to view metrics, as shown in the figure.



Viewing metrics on Windows

Amazon Data Provider for SAP now starts automatically each time the operating system starts. You can also manually stop and restart the Amazon Data Provider for SAP, just as you would stop and restart any other Windows service.



Stopping and restarting the Amazon Data Provider for SAP on Windows

In order to configure proxy settings, you can place a customized proxy.properties file in Window's temp directory, which is designated by the windows system variable %TEMP%.

Subscribe to Amazon Data Provider Agent for notifications

Amazon Simple Notification Service can notify you when new versions of Amazon Data Provider Agent are released. Use the following steps to setup this subscription.

- 1. Open https://console.amazonaws.cn/sns/v3/home.
- 2. Ensure that you are in **US East N. Virginia** (us-east-1) Region.
- 3. In the left navigation pane, select **Subscriptions** > **Create subscription**.
- 4. Add a **Topic ARN** based on the Amazon Web Services Region in which you are using Amazon Data Provider Agent.

Region	ARN
Default	arn:aws:sns:us-east-1:80484 5276281:AWS-DataProvider-SAP- Update

Region	ARN
Amazon GovCloud (US-West) and Amazon GovCloud (US-East)	arn:aws-us-gov:sns:us-gov-w est-1:140982767562:AWS-Data Provider-SAP-Update
China (Beijing) Region and China (Ningxia) Region	arn:aws-cn:sns:cn-north-1:0 01645243879:AWS-DataProvider- SAP-Update

- **Protocol** choose Email or SMS.
 - Email enter an email address where you would like to receive the notification in the **Endpoint** field.



Note

To enable email notifications, you must confirm your email subscription by following the instructions you receive on the provided email address.

- SMS enter a phone number where you would like to receive the notification in the **Endpoint** field.
- Choose Create subscription. You can now receive notifications whenever a new version of Amazon Data Provider Agent is released.

To unsubscribe from notifications, use the following steps.

- 1. Open https://console.amazonaws.cn/sns/v3/home.
- 2. In the left navigation pane, select **Subscriptions**.
- Select the subscription from your list of subscriptions and choose **Delete**.

Updating to DataProvider 4.3

If you have previously installed DataProvider 2.0 or 3.0 and want to update to DataProvider 4.3, you need to uninstall the running version first, and then install DataProvider 4.3.

Updating to DataProvider 4.3 using the SSM Distributor package

When you install DataProvider 4.3 through the SSM distributor, it will auto-update the installed package on a new version release. For more information, see Install or update packages.

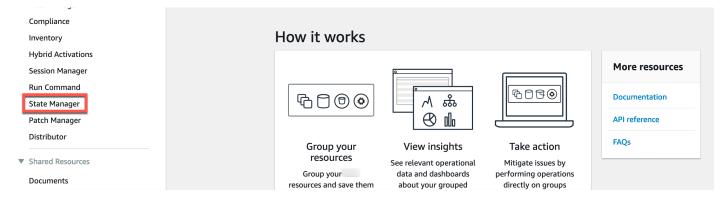
To update the DataProvider 4.3 manually, you must first uninstall the running version and then install the updated version.

DataProvider 4.3 does not support the following actions:

- Manual update of the RPM package if the DataProvider has been installed using the SSM distributor.
- Automatic update through the SSM distributor if the DataProvider has been manually installed using RPM.
- Manual update of the RPM package in any scenario.

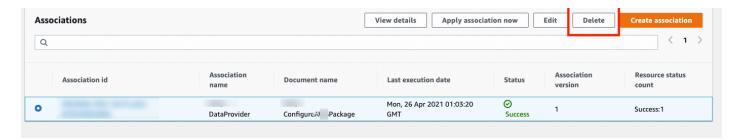
Uninstall DataProvider 4.3 using the SSM distributor

Open the <u>Amazon Systems Manager</u> console, on the left navigation pane, choose **State** Manager.

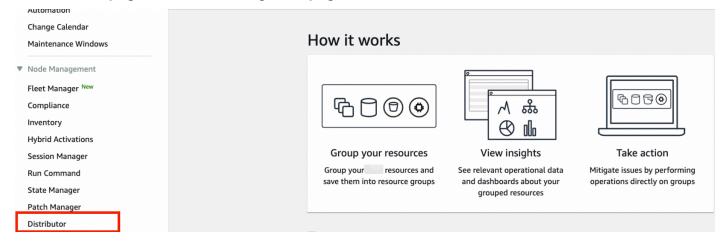


2. On the Associations page, and choose the Association id. Then, choose Delete.

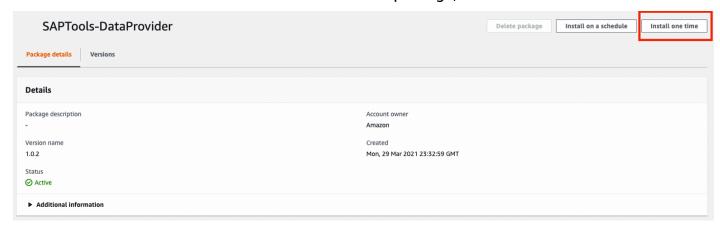
After the delete is successful, the auto-update stops.



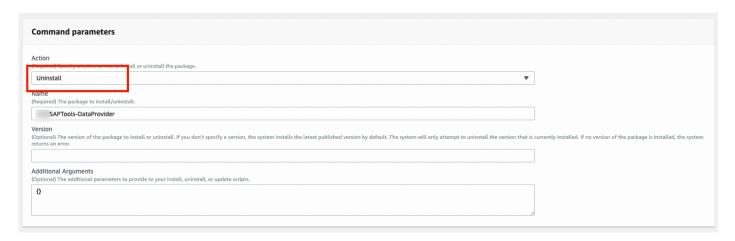
3. On the main page, in the left navigation page, choose **Distributor**.



4. Choose the AWSSAPTools-DataProvider distributor package, and choose Install one time.

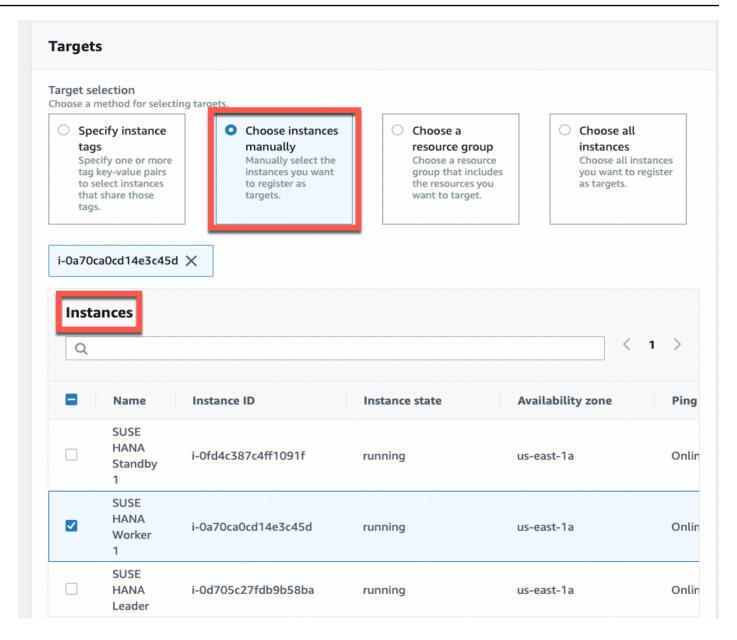


5. In the Command parameters section, choose Uninstall.



6. In the **Targets** section, select **Choose instances manually**. Then choose the **instance** to uninstall the DataProvider.

Updating to DataProvider 4.3 102



7. Choose **Run** to begin the uninstall.

Uninstall DataProvider 4.3 manually

RedHat

yum erase aws-dataprovider-standalone

SUSE

zypper rm aws-dataprovider-standalone

Updating to DataProvider 4.3 103

Windows

Run the uninstaller.

```
C:\Program Files\AmazonA\DataProvider\uninstall.exe
```

2. When prompted, choose **Uninstall**.

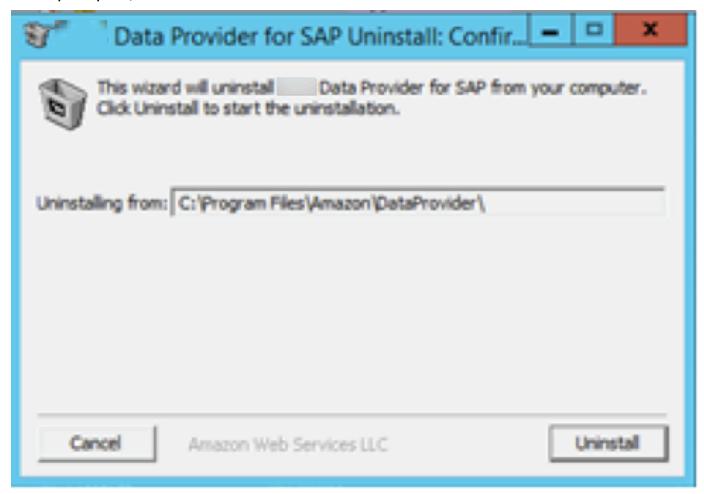


Figure 26: Uninstalling the Amazon Data Provider for SAP on Windows

Uninstalling older versions

Uninstalling the Amazon Data Provider for SAP does not require SAP downtime and can be done online. The only impact will be a gap in metric monitoring information for the time that the DataProvider was installed on your system.

Uninstalling older versions 104

Uninstall DataProvider 3.0

Linux

- 1. Log in to Linux as a superuser, like root.
- 2. Stop and remove the DataProvider using the following command.

SLES

```
zypper remove -y aws-sap-dataprovider
```

RHEL/OEL

```
yum -y erase aws-sap-dataprovider
```

Windows

```
"C:\Program Files\Amazon\DataProvider\uninstall.exe"
```

Uninstall DataProvider 2.0

Linux

- 1. Log in to Linux as a superuser, like root.
- 2. Stop and remove the DataProvider using the following command.

```
/usr/local/ec2/aws-agent/bin/aws-agent_uninstall
```

Windows

```
"C:\Program Files\Amazon\DataProvider\uninstall.exe"
```

Troubleshooting

This section provides help to analyze installation problems.

Troubleshooting 105

Troubleshooting on Linux

Problem: The installation failed, and I'm not sure if my files are in a consistent state.

Stop and remove the Data Provider with the following command.

SLES:

```
zypper remove -y aws-sap-dataprovider
```

RHEL / OEL:

```
yum -y erase aws-sap-dataprovider
```

Problem: The Amazon Data Provider for SAP failed to start at the end of the installation process.

Check the log files in /var/log/aws-dataprovider for hints on what is not going as expected. If needed, uninstall and reinstall the Data Provider. If reinstalling the Amazon Data Provider for SAP doesn't solve the problem, you can gather debug information about the Amazon Data Provider for SAP by editing the /usr/local/ec2/aws-dataprovider/bin/aws-dataprovider file.

```
aws-dataprovider
$local_fs $network
$local_fs
  Provides:
  Required-Start:
  Required-Stop:
  Default-Stop:
Default-Stop:
                           3 5
0 1 6
  Short-Description: aws-dataprovider
### END INIT INFO
## Determine Download bucket by region
REGION="$(wget --no-proxy -q -0 /dev/stdout <u>http://169.254.169.254/latest/dynamic/instance-identity/document</u> | awk -F\" '/region/ {
print $4}' )"
 f [[ $REGION == *"cn"* ]]
   REGIONENDPOINT="s3.cn-northwest-1.amazonaws.com.cn"
   MYBUCKET="aws-sap-data-provider-china"
   REGIONENDPOINT="s3.amazonaws.com"
   MYBUCKET="aws-sap-data-provider
MYS3="https://${REGIONENDPOINT}/${MYBUCKET}"
DAEMON="/usr/local/ec2/aws-dataprovider/bin/jsvc"
DAEMONOPTS="-jvm server -home /usr/local/ec2/aws-dataprovider/jre/amazon-corretto-8.242.08.1-linux-x6
der.pid -cp /usr/local/ec2/aws-dataprovider/lib/installer.jar:/usr/local/ec2/aws-dataprovider/lib/comm
                                                                                                                                            -pidfile /tmp/aws-dataprovi
                                                                                                                                             s-daemon-1.0.10/commons-dae
                                                                                                                                             ws.aws-dataprovider.access
-${MYS3} com.amazon......dataprovider.AwsDataProvider /usr/local/ec2/aws-dataprovider -vhostmd -LogLevel=INFO
```

Figure 27: Debugging the installation on Linux

Now if you run service aws-dataprovider-start or systemctl start aws-dataprovider, you will get a lot of debugging output that might help you diagnose the root cause of the problem.

```
Java VM created successfully
Class org/apache/commons/daemon/support/DaemonLoader found
Native methods registered
java init done
Daemon loading...
Daemon loaded successfully
java load done
18:05:28.561 I 06002 Starting the data collector engine
18:05:29.230 I 00001 *****
                             SAP Data Collector Agent is Starting *****
18:05:29.237 I 00002 Software Version: 1.0.47 : Fri Sep 21 22:22:00 UTC 20:
18:05:29.239 I 0000C Agent log level has been set to FINE
18:05:29.241 I 08001 ** Running Diagnostics **
18:05:29.241 I 08002 Diagnostic : Connectivity
18:05:30.078 I 08005 Diagnostic : Passed
18:05:30.078 I 08006 Diagnostic : Amazon CloudWatch Connectivity & Access
18:05:30.803 I 08009 Diagnostic : Passed
18:05:30.803 I 0800A Diagnostic : EC2 API Connectivity & Access
18:05:31.082 I 0800D Diagnostic : Passed
18:05:31.083 I 0800E ** Diagnostics Complete **
18:05:31.282 I 03002 vhostmd agent is listening on localhost
18:05:31.282 I 0000D Agent is starting the vhostmd provider
```

Figure 28: Debugging information on Linux

Problem: When I looked at my logs I noticed that my installation failed all diagnostics.

```
14:32:15.862 I 08001 ** Running Diagnostics **

14:32:15.862 I 08002 Diagnostic : Connectivity

14:33:19.362 W 08003 Diagnostic : Failed

14:33:19.362 I 08006 Diagnostic : Amazon CloudWatch Connectivity & Access

14:33:19.515 W 08007 Diagnostic : Failed

14:33:19.516 I 0800A Diagnostic : EC2 API Connectivity & Access

14:33:19.542 W 0800B Diagnostic : Failed

14:33:19.542 I 0800E ** Diagnostics Complete **
```

Figure 29: Symptoms of internet connectivity problems on Linux

Failing *all* diagnostics indicates that there's a problem with your outbound connection to the internet. You can confirm this by pinging a well-known internet location, like <u>www.amazon.com</u>. The most common cause of routing issues is in the VPC network configuration, which needs to have either an internet gateway in place or a VPN connection to your data center with a route to the internet. For details, see <u>Amazon VPC Network Topologies</u>, earlier in this guide.

Problem: When I looked at my logs I noticed that I don't have access to CloudWatch and Amazon EC2, but I did pass the first diagnostic for Amazon connectivity.

```
14:38:57.467 I 08001 ** Running Diagnostics **

14:38:57.468 I 08002 Diagnostic : Connectivity

14:38:58.182 I 08005 Diagnostic : Passed

14:38:58.182 I 08006 Diagnostic : Amazon CloudWatch Connectivity & Access

14:38:58.325 W 08007 Diagnostic : Failed

14:38:58.325 I 0800A Diagnostic : EC2 API Connectivity & Access

14:38:58.357 W 0800B Diagnostic : Failed

14:38:58.357 I 0800E ** Diagnostics Complete **
```

Figure 30: Symptoms of authorization issues on Linux

This is a clear indicator that you have an authorization issue when trying to access CloudWatch and Amazon EC2. The common cause for this problem is not having an IAM role associated with your instance that contains the IAM policy, as specified in <u>IAM Roles</u>, earlier in this guide. You can quickly diagnose this issue by looking at the Amazon EC2 instance in question in the Amazon EC2 console and verifying the IAM role.

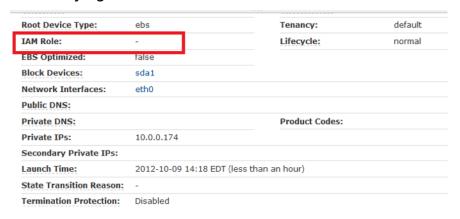


Figure 31: Verifying the IAM role for an EC2 instance

If the IAM role doesn't exist, then create it as specified in IAM Roles, which is described earlier in this guide.

If you do have an IAM role assigned to the instance, go to the IAM console, select the IAM role name, and then expand the policy. Verify that you have the required policy that is specified in IAM Roles, earlier in this guide.

```
Permissions
              Policy usage
                             Policy versions
                                              Access Advisor
  Policy summary { } JSON
                                Edit policy
            "Version": "2012-10-17",
            "Statement": [
                     "Sid": "VisualEditor0",
                     "Effect": "Allow",
                     "Action": [
                         "EC2:DescribeInstances",
                         "cloudwatch:GetMetricStatistics",
                         "EC2:DescribeVolumes"
                     "Resource": "*"
                    "Sid": "VisualEditor1",
                    "Effect": "Allow",
"Action": "s3:GetObject",
                    "Resource": [
                          arn:aws:s3:::aws-data-provider/config.properties"
```

Figure 32: Verifying the policy for the IAM role

Problem: I want to configure/update JAVA_HOME for Data Provider.

Open the /usr/local/ec2/aws-dataprovider/env file and update the JAVA_HOME variable. Restart Data Provider after the update, with the following command.

```
sudo systemctl daemon-reload
sudo systemctl start aws-dataprovider
```

Troubleshooting on Windows

Problem: The installation failed, and I'm not sure if my files are in a consistent state.

Based on the DataProvider version on your system, follow the procedure in the section called "Updating to DataProvider 4.3" or the section called "Uninstalling older versions".

Problem: The Amazon Data Provider for SAP failed to start at the end of the installation process.

If reinstalling the Amazon Data Provider for SAP doesn't solve the problem, you can gather debugging information about the Amazon Data Provider for SAP by reviewing the log files in the C:\Program Files\Amazon\DataProvider directory.

These log files include an installation log, a log of the service installation, and the output of the Amazon Data Provider for SAP itself.

```
LastWriteTime
                          Length Name
10/9/2012
             4:20
                                 aws-agent
10/9/2012
             4:26
                  PM
                                 agentinstallog.txt
                            2203
201
10/9/2012
             4:26
                  PM
                                     agent-stdout.2012-10-09.log
10/9/2012
             4:26
                                     DataCollectorLibraryUpdateService.txt
10/9/2012
                            1267
                                 commons-daemon.2012-10-09.log
```

Figure 33: Log files on Windows

Problem: I want to get more detailed log information from the data provider.

Start by stopping the data provider service.

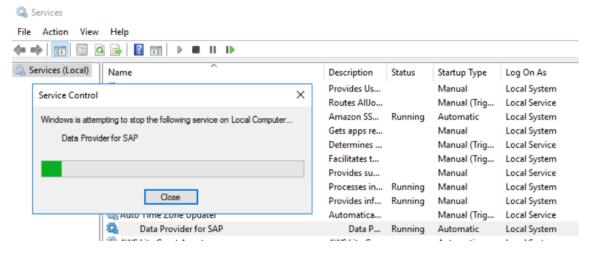


Figure 34: Stop Service on Windows

Open the registry editor by clicking on the Windows logo in the bottom left and typing regedit and then clicking the option that is shown on screen:



Figure 35: Start regedit

In the registry, navigate to the key:

HKEY_LOCAL_MACHINE\SOFTWARE\WOW6432Node\Apache Software Foundation\Procrun 2.0\awsDataProvider\Start Registry Editor File Edit View Favorites Help Policies Name Data RegisteredApplications (Default) REG_SZ (value not set) Volatile ab Class REG_SZ DataProvider WOW6432Node **ab** Method REG_SZ windowsService Apache Software Foundation ab Mode REG_SZ jvm Procrun 2.0 **ab** Params REG_MULTI_SZ start C:\PROGRA~1\Amazon\DATAPR~1 -Xmx64m -Xms64m -vhostmd -LogLevel=INFO awsDataProvider Parameters Java Log Start Stop

Figure 36: Logging setting

awsDataProvider

The data provider accepts two log levels: INFO and FINE. FINE will generate more detailed logging which can be useful when debugging a problem. The recommendation is to set it back to INFO when you have finished troubleshooting to avoid the logs consuming unnecessary disk space.

Problem: I want to reinstall the Amazon Data Provider for SAP from scratch.

Based on the DataProvider version on your system, follow the procedure in the section called "Updating to DataProvider 4.3" or the section called "Uninstalling older versions".

Problem: When I looked at my logs, I noticed that my installation failed all diagnostics.

```
14:32:15.862 I 08001 ** Running Diagnostics **

14:32:15.862 I 08002 Diagnostic : Connectivity

14:33:19.362 W 08003 Diagnostic : Failed

14:33:19.362 I 08006 Diagnostic : Amazon CloudWatch Connectivity & Access

14:33:19.515 W 08007 Diagnostic : Failed

14:33:19.516 I 0800A Diagnostic : EC2 API Connectivity & Access

14:33:19.542 W 0800B Diagnostic : Failed

14:33:19.542 I 0800E ** Diagnostics Complete **
```

Figure 37: Symptoms of internet connectivity problems on Windows

Failing *all* diagnostics indicates that there's a problem with your outbound connection to the internet. You can confirm this by pinging a well-known internet location, like <u>www.amazon.com</u>. The most common cause of routing issues is in the VPC network configuration, which needs to have either an internet gateway in place or a VPN connection to your data center with a route to the internet.

Problem: When I looked at my logs, I noticed that I don't have access to CloudWatch and Amazon EC2, but I did pass the first diagnostic for Amazon connectivity.

```
14:38:57.467 I 08001 ** Running Diagnostics **

14:38:57.468 I 08002 Diagnostic : Connectivity

14:38:58.182 I 08005 Diagnostic : Passed

14:38:58.182 I 08006 Diagnostic : Amazon CloudWatch Connectivity & Access

14:38:58.325 W 08007 Diagnostic : Failed

14:38:58.325 I 0800A Diagnostic : EC2 API Connectivity & Access

14:38:58.357 W 0800B Diagnostic : Failed

14:38:58.357 I 0800E ** Diagnostics Complete **
```

Figure 38: Symptoms of authorization issues on Windows

This is a clear indicator that you have an authorization issue when trying to access Amazon CloudWatch and Amazon EC2. The common cause for this problem is not having an IAM role associated with your instance that contains the IAM policy, as specified in IAM Roles earlier in this guide. You can quickly diagnose this issue by looking at the specific EC2 instance in the Amazon EC2 console and verifying the IAM role.

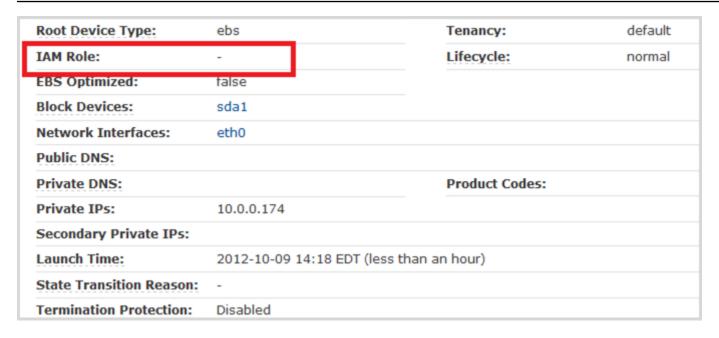


Figure 39: Verifying the IAM role for an EC2 instance

If the IAM role doesn't exist, then create is as specified in IAM Roles, described earlier in this guide.

If you do have an IAM role assigned to the instance, go to the IAM console, select the IAM role name, and then choose **Show**. Verify that you have the required policy that is specified in IAM Roles.

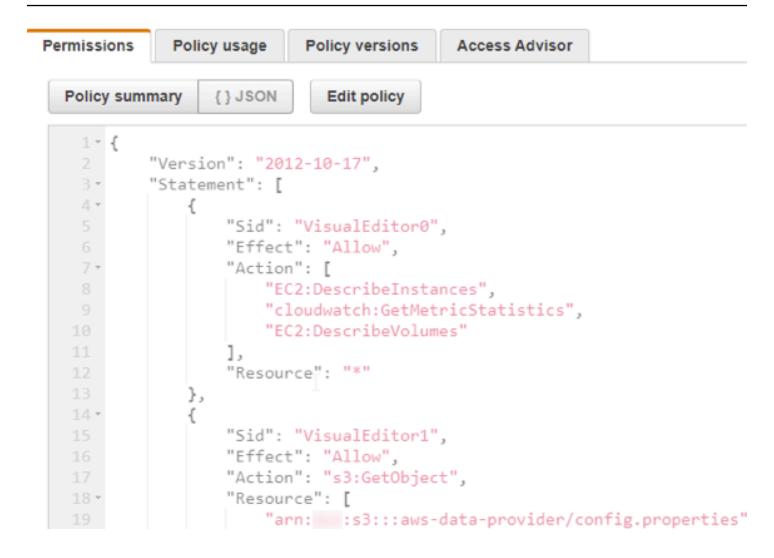


Figure 40: Verifying the policy for the IAM role

Customizing the Amazon Data Provider for SAP

Some settings are hard coded in the Amazon Data Provider for SAP. You can override existing settings or add new settings. For example, when Amazon adds new instance types, you can add these to the Amazon Data Provider for SAP configuration.

The Amazon Data Provider for SAP creates a database by reading the configuration information from the following files, in this sequence:

• It reads the config.properties file from the JAR (Java Archive) file of the data provider application.

• It reads the file from the installation directory in which the proxy configuration file is located.

This file is required only if a user wants to override or extend the current configuration. Default locations are:

- On Linux: /usr/local/ec2/aws-dataprovider/config.properties
- On Windows: C:\Program Files\Amazon\DataProvider\config.properties
- It reads the file from https://s3.amazonaws.com/aws-data-provider/config.properties.

Syntax Rules for Configuration Files

- The configuration files require a comma after the last value in every row.
- Spaces are not ignored in strings. The entire string between the commas, including any spaces, is accepted as the value.
- If there are multiple rows with the same instance type, the existing value for that type is overwritten.
- Capitalization in strings is case sensitive.

User-Configurable EC2 Instance Types

The Amazon Data Provider for SAP maintains a database of all relevant Amazon EC2 instance types for SAP.

Entries for EC2 instance types must be in a comma-separated list, as follows:

ec2type,i-type,cpu,core,threads,t-ecu,ecu,hthread,l-map,w-map,speed,p-ecu,

For example:

```
ec2type,r3.8xlarge,2,16,2,32,1,thread,eth0,lan2,10000,true,
```

where the following applies:

Field name	Content	Example	Туре	Description
keyword	ec2type	_	String	A token to identify a

Field name	Content	Example	Туре	Description
				record with an EC2 instance description
i-type (instance -type)	See list	r3.8xlarge	String	Instance type, which must match the EC2 instance metadata string
cpu (CPUs)	1 2	2	Integer	Number of sockets
core (Cores)	integer	16	Integer	Total number of processor cores
threads (threads per core)	1 2	2	Integer	Threads per core
t-ecu (total ECU value)	integer	32	Double	ECU value for previous- generation instance types that have ECU ratings; number of cores for post-ECU instance types

Field name	Content	Example	Туре	Description
ecu (ECU per core)	double	1	Double	1 for all post- ECU instance types; total ECU divided by cores for previous- generation instance types that have ECU ratings
hthread (hyperthr eading)	thread core	thread	String	thread for hyperthreaded instance types; core for non- hyperthreaded instance types
l-map (Linux NIC mapping)	eth0	eth0	String	Linux mapping of network interface
w-map (Windows NIC mapping)	eth0	lan2	String	Windows mapping of network interface
speed (network interface speed)	1000 2000 10000	100000	Integer	Maximum speed of network interface, in KB
p-ecu (post ECU)	true false	true	Boolean	true for modern instances that don't have ECU ratings

User-Configurable EBS Volume Types

The Amazon Data Provider for SAP maintains a database of all relevant EBS volume types for SAP.

Entries for EBS volume types must be in a comma-separated list, as follows:

voltype,ebs-type,sample-time,

For example:

voltype, io1,60,

where the following applies:

Field name	Content	Example	Туре	Description
keyword	voltype		String	A token to identify a record with an EBS volume description
ebs-type (EBS- type)	io1 gp2 sc1 st1	lo1	String	EBS type, which must match the EBS volume metadata string
sample-time	60 300	60	Integer	CloudWatch sample time, in seconds

Important

The sample time is required to calibrate the EBS metrics to the SAP monitoring requirements. Changes in the sample time will lead to incorrect EBS metrics in the SAP monitoring system.

Verification of Amazon Data Provider for SAP in SAP system monitoring

The Amazon Data Provider for SAP exposes Amazon-specific metrics through an XML page at http://localhost:8888/vhostmd of the given system.

This section explains which metrics get exposed to the SAP system and how you can access them for SAP system monitoring.

Checking Metrics with the SAP Operating System Collector (SAPOSCOL)

The information provided by the Amazon Data Provider for SAP is read by the SAP Operating System Collector (SAPOSCOL). You can use the interactive mode of SAPOSCOL to verify that the two tools are working together correctly. The following example shows a lookup under Windows. A lookup under Linux is very similar.

 Open a Windows command shell and direct the shell to the directory C:\Program Files \SAP\hostctrl\exe. Start saposcol.exe with the -d option.

```
PS C:\Users\Administrator> cd 'C:\Program Files'
PS C:\Program Files> cd .\SAP\hostctrl\exe
PS C:\Program Files> cd .\SAP\hostctrl\exe
PS C:\Program Files> SAP\hostctrl\exe> .\saposcol.exe -d
PROGRAM Files SAP\hostctrl\exe> .\saposcol.exe -d
PROGRAM FILES SAP\hostctrl\exe> .\saposcol.exe -d
PROGRAM FILES SAP\hostcrl\execute SAP\hostcrl\execute
```

Figure 41: Starting SAPOSCOL

2. SAPOSCOL is now in interactive mode. Type dump ccm and press **Enter** to list all values gathered. SAPOSCOL will display a lengthy list of metrics, as shown here.

Figure 42: Metrics from SAPOSCOL

Verification of monitoring 121

The following two metrics indicate that SAPOSCOL is collaborating successfully with the Amazon Data Provider for SAP:

- Enhanced Monitoring Access TRUE
- Enhanced Monitoring Details ACTIVE

The Amazon-specific metrics start with the following strings:

- Virtualization_Configuration
- CPU_Virtualization_Virtual_System
- Memory_Virtualization_Virtual_System
- System_Info_Virtualization_System

Figure 43: Amazon-specific metrics

Checking Metrics with the SAP CCMS Transactions

SAPOSCOL hands the Amazon-enhanced statistics with other operating system-specific metrics to the SAP system. You can also check the Amazon-enhanced statistics in the SAP CCMS. You can enter the transaction *st06* (or */nst06*) in the upper-left transaction field of the SAP GUI for quick access to this data.



Note

You will need the appropriate authorizations to look up this information.

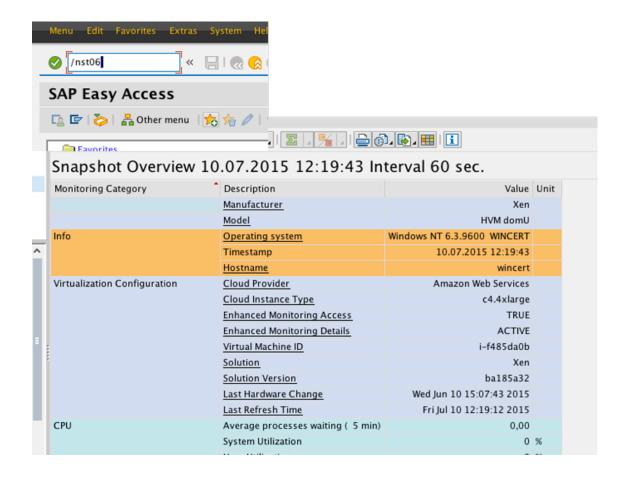


Figure 44: Statistics in the SAP CCMS (standard view)

On this screen, you can verify core Amazon information such as:

- Cloud provider
- Instance type
- Status of enhanced monitoring access (must be TRUE)
- Status of enhanced monitoring details (must be ACTIVE)
- Virtual machine identifier

▲ Important

The enhanced Amazon metrics aren't shown in standard view.

To view enhanced Amazon statistics, choose the **Standard View** button in the upper-left corner. It changes to **Expert View** and displays the enhanced Amazon statistics. The list that appears is comprehensive. It shows the processor details.

CPU Virtualization Host	Overprovisioning	no	
	Processor	Intel(R) Xeon(R) @ 2900MHz	
	Number of Cores per Physical CPU	8	Core
	Number of Threads per Core	2	Thds
	Current Processor Frequency	2.900	MHz
	Maximum Processor Frequency	2.900	MHz
CPU Virtualization Virtual System	Physical Reference Compute Unit (CU)	Intel(R) Xeon(R) @ 2900MHz	
	CPU Physical Equivalent	thread @ 1CUs	
	Guaranteed Capacity	16,00	CPUs
	Capacity Consumed	0,00	CPUs
	Additional Capacity Available	16,00	CPUs
	Available Capacity	16,00	CPUs
	Available Capacity Consumed	0,0	%
	Capacity Maximum	16,00	CPUs

Figure 45: Enhanced Amazon statistics (expert view)

It also shows details about the memory subsystem (main memory and disks) and network interfaces.

Memory Virtualization Host	Overprovisioning	no	
	Network Read Throughput	5	kB/s
	Network Write Throughput	4	kB/s
	Network TCP Retransmission Rate	0	/s
	lan2\Network Device Id	eni-8235b5d8	
	lan2\Minimum Network Bandwith	10.000	Mb
	lan2\Maximum Network Bandwith	10.000	Mb
	disk0\Volume Id	vol-f4b78f0c	
	disk0\Refresh Interval	300	s
	disk0\Volume Utilization	0,0	%
	disk0\Guaranteed Disk IOPS	3.000	
	disk0\Volume Queue Length	0	
	disk0\Volume Read Response Time	0	msec
	disk0\Volume Write Response Time	0	msec
	disk0\Volume Read Throughput	1	kB/s
	disk0\Volume Write Throughput	15	kB/s
	disk0\Volume Read Ops	0	/s
	disk0\Volume Write Ops	2	/s
	disk1\Volume Id	vol-89676771	
	disk1\Refresh Interval	300	s
	disk1\Volume Utilization	0,0	%
	disk1\Guaranteed Disk IOPS	300	
	disk1\Volume Queue Length	0	
	disk1\Volume Read Response Time	0	msec
	disk1\Volume Write Response Time	0	msec
	disk1\Volume Read Throughput	0	kB/s
	disk1\Volume Write Throughput	0	kB/s
	disk1\Volume Read Ops	0	/s
	disk1\Volume Write Ops	0	/s

Figure 46: Memory and networking statistics (expert view)



The screen illustrations in Figures 37–39 were taken from SAP NetWeaver 7.4 SP08. This version shows the enhanced Amazon statistics in the **Memory Virtualization** section. This problem has been fixed by SAP in later versions of NetWeaver.

Example of captured metrics

This following show example metrics. Your system metrics may slightly differ.

```
<metrics>
<metric context="host" category="config" type="long" unit="posixtime">
<name>Time Stamp</name>
```

```
<value>1584376572</value>
</metric>
<metric context="host" category="config" type="int64" unit="sec">
<name>Refresh Interval
<value>60</value>
</metric>
<metric context="vm" category="config" type="string" unit="none">
<name>Data Provider Version
<value>3.0.139</value>
</metric>
<metric context="host" category="config" type="string" unit="none">
<name>Cloud Provider</name>
<value>Amazon Web Services</value>
</metric>
<metric context="vm" category="config" type="string" unit="none">
<name>Instance Type</name>
<value>m5.large</value>
</metric>
<metric context="host" category="config" type="string" unit="none">
<name>Virtualization Solution</name>
<value>KVM</value>
</metric>
<metric context="host" category="config" type="string" unit="none">
<name>Virtualization Solution Version
<value>ba185a32</value>
</metric>
<metric context="host" category="config" type="long" unit="none">
<name>CloudWatch Calls</name>
<value>12</value>
</metric>
<metric context="host" category="config" type="long" unit="none">
<name>EC2 Calls</name>
<value>4</value>
</metric>
<metric context="vm" category="config" type="string" unit="none">
<name>CPU Over-Provisioning</name>
<value>no</value>
</metric>
<metric context="vm" category="config" type="string" unit="none">
<name>Memory Over-Provisioning</name>
<value>no</value>
</metric>
<metric context="vm" category="config" type="string" unit="none">
<name>Virtualization Type</name>
```

```
<value>default-hvm</value>
</metric>
<metric context="vm" category="config" type="string" unit="none">
<name>Virtual Machine ID
<value>i-############</value>
</metric>
<metric context="vm" category="config" type="long" unit="posixtime">
<name>Last Hardware Change</name>
<value>1572284007</value>
</metric>
<metric context="host" category="cpu" type="string" unit="none">
<name>Processor Type</name>
<value>Intel(R) Xeon(R) @ 2500MHz</value>
</metric>
<metric context="host" category="cpu" type="int64" unit="none">
<name>Number of Cores per CPU</name>
<value>1</value>
</metric>
<metric context="host" category="cpu" type="int64" unit="none">
<name>Number of Threads per Core</name>
<value>2</value>
</metric>
<metric context="host" category="cpu" type="int64" unit="MHz">
<name>Max HW Frequency</name>
<value>2500</value>
</metric>
<metric context="host" category="cpu" type="int64" unit="MHz">
<name>Current HW Frequency</name>
<value>2500</value>
</metric>
<metric context="vm" category="cpu" type="string" unit="none">
<name>Reference Compute Unit (CU)</name>
<value>Intel(R) Xeon(R) @ 2500MHz</value>
</metric>
<metric context="vm" category="cpu" type="string" unit="none">
<name>vCPU Mapping</name>
<value>thread</value>
</metric>
<metric context="vm" category="cpu" type="long" unit="cu">
<name>Phys. Processing Power per vCPU</name>
<value>1</value>
</metric>
<metric context="vm" category="cpu" type="int64" unit="cu">
<name>Guaranteed VM Processing Power</name>
```

```
<value>2</value>
</metric>
<metric context="vm" category="cpu" type="int64" unit="cu">
<name>Current VM Processing Power</name>
<value>2</value>
</metric>
<metric context="vm" category="cpu" type="int64" unit="cu">
<name>Max. VM Processing Power</name>
<value>2</value>
</metric>
<metric context="vm" category="cpu" type="double" unit="percent">
<name>VM Processing Power Consumption
<value>3.00</value>
</metric>
<metric context="vm" category="memory" type="long" unit="MB">
<name>Guaranteed Memory assigned</name>
<value>8274</value>
</metric>
<metric context="vm" category="memory" type="long" unit="MB">
<name>Current Memory assigned
<value>8274</value>
</metric>
<metric context="vm" category="memory" type="long" unit="MB">
<name>Max Memory assigned</name>
<value>8274</value>
</metric>
<metric context="vm" category="memory" type="double" unit="percent">
<name>VM Memory Consumption</name>
<value>29.00</value>
</metric>
<metric context="vm" category="memory" type="int64" unit="KB/sec">
<name>Memory SwapIn Rate
<value>0</value>
</metric>
<metric context="vm" category="memory" type="int64" unit="MB">
<name>Memory Swapped Out</name>
<value>0</value>
</metric>
<metric context="vm" category="memory" type="int64" unit="MB">
<name>Memory Lent</name>
<value>0</value>
</metric>
<metric context="vm" category="memory" type="int64" unit="MB">
<name>Total Visible Memory</name>
```

```
<value>-1</value>
</metric>
<metric context="vm" category="memory" type="int64" unit="percent">
<name>Visible Memory Consumed</name>
<value>-1</value>
</metric>
<metric context="vm" category="memory" type="int64" unit="KB/sec">
<name>Visible Memory SwapIn Rate
<value>0</value>
</metric>
<metric context="vm" category="memory" type="int64" unit="MB">
<name>Visible Memory Swapped Out</name>
<value>0</value>
</metric>
<metric context="vm" category="network" type="int64" unit="bytes">
<name>Network Read Bytes
<value>54110386
</metric>
<metric context="vm" category="network" type="int64" unit="bytes">
<name>Network Write Bytes</name>
<value>1330726</value>
</metric>
<metric context="vm" category="network" type="int64" unit="none">
<name>TCP Packets Retransmitted</name>
<value>396480</value>
</metric>
<metric context="vm" category="network" type="int64" unit="Mbps" device-id="eni--</pre>
############
<name>Minimum Network Bandwidth</name>
<value>10000</value>
</metric>
<metric context="vm" category="network" type="int64" unit="Mbps" device-id="eni--</pre>
############*
<name>Maximum Network Bandwidth</name>
<value>10000</value>
</metric>
<metric context="vm" category="network" type="string" unit="none" device-id="eni--</pre>
############
<name>Mapping</name>
<value>lan2</value>
</metric>
<metric context="vm" category="disk" type="int64" unit="msec" device-id="vol--</pre>
############*
<name>Volume Idle Time</name>
```

```
<value>58489</value>
</metric>
<metric context="vm" category="disk" type="int64" unit="none" device-id="vol--</pre>
############*
<name>Volume Queue Length</name>
<value>0</value>
</metric>
<metric context="vm" category="disk" type="int64" unit="bytes" device-id="vol--</pre>
############
<name>Volume Read Bytes</name>
<value>9878528</value>
</metric>
<metric context="vm" category="disk" type="int64" unit="none" device-id="vol--</pre>
############
<name>Volume Read Ops</name>
<value>144</value>
</metric>
<metric context="vm" category="disk" type="int64" unit="msec" device-id="vol--</pre>
############
<name>Volume Read Time</name>
<value>246</value>
</metric>
<metric context="vm" category="disk" type="int64" unit="msec" device-id="vol--</pre>
############
<name>Volume Write Time</name>
<value>8266</value>
</metric>
<metric context="vm" category="disk" type="int64" unit="bytes" device-id="vol--</pre>
############
<name>Volume Write Bytes</name>
<value>282332160</value>
</metric>
<metric context="vm" category="disk" type="int64" unit="none" device-id="vol--</pre>
############
<name>Volume Write Ops</name>
<value>3090</value>
</metric>
<metric context="vm" category="disk" type="string" unit="none" device-id="vol--</pre>
############
<name>Volume Type</name>
<value>gp2</value>
</metric>
<metric context="vm" category="disk" type="int64" unit="none" device-id="vol--</pre>
############*
```

```
<name>Guaranteed IOps</name>
<value>180</value>
</metric>
<metric context="vm" category="disk" type="int64" unit="sec" device-id="vol--
###############/">
<name>Interval</name>
<value>300</value>
</metric>
<metric context="vm" category="disk" type="string" unit="none" device-id="vol--
################/">
<name>Mapping</name>
<value>disk0</value>
</metric>
</metric>
</metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric></metric
```

Version history

Version 4.3.2 (August, 2023)

Bug fix: Security updates to address CVE-2022-45688.

Version 4.3.1 (June, 2023)

• Bug fix: Data Provider is now setup for successful installation of SAP JVM.

Version 4.3 (January, 2023)

- Added support for JDK 17
- Added reading configuration information from remote Amazon S3 bucket

Version 4.2 (November, 2022)

- Added support for Oracle and Linux
- Added integration with Linux logrotate feature
- Updates to the RPM package build.

Version 4.1.1 (September, 2022)

Added support for new Amazon EC2 instance types.

Version 4.1.0 (January, 2022)

- Added support for JDK 11.
- Added support for new Amazon EC2 instance types.

Version 4.0.3 (December, 2021)

• Bug fixes: Removed Log4j dependency.

Version 4.0.2 (December, 2021)

Bug fixes: Security updates for Log4j2 issue (CVE-2021-44228).

Version 4.0 (April, 2021)

- Initial release of the 4.0 version.
- Support for SSM package installation.
- Support for IMDSv2.

Version 3.0 (April, 2020)

- Initial release of the 3.0 version.
- Switched the Java Runtime from Oracle to Amazon Corretto.

Version 2.9 (August 30, 2017)

- Added support for China Regions.
- · Added Linux uninstaller.
- Linux installer can be customized to install from a custom S3 bucket.
- Silent installer for Windows (does not require any input).
- Improvements in determination of access points.
- Support for X1E instance family.

Version 2.8 (March 1, 2017)

• SLES 12, Red Hat 7, and Oracle Linux 7 will now use SYSTEMD to manage the daemon.

- Support for SLES and SLES for SAP 12 SP2.
- SLES 12 SP1 systems will get migrated from Linux services to SYSTEMD when trying to install the Amazon Data Provider without having it de- installed first.
- Minor changes in logging texts.
- Support for R4 and M4 instance types.
- Updated Windows installation verification.

Version 2.7 (December 21, 2016)

- Support for Canada (Central), US East (Ohio), and EU (London) Regions.
- Default access point resolution for common Amazon Regions is added.

Version 2.6 (September 1, 2016)

- Bug fixes: Installation script checks for existence of wget
- Support for Oracle Linux.

Version 2.5 (May 2, 2016)

- Bug fixes: Security and stability fixes in versions 2.2-2.4.
- New: Support for new Amazon EBS volume types:
 - Throughput Optimized HDD (st1)
 - Cold HDD (sc1)
- New: Support for the Amazon EC2 X1 instance family.

Version 2.1 (January 20, 2016)

- Support for Asia Pacific (Seoul) Region.
- Bug fix: Version 2.0 pulled files from an incorrect S3 bucket for installation. Version 2.0 needs to be uninstalled before version 2.1 is installed.

Version 2.0 (December 22, 2015)

• New: Windows devices in the range sdb to sdzz get correct SCSI device IDs assigned.

• New: Java VM consumption is now limited to 64 MB maximum heap size.

Version 1.3.1 (July 14, 2015)

- Bug fixes: Security fixes.
- New: Support for C4, D2, and M4 instance types. Users who migrate instances with installed
 1.3 agents will automatically receive support for the new instance types through an updated configuration database on the web.

Version 1.3 (February 17, 2015)

- New: Support for new Amazon EC2 C4 instance family.
- Security fix: Upgraded Linux and Windows versions to JRE 8u31.
- Bug fix: Relative performance of c3.8xlarge instances is now reported correctly.
- New: CloudWatch and Amazon EC2 metrics access points:
 - Support for the EU (Frankfurt) Region was added.
 - Access points are user configurable. You can add information about new Amazon Regions without having to install a new product version.
 - Access points are now updated from an internet-based database file. You can add new Amazon Regions by updating a web-based configuration file and then restarting the daemon/service.
- New: Message log files with fixed disk space consumption are provided on Linux.
- New: User-configurable EC2 instance types are available.
- New: Web update support was added for future EC2 instance types without product updates.
- Bug fix: GP2 volumes now report the correct sample interval time.
- New: User-configurable sample times for new EBS volume types are now available.
- New: The Amazon Data Provider for SAP now reports the virtualization type of the EC2 instance.

Version 1.2.2 (October 1, 2014)

Windows bug fix: Installer executable pulls installation from correct Amazon S3 bucket.

 Windows bug fix: Amazon Data Provider for SAP now reports the correct disk mapping for Windows EBS volumes with the following names: xvd[a-z][a-z].

Version 1.2.1 (September 29, 2014)

• Bug fix: EBS volumes now report correct attribute type ("string") for volume type.

Version 1.2 (September 16, 2014)

- New: Support for the T2, R3, and C3 instance families.
- New: Support for post-ECU (EC2 Compute Unit) instance types:
 - New instance types no longer have ECU values.
 - The reference compute power for these instance types is a hardware thread of the given processor. The total CPU power is equal to the number of the vCPUs of a given instance type.
- New: Support for the new EBS GP2 volume type.
 - Every volume is now tagged with the EBS volume type.
- New: Report of EBS one-minute volume statistics.
 - EBS volumes now report their individual sample interval in a separate attribute.
- Bug fix: EBS volume mapping for Windows devices now reports the correct name.
- Bug fix: Installation, update, and operation through HTTP/HTTPS proxies has been fixed.
- New: JRE 8 support has been added on Linux.

SAP on Amazon cost estimation

Last updated: May 2023

Amazon offers pay-as-you-go pricing. You only pay for the services you use, for as long as you use them. There are no long-term contracts or complex licensing requirements. For more information, see Amazon Pricing and https://calculator.aws/#/.

The following is an overview of the pricing characteristics of Amazon services that are frequently used for the deployment and operation of SAP systems on Amazon.

Topics

- Amazon Region
- Compute
- Storage
- Network
- Automation
- · Backup, restore, and recovery
- Migration
- Monitoring
- Operating System licenses
- Amazon Marketplace
- Amazon Support

Amazon Region

Amazon service pricing varies between different Amazon Regions. You must select the Region in which you want to deploy your SAP system to begin creating an estimate. For more information, see Regions and Availability Zones.

Compute

<u>Amazon Elastic Compute Cloud (Amazon EC2)</u> provides a wide selection of instance types that provide varying combinations of CPU, memory, storage, I/O, and networking capabilities. Each running instance is charged by the hour. For more information, see <u>Amazon EC2 pricing</u>.

Amazon Region 136

Amazon EC2 offers multiple purchasing options that give you flexibility to optimize your costs. For more information, see Instance purchasing options.

Storage

The following Amazon services are flexible, cost-effective, and easy-to-use data storage options available for your SAP systems. Each option has a unique combination of performance and durability. For more information, see Cloud storage on Amazon.

Services

- Amazon EBS
- Amazon EFS
- Amazon FSx for Windows File Server
- Amazon FSx for NetApp ONTAP
- Amazon S3

Amazon EBS

Amazon Elastic Block Store (Amazon EBS) provides persistent, block-level storage volumes for Amazon EC2 instances. Each Amazon EC2 instance that runs an SAP environment requires one or more Amazon EBS volumes to store system components, such as operating system, SAP software, SAP database data and log files, and local backup storage.

With Amazon Elastic Block Store, you only pay for what you provision. For more information, see Amazon EBS pricing.

Amazon EBS snapshots

Amazon EBS snapshots are point-in-time copies of your block data stored in Amazon EBS volumes. Amazon EBS snapshots in the Standard tier are stored incrementally, which means you are billed only for the changed blocks stored. Amazon EBS snapshots in the Archive tier are full copies of your block data, which means you are billed for all the blocks stored and not just the changed blocks.

You can also enable a Recycle Bin feature to protect against accidental deletion. Amazon EBS snapshots in the recycle bin are billed at the same rate.

Storage 137

General SAP Guides SAP Guides

Another feature of Amazon EBS snapshots which is applicable to SAP workloads is the Fast Snapshot Restore (FSR). This feature enables you to promptly restore fully provisioned Amazon EBS volumes from snapshots, regardless of the size of the volume or snapshot. FSR is charged in Data Services Unit-Hours (DSU-Hours) for each snapshot and each Availability Zone in which it is enabled. DSUs mean that you are billed per minute with a one-hour minimum.

Amazon EBS snapshots can be used for both root and binary volumes of your SAP system as well as database volumes.

Amazon EFS

<u>Amazon Elastic File System (Amazon EFS)</u> provides serverless file storage. You can share file data without provisioning or managing storage capacity and performance. Amazon EFS is built to scale on demand to petabytes without disrupting applications, growing and shrinking automatically as you add and remove files. You can create and configure file systems quickly and easily.

Amazon EFS stores every object across multiple Availability Zones, making it highly available. It supports the Network File System version 4 (NFSv4.1 and NFSv4.0) protocol. It is certified for SAP file shares, and it can also be used for storing data files in your SAP landscape.

With Amazon EFS, you pay only for the storage used by your file system, and there is no minimum fee or setup cost. For more information, see Amazon EFS pricing.

Amazon FSx for Windows File Server

<u>Amazon FSx for Windows File Server</u> provides fully managed Microsoft Windows file servers, backed by a fully native Windows file system. It has support for Windows file system features and for the industry-standard Server Message Block (SMB) protocol to access file storage over a network.

FSx for Windows File Server is certified for SAP workloads on Amazon, and can also be used for Windows based data file sharing in your SAP landscape.

With FSx for Windows File Server, you only pay for the resources you use, and there is no minimum fee or setup cost. For more information, see Amazon FSx for Windows File Server pricing.

Amazon FSx for NetApp ONTAP

<u>Amazon FSx for NetApp ONTAP</u> is a fully managed service that provides highly reliable, scalable, high-performing, and feature-rich file storage built on NetApp's popular ONTAP file system.

Amazon EFS 138

FSx for ONTAP is certified for SAP workloads on Amazon.

You are billed for the file systems you use, based on the following categories.

- SSD storage capacity (per gigabtye-month, or GB-month)
- SSD IOPS that you provision above three IOPS/GB (per IOPS-month)
- Throughput capacity (per megabytes per second [MBps]-month)
- Capacity pool storage consumption (per GB-month)
- Capacity pool requests (per read and write)
- Backup storage consumption (per GB-month)

For more information, see Amazon FSx for NetApp ONTAP pricing.

Amazon S3

Amazon Simple Storage Service (Amazon S3) is an object storage service that offers industry-leading scalability, data availability, security, and performance. You can use Amazon S3 to stage media, store backups, and archive data. Amazon S3 offers a range of storage classes designed for different use cases.

Amazon S3 charges you only for what you actually use, with no hidden fees and no overage charges. This model gives you a variable-cost service that can grow with your business while giving you the cost advantages of Amazon infrastructure. For more information, see Amazon S3 pricing.

Network

Amazon offers multiple strong and secure networking services.

Services

- Amazon VPC
- Amazon Site-to-Site VPN
- Amazon Direct Connect
- Elastic Load Balancing
- Data transfer pricing

Amazon S3 139

Amazon VPC

With <u>Amazon Virtual Private Cloud (Amazon VPC)</u>, you can launch Amazon resources in a logically isolated virtual network that you've defined. This virtual network closely resembles a traditional network that you'd operate in your own data center, with the benefits of using the scalable infrastructure of Amazon.

There's no additional charge for using Amazon VPC. There are charges for some components, such as NAT gateways, IP Address Manager, traffic mirroring, Reachability Analyzer, and Network Access Analyzer. For more information, see Amazon VPC pricing.

Use the following options for a secure connection between your on-premises network and Amazon VPC.

Amazon Transit Gateway is a network transit hub that you can use to interconnect your virtual
private clouds (VPCs) and on-premises networks. As your cloud infrastructure expands globally,
inter-Region peering connects transit gateways together using the Amazon Global Infrastructure.
Your data is automatically encrypted and never travels over the public internet.

You are charged hourly for each attachment on a transit gateway, and you are charged for the amount of traffic processed on the transit gateway. For more information, see Amazon Transit Gateway pricing.

• <u>NAT gateway</u> is a Network Address Translation (NAT) service. You can use a NAT gateway so that instances in a private subnet can connect to services outside your VPC but external services cannot initiate a connection with those instances.

When you provision a NAT gateway, you are charged for each hour that your NAT gateway is available and each Gigabyte of data that it processes. For more information, see Amazon VPC pricing.

Amazon PrivateLink is a highly available, scalable technology that you can use to privately connect your VPC to services as if they were in your VPC. You do not need to use an internet gateway, NAT device, public IP address, Amazon Direct Connect connection, or Amazon Siteto-Site VPN connection to allow communication with the service from your private subnets. Therefore, you control the specific API endpoints, sites, and services that are reachable from your VPC.

For information about the pricing for VPC endpoints, see Amazon PrivateLink pricing.

Amazon VPC 140

Amazon Site-to-Site VPN

By default, instances that you launch into an Amazon VPC can't communicate with your own (remote) network. You can enable access to your remote network from your VPC by creating an <u>Amazon Site-to-Site VPN</u> (Site-to-Site VPN) connection, and configuring routing to pass traffic through the connection.

You are charged for each VPN connection hour that your VPN connection is provisioned and available. For more information, see <u>Amazon Site-to-Site VPN and Accelerated Site-to-Site VPN Connection pricing</u>.

You are charged for data transfer out from Amazon EC2 to the internet. For more information, see Data Transfer.

When you create an accelerated VPN connection, we create and manage two accelerators on your behalf. You are charged an hourly rate and data transfer costs for each accelerator. For more information, see Amazon Global Accelerator pricing.

Amazon Direct Connect

Amazon Direct Connect links your internal network to an Amazon Direct Connect location over a standard Ethernet fiber-optic cable. One end of the cable is connected to your router, the other to an Amazon Direct Connect router. With this connection, you can create virtual interfaces directly to public Amazon services (for example, to Amazon S3 or Amazon VPC), bypassing internet service providers in your network path. An Amazon Direct Connect location provides access to Amazon in the Region with which it is associated. You can use a single connection in a public Region or Amazon GovCloud (US) to access public Amazon services in all other public Regions.

Amazon Direct Connect has two billing elements: port hours and outbound data transfer. For more information, see Amazon Direct Connect pricing.

Elastic Load Balancing

<u>Elastic Load Balancing</u> automatically distributes your incoming traffic across multiple targets, such as Amazon EC2 instances, containers, and IP addresses, in one or more Availability Zones. It monitors the health of its registered targets, and routes traffic only to the healthy targets. Elastic Load Balancing scales your load balancer capacity automatically in response to changes in incoming traffic.

Use Elastic Load Balancing for setting up highly available SAP environments on Amazon.

Amazon Site-to-Site VPN 141

General SAP Guides SAP Guides

With your load balancer, you pay only for what you use. For more information, see <u>Elastic Load</u> Balancing pricing.

Data transfer pricing

Data transfer pricing varies based on the architecture pattern and use case. It only amounts to a small percentage of the overall cost of SAP workloads on Amazon, 1-5%.

The following table summarizes the common data transfer scenarios that can apply to your SAP landscape.

Scenario	Architecture	Data transfer cost	Pricing guidance	Other costs
Inbound to Amazon	All	No		
Outbound from Amazon to Internet	All	Yes	Based on the services used	
Services in the same Amazon Region	Internet gateway	No		
Services in the same Amazon Region	NAT gateway	No		Per GB processin g charge, see Amazon VPC pricing
Services in the same Amazon Region	Amazon PrivateLink/ VPC endpoint	No		See <u>Amazon</u> <u>PrivateLink</u> <u>pricing</u>
Services in different Amazon Regions	All	Yes	Per GB charge for inter-Region data transfer, see <u>Amazon EC2</u> <u>pricing</u>	

Data transfer pricing 142

General SAP Guides SAP Guides

Scenario	Architecture	Data transfer cost	Pricing guidance	Other costs
Components in same Availability Zone	All	No		
Component s in different Availability Zones	Amazon Transit Gateway	No		Amazon Transit Gateway processing charges, see Amazon Transit Gateway pricing
Component s in different Availability Zones	Amazon VPC peering	Yes	Per GB charge for inter-Region data transfer, see <u>Amazon EC2</u> <u>pricing</u>	
Component s in different Amazon Regions	Amazon Transit Gateway	Yes	Per GB charge for inter-Region data transfer, see <u>Amazon EC2</u> <u>pricing</u>	Amazon Transit Gateway processing charges
Component s in different Amazon Regions	Amazon VPC peering	Yes	Per GB charge for inter-Region data transfer, see <u>Amazon EC2</u> <u>pricing</u>	
Transfer to on-premises or corporate network	Amazon VPN	Yes	Per GB charge for data transfer out, see <u>Amazon</u> <u>EC2 pricing</u>	Amazon VPN and Amazon Transit Gateway charges

Data transfer pricing 143

Scenario	Architecture	Data transfer cost	Pricing guidance	Other costs
Transfer to on-premises or corporate network	Amazon Direct Connect	Yes	Per GB charge for data transfer out based on location and provider, see Amazon EC2 pricing	Amazon Direct Connect and Amazon Transit Gateway charges

Automation

With Amazon Systems Manager for SAP, you can backup and restore SAP HANA databases on Amazon EC2 with Amazon Backup.

Amazon Systems Manager for SAP is available to you at no additional cost. You only pay for the Amazon resources that you provision to manage and operate your SAP environments.

Backup, restore, and recovery

With these services, you can quickly and effectively backup, restore, and recovery your SAP workloads.

Services

- Amazon Backint Agent for SAP HANA
- Amazon Elastic Disaster Recovery
- Amazon EBS snapshots

Amazon Backint Agent for SAP HANA

Amazon Backint Agent for SAP HANA (Amazon Backint agent) is an SAP-certified backup and restore application for SAP HANA workloads running on Amazon EC2 instances in the cloud. Amazon Backint agent runs as a standalone application that integrates with your existing workflows to back up your SAP HANA database to Amazon S3 and Amazon Backup.

Automation 144

Amazon Backint agent is a free service. You pay for only the underlying Amazon services that you use, for example Amazon S3 or Amazon Backup. See the following references for more information

- Amazon S3 pricing
- Amazon Backup pricing

Amazon Elastic Disaster Recovery

<u>Amazon Elastic Disaster Recovery (Elastic Disaster Recovery)</u> minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications using affordable storage, minimal compute, and point-in-time recovery.

You only pay for the servers you are actively replicating to Amazon. For more information, see Amazon Elastic Disaster Recovery pricing.

Amazon EBS snapshots

See the the section called "Amazon EBS" section.

Migration

The following services enable you to quickly move application and files.

Services

- Migration Hub Orchestrator
- Amazon DataSync

Migration Hub Orchestrator

<u>Amazon Migration Hub Orchestrator</u> simplifies and automates the migration of servers and enterprise applications to Amazon. It provides a single location to run and track your migrations.

Amazon Migration Hub Orchestrator is available to you at no additional cost. You only pay for the Amazon resources that you provision for migrations.

Amazon DataSync

<u>Amazon DataSync</u> is an online data movement and discovery service that simplifies data migration and helps you quickly, easily, and securely move your file or object data to, from, and between Amazon storage services.

You pay only for the amount of data that you migrate based on a flat, per-gigabyte fee according to your Amazon Region. For more information, see Amazon DataSync pricing.

Monitoring

With the use of following services, you can monitor your SAP workloads on Amazon.

Services

- Amazon Data Provider
- Amazon CloudWatch Application Insights for SAP HANA
- Amazon CloudWatch
- Amazon CloudTrail
- VPC Flow Logs

Amazon Data Provider

Amazon Data Provider for SAP is a tool that collects performance-related data from Amazon services. It makes this data available to SAP applications to help monitor and improve the performance of business transactions.

For information about costs, see Amazon Data Provider for SAP pricing.

Amazon CloudWatch Application Insights for SAP HANA

<u>Amazon CloudWatch Application Insights</u> helps you monitor your applications that use Amazon EC2 instances along with other application resources.

CloudWatch Application Insights sets up recommended metrics and logs for selected application resources using CloudWatch metrics, Logs, and Events for notifications on detected problems. These features are charged to your Amazon account according to <u>Amazon CloudWatch pricing</u>. For more information, see CloudWatch Application Insights pricing.

Amazon DataSync 146

Amazon CloudWatch

<u>Amazon CloudWatch</u> monitors your resources and applications running on Amazon in real time. You can collect and track metrics for your resources and applications.

For information about CloudWatch pricing, refer the following resources.

- CloudWatch billing and cost
- · Amazon CloudWatch pricing

Amazon CloudTrail

<u>Amazon CloudTrail</u> is an Amazon service that helps you enable operational and risk auditing, governance, and compliance of your Amazon account. Actions taken by a user, role, or an Amazon service are recorded as events in CloudTrail.

CloudTrail is charged pay per usage with no minimum fee. For more information, see <u>Amazon</u> CloudTrail pricing.

VPC Flow Logs

<u>VPC Flow Logs</u> is a feature that enables you to capture information about the IP traffic going to and from network interfaces in your VPC.

Data ingestion and archival charges for vended logs apply when you publish flow logs. For more information about pricing when publishing vended logs, open <u>Amazon CloudWatch pricing</u>, select **Logs** > **Vended Logs**.

Operating System licenses

You can bring your own licenses for the operating system of your choice or purchase from <u>Amazon</u> <u>Web Services Marketplace</u>.

Operating Systems

- Red Hat
- SUSE
- Windows
- Oracle Enterprise Linux

Amazon CloudWatch 147

General SAP Guides **SAP Guides**

Red Hat

Red Hat offers two Linux distributions to run SAP workloads. For more details, you can check Red Hat documentation – Overview of the Red Hat Enterprise Linux (RHEL) for SAP Solutions subscription.

You can avail these options from Amazon Web Services Marketplace or Red Hat Cloud Access. When you purchase Red Hat operating systems from Amazon, your Amazon Web Services Support plan includes operating system support.

If you want to launch an Amazon EC2 instance with RHEL for SAP Applications, then you must have a subscription for the Red Hat Cloud Access program. With RHEL 8, RHEL for SAP Solutions or RHEL for SAP Applications is required for running SAP applications in production environments.

RHEL for SAP Solutions on Amazon Web Services Marketplace is available at an hourly rate or as an annual commitment. RHEL for SAP Solutions is designed specifically to run SAP workloads. It includes a longer lifecycle support with Extended Update Support (E4S) that provides support on specific minor releases for four years from general availability. It also provides all the necessary packages for configuring the Pacemaker based cluster, ensuring reliability and availability of critical production services.



Note

Amazon Web Services Marketplace pricing displays the software cost for RHEL. The additional cost of RHEL is included in Amazon EC2 pricing.

SUSE

SUSE offers two Linux distributions to run SAP workloads – SUSE Linux Enterprise Server (SLES) and SUSE Linux Enterprise Server for SAP Applications (SLES for SAP)

You can avail these options from Amazon Web Services Marketplace or from SUSE. When you purchase SUSE operating systems from Amazon, your Amazon Web Services Support plan includes operating system support.

When you bring your own subscription, the support for Amazon EC2 is based on your SUSE purchasing agreement.

Red Hat 148 General SAP Guides **SAP Guides**

SLES for Amazon EC2 is available at an hourly rate or as an annual commitment. RHEL for SAP Solutions is designed specifically to run SAP workloads. It includes a longer lifecycle support with Extended Update Support (E4S) that provides support on specific minor releases for four years from general availability. It also provides all the necessary packages for configuring the Pacemaker based cluster, ensuring reliability and availability of critical production services.

SLES for SAP on Amazon Web Services Marketplace is available at an hourly rate or as an annual commitment. The price of the SLES subscription is included with your Amazon EC2 instance cost, and is based on the vCPUs of the Amazon EC2 instance. SLES for SAP is designed specifically to run SAP workloads. It includes a longer lifecycle support with Extended Service Pack Overlap Support that provides 4.5 years of total support. SLES for SAP also offers software components and service offerings like SAP HANA high availability resource agents and cluster connector.



Note

SLES for SAP pricing is the cost of the software and there is not an additional cost for SLES.

Windows

Windows server on Amazon EC2 can be availed at a flat, hourly rate with no commitment (On-Demand) or through a one-time payment (Savings Plan or Reserved Instances). There is no difference in cost in terms of a Windows operating system with either of these options. You can also bring your own licence. For more information, see Microsoft Licensing on Amazon.

Oracle Enterprise Linux

SAP requires you to have Oracle Linux Premier Support subscription to use Oracle Enterprise Linux operating system. For additional information, review the following resources from Oracle and SAP.

- Oracle Store
- SAP Note 2069760 Oracle Linux 7.x SAP Installation and Upgrade (requires SAP portal access)

Amazon Marketplace

Amazon Marketplace is a curated digital catalog that customers can use to find, buy, deploy, and manage third-party software, data, and services to build solutions and run their businesses.

Windows 149 General SAP Guides SAP Guides

In Amazon Marketplace, products can be free to use or can have associated charges. For more information, see Product pricing.

Amazon Support

<u>Amazon Support</u> offers different levels of support. For more information, see <u>Amazon Support</u> Plan Pricing.

SAP requires you to have at least a Business level of support when running SAP workloads on Amazon. To learn more about the SAP prerequisite, see SAP Note 1656250 - SAP on Amazon: Support Prerequisites (requires SAP portal access).

Amazon Support 150

Architecture guidance for availability and reliability of SAP on Amazon

August 2021

This guide is part of a content series that provides detailed information about hosting, configuring, and using SAP technologies in the Amazon Web Services (Amazon) Cloud. For more information, see SAP on Amazon Technical Documentation.

Overview

This guide provides a set of architecture guidelines, strategies, and decisions for deploying SAP NetWeaver-based systems with a highly available and reliable configuration on Amazon.

In this guide we cover:

- Introduction to SAP high availability and reliability
- Architecture guidelines and decision consideration
- Architecture patterns and recommended usage

This guide is intended for users who have previous experience designing high availability and disaster recovery (HADR) architectures for SAP.

This guide does not cover the business requirements determining the need for HADR and/or the implementation details for a specific partner or customer solution.

Prerequisites

Specialized knowledge

Before following the configuration instructions in this guide, we recommend familiarizing yourself with the following Amazon services. (If you are new to Amazon, see <u>Getting Started with Amazon</u>.)

- Amazon EC2
- Amazon EBS
- Amazon VPC

Overview 151

General SAP Guides SAP Guides

- Amazon EFS
- Amazon S3

Recommended reading

Before reading this document, we recommend understanding key concepts and best practices from these guides:

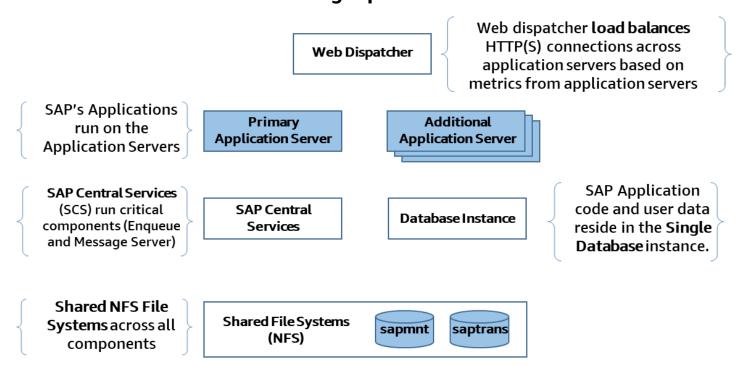
- SAP on Amazon Overview and Planning
- Getting Started with Architecting SAP on the Amazon Cloud

Introduction

For decades, SAP customers protected SAP workloads on premise with two common patterns: high availability and disaster recovery. The advent of cloud computing provided an opportunity to rethink HADR capabilities for SAP, using modern architectures and technologies.

Let's recap the SAP system design and single points of failure that are part of the SAP n-tier architecture.

SAP NetWeaver architecture single points of failure



Recommended reading 152

Figure 1: SAP single points of failure

Figure 1 shows the typical SAP NetWeaver architecture, which has several single points of failure which are listed below:

- SAP Central Services (message server and enqueue processes)
- SAP Application Server
- NFS (shared storage)
- Database
- SAP Web Dispatcher

For the SAP Central Services and Database, protection can be added by deploying additional hosts. For example, an additional host running the SAP replicated enqueue can protect the loss of application level locks (enqueue locks) and an additional host running a secondary database instance can protect against data loss.

However, the inherent design of these single points of failure limits the ability to easily take advantage of cloud native features to provide high availability and reliability.

Amazon Elastic File Service (Amazon EFS) is a highly available and durable managed NFS service that runs actively across multiple physical locations (Amazon Availability Zones). This service can help protect one of the SAP single points of failure.

High availability and disaster recovery

High availability (HA) is the attribute of a system to provide service during defined periods, at acceptable or agreed upon levels and to mask unplanned outages from end users. This is often achieved by using clustered servers. These servers provide automated failure detection, recovery or highly resilient hardware, robust testing, and problem and change management.

Disaster recovery (DR) protects against unplanned major outages, such as site disasters, through reliable and predictable recovery on a different hardware and/or physical location. The loss of data due to corruption or malware is considered a logical disaster event. It is normally resolved in a separate solution, such as recovery from the latest backup or storage snapshot. Logical DR does not necessarily imply a fail over to another facility.

From the perspective of documented and measurable data points, HADR requirements are often defined in terms of the following:

- **Percentage uptime** is the percentage of uptime in a given period (monthly or annual).
- Mean time to recovery (MTTR) is the average time required to recover from failure.
- Return to service (RTS) is the time it takes to bring the system back to service for the users.
- Recovery time objective (RTO) is the maximum length of time that a system or service can be down, how long a solution takes to recover, and the time it takes for a service to be available again.
- **Recovery point objective (RPO)** is how much data a business is willing to lose, expressed in time. It's the maximum time between a failure and the recovery point.

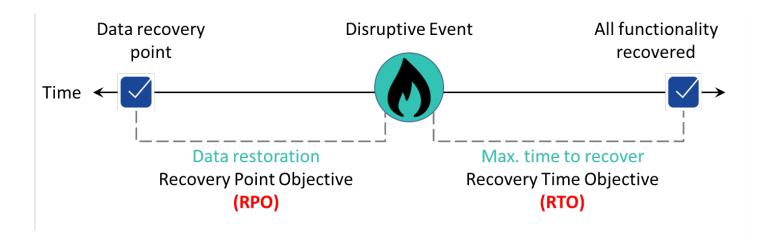


Figure 2: Recovery from a disruptive event

On premises vs. cloud deployment patterns

Traditionally, customers with high availability requirements would deploy their primary compute capabilities in a single data center or hosting facility, often in two separate rooms or data center halls with disparate cooling and power, and high-speed network connectivity. Some customers would run two hosting facilities in close proximity, with a separation of compute capabilities, yet close enough to not be impacted by network latency.

To meet disaster recovery requirements (the preceding scenarios represent an elevated risk to unforeseen location failure), many customers would extend their architecture to include a secondary location where a copy of their data resided, with additional idle compute capacity. The distance between the primary and secondary locations often created the need for asynchronous transfer of data which impacted the recovery point objective. This was the standard and generally accepted architecture pattern for high availability and disaster recovery for many industries and companies running SAP.



Figure 3: On premise disaster recovery

In Figure 3, we give an example of an approach that customers often take on premises. In **Location** 1, the customer has two hosting facilities often separate rooms or halls in the same data center where they deploy a high availability architecture for the SAP single point of failure. **Location 2** is the disaster recovery location in which the SAP systems are recovered, in the event of a significant failure of both hosting facilities in **Location 1**.

Customers migrating their SAP workloads to cloud providers still revert to this architecture and map it to Amazon Regions and Availability Zones (AZs) as depicted in Figure 4. While this architecture can work in your environment, it does not follow the <u>Amazon Well-Architected</u>
<u>Framework</u> which helps cloud architects build secure, high-performing, resilient, and efficient infrastructure for their applications.

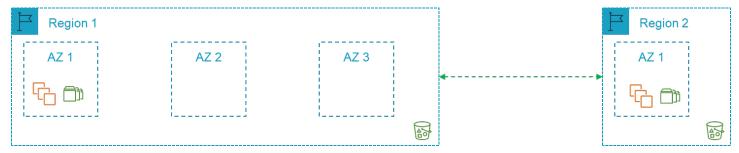


Figure 4: On premises to Amazon region mapping approach

Amazon isolates facilities geographically in Regions and Availability Zones. A Multi-AZ approach provides distance while maintaining performance for the primary compute capacity. This approach (Figure 5) greatly reduces the risk of location failure.

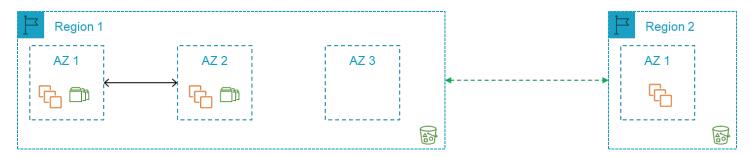


Figure 5: Alternative approach for on premises to Amazon region mapping

With the risk of location failure significantly reduced for the primary compute capacity, the requirements for a second Region can be evaluated based on business requirements. You can rapidly deploy required capacity in the same or different Region with Amazon. Idle hardware is no longer an issue. Data backups can be stored on Amazon Simple Storage Service (Amazon S3) in a single Amazon Region or in multiple Amazon Regions by leveraging cross-Region replication. This architecture can be simplified and be made readily available (Figure 6).

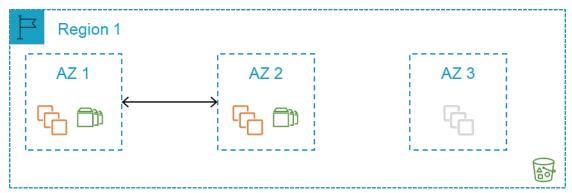


Figure 6: Single Amazon Region approach

In addition to considering the impact of infrastructure or hosting facility failure, another scenario to consider is the loss of business data due to accidental or malicious technical activity.

Loss of business data due to accidental or malicious technical activity is referred to as *logical* disaster recovery. It requires a decision to restore the business data from a good local copy. To enable this, decisions need to be made with regard to the storage location of the data and how it will be used in the event of a *logical disaster recovery*.

Further in this guide, we detail the key architecture guidelines, architecture patterns, and decisions to consider for your availability and reliability requirements.

Architecture guidelines and decisions

This section will provides a brief overview of the Amazon services typically used for SAP workloads and some of the key points to understand when designing your architecture for hosting SAP on Amazon. If you are already familiar with these Amazon services, you can skip this section.

Regions and Availability Zones

The <u>Amazon Global Infrastructure</u> consists of <u>Amazon Regions</u> and <u>Availability Zones</u> (AZs). For more details on the Amazon Global Infrastructure, see Regions and Availability Zones.

Regions

Amazon has a global footprint and ensures that customers are served across the world. Amazon maintains multiple Regions in North America, South American, Europe, Asia Pacific, and the Middle East.

An Amazon Region is a collection of Amazon resources in a geographic area. Each Region is isolated and independent. For a list of Region names and codes, see <u>Regional endpoints</u>.

Regions provide fault tolerance, stability, and resilience. They enable you to create redundant resources that remain available and unaffected in the unlikely event of an outage.

Amazon Regions consist of multiple Availability Zones (AZs), typically 3. An Availability Zone is a fully isolated partition of the Amazon infrastructure. It consists of discrete data centers housed in separate facilities, with redundant power, networking, and connectivity.

You retain complete control and ownership over the Amazon Region in which your data is physically located, making it easy to meet Regional compliance and data residency requirements.

Availability Zones

Availability Zones (AZs) enable customers to operate production applications and databases that are more highly available than would be possible from a single data center. Distributing your applications across multiple zones provides the ability to remain resilient in the face of most failure modes, including natural disasters or system failures.

Each Availability Zone can be multiple data centers. At full scale, it can contain hundreds of thousands of servers. They are fully isolated partitions of the Amazon global infrastructure. With its own powerful infrastructure, an Availability Zone is physically separated from any other zones.

There is a distance of several kilometers, although all are within 100 km (60 miles of each other). This distance provides isolation from the most common disasters that could affect data centers (i.e. floods, fire, severe storms, earthquakes, etc.).

All Availability Zones (AZs) within a Region are interconnected with high-bandwidth and low-latency networking, over fully redundant and dedicated metro fiber. This ensures high-throughput, low-latency networking between zones. The network performance is sufficient to accomplish synchronous replication between zones.

Amazon Availability Zones (AZs) enable our customers to run their applications in a highly-available manner. To be highly available, an application needs to run in more than one location simultaneously with the exact same data, thus allowing for a seamless fail over with minimal downtime, in the event of a disaster.

Services

Our general policy is to deliver Amazon services, features, and instance types to all Amazon Regions within 12 months of general availability, based on customer demand, latency, data sovereignty, and other factors. You can share your interest for local Region delivery, request service roadmap information, or gain insight on service interdependency (under NDA) by contacting your Amazon sales representative.

Due to the nature of the service, some Amazon services are delivered globally rather than Regionally, such as Route 53, Amazon Chime, Amazon WorkDocs, Amazon WorkMail, WorkSpaces, and Amazon WorkLink.

Other services, such as Amazon Elastic Compute Cloud (Amazon EC2) and Amazon Elastic Block Store (Amazon EBS) are zonal services. When you create an Amazon EC2 or Amazon EBS resource for launch, you need to specify the required Availability Zone within a Region.

Selecting the Amazon Regions

When selecting the Amazon Region(s) for your SAP environment deployment, you should consider the following:

- Proximity to on-premises data centers, systems, and end users to minimize network latency.
- Data residency and compliance requirements.
- Availability of the Amazon products and services that you plan to use in the Region. For more details, see Region Table.

• Availability of the Amazon EC2 instance types that you plan to use in the Region. For more details, see Amazon EC2 Instance Types for SAP.

Pricing variation between different Amazon Regions. For more details, see <u>SAP on Amazon</u>
 Pricing and Optimization guide .

Multi-Region considerations

When deploying across multiple Regions, an important consideration is the associated cost and management effort for core services required in each Region such as networking, security, and audit services.

Network latency

If you decide on a multiple Region approach, you should consider the impact of any increase in the network latency to the secondary Region from your on-premises locations.

Cross-Regional data transfer

Amazon provides several methods of data transfer between Regions. These methods are relevant when designing an SAP Architecture for disaster recovery. You should consider any data residency requirements when transferring data to another Amazon Region, the costs associated with the data transfer (cross-Region peering and/or Amazon S3 replication), and storage in the secondary Region.

Tier 0 services

When using an Amazon Region, there are a number of Tier 0 services that you need before deploying an SAP workload. These include DNS, Active Directory, and/or LDAP as well as any Amazon or ISV-provided security and compliance products and services.

Amazon accounts

While there is no one-size-fits-all answer for how many Amazon accounts a particular customer should have, most organizations want to create more than one Amazon account. Multiple accounts provide the highest level of resource and billing isolation.

In the context of SAP workloads, it is common for customers to deploy the Production environment in a separate Amazon account. It helps isolate the production environment from the rest of the SAP landscape.

Amazon accounts 159

Amazon Organizations is an account management service that enables you to consolidate multiple Amazon accounts into an *organization* that you create and centrally manage. Amazon Organizations includes account management and consolidated billing capabilities. It enables you to better meet the budgetary, security, and compliance needs of your business. As an administrator of an organization, you can create accounts in your organization and invite existing accounts to join the organization.

Amazon Landing Zone is a solution that helps customers more quickly set up a secure, multi-account Amazon environment based on Amazon best practices. You can save time by automating the setup of an environment for running secure and scalable workloads while implementing an initial security baseline through the creation of core accounts and resources. It also provides a baseline environment to get started with a multi-account architecture, Amazon Identity and Access Management, governance, data security, network design, and logging.

Note: The Amazon Landing Zone solution is delivered by Amazon Solutions Architects or Professional Services consultants to create a customized baseline of Amazon accounts, networks, and security policies.

Consider using the Amazon Landing Zone solution if you are looking to set up a configurable landing zone with rich customization options through custom add-ons such as, Active Directory, and change management through a code deployment and configuration pipeline.

Amazon Control Tower provides the easiest way to set up and govern a secure, compliant, multi-account Amazon environment based on best practices established by working with thousands of enterprises. With Amazon Control Tower, your distributed teams can provision new Amazon accounts quickly. Meanwhile your central cloud administrators will know that all accounts are aligned with centrally established, company-wide compliance policies.

Consider using the Amazon Control Tower to set up a new Amazon environment based on a landing zone with pre-configured blueprints. You can interactively govern your accounts with pre-configured guardrails.

Compute

Amazon Elastic Compute Cloud (Amazon EC2) provides scalable computing capacity in the Amazon Web Services (Amazon) cloud. An Amazon EC2 instance is launched in a specific Availability Zone within a specified Amazon Virtual Private Cloud (Amazon VPC).

When the Amazon EC2 instances are deployed across two or more Availability Zones within a single Region then Amazon offers an SLA of 99.99%.

Instance types

A range of <u>Amazon EC2 instance types</u> are supported by SAP. When selecting the instance type for your SAP workload, you should consider which tiers allow flexibility on the instance used (application tier). Also consider which tiers will require the use of a specific instance type (database tier) based on compute, memory, storage throughput, and license compliance requirements.

For the tiers with specific instance type requirements and without flexibility to change during a failure scenario, consider having a capacity reservation using <u>Reserved Instances</u> or <u>On-Demand Capacity Reservations</u> within the required Availability Zones and Regions where the instance will run. This approach is called static stability. For more information, see <u>Static stability using</u> Availability Zones.

Reserved Instances

<u>Reserved Instances</u> provide significant savings on your Amazon EC2 costs compared to on-demand instance pricing. Reserved Instances are not physical instances. They are a billing discount applied to the use of On-Demand Instances in your account. To avail the discount benefit, these ondemand instances must match certain attributes, such as instance type and Region.

When you deploy Amazon EC2 across multiple Availability Zones for high availability, we recommend that you use zonal Reserved Instances. In addition to the savings over the on-demand instance pricing, a zonal Reserved Instance provides a capacity reservation in the specified Availability Zone. This ensures that the required capacity is readily available as and when you need it.

For billing purposes, the <u>consolidated billing</u> feature of Amazon Organizations treats all of the accounts in the organization as one account. This means that all accounts in the organization can receive the hourly cost benefit of Reserved Instances that are purchased by any other account.

Savings Plans

<u>Savings Plans</u> is a flexible pricing model that provides savings of up to 72% on your Amazon compute usage. It offers lower prices on Amazon EC2 instance usage, regardless of instance family, size, tenancy or Amazon Region. The Savings Plan model also applies to Amazon Fargate and Amazon Lambda usage.

Savings Plans offer significant savings over on-demand, just like the Amazon EC2 Reserved Instances, in exchange for a commitment to use a specific amount of compute power (measured in \$/hour) for a one- or three-year period.

On-Demand Capacity Reservations

On-Demand Capacity Reservations enable you to reserve capacity for your Amazon EC2 instances in a specific Availability Zone for any duration. This gives you the ability to create and manage Capacity Reservations independently, with the billing discounts offered by Savings Plans or Regional Reserved Instances. You can create Capacity Reservations at any time, you ensure that you always have access to Amazon EC2 capacity when you need it, for as long as you need it. You can create Capacity Reservations at any time, without entering into a one-year or three-year term commitment, and the capacity is available immediately. When you no longer need the reservation, we recommend that you cancel the Capacity Reservation to stop incurring charges for it.

Instance Family Availability across Availability Zones

Certain Amazon EC2 instance families (for example, X1 and High Memory) are not available across all Availability Zones within a Region. You should confirm the instance types required for your SAP workload and check if they are available in your target Availability Zones.

Amazon EC2 auto recovery

<u>Amazon EC2 auto recovery</u> is an Amazon EC2 feature that automatically recovers the instance within the same Availability Zone, if it becomes impaired due to an underlying hardware failure or a problem that requires Amazon involvement to repair.

You can enable auto recovery for Amazon EC2 instances by creating an Amazon CloudWatch alarm which monitors the instance status. Examples of problems that cause system status checks to fail include:

- · Loss of network connectivity
- Loss of system power
- Software issues on the physical host
- Hardware issues on the physical host that impact network reachability

Though it typically takes under 15 minutes for a failed instance to restart, <u>Amazon EC2 auto recovery</u> does not offer an SLA. Therefore, if the recovery of the application that's running on the failed host is critical (for example, SAP Database or SAP Central Services), you should consider using <u>clustering across two Availability Zones</u> to help ensure high availability.

High Memory Bare Metal Dedicated Hosts

<u>Amazon EC2 High Memory Instances</u> are specifically designed to run large in-memory databases, such as SAP HANA. High Memory Bare Metal instances are available on Amazon EC2 <u>Dedicated</u> Hosts on a one- or three-year reservation.

High Memory instances support <u>Dedicated Host Recovery</u>. Host recovery automatically restarts your instances on a new replacement host if failures are detected on your Dedicated Host. Host recovery reduces the need for manual intervention and lowers the operational burden in case of an unexpected Dedicated Host failure.

We recommend a second-High Memory instance in a different Availability Zone of your chosen Region to protect against zone failure.

Amazon EC2 maintenance

When Amazon maintains the underlying host for an instance, it schedules the instance for maintenance. There are two types of maintenance events:

- During network maintenance, scheduled instances lose network connectivity for a brief period of time. Normal network connectivity to your instance is restored after maintenance is complete.
- During power maintenance, scheduled instances are taken offline for a brief period, and then rebooted. When a reboot is performed, all of your instance's configuration settings are retained.

Additionally, we frequently upgrade our Amazon EC2 fleet with many patches and upgrades being applied to instances transparently. However, some updates require a short reboot. Reboots such as these should be infrequent but necessary to apply upgrades that strengthen our security, reliability, and operational performance.

There are two kinds of reboots that can be required as part of Amazon EC2 scheduled maintenance:

- Instance reboots are reboots of your virtual instance and are equivalent to an operating system reboot.
- System reboots require reboots of the underlying physical server hosting an instance.

You can view any upcoming scheduled events for your instances in the Amazon Web Services Management Console or using the API tools or command line.

If you do not take any action, the impact on your instance is the same in both cases: during your scheduled maintenance window your instance will experience a reboot that in most cases takes a few minutes.

Alternatively, you can migrate your instance to a new host by performing a stop and start on your instance. For more information, see Stop and start your instance. You can automate an immediate stop and start in response to a scheduled maintenance event.

Networking

Amazon Virtual Private Cloud and subnets

An <u>Amazon Virtual Private Cloud</u> (Amazon VPC) is a virtual network dedicated to your Amazon account. It is logically isolated from other virtual networks in the Amazon Cloud. You can launch your Amazon resources, such as Amazon EC2 instances, into your VPC.

When you create a VPC, you must specify a range of IPv4 addresses for the VPC in the form of a Classless Inter-Domain Routing (CIDR) block, for example, 10.0.0.0/16. This is the primary CIDR block for your VPC.

You can create a VPC within your chosen Amazon Region and it will be available across all Availability Zones within that Region.

To add a new <u>subnet</u> to your VPC, you must specify an IPv4 CIDR block for the subnet from the range of your VPC. You can specify the Availability Zone in which you want the subnet to reside. You can have multiple subnets in the same zone but a single subnet cannot span across multiple zones.

To provide future flexibility, we recommend that your subnet and connectivity design support all of the available Availability Zones in your account within the Region, regardless of the number of zones that you initially plan to use within a Region.

Latency across Availability Zones

All Availability Zones (AZs) are interconnected with high-bandwidth, low-latency networking, over fully redundant and dedicated metro fiber. This results in single-digit millisecond latency between resources in different Availability Zones in the same Region.

For high availability, we recommend deploying production SAP workloads across multiple Availability Zones, including the SAP Application Server Layer. If you have SAP transactions or

Networking 164

batch jobs that involve significant database calls, we recommend that you run these transactions on SAP Application Servers located in the same Availability Zone as the database and that you use SAP Logon Groups (SMLG) for end users and batch server group (SM61) for background processing jobs. This ensures that latency-sensitive parts of the SAP workload run on the correct application servers.

On premises to Amazon connectivity

You can connect to your VPC through a Site-to-Site <u>virtual private network</u> (VPN) or <u>Amazon Direct</u> <u>Connect</u> from on premises. Amazon Direct Connect provides and <u>SLA</u> of up to 99.99% and Site-to-Site VPN provides an SLA of 99.95%

Site-to-Site VPN connections are to specific Regions. For Direct Connect-based connections, <u>Direct Connect Gateway</u> allows you to connect to multiple Regions.

When establishing connectivity to Amazon from on premises, ensure that you have resilient connections either through the use of multiple Direct Connect Links, multiple VPN connections, or a combination of the two.

The <u>Amazon Direct Connect Resiliency Toolkit</u> provides a connection wizard with multiple resiliency models. These models help you order dedicated connections to achieve your SLA objective.

VPC endpoints

A <u>VPC endpoint</u> privately connects your VPC to supported Amazon services and VPC endpoint services powered by <u>Amazon PrivateLink</u>. It doesn't require internet access via an internet gateway, NAT device, VPN connection, or Amazon Direct Connect connection. Instances in your VPC do not require public IP addresses to communicate with resources in the Amazon service. Traffic between your VPC and other services does not leave the Amazon network.

VPC endpoints are available for all of the core Amazon services cthat are required to support an SAP-based workload, including Amazon EC2 API, Amazon S3, and Amazon Elastic File System.

Cross-Region peering

<u>Amazon Virtual Private Cloud</u> (Amazon VPC) supports <u>Inter-Region peering</u> between two VPCs in different Regions. This can be used to allow network traffic, such as database replication traffic to flow between two <u>Amazon EC2</u> instances in different Regions. Inter-Region peering incurs data transfer costs.

Networking 165

Amazon Transit Gateway is a network transit hub that you can use to interconnect your virtual private clouds (VPC) within an Amazon Region to other VPCs in other Amazon Regions and to on premises networks using Amazon Direct Connect or VPN. Use of Transit Gateway will incur Transit Gateway costs. Amazon Transit Gateway provides an SLA of 99.95% within a Region.

Load balancing

<u>Elastic Load Balancing</u> supports four types of load balancing: Application Load Balancers, Network Load Balancers, Gateway Load Balancers, and Classic Load Balancers.

A <u>Network Load Balancer</u> can be used to support a high-availability deployment of SAP Web Dispatchers and/or SAP Central Services across multiple Availability Zones. For more details, see Overlay IP Routing with Network Load Balancer.

A *load balancer* serves as the single point of contact for clients. The load balancer distributes incoming traffic across multiple targets, such as Amazon EC2 instances.

A *listener* checks for connection requests from clients, using the protocol and port that you configure, and forwards requests to a target group.

Each *target group* routes requests to one or more registered targets, such as Amazon EC2 instances, using the TCP protocol and the specified port number. You can configure health checks on a per target group basis. Health checks are performed on all targets registered to a target group that is specified in a listener rule for your load balancer.

For TCP traffic, the Network Load Balancer selects a target using a flow hash algorithm based on the protocol, source IP address, source port, destination IP address, destination port, and TCP sequence number. Each individual TCP connection is routed to a single target for the life of the connection.

DNS

<u>Amazon Route 53</u> is a highly available and scalable Domain Name System (DNS) web service. You can use Route 53 to perform three main functions in any combination: domain registration, DNS routing, and health checking. Route 53 offers an SLA of 100%.

<u>Amazon Route 53 Resolver</u> provides a set of features that enable bi-directional querying between on premises and Amazon over private connections.

Networking 166

Storage

Object storage

To protect against data loss, you can perform backups (such as database backups or file backups) to Amazon S3. Additionally, <u>Amazon EBS Snapshots</u> and <u>Amazon Machine Images</u> (AMIs) are stored in Amazon S3.

Amazon S3 Replication enables automatic, asynchronous copying of objects across Amazon S3 buckets. Buckets that are configured for object replication can be owned by the same Amazon account or by different accounts.

Amazon S3 replication

You can replicate objects between the same or different Amazon Regions.

- Cross-Region replication (CRR) is used to copy objects across Amazon S3 buckets in different Amazon Regions.
- Same-Region replication (SRR) is used to copy objects across Amazon S3 buckets in the same Amazon Region.

Cross-Region replication incurs the following costs:

- Data Transfer charges for the data transferred between the first and second Amazon Regions
- Amazon S3 charges for the data stored in Amazon S3 in the two different Amazon Regions

Additionally, you can enable Amazon S3 Replication Time Control with cross-Region replication. Amazon S3 Replication Time Control (Amazon S3 RTC) helps you meet compliance or business requirements for data replication and provides visibility into Amazon S3 replication times. Amazon S3 RTC replicates most objects that you upload to Amazon S3 in seconds, and 99.99 percent of those objects within 15 minutes.

Amazon S3 RTC incurs the following costs in addition to the costs listed above for cross-Region replication:

- Amazon S3 RTC Management Feature priced per GB
- Amazon CloudWatch Amazon S3 Metrics priced by number of metrics

Same-Region replication incurs the following costs:

Charges for the data stored in Amazon S3

Block storage

Amazon Elastic Block Store (<u>Amazon EBS</u>) provides block level storage volumes for use with Amazon EC2 instances. Amazon EBS volumes behave like raw, unformatted block devices. You can mount these volumes as devices on your instances. You can create a file system on top of these volumes or use them in any way that you would use a block device (like a hard drive). You can dynamically change the configuration of a volume that's attached to an instance.

Amazon EBS volumes are placed in a specific Availability Zone where they are automatically replicated to protect you from the failure of a single component. All Amazon EBS volume types offer durable snapshot capabilities and are designed for 99.999% availability per volume and 99.99% service availability with Multi-AZ configuration. The use of a database replication capability, block level replication solution or Amazon EBS Snapshots is required to provide durability of the SAP data stored on Amazon EBS across multiple Availability Zones.

Amazon EBS volumes are designed for an annual failure rate (AFR) of between 0.1% - 0.2%, where failure refers to a complete or partial loss of the volume, depending on the size and performance of the volume. This makes Amazon EBS volumes 20 times more reliable than typical commodity disk drives, which fail with an AFR of around 4%. For example, if you have 1,000 Amazon EBS volumes running for 1 year, you should expect 1 to 2 will have a failure.

Amazon EBS offers a number of different <u>volume types</u>. It must be used for the SAP database-related data, General Purpose SSD (gp2) or Provisioned IOPS SSD (io1) must be used. The throughput and IOPS requirement will determine if gp2 or io1 is required.

Amazon EBS Multi-Attach enables you to attach a single Provisioned IOPS SSD (io1) volume to up to 16 <u>Amazon Nitro-based instances</u> that are in the same Availability Zone. You can attach multiple Multi-Attach enabled volumes to an instance or set of instances. Each instance to which the volume is attached has full read and write permission to the shared volume. Multi-Attach enabled volumes do not support I/O fencing. I/O fencing protocols control write access in a shared

storage environment to maintain data consistency. Your applications must provide write ordering for the attached instances to maintain data consistency.

Amazon EBS snapshots

You can back up the data on your Amazon EBS volumes to Amazon S3 by taking point-in-time snapshots. Snapshots are *incremental* backups, which means that only the blocks on the device that have changed after your most recent snapshot are saved. This minimizes the time required to create the snapshot and saves on storage costs by not duplicating data. When you delete a snapshot, only the data unique to that snapshot is removed. Each snapshot contains all of the information that is needed to restore your data (from the moment when the snapshot was taken) to a new Amazon EBS volume.

Amazon EBS Snapshots can be <u>copied</u> (replicated) to a different Region and/or shared with a different Amazon Account.

Copying Snapshots across Regions incurs the following costs:

- Data Transfer charges for the data transferred between the first and second Amazon Regions
- Amazon EBS Snapshot charges for the data stored in Amazon S3 in the two different Amazon Regions

Restoring snapshots

New volumes created from existing Amazon EBS snapshots load lazily in the background. This means that after a volume is created from a snapshot, there is no need to wait for all of the data to transfer from Amazon S3 to your Amazon EBS volume before your attached instance can start accessing the volume and all its data.

This preliminary action takes time and can significantly increase the latency of I/O operations. If your instance accesses data that hasn't yet been loaded, the volume immediately downloads the requested data from Amazon S3, and then continues loading the rest of the volume data in the background.

Fast snapshot restore

Amazon EBS <u>fast snapshot restore</u> enables you to create a volume from a snapshot that is fully-initialized at creation. This eliminates the latency of I/O operations on a block when it is accessed for the first time. Volumes created using fast snapshot restore instantly deliver all of their

provisioned performance. To use fast snapshot restore, enable it for specific snapshots in specific Availability Zones. Fast Snapshot Restore is <u>charged</u> in Data Services Unit-Hours (DSUs) for each zone in which it is enabled. DSUs are billed per minute with a 1-hour minimum.

File storage

Amazon EFS

<u>Amazon Elastic File System</u> (Amazon EFS) provides scalable NFS version 4 based file storage for use with Linux-based Amazon EC2 (Windows-based Amazon EC2 instances do not support Amazon EFS). The service is designed to be highly scalable, available, and durable. Amazon EFS file systems store data and metadata across multiple Availability Zones in an Amazon Region. Amazon EFS offers an SLA of 99.99%.

Amazon EFS file systems can be shared across <u>accounts and VPCs</u> within the same Region or a different Region, enabling Amazon EFS to be an ideal choice for SAP global file system (/sapmnt) and SAP transport directory (/usr/sap/trans).

<u>Amazon DataSync</u> supports <u>Amazon EFS to Amazon EFS transfer</u> between Regions and different Amazon Accounts, allowing the replication of key SAP file based data across Regions. <u>Amazon Backup</u> can also be used to replicate backups of Amazon EFS file systems across Regions.

Amazon FSx

<u>Amazon FSx for Windows File Server</u> provides fully managed Microsoft Windows file servers, backed by a fully native Windows file system. Amazon FSx offers an <u>SLA</u> of 99.9% and supports both Single-AZ and Multi-AZ File Systems.

With Single-AZ file systems, Amazon FSx automatically replicates your data within an Availability Zone, continuously monitors for hardware failures and automatically replaces infrastructure components in the event of a failure. Amazon FSx also takes highly durable backups of your file system daily using Windows' Volume Shadow Copy Service that are stored in Amazon S3. You can take additional backups at any point.

Multi-AZ file systems support all the availability and durability features of Single-AZ file systems. In addition, they are designed to provide continuous availability to data, even when an Availability Zone is unavailable. In a Multi-AZ deployment, Amazon FSx automatically provisions and maintains a standby file server in a different zone. Any changes written to disk in your file system are synchronously replicated across Availability Zones to the standby.

Amazon FSx File systems can be <u>shared across Accounts and VPCs</u> within the same Region or a different Region, enabling Amazon FSx to be used not only for the SAP Global File System but also the SAP Transport Directory.

Additionally, Amazon FSx can also be used for providing <u>Continuously Available (CA) File Shares for Microsoft SQL Server.</u>

Monitoring and audit

Amazon CloudWatch

Amazon CloudWatch is a monitoring and observability service built for DevOps engineers, developers, site reliability engineers (SREs), and IT managers. CloudWatch provides you with data and actionable insights to monitor your applications, respond to system-wide performance changes, optimize resource utilization, and get a unified view of operational health. CloudWatch collects monitoring and operational data in the form of logs, metrics, and events, providing you with a unified view of Amazon resources, applications, and services that run on Amazon and onpremises servers. You can use CloudWatch to detect anomalous behavior in your environments, set alarms, visualize logs and metrics side by side, take automated actions, troubleshoot issues, and discover insights to keep your applications running smoothly.

Amazon CloudTrail

Amazon CloudTrail is a service that enables governance, compliance, operational auditing, and risk auditing of your Amazon account. With CloudTrail, you can log, continuously monitor, and retain account activity related to actions across your Amazon infrastructure. CloudTrail provides event history of your Amazon account activity, including actions taken through the Amazon Web Services Management Console, Amazon SDKs, command line tools, and other Amazon services. This event history simplifies security analysis, resource change tracking, and troubleshooting. In addition, you can use CloudTrail to detect unusual activity in your Amazon accounts. These capabilities help simplify operational analysis and troubleshooting.

Architecture patterns

In this section, we elaborate on the architecture patterns that you can select based on your availability and recovery requirements. We also analyze failure scenarios that can help you select the right pattern for your SAP system(s).

Monitoring and audit 171

General SAP Guides SAP Guides

Failure scenarios

For the failure scenarios below, the primary consideration is the physical unavailability of the compute and/or storage capacity within the Availability Zones.

Availability Zone failure

An Availability Zone failure can be caused by a significant availability degradation of one or more Amazon services utilized by your resources within that Availability Zone. For example:

- Several Amazon EC2 instances have failed with <u>System Status Check errors</u> or are unreachable and cannot be restarted.
- Several Amazon Elastic Block Store (Amazon EBS) volumes with <u>Volume Status Check errors</u> have failed.

Amazon Elastic Block Store failure

Loss of one or more Amazon EBS volumes attached to a single Amazon EC2 instance may result in the unavailability of a critical component (i.e. the database) of the SAP system.

Amazon EC2 failure

Loss of a single Amazon EC2 instance may result in the unavailability of a critical component (i.e. the database or SAP Central Services) of the SAP system.

Logical data loss

You should also consider the potential for logical data loss where the underlying hardware capacity still exists but the primary copies of the data have been corrupted or lost. This data loss could be due to malicious activity within your Amazon account or due to human error.

To protect against logical data loss, it is recommended that regular copies of the data are backed up to an Amazon S3 bucket. This bucket is replicated (using Single-Region or Cross-Region replication) to another Amazon S3 bucket owned by a separate Amazon account. With the appropriate Amazon Identity and Access Management (IAM) controls between the two Amazon accounts, this strategy ensures that not all copies of the data are lost due to malicious activity or human error.

Failure scenarios 172

Patterns

In this section, we examine the architecture patterns available to handle the failure scenarios detailed above.

There are two key parameters to consider when selecting a pattern to meet your organization's specific business requirements:

- · Availability of compute for the SAP single points of failure
- Availability of the SAP data persisted on Amazon EBS

These parameters determine the time taken to recover from a failure scenario, that is, the time taken by your SAP system to return to service.

Types of architecture patterns

The architecture patterns are grouped into single Region and multi-Region patterns. The distinguishing factor would be:

1. If you require the data to only reside in a specific geographical location (Amazon Region) at all times (for example, data residency requirements).

or

2. If you require the data to reside in two specific geographical locations (Amazon Regions) at all times (for example, two copies of SAP data must reside at least 500 miles apart for compliance).

If your production systems are critical to your business and you require minimal downtime in the event of failure, you should select a Multi-Region pattern to ensure that your production systems are highly available at all times. When deploying a Multi-Region pattern, you can benefit from using an automated approach (such as, Cluster solution) for fail over between Availability Zones to minimize the overall downtime and remove the need for human intervention. Multi-Region patterns not only provide high availability but also disaster recovery, thereby lowering overall costs.

Single Region architecture patterns

Select a single Region pattern if:

You require the data to reside only in a specific geographical Region (Amazon Region) at all times

Patterns 173

 You want to avoid the <u>potential network latency</u> considerations associated with a Multi-Region approach

- You want to avoid the cost implications or differences associated with a Multi-Region approach including:
 - Amazon service pricing in different Amazon Regions
 - Cross-Region data transfer costs

Pattern 1: A single Region with two AZs for production

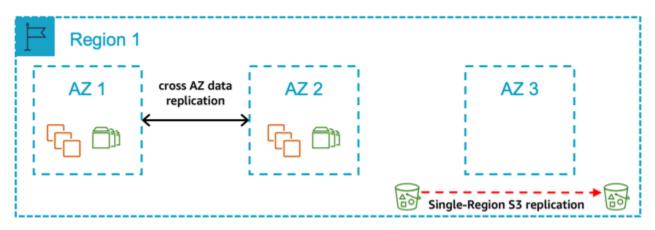


Figure 7: A single Region with two Availability Zones for production

In this pattern, you deploy all your production systems across two Availability Zones. The compute deployed for the production SAP database and central services tiers are the same size in both Availability Zones, with automated fail over in the event of a zone failure. The compute required for the SAP application tier is split 50/50 between two zones. Your non-production systems are **not** an equivalent size to your production and are deployed in the same zones or a different Availability Zone within the Region.

Select this pattern if:

- You require a defined time window to complete recovery of production as well as assurance of the availability of compute capacity in another Availability Zone for the production SAP database and central services tiers.
- You can accept the additional cost of deploying the required compute and storage for production SAP database and central services tiers across two Availability Zones.

 Your non-production environment is not of equivalent size as production and therefore cannot be used as sacrificial capacity for production in the event of an Availability Zone failure or significant Amazon EC2 service degradation.

- You can accept data replication across Availability Zones (database replication capability or a block level replication solution required) and the associated cost.
- You can accept that automated fail over between Availability Zones requires a third-party cluster solution.
- You can accept the variable time duration required (including any delay in availability of the required compute capacity in the remaining Availability Zones) to return the application tier to 100% capacity in the event of a zone failure.

Key design principles

- 100% compute capacity deployed in Availability Zone 1 and Availability Zone 2 for production SAP database and central services tiers.
- Compute capacity is deployed in Availability Zone 1 and Availability Zone 2 for production application tier (Active/Active). In the event of an Availability Zone failure, the application tier needs to be scaled to return to 100% capacity within the remaining zone.
- The SAP Database is persisted on Amazon EBS in two Availability Zones using either a database replication capability or a block level replication solution.
- Amazon EC2 auto recovery is configured for all instances to protect against underlying hardware failure, with the exception of instances protected by a third-party cluster solution.
- Amazon EFS is used for the SAP Global File Systems.
- SAP Database is backed up regularly to Amazon S3.
- Amazon S3 single-Region replication is configured to protect logical data loss.
- Amazon Machine Image/Amazon EBS Snapshots are taken for all servers on a regular basis.

Benefits

- Low Mean Time to Recovery (MTTR)
- Predictable Return to Service (RTS)
- Ability to protect against significant degradation or total Availability Zone failure through fail over of database and central services tiers to Availability Zone 2

 No requirement to restore data from Amazon S3 in the event of an Availability Zone or Amazon EBS failure

Considerations

- Well documented and tested processes are required for the automated fail over between Availability Zones.
- Well documented and tested processes are required for maintaining the automated fail over solution.
- Well documented and tested processes are required for scaling the Amazon resources to return the application tier to required capacity in the event of an Availability Zone failure or significant Amazon EC2 service degradation.

Pattern 2: A single Region with two AZs for production and production sized non-production in a third AZ

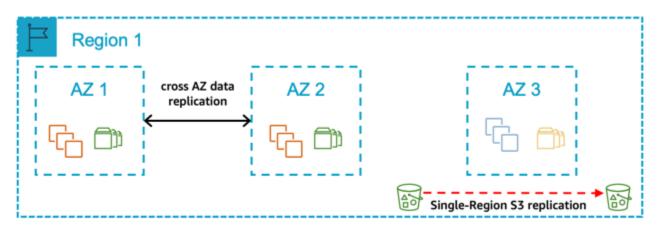


Figure 8: A single Region with two Availability Zones for production and production sized non-production in a third Availability Zone

In this pattern, you deploy all your production systems across two Availability Zones. The compute deployed for the production SAP database and central services tiers are the same size in both Availability Zones, with automated fail over in the event of a zone failure. The compute required for the SAP application tier is split 50/50 between two Availability Zones. Your non-production systems are an equivalent size to your production and deployed in a third Availability Zone. In the event of an Availability Zone failure where your production systems are deployed, the non-production capacity is reallocated to enable production to be returned to a Multi-AZ pattern.

Select this pattern if:

• You require the ability to continue to have a Multi-AZ configuration for production in the event of an Availability Zone failure within the Region.

- You require a defined time window to complete recovery of production and assurance of the
 availability of the compute capacity in another Availability Zone for the production SAP database
 and central services tiers.
- You can accept the additional cost of deploying the required compute and storage for production SAP database and central services tiers across two Availability Zones.
- You can accept data replication across Availability Zones (database replication capability or a block level replication solution required) and the associated cost.
- You can accept that automated fail over between Availability Zones requires a third-party cluster solution.
- You can accept the variable time duration required (including any delay in availability of the required compute capacity in the remaining Availability Zones) to return the application tier to 100% capacity in the event of an Availability Zone failure.

Key design principles

- 100% compute capacity is deployed in Availability Zone 1 and Availability Zone 2 for production SAP database and central services tiers.
- 100% production compute capacity (database and central services) is deployed in the third Availability Zone for use by non-production in normal operations.
- Compute capacity is deployed in Availability Zone 1 and Availability Zone 2 for production application tier (Active/Active). In the event of an Availability Zone failure, the application tier needs to be scaled to return to 100% capacity within the remaining zone.
- <u>Amazon EC2 auto recovery</u> is configured for all instances to protect against underlying hardware failure, with the exception of instances protected by a third-party cluster solution.
- The SAP Database is persisted on Amazon EBS in two Availability Zones using either a database replication capability or a block level replication solution.
- Amazon EFS is used for the SAP Global File Systems.
- SAP Database is backed up regularly to Amazon S3.
- Amazon S3 <u>single-Region replication</u> is configured to protect against <u>logical data loss</u>.
- Amazon Machine Image/Amazon EBS Snapshots for all servers are taken on a regular basis.

Benefits

- Low Mean Time to Recovery (MTTR)
- Predictable Return to Service (RTS)
- Ability to protect against significant degradation or total Availability Zone failure through fail over of database and central services tiers to Availability Zone 2
- No requirement to restore data from Amazon S3 in the event of an Availability Zone failure or Amazon EBS failure
- Option for data to be persisted on Amazon EBS in three different Availability Zones, dependent on capabilities of database or block level replication solution
- Use of non-production compute capacity to return to production run across two Availability Zones in the event of a significant degradation or total Availability Zone failure

Considerations

- Well documented and tested processes are required for the automated fail over between Availability Zones.
- Well documented and tested processes are required for maintaining the automated fail over solution.
- Well documented and tested processes are required for scaling the Amazon resources to return the application tier to required capacity in the event of an Availability Zone failure or significant Amazon EC2 service degradation.
- Well documented and tested processes are required for re-allocating the compute capacity from non-production to return production to run across two Availability Zones in the event of an Availability Zone failure impacting production.

Pattern 3: A single Region with one AZ for production and another AZ for non-production

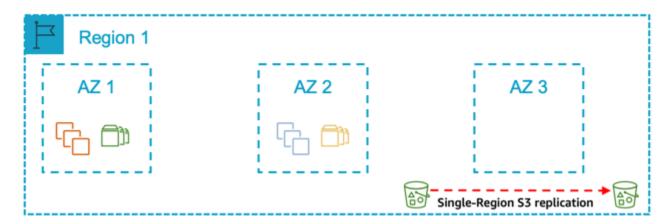


Figure 9: A single Region with one Availability Zone for production and another Availability Zone for non-production

In this pattern, you deploy all your production systems in one Availability Zone and all your non-production systems in another Availability Zone. Your non-production systems are an equivalent size to your production.

Select this pattern if:

- You require a defined time window to complete recovery of production and assurance of the availability of compute capacity in another Availability Zone for the SAP database and central services tiers.
- You can accept the additional time required to re-allocate compute capacity from non-production to production as part of the overall time window to recover production.
- You can accept the time required to restore data to Amazon EBS from Amazon S3 in another Availability Zone as part of the overall time window to recover production.
- You can accept the variable time duration required to return the application tier to 100% capacity following an Availability Zone failure (including any delay in availability of the required compute capacity in the remaining Availability Zones).
- You can accept a period of time where there is only one set of computes deployed for the
 production SAP database and central services tiers in the event of an Availability Zone failure or
 significant Amazon EC2 service degradation.

Key design principles

- 100% compute capacity is deployed in Availability Zone 1 for production SAP database and central services tiers.
- 100% compute capacity is deployed in Availability Zone 1 for production SAP application tier.
- 100% of production compute capacity (SAP database and central services) is deployed in Availability Zone 2 for use by non-production in normal operations.
- Amazon EC2 auto recovery is configured for all instances to protect against underlying hardware failure.
- The SAP database is persisted on Amazon EBS in a single Availability Zone only and not replicated on another Availability Zone.
- Amazon EFS is used for the SAP Global File Systems.
- SAP Database Data is backed up regularly to Amazon S3.

- Amazon S3 single-Region replication is configured to protect against logical data loss.
- Amazon Machine Image/Amazon EBS Snapshots are taken for all servers on a regular basis.

Benefits

- Cost optimized through use of non-production capacity in the event of production Availability
 Zone failure
- Required compute capacity deployed in two Availability Zones to allow a more predictable recovery time duration

Considerations

- Well documented and tested processes for re-allocating the required compute capacity from non-production to production and restoring the data in a different Availability Zone are required to ensure recoverability.
- There may be loss of non-production environments in the event of an Availability Zone failure impacting production.
- Due to the lack of high availability across two Availability Zones, the time required to recover production in the event of compute failure or Availability Zone failure increases.

Pattern 4: A single Region with a single AZ for production

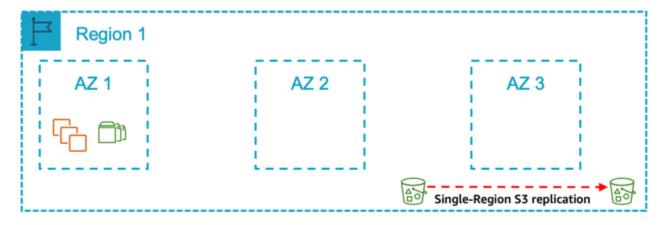


Figure 10: A single Region with a single Availability Zone for production

In this pattern, you deploy all your production systems in one Availability Zone and all your non-production systems in either the same Availability Zone or another Availability Zone. Your non-production systems are **not** a similar size to your production.

Select this pattern if:

In the event of an Availability Zone failure or significant Amazon EC2 service degradation,
you can accept the risks related to the variable time duration required (including any delay in
availability of the required compute capacity in the remaining Availability Zones) to re-create the
Amazon resources in a different Availability Zone and restore the persistent data to Amazon EBS.

 You want to avoid the cost implications with a Multi-AZ approach and accept the related risks of downtime of your production SAP systems.

Key design principles

- 100% compute capacity is deployed in Availability Zone 1 for production SAP database and central services tiers.
- 100% compute capacity is deployed in Availability Zone 1 for production SAP application tier.
- Amazon EC2 is configured for all instances to protect against underlying hardware failure.
- Deployed non-production compute capacity is less than 100% the compute capacity deployed for production SAP database and central services tiers.
- The SAP database is persisted on Amazon EBS in a single Availability Zone only and not replicated on another Availability Zone.
- Amazon EFS is used for the SAP Global File Systems.
- SAP Database is backed up regularly to Amazon S3.
- Amazon S3 single-Region replication is configured to protect against logical data loss.
- Amazon Machine Image/Amazon EBS Snapshots for all servers are taken on a regular basis.

Benefits

- Lowest cost
- Simplest design
- Simplest operation

Considerations

• Well documented and tested processes for scaling the Amazon resources and restoring data in a different Availability Zone are required to ensure recoverability.

Multi-Region Architecture Patterns

You should select a multi-Region architecture if you require the following:

- You require the data to reside in two specific geographical Amazon Regions at all times.
- You can accept the potential network latency considerations associated with a multi-Region approach.
- You can accept the increased complexity associated with multi-Region approach.
- You can accept the cost implications / differences associated with a multi-Region approach including:
 - Amazon service pricing (e.g. Amazon EC2) in different Amazon Regions
 - Cross-Region data transfer costs
- Additional compute and/or storage costs in the second Region.

Pattern 5: A primary Region with two AZs for production and a secondary Region containing a replica of backups/AMIs

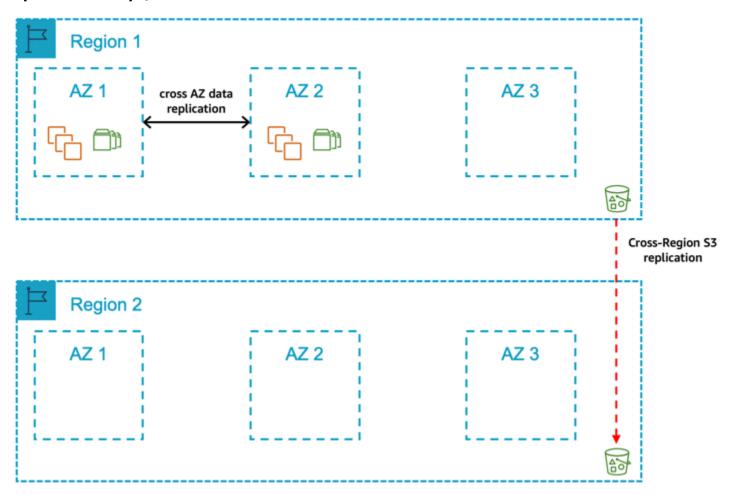


Figure 11: A primary Region with two Availability Zones for production and a secondary Region containing a replica of backups/AMIs

In this pattern, you deploy your production system across two Availability Zones in the primary Region. The compute deployed for the production SAP database and central services tiers are the same size in both Availability Zones with automated fail over in the event of an Availability Zone failure. The compute required for the SAP application tier is split 50/50 between two Availability Zones. Additionally, the production database backups stored in Amazon S3, Amazon EBS Snapshots, and Amazon Machine Images are replicated on the secondary Region. In the event of a complete Region failure, the production systems would be restored from the last set of backups in the second Region.

Select this pattern if:

• You require a defined time window to complete recovery of production and assurance of the availability of compute capacity in another Availability Zone within the primary Region for the production SAP database and central services tiers.

- You can accept the additional cost of deploying the required compute and storage for production SAP database and central services tiers across two Availability Zones within the primary Region.
- You can accept the cross-Availability Zone related data transfer costs for data replication.
- You can accept that automated fail over between Availability Zones requires a third-party cluster solution.
- You can allow for a period of time where there is only one set of computes deployed for the SAP database and central services in the event of an Availability Zone failure or significant Amazon EC2 failure.
- You can accept that data replication across Availability Zones requires either a database replication capability or a block level replication solution.
- You can accept the variable time duration required (including any delay in availability of the required compute capacity in the remaining Availability Zones) to return the application tier to 100% capacity.
- You can accept the variable time duration required to complete recovery of production in the event of a Region failure.
- You can accept the increased complexity and costs associated with multi-Region approach.
- You can accept that manual actions are required to restore production in the second Region.

Key design principles

- 100% compute capacity is deployed in Availability Zone 1 and Availability Zone 2 for production SAP database and central services tiers.
- Compute capacity is deployed in Availability Zone 1 and Availability Zone 2 for production SAP
 application tier (Active/Active) and needs to be scaled in the event of an Availability Zone failure
 or significant Amazon EC2 service degradation.
- <u>Amazon EC2 auto recovery</u> is configured for all instances to protect against underlying hardware failure with the exception of instances protected by a third-party cluster solution.
- The SAP database-related data on Amazon EBS is replicated between Availability Zones using either a database replication capability or a block level replication solution.

General SAP Guides SAP Guides

• Amazon EFS is used for the SAP Global File Systems and is replicated on the secondary Region.

- SAP Database data is backed up regularly to Amazon S3.
- Amazon Machine Image/Amazon EBS Snapshots are taken for all servers on a regular basis

 Amazon S3 Data (database backups), Amazon EBS Snapshots and Amazon Machine Images are replicated/copied to a secondary Region to protect <u>logical data loss</u>.

Benefits

- Low Mean Time to Recovery (MTTR) in the event of Amazon EC2 or Availability Zone failure
- Predictable Return to Service (RTS) in the event of Amazon EC2 or Availability Zone failure
- Database-related data persisted on different sets of Amazon EBS volumes in two Availability
 Zones via database replication capability or a block level replication solution
- Required compute capacity deployed in two Availability Zones in primary Region
- No dependency on restoring data from Amazon S3 in the event of an Availability Zone failure in the primary Region
- Ability to protect against significant degradation or total Availability Zone failure through fail over to Availability Zone 2 of database and central services tiers
- Ability to protect against significant degradation or total Region failure through fail over to secondary Region

Considerations

- Well documented and tested processes are required for the automated fail over between Availability Zones.
- Well documented and tested processes are required for maintaining the automated fail over solution.
- Well documented and tested processes are required for scaling the Amazon resources to return the application tier to full capacity in the event of an Availability Zone failure or significant Amazon EC2 service degradation.
- Well documented and tested processes are required for scaling the Amazon resources, restoring the data, and moving production to the secondary Region.
- Higher network latency from your on-premises locations to the secondary Amazon Region may impact end user performance.

General SAP Guides SAP Guides

Pattern 6: A primary Region with two AZs for production and a secondary Region with compute and storage capacity deployed in a single AZ

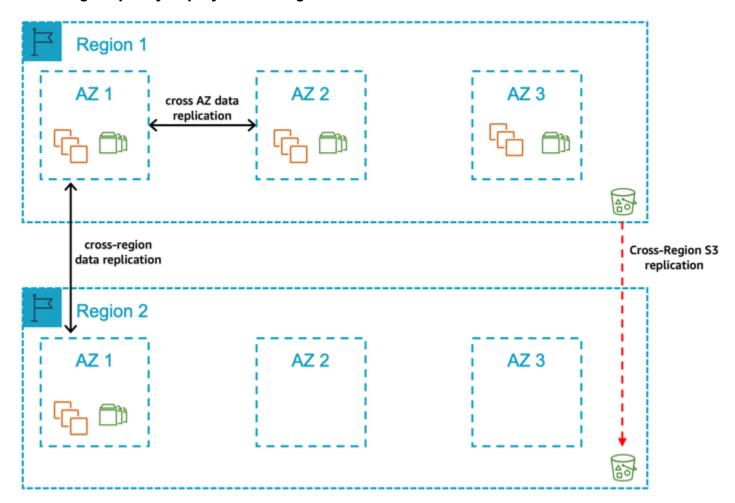


Figure 12: A primary Region with two Availability Zones for production and a secondary Region with compute and storage capacity deployed in a single Availability Zone

In this pattern, you deploy all of your production systems across two Availability Zones in the primary Region. The compute deployed for the production SAP database and central services tiers are the same size in both Availability Zones with automated fail over in the event of Availability Zone failure. The compute required for the SAP application tier is split 50/50 between two Availability Zones. Your non-production systems are **not** an equivalent size to your production and are deployed in a different Availability Zone within the Region. Additionally, compute capacity is deployed in Availability Zone 1 in secondary Region for production SAP database and central services tiers. The production database is replicated to the secondary Region using a database replication capability or a block level replication solution.

The Production database backups stored in Amazon S3, Amazon EBS Snapshots, and Amazon Machine Images are replicated to the secondary Region. In the event of a complete Region failure, the production systems would be restored in the secondary Region using the replicated data for the database tier and the last set of backups for the SAP central services and application tiers.

Select this pattern if:

- You require a defined time window to complete recovery of production and assurance of the availability of compute capacity in another Availability Zone within the primary Region for the production SAP database and central services tiers.
- You can accept the additional cost of deploying the required compute and storage for production SAP database and central services tiers across two Availability Zones within the primary Region.
- You can accept the increased cost of deploying the required compute and storage for production SAP database and central services tiers across two Availability Zones in the primary Region.
- You can accept the cross-Availability Zone related data transfer costs for data replication.
- You can accept that automated fail over between Availability Zones requires a third-party cluster solution.
- You can allow for a period of time where there is only one set of computes deployed for the SAP database and central services in the event of an Availability Zone failure or significant Amazon EC2 failure.
- You can accept that data replication across Availability Zones of the database-related data requires either a database replication capability or a block level replication solution.
- You can accept the variable time duration required (including any delay in availability of the required compute capacity in the remaining Availability Zones) to return the application tier to 100% capacity.
- You require a defined time window to complete recovery of production in the event of a Region failure.
- You can accept the increased complexity and costs associated with multi-Region approach.
- You require assurance of availability of compute capacity in a single Availability Zone in the secondary Region for the production SAP database and central services tiers.
- You can accept the increased cost of deploying the required compute and storage for production SAP database and central services tiers across two Availability Zones in one Availability Zone in the secondary Region.
- You can accept that manual actions are required to fail over between Regions.

Key design principles

• 100% compute capacity is deployed in Availability Zone 1 and Availability Zone 2 for production SAP database and central services tiers.

- 100% compute capacity is deployed in Availability Zone 1 in the secondary Region for production SAP database and central services tiers.
- Compute capacity is deployed in Availability Zone 1 and Availability Zone 2 for production SAP
 application tier (Active/Active) and needs to be scaled in the event of an Availability Zone failure
 or significant Amazon EC2 service degradation.
- <u>Amazon EC2 auto recovery</u> is configured for all instances to protect against underlying hardware failure with the exception of those instances protected by a third-party cluster solution.
- The database-related data on Amazon EBS is replicated between Availability Zones using either a database replication capability or a block level replication solution.
- The SAP database-related data on Amazon EBS is replicated between Regions using either a
 database replication capability or a block level replication solution.
- Amazon EFS is used for the SAP Global File Systems and replicated to the secondary Region.
- SAP database data is backed up regularly to Amazon S3.
- Amazon Machine Image/Amazon EBS Snapshots are taken for all servers on a regular basis
- Amazon S3 data (database backups), Amazon EBS Snapshots, and Amazon Machine Images are replicated/copied to a secondary Region to protect logical data loss.

Benefits

- Low Mean Time to Recovery (MTTR) in the event of an Amazon EC2, Availability Zone or Region failure
- Predictable Return to Service (RTS)
- Database-related data persisted on different sets of Amazon EBS volumes in two Availability
 Zones in primary Region and one set of volumes in an Availability Zone in secondary Region via database replication capability or a block level replication solution
- Required compute capacity deployed in two Availability Zones in primary Region and one Availability Zone in secondary Region
- No dependency on restoring data from Amazon S3 in the event of an Availability Zone failure or Region failure

• Ability to protect against significant degradation or total Availability Zone failure through fail over to Availability Zone 2 of database and central services tiers

 Ability to protect against significant degradation or total Region failure through fail over to secondary Region

Considerations

- Well documented and tested processes are required for the automated fail over between Availability Zones.
- Well documented and tested processes are required for maintaining the automated fail over solution.
- Well documented and tested processes are required for scaling the Amazon resources to return the application tier to full capacity in the event of an Availability Zone failure or significant Amazon EC2 service degradation.
- Well documented and tested processes are required for moving production to the secondary Region.
- Higher network latency from your on-premises locations to the secondary Amazon Region may impact end user performance.
- There is an overhead of maintaining the same software version and patch levels (OS, Database, SAP) across two different Regions.

General SAP Guides SAP Guides

Pattern 7: A primary Region with two AZs for production and a secondary Region with compute and storage capacity deployed and data replication across two AZs

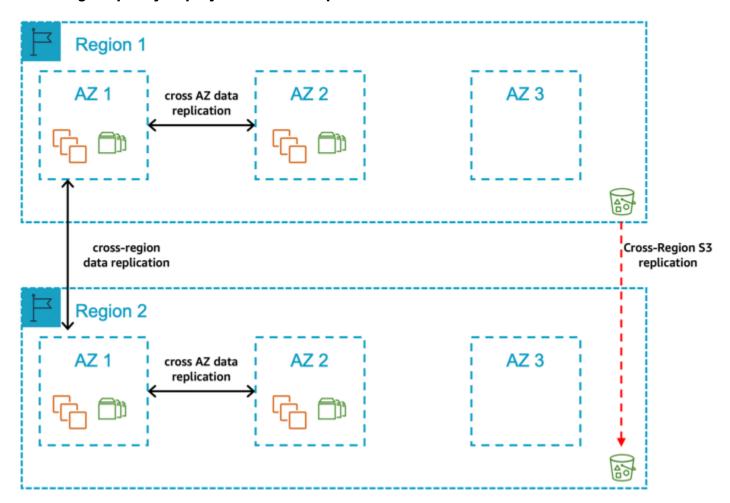


Figure 13: A primary Region with two Availability Zones for production and a secondary Region with compute and storage capacity deployed and data replication across two Availability Zones

In this pattern, you deploy all of your production systems across two Availability Zones in the primary Region. The compute deployed for the production SAP database and central services tiers are the same size in both Availability Zones with automated fail over in the event of Availability Zone failure. The compute required for the SAP application tier is split 50/50 between two Availability Zone. Additionally, you have compute capacity deployed in Availability Zone 1 and Availability Zone 2 in secondary Region for production SAP database and central services tiers and the production database is replicated to the secondary Region using either a database replication capability or a block level replication solution. The production database backups stored in Amazon S3, Amazon EBS Snapshots, and Amazon Machine Images are replicated on a secondary Region. In the event of a complete Region failure, the production systems would be moved over to the secondary Region manually.

Select this pattern if:

• You require a defined time window to complete recovery of production and assurance of the availability of compute capacity in another Availability Zone within the primary Region for the production SAP database and central services tiers.

- You can accept the additional cost of deploying the required compute and storage for production SAP database and central services tiers across two Availability Zones within the primary Region.
- You can allow for a period of time where there is only one set of computes deployed for the SAP database and central services in the event of an Availability Zone failure or significant Amazon EC2 failure.
- You can accept that data replication across Availability Zones of the database-related data requires either a database replication capability or a block level replication solution.
- You can accept the cross-Availability Zone related data transfer costs for data replication.
- You can accept that automated fail over between Availability Zones requires a third-party cluster solution.
- You can accept the variable time duration required (including any delay in availability of the required compute capacity in the remaining Availability Zones) to return the application tier to 100% capacity.
- You require a defined time window to complete recovery of production in the event of a Region failure.
- You require assurance of availability of compute capacity in two Availability Zones in the secondary Region for the production SAP database and central services tiers.
- You can accept the additional cost of deploying the required compute and storage for production SAP database and central services tiers across two Availability Zones in the secondary Region.
- You can accept the increased complexity and costs associated with multi-Region approach.
- You can accept that manual actions are required to fail over between Regions.

Key design principles

- 100% compute capacity is deployed in Availability Zone 1 and Availability Zone 2 in the primary Region for production SAP database and central services tiers.
- 100% compute capacity is deployed in Availability Zone 1 and Availability Zone 2 in the secondary Region for production SAP database and central services tiers.

Compute capacity is deployed in Availability Zone 1 and Availability Zone 2 in the primary
Region for production SAP application tier (Active/Active) and needs to be scaled in the event of
an Availability Zone failure or significant Amazon EC2 service degradation.

- Amazon EC2 auto recovery is configured for all instances to protect against underlying hardware failure with the exception of instances protected by a third-party cluster solution.
- The SAP database-related data on Amazon EBS is replicated between Availability Zones using either a database replication capability or a block level replication solution.
- The SAP database-related data on Amazon EBS is replicated between Regions using either a database replication capability or a block level replication solution.
- Amazon EFS is used for the SAP Global File Systems and replicated on the secondary Region.
- SAP database data is backed up regularly on Amazon S3.
- Amazon Machine Image/Amazon EBS Snapshots for all servers are taken on a regular basis.
- Amazon S3 data (database backups), Amazon EBS Snapshots, and Amazon Machine Images are replicated/copied to a secondary Region to protect logical data loss.

Benefits

- Low Mean Time to Recovery (MTTR) in the event of Amazon EC2, Availability Zone or Region failure
- Predictable Return to Service (RTS)
- Database-related data persisted on different sets of Amazon EBS volumes in two Availability
 Zones in the primary Region and different sets of Amazon EBS volumes in two Availability Zones in the secondary Region via database replication capability or a block level replication solution
- Required compute capacity deployed in two Availability Zones in primary Region and two Availability Zones in secondary Region
- No dependency on restoring data from Amazon S3 in the event of an Availability Zone or Region failure
- Ability to protect against significant degradation or total Availability Zone failure through fail over to Availability Zone 2 of database and central services tiers
- Ability to protect against significant degradation or total Region failure through fail over to secondary Region

Considerations

 Well documented and tested processes are required for the automated fail over between Availability Zones.

- Well documented and tested processes are required for maintaining the automated fail over solution.
- Well documented and tested processes are required for scaling the Amazon resources to return the application tier to full capacity in the event of an Availability Zone failure or significant Amazon EC2 service degradation.
- Well documented and tested processes are required for moving production to the secondary Region.
- Higher network latency from your on-premises locations to the secondary Amazon Region may impact end user performance.
- There is an overhead of maintaining the same software version and patch levels (OS, Database, SAP) across two different Regions.

General SAP Guides SAP Guides

Pattern 8: A primary Region with one AZ for production and a secondary Region containing a replica of backups/AMIs

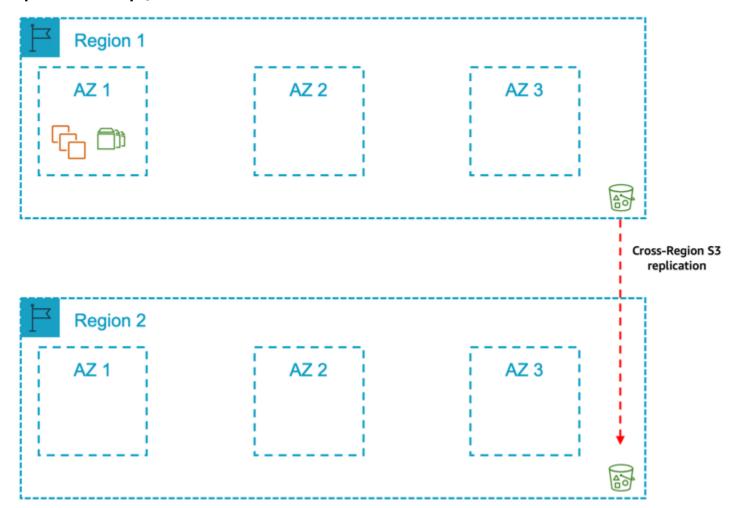


Figure 14: A primary Region with one Availability Zone for production and a secondary Region containing a replica of backups/AMIs

In this pattern, you deploy your production systems in the primary Region in one Availability Zone. Your non-production systems are **not** an equivalent size to your production and are deployed in the same Availability Zones or a different Availability Zone within the Region.

Additionally, the production database backups stored in Amazon S3, Amazon EBS Snapshots, and Amazon Machine Images are replicated to a secondary Region. In the event of a complete Region failure, the production systems would be restored from the last set of backups in the second Region.

Select this pattern if:

In the event of an Availability Zone failure or significant Amazon EC2 service degradation,
you can accept the risks related to the variable time duration required (including any delay in
availability of the required compute capacity in the remaining Availability Zones) to re-create the
Amazon resources in a different Availability Zone and restore the persistent data to Amazon EBS.

- You can accept the risks related to variable time duration required to complete recovery of production in the event of a Region failure.
- You want to avoid the cost implications with a Multi-AZ approach and accept the related risks of downtime of your production SAP systems.
- You can accept the increased complexity and costs associated with multi-Region approach.
- You can accept that manual actions are required to restore production in the secondary Region.

Key design principles

- 100% compute capacity is deployed in Availability Zone 1 for production SAP database and central services tiers.
- 100% compute capacity is deployed in Availability Zone 1 for production SAP application tier.
- <u>Amazon EC2 auto recovery</u> is configured for all instances to protect against underlying hardware failure.
- Deployed non-production compute capacity is less than 100% of the compute capacity deployed for production SAP database and central services tiers.
- The SAP database is persisted on Amazon EBS in a single Availability Zone only and not replicated to another Availability Zone using either a database replication capability or a block level replication solution.
- Amazon EFS is used for the SAP global file systems.
- SAP database is backed up regularly to Amazon S3.
- Amazon S3 is configured to protect against logical data loss.
- Amazon Machine Image/Amazon EBS Snapshots are taken for all servers on a regular basis.
- Amazon S3 data (database backups), Amazon EBS Snapshots, and Amazon Machine Images are replicated/copied to a secondary Region to protect logical data loss.

Benefits

Reduced cost compared to Multi-AZ

 Ability to protect against significant degradation or total Region failure through fail over to secondary Region

Considerations

- Well documented and tested processes are required for scaling the Amazon resources to return the SAP application tier to full capacity in the event of an Availability Zone failure or significant Amazon EC2 service degradation.
- Well documented and tested processes are required for scaling the Amazon resources, restoring the data, and moving production to the secondary Region.
- Higher network latency from your on-premises locations to the secondary Amazon Region may impact end user performance.
- In the event of compute, Availability Zone or Region failure due to lack of high availability across two Availability Zones, there is an increased time required to recover production.

Summary

The table below summarizes the patterns and their key characteristics.

#	Single Region	Multi Region	Single AZ Primary	Multi AZ Primary	Single AZ Second Region	Multi AZ Second Region	Prod Capacity in 2nd AZ	Use of non- prod capacity	Cross- Reg ion data replicati on
1	Yes	No	No	Yes	No	No	Yes	No	No
2	Yes	No	No	Yes	No	No	Yes	Yes	No
3	Yes	No	Yes	No	No	No	No	Yes	No
4	Yes	No	Yes	No	No	No	No	No	No
5	No	Yes	No	Yes	Yes	No	Yes	No	Yes
6	No	Yes	No	Yes	Yes	No	Yes	No	Yes

Summary 196

#	Single Region	Multi Region	Single AZ Primary	Multi AZ Primary	Single AZ Second Region	Multi AZ Second Region	Prod Capacity in 2nd AZ	Use of non- prod capacity	Cross- Reg ion data replicati on
7	No	Yes	No	Yes	No	Yes	Yes	No	Yes
8	No	Yes	Yes	No	Yes	No	No	No	Yes

Table 1: Summary of patterns

With the flexibility and agility of the Amazon Cloud, you have the ability to select any of the patterns described in this guide. You can select the pattern that best meets the business requirements for your SAP systems. It saves you from the trouble of selecting the highest requirement and applying it to all production systems.

For example, if you require highly-available compute capacity in another Availability Zone for the production SAP database and central services tiers of your core ERP system and for your BW system, you can accept the variable time duration required to re-create the Amazon resources in a different Availability Zone and restore the persistent data. In this case, you would select Pattern 3 for ERP and Pattern 1 for BW to reduce the overall TCO.

If your requirements change over time, it is possible to move to a different pattern without significant re-design. For example, during the earlier phases of an implementation project, you may not require highly-available compute capacity in another Availability Zone but you can deploy the capacity into a second Availability Zone a few weeks before go-live.

You should consider the following when selecting an architecture pattern to run your SAP system in Amazon:

- The geographical residency of the data
- The impact of your production SAP systems' downtime on your organization
- The recovery time objective
- The recovery point objective
- The cost profile

Summary 197

SAP on Amazon architecture patterns for Microsoft SQL server

This document provides information about architecture patterns for deploying SAP workloads in Amazon Cloud on Microsoft SQL servers. These patterns offer highly available and resilient implementation options while considering your recovery time and point objectives.

Work backwards from your business requirements to define an approach that meets the availability goals of your SAP systems and data. For each failure scenario, the resiliency requirements, acceptable data loss, and mean time to recover need to be proportionate to the criticality of the component and the supported business applications.

You can customize these patterns for your specific business criteria. You should consider the risk and impact of each failure type, and the cost of mitigation when choosing a pattern.

Topics

- Patterns
- · Comparison matrix
- Single Region architecture patterns for Microsoft SQL server
- Multi-Region patterns for Microsoft SQL server

Patterns

The architecture patterns are divided into two categories.

- Single Region patterns
- Multi-Region patterns

Comparison matrix

The following table provides a comparison of all the architecture patterns discussed further.

Microsoft SQL 198

Patterns	Business re	equirements		Solution cl	haracter	Implementation details	
	Resilienc e type	Recovery point objective	Recovery time objective	Cost	Complexit y	SQL AlwaysOn	Amazon S3 replicati on
Pattern 1	Single Region disaster recovery	Near zero*	Low	Medium	Medium	2-tier	N/A
Pattern 2		Medium	High	Very low	Very low	N/A	N/A
Pattern 3	Multi- Region disaster recovery	Medium	High	Medium	Medium	2-tier	Cross Region
Pattern 4		Near zero*	Low	High	High	3-tier	Cross Region
Pattern 5		Medium	High	Low	Low	N/A	Cross Region
Pattern 6		Low	Low	Medium	Medium	N/A	N/A

^{*}To achieve near zero recovery point objective, database replication must be setup in synchronous data commit mode within the same Amazon Region.

Single Region architecture patterns for Microsoft SQL server

Single Region architecture patterns help you avoid network latency as your SAP workload components are located in a close proximity within the same Region. Every Amazon Region generally has three Availability Zones. For more information, see <u>Amazon Global Infrastructure Map</u>.

You can choose these patterns when you need to ensure that your SAP data resides within regional boundaries stipulated by the data sovereignty laws.

The following are the two single Region architecture patterns.

Single Region patterns 199

Patterns

- Pattern 1: Single Region with two Availability Zones for production
- Pattern 2: Single Region with one Availability Zone for production

Pattern 1: Single Region with two Availability Zones for production

In this pattern, your Microsoft SQL server is deployed across two Availability Zones with AlwaysOn configured on both the instances. The primary and secondary instances are of the same instance type. The secondary instance can be deployed in active/passive or active/active mode. We recommend using the sync mode of replication for the low-latency connectivity between the two Availability Zones.

This pattern is foundational if you are looking for high availability cluster solutions for automated failover to fulfill near-zero recovery point and time objectives. SQL AlwaysOn with Windows cluster for automatic failover provides resiliency against failure scenarios, including the rare occurrence of loss of Availability Zone.

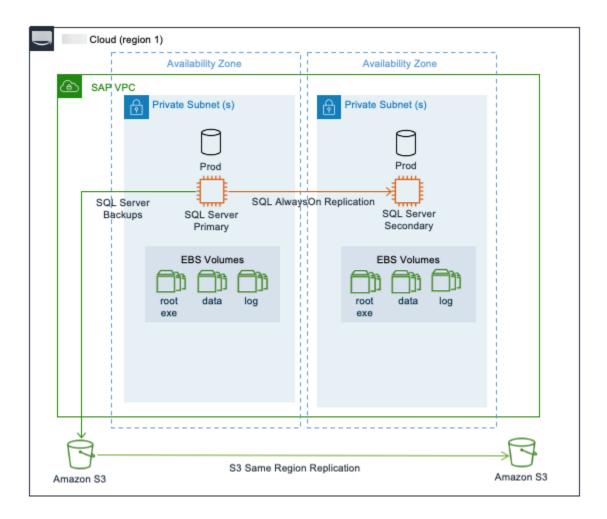
You need to consider the additional cost of licensing for AlwaysOn configuration. Also, provisioning production equivalent instance type as standby adds to the total cost of ownership.

Microsoft SQL server backups can be stored in Amazon S3 buckets. Amazon S3 objects are automatically stored across multiple devices, spanning a minimum of three Availability Zones across a Region. To protect against logical data loss, you can use the Same-Region Replication feature of Amazon S3.

With Same-Region replication, you can setup automatic replication of an Amazon S3 bucket in a separate Amazon account. This strategy ensures that not all copies of data are lost due to malicious activity or human error. To setup Same-Region replication, see Setting up replication.

Single Region patterns 200

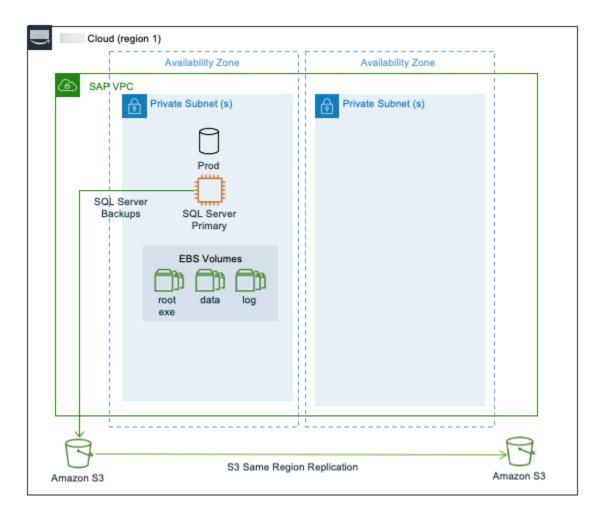
General SAP Guides SAP Guides



Pattern 2: Single Region with one Availability Zone for production

In this pattern, Microsoft SQL server is deployed as a standalone installation with no target systems to replicate data. This is the most basic and cost-efficient deployment option. The options available to restore business operations during a failure scenario are by Amazon EC2 auto recovery, in the event of an instance failure or by restoration and recovery from most recent and valid backups, in the event of a significant issue impacting the Availability Zone.

Single Region patterns 201



Multi-Region patterns for Microsoft SQL server

Amazon Global Infrastructure spans across multiple Regions around the world and this footprint is constantly increasing. For the latest updates, see Amazon Global Infrastructure. If you are looking for your SAP data to reside in multiple regions at any given point to ensure increased availability and minimal downtime in the event of failure, you should opt for multi-Region architecture patterns.

When deploying a multi-Region pattern, you can benefit from using an automated approach such as, cluster solution, for fail over between Availability Zones to minimize the overall downtime and remove the need for human intervention. Multi-Region patterns not only provide high availability but also disaster recovery, thereby lowering overall costs. Distance between the chosen regions have direct impact on latency and hence, in a multi-Region pattern, this has to be considered into the overall design.

There are additional cost implications from cross-Region replication or data transfer that also need to be factored into the overall solution pricing. The pricing varies between Regions.

The following are the four multi-Region architecture patterns.

Patterns

- Pattern 3: Primary Region with two Availability Zones for production and secondary Region with a replica of backups/AMIs
- Pattern 4: Primary Region with two Availability Zones for production and secondary Region with compute and storage capacity deployed in a single Availability Zone
- Pattern 5: Primary Region with one Availability Zone for production and a secondary Region with a replica of backups/AMIs
- Pattern 6: Primary Region with one Availability Zone for production and a secondary Region replicated at block level using Amazon Elastic Disaster Recovery

Pattern 3: Primary Region with two Availability Zones for production and secondary Region with a replica of backups/AMIs

This pattern is similar to pattern 1 where your Microsoft SQL server is highly available. You deploy your production instance across two Availability Zones in the primary Region using AlwaysOn. You can restore your SQL database in a secondary Region with a replica of backups stores in Amazon S3, Amazon EBS, and Amazon Machine Images (AMIs).

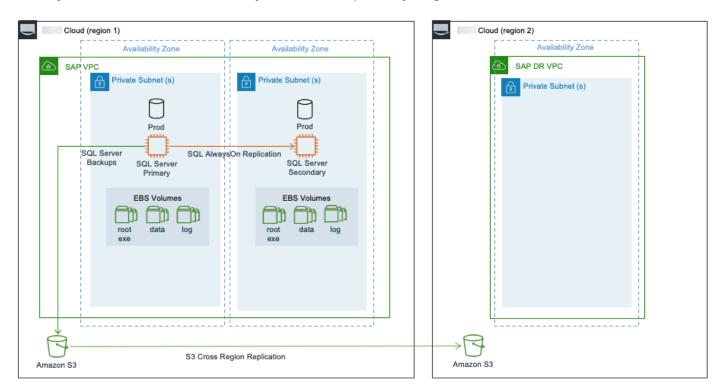
With cross-Region replication of files stored in Amazon S3, the data stored in a bucket is automatically (asynchronously) copied to the target Region. Amazon EBS snapshots can be copied between Regions. For more information, see Copy an AMI. You can copy an AMI within or across Regions using Amazon CLI, Amazon Web Services Management Console, Amazon SDKs or Amazon EC2 APIs. For more information, see Copy an AMI. You can also use Amazon Backup to schedule and run snapshots and replications across Regions.

In the event of a complete Region failure, the production SQL server needs to be built in the secondary Region using AMI. You can use Amazon CloudFormation templates to automate the launch of a new SQL server. Once your instance is launched, you can then download the last set of backup from Amazon S3 to restore your SQL server to a point-in-time before the disaster event. After restoring and recovering your SQL server in the secondary Region, you can redirect your client traffic to the new instance using DNS.

General SAP Guides SAP Guides

This architecture provides you with the advantage of implementing your SQL server across multiple Availability Zones with the ability to failover instantly in the event of a failure. For disaster recovery that is outside the primary Region, recovery point objective is constrained by how often you store your SQL backup files in your Amazon S3 bucket, and the time it takes to replicate your Amazon S3 bucket to the target Region. You can use Amazon S3 replication time control for a time-bound replication. For more information, see Enabling Amazon S3 Replication Time Control.

Your recovery time objective depends on the time it takes to build the system in the secondary Region and restore operations from backup files. The amount of time will vary depending on the size of the database. Also, the time required to get the compute capacity for restore procedures may be more in the absence of a reserved instance capacity. This pattern is suitable when you need the lowest possible recovery time and point objectives within a Region and high recovery point and time objectives for disaster recovery outside the primary Region.



Pattern 4: Primary Region with two Availability Zones for production and secondary Region with compute and storage capacity deployed in a single Availability Zone

In addition to the architecture of pattern 3, this pattern has SQL AlwaysOn setup between the SQL server in the primary Region and an identical third instance in one of the Availability Zones in the

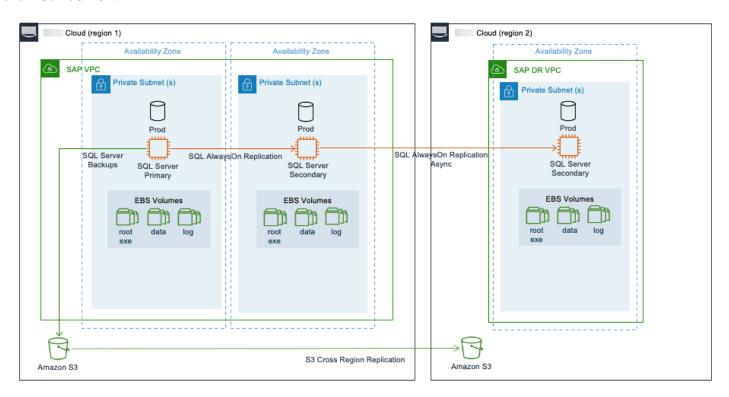
secondary Region. We recommend using the asynchronous (async) mode for SQL AlwaysOn when replicating between Amazon Regions due to increased latency.

In the event of a failure in the primary Region, the production workloads are failed over to the secondary Region manually. This pattern ensures that your SAP systems are highly available and are disaster-tolerant. This pattern provides a quicker failover and continuity of business operations with continuous data replication.

There is an increased cost of deploying the required compute and storage for the production SQL server in the secondary Region and of data transfers between Regions. This pattern is suitable when you require disaster recovery outside of the primary Region with low recovery point and time objectives.

This pattern can be deployed in a multi-tier as well as multi-target replication configuration.

The following diagram shows a multi-tier replication where the replication is configured in a chained fashion.



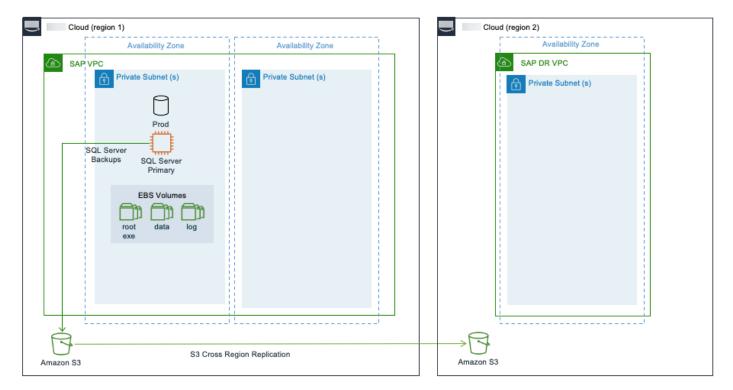
Pattern 5: Primary Region with one Availability Zone for production and a secondary Region with a replica of backups/AMIs

This pattern is similar to pattern 2 with additional disaster recovery in a secondary Region containing replicas of the SQL server backups stored in Amazon S3, Amazon EBS snapshots, and

AMIs. In this pattern, the SQL server is deployed as a standalone installation in the primary Region in one Availability Zone with no target SQL systems to replicate data.

With this pattern, your SQL server is not highly available. In the event of a complete Region failure, the production SQL server needs to be built in the secondary Region using AMI. You can use Amazon CloudFormation templates to automate the launch of a new SQL server. Once your instance is launched, you can then download the last set of backup from Amazon S3 to restore your SQL server to a point-in-time before the disaster event. You can then redirect your client traffic to the new instance in the secondary Region using DNS.

For disaster recovery that is outside the primary Region, recovery point objective is constrained by how often you store your SQL backup files in your Amazon S3 bucket and the time it takes to replicate your Amazon S3 bucket to the target Region. Your recovery time objective depends on the time it takes to build the system in the secondary Region and restore operations from backup files. The amount of time will vary depending on the size of the database. This pattern is suitable for non-production or non-critical production systems that can tolerate a downtime required to restore normal operations.



Pattern 6: Primary Region with one Availability Zone for production and a secondary Region replicated at block level using Amazon Elastic Disaster Recovery

Amazon Elastic Disaster Recovery provides organizations with a modern approach to protecting Microsoft SQL server environments by enabling cloud-based disaster recovery on Amazon Cloud. For more information, see What is Elastic Disaster Recovery?

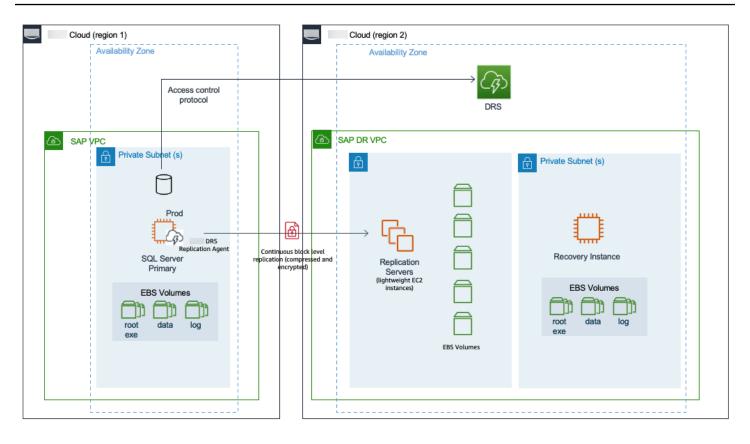
Elastic Disaster Recovery uses block level replication and replicates the operating system, databases, application, and system files for supported Windows and Linux operating system versions. To learn more, see Supported operating systems. An initial setup of the Amazon Replication Agent is required on the source systems for Elastic Disaster Recovery to initiate secure data replication. The agent runs in memory and recognizes write operations to locally attached disks. These writes are captured and asynchronously replicated into a staging area in your Amazon Web Services account. During this ongoing replication process, Elastic Disaster Recovery maintains the write order among all disks in the same source server. The replicated Amazon EC2 instances can be run in a *test mode* to perform drills in a segregated environment.

Elastic Disaster Recovery allows you to monitor the data replication status of your recovery instances, view recovery instance details, add recovery instances to Elastic Disaster Recovery, edit recovery instance failback settings, and terminate recovery instances.

With Elastic Disaster Recovery, you can perform a failover by launching recovery instances on Amazon Cloud. Once the recovery instance is launched, you must redirect the traffic from your primary site to the recovery site.

Amazon Elastic Disaster Recovery uses Amazon EBS snapshots to take point-in-time snapshots of data held within the staging area. To learn more, see Amazon EBS snapshots. It then provides crash consistent point-in-time recovery options that can be used in the event of a disaster or drill. Elastic Disaster Recovery can protect individual nodes of the SQL Server Always On availability group. During disaster recovery, the group is launched as individual SQL server instances on Amazon. This solution works for both the SQL Server Standard edition and SQL Server Enterprise edition for any supported version of the SQL server.

General SAP Guides SAP Guides



Disaster recovery for SAP workloads on Amazon using Amazon Elastic Disaster Recovery

Disasters due to natural events (earthquakes, hurricanes, or floods), application failures, technical failures or human actions cause application downtime and potential data loss, impacting revenue. To mitigate such scenarios, you can create a business continuity plan with the key element of disaster recovery. Designing, implementing, and maintaining a disaster recovery plan is critical for organizations running mission-critical applications, such as SAP. For more information, see Business Continuity Plan (BCP).

Amazon Elastic Disaster Recovery enables organizations to quickly and easily implement a new or migrate an existing disaster recovery plan to Amazon. The source servers can be hosted on Amazon, existing physical or virtual data centers, private cloud or with other cloud providers. We recommend using Elastic Disaster Recovery to implement a disaster recovery plan for your SAP workloads, where Amazon is the disaster recovery environment, and the source environment may or may not be on Amazon. You can access Elastic Disaster Recovery from console.amazonaws.cn.

An initial setup of the Amazon Replication Agent is required on the source systems for Elastic Disaster Recovery to initiate secure data replication. Your data is replicated using secure protocols, either directly over the internet, or via an encrypted and/or dedicated network connection, to any Amazon Region supported by Elastic Disaster Recovery. By replicating the source systems to replication servers in a staging area, the cost of disaster recovery is optimized by using affordable storage, shared servers, and minimal compute resources to maintain ongoing replication.

You can perform non-disruptive tests, known as drills, to confirm that your Elastic Disaster Recovery implementation is ready for a disaster recovery scenario. Elastic Disaster Recovery automatically converts your servers to boot and run natively on Amazon when you launch instances for drills or recovery. The service also automatically creates point in time (PIT) snapshots of your server state as it replicates. If you need to recover applications, you can launch recovery instances on Amazon within minutes, using the latest snapshot or an earlier PIT snapshot. Once your applications are running on Amazon, you can choose to keep them there or initiate data replication back to your primary site when the issue is resolved. You can fail back to your primary site with Elastic Disaster Recovery tools, such as Failback Client.

For more information, see What is Elastic Disaster Recovery?

Topics

- Scenarios
- References
- Service-level agreements and SAP licenses
- Network, storage, and compute
- · Disaster recovery scenarios
- Shared storage resiliency
- Implementing disaster recovery on Amazon cloud for SAP workloads

Scenarios

The following disaster recovery scenarios are covered in this document.

- in-region source workload is running on Amazon cloud and disaster recovery implementation uses a second Availability Zone in the same Amazon Region.
- cross-region source workload is running on Amazon cloud and disaster recovery implementation uses a different Amazon Region. The choice of another Region can be for compliance reasons.
- outside of Amazon source workload is running outside of Amazon (on-premises, public or private cloud) and disaster recovery is implemented with Amazon.

References

This document does not provide detailed steps for setting up and using Amazon Elastic Disaster Recovery. For more information, see <u>Amazon Elastic Disaster Recovery</u>.

It is important to understand the key business requirements that guide a disaster recovery solution design and implementation, including recovery point objectives, recovery time objectives, along with the disaster recovery plan and disaster recovery drill. Check the following resources for concepts related to a disaster recovery implementation on Amazon.

- Amazon Elastic Disaster Recovery Core concepts
- Amazon Well-Architected Framework: Best Practice 10.1
- Architecture guidance for availability and reliability of SAP on Amazon

If you are new to Amazon, see the following documents.

Scenarios 210

- Getting started with Amazon
- What is Amazon EC2?
- What is Amazon VPC?
- Amazon Elastic Block Store (Amazon EBS)

To use this information provided here effectively, you must have previous experience installing, migrating, and operating SAP environments and systems on Amazon, along with high availability and disaster recovery solution implementation.

Service-level agreements and SAP licenses

For a disaster recovery implementation, Service Level Agreements (SLA) are used to define how resilient your system is at avoiding loss of data and reducing downtime when your workload becomes unavailable due to a disaster event.

An SAP system disaster recovery approach requires replication of the application tier, database tier, and any file shares, such as NFS mounts. The following are some of the factors to consider for your disaster recovery implementation.

Topics

- Recovery time objective (RTO)
- Recovery point objective (RPO)
- Recovery consistency objective (RCO)
- SAP licenses

Recovery time objective (RTO)

Recovery Time Objective (RTO) refers to how quickly can your application recover after an outage. In the event of a disaster, Elastic Disaster Recovery enables you to launch your replicated servers to a fully provisioned state at the target Region within minutes and continue operations. This automated approach supports a low RTO. It can be faster and more effective than a manual approach.

Consideration

SLAs and licenses 211

As RTO is usually evaluated in the impact to business processes, other factors such as Domain Name System (DNS) propagation, environmental factors, including your disaster recovery team's reaction time, your target environment's storage architecture, operating system boot, and application startup times, influence this target value.

Recovery point objective (RPO)

Elastic Disaster Recovery continuously replicates the changes to the disk at the block level asynchronously, to the target site. The RPO of Elastic Disaster Recovery is typically in the subsecond range. RPO can be influenced by external factors such as, the time taken by the source system to send changes to the staging area. This is further impacted by the volume of transactions on the source system. Other factors include network throughput and latency, source and replication server performance, etc. These factors should be measured to calculate the potential amount of data loss during a disaster recovery event.

Consideration

SAP workloads may from time to time may observe longer amounts of data loss than what is seen with a sub-second RPO due to how Elastic Disaster Recovery manages certain scenarios.

In the event of hard reboots, disk changes, and crashes, Elastic Disaster Recovery triggers a rescan of the disks. During the rescan, the Replication Agent does not replicate the changes of the source server to the target. This creates a lag between the two servers. If the primary system fails during this time, customers may experience a longer amount of data loss (measured in RPO) than expected.

The rescan time depends on multiple factors, and cannot be predicted without testing. A rescan may occur after a reboot of the source server. The rescan time will vary depending on the size of the source disks. The time depends on the performance of the disks (linear read), staging area disk performance, and the rate of write operations on the source server (which are sent in parallel with the rescan). The rescan is functioning normally as long as its moving forward, and is not "stuck".

SAP databases can have large disk sizes and high change rates. We recommend conducting tests to ensure that your SLA requirements are met in such events. Additionally, you must ensure that the primary and target databases are in sync during peak activity cycles.

Recovery consistency objective (RCO)

Many disaster recovery solutions consider only RTO and RPO as SLAs for resiliency. You must also consider Recovery Consistency Objective (RCO) for your SAP workloads. RCO is a measurement

for the consistency of distributed business data within interlinked systems. In a typical customer environment, SAP systems are tightly integrated and data is frequently exchanged between these systems, like SAP ECC or SAP S/4HANA, SAP BW or SAP BW/4HANA, SAP CRM, SAP SRM, SAP GTS etc. This group of tightly integrated systems is called a *system group*. In case of disaster recovery failover, you may have zero RCO requirement within the system group. This means that in case of disaster recovery failover, all of the databases within the SAP system group must be recovered to the same point-in-time.

Consideration

Elastic Disaster Recovery does not guarantee consistency across multiple source instances. If you have zero RCO requirement, you can use database native replication technology with point-in-time recovery or backtrack with secondary time travel.

For more information, see <u>SAP Note 434647 - Point-in-time recovery in an SAP system group</u> (requires SAP portal access).

SAP licenses

The SAP system is secured by a license using a hardware key. On Amazon, the hardware key is based on your Amazon EC2 instance ID. Your Amazon EC2 instance must be launched before you can generate your SAP license. When you recover your SAP system in the disaster recovery site, the SAP license becomes invalid as the disaster recovery site is a new Amazon EC2 instance. The hardware key will no longer match. A temporary SAP license is created when the recovery instance is launched, and it is valid for 28 days. You do not need to create a new SAP license. If you need the disaster recovery instance to continue running after 28 days, you can request a new SAP license with the recovery Amazon EC2 instance ID.

Network, storage, and compute

This section provides information about configuring network, storage, and compute for staging and target environments to achieve disaster recovery goals for your SAP workloads on Amazon with Elastic Disaster Recovery.

Topics

- Network
- Storage
- Compute

SAP licenses 213

Network

Your network architecture and configuration used for disaster recovery can play a significant role in supporting an effective RTO and RPO SLA. You must consider network design and redirecting traffic to recovery instance when disaster recovery is triggered.

The following are the four steps to design network for disaster recovery.

- the section called "Connecting the source and target network"
- the section called "Defining the staging and recovery subnets"
- the section called "Configuring the network security settings"
- the section called "SAP end user and integration traffic"

Connecting the source and target network

The first step is to choose and configure the network connection method from the source network to the replication servers. You can choose between private or public. For more information, see Data routing and throttling.

Regardless of the method, transferred data is always encrypted in transit. The default method is public, where data is routed over the internet to a public network interface on the replication servers. In the private method, the data is replicated over a private network. A private network selection depends on the disaster recovery scenario in use.

- the section called "Amazon In-Region disaster recovery" Private networks are generally between VPCs, using either Amazon VPC peering or Amazon Transit Gateway for connectivity. We recommend using a different Amazon account, and separate Amazon VPC for disaster recovery. For more information, see <u>What is Amazon VPC peering?</u> and <u>What is a transit</u> gateway?
- the section called "Amazon Cross-Region disaster recovery" We recommend using the fully redundant Amazon network backbone that connects different Amazon Regions together.
 Amazon VPC peering and Amazon Transit Gateway enable connectivity between Regions. For more details, see <u>Introduction to Network Transformation on Amazon</u>.
- the section called "Outside of Amazon to Amazon disaster recovery" In this scenario, your
 physical network between your source network and Amazon are provided through various
 telecommunications or internet services providers (ISP). The following solutions are available on
 Amazon.

Network 214

- Amazon Direct Connect
- Amazon Site-to-Site VPN
- SD-WAN available on Amazon marketplace

Amazon Direct Connect is commonly used by SAP on Amazon customers. It provides more predictable performance against service level agreement (SLA) based targets such as throughput, jitter, and latency, versus VPN or SD-WAN based solutions. You can work with Amazon Direct Connect Delivery Partners for guidance on which options are the best fit for your environment.

Defining the staging and recovery subnets

One subnet is recommended to host the replication servers, called the *staging area subnet*. Additional subnets, called the recovery subnets, are necessary as the target of your disaster recovery action. For scenarios where the source network is on Amazon, consider how your subnets should be allocated based on your selected Amazon Web Services account strategy and landing zone. Often this may mean that the staging area subnets should be in a different Amazon VPC than your source servers. For a simplified environment, this may just use different subnets in the same Amazon VPC. This would mean reduced isolation between your production and non-production disaster recovery environments. For more information, see <u>Amazon Well-Architected Framework : Best Practice 5.3.</u>

Ultimately, the number and design of these subnets should follow similar concepts as your source environment. For more information, see Network diagrams.

For <u>the section called "Amazon In-Region disaster recovery"</u> scenario, we recommend hosting the staging area subnet in a different Availability Zone than the recovery subnets. This design enables an additional redundancy for disaster recovery. The launched recovery instances are protected by a staging area in a separate Availability Zone. This follows the design principle of using multiple Availability Zones to maintain resiliency.

Configuring the network security settings

Ensure that the required network security settings are configured. This includes enabling access through a number of ports in your on-premises firewall, network security devices, security groups, or network access control lists (network ACL), and possibly other tasks depending on the location of your source environment. For more information, see Replication network requirements.

Network 215

SAP end user and integration traffic

The following are some of the factors that affect how the end user and integration related network traffic can affect your RTO and RPO.

- DNS propagation time for clients to identify and resolve to new IP
- Delays in network components (if any used) to reroute traffic, such as global or local load balancers, including Amazon Application Load Balancers, Amazon Global Accelerator, or Amazon Route 53 Public Data Plane

For more information, see Disaster recovery options in the cloud.

Storage

Amazon Elastic Disaster Recovery is designed to evaluate and define the optimal Amazon EBS volume settings for your staging environment based on the source server performance. A default performance setting is used for drill and recovery servers. These volumes are sized to match the capacity needs of the source systems. You must review these settings with the specific requirements of your SAP workloads. This ensures an efficient and disaster recovery SLA compliant environment. These different server types have different requirements, and methods of managing storage.

Topics

- Replication servers
- Drill and recovery instances
- Point in time recovery

Replication servers

The staging area requires storage to support ongoing replication from source machines. These Amazon EBS volumes are usually low-cost, hard disk drive (HDD) type storage volumes. However, if the replicated disk write throughput is high, the default Replication server settings dynamically change to a higher performance, solid state drive (SSD) storage type. The default Amazon EBS volume type setting – **Auto volume type selection** for replication servers, is the recommended setting for SAP workloads. It automatically chooses the high-performing, cost-efficient Amazon EBS volumes for your workload requirements.

Storage 216

You have the option to increase the performance of the staging area by selecting solid state drives (SSD). This can help SAP workloads, such as bursty or consistently high transaction rate databases which have a high rate of create, update, and/or delete operations that must be applied to its storage. For such workloads, we recommend monitoring Amazon CloudWatch metrics and check for any persistent or increasing delays. You can use the following CloudWatch metrics for Elastic Disaster Recovery.

- LagDuration the age of the latest consistent snapshot, in seconds
- Backlog the amount of data yet to be synced, in bytes

If Amazon EBS metrics on the replication server also indicate performance issues, you can change Amazon EBS volume type. See the following resources to learn more.

- Amazon EBS volume performance on Linux instances
- Volumes
- Disk settings

Drill and recovery instances

SAP workloads require at least the gp3 volume type for 90% or more of the use cases, including SAP applications and databases (SAP HANA and any other). If you have a higher per-volume IOPS requirement of more than 16,000 IOPS, or per-volume throughput requirement greater than 1,000 MiB/s, consider io2 or io2 Block Express volumes.

When you launch drill or recovery instances, Elastic Disaster Recovery creates Amazon EBS storage volumes based on the types defined in the launch template. For more information, see Amazon EC2 launch template. The launch template is automatically generated by Elastic Disaster Recovery, with default values for storage performance, using general purpose SSD (volumes sized to match the source system capacity requirements). Review the launch template to confirm that your workload's storage requirements are being met by the default allocations of the launch template.

You can modify the launch template for a different volume type or performance setting. Before modifying, confirm that your target Amazon EC2 instance type supports higher storage. For more details, see <u>Supported instance types</u>. For SAP HANA databases, see <u>Storage configuration</u>. Define the modified version as the default launch template for your server once your changes are applied to the template. We do not recommend adding or removing Amazon EBS volumes in the template when using it with Elastic Disaster Recovery.

Storage 217

For servers that require loading larger amounts of data before they become active, such as database servers, you can configure higher performance settings and types of storage in the launch template. For example, if your server is configured with gp3 storage, then defining more provisioned throughput and IOPS for your storage, and/or using a higher performance scaling storage such as io2 Block Express (with a supported Amazon EC2 instance type), can reduce the time it takes for your drill or recovery instance to handle the expected workload quantity. Once your drill or recovery instance is fully online, you can change revert your storage settings. For more information, see Amazon EBS Elastic Volumes. You can increase the volume size, change the volume type, or adjust the performance of your Amazon EBS volumes, without detaching the volume or restarting the instance.

Point in time recovery

Amazon Elastic Disaster Recovery uses Amazon EBS snapshots to give Point in Time (PiT) recovery options that can be used during a drill or recovery. Amazon EBS snapshots of the staging are volumes are continuously taken to provide recovery points of latest (sub-second RPO), 10-minute increments for the first hour, in one hour increments for 24 hours. A daily PiT is retained for the amount of days specified in your Point in Time (PiT) policy. You can specify between 1 to 365 days, with 7 days being the default. For more information, see <u>Understanding Point In Time states</u>.

Compute

You must choose an Amazon EC2 instance type for both the replication server and the recovery server.

Topics

- Replication servers
- Drill and recovery instances
- Source server

Replication servers

The replication server is normally smaller than the source system. t3.small is the default instance type, and it can replicate up to 15 volumes. You can use a shared replication server between SAP application servers, or other servers with low change rates.

If you have a workload that is bursty or has consistently high transaction rate databases, with a high rate of create, update, and/or delete operations that must be applied to its storage, you

Compute 218

may require different configurations for the staging area. If you see lag in the replication for your workload, change the default replication server to a different instance family. For example, General Purpose Amazon EC2 instance family or use a dedicated replication server. This change can impact cost. For more information, see Replication server configuration.

Drill and recovery instances

For recovery instances, configure the Amazon EC2 launch template settings to match Amazon target instances with source. See the following resources for a list of SAP certified instances.

- SAP NetWeaver certified instances
- SAP HANA certified instances

The following are some of the compute-related factors impacting the RTO of your disaster recovery solutions.

- Server startup time
- SAP running on Microsoft Windows Server operating system
- Large SAP HANA database that takes more than 10 minutes to start up
- SAP application(s) installed on the server, and their startup times
- Mismatch in the source and target server and storage configurations configuring a lesser compute power or storage performance at the target side increases the RTO

You must consider application startup times as a factor in recovery. We recommend choosing an Amazon EC2 instance type and storage configuration that provides an effective startup time. This helps your optimize the RTO for your disaster recovery solutions. Also, performing a disaster recovery test or drill enables you to measure the RTO based on your operating system and database.

SAP systems can run on a variety of operating systems, infrastructure platforms, and processor instruction sets. If your source servers is on-premises or with another cloud provider, it must be compatible with Amazon EC2 and Elastic Disaster Recovery. The source server must have a 64-bit based operating system built for the x86 system architecture. Various x86 based CPUs are available on Amazon, being used on source servers, especially if the servers are old models. Using an SAP sizing-based approach to map the source system to an Amazon EC2 instance type is recommended. To learn more, see SAP's Sizing information.

Compute 219

Source server

While the system requirements for the Replication Agent are relatively low, consider the constraints on the source server for CPU, memory, network, storage, and other resources that can impact the performance of your disaster recovery solution. Size the source server based on these factors. For more information, see <u>Source server requirements</u>.

Disaster recovery scenarios

The following are the three disaster recovery scenarios.

Scenarios

- Amazon In-Region disaster recovery
- Amazon Cross-Region disaster recovery
- Outside of Amazon to Amazon disaster recovery

Amazon In-Region disaster recovery

In Amazon cloud, Availability Zones are separated by a meaningful distance, within 100 km (60 miles away from each other). This distance provides isolation from the most common disasters that could affect data centers, for instance, floods, fire, severe storms, earthquakes, etc. It is used by many Amazon customers today to support their resiliency requirements for SAP workloads. Based on your business continuity requirements, in-Region disaster recovery maybe suitable for you. For more information, see <u>Single Region architecture patterns</u>.

Best practices

- Separate Amazon VPCs for the source and recovery areas in the same Region
- Separate Amazon Web Services accounts for source and recovery areas
- Amazon Transit Gateway or Amazon VPC peering for supporting replication traffic and end user connectivity

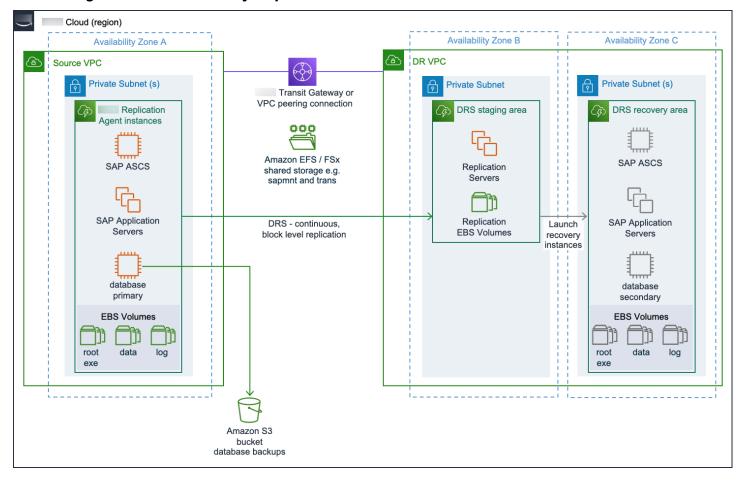
For more information, see the section called "Network".

- Multiple Availability Zones resiliency of Amazon S3 buckets and Amazon EFS for data protection
- Separate Availability Zones for source, staging, and recovery areas

The following two sections cover the reference architectures for this scenario.

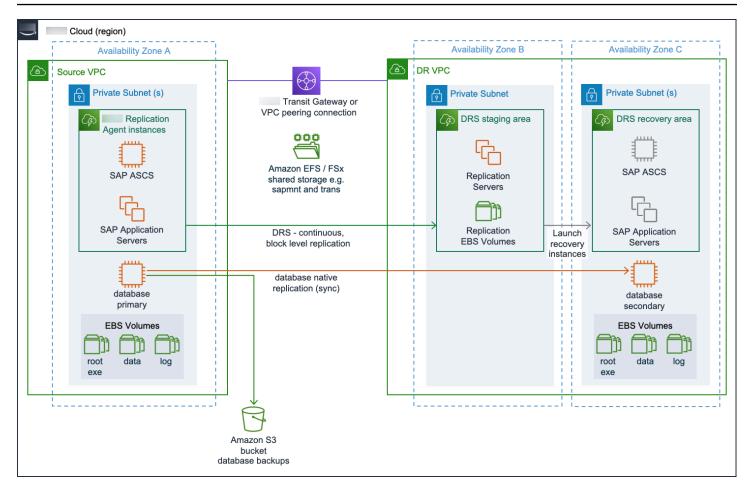
Scenarios 220

Full in-Region disaster recovery implementation



In full in-Region disaster recovery implementation, the source servers running SAP application components, such as central services instance ((A)SCS), primary application server (PAS), additional application server (AAS), and the database, are replicated using Elastic Disaster Recovery.

Hybrid in-Region disaster recovery implementation



In hybrid in-Region disaster recovery implementation, the source servers running SAP application components, such as central services instance [(A)SCS], primary application server (PAS), and additional application server (AAS) are replicated using Elastic Disaster Recovery. The database is replicated using a database native replication method.

Amazon Cross-Region disaster recovery

A disaster recovery scenario with multiple Amazon Regions enables business continuity with your data storage in two separate geographical locations. For more information, see <u>Multi-Region</u> architecture patterns.

Best practices

- Separate Amazon VPCs for the source and recovery areas in the same Region
- A shared Amazon Web Services account for source and recovery areas
- Amazon Transit Gateway or Amazon VPC peering for supporting replication traffic and end user connectivity

For more information, see the section called "Network".

• Replication via Amazon EFS or other file systems to protect shared storage between Regions

For more information, see the section called "Amazon Cross-Region disaster recovery".

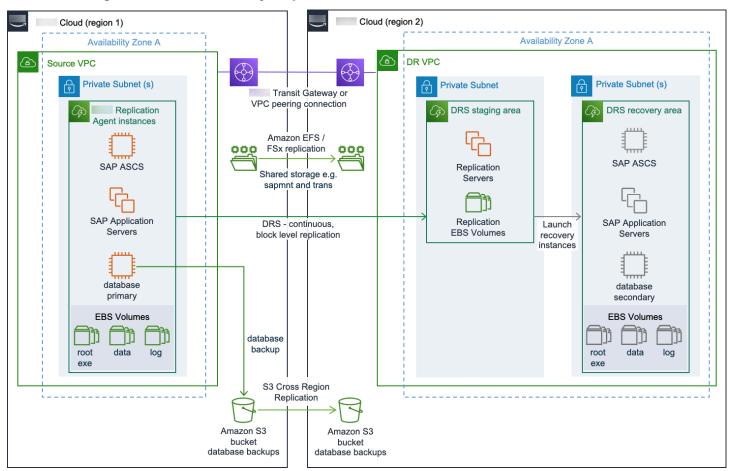
Note

You can only replicate in the same Amazon Web Services account with Amazon EFS.

- Amazon S3 cross-Region replication to provide a copy of your database backups and other
 Amazon S3 bucket data to disaster recovery Amazon VPC
- Separate subnets for the source, staging, and recovery areas

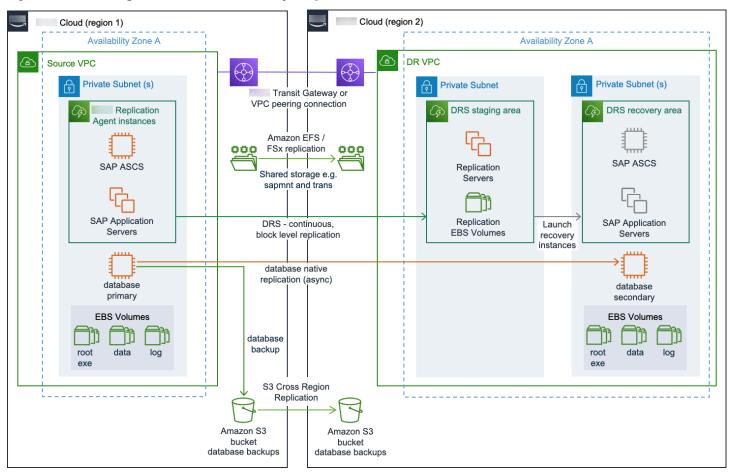
The following two sections cover the reference architectures for this scenario.

Full cross-Region disaster recovery implementation



In full cross-Region disaster recovery implementation, the source servers running SAP application components, such as central services instance ((A)SCS), primary application server (PAS), additional application server (AAS), and the database, are replicated using Elastic Disaster Recovery.

Hybrid cross-Region disaster recovery implementation



In hybrid cross-Region disaster recovery implementation, the source servers running SAP application components, such as central services instance [(A)SCS], primary application server (PAS), and additional application server (AAS) are replicated using Elastic Disaster Recovery. The database is replicated using a database native replication method.

Outside of Amazon to Amazon disaster recovery

In this scenario, the source systems are running in a non-Amazon environment. A hybrid disaster recovery solution like this is can be implemented to quickly add resiliency to your existing production environments on other platforms.

Best practices

• Amazon Direct Connect for supporting replication traffic and end user connectivity

For more information, see the section called "Network".

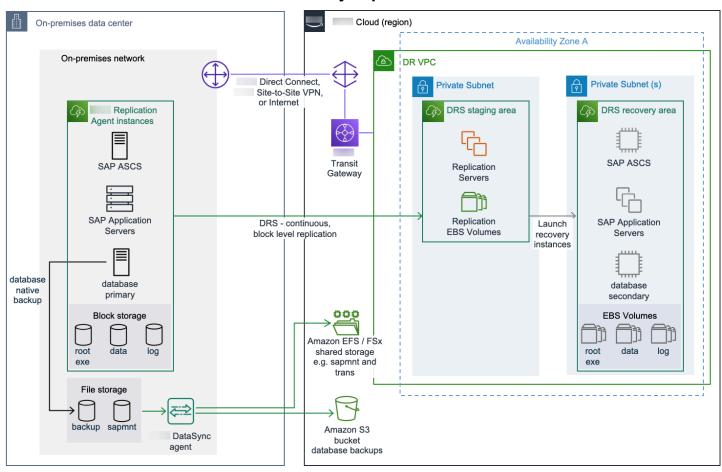
Amazon DataSync to protect shared storage

For more information, see the section called "Outside of Amazon to Amazon disaster recovery".

· Separate subnets for the staging and recovery areas

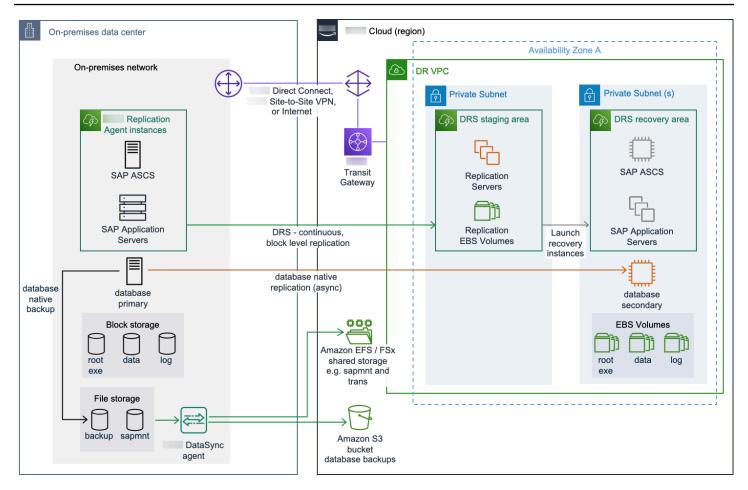
The following two sections cover the reference architectures for this scenario.

Full non-Amazon to Amazon disaster recovery implementation



In full non-Amazon to Amazon disaster recovery implementation, the source servers running SAP application components, such as central services instance ((A)SCS), primary application server (PAS), additional application server (AAS), and the database, are replicated using Elastic Disaster Recovery.

Hybrid non-Amazon to Amazon disaster recovery implementation



In hybrid non-Amazon to Amazon disaster recovery implementation, the source servers running SAP application components, such as central services instance [(A)SCS], primary application server (PAS), and additional application server (AAS) are replicated using Elastic Disaster Recovery. The database is replicated using a database native replication method.

For more disaster recovery options and information, you can reach out to **Amazon Support**.

Shared storage resiliency

File systems on an SAP server can be created on block type storage, for instance, on locally attached disks or Enterprise Storage Area Network (SAN) devices, or may be based on shared file systems such as SMB or NFS shared volumes from servers or Network Attached Storage (NAS) devices.

As Elastic Disaster Recovery is a block level replication service, it only replicates the disks if they are represented as block storage devices. Other tools and processes must be used to provide resiliency for shared file systems. To address these requirements, we recommend using the fully managed shared storage services of Amazon that are easy and cost-effective to launch, run, and

Shared storage resiliency 226

scale feature-rich, high performance, and resilient file systems in the cloud. The choice of your file system depends on the operating system of your disaster recovery scenario.

- Linux Amazon Elastic File System (Amazon EFS)
- Microsoft Windows Server Amazon FSx for Windows File Server (Amazon FSx)
- Mixed Amazon FSx for Windows File Server or Amazon FSx for NetApp ONTAP (FSx for ONTAP)

The following sections provide guidance on file systems based on your disaster recovery scenario.

Topics

- Amazon In-Region disaster recovery
- Amazon Cross-Region disaster recovery
- Outside of Amazon to Amazon disaster recovery

Amazon In-Region disaster recovery

When using managed services such as Amazon EFS, FSx for ONTAP or FSx for Windows File Server to host your shared file systems, the built-in resiliency offered through their multi-Availability Zone design means that your shared storage is already disaster recovery ready. For further resiliency, ensure that your shared storage is backed up regularly, to protect against potential data corruption.

If you are sharing a file system using NFS or SMB protocols directly from one of your Amazon EC2 instances, you may not need additional steps if it is on Amazon EBS and attached to a server with the Replication Agent. This ensures replication via Elastic Disaster Recovery. If the shared file system is hosted on another Amazon EC2 instance along with other content that is not part of your SAP workload, use OS native tools like rsync to manage the replication of this file system to the recovery area.

You can also use Amazon DataSync to provide selective replication. It can be scheduled to run once an hour at minimum, and replicate these files to the target storage in the recovery area. You must have an additional agent installed on an Amazon EC2 instance that has access to the file system. For more information, see How Amazon DataSync works.

Amazon Cross-Region disaster recovery

To support cross-Region disaster recovery, another shared file system must be available in the second Region. The data from the primary shared file system must be replicated on the shared

file system in the second Region. Your implementation will differ based on your choice of Amazon service.

- Amazon Elastic File System Amazon EFS native replication can support cross-Region replication within a single Amazon Web Services account.
- Amazon FSx for Windows File Server You can also use Amazon DataSync to replicate data between your primary to secondary shared storage. For more information, see How Amazon DataSync works.
- Amazon FSx for NetApp ONTAP You can use NetApp SnapMirror to copy your files between FSx for ONTAP file systems on your source and target instances, as frequently as every 5 minutes, to maintain a current copy of your shared file systems. For more information, see SnapMirror.

Outside of Amazon to Amazon disaster recovery

Depending on your source area design for shared storage, you must consider replicating these files on your disaster recovery instance in Amazon. We recommend using <u>Amazon DataSync</u>. It can copy data to and from services, such as NFS and SMB shares, along with file systems using Amazon EFS, FSx for Windows File Server, and FSx for ONTAP.

In certain scenarios, you can consider using other options to protect your source area SAP shared file systems, such as if the following is being used on your source environment.

- FSx for ONTAP You can use NetApp SnapMirror to copy your files between FSx for ONTAP file
 systems on your source and target instances, as frequently as every 5 minutes, to maintain a
 current copy of your shared file systems. For more information, see Scheduled replication using
 NetApp SnapMirror.
- Local storage Elastic Disaster Recovery will replicate it to your disaster recovery environment on Amazon, if Replication Agent can be configured on the source server hosting the local storage.

Implementing disaster recovery on Amazon cloud for SAP workloads

Using Elastic Disaster Recovery to implement a disaster recovery solution for SAP workloads on Amazon follows different considerations for different parts of a typical SAP workload, such as

S/4HANA deployment. The following sections provide guidance on the differences in how to design, implement, and manage Elastic Disaster Recovery when used for the application and the database layers.

Topics

- SAP application layer
- SAP database layer

SAP application layer

We recommend using Amazon Elastic Disaster Recovery to protect your SAP application servers, such as SAP ASCS/SCS, PAS, AAS etc. Elastic Disaster Recovery supports the SAP application layer based on SAP NetWeaver, ABAP foundation, and stand-alone applications like TREX, content servers, and so on. You can use Elastic Disaster Recovery for Amazon EBS backed storages, such as SAP instance binaries, local files stored on an Amazon EBS volume.

The application layer also contains shared file systems, such as SAP mount, transport, and interface directories. These file systems usually need to be managed separately. For more information, see the section called "Shared storage resiliency".

To setup, install Elastic Disaster Recovery agent on the application servers. Create an IAM user with required permissions. Provide Elastic Disaster Recovery agent with the user information to establish a connection with Elastic Disaster Recovery APIs. Once the agent is configured, it engages in an authentication handshake with the TLS 1.3-encrypted Elastic Disaster Recovery API endpoint. The service produces identically sized Amazon EBS volumes in the staging area subnet, for each source volume that is duplicated, for data synchronization. The type of Amazon EBS volumes can be configured in the replication server settings. Replication starts after the staging area subnet resources are generated and the agent is installed. The data is transported with encryption from the source servers to the replication server directly. The service automatically manages the subnet resources for the staging area, scaling them up or down based on the concurrent replication of the source servers and disks.

SAP database layer

Amazon Elastic Disaster Recovery is fully supported as disaster recovery solution for SAP applications running on any database, and for SAP applications running on SAP HANA database in scale-up configuration. It is not supported for replication of multi-node SAP databases, such as SAP HANA scale-out cluster.

SAP application layer 229

The data in an SAP system is stored in a database. This data includes master data, transactional data, and ABAP artifacts. You must consider your business RPO and RTO requirements when evaluation Elastic Disaster Recovery for a disaster recovery solution. The service is not application aware but works at the operating system layer by replicating the attached storage to the target staging environment. Based on your RTO and RPO requirements, you can select Elastic Disaster Recovery or database native replication methods, such as SAP HANA System Replication (HSR) for SAP HANA.

The following are the important considerations to choose your database replication method.

Topics

- · Network bandwidth
- RPO
- Change rate
- RTO
- Cost
- RCO
- Storage limits

Network bandwidth

Amazon Elastic Disaster Recovery works at the operating system layer, with block level replication of attached storage devices. Depending on the change rate at the source, you may need higher network bandwidth to stay current with replication. Database aware technologies such as SAP HSR require lesser network bandwidth, enabling faster replications for systems with high rate of change.

RPO

Elastic Disaster Recovery supports sub-second RPO. For SAP workloads, ensure that your network can support peaks in change rate. If your RPO is very small, we recommend testing database native replication methods along with Elastic Disaster Recovery.

Actions that lead to significant changes to the data of your database cause delays data replication on staging area. It can include a partial or full recovery of a backup to protected volumes for a database on source server. The changes made to your storage volumes are much higher than your usual change rates on source server. Data restored from backup to protected volumes on

SAP database layer 230

the source server is treated as changed blocks and is replicated by Elastic Disaster Recovery. The replication servers need additional time to receive and write this larger amount of changed data from the source system. This can impact your business RPO.

It is recommended to manage actions, such as recovery from backups, at less critical workload times. This way, longer RPO values won't impact your workload. You can track the amount of changed data still waiting to be replicated with Elastic Disaster Recovery. For more information, see Recovery dashboard.

Change rate

For databases with high change rates, you can meet the performance requirements with sufficiently performing networks, along with the storage and compute configuration of the replication server. If these changes are insufficient to meet the business performance requirements, you can choose database native replication methods to optimize your RPO.

RTO

With Elastic Disaster Recovery, the target disaster recovery environment is provisioned once the disaster recovery event is triggered. The total time depends on the size of your database, and the chosen Point-in-Time (PiT). You must test your disaster recovery scenario before implementation on production environments.

Cost

As Elastic Disaster Recovery is not using a warm or hot standby approach, the compute costs are minimized for your disaster recovery environment as compared to many other disaster recovery options. For more information, see <u>Amazon Elastic Disaster Recovery pricing</u>. With database native replication methods, costs can increase with the compute resources in the disaster recovery area.

RCO

If you have multiple tightly coupled systems, then you need to use database native replication methods.

Storage limits

In most cases, the available Amazon EBS volume types are sufficient to address any storage capacity and performance needs. Depending on the source environment architecture, in some cases, the storage volume on the recovery instance exceeds the capacity and/or performance limits

SAP database layer 231

of individual Amazon EBS volumes. This can happen in a non-Amazon to Amazon disaster recovery implementation with data and log volumes attached to high workload database servers. For more information, see Amazon EBS volume types.

When migrating servers to Amazon, such storage volumes must be refactored to a new storage architecture, for instance, creating striped volume sets. Striped volume sets are defined and maintained using logical volume manager tools in your recovery instance's operating system. For more information, see RAID configuration on Linux. These volume sets will span two or more Amazon EBS volumes, up to the total needed to meet the required volume size and performance. The storage volume data is then copied to the new striped volume set. While it may be possible to automate this process through Elastic Disaster Recovery post-launch scripts or alarm events which trigger code through Amazon EventBridge event rules, the additional steps can cause longer recovery time.

In these cases, implementing a hybrid disaster recovery solution is suitable. Most of the servers are managed by Elastic Disaster Recovery and select servers (with storage performance considerations) use alternative disaster recovery approaches, such as native database replication technologies. The storage architecture refactoring is done when the standby replication server is set up during the initial disaster recovery environment implementation. As the replication now happens at an application level, the disaster recovery server is able to write to a storage architecture that is different from what is on the source server.

SAP database layer 232

RISE with SAP on Amazon Cloud

RISE with SAP S/4HANA Cloud, private edition is a cloud ERP offering from SAP. Along with ERP, it includes Business Process Intelligence, Business Platform and Analytics, and Business Networks. SAP maintains responsibility for the holistic service level agreement, cloud operations, and technical support for RISE. You can choose your own cloud service provider in RISE with SAP.

SAP S/4 HANA Cloud, private edition is a single-tenant setup where different customer environments are isolated by Amazon accounts and a dedicated Virtual Private Cloud (VPC).

Important

SAP owns and manages the Amazon Web Services account where RISE with SAP is deployed, and is responsible for the Amazon Web Services used to serve your SAP landscape on Amazon.

SAP is responsible for security in the cloud in RISE with SAP. For more information, see Amazon Cloud Security – Shared Responsibility Model and SAP and Hyperscalers: Clarifying Security in the Cloud. In addition to the security provided by SAP, you can also implement additional security for your SAP landscape. See the the section called "Security" section for more details.

In your Amazon Web Services account managed by SAP, SAP manages the Amazon services required to run your SAP landscape on Amazon. You can still utilize Amazon services to extend RISE with SAP in your own Amazon Web Services account that is not managed by SAP. For example, you can create a data lake with Amazon AppFlow or Amazon Glue. See the the section called "Extensions" section for more details.



Note

You must create a separate Amazon Web Services account or use your existing Amazon Web Services account that is not managed by SAP for creating extensions with Amazon services.

SAP avails Amazon Web Services Support for Amazon Web Services account that is managed by SAP. You are not required to establish additional Amazon Web Services Support for the Amazon Web Services account managed by SAP.

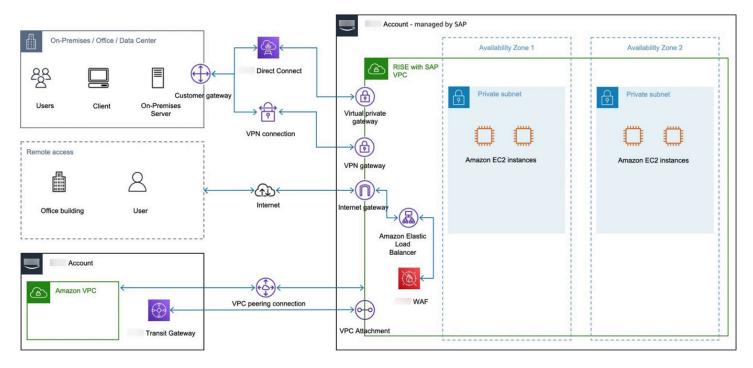
This documentation is focused on RISE with SAP S/4HANA Cloud, private edition and SAP S/4HANA Cloud, private edition, tailored option. The following topics are covered in this document.

Topics

- Connectivity
- Security
- Reliability
- Extensions

Connectivity

You must establish connectivity between Amazon cloud where your RISE with SAP solution is running and on-premises data centers. You also need a connection for direct data transfer (to avoid routing data via your on-premises locations) and communication between SAP systems and your applications running on Amazon cloud. The following images provides an overview of options for connectivity to RISE with SAP VPC.



See the following topics for more details.

Topics

• On-premises to Amazon connectivity

Connectivity 234

- Connecting to RISE with SAP VPC using Amazon VPN
- Connecting to RISE with SAP VPC using Amazon Direct Connect
- Connectivity between Amazon Web Services accounts
- Connecting to RISE with SAP using your Amazon Web Services account
- SAP Business Technology Platform (BTP) with RISE with SAP on Amazon

On-premises to Amazon connectivity

RISE with SAP on Amazon connectivity is supported by Amazon VPN and Amazon Direct Connect. You can connect your Amazon VPC and RISE with SAP VPC via your Amazon Web Services account or by establishing a direct connection between the two.

For more information, see the following resources.

- Connect your VPC to remote networks using Amazon Virtual Private Network
- Amazon Direct Connect connections

Connecting to RISE with SAP VPC using Amazon VPN

Enable access to your remote network from RISE with SAP VPC using <u>Amazon Site-to-Site VPN</u>. Traffic between Amazon cloud and your on-premises location is encrypted via Internet Protocol security (IPsec) and transferred through a secure tunnel on internet. This option is efficient, fast, and more cost-optimized when compared to Amazon Direct Connect. You can get a maximum bandwidth of up to 1.25 Gbps per VPN tunnel. For more information, see <u>Site-to-Site VPN</u> quotas.

Connecting to RISE with SAP VPC using Amazon Direct Connect

Use Amazon Direct Connect if you require a higher throughput and more consistent network experience than an internet-based connection. Amazon Direct Connect links your internal network to an Amazon Direct Connect location over a standard Ethernet fiber-optic cable. You can create virtual interfaces to public Amazon services. For example, you can create interfaces to Amazon S3 or Amazon VPC while bypassing the internet service providers in your network path.

You can choose from a dedicated connection of 1 Gbps, 10 Gbps or 100 Gbps Ethernet port dedicated to a single customer, or a Amazon Direct Connect Partner's hosted connection where the Partner has an established network link with Amazon cloud. Hosted connections are available from 50 Mbps up to 10 Gbps. You can order hosted connections from an Amazon Direct Connect

Delivery Partner approved to support this model. For more information, see Amazon Direct Connect Delivery Partners.

To connect, use a virtual private gateway in Amazon Web Services account managed by SAP or a Direct Connect gateway in your Amazon Web Services account associated with a virtual private gateway in Amazon Web Services account managed by SAP. For more information, see Direct Connect gateways. Direct Connect gateway can also connect to a Amazon Transit Gateway. For more information, see the section called "Connecting to RISE with SAP using your Amazon Web Services account".

You must acquire a *Letter of Authorization* from SAP to setup a Amazon Direct Connect connection in the Amazon Web Services account managed by SAP.

Amazon Direct Connect internet access

Inbound internet traffic can be implemented to access SAP systems via internet gateway. Traffic is inspected by Amazon WAF and protected by Amazon Shield. An elastic load balancer in front of the SAP system enables the use of private IPs for the SAP application server.



Note

SAP only allows HTTPS or TLS1.2 for inbound internet connections.

Connectivity between Amazon Web Services accounts

The following two options are available to connect your Amazon Web Services account with Amazon Web Services account managed by SAP on the third layer.

Options

- Amazon VPC peering
- Amazon Transit Gateway

Amazon VPC peering

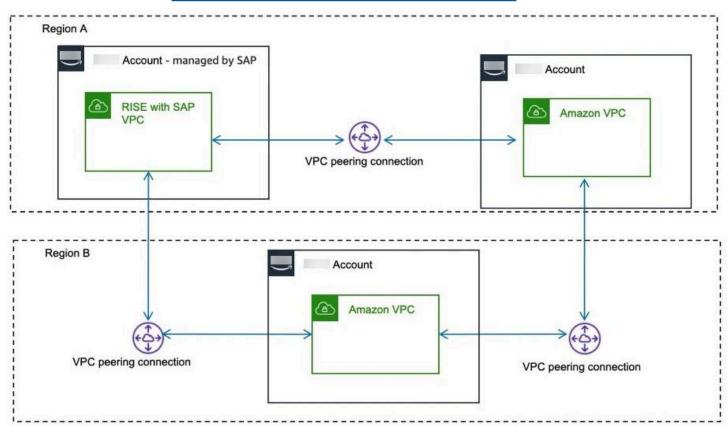
VPC peering enables network connection between two VPCs using private IPv4 and IPv6 addresses. Instances can communicate over the same network. For more information, see What is VPC peering?

Before setup a peering connection, you need to create a request for SAP's approval. For a successful VPC peering, the defined IPv4 Classless Inter-Domain Routing (CIDR) block must not overlap. Check with SAP for the CIDR ranges that can be used in RISE with SAP VPC.

VPC peering is one-on-one connection between VPCs, and is not transitive. Traffic cannot transit from one VPC to another via an intermediary VPC. You must setup multiple peering connections to establish direct communication between RISE with SAP VPC and multiple VPCs. For more details, see the section called "Amazon Transit Gateway".

VPC peering works across Amazon Web Services Regions. All inter-Region traffic is encrypted with no single point of failure or bandwidth bottleneck. Traffic stays on Amazon Global Network and never traverses the public internet, reducing threats of common exploits and DDos attacks. Traffic is encrypted using AES-256 encryption at the virtual network layer.

Data transfer for VPC peering within an Availability Zone is free, and for across Availability Zones is charged per-GB. For more information, see Amazon EC2 pricing. In your Amazon Web Services account, use the Availability Zone ID of Amazon Web Services account managed by SAP to avoid cross-Availability Zone data transfer charges. You can ask for the Availability Zone ID from SAP. For more information, see Availability Zone IDs for your Amazon resources.

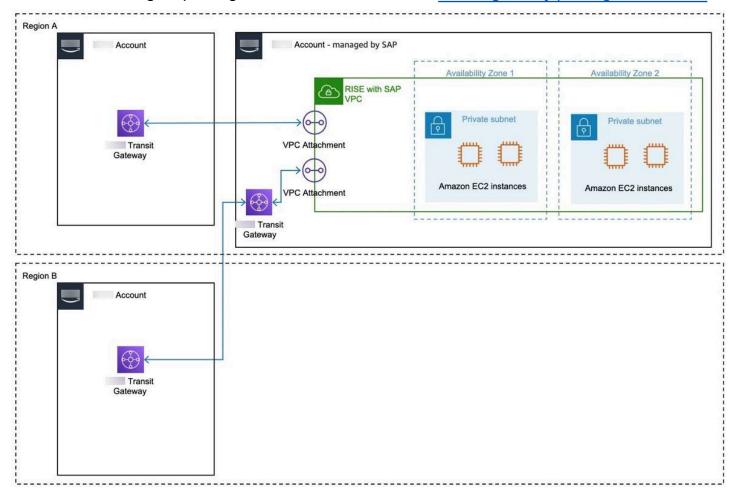


Amazon Transit Gateway

Amazon Transit Gateway is a network transit hub to interconnect Amazon VPCs. It acts as a cloud router, resolving complex peering setup issues by acting as the central communication hub. You need to establish this connection with Amazon Web Services account managed by SAP only once.

To establish connection with Amazon Web Services account managed by SAP, create and share Amazon Transit Gateway in your Amazon Web Services account. SAP then creates an attachment to enable traffic flow through an entry in route table. As Amazon Transit Gateway resides in your Amazon Web Services account, you can retain control over traffic routing.

For cross-Region peering, connect Amazon Transit Gateway in Amazon Web Services account managed by SAP with Amazon Transit Gateway in a different Region in your Amazon Web Services account. Amazon Transit Gateway in Amazon Web Services account managed by SAP is currently limited to cross-Region peering. For more information, see Transit gateway peering attachments.



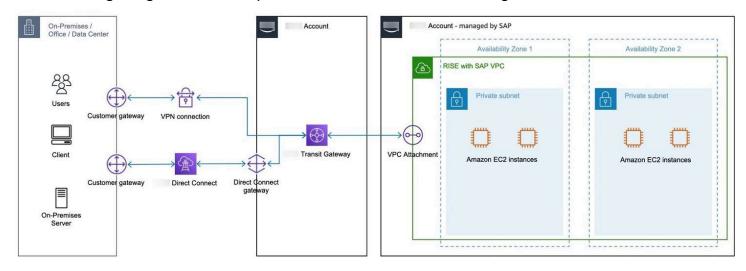
Connecting to RISE with SAP using your Amazon Web Services account

You can establish connectivity between on-premises and RISE with SAP VPC using you Amazon Web Services account. This method provides you with more control but also requires managing Amazon Web Services in your Amazon Web Services account. You can use any one of the following options.

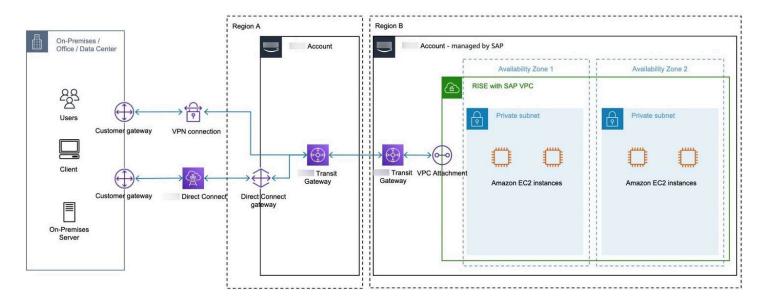
- Amazon Transit Gateway Share Amazon Transit Gateway resource in you Amazon Web Services account with Amazon Web Services account managed by SAP.
 - To scale beyond the default maximum limit of 1.25 Gbps of a single VPN tunnel, see How can I achieve ECMP routing with multiple Site-to-Site VPN tunnels that are associated with a transit gateway?
 - For inter-Region connection, use VPC peering. See the section called "Amazon VPC peering".
- Amazon VPN with Amazon Transit Gateway Create an IPsec VPN connection between your remote network and transit gateway over the internet. For more information, see How Amazon Site-to-Site VPN works and Transit gateway VPN attachments.
- Direct Connect gateway Create a Direct Connect gateway with a transit or virtual private gateway. For more information, see Transit gateway attachments to a Direct Connect gateway.

To strengthen the security, see <u>How do I establish an Amazon VPN over an Amazon Direct</u> Connect connection?

The following image shows this option within the same Amazon Region.



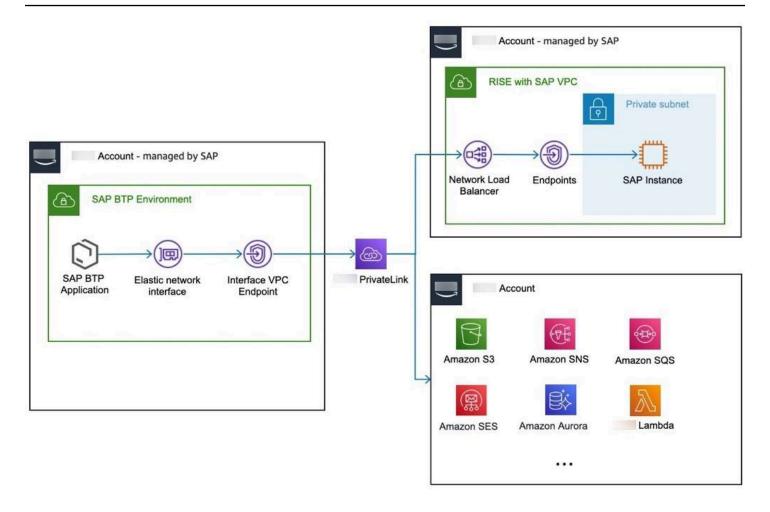
The following image shows this option across different Amazon Regions.



SAP Business Technology Platform (BTP) with RISE with SAP on Amazon

You can use SAP Business Technology Platform BTP services on Amazon to extend the functionality of the RISE with SAP. SAP recommends SAP Cloud Connector to connect RISE with SAP VPC with SAP BTP via internet. When both RISE with SAP and SAP BTP run on Amazon, the network traffic is encrypted and contained within Amazon Global Network, without going through the internet. This provides better security and performance for any integration use-cases between RISE with SAP and SAP BTP. For more information, see Amazon VPC FAQs.

SAP also offers SAP Private Link Service for SAP BTP on Amazon. SAP Private Link connects SAP BTP on Amazon with a secure connection without using public IPs in your Amazon Web Services account.



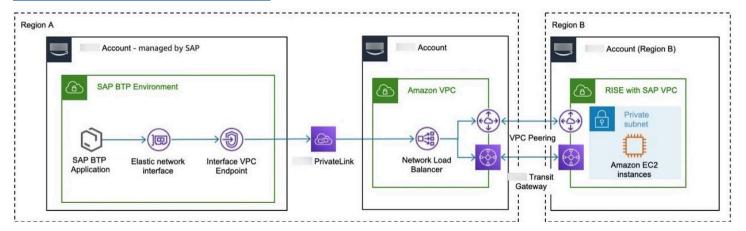
You can connect to an Amazon endpoint service from an SAP BTP application running on Cloud Foundry. By establishing this connection, you can directly connect to Amazon Web Services or for example. to an S/4HANA system. For a complete list of supported Amazon Web Services, see Consume Amazon Web Services in SAP BTP.

You can establish a secure and private communication between SAP BTP and Amazon Web Services with <u>SAP Private Link Service</u>. By using private IP address ranges (RFC 1918), you reduce the attack surface of the application. The connection does not require an internet gateway. If you do not require this extra layer of security, you can still connect via the public APIs of SAP BTP without SAP Private Link, and benefit from Amazon global network. For more information, see <u>Amazon VPC FAQs</u>.

SAP Private Link for Amazon currently supports connections initiated from SAP BTP Cloud Foundry to Amazon.

For Amazon Web Services across Amazon Regions, you can create a VPC in the same Amazon Region as your SAP BTP Cloud Foundry Runtime, and connect these VPCs via VPC peering or

Amazon Transit Gateway. For a list of supported Regions, see <u>Regions and API Endpoints Available</u> for the Cloud Foundry Environment.



Security

SAP manages the security in Amazon Web Services account managed by SAP. You can implement additional security mechanisms in your own Amazon Web Services account.

Topics

- Single Sign-On SAP Business Technology Platform (BTP) and Amazon IAM Identity Center
- Use Amazon services to enhance security in RISE with SAP

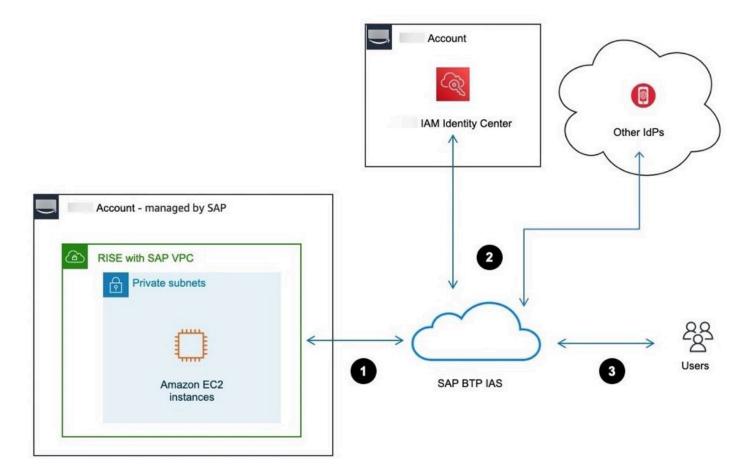
Single Sign-On – SAP Business Technology Platform (BTP) and Amazon IAM Identity Center

In RISE with SAP, you can integrate your own Identity Provided (IdP), such as Amazon IAM, Okta, Ping, Microsoft Windows Azure Active Directory, and others using Identity Authentication of SAP BTP.

- For more information about Amazon IAM Identity Center, see What is IAM Identity Center?
- For more information about Identity Authentication, see <u>SAP Cloud Identity Services Identity</u> Authentication on SAP Help Portal.

The following image shows the integration between Identity Authentication from SAP BTP and Amazon IAM Identity Center.

Security 242



Authentication flow

- User accesses SAP Fiori via an Internet browser.
- 2. Identity Authentication intercepts and routes request to IAM Identity Center.
- 3. Access to SAP S/4HANA in RISE with SAP VPC is enabled on successful authentication.

Use Amazon services to enhance security in RISE with SAP

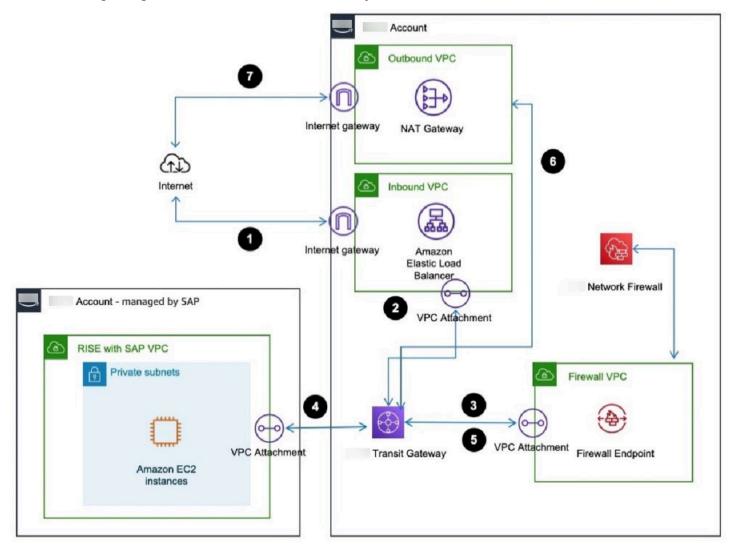
Integrate RISE with SAP VPC with Amazon VPC in your own Amazon Web Services account that is not managed by SAP at the network layer to implement security mechanisms. For more information, see VPC to VPC connectivity.

The network traffic can be routed through Amazon services that strengthen security and operate on a network layer, coming in and out of the RISE with SAP VPC. For more information, see Infrastructure OU – Network account.

You can use the following Amazon services to enhance RISE with SAP security.

- Amazon Shield to mitigate risks against denial of service (DDoS) attacks.
- Amazon Network Firewall to implement intrusion detection and prevention.
- Amazon WAF to protect web applications such as, SAP Fiori, against attacks.
- Amazon Certificate Manager to handle management of public SSL certificates.
- <u>Amazon CloudFront</u> (including Amazon Shield, and optionally Amazon WAF) to implement security at edge for protecting SAP workloads.
- Amazon Partner Solutions on <u>Amazon Web Services Marketplace</u> for more advanced security requirements.

The following image shows RISE with SAP security enhancement with Amazon services.



Traffic flow

- User accesses SAP Fiori via an Internet browser.
- 2. The traffic is routed through the inbound VPC managed by you. The entire traffic flow is via Amazon Transit Gateway.
- 3. The traffic is routed to the firewall VPC managed by you. Based on your company's security policies, the traffic may be filtered by a Network Firewall.
- 4. The traffic is now passed to an Amazon EC2 instance running in Amazon Web Services account managed by SAP.
- 5. Amazon EC2 instance response is routed back to customer managed firewall VPC.
- 6. The response traffic is now passed on to the outbound VPC managed by you.
- 7. The response traffic reaches the user.

Reliability

Reliability is one of the six pillars of SAP Lens - Amazon Well-Architected Framework. For more information, see Reliability.

Amazon cloud offers reliability with multiple Availability Zones within an Amazon Region. This enables your SAP applications on Amazon to be more resilient. Each Region is further isolated from other Regions, providing the greatest possible fault tolerance and stability. Within each Amazon Region, there are a minimum of three, isolated, physically separate Availability Zones. For more information, see Regions and Availability Zones.

Availability Zones enable you to operate production applications and databases that are more highly available than would be possible from a single data center. Distributing your applications across multiple Availability Zones provides the ability to remain resilient in the face of most failure modes, including natural disasters or system failures.

Each Availability Zone can be multiple data centers. At full scale, it can contain hundreds of thousands of servers. They are fully isolated partitions of Amazon Global Infrastructure. An Availability Zone is physically separated from any other zones. There is a distance of several kilometers, although all are within 100 km (60 miles of each other). This distance provides isolation from the most common disasters that could affect data centers, such as floods, fire, severe storms, earthquakes, etc.

All Availability Zones within a Region are interconnected with high-bandwidth and low-latency networking, over fully redundant and dedicated metro fiber. This ensures high-throughput,

Reliability 245

low-latency networking between Availability Zones. The network performance is sufficient to accomplish synchronous replication.

Availability Zones enable you to run your applications in a highly-available manner, with synchronous data replication and automated failover between Availability Zones. RISE with SAP can offer such high available designs for your workload in every Amazon Region.

Considerations

SAP has options available for RISE to meet different resiliency requirements. The following key requirements are adjustable for RISE via option packages available from SAP.

- Service Level Agreement (SLA) describes the targeted availability of the solution.
- Recovery Time Objective (RTO) describes the targeted duration within which a recovery from a disaster event should be completed.
- Recovery Point Objective (RPO) describes the targeted level of data loss that may occur during recovery from a disaster event.

For more details, refer to the definitions provided by SAP as part of RISE agreement for specific definitions, clauses, impacts, and penalties in the event of a breach.

The impact of an outage on your organisation and loss of data can cause loss of productivity, loss of income, and can damage reputation. Weighing the trade-off between cost and resiliency can help assess the risk to your organisation.

Disaster recovery options

You can implement a disaster recovery solution by replicating data into a second Amazon Region. Your SAP workloads are protected in the event of rare occurrence of local or regional failures.

RISE with SAP S/4HANA Cloud, private edition offers the following two options.

- Short distance disaster recovery or Metro disaster recovery RISE with SAP uses multiple Availability Zones in an Amazon Region. Isolated Regions and low latency between Availability Zones enables a synchronous replication between the primary and standby instances.
- Long distance disaster recovery or Regional disaster recovery RISE with SAP uses a secondary Amazon Region as standby for failover systems. Owing to the physical distance between two Amazon Regions, data is replicated asynchronously between two Amazon Regions.

Reliability 246

For more details, see SAP documentation <u>SAP Service Description</u>: <u>Disaster Recovery and Customer</u> Invoked Failover.

Extensions

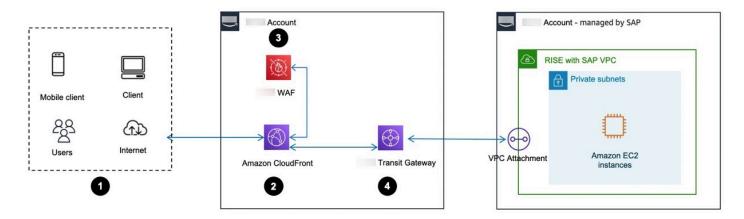
You can extend RISE with SAP by using Amazon services to improve performance, security, agility, and reduce costs. The following table provides recommended Amazon services based on use case.

Category	Use case	Amazon services
the section called "Performance"	SAP Fiori launchpad and global access	Amazon CloudFront
the section called "Data lake"	Analytics	Amazon AppFlow, Amazon Glue, and Amazon QuickSight
the section called "Applicat ion integration"	Integration	Amazon Lambda and Amazon API Gateway
the section called "Document management"	Archiving	Amazon S3 Glacier and Amazon S3 File Gateway
the section called "Developm ent and extension"	Development	Amazon SDK for SAP ABAP
the section called "User front- end"	Front-end	Amazon WorkSpaces and Amazon AppStream 2.0

Performance

Deploy <u>Amazon CloudFront</u> in your VPC to increase performance and reduce latency of SAP Fiori launchpad in RISE with SAP. CloudFront create a cache for the static content and accelerates dynamic content through edge computing. For more information, see <u>Improving SAP Fiori Performance with Amazon CloudFront and Amazon Global Accelerator</u>. You can attach Amazon WAF to strengthen the security at edge. The following image shows this scenario.

Extensions 247

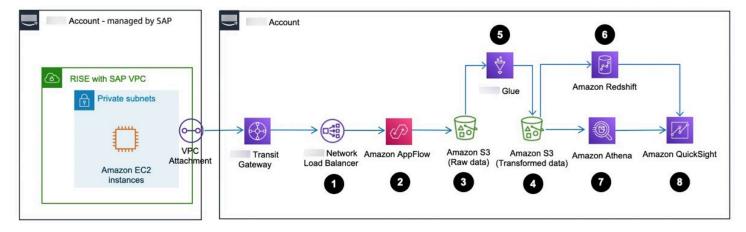


User flow

- 1. User accesses SAP Fiori launchpad via Internet browser or mobile device.
- 2. The request is routed through Amazon CloudFront.
- 3. The request is filtered by Amazon WAF to prevent passage of malicious traffic, before it is processed by Amazon CloudFront.
- 4. SAP Fiori launchpad is served from RISE with SAP VPC and presented to the user via Amazon Transit Gateway.

Data lake

Deploy <u>Amazon AppFlow</u> to extract data out of SAP S/4HANA via 0Data protocol which can also be based ODP framework. The extraction result is stored in Amazon S3 data lake. This data can be further processed with <u>Amazon Glue</u>, Amazon Redshift, and Amazon Athena. Users can consume this data with Amazon QuickSight. The following image shows this scenario.



Data flow

Data lake 248

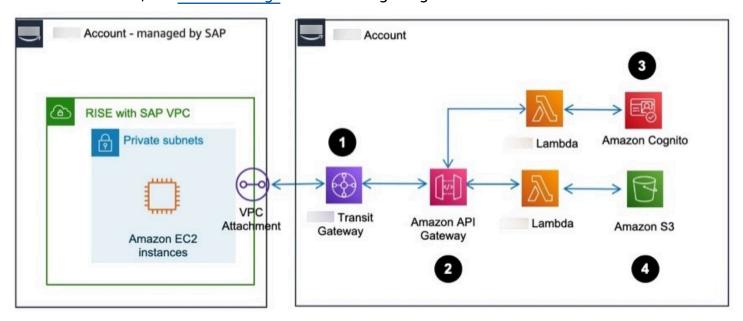
 RISE with SAP VPC is connected to your Amazon Web Services account not managed by SAP, via Amazon Transit Gateway and Network Load Balancer.

- 2. Amazon AppFlow extracts data out of SAP S/4HANA via OData protocol.
- 3. Raw data is stored in an Amazon S3 bucket.
- 4. Amazon Glue performs transformation and cleansing of data.
- 5. Transformed results are stored in another Amazon S3 bucket.
- 6. Amazon Redshift is used to further process the data through its data warehousing capability.
- 7. Amazon Athena is used to query the *transformed* data in Amazon S3.
- 8. User accesses data through Amazon QuickSight.

For more information, see Guidance for DataLake with SAP and non-SAP data on Amazon.

Application integration

Deploy <u>Amazon API Gateway</u> to extract data out of SAP S/4HANA via HTTP API. API Gateway can consume data from IDOC, BAPI, and RFC. These need to be translated to a web service call. For more information, see <u>Amazon blogs</u>. The following image shows this scenario.



Data flow

1. RISE with SAP VPC is connected to your Amazon Web Services account not managed by SAP, via Amazon Transit Gateway.

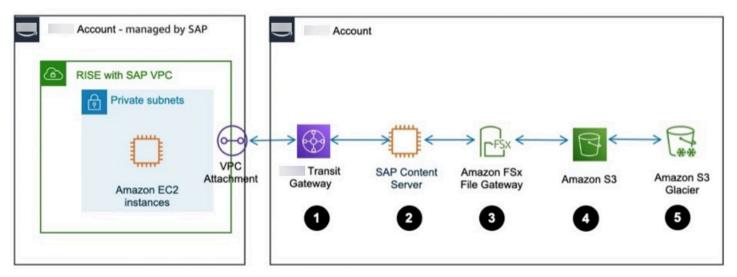
Application integration 249

2. Amazon API Gateway is configured to route the authentication to Amazon Lambda and Amazon Cognito

- 3. Amazon Cognito authenticates the session.
- 4. Once authenticated, Amazon API Gateway routes the package to Amazon Lambda.
- 5. Amazon Lambda stores the data in an Amazon S3 bucket.

Document management

Deploy an SAP Content Server that is integrated with Amazon S3, to archive SAP documents and data. The following image shows this scenario with Amazon services.



Data flow

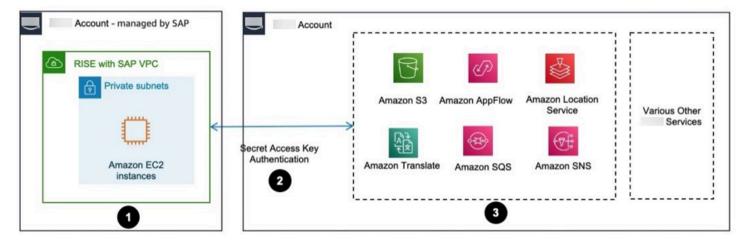
- RISE with SAP VPC is connected to your Amazon Web Services account not managed by SAP, via Amazon Transit Gateway.
- 2. SAP Content Server is installed in SAP S/4HANA as target storage for document and data archiving.
- 3. Amazon S3 File Gateway enables Amazon S3 to be mounted as NFS on SAP Content Server.
- 4. Amazon S3 bucket stores the required archive files.
- 5. You can move files to different Amazon S3 storage classes. For more information, see <u>Using</u> Amazon S3 storage classes.

Document management 250

Development and extension

Deploy Amazon SDK for SAP ABAP on RISE with SAP VPC to avail Amazon services using the ABAP language. For more information, see What is Amazon SDK for SAP ABAP?

You can authenticate Amazon SDK for SAP ABAP with IAM access key. The following image shows this scenario.



Data flow

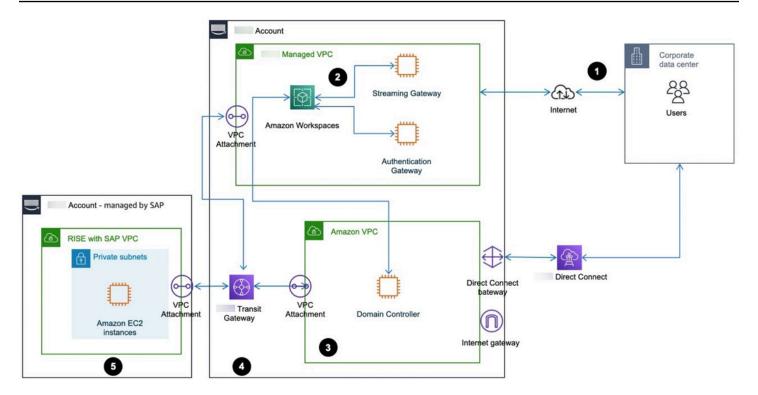
- 1. Amazon SDK for SAP ABAP is installed via a set of transports in SAP S/4HANA within RISE with SAP VPC.
- 2. SAP S/4HANA is configured with IAM access key for authenticating access to Amazon services. For more information, see Managing access keys for IAM users.
- Access to Amazon services with Amazon SDK for SAP ABAP has been established.

User front-end

Deploy <u>Amazon WorkSpaces</u> and/or <u>Amazon AppStream 2.0</u> to enhance user experience. With these Amazon services, you can improve user experience by tackling the high network latency for SAPGUI and browsers to access SAP S/4HANA within RISE with SAP VPC.

The following image shows this scenario with Amazon WorkSpaces.

Development and extension 251

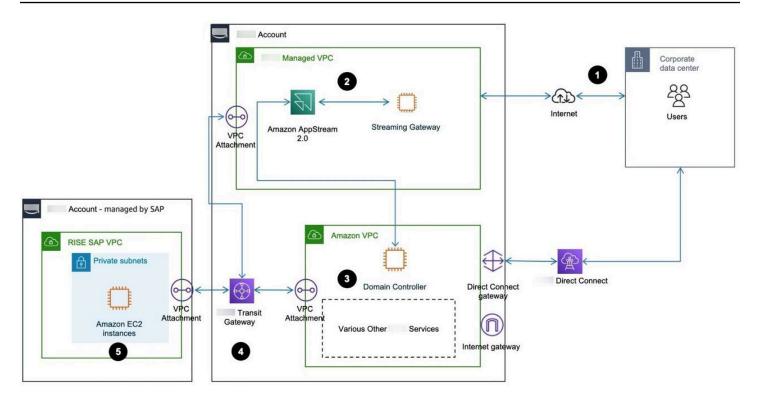


User flow

- 1. User accesses WorkSpaces via Internet.
- 2. WorkSpaces handles user access with Streaming and Authentication gateways.
- 3. WorkSpaces is integrated with Domain Controller.
- 4. User's SAPGUI request is routed through Amazon Transit Gateway.
- 5. User's request reaches SAP S/4HANA within RISE with SAP VPC.

The following image shows this scenario with Amazon AppStream 2.0.

User front-end 252



User flow

- 1. User accesses Amazon AppStream 2.0 via Internet.
- 2. AppStream 2.0 handles user access with Streaming gateway.
- 3. AppStream 2.0 is integrated with Domain Controller.
- 4. User's SAPGUI request is routed through Amazon Transit Gateway.
- 5. User's request reaches SAP S/4HANA within RISE with SAP VPC.

User front-end 253