

---

# AWS Launch Wizard

## User Guide



## **AWS Launch Wizard: User Guide**

Copyright © 2020 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

## Table of Contents

AWS Launch Wizard for SQL Server .....	1
What Is AWS Launch Wizard for SQL Server? .....	1
Supported versions .....	1
Features .....	2
Components (Windows) .....	4
Components (Linux) .....	6
Related services .....	7
How it works .....	8
Deployment options .....	11
Application launch limits .....	13
Get started .....	13
Set Up .....	13
Deploy on Windows .....	18
Deploy on Linux .....	24
Managing application resources .....	30
Monitoring .....	31
Best practices .....	32
High availability .....	32
Automatic failover .....	32
Security groups and firewalls .....	33
Troubleshoot .....	33
Active Directory objects and DNS record clean up (deployment on Windows) .....	34
Launch Wizard provisioning events .....	35
CloudWatch Logs .....	35
SSM Automation execution .....	35
AWS CloudFormation stack .....	35
Pacemaker on Ubuntu (deployment on Linux) .....	36
Errors .....	36
AWS Launch Wizard for SAP .....	38
What is AWS Launch Wizard for SAP? .....	38
Supported deployments and features .....	38
Components .....	40
Related services .....	41
How it works .....	43
Get started .....	47
Set Up .....	47
Deploy an application with Launch Wizard .....	50
Manage application resources .....	68
Troubleshoot .....	69
Launch Wizard provisioning events .....	69
CloudWatch Logs .....	69
AWS CloudFormation stack .....	69
Pre- and post-deployment configuration scripts .....	70
Application launch quotas .....	70
Errors .....	70
Make SAP HANA software available to Launch Wizard .....	70
Download SAP software .....	71
Upload SAP HANA to Amazon S3 .....	71
OS version support for SAP deployments .....	72
Security groups .....	73
Security groups .....	73
Connectivity to external systems and users .....	75
Security .....	33
Infrastructure Security .....	76

Resilience .....	76
Data Protection .....	77
Identity and Access Management .....	77
Update Management .....	78
Document History .....	79

# AWS Launch Wizard for SQL Server

This section of the AWS Launch Wizard documentation contains guidance specific to the deployment of Microsoft SQL Server applications on AWS using the Launch Wizard service. AWS Launch Wizard supports both single instance and high availability (HA) application deployments.

## Topics

- [What Is AWS Launch Wizard for SQL Server? \(p. 1\)](#)
- [Get started with AWS Launch Wizard for SQL Server \(p. 13\)](#)
- [Managing application resources with AWS Launch Wizard for SQL Server \(p. 30\)](#)
- [Monitoring SQL Server Always On deployments \(p. 31\)](#)
- [High availability and security best practices for AWS Launch Wizard for SQL Server \(p. 32\)](#)
- [Troubleshoot AWS Launch Wizard for SQL Server \(p. 33\)](#)

## What Is AWS Launch Wizard for SQL Server?

AWS Launch Wizard is a service that guides you through the sizing, configuration, and deployment of Microsoft SQL Server applications on AWS, which follow [AWS cloud application best practices](#). AWS Launch Wizard supports both single instance and high availability (HA) application deployments.

AWS Launch Wizard reduces the time it takes to deploy SQL Server solutions to the cloud. You input your application requirements, including performance, number of nodes, and connectivity on the service console, and AWS Launch Wizard identifies the right AWS resources to deploy and run your SQL Server application. AWS Launch Wizard provides an estimated cost of deployment, and gives you the ability to modify your resources and instantly view the updated cost assessment. When you approve, AWS Launch Wizard provisions and configures the selected resources in a few hours to create a fully-functioning production-ready SQL Server application. It also creates custom AWS CloudFormation templates, which can be reused and customized for subsequent deployments.

Once deployed, your SQL Server application is ready to use and can be accessed on the EC2 console. You can manage your SQL Server application with [AWS Systems Manager](#).

## Contents

- [Supported operating systems and SQL versions \(p. 1\)](#)
- [Features of AWS Launch Wizard \(p. 2\)](#)
- [Components \(deployment on Windows\) \(p. 4\)](#)
- [Components \(deployment on Linux\) \(p. 6\)](#)
- [Related services \(p. 7\)](#)
- [How AWS Launch Wizard works \(p. 8\)](#)
- [Deployment options \(p. 11\)](#)
- [Application launch limits \(p. 13\)](#)

## Supported operating systems and SQL versions

AWS Launch Wizard supports the following operating systems and SQL Server versions:

### Deployments on Windows

- Windows Server 2019/2016/2012 R2
- Enterprise and Standard Editions of Microsoft SQL Server 2019/2017/2016

### Deployments on Linux

- Ubuntu 18.04
- Enterprise and Standard Edition of Microsoft SQL Server 2019

## Features of AWS Launch Wizard

### AWS Launch Wizard provides the following features:

- [Simple application deployment \(p. 2\)](#)
- [AWS resource selection \(p. 2\)](#)
- [Cost estimation \(p. 2\)](#)
- [Reusable code templates \(p. 2\)](#)
- [SNS notification \(p. 3\)](#)
- [Always On Availability Groups \(SQL Server\) \(p. 3\)](#)
- [Dedicated Hosts \(deployment on Windows\) \(p. 3\)](#)
- [Early input validation \(p. 3\)](#)
- [Application resource groups for easy discoverability \(p. 4\)](#)
- [One-click monitoring \(p. 4\)](#)

## Simple application deployment

AWS Launch Wizard makes it easy for you to deploy third-party applications on AWS, such as Microsoft SQL Server. When you input the application requirements, AWS Launch Wizard deploys the necessary AWS resources for a production-ready application. This means that you do not have to manage separate infrastructure pieces or spend time provisioning and configuring your SQL Server application.

## AWS resource selection

Launch Wizard considers performance, memory, bandwidth, and other application features to determine the best instance type, EBS volumes, and other resources for your SQL Server application. You can modify the recommended defaults.

## Cost estimation

Launch Wizard provides a cost estimate for the complete deployment that is itemized for each individual resource being deployed. The estimated cost automatically updates each time you change a resource type configuration in the wizard. However, note that the provided estimates are only for general comparisons. They are based on On-Demand costs and actual costs may be lower.

## Reusable code templates

Launch Wizard creates a CloudFormation stack that can be reused to customize and replicate your infrastructure in multiple environments. Code in the template helps you provision resources. You can

access and use the templates created by your Launch Wizard deployment from the CloudFormation console. For more information about CloudFormation stacks, see [Working With Stacks](#).

## SNS notification

You can provide an [SNS topic](#) that allows Launch Wizard to send you notifications and alerts about the status of a deployment.

## Always On Availability Groups (SQL Server)

Always On Availability Groups (AG) is a Microsoft SQL Server feature that is supported by the AWS SQL Server installation. AG augments the availability of a set of user databases. An availability group supports a failover environment for a discrete set of user databases, known as availability databases. If one of these databases fails, another database takes over its workload with no impact on availability. Always On Availability improves database availability, enabling more efficient resource usage. For more information about the concepts and benefits of Always On Availability, see [Always On Availability Groups \(SQL Server\)](#).

## Dedicated Hosts (deployment on Windows)

You can deploy SQL Server Always On Availability Groups (AG) or basic availability groups on your Dedicated Hosts to leverage your existing SQL Server Licenses (BYOL). From the Launch Wizard console, select **Dedicated Host** tenancy, and select the Dedicated Hosts for your VPC. For more information about Amazon EC2 Dedicated Hosts, see [Dedicated Hosts](#).

## Early input validation

You can leverage your existing infrastructure (such as VPC or Active Directory) with Launch Wizard. This may lead to deployment failures if your existing infrastructure does not meet certain deployment prerequisites. For example, for a SQL Server Always On deployment in your existing VPC, the VPC must have at least one public subnet and two private subnets. It must also have outbound connectivity to Amazon S3, Systems Manager, and AWS CloudFormation service endpoints. If these requirements are not met, the deployment will fail. If you are in a later stage of a deployment, this failure can take more than an hour to detect. To detect these types of issues early in the application deployment process, Launch Wizard's validation framework verifies key application and infrastructure specifications before provisioning. Verification takes approximately 15 minutes. If necessary, you can take appropriate actions to adjust your VPC configuration.

Launch Wizard performs the following infrastructure validations:

- Resource limit validations at the AWS account level:
  - VPC
  - Internet gateway
  - Number of CloudFormation stacks
- Additionally, Launch Wizard performs the following application-specific validations:
  - Active Directory credentials (deployment on Windows)
  - Public subnet outbound connectivity
  - Private subnet outbound connectivity
  - Custom Windows AMIs:
    - SQL Server installed and running on instance
    - Compliant versions of Windows and SQL Server
  - Dedicated Hosts (deployment on Windows)
    - AMIs are filtered according to the billing code. When you select Dedicated Host tenancy in the application, the AMI selection dropdown filters out AMIs for which the usage operation is set to

include SQL Server Enterprise or SQL Server Standard, per the [details and usage operation values](#). This filtering behavior is the result of restrictions described in the [Dedicated Host restrictions](#) page.

- Supported instance type
- Sufficient capacity to launch number of nodes and instances
- Selected subnet and corresponding Dedicated Host are in the same Availability Zone for any additional nodes beyond the primary and first secondary nodes

**Note**

Some validations, for example for valid Active Directory credentials, require Application Wizard to launch a t2.large EC2 instance in your account for a few minutes. After it runs the necessary validations, Launch Wizard terminates the instance.

## Application resource groups for easy discoverability

Launch Wizard creates a resource group for all of the AWS resources created for your SQL Server application. You can manage the resources through the EC2 console or with Systems Manager. When you access Systems Manager through Launch Wizard, the resources are automatically filtered for you based on your resource group. You can manage, patch, and maintain your SQL Server applications in Systems Manager.

## One-click monitoring

Launch Wizard integrates with [CloudWatch Application Insights](#) to provide a one-click monitoring setup experience for deploying SQL Server HA workloads on AWS. When you select the option to set up monitoring and insights with Application Insights on the Launch Wizard console, Application Insights automatically sets up relevant metrics, logs, and alarms on CloudWatch, and starts monitoring newly deployed workloads. You can view automated insights and detected problems, along with the health of your SQL Server HA workloads, on the CloudWatch console.

Counters that you can configure using Application Insights include:

- Mirrored Write Transaction/sec
- Recovery Queue Length
- Transaction delay
- Windows Event Logs on CloudWatch

You can also get automated insights when a failover event or problem, such as a restricted access to query a target database, is detected on your workload.

## Components (deployment on Windows)

A SQL Server application deployed on Windows with Launch Wizard includes the following components:

- A **virtual private cloud (VPC)** configured with [public and private subnets](#) across two Availability Zones. A public subnet is a subnet whose traffic is routed to an internet gateway. If a subnet does not have a route to the internet gateway, then it is a private subnet. The VPC provides the network infrastructure for your SQL Server deployment. You can choose an optional third Availability Zone for additional SQL cluster nodes, as shown below.
- An **internet gateway** to provide access to the internet.
- In the public subnets, **Windows Server-based Remote Desktop Gateway (RDGW) instances and network address translation (NAT) gateways** for outbound internet access. If you are deploying in

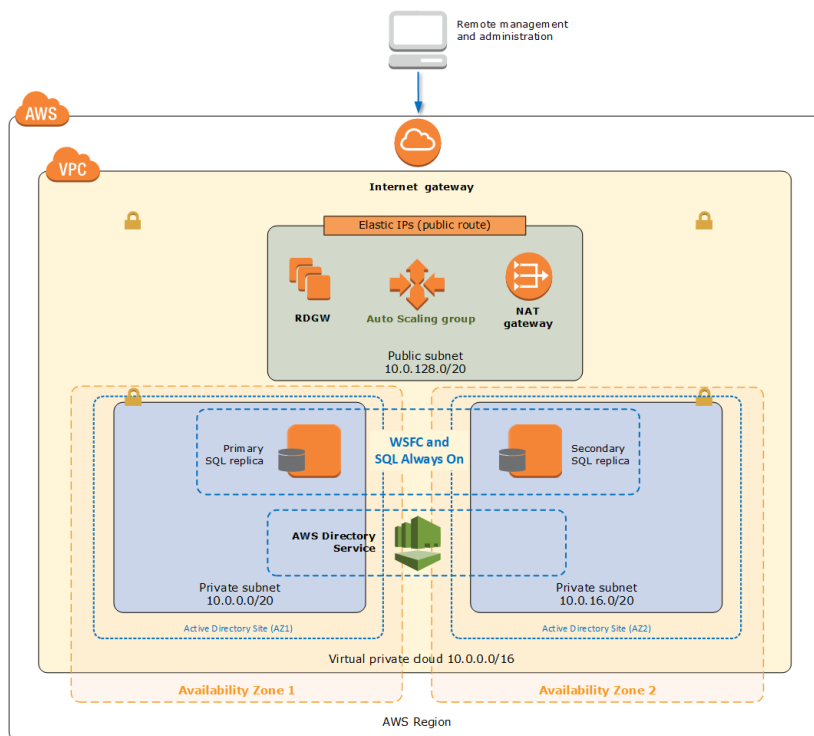


your preexisting VPC, Launch Wizard uses the existing NAT gateway in your VPC. For more information about NAT gateways, see [NAT Gateways](#).

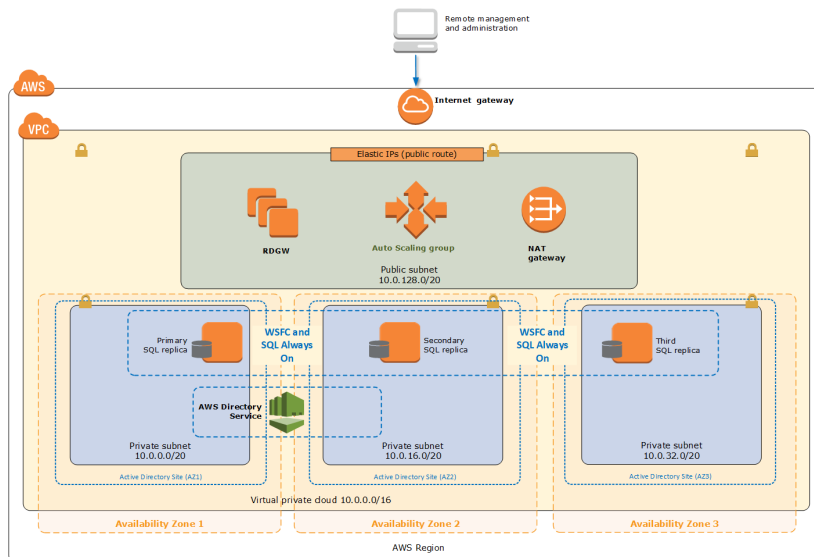
- **Elastic IP addresses** associated with the NAT gateway and RDGW instances. For more information about Elastic IP addresses, see [Elastic IP Addresses](#).
- In the private subnets, **Active Directory domain controllers**.
- In the private subnets, **Windows Server-based instances as Windows Server Failover Clustering (WSFC) nodes**. For more information, see [Windows Server Failover Clustering with SQL Server](#).
- **SQL Server Enterprise edition with SQL Server Always On Availability Groups on each WSFC node**. This architecture provides redundant databases and a witness server to ensure that a quorum can vote for the node to be promoted to the controlling resource. The default architecture mirrors an on-premises architecture of two SQL Server instances spanning two subnets placed in two different Availability Zones. For more information about SQL Server Always On Availability Groups, see [Overview of Always On Availability Groups \(SQL Server\)](#).
- **Security groups** to ensure the secure flow of traffic between the instances deployed in the VPC. For more information, see [Security Groups for Your VPC](#).

### Note

If you choose to deploy SQL Server Always On through Launch Wizard into your existing VPC, there is an additional mandatory check box on the console for you to indicate whether VPC and public/private subnet requirements have been met.



You can also choose to build an architecture with three Availability Zones, as shown in the following diagram.

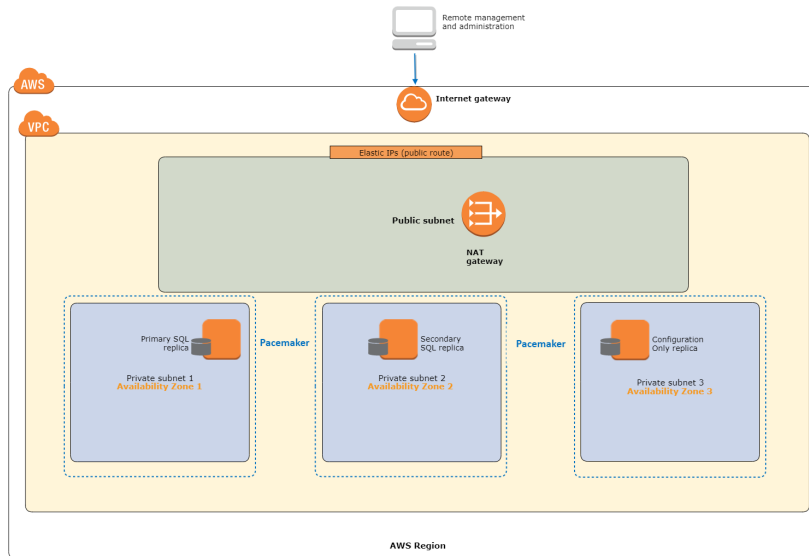


## Components (deployment on Linux)

A SQL Server application deployed on Linux with Launch Wizard includes the following components:

- A **virtual private cloud (VPC)** configured with **public and private subnets** across three Availability Zones. A public subnet is a subnet whose traffic is routed to an internet gateway. If a subnet does not have a route to the internet gateway, then it is a private subnet. The VPC provides the network infrastructure for your SQL Server deployment.
- An **internet gateway** to provide access to the internet.
- In the public subnets, **network address translation (NAT)** for outbound internet access. If you are deploying in your preexisting VPC, Launch Wizard uses the existing NAT gateway in your VPC. For more information about NAT gateways, see [NAT Gateways](#).
- Two of the private subnets each run a SQL Server **replica node**. One acts as the primary node, and the other as secondary node. The third private subnet is used to run the configuration replica. Launch Wizard deployments on Linux use [Pacemaker](#) as the cluster resource manager. Pacemaker differs from Windows Server Failover Cluster (WSFC), which is used for Windows deployments, in terms of how it handles quorum. For Always On availability groups (AG) on Linux, arbitration happens in SQL Server where the metadata is stored. This is where the configuration-only replica is relevant. In order to maintain quorum and enable automatic failovers, Launch Wizard sets up a third node that acts as the configuration-only replica.
- **Security groups** to ensure the secure flow of traffic between the instances deployed in the VPC. For more information, see [Security Groups for Your VPC](#).

The high-level architecture of a SQL Server high availability solution on Linux is similar to the architecture for deployment on Windows. The main differences are the low-level components and technologies. The architecture for Linux deployments provides redundant databases and a configuration-only replica node to verify that a quorum can vote for the node to be promoted to the controlling resource. The default architecture mirrors an on-premises architecture of two SQL Server instances spanning two subnets placed in two different Availability Zones. For more information about SQL Server Always On Availability Groups (AG), see [Overview of Always On Availability Groups \(SQL Server\)](#) in the Microsoft documentation.



## Related services

The following services are used when you deploy a SQL Server application with Amazon Launch Wizard.

- [AWS CloudFormation](#) (p. 7)
- [AWS Systems Manager](#) (p. 7)
- [Amazon Simple Notification Service \(SNS\)](#) (p. 8)
- [Amazon CloudWatch Application Insights](#) (p. 8)
- [Linux-only technologies](#) (p. 8)

## AWS CloudFormation

[AWS CloudFormation](#) is a service for modeling and setting up your AWS resources, enabling you to spend more time focusing on your applications that run in AWS. You create a template that describes all of the AWS resources that you want to use (for example, Amazon EC2 instances or Amazon RDS DB instances), and AWS CloudFormation takes care of provisioning and configuring those resources for you. With Launch Wizard, you don't have to sift through CloudFormation templates to deploy your application. Instead, Launch Wizard combines infrastructure provisioning and configuration (with a CloudFormation template) and application configuration (with code that runs on EC2 instances to configure the application) into a unified SSM Automation document. The SSM document is then invoked by Launch Wizard's backend service to provision a SQL Server application in your account. For more information, see the [AWS CloudFormation User Guide](#).

## AWS Systems Manager

[AWS Systems Manager](#) is a collection of capabilities for configuring and managing your Amazon EC2 instances, on-premises servers and virtual machines, and other AWS resources at scale. Systems Manager includes a unified interface that enables you to centralize operational data and automate tasks across your AWS resources. Systems Manager shortens the time to detect and resolve operational problems in your infrastructure. You have the option of managing your application with Systems Manager after deploying with Launch Wizard. For more information, see the [AWS Systems Manager User Guide](#).

## Amazon Simple Notification Service (SNS)

[Amazon Simple Notification Service \(SNS\)](#) is a highly available, durable, secure, fully managed pub/sub messaging service that provides topics for high-throughput, push-based, many-to-many messaging. Using Amazon SNS topics, your publisher systems can fan out messages to a large number of subscriber endpoints and send notifications to end users using mobile push, SMS, and email. You can use SNS topics for your Launch Wizard deployments to stay up-to-date on deployment progress. For more information, see the [Amazon Simple Notification Service Developer Guide](#).

## Amazon CloudWatch Application Insights

[Amazon CloudWatch Application Insights](#) facilitates observability for .NET and SQL Server applications. It can help you set up the best monitors for your application resources to continuously analyze data for signs of problems with your applications. Application Insights, which is powered by Sagemaker and other AWS technologies, provides automated dashboards that show potential problems with monitored applications, helping you to quickly isolate ongoing issues with your applications and infrastructure. The enhanced visibility into the health of your applications that Application Insights provides can help you reduce your mean time to repair (MTTR) so that you don't have to pull in multiple teams and experts to troubleshoot your application issues.

## Linux-only technologies

The following key technologies are used when you deploy a SQL Server application with Amazon Launch Wizard to the Linux platform.

- **Pacemaker** is an open source cluster resource manager (CRM), which is a system that coordinates managed resources and services made highly available by a cluster.
- **Corosync** is an open source program that provides cluster membership and messaging capabilities, often referred to as the messaging layer, to client servers. In contrast to Pacemaker, which allows you to control cluster behavior, Corosync makes it possible for servers to communicate as a cluster.
- **Transact-SQL** is an extension to the SQL language. It is used to interact with relational databases. Transact-SQL is platform-agnostic and can be used to configure the AlwaysOn Availability Group and listener.
- **Fencing** is used to isolate a malfunctioning server from the cluster in order to protect and secure the synced resources. The recommended solution to use in the case of a malfunctioning server is the "Shoot the other node in the head" (STONITH) method. STONITH is a fencing technique that isolates a failed node so that it does not disrupt a computer cluster. The STONITH method fences failed nodes by resetting or powering down the failed node. Fencing is also used when a clustered service cannot be stopped. In this case, the cluster uses fencing to force the whole node offline, which makes it safe to start the service from a different server. Fencing can be performed at two levels: the node or resource level. Launch Wizard supports only node-level fencing.

## How AWS Launch Wizard works

AWS Launch Wizard provides a complete solution to provision popular third-party applications on AWS. Currently, Launch Wizard supports Microsoft SQL Server application deployments across multiple Availability Zones or on a single instance. You provide the specifications, such as for performance, throughput, and networking requirements. Based on the application requirements that you enter, Launch Wizard automatically provisions the right AWS resources in the cloud. For example, Launch Wizard determines the best instance type and EBS volume for your CPU, memory, and bandwidth specifications, then deploys and configures them.

Launch Wizard provides an estimated cost of deployment. You can modify your resources and instantly view an updated cost assessment. Once you approve, Launch Wizard validates the inputs and flags

inconsistencies. After you resolve the inconsistencies, Launch Wizard provisions the resources and configures them. The result is a ready-to-use SQL Server Always On application.

Launch Wizard creates a CloudFormation stack according to your infrastructure needs. You can reuse this template as a baseline for future infrastructure provisioning.

For deployments on Windows, Launch Wizard supports AWS Managed Microsoft Active Directory (AD) as well as connecting to Active Directory on-premises through AWS Direct Connect.

#### Topics

- [Implementation details for deployment on Windows \(p. 9\)](#)
- [Implementation details for deployment on Linux \(p. 10\)](#)

## Implementation details for deployment on Windows

AWS Launch Wizard implements SQL Server deployments on Windows as follows.

### SQL Server Enterprise Edition

Launch Wizard supports installation of SQL Server Enterprise and Standard Editions of 2016 and 2017 on Windows Server 2012 R2, 2016, and 2019 through License Included Amazon Machine Images (AMIs). Launch Wizard allows you to bring your own SQL licenses through a custom AMI. If you use a custom AMI, ensure that your AMI meets the requirements listed in [Requirements for using custom AMIs \(deployment on Windows\) \(p. 15\)](#).

### Storage on WSFC nodes

Storage capacity and performance are key aspects of any production SQL Server installation. Launch Wizard lets you choose capacity and performance based on your deployment needs.

Amazon Elastic Block Store (Amazon EBS) volumes are included in the architecture to provide durable, high-performance storage. EBS volumes are network-attached disk storage, which you can create and attach to EC2 instances. When attached, you can create a file system on top of these volumes, run a database, or use them in any way that you would use a block device. EBS volumes are placed in a specific Availability Zone, where they are automatically replicated to protect you from the failure of a single component. EBS volume type `io1` is not supported.

[Provisioned IOPS EBS volumes](#) offer storage with consistent and low-latency performance. They are backed by solid state drives (SSDs) and designed for applications with I/O intensive workloads, such as databases. Amazon EBS-optimized instances, such as the R4 instance type, deliver dedicated throughput between Amazon EC2 and Amazon EBS.

By default, Launch Wizard deploys three 500 GiB general purpose SSD volumes to store databases, logs, tempdb, and backups on each WSFC node. These general purpose SSD volumes are in addition to the root general purpose SSD volume used by the operating system, which delivers a consistent baseline of 3 IOPS/GiB and provides a total of 1,500 IOPS per volume for SQL Server database and log volumes. You can customize the volume size and switch to using dedicated IOPS volumes with the volume you specify. If you need more IOPS per volume, consider using Provisioned IOPS SSD volumes by changing the SQL Server Volume Type and SQL Server Volume IOPS parameters.

The default disk layout for SQL Server deployed by Launch Wizard is:

- One general purpose SSD volume (100 GiB) for the operating system (C:)
- One general purpose SSD volume (500 GiB) to host the SQL Server database files (D:)
- One general purpose SSD volume (500 GiB) to host the SQL Server log files (E:)
- One general purpose SSD volume (500 GiB) to host the SQL Server tempdb and backup files (F:)

## IP Addressing on the Windows Server Failover Clustering (WSFC) Nodes

In order to support WSFC and Always On Availability Group listeners, each node that hosts the SQL Server instances that participate in the cluster must have three IP addresses assigned, as follows:

- One IP address as the primary IP address for the instance
- A second IP address as the WSFC IP resource
- A third IP address to host the Always On Availability Group listener

When you launch the AWS CloudFormation template, you can specify the addresses for each node. By default, the underlying CloudFormation templates of Launch Wizard use 10.0.0.0/20, 10.0.16.0/20, and 10.0.32.0/20 as CIDR blocks for the private subnets. This is true only when you use Launch Wizard to deploy SQL Server Always On clusters in a new VPC.

## Windows Server Failover Clustering (WSFC)

You can build the failover cluster after your Windows Server instances have been deployed and domain-joined. Launch Wizard's underlying AWS CloudFormation templates build the cluster when deploying the second node. If you use the default template parameter settings, Launch Wizard executes the following Windows PowerShell commands to complete this task.

```
PS C:\> Install-WindowsFeature failover-clustering -IncludeManagementTools
```

```
PS C:\> New-Cluster -Name WSFCclusterName -Node $nodes -StaticAddress $addr
```

The first command runs on each instance during the bootstrapping process. It installs the required components and management tools for the failover clustering services. The second command runs near the end of the bootstrapping process on the second node and is responsible for creating the cluster and for defining the server nodes and IP addresses.

If you set the optional third Availability Zone, Launch Wizard keeps the quorum settings to the default node majority.

```
PS C:\> Set-ClusterQuorum -NodeMajority
```

## Always On configuration

After SQL Server Enterprise edition has been installed and the Windows Server failover cluster has been built, Launch Wizard enables SQL Server Always On with the following PowerShell command.

```
PS C:\> Enable-SqlAlwaysOn -ServerInstance $ServerInstance
```

Launch Wizard runs this command on each node, and the proper server name is provided as a value for the `ServerInstance` parameter.

When the deployment is complete, Launch Wizard creates your databases and make them highly available by creating an Always On Availability Group.

When you create an availability group, you provide a network share that is used to perform an initial data synchronization. As you progress through the New Availability Group wizard, a full backup for each selected database is taken and placed in the share. The secondary node connects to the share and restores the database backups before joining the availability group.

## Implementation details for deployment on Linux

AWS Launch Wizard implements SQL Server deployments on Linux as follows.

## Always On availability group configuration

- **Availability group enablement.** Launch Wizard enables availability groups using the following commands.

```
sudo /opt/mssql/bin/mssql-conf set hadr.hadrenabled 1
```

```
sudo systemctl restart mssql-server
```

- **Authentication of TCP endpoints.** An availability group uses TCP endpoints for communication. Launch Wizard uses certificate-based authentication to support endpoints for availability groups. All nodes create a self-signed certificate and upload their certificates to an Amazon S3 location that you specify. Each node will then locally sync the certificates from all of the other nodes.
- **Creation of availability group endpoints.** Launch Wizard performs the backup, create, and restore of the authentication certificates with Transact-SQL on Linux. Transact-SQL on Linux is also used to create the availability group endpoints and availability groups with automatic seeding. The listener and read-only routing are configured after an availability group is created.
- **Creation of availability group resources in Pacemaker cluster.** After an availability group is created in SQL Server, the corresponding resources must be created in Pacemaker. There are two resources associated with an availability group: the availability group and a floating IP address. The created availability group resource is a unique resource called a clone. The availability group resource contains copies on each node, with one controlling resource. The controlling resource is associated with the server that hosts the primary replicas. The secondary replicas (except for configuration-only replicas) can be promoted to controlling resource during failover.

The floating IP address is an unused IP address in the VPC that is not part of any of the subnet CIDR ranges that are part of this cluster. Launch Wizard creates a DNS entry in the host files to map `AGListenerName` to a floating IP Address. Launch Wizard also creates a route for Floating IP/32 to the Primary Node Network Interface so that all internal VPC network traffic is routed to the primary node when any traffic reaches the floating IP. In case of failover, the resource agent will update the routing table to point the floating IP to the new primary node.

- **Creation of availability group listener.** The creation of the availability group listener is performed using the following Transact-SQL commands:

```
ALTER AVAILABILITY GROUP MyAg2
```

```
ADD LISTENER ...
```

## Deployment options

**AWS Launch Wizard provides the following deployment paths:**

- [Deployment on Windows \(p. 11\)](#)
- [Deployment on Linux \(p. 12\)](#)

### Deployment on Windows

1. **Deploy SQL Server into a new VPC across multiple Availability Zones.** When you choose this configuration option, Launch Wizard builds a new AWS environment consisting of the VPC, subnets, NAT gateways, security groups, domain controllers, and other infrastructure components. It then deploys Windows Server Failover Clustering (WSFC) with SQL Server across multiple Availability Zones into this new VPC.



2. **Deploy SQL Server into a new VPC on a single node.** When you choose this configuration option, Launch Wizard builds a new AWS environment consisting of the VPC, subnet, NAT gateway, security groups, domain controllers, and other infrastructure components. It then deploys SQL Server on a single node into this new VPC.
3. **Deploy SQL Server into an existing VPC and create a new AWS Managed Active Directory.** When you choose this configuration option, Launch Wizard builds a new AWS environment that consists of security groups, domain controllers, and other infrastructure components, and then deploys WSFC with SQL Server in to the specified VPC and subnets. Your AWS environment must include a VPC with two or three Availability Zones, private subnets in each Availability Zone, and at least one public subnet in the VPC. Currently, Launch Wizard only supports AWS Managed Microsoft Active Directory for this scenario.
4. **Deploy a SQL Server into an existing VPC with an existing AWS Managed Active Directory across multiple Availability Zones.** When you choose this configuration option, Launch Wizard provisions WSFC across multiple Availability Zones in your existing AWS infrastructure. Your AWS environment must include a VPC with two or three Availability Zones, private subnets in each Availability Zone, at least one public subnet in the VPC, and an AWS Active Directory in the VPC (this is the Active Directory on which you deploy your SQL nodes).
5. **Deploy a SQL Server into an existing VPC with an existing AWS Managed Active Directory on a single node.** When you choose this configuration option, Launch Wizard provisions SQL Server on a single node in your existing AWS infrastructure. Your AWS environment must include a VPC in one Availability Zone, a private subnet, a public subnet in the VPC, and an AWS Active Directory in the VPC (this is the Active Directory on which you deploy your SQL nodes).
6. **Deploy a SQL Server into an existing VPC across multiple Availability Zones and connect to an on-premises Active Directory.** When you choose this configuration option, Launch Wizard provisions WSFC across multiple Availability Zones in your existing AWS infrastructure. Your AWS environment must include a VPC with two or three Availability Zones, private subnets in each Availability Zone, at least one public subnet in the VPC, and an [AWS Direct Connect](#) connection to your on-premises Active Directory.
7. **Deploy a SQL Server into an existing VPC on a single node and connect to an on-premises Active Directory.** When you choose this configuration option, Launch Wizard provisions SQL Server on a single node in your existing AWS infrastructure. Your AWS environment must include a VPC in one Availability Zone, a private subnet, a public subnet in the VPC, and an [AWS Direct Connect](#) connection to your on-premises Active Directory.
8. **Deploy SQL HA on Dedicated Hosts with your Windows BYOL- or SQL Server BYOL-licensed AMIs.** When you choose this configuration option, Launch Wizard provisions SQL Server Always On Availability Groups on your existing Dedicated Hosts using your existing SQL Server licenses (BYOL) for deploying SQL Server High Availability solutions.

Launch Wizard allows you to configure additional settings, such as the version of SQL Server (by your choice of AMI), in addition to instance types and Amazon EBS volume types based on the infrastructure requirements that you specify.

## Deployment on Linux

1. **Deploy SQL Server into a new VPC across multiple Availability Zones.** When you choose this configuration option, Launch Wizard builds a new AWS environment consisting of the VPC, subnets, NAT gateways, security groups, and other infrastructure components. It then deploys SQL Server AG across multiple Availability Zones into the VPC using Pacemaker and fencing agents for cluster management.
2. **Deploy SQL Server into a new VPC on a single node.** When you choose this configuration option, Launch Wizard builds a new AWS environment consisting of the VPC, subnet, NAT gateways, security groups, and other infrastructure components. It then deploys SQL Server AG on a single node into the VPC using Pacemaker and fencing agents for cluster management.



3. **Deploy SQL Server into an existing VPC across multiple Availability Zones.** When you choose this configuration option, Launch Wizard builds a new AWS environment that consists of security groups and other infrastructure components, and then deploys SQL Server AG across multiple Availability Zones in to the specified VPC and subnets. Your AWS environment must include a VPC with three Availability Zones, private subnets in each Availability Zone, and at least one public subnet in the VPC.
4. **Deploy SQL Server into an existing VPC on a single node.** When you choose this configuration option, Launch Wizard builds a new AWS environment that consists of security groups and other infrastructure components, and then deploys SQL Server AG on a single node in to the specified VPC and subnet. Your AWS environment must include a VPC with one Availability Zone, a private subnet, and a public subnet in the VPC.

## Application launch limits

Launch Wizard allows for a maximum of 50 active applications (with status in progress or completed) for any given application type. If you want to increase this limit, contact [AWS Support](#). Launch Wizard supports 3 parallel in-progress deployment per account.

# Get started with AWS Launch Wizard for SQL Server

This section contains information you need to set up your environment for Launch Wizard, including:

- Active Directory permissions
- How to create an IAM policy and attach it to your IAM user identity
- OS and SQL version requirements
- Configuration settings

When your environment is set up, you can deploy a SQL Server Always On application with Launch Wizard by following the [steps and parameter specification details \(p. 18\)](#) provided in this section.

### Contents

- [Set up for AWS Launch Wizard for SQL Server \(p. 13\)](#)
- [Access and deploy an application with AWS Launch Wizard for SQL Server on Windows \(p. 18\)](#)
- [Access and deploy an application with AWS Launch Wizard for SQL Server on Linux \(p. 24\)](#)

## Set up for AWS Launch Wizard for SQL Server

The following prerequisites must be verified in order to deploy a SQL Server Always On application with AWS Launch Wizard.

### Contents

- [Active Directory \(Windows deployment\) \(p. 14\)](#)
- [AWS Identity and Access Management \(IAM\) \(p. 15\)](#)
- [Requirements for using custom AMIs \(deployment on Windows\) \(p. 15\)](#)
- [Requirements for using custom AMIs \(deployment on Linux\) \(p. 16\)](#)
- [Configuration settings \(deployment on Windows\) \(p. 17\)](#)

## Active Directory (Windows deployment)

### AWS Managed Active Directory

If you are [deploying SQL Server into an existing VPC with an existing Active Directory](#), Launch Wizard uses your Managed Active Directory (AD) domain user credentials to set up a fully functional SQL Server Always On Availability Group in the Active Directory. Currently, Launch Wizard only supports this deployment option for an AWS Managed Active Directory. Your Managed Active Directory does not have to be in the same VPC as the one in which SQL Server Always On is deployed. If it is in a different VPC than the one in which SQL Server Always On is deployed, ensure that you set up connectivity between the two VPCs. The domain user requires the following permissions in the [Active Directory Default organizational unit \(OU\)](#) to enable Launch Wizard to perform the deployment successfully.

- Join machines to the domain
- Create user accounts
- Create computer objects
- Read all properties
- Modify permissions

The following key operations are performed against your Active Directory by Launch Wizard. These operations result in the creation of new records or entries in Active Directory.

- SQL Server service user added as a new Active Directory user if it does not already exist in Active Directory.
- SQL Server instance and Remote Desktop Gateway Access instance joined to the Active Directory domain.
- CreateChild role added to Windows Server Failover Cluster as part of ActiveDirectoryAccessRule.
- FullControl role added to SQL Server Service user as part of FileSystemRights.

### On-premises Active Directory through AWS Direct Connect

If you are [deploying SQL Server into an existing VPC and connecting to an on-premises Active Directory](#), ensure the following prerequisites.

- Make sure you have connectivity between your AWS account and your on-premises network. You can establish a dedicated network connection from your on-premises network to your AWS account with AWS Direct Connect. For more information, see [the AWS Direct Connect documentation](#).
- The domain functional level of your Active Directory domain controller must be Windows Server 2012 or later.
- The IP addresses of your DNS server must be either in the same VPC CIDR range as the one in which your Launch Wizard SQL Server Always On deployment will be created, or in the private IP address range.
- The firewall on the Active Directory domain controllers should allow the connections from the Amazon VPC from which you will create the Launch Wizard deployment. At a minimum, your configuration should include the ports mentioned in [How to configure a firewall for Active Directory domains and trusts](#).

You can optionally perform the following step.

- Establish DNS resolution across your environments. For options on how to set this up, see [How to Set Up DNS Resolution Between On-Premises Networks and AWS Using AWS Directory Service and Amazon Route 53](#) or [How to Set Up DNS Resolution Between On-Premises Networks and AWS Using AWS Directory Service and Microsoft Active Directory](#).

## AWS Identity and Access Management (IAM)

The following steps for establishing the AWS Identity and Access Management (IAM) role and setting up the IAM user for permissions are typically performed by an IAM Administrator for your organization.

### One-time creation of IAM Role

On the **Choose Application** page of Launch Wizard, under **Permissions**, Launch Wizard displays the IAM role required for the Amazon EC2 instances that have been created by Launch Wizard to access other AWS services on your behalf. When you select **Next**, Launch Wizard attempts to discover the IAM role in your account. If the role exists, the role is attached to the instance profile for the EC2 instances that Launch Wizard will launch into your account. If the role does not exist, Launch Wizard attempts to create the role with the same name, `AmazonEC2RoleForLaunchWizard`. This role is comprised of two IAM managed policies: `AmazonSSMManagedInstanceCore` and `AmazonEC2RolePolicyForLaunchWizard`. After the role is created, the IAM Administrator can delegate the application deployment process to another IAM user who, in turn, must have the Launch Wizard IAM managed policy described in the following section.

### IAM user setup

To deploy a SQL Server Always On application with Launch Wizard, you must create an [Identity and Access Management \(IAM\) policy](#) and attach it to your IAM user identity. The IAM policy defines the user permissions. If you do not already have an IAM user in your account, follow the steps listed in [Create an IAM User in Your AWS Account](#).

When you have an IAM user in your account, create an IAM policy.

1. Go to the IAM console at <https://console.aws.amazon.com/iam/>. In the left navigation pane, choose **Policies**.
2. Choose **Users** from the left navigation pane.
3. Select the **User name** of the user to which you want to attach the policy.
4. Select **Add permissions**.
5. Select **Attach existing policies directly**.
6. Search for the policy named **AmazonLaunchWizard\_Fullaccess** and select the check box to the left of the policy name.
7. Select **Next: Review**.
8. Verify that the correct policy is listed, and then select **Add permissions**.

#### **Important**

Make sure that you log in with the user associated with the above policy when you use Launch Wizard.

## Requirements for using custom AMIs (deployment on Windows)

We recommend that you use Amazon Windows license-included AMIs whenever possible. There are occasions when you may want to use a custom Windows AMI. For example, you may have existing licenses (BYOL) or you may have made changes to one of our public images and re-imaged it.

If you use Amazon Windows license-included AMIs, you are not required to perform any pre-checks on the AMI to ensure that it meets Launch Wizard requirements.

Launch Wizard relies on user data to begin the process of configuring SQL Server or RGW instances the service launches in your accounts. For more information, see [User Data Scripts](#). By default, all AWS Windows AMIs have user data execution enabled for the initial launch. To ensure that your custom AMIs are set up to run the User Data script at launch, follow the AWS recommended method to prepare your AMIs using either EC2Launch (Windows Server 2016 and later) or the EC2Config service (Windows

2012 R2 and earlier). For more information about how to prepare your custom AMI using the options to Shutdown with Sysprep or Shutdown without Sysprep, see [Create a Standard Amazon Machine Image Using Sysprep](#). For Windows Server 2016 and later, see [Using Sysprep with EC2Launch](#). If you want to directly enable user data as part of the custom AMI creation process, follow the steps for Subsequent Reboots or Starts under [Running Commands on Your Windows Instance at Launch](#).

If you use a custom AMI, the volume drive letter for the root partition should be c :, because EC2Launch and EC2Config rely on this configuration to install the components.

While not exhaustive, the following requirements cover most of the configurations whose alteration might impact the successful deployment of a SQL Server Always On application using Launch Wizard.

SQL Server Version	Windows Server 2012 R2	Windows Server 2016	Windows Server 2019
SQL Server 2016	YES	YES	YES
SQL Server 2017	YES	YES	YES
SQL Server 2019	Currently not supported.	YES	YES

### OS and SQL requirements

- Microsoft Windows Server 2012 R2 (Datacenter) (64-bit only)
- Microsoft Windows Server 2016 (Datacenter) (64-bit only)
- Microsoft Windows Server 2019 (Datacenter) (64-bit only)
- MBR-partitioned volumes and GUID Partition Table (GPT) partitioned volumes that are formatted using the NTFS file system
- English language pack only
- SQL Server Enterprise Edition 2017/2016 or Standard Edition 2017/2016
- The root volume drive for the custom AMI should be c :
- SQL Server is installed in the root drive

### AWS software and drivers

- EC2Config service (Windows Server 2012 R2)
- EC2Launch (Windows Server 2016)
- AWS Systems Manager ([SSM agent must be installed](#))
- AWS Tools for Windows PowerShell
- Network drivers (SRIOV, ENA)
- Storage drivers (NVMe, AWS PV)

## Requirements for using custom AMIs (deployment on Linux)

There are occasions when you may want to use a custom Linux AMI. For example, you may have existing licenses (BYOL) or you may have made changes to one of our public images and re-imaged it.

If you use a custom AMI, you must adhere to the following requirements.

- The operating system must be Ubuntu version 18.04 LTS.
- The system installer and administrator must be a sudo user and be able to log in to the cluster nodes using SSH.

- SQL Server for Linux must be a default installation.
- The SQL Server for Linux version must be 2019.
- The latest Microsoft SQL tools must be installed.

## Configuration settings (deployment on Windows)

The following configuration settings are applied when deploying a SQL Server Always On application with Launch Wizard.

Setting	Applies to
Current EC2Config and SSM Agent	Windows Server 2012 R2
Current EC2Launch and SSM Agent	Windows Server 2016 and 2019
Current AWS PV, ENA, and NVMe drivers	Windows Server 2012 R2, 2016, and 2019
Current SRIOV drivers	Windows Server 2012 R2, 2016, and 2019
Microsoft SQL Server: Latest service pack SQL Service configured to start automatically SQL Service running BUILTIN\Administrators added to the SysAdmin server role TCP port 1433 and UDP port 1434 open	Windows Server 2012 R2, 2016, and 2019
Allow ICMP traffic through the firewall	Windows Server 2012 R2, 2016, and 2019
Allow RDP traffic through host firewall	Windows Server 2012 R2, 2016, and 2019
Enable file and printer sharing	Windows Server 2012 R2
RealTimeIsUniversal registry key set	Windows Server 2012 R2, 2016, and 2019

### The following AMI settings can impact the Launch Wizard deployment:

#### System Time

**RealTimeIsUniversal.** If disabled, Windows system time drifts when the time zone is set to a value other than UTC.

#### Windows Firewall

In most cases, Launch Wizard configures the correct protocols and ports. However, custom Windows Firewall rules could impact the cluster service. To ensure that your custom AMI works with Launch Wizard, see [Service overview and network port requirements for Windows](#).

#### Remote Desktop

**Service Start.** Remote Desktop service must be enabled.

**Remote Desktop Connections.** Must be enabled.

### EC2Config (Server 2012 R2)

**Installation.** We recommend using the latest version of EC2Config.

**Service Start.** EC2Config service should be enabled.

### Network Interface

**DHCP Service Startup.** DHCP service should be enabled.

**DHCP on Ethernet.** DHCP should be enabled.

### Microsoft SQL Server

**TCPIP.** Must be enabled for protocols in SQL Configuration Manager.

### PowerShell

**Execution Policy.** The execution policy in all AWS license-included AMIs is set to `Unrestricted`. We recommend that you set this policy to `Unrestricted` when you set up SQL Server Always On Availability Groups using Launch Wizard. You can change the policy when setup is complete.

## Access and deploy an application with AWS Launch Wizard for SQL Server on Windows

### Access AWS Launch Wizard

You can launch AWS Launch Wizard from the following locations.

- **AWS Console.** From the [AWS Management Console](#) under **Management and Governance**.
- **AMI launch page in EC2 console.** From the Launch Wizard banner that appears when you select **AMIs**, under **Images**, in the EC2 console. The banner appears when you search for SQL AMIs.
- **AWS Launch Wizard landing page.** From the AWS Launch Wizard page, located at <https://aws.amazon.com/launchwizard/>.

### Deploy AWS Launch Wizard on Windows

The following steps guide you through a SQL Server Always On application deployment with AWS Launch Wizard after you have launched it from the console.

1. When you select **Create deployment** from the AWS Launch Wizard landing page, you are directed to the **Choose application** page where you are prompted to select the type of application that you want to deploy. Select **Microsoft SQL Server Always On**.
2. Select the operating system on which you want to install the SQL Server Always On application — in this case, **Windows**.
3. After you select the operating system, under **Permissions**, Launch Wizard displays the AWS Identity and Access Management (IAM) policy required for Launch Wizard to access other AWS services on your behalf. For more information about setting up IAM for Launch Wizard, see [AWS Identity and Access Management \(IAM\) \(p. 15\)](#). Select **Next**.
4. After selecting the type of application to deploy, you are prompted to enter specifications for the new deployment on the **Configure application settings** page. The following tabs provide information about the specification fields.

## General

- **Deployment model.** Choose **High availability deployment** to deploy your SQL Server Always On application across multiple Availability Zones or **Single instance deployment** to deploy your SQL Server application on a single node.
- **Deployment name.** Enter a unique application name for your deployment.
- **Simple Notification Service (SNS) topic ARN — optional.** Specify an SNS topic where AWS Launch Wizard can send notifications and alerts. For more information, see the [Amazon Simple Notification Service Developer Guide](#).
- **CloudWatch application monitoring (Optional for HA deployments).** Select the checkbox to set up monitors and automated insights for your deployment using CloudWatch Application Insights. For more information, see the [Amazon CloudWatch User Guide](#).

## Connectivity

Enter specifications for how you want to connect to your instance and configure your Virtual Private Cloud (VPC).

### Key pair name

- Select an existing key pair from the dropdown list or create a new one. If you select **Create new key pair name** to create a new key pair, you are directed to the Amazon EC2 console. From there, under **Network and Security**, choose **Key Pairs**. Choose **Create a new key pair**, enter a name for the key pair, and then choose **Download Key Pair**.

#### Important

This is the only chance for you to save the private key file, so be sure to download it and save it in a safe place. You must provide the name of your key pair when you launch an instance, and provide the corresponding private key each time that you connect to the instance.

Return to the Launch Wizard console and choose the refresh button next to the **Key Pairs** dropdown list. The newly created key pair appears in the dropdown list. For more information about key pairs, see [Amazon EC2 Key Pairs and Windows Instances](#).

### Tenancy (HA deployments only)

Select your preferred tenancy. Each instance that you launch into a VPC has a tenancy attribute. The **Shared** tenancy option means that the instance runs on shared hardware. The **Dedicated Host (HA deployments)** tenancy option means that the instance runs on a Dedicated Host, which is an isolated server with configurations that you can control. For more information, see [Dedicated Hosts](#).

**Virtual Private Cloud (VPC).** Choose whether you want to use an existing VPC or create a new VPC.

- **Select Virtual Private Cloud (VPC) option.** Choose the VPC that you want to use from the dropdown list. If you choose to enable Remote Desktop Gateway access on single-node deployments, then your VPC must include one public subnet and one private subnet. It must include at least two private subnets for HA deployments. Your VPC must be associated with a [DHCP Options Set](#) to enable DNS translations to work. The private subnets must have outbound connectivity to the internet and other AWS services (S3, CFN, SSM, Logs). We recommend that you enable this connectivity with a NAT Gateway. For more information about NAT Gateways, see [NAT Gateways](#) in the Amazon VPC User Guide.

- **Remote Desktop Gateway preferences (single-node deployments only).** When you select **Yes**, enter the public subnet into which to deploy the RDGW instance, and the Remote Desktop Gateway CIDR block.
- **Public Subnet.** If you choose to enable Remote Desktop Gateway access on single-node deployments, then your VPC must include one public subnet and one private subnet. It must include at least two private subnets for HA deployments. Choose a public subnet for your VPC from the dropdown list. To continue, you must select the check box that indicates that the public subnet has been set up and the private subnets have outbound connectivity enabled.

### To add a new public subnet

If a subnet's traffic is routed to an internet gateway, the subnet is known as a public subnet. If, however, a subnet doesn't have a route to the internet gateway, the subnet is known as a private subnet. To use an existing VPC that does not have a public subnet, you can add a new public subnet using the following steps.

- Follow the steps in [Creating a Subnet in the Amazon VPC User Guide](#) using the existing VPC you intend to use AWS Launch Wizard.
- To add an internet gateway to your VPC, follow the steps in [Attaching an Internet Gateway](#) in the Amazon VPC User Guide.
- To configure your subnets to route internet traffic through the internet gateway, follow the steps in [Creating a Custom Route Table](#) in the Amazon VPC User Guide. Use IPv4 format (0.0.0.0/0) for Destination.
- The public subnet should have the "auto-assign public IPv4 address" setting enabled. To enable this setting, follow the steps in [Modifying the Public IPv4 Addressing Attribute for Your Subnet](#) in the Amazon VPC User Guide.
- **Availability Zone (AZ) configuration.** You must choose at least two Availability Zones for High Availability (HA) deployments and one Availability Zone for single-node deployments, with one private subnet for each zone that you select. For HA deployments, select the **Availability Zones** within which you want to deploy your **primary** and **secondary** SQL nodes. Depending on the number of secondary nodes that you plan to use to set up a SQL Server Always On deployment, you may have to specify a **private subnet** for each of them. Cross-Region replication is not supported.

### To create a private subnet

If a subnet doesn't have a route to an internet gateway, the subnet is known as a private subnet. To create a private subnet, you can use the following steps. We recommend that you enable the outbound connectivity for each of your selected private subnets using a NAT Gateway. To enable outbound connectivity from private subnets with public subnet, see the steps in [Creating a NAT Gateway](#) to create a NAT Gateway in your chosen public subnet. Then, follow the steps in [Updating Your Route Table](#) for each of your chosen private subnets.

- Follow the steps in [Creating a Subnet](#) in the Amazon VPC User Guide using the existing VPC you will use in AWS Launch Wizard.
- When you create a VPC, it includes a main route table by default. On the **Route Tables** page in the Amazon VPC console, you can view the main route table for a VPC by looking for **Yes** in the **Main** column. The main route table controls the routing for all subnets that are not explicitly associated with any other route table. If the main route table for your VPC has an outbound route to an internet gateway, then any subnet created using the previous step, by default, becomes a public subnet. To ensure the subnets are private, you may need to create separate route table(s) for your private subnets. These route tables must not contain any routes to an internet gateway. Alternatively, you can create a custom route table for your public subnet and remove the internet gateway entry from the main route table.



If you selected **Dedicated host** tenancy, you must select a Dedicated Host for each Availability Zone. If you have not allocated any dedicated hosts to your account, you can choose **Create new dedicated host** to do so from the EC2 console.

- **Remote Desktop Gateway access — Optional.** Select **Custom IP** from the dropdown list. Enter the CIDR block. If you do not specify any value for the Custom IP parameter, Launch Wizard does not set the inbound RDP access (Port 3389) from any IP. You can choose to do this later by modifying the security group settings via the Amazon EC2 console. See [Adding a Rule for Inbound RDP Traffic to a Windows Instance](#) for instructions on adding a rule that allows inbound RDP traffic to your RDGW instance.
- **Create new Virtual Private Cloud (VPC) option.** Launch Wizard creates your VPC. You can optionally enter a VPC name tag. If you selected **Dedicated Host** tenancy for high availability deployments, select a primary and secondary Dedicated Host. If you haven't allocated any Dedicated Hosts to your account, select **Create a new dedicated host** to be directed to the EC2 console to do so.
- **Remote Desktop Gateway preferences.** When you select **Yes**, only the Remote Desktop Gateway access information will be taken from the VPC.
- **Remote Desktop Gateway access — Optional.** Select **Custom IP** from the dropdown list. Enter the CIDR block. If you do not specify any value for the Custom IP parameter, Launch Wizard does not set the inbound RDP access (Port 3389) from any IP. You can choose to do this later by modifying the security group settings via the Amazon EC2 Console. See [Adding a Rule for Inbound RDP Traffic to a Windows Instance](#) for instructions on adding a rule that allows inbound RDP traffic to your RDGW instance.

## Active Directory

You can connect to an existing Active Directory or, for high availability deployments, you can create a new one. If you selected the **Create new Virtual Private Cloud (VPC)** option for high availability deployments, you must select **Create a new Active Directory**.

### Connecting to existing AWS Managed Active Directory

Follow the steps for granting permissions in the Managed Active Directory Default Organizational Unit (OU).

- **Domain user name and password.** Enter the user name and password for your directory. For required permissions for the domain user, see [Active Directory \(Windows deployment\) \(p. 14\)](#). Launch Wizard stores the password in the Systems Manager Parameter Store of your account as a secure string parameter. It does not store the password on the service side. To create a functional SQL Server Always On deployment, it reads from the Parameter Store.
- **DNS address.** Enter the IP address of the DNS servers to which you are connecting. These servers must be reachable from within the VPC that you selected.
- **Optional DNS address.** If you would like to use a backup DNS server, enter the IP address of the DNS server that you want to use as backup. These servers must be reachable from within the VPC that you selected.
- **Domain DNS name.** Enter the Fully Qualified Domain Name (FQDN) of the [forest root domain](#) used for the Active Directory. When you choose to create a new Active Directory, Launch Wizard creates a domain admin user on your Active Directory.

## Creating a new AWS Managed Active Directory through Launch Wizard

- **Domain user name and password.** The domain user name is preset to “admin.” Enter a password for your directory. Launch Wizard stores the password in the Systems Manager Parameter Store of your account as a secure string parameter. It does not store the password on the server side. To create a functional SQL Server Always On deployment, it reads from the Parameter Store.
- **Domain DNS name.** Enter a Fully Qualified Domain Name (FQDN) of the forest root domain used for the Active Directory. When you choose to create a new Active Directory, Launch Wizard creates a domain admin user on your Active Directory.

## Creating an on-premises Active Directory through Launch Wizard

Launch Wizard allows you to connect to your on-premises environment with [AWS Direct Connect](#).

### SQL Server

When you use an existing Active Directory, you have the option of using an existing SQL Server service account or creating a new account. If you create a new Active Directory account, you must create a new SQL Server account.

- **User name and password.** If you are using an existing SQL Server service account, provide your user name and password. This SQL Server service account should be part of the Managed Active Directory in which you are deploying. If you are creating a new SQL Server service account through Launch Wizard, enter a user name for the SQL Server service account. Create a complex Password that is at least 8 characters long, and then reenter the password to verify it. See [Password Policy](#) for more information.
  - **SQL Server install type.** Select the version of SQL Server Enterprise that you want to deploy. You can select an AMI from either the License-included AMI or Custom AMI dropdown lists.
  - **Additional SQL Server settings (Optional).** You can optionally specify additional nodes and their subnets.
    - **Nodes.** Enter a **Primary SQL node name** and a **Secondary SQL node name (HA deployments only)**.
    - **Additional secondary SQL node (HA deployments only, maximum of 5).** Enter a secondary **Node name**, and select the **Access type**, the **Private subnet**, and the **Dedicated host** for the additional secondary SQL node from the dropdown lists. You can add more secondary nodes by selecting **Add an additional secondary node**.
    - **Additional naming.** Enter a **Database name**. For HA deployments, enter an **Availability group name**, a **Listener name**, or a **Cluster name**.
5. When you are satisfied with your configuration selections, select **Next**. If you don't want to complete the configuration, select **Cancel**. When you select **Cancel**, all of the selections on the specification page are lost and you are returned to the landing page. To go to the previous screen, select **Previous**.
  6. After configuring your application, you are prompted to define the infrastructure requirements for the new deployment on the **Define infrastructure requirements** page. The following tabs provide information about the specification fields.

### Storage and Compute

You can choose to select your instance and volume types, or to use AWS recommended resources. If you choose to use AWS recommended resources, you have the option of defining your high availability cluster needs. If no selections are made, default values are assigned.

- **Number of instance cores.** Choose the number of CPU cores for your infrastructure. The default value assigned is 4.
- **Network performance.** Choose your preferred network performance in Gbps.
- **Expected RAM size (Memory).** Choose the amount of RAM that you want to attach to your EC2 instances. The default value assigned is 4 GB.
- **Type of storage drive.** Select the storage drive type for the SQL data, logs, and tempdb volumes. The default value assigned is SSD.
- **SQL Server throughput.** Select the sustained SQL Server throughput that you need.
- **Recommended resources.** Launch Wizard displays the system-recommended resources based on your infrastructure selections. If you want to change the recommended resources, select different infrastructure requirements.
- **Drive letters and volume size**
  - **Drive letter.** Select the storage drive letter for **Root drive**, **Logs**, **Data**, and **Backup** volumes.

#### **Important**

For custom AMIs, Launch Wizard assumes the root volume drive is C:.

- **Volume size.** Select the size of the SQL Server data volume in Gb for **Root drive**, **Logs**, **Data**, and **Backup** volumes.
- **Provisioned IOPS.** Select the IOPS value of the SQL Server data volume in IOPS for **Root drive**, **Logs**, **Data**, and **Backup** volumes.

For throughput limits and volume characteristics, see [Amazon EBS Volume Types](#) in the Amazon EC2 User Guide.

#### Tags-Optional

You can provide optional custom tags for the resources Launch Wizard creates on your behalf. For example, you can set different tags for EC2 instances, EBS volumes, VPC, and subnets. If you select **All**, you can assign a common set of tags to your resources. Launch Wizard assigns tags with a fixed key `LaunchWizardResourceGroupID` and value that corresponds to the ID of the AWS Resource Group created for a deployment. Launch Wizard does not support custom tagging for root volumes.

#### Estimated On-Demand Cost to Deploy Additional Resources

AWS Launch Wizard provides an estimate for application charges incurred to deploy the selected resources. The estimate updates each time you change a resource type in the Wizard. The provided estimates are only for general comparisons. They are based upon On-Demand costs and your actual costs may be lower.

7. When you are satisfied with your infrastructure selections, select **Next**. If you don't want to complete the configuration, select **Cancel**. When you select **Cancel**, all of the selections on the specification page are lost and you are returned to the landing page. To go to the previous screen, select **Previous**.
8. On the **Review and deploy** page, review your configuration details. If you want to make changes, select **Previous**. To stop, select **Cancel**. When you select **Cancel**, all of the selections on the specification page are lost and you are returned to the landing page. If you are ready to deploy, read and select the check box next to the Acknowledgment. Then choose **Deploy**.
9. Launch Wizard validates the inputs and notifies you if something must be addressed.
10. When validation is complete, Launch Wizard deploys your AWS resources and configures your SQL Server Always On application. Launch Wizard provides you with status updates about the progress of the deployment on the Deployments page. From the Deployments page, you can also view the list of current and previous deployments.

11. When your deployment is ready, you see notification that your SQL Server Always On application was successfully deployed. If you have set up SNS notification, you are also alerted through SNS. You can manage and access all of the resources related to your SQL Server Always On application by selecting **Manage**.
12. When the SQL Server Always On application is deployed, you can access your Amazon EC2 instances through the EC2 console. You can also use [AWS Systems Manager](#) to manage your SQL Server Always On application for future updates and patches through built-in integration via resource groups.

## Access and deploy an application with AWS Launch Wizard for SQL Server on Linux

### Topics

- [Access AWS Launch Wizard \(p. 24\)](#)
- [Deploy AWS Launch Wizard on Linux \(p. 24\)](#)
- [Post-deployment cluster tasks \(p. 29\)](#)

## Access AWS Launch Wizard

You can launch AWS Launch Wizard from the following locations.

- **AWS Console.** From the [AWS Management Console](#) under **Management and Governance**.
- **AMI launch page in EC2 console.** From the Launch Wizard banner that appears when you select **AMIs**, under **Images**, in the EC2 console. The banner appears when you search for SQL AMIs.
- **AWS Launch Wizard landing page.** From the AWS Launch Wizard page, located at <https://aws.amazon.com/launchwizard/>.

## Deploy AWS Launch Wizard on Linux

The following steps guide you through a SQL Server Always On application deployment with AWS Launch Wizard on the Linux platform after you have launched it from the console. For SQL Server deployments on Linux, you must use an instance type built on the [Nitro System](#). EBS volumes are exposed as NVMe block devices on instances built with the Nitro System. Device names that are specified for NVMe EBS volumes in a block device mapping are renamed using NVMe device names (`/dev/nvme[ [0-26 ]n1`). Launch Wizard deployments on Linux do not support block devices on Xen-virtualized instances.

1. When you select **Create deployment** from the AWS Launch Wizard landing page, you are directed to the **Choose application** page where you are prompted to select the type of application that you want to deploy. Select **Microsoft SQL Server Always On**.
2. Select the operating system on which you want to install the SQL Server Always On application — in this case, **Linux**.
3. After you select the operating system, under **Permissions**, Launch Wizard displays the AWS Identity and Access Management (IAM) policy required for Launch Wizard to access other AWS services on your behalf. For more information about setting up IAM for Launch Wizard, see [AWS Identity and Access Management \(IAM\) \(p. 15\)](#). Select **Next**.
4. After selecting the type of application to deploy and the operating system, you are prompted to enter specifications for the new deployment on the **Configure application settings** page. The following tabs provide information about the specification fields.

## General

- **Deployment model.** Choose **High availability deployment** to deploy your SQL Server Always On application across multiple Availability Zones or **Single instance deployment** to deploy your SQL Server application on a single node.
- **Deployment name.** Enter a unique application name for your deployment.
- **Simple Notification Service (SNS) topic ARN (Optional).** Specify an SNS topic where AWS Launch Wizard can send notifications and alerts. For more information, see the [Amazon Simple Notification Service Developer Guide](#).

## Connectivity

Enter specifications for how you want to connect to your application instance and what kind of Virtual Private Cloud (VPC) you want to set up.

### Key pair name

- Select an existing key pair from the dropdown list or create a new one. If you select **Create new key pair name** to create a new key pair, you are directed to the Amazon EC2 console. From there, under **Network and Security**, choose **Key Pairs**. Choose **Create a new key pair**, enter a name for the key pair, and then choose **Download Key Pair**.

#### Important

This is the only chance for you to save the private key file, so be sure to download it and save it in a safe place. You must provide the name of your key pair when you launch an instance, and provide the corresponding private key each time that you connect to the instance.

Return to the Launch Wizard console and choose the refresh button next to the **Key Pairs** dropdown list. The newly created key pair appears in the dropdown list. For more information about key pairs, see [Amazon EC2 Key Pairs and Windows Instances](#).

**Virtual Private Cloud (VPC).** Choose whether you want to use an existing VPC or create a new VPC.

- **Select Virtual Private Cloud (VPC) option.** Choose the VPC that you want to use from the dropdown list. Your VPC must contain one public subnet. For HA deployments, it must also contain at least three private subnets. For single node deployments, it must contain one private subnet. The private subnets must have outbound connectivity to the internet and other AWS services (S3, CFN, SSM, Logs). We recommend that you enable this connectivity with a NAT Gateway. For more information about NAT Gateways, see [NAT Gateways](#) in the Amazon VPC User Guide.
- **Public Subnet.** Your VPC must contain one public subnet. For HA deployments it must also contain three private subnets. For single node deployments, it must contain one private subnet. Choose a public subnet for your VPC from the dropdown list. To continue, you must select the check box that indicates that the Public subnet has been set up and each of the selected private subnets have outbound connectivity enabled.

### To add a new public subnet

If the traffic of a subnet is routed to an internet gateway, the subnet is known as a public subnet. If, however, a subnet doesn't have a route to the internet gateway, the subnet is known as a private subnet. To use an existing VPC that does not have a public subnet, you can add a new public subnet using the following steps.

- Follow the steps in [Creating a Subnet in the Amazon VPC User Guide](#) using the existing VPC you intend to use AWS Launch Wizard.
- To add an internet gateway to your VPC, follow the steps in [Attaching an Internet Gateway](#) in the Amazon VPC User Guide.
- To configure your subnets to route internet traffic through the internet gateway, follow the steps in [Creating a Custom Route Table](#) in the Amazon VPC User Guide. Use IPv4 format (0.0.0.0/0) for Destination.
- The public subnet should have the “auto-assign public IPv4 address” setting enabled. To enable this setting, follow the steps in [Modifying the Public IPv4 Addressing Attribute for Your Subnet](#) in the Amazon VPC User Guide.
- **Availability Zone (AZ) configuration.** You must choose at least three Availability Zones for High Availability (HA) deployments and one Availability Zone for single-node deployments, with one private subnet for each zone that you select. From the dropdown lists, select the **Availability Zones** within which you want to deploy your **primary** and **secondary**, and **configuration** nodes.

### To create a private subnet

If a subnet doesn't have a route to an internet gateway, the subnet is known as a private subnet. To create a private subnet, you can use the following steps. We recommend that you enable the outbound connectivity for each of your selected private subnets using a NAT Gateway. To enable outbound connectivity from private subnets with public subnet, see the steps in [Creating a NAT Gateway](#) to create a NAT Gateway in your chosen public subnet. Then, follow the steps in [Updating Your Route Table](#) for each of your chosen private subnets.

- Follow the steps in [Creating a Subnet](#) in the Amazon VPC User Guide using the existing VPC you will use in AWS Launch Wizard.
- When you create a VPC, it includes a main route table by default. On the **Route Tables** page in the Amazon VPC console, you can view the main route table for a VPC by looking for **Yes** in the **Main** column. The main route table controls the routing for all subnets that are not explicitly associated with any other route table. If the main route table for your VPC has an outbound route to an internet gateway, then any subnet created using the previous step, by default, becomes a public subnet. To ensure the subnets are private, you may need to create one separate route table for all of your private subnets. This route table must not contain any routes to an internet gateway. Verify that all of the private subnets have the same route table association.
- **Create new Virtual Private Cloud (VPC)** option. Launch Wizard creates your VPC. You can optionally enter a VPC name tag.

## SQL Server

### SQL Server configuration

- **User name and password.** By default, Launch Wizard applies the user name `sa`. This system administrator account is used for SQL Server management. Create a complex password that is at least 8 characters long, and then reenter the password to verify it. See [Password Policy](#) for more information.
- **Floating IP Address (HA deployments only).** This field is available when you select a Virtual Private Cloud (VPC). The IP address that you enter is used as the endpoint for your SQL Server Availability Group listener. Launch Wizard creates a route from this IP address to the SQL primary node in your route table. Verify that the IP address is not already used within your VPC and is outside of all of the provided subnet CIDRs.

- **Amazon Machine Image (AMI).** Select the version of Microsoft SQL Server Enterprise to deploy. You can select an AMI from the lists of either license-included or custom AMIs.
- **SQL Server Edition.** This field is available when you select a custom AMI. Choose the edition of SQL Server for the custom AMI: **SQL Enterprise** or **SQL Standard**.

### Pacemaker cluster configuration (HA deployments only)

Pacemaker is a high-availability cluster resource manager. This software runs on a set of hosts, or cluster of nodes, to preserve integrity and minimize the downtime of selected services, or resources. Pacemaker is maintained by the [ClusterLabs](#) community.

- **Pacemaker cluster name.** Enter a name to identify your pacemaker cluster.
- **Pacemaker cluster user name.** By default, Launch Wizard applies the pacemaker user name `hacluster`. This user name is used to securely communicate between cluster nodes.
- **Pacemaker cluster password.** Create a complex password that is at least 8 characters long, and then reenter the password to verify it. See [Password Policy](#) for more information.

### SQL - Pacemaker cluster connection settings (HA deployments only)

After you configure Pacemaker cluster and SQL Server, you must create a user in SQL Server to communicate with Pacemaker.

- **SQL Pacemaker user name and password.** Enter a user name for SQL Server to communication with the Pacemaker cluster. Create a complex password that is at least 8 characters long, and then reenter the password to verify it. See [Password Policy](#) for more information.
- **S3 location for node certificates.** An Amazon S3 bucket location is required by the SQL nodes to share self-signed certificates with each other. Provide the bucket or object locations and verify that the names begin with `launchwizard`.

### Additional SQL Server settings (optional)

- **Nodes.** Enter a **Primary SQL node name**, a **Secondary SQL node name**, and a **Configuration node name**.
  - **Additional naming.** Enter a **Database name** and an **Availability group name**.
5. When you are satisfied with your configuration selections, select **Next**. If you don't want to complete the configuration, select **Cancel**. When you select **Cancel**, all of the selections on the specification page are lost and you are returned to the landing page. To go to the previous screen, select **Previous**.
  6. After configuring your application, you are prompted to define the infrastructure requirements for the new deployment on the **Define infrastructure requirements** page. The following tabs provide information about the specification fields.

#### Storage and Compute

You can choose to select your instance and volume types, or to use AWS recommended resources. If you choose to use AWS recommended resources, you have the option of defining your high availability cluster needs. If no selections are made, default values are assigned.

- **Number of instance cores.** Choose the number of CPU cores for your infrastructure. The default value assigned is 4.
- **Network performance.** Choose your preferred network performance in Gbps.



- **Expected RAM size (Memory).** Choose the amount of RAM that you want to attach to your EC2 instances. The default value assigned is 4 GB.
- **Type of storage drive.** Select the storage drive type for the SQL data, logs, and tempdb volumes. The default value assigned is SSD.
- **SQL Server throughput.** Select the sustained SQL Server throughput that you need.
- **Recommended resources.** Launch Wizard displays the system-recommended resources based on your infrastructure selections. If you want to change the recommended resources, select different infrastructure requirements.
- **Volume sizes**
  - **Volume size.** Select the size of the SQL Server data volume in Gb for **Temporary database, Logs, Data, and Backup** volumes.
  - **Provisioned IOPS.** Select the IOPS value of the SQL Server data volume in IOPS for **Temporary database, Logs, Data, and Backup** volumes.

For throughput limits and volume characteristics, see [Amazon EBS Volume Types](#) in the Amazon EC2 User Guide.

#### Tags-Optional

You can provide optional custom tags for the resources Launch Wizard creates on your behalf. For example, you can set different tags for EC2 instances, EBS volumes, VPC, and subnets. If you select **All**, you can assign a common set of tags to your resources. Launch Wizard assigns tags with a fixed key `LaunchWizardResourceGroupID` and value that corresponds to the ID of the AWS Resource Group created for a deployment. Launch Wizard does not support custom tagging for root volumes.

#### Estimated On-Demand Cost to Deploy Additional Resources

AWS Launch Wizard provides an estimate for application charges incurred to deploy the selected resources. The estimate updates each time you change a resource type in the Wizard. The provided estimates are only for general comparisons. They are based upon On-Demand costs and your actual costs may be lower.

7. When you are satisfied with your infrastructure selections, select **Next**. If you don't want to complete the configuration, select **Cancel**. When you select **Cancel**, all of the selections on the specification page are lost and you are returned to the landing page. To go to the previous screen, select **Previous**.
8. On the **Review and deploy** page, review your configuration details. If you want to make changes, select **Previous**. To stop, select **Cancel**. When you select **Cancel**, all of the selections on the specification page are lost and you are returned to the landing page. If you are ready to deploy, read and select the check box next to the Acknowledgment. Then choose **Deploy**.
9. Launch Wizard validates the inputs and notifies you if something must be addressed.
10. When validation is complete, Launch Wizard deploys your AWS resources and configures your SQL Server Always On application. Launch Wizard provides you with status updates about the progress of the deployment on the Deployments page. From the Deployments page, you can also view the list of current and previous deployments.
11. When your deployment is ready, you see notification that your SQL Server Always On application was successfully deployed. If you have set up SNS notification, you are also alerted through SNS. You can manage and access all of the resources related to your SQL Server Always On application by selecting **Manage**.
12. When the SQL Server Always On application is deployed, you can access your Amazon EC2 instances through the EC2 console. You can also use [AWS Systems Manager](#) to manage your SQL Server Always On application for future updates and patches through built-in integration via resource groups.



## Post-deployment cluster tasks

The Launch Wizard Pacemaker implementation includes three cluster nodes: primary, secondary, and configuration only. The primary node provides the Microsoft SQL Server for Linux resource and the floating IP address. To ensure that the cluster operates correctly, some administrative tasks must be performed in a specific way. If these tasks are performed incorrectly, then Pacemaker may identify the activity as a resource failure and attempt to fail over the resources to the secondary node. If the resources are failed over to the secondary node, the cluster can remain in an unknown state, which can impact user access.

There are four primary tasks: Start Cluster, Stop Cluster, Move Resources, and Recovery. These tasks must be carried out by a sudo user with an SSH connection to any of the cluster nodes. Before performing any of these tasks, verify the cluster status using `pcs resource status --all`. Running this command returns all cluster issues. All issues must be addressed prior to performing any administrative tasks.

### Start cluster

1. Log in to a cluster node using a sudo user over an SSH connection.
2. Verify that all cluster nodes are available.
3. Verify cluster status using the following command: `pcs resource --all`.

Address all issues before attempting to start the cluster.

4. Start all cluster nodes using the following command: `pcs cluster start --all --wait`.
5. Verify that the cluster has started using the following command: `pcs resource --all`.

The output provides information about the cluster nodes and cluster resources. All cluster nodes should be online and all resource agents should be visible and allocated to their assigned cluster nodes.

6. Verify that the availability group listener is available by pinging the floating IP address.

### Manually move cluster resources

1. Log in to a cluster node using a sudo server over an SSH connection.
2. Verify that all cluster nodes are available.
3. Verify cluster status using the following command: `pcs resource --all`.

Address all issues before attempting to start the cluster.

4. Run the following command: `pcs resource move <RESOURCE_NAME> <NODE_NAME> --force`.

This command moves the resource agent to `<NODE_NAME>` and starts the resource. All cluster constraints will be applied. If the Microsoft SQL Server resource agent is moved, then the Availability Group listener will follow.

5. Verify cluster status using the following command: `pcs resource --all`.

The resource that was moved should be located on the `<NODE_NAME>`.

6. Clear temporary constraints using the following command: `pcs resource clear <RESOURCE_NAME>`.

### Stop cluster

1. Log in to a cluster node using a sudo server over an SSH connection.
2. Verify that all cluster nodes are available.
3. Verify cluster status using the following command: `pcs resource --all`.

Address all issues before attempting to start the cluster.

4. Stop the cluster using the following command: `pcs cluster stop --ALL`. This will gracefully shut down all of the cluster nodes.
5. Verify the shut down status using the following command: `pcs status --all`.

This command should return that the cluster is no longer running.

## Recovery

If a node is restarted from the operating system or the AWS Management Console, the Pacemaker node and its related services will not automatically start. This prevention protects the high availability database replicas from split-brain corruption.

The following steps are required to restore the cluster to normal operations.

1. Log in to a cluster node using a sudo server over an SSH connection.
2. Determine the node that was restarted using the following command: `pcs resource --ALL`. The restarted node will be offline.
3. Verify cluster status using the following command: `pcs resource --all`.

Address all issues before attempting to start the cluster.

4. Start the restarted node using the following command: `pcs cluster start --<NODE_NAME>`.
5. Verify cluster status using the following command: `pcs resource --all`.

Address all issues before attempting to start the cluster.

6. If the restarted node is the primary node of the cluster, then the Availability Group resource must be returned to the primary node.
7. Remove all temporary constraints using the following commands: `pcs resource clear <AG_RESOURCE>` and `pcs resource clear <AG_LISTENER>`.
8. Run the following command: `pcs resource move <RESOURCE_NAME> <PRI_NODE_NAME> --force`.

This command moves the resources to `<PRI_NO_NAME>` and starts the resource. Any cluster constraints are applied. In this scenario, if the Microsoft SQL Server resource agent is moved, then the Availability Group listener follows.

9. Verify cluster status using the following command: `pcs resource --all`. The restarted node will be located on `<PRI_NO_NAME>`.

# Managing application resources with AWS Launch Wizard for SQL Server

After your SQL Server Always On application is deployed, you can manage it by following these steps.

1. From the navigation pane, choose **Deployments**.
2. From the **Deployments** page, select **Actions**. You can select to do the following:
  1. **Manage resources on the EC2 console.** You are taken to the Amazon EC2 console, where you can view and manage your SQL Server Always On application resources. For example, you can view and manage EC2, Amazon EBS, Active Directory, Amazon VPC, Subnets, NAT Gateways, and Elastic IPs. For SQL Server on Linux deployments, you can use AWS Systems Manager Session Manager to manage your deployed EC2 instances. For more information about SSM Session Manager, see [AWS Systems Manager Session Manager](#).

2. **Access SQL Server using RDGW instance (Windows deployments).** Connect to SQL Server via Remote Desktop Protocol. For more information, see [Connecting to your Windows Instance](#) in the *User Guide for Windows Instances*.
3. **View resource group with SSM.** You are taken to the Systems Manager console to view your resource groups.
4. **View SSM deployment template (Windows deployments).** You are taken to the Systems Manager console to view your documents.
5. **View CloudWatch application Logs.** You are taken to CloudWatch Logs, where you can monitor, store, and access your SQL Server Always On application log files.
6. **View your CloudFormation template.** This is the CloudFormation template created by your most recent deployment, and it can be accessed through the CloudFormation console. For help with finding and using your CloudFormation template, see [Viewing AWS CloudFormation Stack Data and Resources on the AWS Management Console](#).
7. If you have not set up monitoring for your application on CloudWatch Application Insights, you have the option to **Set up monitoring on CloudWatch Application Insights**. You are taken to the CloudWatch Application Insights console to set up monitoring for your application.

If you have set up monitoring for your application on CloudWatch Application Insights, you can **View insights on Amazon CloudWatch**. You are taken to the application monitoring dashboards on the CloudWatch Application Insights console.

3. To delete a deployment, select the application that you want to delete and select **Delete**. You are prompted to confirm your action.

#### **Important**

You lose all specification settings for the SQL Server Always On application when you delete a deployment. Launch Wizard attempts to delete only the AWS resources that it created in your account as part of the deployment. If you created resources outside of Launch Wizard, for example resources that reside in a VPC created by Launch Wizard, the deletion may fail. Launch Wizard does not delete any Active Directory objects in your Active Directory, nor any of the records in your DNS server. Launch Wizard has no control over your Active Directory domain user password over time, which is required to clean up Active Directory objects or DNS records. We recommend that you remove these entries from your Active Directory after Launch Wizard deletes the deployment. For key operations performed against your Active Directory resulting in new records or entries, see [AWS Managed Active Directory \(p. 14\)](#).

4. To drill down into details regarding your SQL Server Always On application resources, select the **Application name**. You can then view the **Deployment events** and **Summary** details for your application by using the tabs at the top of the page.

## Monitoring SQL Server Always On deployments

You can monitor your SQL Server Always On deployments using Amazon CloudWatch Application Insights. When you [select the option to monitor your deployment \(p. 18\)](#) using the Launch Wizard console, Application Insights identifies and sets up key metrics, logs, and alarms across your application resources and technology stack for your Microsoft SQL Server database. Anomalies and errors are detected and correlated as Application Insights continuously monitors metrics and logs. When errors and anomalies are detected, Application Insights generates CloudWatch Events that you can use to set up notifications or take action. To help with troubleshooting, Application Insights creates automated dashboards for detected problems, which include correlated metric anomalies and log errors, along with additional insights to point you to a possible root cause. Use the automated dashboards to take remedial actions to keep your applications healthy and prevent end-user impact. You can also resolve problems with [AWS Systems Manager OpsCenter](#) using generated OpsItems.

For Microsoft SQL Server High Availability (HA) workloads, you can use CloudWatch Application Insights to configure important counters, such as Mirrored Write Transaction/sec, Recovery Queue Length,

and Transaction Delay, and Windows Event Logs on CloudWatch. You can also get automated insights whenever a failover event or problem, such as restricted access to query a target database, is detected with SQL HA workloads. See the Amazon CloudWatch Application Insights documentation for a complete list of [Logs and metrics supported by Application Insights](#).

## High availability and security best practices for AWS Launch Wizard for SQL Server

The application architecture created by AWS Launch Wizard supports AWS best practices for high availability and security as promoted by the [AWS Well-Architected Framework](#).

### Topics

- [High availability \(p. 32\)](#)
- [Automatic failover \(p. 32\)](#)
- [Security groups and firewalls \(p. 33\)](#)

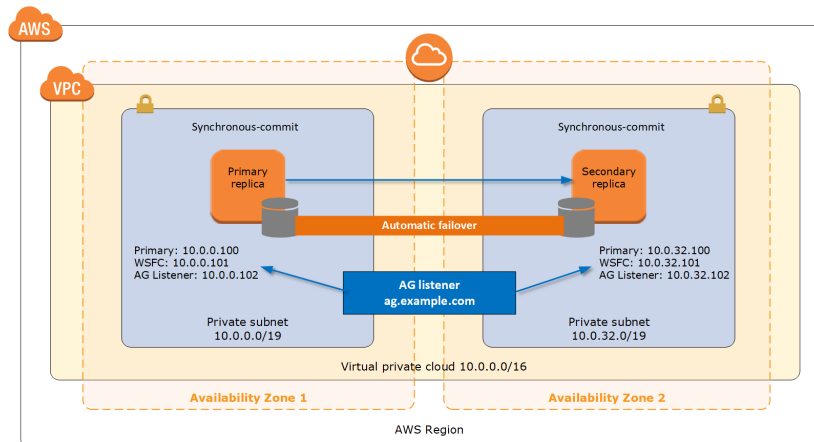
## High availability

With Amazon EC2, you can set the location of instances in multiple locations composed of AWS Regions and Availability Zones. Regions are dispersed and located in separate geographic areas. Availability Zones are distinct locations within a Region that are engineered to be isolated from failures in other Availability Zones. Availability Zones provide inexpensive, low-latency network connectivity to other Availability Zones in the same Region.

When you launch your instances in different Regions, you can set your SQL Server Always On application to be closer to specific customers, or to meet legal or other requirements. When you launch your instances in different Availability Zones, you can protect your SQL Server Always On applications from the failure of a single location. [Windows Server Failover Clustering \(WSFC\)](#) provides infrastructure features that complement the high availability and disaster recovery scenarios supported in the AWS Cloud.

## Automatic failover

When you deploy AWS Launch Wizard with the default parameters, it configures a two-node automatic failover cluster with a file share witness. An Always On Availability Group is deployed on this cluster with two availability replicas, as shown in the following diagram.



Launch Wizard implementation supports the following scenarios:

- Protection from the failure of a single instance
- Automatic failover between the cluster nodes
- Automatic failover between Availability Zones

The default implementation of Launch Wizard does not provide automatic failover in every case. For example, the failure of Availability Zone 1, which contains the primary node and file share witness, would prevent automatic failover to Availability Zone 2 because the cluster would fail as it loses quorum. In this scenario, you could follow manual disaster recovery steps that include restarting the cluster service and forcing quorum on the second cluster node (for example, WSFCNode2) to restore application availability. Launch Wizard also provides an option to deploy to three Availability Zones. This deployment option can mitigate the loss of quorum if a single node fails. However, you can select this option only in AWS Regions that include three or more Availability Zones. For a current list of supported Regions, see [AWS Global Infrastructure](#).

## Security groups and firewalls

Launch Wizard creates a number of security groups and rules for you. When Amazon EC2 instances are launched, they must be associated with a security group, which acts as a stateful firewall. You have complete control over the network traffic entering or leaving the security group. You can also build granular rules that are scoped by protocol, port number, and source or destination IP address or subnet. By default, all outbound traffic from a security group is permitted. Inbound traffic, on the other hand, must be configured to allow the appropriate traffic to reach your instances.

The [Securing the Microsoft Platform on Amazon Web Services](#) whitepaper discusses the different methods for securing your AWS infrastructure. Recommendations include providing isolation between application tiers using security groups. We recommend that you tightly control inbound traffic in order to reduce the attack surface of your EC2 instances.

Domain controllers and member servers require several security group rules to allow traffic for services such as AD DS replication, user authentication, Windows Time services, and Distributed File System (DFS), among others. The WSFC nodes running SQL Server must permit several additional ports to communicate with each other. Finally, instances launched into the application server tier must establish SQL client connections to the WSFC nodes.

For a detailed list of port mappings, see the [Security](#) section of the *Active Directory DS Quick Start guide*.

In addition to security groups, the Windows Firewall must also be modified on the SQL Server instances. During the bootstrapping process, a script runs on each instance that opens the TCP ports 1433, 1434, 4022, 5022, 5023, and 135 on the Windows Firewall.

## Troubleshoot AWS Launch Wizard for SQL Server

Each application in your account in the same AWS Region can be uniquely identified by the application name specified at the time of a deployment. The application name can be used to view the details related to the application launch.

For SQL Server deployments on Linux, you must use an instance type built on the [Nitro System](#). EBS volumes are exposed as NVMe block devices on instances built with the Nitro System. Device names that are specified for NVMe EBS volumes in a block device mapping are renamed using NVMe device names (`/dev/nvme[ [0-26 ]n1`). Launch Wizard deployments on Linux do not support block devices on Xen-virtualized instances.

### Contents

- [Active Directory objects and DNS record clean up \(deployment on Windows\) \(p. 34\)](#)
- [Launch Wizard provisioning events \(p. 35\)](#)
- [CloudWatch Logs \(p. 35\)](#)
- [SSM Automation execution \(p. 35\)](#)
- [AWS CloudFormation stack \(p. 35\)](#)
- [Pacemaker on Ubuntu \(deployment on Linux\) \(p. 36\)](#)
- [Errors \(p. 36\)](#)

## Active Directory objects and DNS record clean up (deployment on Windows)

When you delete a deployment, you lose all specification settings for the SQL Server Always On application. Launch Wizard attempts to delete only the AWS resources that it created in your account as part of the deployment. If you created resources outside of Launch Wizard, for example, resources in a VPC created by Launch Wizard, the deletion can fail. Launch Wizard does not delete Active Directory objects in your Active Directory, nor does it delete any of the records in your DNS server. Launch Wizard has no control over your Active Directory domain user password over time, which is required to clean up Active Directory objects or DNS records. We recommend that you remove these entries from your Active Directory after Launch Wizard deletes the deployment.

If the initial Active Directory objects or DNS records are not cleaned up, then when you attempt to deploy Launch Wizard on an existing Active Directory using a deployment name that has already been used or Availability Group/Listener/Cluster name that has already been used, the deployment may fail with the following error.

### Error message

```
System.Management.Automation.Remoting.PSRemotingTransportException: Connecting to remote server xxxxxx failed with the following error message : WinRM cannot complete the operation. Verify that the specified computer name is valid, that the computer is accessible over the network, and that a firewall exception for the WinRM service is enabled and allows access from this computer. By default, the WinRM firewall exception for public profiles limits access to remote computers within the same local subnet.
```

To address this error, we recommend that you remove the initial entries from your Active Directory.

To clean up Active Directory Objects, run the following example PowerShell commands as a domain user with the appropriate authorization to perform these operations.

```
$Pwd = ConvertTo-SecureString $password -AsPlainText -Force  
  
$cred = New-Object System.Management.Automation.PSCredential $domainUser, $Pwd  
  
$ADObject = Get-ADObject -Filter 'DNSHostName -eq "SQLnode.example.com"  
  
Remove-ADObject -Recursive -Identity $ADObject -Credential $cred
```

To remove a DNS Record, the name of the record that you want to delete (SQL Server node name), the DNS server FQDN, and the DNS zone within which the record is residing are required. The following are example PowerShell commands to perform the DNS record removal.

```
$NodeDNS = Get-DnsServerResourceRecord -ZoneName $ZoneName -ComputerName  
$DNSServer -Node $NodeToDelete -RRType A -ErrorAction SilentlyContinue
```

```
Remove-DnsServerResourceRecord -ZoneName $ZoneName -ComputerName $DNSServer -  
InputObject $NodeDNS -Force
```

## Launch Wizard provisioning events

Launch Wizard captures events from SSM Automation and AWS CloudFormation to track the status of an ongoing application deployment. If an application deployment fails, you can view the deployment events for this application by selecting **Deployments** from the navigation pane. A failed event shows a status of **Failed** along with a failure message.

## CloudWatch Logs

Launch Wizard streams provisioning logs from all of the AWS log sources, such as AWS CloudFormation, SSM, and CloudWatch Logs. CloudWatch Logs for a given application name can be viewed on the CloudWatch console for the log group name `LaunchWizard-APPLICATION_NAME` and log stream `ApplicationLaunchLog`.

## SSM Automation execution

Launch Wizard uses SSM Automation to provision SQL Server Always On applications. SSM Automation execution can be found in your account using the `ssm describe-automation-executions` API, and adding document name prefix filters. Launch Wizard launches various automation documents in your account for validation and application provisioning. The following are the relevant filters for the `ssm describe-automation-executions` API.

- **Validation: Validate VPC connectivity**

`LaunchWizard-Validate-VPC-Connectivity-APPLICATION_NAME-Subnet_id`, where `Subnet_id` is the subnet on which to perform the validation.

- **Validation: Validate credentials**

`LaunchWizard-Validate-Credentials-APPLICATION_NAME`

- **Application Provisioning: Provisioning resources and Post Configuration**

`LaunchWizard-SQLHAAwaysOn-APPLICATION_NAME-Provision`

You can view the status of these SSM Automation executions. If any of them fail, you can view the cause of the failure.

## AWS CloudFormation stack

Launch Wizard uses AWS CloudFormation to provision the infrastructure resources of an application. CloudFormation stacks can be found in your account using the CloudFormation `describe-stacks` API. Launch Wizard launches various stacks in your account for validation and application resource creation. The following are the relevant filters for the `describe-stacks` API.

- **Validation**

`LaunchWizard-APPLICATION_NAME-checkCredentials-SSM_execution_id`

- **Validation**

`LaunchWizard-APPLICATION_NAME-checkVPCConnectivity-SSM_execution_id`

- **Application resources**



LaunchWizard-*APPLICATION\_NAME*. This stack also has nested stacks for VPC, AD, the RDGW node, and SQL nodes.

You can view the status of these CloudFormation stacks. If any of them fail, you can view the cause of failure.

## Pacemaker on Ubuntu (deployment on Linux)

To troubleshoot Pacemaker cluster resource issues, take the following actions as an administrator.

- Inspect the system log files for operating system errors and address the errors, as needed.
- Inspect the cluster log files for errors, including for errors that relate to Pacemaker, Corosync, or SQL Server. Check the log files carefully because the related services may provide only one or two related log entries.
- Verify resource configuration, and configuration of cluster-related functions.
  - The following commands display the configuration details:
    - To display all resources, use:`pcs resource show -full`
    - Or, you can use:`pcs resource show <resource name>`
  - The following command will display the cluster constraints:`pcs constraints -full`
  - The following command displays the cluster properties:`pcs property list -all`
- Manually start the resource with `debug-start`.
- Clear failed actions with the following command: `pcs resource cleanup <resource <resource name`

## Errors

### Directory fails to create

- **Cause:** An internal service error has been encountered during the creation of the directory.
- **Solution:** Retry the operation. For this scenario, you must retry the deployment from the initial page of the Launch Wizard console.

### Your requested instance type is not supported in your requested Availability Zone

- **Cause:** This failure might happen during the launch of either your RDGW instance or your SQL Server instance, or during the validation of the instances that Launch Wizard launches in your selected subnets.
- **Solution:** For this scenario, you must choose a different Availability Zone and retry the deployment from the initial page of the Launch Wizard console.

### Validate connectivity for subnet. The following resource(s) failed to create: [ValidationNodeWaitCondition]

This failure can occur for multiple reasons. The following list shows known causes and solutions.

- **VPC or subnet configuration does not meet prerequisites**
  - **Cause:** This failure occurs when your VPC or subnet configuration does not meet the prerequisites documented in the VPC Connectivity Section under [Access and deploy an application with AWS Launch Wizard for SQL Server on Windows \(p. 18\)](#). If the failure message points to your selected



public subnet, then the public subnet is not configured for outbound internet access. If the failure message points to one of your selected private subnets, then the specified private subnet does not have outbound connectivity.

- **Solution:** Check that your VPC includes one public subnet and, at least, two private subnets. Your VPC must be associated with a DHCP Options Set to enable DNS translations to work. The private subnets must have outbound connectivity to the internet and other AWS services (S3, CFN, SSM, and Logs). We recommend that you enable this connectivity with a NAT Gateway. Note that, in the console, when you select a private subnet for the public subnet dropdown or you select a public subnet for the private subnet dropdown, you will encounter the same error. Please refer to the VPC Connectivity section under [Access and deploy an application with AWS Launch Wizard for SQL Server on Windows \(p. 18\)](#) for more information about how to configure your VPC.
- **EC2 instance stabilization error**
  - **Cause:** Failure can occur if the EC2 instance used for validation fails to stabilize. When this happens, the EC2 instance is unable to communicate to the CloudFormation service to signal completions, resulting in `waitCondition` errors.
  - **Solution:** Please contact [AWS Support](#) for assistance.

# AWS Launch Wizard for SAP

This section of the AWS Launch Wizard documentation contains guidance specific to the deployment of SAP applications on AWS using the Launch Wizard service.

## Topics

- [What is AWS Launch Wizard for SAP? \(p. 38\)](#)
- [Get started with AWS Launch Wizard for SAP \(p. 47\)](#)
- [Manage application resources with AWS Launch Wizard for SAP \(p. 68\)](#)
- [Troubleshoot AWS Launch Wizard for SAP \(p. 69\)](#)
- [Make SAP HANA software available for AWS Launch Wizard to deploy a HANA database \(p. 70\)](#)
- [Supported operating system versions for SAP deployments \(p. 72\)](#)
- [Security groups in AWS Launch Wizard for SAP \(p. 73\)](#)

## What is AWS Launch Wizard for SAP?

AWS Launch Wizard for SAP is a service that guides you through the sizing, configuration, and deployment of SAP applications on AWS, and follows [AWS cloud application best practices](#).

AWS Launch Wizard reduces the time it takes to deploy SAP applications on AWS. You input your application requirements, including SAP HANA settings, SAP landscape settings, and deployment details on the service console, and Launch Wizard identifies the AWS resources to deploy and run your SAP application. Launch Wizard provides an estimated cost of deployment, and you can modify your resources and instantly view the updated cost. When you approve your settings, Launch Wizard provisions and configures the selected resources. It then optionally installs an SAP HANA database to deploy and run SAP HANA and SAP Netweaver-based applications. For subsequent deployments, Launch Wizard creates custom AWS CloudFormation templates that can be reused and customized.

After you deploy an SAP application, you can access it from the Amazon EC2 console. You can manage your SAP applications with [AWS Systems Manager](#).

## Contents

- [Supported deployments and features of AWS Launch Wizard \(p. 38\)](#)
- [Components \(p. 40\)](#)
- [Related services \(p. 41\)](#)
- [How AWS Launch Wizard for SAP works \(p. 43\)](#)

## Supported deployments and features of AWS Launch Wizard

### Supported deployments

AWS Launch Wizard currently supports the deployment of AWS resources for the following SAP systems and patterns. SAP HANA database software is optionally installed and customer provided.

- **SAP HANA database on a single Amazon EC2 instance.** Deploy SAP HANA in a single-node, scale-up architecture, with up to 24TB of memory.

- **SAP Netweaver on SAP HANA system on a single Amazon EC2 instance.** Deploy an SAP application on the same Amazon EC2 instance as your SAP HANA Database.
- **SAP HANA database on multiple EC2 instances.** Deploy SAP HANA in a multi-node, scale-out architecture.
- **SAP Netweaver system on multiple EC2 instances.** Deploy an SAP Netweaver system using a distributed deployment model, which includes an ASCS/PAS server, single/multiple SAP HANA servers running SAP HANA databases, and multiple application servers.
- **Cross-AZ SAP HANA database high availability setup.** Deploy SAP HANA with high availability configured across two Availability Zones.
- **Cross-AZ SAP Netweaver system setup.** Deploy Amazon EC2 instances for ASCS/ERS and SAP HANA databases across two Availability Zones, and spread the deployment of application servers across them.

AWS Launch Wizard provides the following features:

#### Features

- [Instance selection and configuration \(p. 39\)](#)
- [AWS resource selection \(p. 39\)](#)
- [Cost estimation \(p. 39\)](#)
- [Reusable infrastructure settings \(p. 39\)](#)
- [SNS notification \(p. 40\)](#)
- [Application resource groups \(p. 40\)](#)
- [AWS Data Provider for SAP \(p. 40\)](#)
- [AWS Backint Agent for SAP HANA \(p. 40\)](#)
- [Custom deployment configuration scripts \(p. 40\)](#)

## Instance selection and configuration

When you input the application requirements, Launch Wizard deploys the necessary AWS resources for a production-ready application. This means that you do not have to figure out how to select the right instances and configure them to run supported SAP applications.

## AWS resource selection

Launch Wizard considers CPU/Memory or SAPS requirements that you provide to determine the most appropriate instance types and other resources for your SAP application. You can modify the recommended defaults.

## Cost estimation

Launch Wizard provides a cost estimate for the complete deployment that is itemized for each individual resource being deployed. The estimated cost automatically updates each time you change a resource type configuration in the wizard. The provided estimates are only for general comparisons. They are based on On-Demand instance costs. Actual costs may be lower.

## Reusable infrastructure settings

You can save the settings for your AWS infrastructure for the SAP landscape to reuse when you want to deploy SAP systems that function similarly within a landscape. For example, a development configuration can be created for the first development instance, which can later be reused to deploy other development systems.

Some example scenarios for which DevOps and SAP architecture teams can create templates include:

- Organize the SAP systems within a landscape.
- Save infrastructure settings, including VPC, subnets, key pairs, and security groups to ensure that systems that must be deployed with the same settings are correctly deployed.
- Set up connectivity between the systems using the same configuration template so they can communicate with each other when security groups are created with Launch Wizard.
- Use the same GID for SPSYS group across different configuration templates to ensure that SAP transport files systems are mounted properly.

## SNS notification

You can provide an [SNS topic](#) so that Launch Wizard will send you notifications and alerts about the status of a deployment.

## Application resource groups

Launch Wizard creates a resource group for all of the AWS resources created for your SAP system. You can manage the resources through the Amazon EC2 console or by using Systems Manager.

## AWS Data Provider for SAP

Deploying and running the Amazon Web Services (AWS) Data Provider for SAP is a prerequisite for running SAP systems on AWS. Launch Wizard automatically deploys AWS Data Provider for SAP on every EC2 instance that it launches. AWS Data Provider for SAP is a tool that collects performance-related data from AWS services. It makes this data available to SAP applications to help monitor and improve the performance of business transactions. AWS 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 EC2 and Amazon CloudWatch.

## AWS Backint Agent for SAP HANA

Launch Wizard deploys and configures AWS Backint Agent for SAP HANA, an SAP-certified backup and restore application for SAP HANA workloads running on Amazon EC2 instances in the cloud. Launch Wizard supports the deployment and configuration of Backint Agent for single-node, multi-node, and high availability deployments for supported HANA and Netweaver on HANA applications.

## Custom deployment configuration scripts

You can provide custom pre-deployment and post-deployment configuration scripts that can run on various instance tiers, such as SAP HANA Database, Primary Application Server, and Enqueue Replication Server during the pre-deployment and post-deployment configuration phases. Launch Wizard uses a standalone component manager application (AWSTOE) to execute the scripts. For more information, see [Custom deployment configuration scripts \(p. 46\)](#).

## Components

An SAP application deployed with Launch Wizard includes the following components.

### **SAP applications:**

- **SAP HANA Database** supports the following:
  - single instance deployment

- distributed instance deployment in a single Availability Zone
- cross-Availability Zone, high-availability deployment
- **SAP applications based on Netweaver on SAP HANA database** supports the following:
  - single instance deployment
  - distributed instance deployment
  - cross-Availability Zone, high-availability deployment

### Security groups

Launch Wizard creates optional security groups to ensure that all of the systems sharing the same configuration template can communicate with each other and with systems and end users who access the SAP systems from an IP CIDR range, an external IP address, or security groups. For more information about how Launch Wizard creates security groups and how they are configured, see [Security groups in AWS Launch Wizard for SAP \(p. 73\)](#).

### SAP transport group configuration

You can create an SAP transport file system, or attach an existing transport file system that was created as part of a previous deployment with AWS Launch Wizard. Transport file systems are created with Amazon Elastic File System. For more information, see [Amazon Elastic File System setup for transport directory \(p. 44\)](#).

## Related services

The following AWS services are used when you deploy an SAP application with AWS Launch Wizard.

### Services

- [AWS CloudFormation \(p. 41\)](#)
- [Amazon Virtual Private Cloud security groups \(p. 41\)](#)
- [Amazon Elastic File System \(p. 42\)](#)
- [Amazon EC2 Systems Manager \(p. 42\)](#)
- [Amazon Simple Notification Service \(SNS\) \(p. 8\)](#)
- [Amazon Route 53 \(p. 42\)](#)
- [AWS Backint Agent for SAP HANA \(p. 42\)](#)
- [AWSTOE \(p. 42\)](#)

## AWS CloudFormation

[AWS CloudFormation](#) is a service that helps you model and set up your AWS resources, and lets you spend more time focusing on your applications that run in AWS. You create a template that describes all of the AWS resources that you want (for example, Amazon EC2 instances or Amazon RDS DB instances), and AWS CloudFormation takes care of provisioning and configuring those resources for you. With AWS Launch Wizard for SAP, you don't need to build AWS CloudFormation templates to deploy your application. Instead, AWS Launch Wizard combines infrastructure provisioning and application configuration (code that runs on EC2 instances to configure the application) into a unified AWS CloudFormation template. The AWS CloudFormation template is then invoked by AWS Launch Wizard's backend service to provision an application in your account.

## Amazon Virtual Private Cloud security groups

[Amazon Virtual Private Cloud security groups](#) act as a virtual firewall for your instance to control inbound and outbound traffic. When you launch an instance in a VPC, you can assign up to five security

groups to the instances. AWS Launch Wizard displays the security groups that will be assigned to the EC2 instances that run the SAP applications. This allows the components to communicate.

## Amazon Elastic File System

[Amazon EFS](#) provides file storage in the AWS Cloud. With Amazon EFS, you can create a file system, mount the file system on an Amazon EC2 instance, and then read and write data to and from your file system. For more information, see [Amazon Elastic File System setup for transport directory \(p. 44\)](#).

## Amazon EC2 Systems Manager

[Amazon EC2 Systems Manager](#) is an AWS service that you can use to view and control your infrastructure on AWS. Using the Amazon EC2 Systems Manager console, you can view operational data from multiple AWS services and automate operational tasks across your AWS resources. SSM helps you maintain security and compliance by scanning your managed instances and reporting on, or taking corrective action on, any policy violations that it detects.

## Amazon Simple Notification Service (SNS)

[Amazon Simple Notification Service \(SNS\)](#) is a highly available, durable, secure, fully managed pub/sub messaging service that provides topics for high-throughput, push-based, many-to-many messaging. Using Amazon SNS topics, your publisher systems can fan out messages to a large number of subscriber endpoints and send notifications to end users using mobile push, SMS, and email. You can use SNS topics for your Launch Wizard deployments to stay up-to-date on deployment progress. For more information, see the [Amazon Simple Notification Service Developer Guide](#).

## Amazon Route 53

[Amazon Route 53](#) 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. Launch Wizard integrates with Route 53 hosted zones, which are containers for records. The records contain information about how you want to route traffic for a specific domain, such as example.com, and its subdomains (acme.example.com, zenith.example.com). There are two types of hosted zones: public and private hosted zones. We recommend that you use private hosted zones for SAP applications unless an application must be directly accessible from the internet.

## AWS Backint Agent for SAP HANA

[AWS Backint Agent for SAP HANA](#) is an SAP-certified backup and restore application for SAP HANA workloads running on Amazon EC2 instances in the cloud. AWS Backint Agent runs as a standalone application that integrates with your existing workflows to back up your SAP HANA database to Amazon S3 and to restore it using SAP HANA Cockpit, SAP HANA Studio, and SQL commands. AWS Backint Agent supports full, incremental, and differential backup of SAP HANA databases. Additionally, you can back up log files and catalogs to Amazon S3. AWS Backint Agent runs on an SAP HANA database server, where backups and catalogs are transferred from the SAP HANA database to the AWS Backint Agent. The AWS Backint Agent stores your files in the S3 bucket that is specified in the agent configuration file. To restore your SAP HANA database server, SAP HANA reads the catalog files stored in your S3 bucket using the AWS Backint Agent. It then initiates a request to restore the required files from S3.

## AWSTOE

[AWSTOE](#) is component management application used to orchestrate complex workflows, modify system configurations, and test your systems without writing code. This application uses a declarative document schema. As a standalone application it does not require additional server setup. It can run on any cloud infrastructure and on premises. AWS Launch Wizard uses this application to orchestrate the download of the pre- and post-configuration scripts, and to run them.

## How AWS Launch Wizard for SAP works

AWS Launch Wizard provisions and configures the infrastructure required to run SAP HANA database- and Netweaver-based SAP applications on AWS. You select the SAP deployment pattern and provide the specifications, such as operating system, instance size, and vCPU/memory. Or, Launch Wizard can make these selections for you according to [SAP Standard Application Benchmarks \(SAPS\)](#). You have the option to manually choose the instance. Based on your selections, Launch Wizard automatically provisions the necessary AWS resources in the cloud.

Launch Wizard recommends Amazon EC2 instances by evaluating the [SAPS](#) or vCPU/memory requirements against the performance of Amazon EC2 instances supported by AWS. When new EC2 instances are released and certified for SAP, the sizing feature of Launch Wizard will take them into consideration when proposing recommendations.

Launch Wizard maintains a mapping rule engine built on the list of certified EC2 instances that are supported by SAP. When you enter your vCPU/memory or SAPS requirements, Launch Wizard recommends an Amazon EC2 instance that is certified for SAP workloads and offers performance that is no less than your input requirements. For certain workloads, such as SAP HANA in a production environment, Launch Wizard recommends instances based on the official SAP recommendations for SAP HANA database workloads. For workloads in a non-production environment, Launch Wizard recommends Amazon EC2 instances that meet SAP recommended requirements; however, the recommended instances are not enforced. You can change the instance types after deployment, or you can override the recommendation by making manual selections.

In addition to launching instances based on the SAP system information that you provide, such as SAP System Number and SAP System Identifier (SAP SID), Launch Wizard performs the following operations:

- Configures the operating system
- Configures hostname
- Attaches security groups so that the systems in the cluster that use the same configuration template, and also external systems, can communicate with the SAP systems that will be deployed on these instances.

Launch Wizard provides an estimated cost of deployment. You can modify your resources and instantly view an updated cost assessment. After you approve the deployment, Launch Wizard validates the inputs and flags inconsistencies. After you resolve the inconsistencies, Launch Wizard provisions and configures the resources. The result is a ready-to-use SAP application.

Launch Wizard creates a CloudFormation stack according to your infrastructure needs. For more information, see [Working With Stacks](#) in the *AWS CloudFormation User Guide*.

## Implementation Details

AWS Launch Wizard implements SAP deployments as follows.

### Deployment aspects

- [Storage for SAP Systems](#) (p. 44)
- [Amazon Elastic File System setup for transport directory](#) (p. 44)
- [Amazon Elastic File System setup for SAP Central Services instances configured for high availability](#) (p. 45)
- [Bring your own image \(BYOI\)](#) (p. 45)
- [Configuration settings](#) (p. 45)
- [Custom deployment configuration scripts](#) (p. 46)
- [Manual cleanup activities](#) (p. 47)

- [Default Quotas \(p. 47\)](#)
- [AWS Regions and Endpoints \(p. 47\)](#)

## Storage for SAP Systems

Storage capacity and performance are key aspects of any SAP system installation. Launch Wizard provides storage type options for the SAP Netweaver Application tier and the SAP HANA database tiers.

Amazon Elastic Block Store (Amazon EBS) volumes are included in the architecture to provide durable, high-performance storage. Amazon EBS volumes are network-attached disk storage, which you can create and attach to EC2 instances. When attached, you can create a file system on top of these volumes, run a database, or use them in any way that you would use a block device. 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.

[General Purpose EBS Volumes](#) offer storage for a broad range of workloads. These volumes deliver single-digit millisecond latencies and the ability to burst to 3,000 IOPS for extended periods of time. Between a minimum of 100 IOPS (at 33.33 GiB and below) and a maximum of 16,000 IOPS (at 5,334 GiB and above), baseline performance scales linearly at 3 IOPS per GiB of volume size.

[Provisioned IOPS Amazon EBS volumes](#) offer storage with consistent and low-latency performance. They are backed by solid state drives (SSDs) and designed for applications with I/O intensive workloads, such as databases. Amazon EBS-optimized instances, such as the R4 instance type, deliver dedicated throughput between Amazon EC2 and Amazon EBS.

By default, Launch Wizard deploys Amazon EBS volumes for the SAP HANA database that meet the storage KPIs for SAP as listed in [Storage Configurations for SAP HANA](#).

For Netweaver database stacks, you can choose between a `gp2` or `io1` volume for the `usr/sap/SAPSID` file system, whereas other configurations are deployed with `gp2` volumes.

In an SAP landscape, development occurs in the development system and is then imported into the QA and follow-on systems. For this import to occur successfully, a shared file system is required for SAP systems in the landscape. Amazon EFS is used to create the SAP Transport file system that is shared between multiple SAP systems in the landscape.

## Amazon Elastic File System setup for transport directory

The SAP transport directory is a shared file system between SAP systems (for example, Development, Quality, and Production) that are part of the same SAP Transport Domain for releasing and importing SAP transports. To avoid a single point of failure, Launch Wizard creates a file system with Amazon Elastic File System or reuses existing file systems. It mounts the file systems on the SAP systems that you select based on the role of the system. The transport file system is mounted on all of the applications servers included in the deployment.

When systems within the same SAP Transport Domain are created in one VPC and need to be attached to SAP systems in other VPCs (for example, if Development and Quality are deployed in a VPC tagged as `Non_Prod`, and Production is deployed in a VPC tagged as `Prod`), a prerequisite for using VPC Peering/ Transit Gateway is that you must enable the VPCs to be able to communicate. This allows Launch Wizard to attach the transport directory created in one VPC to instance(s) in other VPCs using a mount target in the same Availability Zone or other Availability Zones, as applicable. If the VPCs are not permitted to communicate, then the deployment will fail when it attempts to mount the transport file system created in one VPC to systems in another VPC.

### Note

When a transport files system is created with Amazon Elastic File System, Launch Wizard considers it a shared resource and will not delete it when you delete the deployment or if the deployment is rolled back.



## Amazon Elastic File System setup for SAP Central Services instances configured for high availability

The SAP Central Services instances that make up a Netweaver high availability deployment, ABAP Central Server (ASCS) and Enqueue Replication Server (ERS) instances, must contain the following file systems to be highly available: `/sapmnt`, `/usr/sap<SAPSID>/ASCS<XX>`, and `/usr/sap/<SAPSID>/ERS<XX>`. These file systems are built with Amazon EFS to avoid a single point of failure for the SAP system. Launch Wizard creates these file systems for the Netweaver high availability pattern using a single Amazon Elastic File System.

The following table contains information about how a single Amazon EFS is configured and mounted on an ASCS, ERS, Primary Application Server (PAS), and Additional Application Server (AAS).

EFS ID	EFS DNS name	Instance mounted on	File System name	Server mounted on
fs-123A456B	fs-123A456B.efs.<AWSRegion>.amazonaws.com:/	fs-123A456B.efs.<AWSRegion>.amazonaws.com:/SAPMNT-<SAPSID>	/sapmnt	SAP ASCS, ERS, Primary and Additional Application servers
fs-123A456B	fs-123A456B.efs.<AWSRegion>.amazonaws.com:/	fs-123A456B.efs.<AWSRegion>.amazonaws.com:/ASCS-<SAPSID>	/usr/sap/<SAPSID>/ASCS<XX>	SAP ASCS Server
fs-123A456B	fs-123A456B.efs.<AWSRegion>.amazonaws.com:/	fs-123A456B.efs.<AWSRegion>.amazonaws.com:/ERS-<SAPSID>	/usr/sap/<SAPSID>/ERS<XX>	SAP ERS Server

## Bring your own image (BYOI)

You can bring your own images to deploy and configure EC2 instances for SAP with AWS Launch Wizard. During launch, in order to continue with a deployment, Launch Wizard verifies whether the operating system version selected on the front end matches the operating system version of the instance. If the versions do not match, the deployment fails with an error.

When building your own image, consider the following:

- Launch Wizard configures the operating systems with OS-level parameters and utilities required by SAP
- Refer to SAP installation documents to ensure that operating system prerequisites are in place so that Launch Wizard deployments do not fail.
- Launch Wizard accesses standard repositories provided by OS vendors. Do not block access to them.
- Deployments by Launch Wizard use OS utilities and programs, such as zipper, yum, grep, printf, awk, sed, autofs, python, saptune, and tuned-profiles in the deployment script to configure SAP application and database servers. We recommend that you do not delete standard utilities.

## Configuration settings

The following configuration settings are applied when deploying an SAP application with Launch Wizard.

Setting	Applies to
SSM Agent	All SAP systems and patterns
EBS Volumes for SAP application tier	All SAP systems and patterns
EBS Volumes for SAP HANA database, log and backup file systems	All SAP systems and patterns
EFS volumes for SAP transport file systems	All SAP systems and patterns
EFS volumes for SAP central services: sapmnt, /usr/sap/<SID>/ASCS<XX> and /usr/sap/<SID>/ERS<XX>	ASCS and ERS systems
OS parameters required based on the operating system chosen for SAP HANA	All SAP systems and patterns
Security groups created and assigned for accessing the SAP system	All SAP systems and patterns
SSM Session Manager to remotely access the server for administrator activities	All SAP systems and patterns
Time zone settings at the OS level	All SAP systems and patterns

## Custom deployment configuration scripts

You can use custom shell scripts during the pre-deployment and post-deployment configuration phases. You provide the scripts stored on Amazon S3 or locally. During provisioning, Launch Wizard installs the AWSTOE application. When there are custom scripts to run, Launch Wizard creates an AWSTOE document that downloads the scripts from the location specified and then runs the scripts. The success of the custom scripts is a customer responsibility. Check the CloudWatch log streams for detailed execution logs or failure information after the scripts are deployed.

The number of configuration scripts you can use depends on the deployment model. For SAP HANA deployments, you can use one script, which runs on all of the HANA instances (both primary and worker nodes). For Netweaver stack on SAP HANA database, the following script limits apply:

- *Netweaver stack on SAP HANA single-instance deployment* — Because all tiers are installed on the same database instance, you can use only one script.
- *Netweaver stack on SAP HANA distributed-instance deployment* — You can use one script per each instance tier selected, including for ABAP System Central Services (ASCS) Server and Primary Application Server (PAS), Database (DB) Server, and Additional App Servers (AAS).
- *Netweaver stack on SAP HANA high availability deployment* — You can use one script per each instance tier selected, including for Primary Application Server (PAS), ABAP System Central Services (ASCS) Server, Database (DB) Server, Additional App Servers (AAS), and Enqueue Replication Server (ERS).

### Pre-deployment configuration scripts

Pre-deployment configuration scripts run after the instances are launched and the baseline Launch Wizard configuration tasks, such as deploying Amazon CloudWatch, Amazon EC2 Systems Manager agents, and the AWS CLI, are complete. If you want to run multiple pre-deployment configuration scripts, Launch Wizard runs them in parallel on each EC2 instance in the order in which they are specified. Pre-deployment configuration scripts can be used to perform tasks such as OS hardening or deploying security and logging software. The maximum runtime for all pre-deployment configuration scripts on a single EC2 instance is 45 minutes.

### Post-deployment configuration scripts

Post-deployment configuration scripts run when Launch Wizard completes configuration tasks specific to the application on all of the instances in a deployment. Before the provisioning process completes, post-configuration scripts run on all of the specified instance tiers. Launch Wizard uses SSM and AWS Lambda to trigger running post-deployment scripts on all selected SAP instances in the order in which they are specified. They can be used to perform tasks such as installing monitoring and management software, and for updating your DNS with entries for the newly deployed SAP servers and the domains joining them. The maximum runtime for all post-deployment configuration scripts on a single instance is 2 hours.

### Manual cleanup activities

If you choose to delete a deployment, or a deployment fails during the deployment phase and rolls back, Launch Wizard deletes the Amazon EC2 and Amazon EBS volumes that it launches as part of the deployment. It also removes the AWSTOE application. The following resources are considered shared resources and are created without the deletion flag.

- The Amazon Elastic File System file system that is created for the SAP transport files system `/usr/sap/trans`
- Security groups that you create

These resources must be verified manually to ensure that they are not being used by other systems in the landscape. They must then be manually deleted from either the Amazon Elastic File System or Amazon EC2 consoles, or using APIs.

### Default Quotas

To view the default quotas for AWS Launch Wizard, see [AWS Launch Wizard Endpoints and Quotas](#).

### AWS Regions and Endpoints

To view the service endpoints for AWS Launch Wizard, see [AWS Launch Wizard Endpoints and Quotas](#).

## Get started with AWS Launch Wizard for SAP

This topic contains information to help you set up your environment and deploy AWS resources with Launch Wizard, such as:

- How to create an IAM policy and attach it to your IAM user identity
- Configuration settings to apply to your environment
- How to deploy an SAP application from the AWS Management Console

### Contents

- [Set up for AWS Launch Wizard for SAP \(p. 47\)](#)
- [Access and deploy an SAP application with AWS Launch Wizard \(p. 50\)](#)

## Set up for AWS Launch Wizard for SAP

This section describes the prerequisites that you must verify to deploy an SAP application with AWS Launch Wizard.

## General prerequisites

The following general prerequisites must be met to deploy an application with Launch Wizard.

- You must create an Amazon VPC that consists of private subnet(s) in a minimum of two Availability Zones. The subnets must have outbound internet access. For more information on how to create and set up a VPC, see [Getting Started with Amazon VPC](#) in the *Amazon VPC User Guide*.
- You must create an IAM user and attach the **AmazonLaunchWizard\_Fullaccess** policy. See the [following sections \(p. 48\)](#) for the steps to attach the policy to the IAM user.
- To run custom pre- and post-configuration deployment scripts, you must add the permissions listed in [Add permissions to run custom pre- and post-deployment configuration scripts \(p. 49\)](#) to the `AmazonEC2RoleForLaunchWizard` role.
- If you want to install an SAP HANA database, you must download the software from the SAP Software Download page and upload it to an Amazon S3 bucket. For steps on how to download the software and upload it to an Amazon S3 bucket, see [Make SAP HANA software available for AWS Launch Wizard to deploy a HANA database \(p. 70\)](#).
- Depending on the operating system version you want to use for the SAP deployment, an SAP Marketplace subscription may be required. For a complete list of supported operating system versions, see [Supported operating system versions for SAP deployments \(p. 72\)](#).

## AWS Identity and Access Management (IAM)

Establishing the AWS Identity and Access Management (IAM) role and setting up the IAM user for permissions are typically performed by an **IAM administrator** for your organization. The steps are as follows:

- A one-time creation of IAM roles that Launch Wizard uses to deploy SAP systems on AWS.
- The creation of IAM users who can grant permission for Launch Wizard to deploy applications.

### One-time creation of IAM role

On the **Choose Application** page of Launch Wizard, under **Permissions**, Launch Wizard displays the IAM roles required for Launch Wizard to access other AWS services on your behalf. These roles are **AmazonEC2RoleForLaunchWizard** and **AmazonLambdaRoleForLaunchWizard**. Select **Next**, and Launch Wizard attempts to discover the IAM roles in your account. If either role does not exist in your account, Launch Wizard attempts to create the roles with the same names, **AmazonEC2RoleForLaunchWizard** and **AmazonLambdaRoleForLaunchWizard**.

The **AmazonEC2RoleForLaunchWizard** role is comprised of two IAM managed policies: **AmazonSSMManagedInstanceCore** and **AmazonEC2RolePolicyForLaunchWizard**. The **AmazonEC2RoleForLaunchWizard** role is used by the instance profile for the Amazon EC2 instances that Launch Wizard launches into your account as part of the deployment.

If you want to deploy AWS Backint Agent as a backup and restore solution for your application, you must attach a policy to the **AmazonEC2RoleForLaunchWizard** so that Launch Wizard can perform Backint Agent operations on your behalf. The required policy and instructions can be found in [Step 2 of the Backint Agent IAM documentation](#). During a deployment, Launch Wizard provides the policy as well as the steps to update the role, taking user specifications into account.

The **AmazonLambdaRoleForLaunchWizard** is also comprised of two IAM managed policies: **AWSLambdaVPCLambdaAccessExecutionRole** and **AmazonLambdaRolePolicyForLaunchWizardSAP**. The **AmazonLambdaRoleForLaunchWizard** role is used by the Lambda function invoked by the service in your account to perform certain deployment-related actions, such as validation of route tables, and pre-configuration and configuration tasks for HA mode enabling.

After the IAM roles are created, the IAM administrator can either continue with the deployment process or optionally delegate the application deployment process to another IAM user, as described in the following section. At this point in the IAM set up process, the IAM administrator can exit the Launch Wizard service.

## Create and enable IAM users to use Launch Wizard

To deploy an SAP system with Launch Wizard, the user must be assigned the **AmazonLaunchWizard\_Fullaccess** policy. The following steps guide the IAM administrator through the process of attaching an IAM policy to an IAM user to grant that user permission to access and deploy applications from Launch Wizard.

1. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the left navigation pane, choose **Policies**.
3. Choose **Users**.
4. Choose the name of the user whose permissions you want to modify, and then choose the **Permissions** tab.
5. Choose **Add permissions**.
6. Choose **Attach existing policies directly to user**.
7. Search for the policy named **AmazonLaunchWizard\_Fullaccess** and select the check box to the left of the policy name.
8. Choose **Next: Review** to see the list of policies to attach to the user.
9. Verify that the correct policy is listed, and then choose **Add permissions**.

### Important

You must log in with the user associated with this IAM policy when you use Launch Wizard.

## Add permissions to run custom pre- and post-deployment configuration scripts

To run custom pre- and post-configuration deployment scripts, you must add the following permissions to the `AmazonEC2RoleForLaunchWizard` role. The following steps guide you through the process of adding the required permissions for using custom scripts to the `AmazonEC2RoleForLaunchWizard` role.

1. Sign in to the AWS Management Console and open the IAM console at <https://console.aws.amazon.com/iam/>.
2. In the left navigation pane, choose **Policies**.
3. Choose **Users**.
4. Choose the name of the user whose permissions you want to modify, and then choose the **Permissions** tab.
5. Choose **Add inline policy**.
6. Copy and paste the following policy into the **JSON** tab. Enter the S3 paths where your scripts are stored.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": [
        "s3:GetObject",
        "s3:GetBucketLocation"
      ]
    }
  ]
}
```

```
    ],  
    "Resource": [  
        "arn:aws:s3:::<S3bucket1>/<S3prefix1>/<script1>",  
        "arn:aws:s3:::<S3bucket2>/<S3prefix2>/<script2>",  
        "arn:aws:s3:::<S3bucket1>",  
        "arn:aws:s3:::<S3bucket2>"  
    ]  
  }  
]  
}
```

7. Choose **Review policy** and enter a name for the policy.
8. Choose **Create Policy**.
9. Verify that the correct policy is listed, and then choose **Policy actions**.
10. Choose **Attach**.
11. Search for the policy named **AmazonEC2RoleForLaunchWizard** and select the check box to the left of the policy name.
12. Choose **Attach policy**.

### Important

You must log in with the user associated with this IAM policy when you use Launch Wizard.

If the pre- or post-deployment configuration deployment scripts are expected to run additional AWS services, the permissions to use the services must also be manually added as policy to the `AmazonEC2RoleForLaunchWizard`.

## Access and deploy an SAP application with AWS Launch Wizard

### Access AWS Launch Wizard

You can launch AWS Launch Wizard from the AWS Launch Wizard console located at <https://console.aws.amazon.com/launchwizard>.

### Deploy an SAP application with AWS Launch Wizard

The following steps guide you through deploying an SAP application with AWS Launch Wizard after you have launched it from the console.

#### Topics

- [Create a deployment \(p. 50\)](#)
- [Define infrastructure \(p. 51\)](#)
- [Application and deployment settings \(p. 54\)](#)

### Create a deployment

1. From the AWS Launch Wizard landing page, choose **Create deployment**.
2. Choose **SAP**.
3. Under **Permissions**, Launch Wizard displays the AWS Identity and Access Management (IAM) roles required for Launch Wizard to access other AWS services on your behalf. For more information about these roles and setting up IAM for Launch Wizard, see [Identity and Access Management for AWS Launch Wizard \(p. 77\)](#). Choose **Next**.

## Define infrastructure

On the **Define infrastructure** page, define your deployment name and infrastructure settings.

1. Under the **General** subheading, define the following:
  - **Deployment name.** Enter a unique application name for your deployment.
  - **Description (Optional).** Provide an optional description of your deployment.
  - **Tags (Optional).** Enter a key and value to assign metadata to your deployment. For help with tagging, see [Tagging Your Amazon EC2 Resources](#).
2. Under the **Infrastructure – SAP landscape** subheading, configure the following infrastructure settings for your SAP landscape.

Configuration options

- Under **Configuration type**, choose whether to **Create new configuration** or **Apply saved configuration**.
- Enter the following information:
  - **Configuration name.** Enter a name or short description to identify your configuration. You can save this configuration for future use.
  - **Deployment environment.** (**Create new configuration**, only) Choose whether to deploy into a **Production** or **Non-Production** environment.

Configuration details

If you choose to create a new configuration, enter the following information.

- **Key pair name.** Choose an existing key pair from the dropdown list or select the link to create one. If you select **Create new key pair name**, you are directed to the Amazon EC2 console. From the Amazon EC2 console, under **Network and Security**, choose **Key Pairs**. Choose **Create a new key pair**, enter a name for the key pair, and then choose **Download Key Pair**.

### Important

This is the only time that you can save the private key file, so download and save it in a safe place. You must specify the name of your key pair when you launch an instance, and provide the corresponding private key each time that you connect to the instance.

Return to the Launch Wizard console, and choose the refresh button next to the **Key Pair name** dropdown list. The new key pair appears in the dropdown list. For more information about key pairs, see [Amazon EC2 Key Pairs](#).

- **Virtual Private Cloud.** Choose a VPC from the dropdown list or select the **Create VPC** link. If you select **Create VPC**, you are redirected to the VPC console to create a VPC.
- **Availability Zone (AZ) and private subnets.** You can deploy into one or two Availability Zones (AZs) with up to two private subnets per Availability Zone. Different requirements are needed for different systems in the landscape. You must select two Availability Zones with a required primary and optional secondary subnet for each Availability Zone. These selections are used for each deployment according to the deployment model that you selected.

From the dropdown lists, choose the **Availability Zones** within which you want to deploy your SAP systems and choose the private subnets. The private subnets must have outbound connectivity to the internet and to other AWS services, such as Amazon S3, AWS CloudFormation, and CloudWatch Logs. They must also be able to access the Linux repositories required for instance configuration.

For high availability deployments, the following subnets must share a common route table:

- subnet 1 in Availability Zone 1 and subnet 1 in Availability Zone 2
- subnet 2 in Availability Zone 1 and subnet 2 in Availability Zone 2

### To create a private subnet

- If a subnet doesn't have a route to an internet gateway, the subnet is known as a private subnet. Use the following procedure to create a private subnet. We recommend that you enable the outbound connectivity for each of your selected private subnets using a NAT gateway. To enable outbound connectivity from private subnets with public subnets, [create a NAT Gateway](#) in your chosen public subnet. Then, follow the steps in [Updating Your Route Table](#) for each of your private subnets.
  - Follow the steps in [Creating a Subnet](#) in the Amazon VPC User Guide using the existing VPC that you will use in Launch Wizard.
  - When you create a VPC, it includes a main route table by default. On the **Route Tables** page in the Amazon VPC console, you can view the main route table for a VPC by looking for **Yes** in the **Main** column. The main route table controls the routing for all subnets that are not explicitly associated with any other route table. If the main route table for your VPC has an outbound route to an internet gateway, then any subnet created using the previous step, by default, becomes a public subnet. To ensure the subnets are private, you may need to create separate route tables for your private subnets. These route tables must not contain any routes to an internet gateway. Alternatively, you can create a custom route table for your public subnet and remove the internet gateway entry from the main route table.
- **Verify Connectivity.** Select the check box to verify that your private subnets have outbound internet connectivity.
- **Security groups.** You can choose already existing security groups or Launch Wizard can create security groups that will be assigned to the EC2 instances that Launch Wizard deploys. If you choose already existing security groups, you must ensure that all of the necessary ports required to access the SAP and SAP HANA databases are open. If you choose to allow Launch Wizard to create the security groups, the security groups are created to enable the components of the cluster to communicate. Systems that are deployed with the same configuration settings can also communicate.

If you choose an existing security group, Launch Wizard displays the security groups that will be assigned to the EC2 instances that Launch Wizard deploys. This enables the components to communicate and systems that are deployed with the same configuration settings to communicate.

- **Connectivity to external systems or users.** If you allowed Launch Wizard to create the security groups, then choose the **Connection type** and **Value** of the IP address or security groups required to access the SAP systems. These values can be a network segment from which the end users access the SAP systems, or downstream/upstream systems assigned a different security group in AWS or on premises.
- **Proxy setting.** During the launch process, the deployed Amazon EC2 instances require outbound internet access in order to:
  - Access the operating system (SUSE/RHEL) repositories.
  - Access AWS services, such as Amazon S3, CloudWatch and Systems Manager.

An [internet gateway](#) is typically configured for outbound internet access. If you want to route internet traffic through a proxy server, enter the proxy server details. When proxy server information is provided, Launch Wizard will make the necessary environment changes to the EC2 instances during launch so that outbound internet traffic is routed through the proxy server.

- **PROXY.** Enter the proxy server name and port, for example `http://10.0.0.140:3128` or `https://10.0.0.140.3128`.



- **NO\_PROXY.** When a proxy server is used for outbound communication, the `NO_PROXY` environment variable is used to route traffic without using the proxy for the following types of traffic:
  - local communication
  - traffic to other instances within the VPC
  - traffic to other AWS services for which VPC endpoints are created

Enter a list of comma-separated values to denote hostnames, domain names, or a combination of both.

We recommend that you add all AWS service endpoints (if defined) to the `NO_PROXY` environment variable so that a private connection between the VPC and the service endpoint can be established in the AWS VPN. For more information on AWS service endpoints, see [AWS service endpoints](#).

`NO_PROXY` is an optional parameter. If no value is entered, the following default URLs are added to the environment. Values entered for `NO_PROXY` at a later time are added to this list.

```
NO_PROXY="localhost,127.0.0.1,169.254.169.254,.internal,{VPC_CIDR_RANGE}"
```

#### Default `NO_PROXY` URL details

- **localhost**—loopback hostname
- **127.0.0.1**—loopback adapter IP
- **169.254.169.254**—EC2 metadata link-local address
- **.internal**—default DNS for the VPC
- **{VPC\_CIDR\_RANGE}**—CIDR block of the VPC, for example, 10.0.0/24
- **Time zone.** Choose the time zone settings to configure the timezone on the instances from the dropdown list.
- **EBS encryption.** From the dropdown list, choose whether or not to enable EBS encryption for all of the EBS volumes that are created for the SAP systems. For more information, see [Amazon EBS Encryption](#).
- **Domain name (DNS) settings (Optional).** Select **Domain Name** or **Route 53** from the **DNS type** dropdown list.
  - If you select **Domain Name**, you have the option to enter a domain name to maintain a Fully Qualified Domain Name (FQDN) in the `/etc/hosts` file for each instance that is launched and configured by Launch Wizard.
  - If you select **Route 53**, select a Route 53 hosted zone from the dropdown list. Launch Wizard will create a DNS entry for each EC2 instance launched.

#### Note

Before you use a Route 53 hosted zone, verify that the hosted zone is [integrated with the VPC](#), and that the [VPC DHCP options](#) are correctly set up.

- **SAP landscape settings.** Enter the system settings for your SAP landscape.
  - **SAP System Admin User ID.** Enter the user ID for the SAP system administrator.
  - **SAP System Admin Group ID.** Enter the group ID for SAPSYS. We recommend that you replicate this number across all of your SAP systems because SAPSYS GID must be the same between systems that are part of the transport domain.
  - **SAPINST Group ID.** Enter the group ID for the SAPINST.
- **Simple Notification Service (SNS) topic ARN (Optional).** Specify an SNS topic where Launch Wizard can send notifications and alerts. For more information, see the [Amazon Simple Notification Service Developer Guide](#). You can also choose **Create SNS topic** and then create

one in the Amazon SNS console. After you create an SNS topic, you can enter it in the Launch Wizard SNS field.

- After you specify the infrastructure settings, choose **Next**.

## Application and deployment settings

The following steps show the deployment paths for **Netweaver stack on SAP HANA database** and **SAP HANA database**. Please follow the deployment steps for your deployment path.

### Topics

- [Netweaver stack on SAP HANA database \(p. 54\)](#)
- [SAP HANA database \(p. 63\)](#)

### Netweaver stack on SAP HANA database

#### Application settings

On the **Configure application settings** page, enter your Netweaver stack on SAP HANA database application settings.

1. **Application type.** Select **Netweaver stack on SAP HANA database**. This configuration includes:
  - Netweaver stack for a single instance , distributed instance, or multi-AZ for high availability (HA) deployment.
  - EC2 instances for the Netweaver application tier
  - EC2 instances for SAP HANA database and optional SAP HANA database install
2. **General settings – SAP system.** Enter the settings for your SAP system.
  - **SAP System ID (SAPSID).** An identifier for your system. The ID must be a three character, alphanumeric string.
  - **EBS Volume Type for Netweaver application stack instances.** Choose which volume type to use for the NW application file system `/usr/sap/SAPSID` from the dropdown list.
  - **Transport Domain Controller.** Specify whether the SAP system will be the domain controller for the SAP landscape. If not, select the transport file system of the domain controller to be mounted.
3. **General Settings – SAP HANA.** Enter the settings for your SAP HANA installation.
  - **SAP HANA System ID.** Enter the identifier for your SAP HANA database. The ID must be a three character, alphanumeric string.
  - **SAP HANA Instance number.** Enter the instance number to be used for the SAP HANA installation and setup. The ID must be a two-digit number.
  - **EBS Volume Type for SAP HANA.** Select the EBS volume types to use for both **SAP HANA Data** and **SAP HANA Logs** from the dropdown lists.
  - **SAP HANA software install.** Choose whether to download the SAP HANA software.
    - If you choose **Yes**, enter the Amazon S3 location to store the SAP software files. The S3 bucket must have the prefix “launchwizard” in the bucket name to ensure that the Launch Wizard IAM role policy for EC2 has read-only access to the bucket. For steps to set up the folder structure for your S3 bucket, see [Make SAP HANA software available for AWS Launch Wizard to deploy a HANA database \(p. 70\)](#).
    - **AWS Backint Agent.** Select the check box if you want to deploy AWS Backint Agent for backup and recover along with the application. For more information about AWS Backint Agent, see [AWS Backint Agent for SAP HANA](#).

- **S3 URI.** Enter the URI of the S3 bucket where you want to store your SAP HANA backup files. For example, `s3://<bucket-name>`.
  - **S3 Encryption (KMS key ARN).** Select the ARN of the KMS key that AWS Backint Agent can use to encrypt the backup files stored in your Amazon S3 bucket.
  - **Agent version.** Select the version number of the agent that you want to install. If you do not enter a version number, the latest published version of the agent is installed.
  - **Additional Backint preferences.**
    - **AWS Launch Wizard provisions the local backup EBS volume.** Select the check box if you want Launch Wizard to provision the local backup EBS volume for storing HANA backups.
    - **AWS Launch Wizard to proceed with deployment when Backint installation fails.** Select the check box if you want Launch Wizard to continue with an application deployment when the Backint installation fails.
    - **Verify that you have attached the required policy for Backint operations to the following role.** Select this check box after you have attached the required policy to the `AmazonEC2RoleForLaunchWizard`. This policy allows Launch Wizard to perform Backint Agent operations on your behalf. The policy and instructions to attach the policy to the role are provided by Launch Wizard during deployment. This information can also be found in [Step 2 of the Backint Agent IAM documentation](#).
  - If you choose **No**, only AWS infrastructure is provisioned.
  - **S3 location for SAP HANA media - optional.** Enter the path for the S3 bucket in which you want to store SAP HANA media.
  - **SAP HANA password.** Enter a password for your SAP HANA installation.
4. After you enter your application settings, choose **Next**.

(Use the tab for **Single instance deployment**, **Distributed instance deployment**, or **High availability deployment**, depending on your configuration)

#### Single instance deployment

On the **Configure deployment model** page, enter the deployment details for a single instance deployment.

1. **Deployment details.** Launch Wizard supports single instance deployments, distributed instance deployments, and high availability deployments. Select **Single instance deployment**.
2. **ASCS, PAS, and DB on one EC2 instance.** Enter the deployment settings for your instance.
  - **Instance details.**
    - Under **Instance sizing**, choose whether to use **AWS/Marketplace/Community images** or **Bring your own images (BYOI)**.
    - **Operating System.** Select a supported operating system version for the ASCS instance. For a complete list of operating system versions supported for ASCS, see [Supported operating system versions for SAP deployments \(p. 72\)](#).
    - **AMI ID.** For BYOI, select the AMI that you want to use from the dropdown.
    - **Host name.** Enter the host name for the EC2 instance.
    - **Auto Recovery.** Auto recovery is an Amazon EC2 feature to increase instance availability. Select the check box to enable EC2 automatic recovery for the instance. For more information, see [Recover Your Instance](#) in the Amazon EC2 User Guide.
  - Under **Instance sizing**, choose whether to **Use AWS recommended resources** or **Choose instance**.
  - **Use AWS recommended resources.**
    - **Infrastructure requirements.** Choose the requirements for your recommended resources from the dropdown list.

- **Based on CPU/Memory.** If you select this option, enter the required number of vCPU **Cores** and **Memory**. Amazon EC2 supports up to 448 logical processors. If the amount of memory required exceeds 4TB, [dedicated hosts](#) are required.
- **SAPS (SAP Application Performance Standard).** If you select this option, enter the **SAPS** rating for the SAP certified instance types.
- **Choose your instance.**
  - **Instance type.** Choose the instance type from the dropdown list.
  - **Recommended Resources.** AWS Launch Wizard displays the **Estimated monthly cost of operation** based on your instance sizing selections and the EBS volumes that will be created and attached to the launched instances. This is an estimate of AWS costs to deploy additional resources and does not include any image costs, EC2 reservations, applicable taxes, or discounts.

### 3. Pre- and post-deployment configuration scripts — optional

You can run pre- and post-deployment configuration scripts during application provisioning. For more information about how Launch Wizard accesses and deploys these scripts, see [Custom deployment configuration scripts \(p. 46\)](#).

#### Pre-deployment configuration script — optional

- **Deployment settings.** In the event of a configuration script failure or time out, select whether to ignore all failures and proceed with the deployment. If you do not select this option, when the configuration scripts fail or time out, Launch Wizard will roll back the workload and delete all of the AWS resources created by Launch Wizard. Pre-deployment configuration scripts must finish running in 45 minutes or less.
- **Configuration script.** Choose to use a script located in Amazon S3 and enter the URL path of the script, or enter a script manually by uploading a script file. To remove the configuration script, choose **Remove script**.

#### Post-deployment configuration script — optional

- **Deployment settings.** In the event of a configuration script failure or time out, select whether to ignore all failures and proceed with the deployment. If you do not select this option, when the configuration scripts fail or time out, Launch Wizard will roll back the workload and delete all of the AWS resources created by Launch Wizard. Post-deployment configuration scripts must finish running in 2 hours or less.
- **Configuration script.** Choose to use a script located in Amazon S3 and enter the URL path of the script, or enter a script manually by uploading a script file. To remove the configuration script, choose **Remove script**.

### 4. After you have entered your deployment settings, choose **Next**.

(See the **Review** tab)

#### Distributed instance deployment

On the **Configure SAP HANA deployment model** page, enter the deployment details for a distributed instance deployment.

1. **Deployment details.** Launch Wizard supports single instance deployments, distributed instance deployments, and high availability deployments. Select **Distributed instance deployment**.
2. **ABAP System Central Services (ASCS) Server and Primary Application Server (PAS).** Enter the deployment settings for your instance.
  - **Instance details.**

- Under **Instance sizing**, choose whether to use **AWS/Marketplace/Community images** or **Bring your own images (BYOI)**.
  - **Operating System.** Select a supported operating system version for the ASCS and PAS instances. For a complete list of operating system versions supported for ASCS, see [Supported operating system versions for SAP deployments \(p. 72\)](#).
  - **AMI ID.** For BYOI, select the AMI that you want to use from the dropdown.
  - **Host name.** Enter the host name for the EC2 instances.
  - **Auto Recovery.** Auto recovery is an Amazon EC2 feature to increase instance availability. Select the check box to enable EC2 automatic recovery for the instance. For more information, see [Recover Your Instance](#) in the Amazon EC2 User Guide.
- Under **Instance sizing**, choose whether to Use **AWS recommended resources** or **Choose your instance**.
  - **Use AWS recommended resources.**
    - **Infrastructure requirements.** Choose the requirements for your recommended resources from the dropdown list.
      - **Based on CPU/Memory.** If you select this option, enter the required number of vCPU **Cores** and **Memory**. Amazon EC2 supports up to 448 logical processors. If the amount of memory required exceeds 4TB, [dedicated hosts](#) are required.
      - **SAPS (SAP Application Performance Standard).** If you select this option, enter the **SAPS** rating for the SAP certified instance types.
  - **Choose your instance.**
    - **Instance type.** Choose the instance type from the dropdown list.
    - **Recommended Resources.** AWS Launch Wizard displays the **Estimated monthly cost of operation** based on your instance sizing selections. This is an estimate of AWS costs to deploy additional resources and does not include any applicable taxes or discounts.
- 3. **Settings for Database (DB) Server.** Enter the deployment settings for your instance.
  - **Instance details.**
    - Under **Instance sizing**, choose whether to use **AWS/Marketplace/Community images** or **Bring your own images (BYOI)**.
      - **Operating System.** Select a supported operating system version for the ASCS and PAS instances. For a complete list of operating system versions supported for ASCS, see [Supported operating system versions for SAP deployments \(p. 72\)](#).
      - **AMI ID.** For BYOI, select the AMI that you want to use from the dropdown.
    - **Scale up and Scale out.** Select an upgrade strategy for your system hardware to upgrade for increased data and workload.
      - **Scale-up deployment.** If you choose this deployment upgrade model, enter the host name for the database
      - **Scale-out deployment.** If you choose this deployment upgrade model, enter the **SAP HANA master host name**, the **Number of worker nodes**, and the **Worker node hostname prefix** under **Instance sizing**.
  - Under **Instance sizing**, choose whether to **Use AWS recommended resources** or **Choose instance**.
    - **Use AWS recommended resources.**
      - **Define requirements.** Choose the requirements for your recommended resources from the dropdown list.
        - **Based on CPU/Memory.** If you select this option, enter the required number of vCPU **Cores** and **Memory**. Amazon EC2 supports up to 448 logical processors. If the amount of memory required exceeds 4TB, [dedicated hosts](#) are required.
        - **SAPS (SAP Application Performance Standard).** If you select this option, enter the **SAPS** rating for the SAP certified instance types.

- **Instance type.** Choose the instance type from the dropdown list.
  - **Auto Recovery.** Auto recovery is an Amazon EC2 feature to increase instance availability. Select the check box to enable EC2 automatic recovery for the instance. For more information, see [Recover Your Instance](#) in the Amazon EC2 User Guide.
  - **Recommended Resources.** AWS Launch Wizard displays the **Estimated monthly cost of operation** based on your instance sizing selections. This is an estimate of AWS costs to deploy additional resources and does not include any applicable taxes or discounts.
4. **Settings for Additional App Servers (AAS) - optional.** Enter the deployment settings for your AAS instances.
- **Instance details.**
    - **Number of Additional App Servers (AAS).** Enter the number of additional application servers.
    - **Naming convention for host name.** Enter the naming convention for the host name.
    - **Auto Recovery.** Auto recovery is an Amazon EC2 feature to increase instance availability. Select the check box to enable EC2 automatic recovery for the instance. For more information, see [Recover Your Instance](#) in the Amazon EC2 User Guide.
  - Under **Instance sizing**, choose whether to Use **AWS recommended resources** or **Choose your instance**.
    - **Use AWS recommended resources.**
      - **Define requirements.** Choose the requirements for your recommended resources from the dropdown list.
        - **Based on CPU/Memory.** If you select this option, enter the required number of vCPU **Cores** and **Memory**. Amazon EC2 supports up to 448 logical processors. If the amount of memory required exceeds 4TB, [dedicated hosts](#) are required.
        - **SAPS (SAP Application Performance Standard).** If you select this option, enter the **SAPS** rating for the SAP certified instance types.
    - **Choose your instance.**
      - **Instance type.** Choose the instance type from the dropdown list.
      - **Recommended Resources.** AWS Launch Wizard displays the **Estimated monthly cost of operation** based on your instance sizing selections. This is an estimate of AWS costs to deploy additional resources and does not include any applicable taxes or discounts.
5. **Pre- and post-deployment configuration scripts — optional**

You can run pre- and post-deployment configuration scripts during application provisioning. For more information about how Launch Wizard accesses and deploys these scripts, see [Custom deployment configuration scripts \(p. 46\)](#).

#### **Pre-deployment configuration script — optional**

- **Deployment settings.** In the event of a configuration script failure or time out, choose whether to proceed with the deployment. If you do not select this option, then when the configuration scripts fail or time out, Launch Wizard will roll back the workload and delete all of the AWS resources created by Launch Wizard. Pre-deployment configuration scripts must finish running in 45 minutes or less.
- **Configuration script.** You can add one or more configuration scripts depending on the number of servers you select to run scripts during the launch phase.
  - For each pre-deployment configuration script that you want to run, choose to use a script located in Amazon S3 and enter the URL path of the script, or upload a script file.
  - Select the servers to run the pre-deployment configuration scripts during the launch phase. You can choose to run pre-deployment scripts on **ABAP System Central Services (ASCS) Server and Primary Application Server (PAS), Database (DB) Server, and Additional App Servers (AAS)**. You can add a script for each server selected.

- To remove a configuration script, choose **Remove script**. To add more configuration scripts, choose **Add another script**.

### Post-deployment configuration script — optional

- **Deployment settings.** In the event of a configuration script failure, choose whether to proceed with the deployment. If you do not select this option, then when the configuration scripts fail or time out, Launch Wizard will roll back the workload and delete all of the AWS resources created by Launch Wizard. Post-deployment configuration scripts must finish running in 2 hours or less.
  - **Configuration script.** You can add one or more configuration scripts depending on the number of servers you select to run scripts during the post-deployment phase.
    - For each post-deployment configuration script that you want to run, choose to use a script located in Amazon S3 and enter the URL path of the script, or upload a script file.
    - Select the servers to run the post-deployment configuration scripts when an EC2 instance has been configured for use. You can choose to run the post-deployment scripts on **ABAP System Central Services (ASCS) Server and Primary Application Server (PAS), Database (DB) Server, and Additional App Servers (AAS)**. You can add a script for each server selected.
    - To remove a configuration script, choose **Remove script**. To add more configuration scripts, choose **Add another script**.
6. After you have entered your additional settings, choose **Next**.

(See the **Review** tab)

#### High availability deployment

On the **Configure SAP HANA deployment model** page, enter the deployment details for the high availability deployment.

1. **Deployment details.** Launch Wizard supports single instance deployments, distributed instance deployments, and high availability deployments. Select **High availability deployment**.
2. **Settings for ABAP System Central Services (ASCS) server.** Enter the deployment settings for your instance.
  - **Instance details.**
    - Under **Image type**, choose whether to use **AWS/Marketplace/Community images** or **Bring your own images (BYOI)**.
      - **Operating System.** Select a supported operating system version for the ASCS instances. For a complete list of operating system versions supported for ASCS, see [Supported operating system versions for SAP deployments \(p. 72\)](#).
      - **AMI ID.** For BYOI, select the AMI that you want to use from the dropdown.
    - **Host name.** Enter the host name for the EC2 instance.
    - **ASCS instance number.** Enter the instance number for the SAP installation and setup, and to open up ports for security groups.
  - Under **Instance sizing**, choose whether to Use **AWS recommended resources** or **Choose your instance**.
    - **Use AWS recommended resources.**
      - **Based on CPU/Memory.** If you select this option, enter the required number of vCPU **Cores** and **Memory**. Amazon EC2 supports up to 448 logical processors. If the amount of memory required exceeds 4TB, [dedicated hosts](#) are required.
      - **SAPS (SAP Application Performance Standard).** If you select this option, enter the **SAPS** rating for the SAP certified instance type.



- **Choose your instance.**
    - **Instance type.** Choose the instance type from the dropdown list.
  - **Recommended Resources.** AWS Launch Wizard displays the **Estimated monthly cost of operation** based on your instance sizing selections. This is an estimate of AWS costs to deploy additional resources and does not include any applicable taxes or discounts.
3. **Settings for Enqueue Replication Server (ERS).** Enter the deployment settings for your ERS.
- **Instance details.**
    - Under **Instance sizing**, choose whether to use **AWS/Marketplace/Community images** or **Bring your own images (BYOI)**.
      - **Operating System.** Select a supported operating system version for the ERS instance.
      - **AMI ID.** For BYOI, select the AMI that you want to use from the dropdown.
    - **Host name.** Enter the host name for the EC2 instance.
    - **ERS instance number.** Enter the instance number for the SAP installation and setup, and to open up ports for security groups.
  - Under **Instance sizing**, choose whether to Use **AWS recommended resources** or **Choose your instance**.
    - **Use AWS recommended resources.**
      - **Based on CPU/Memory.** If you select this option, enter the required number of vCPU **Cores** and **Memory**. Amazon EC2 supports up to 448 logical processors. If the amount of memory required exceeds 4TB, [dedicated hosts](#) are required.
      - **SAPS (SAP Application Performance Standard).** If you select this option, enter the **SAPS** rating for the SAP certified instance type.
    - **Choose your instance.**
      - **Instance type.** Choose the instance type from the dropdown list.
  - **Recommended Resources.** AWS Launch Wizard displays the **Estimated monthly cost of operation** based on your instance sizing selections. This is an estimate of AWS costs to deploy additional resources and does not include any applicable taxes or discounts.
4. **Settings for database (DB) Server.** Enter the deployment settings for your database.
- Under **Instance sizing**, choose whether to use **AWS/Marketplace/Community images** or **Bring your own images (BYOI)**.
    - **Instance details.**
      - **Operating System.** Select a supported operating system version for the ERS instance.
      - **AMI ID.** For BYOI, select the AMI that you want to use from the dropdown.
  - **Primary and secondary instance details.** Enter details for both the primary and secondary instances.
    - **SAP HANA host name.** Enter the host name for the SAP HANA primary and secondary instances.
    - **Server site name.** Enter the primary and secondary site name for the SAP HANA system replication.
  - **Overlay IP address.** Enter the overlay IP address to assign to the active node. The IP address should be outside of the VPC CIDR and must not be used by any other HA cluster. It is configured to always point to the active SAP HANA node.
  - **Pacemaker tag name.** Enter the tag to assign to each EC2 instance. This tag is used by the pacemaker component of SLES HAE and RHEL for SAP high availability solutions and must not be used by any other EC2 instance in your account.
  - Under **Instance sizing**, choose whether to Use **AWS recommended resources** or **Choose your instance**.
    - **Use AWS recommended resources.**



- **Based on CPU/Memory.** If you select this option, enter the required number of vCPU **Cores** and **Memory**. Amazon EC2 supports up to 448 logical processors. If the amount of memory required exceeds 4TB, [dedicated hosts](#) are required.
  - **SAPS (SAP Application Performance Standard).** If you select this option, enter the **SAPS** rating for the SAP certified instance type.
  - **Choose your instance.**
    - **Instance type.** Choose the instance type from the dropdown list.
  - **Recommended Resources.** AWS Launch Wizard displays the **Estimated monthly cost of operation** based on your instance sizing selections. This is an estimate of AWS costs to deploy additional resources and does not include any applicable taxes or discounts.
5. **Primary Application Server (PAS).** Enter the deployment settings for your instance.
- **Instance details.**
    - Under **Image type**, choose whether to use **AWS/Marketplace/Community images** or **Bring your own images (BYOI)**.
      - **Operating System.** Select a supported operating system version for the ERS instance.
      - **AMI ID.** For BYOI, select the AMI that you want to use from the dropdown.
    - **Host name.** Enter the host name for the EC2 instance.
    - **Auto Recovery.** Auto recovery is an Amazon EC2 feature to increase instance availability. Select the check box to enable EC2 automatic recovery for the instance. For more information, see [Recover Your Instance](#) in the Amazon EC2 User Guide.
  - Under **Instance sizing**, choose whether to Use **AWS recommended resources** or **Choose your instance**.
    - **Use AWS recommended resources.**
      - **Define requirements.** Choose the requirements for your recommended resources from the dropdown list.
        - **Based on CPU/Memory.** If you select this option, enter the required number of vCPU **Cores** and **Memory**. Amazon EC2 supports up to 448 logical processors. If the amount of memory required exceeds 4TB, [dedicated hosts](#) are required.
        - **SAPS (SAP Application Performance Standard).** If you select this option, enter the **SAPS** rating for the SAP certified instance types.
    - **Choose your instance.**
      - **Instance type.** Choose the instance type from the dropdown list.
    - **Recommended Resources.** AWS Launch Wizard displays the **Estimated monthly cost of operation** based on your instance sizing selections. This is an estimate of AWS costs to deploy additional resources and does not include any applicable taxes or discounts.
6. **Settings for Additional App Servers (AAS) - optional.** Enter the deployment settings for your AAS instances.
- **Instance details**
    - **Number of Additional App Servers (AAS).** Enter the number of additional application servers.
    - **Naming convention for host name.** Enter the naming convention for the host name.
    - **Auto Recovery.** Auto recovery is an Amazon EC2 feature to increase instance availability. Select the check box to enable EC2 automatic recovery for the instance. For more information, see [Recover Your Instance](#) in the Amazon EC2 User Guide.
  - Under **Instance sizing**, choose whether to Use **AWS recommended resources** or **Choose your instance**.
    - **Use AWS recommended resources.**
      - **Infrastructure requirements.** Choose the requirements for your recommended resources from the dropdown list.

- **Based on CPU/Memory.** If you select this option, enter the required number of vCPU **Cores** and **Memory**. Amazon EC2 supports up to 448 logical processors. If the amount of memory required exceeds 4TB, [dedicated hosts](#) are required.
- **SAPS (SAP Application Performance Standard).** If you select this option, enter the **SAPS** rating for the SAP certified instance types.
- **Choose your instance.**
  - **Instance type.** Choose the instance type from the dropdown list.
  - **Recommended Resources.** AWS Launch Wizard displays the **Estimated monthly cost of operation** based on your instance sizing selections. This is an estimate of AWS costs to deploy additional resources and does not include any applicable taxes or discounts.

## 7. Pre- and post-deployment configuration scripts — optional

You can run pre- and post-deployment configuration scripts during application provisioning. For more information about how Launch Wizard accesses and deploys these scripts, see [Custom deployment configuration scripts \(p. 46\)](#).

### Pre-deployment configuration script — optional

- **Deployment settings.** In the event of a configuration script failure or time out, choose whether to proceed with the deployment. If you do not select this option, then when the configuration scripts fail or time out, Launch Wizard will roll back the workload and delete all of the AWS resources created by Launch Wizard. Pre-deployment configuration scripts must finish running in 45 minutes or less.
- **Configuration script.** You can add one or more configuration scripts depending on the number of servers you select to run scripts during the launch phase.
  - For each pre-deployment configuration script that you want to run, choose to use a script located in Amazon S3 and enter the URL path of the script, or upload a script file.
  - Select the servers to run the pre-deployment configuration scripts during the launch phase. You can choose to run pre-deployment scripts on **Primary Application Server (PAS)**, **ABAP System Central Services (ASCS) Server**, **Database (DB) Server**, **Additional App Servers (AAS)**, and **Enqueue Replication Server (ERS)**. You can add a script for each server selected.
  - To remove a configuration script, choose **Remove script**. To add more configuration scripts, choose **Add another script**.

### Post-deployment configuration script — optional

- **Deployment settings.** In the event of a configuration script failure, choose whether to proceed with the deployment. If you do not select this option, then when the configuration scripts fail or time out, Launch Wizard will roll back the workload and delete all of the AWS resources created by Launch Wizard. Post-deployment configuration scripts must finish running in 2 hours or less.
- **Configuration script.** You can add one or more configuration scripts depending on the number of servers you select to run scripts during the post-deployment phase.
  - For each post-deployment configuration script that you want to run, choose to use a script located in Amazon S3 and enter the URL path of the script, or upload a script file.
  - Select the servers to run the post-deployment configuration scripts when an EC2 instance has been configured for use. You can choose to run the post-deployment scripts on **Primary Application Server (PAS)**, **ABAP System Central Services (ASCS) Server**, **Database (DB) Server**, **Additional App Servers (AAS)**, and **Enqueue Replication Server (ERS)**. You can add a script for each server selected.
  - To remove a configuration script, choose **Remove script**. To add more configuration scripts, choose **Add another script**.

8. After you have entered all of your deployment settings, choose **Next**.

(See the **Review** tab)

#### Review

- On the **Review** page, review your infrastructure, application, and deployment model settings. If you are satisfied with your selections, choose **Deploy**. If you want to change settings, choose **Previous**.
- When you choose **Deploy**, you are redirected to the **Deployments** page, where you can view the status of your deployment, and also the deployment details.

### SAP HANA database

#### Application settings

On the **Configure application settings** page, enter your SAP HANA database application settings.

1. **Application type.** Select **SAP HANA database**. This configuration includes:
  - EC2 instances for an SAP HANA database
  - Optional installation of SAP HANA database software
2. **General Settings – SAP HANA.** Enter the settings for your SAP HANA database installation.
  - **SAP HANA System ID (SID).** Enter the SAP HANA system ID for your system. The HANASID must be different from SAPSID if you are configuring a single instance deployment.
  - **SAP HANA Instance number.** Enter the instance number to use for your SAP HANA system. This must be a two-digit number from 00 through 99.
  - **EBS Volume Type for SAP HANA.** Select the EBS volume types that you want to use for both **SAP HANA Data** and **SAP HANA Logs** from the dropdown lists.
  - **SAP HANA software install.** Select whether you want to download the SAP HANA software.
    - If you select **Yes**, enter the Amazon S3 location where the SAP HANA software is located. The S3 bucket must have the prefix "launchwizard" in the bucket name to ensure that the Launch Wizard IAM role policy for EC2 has read-only access to the bucket. For steps to set up the folder structure for your S3 bucket, see [Make SAP HANA software available for AWS Launch Wizard to deploy a HANA database \(p. 70\)](#). Enter a password to use for your SAP HANA installation.
  - **AWS Backint Agent.** Select the check box if you want to deploy AWS Backint Agent for backup and recover along with the application. For more information about AWS Backint Agent, see [AWS Backint Agent for SAP HANA](#).
    - **S3 URI.** Enter the URI of the S3 bucket where you want to store your SAP HANA backup files. For example, `s3://<bucket-name>`.
    - **S3 Encryption (KMS key ARN).** Select the ARN of the KMS key that AWS Backint Agent can use to encrypt the backup files stored in your Amazon S3 bucket.
    - **Agent version.** Select the version number of the agent that you want to install. If you do not enter a version number, the latest published version of the agent is installed.
  - **Additional Backint preferences.**
    - **AWS Launch Wizard provisions the local backup EBS volume.** Select the check box if you want Launch Wizard to provision the local backup EBS volume for storing HANA backups.
    - **AWS Launch Wizard to proceed with deployment when Backint installation fails.** Select the check box if you want Launch Wizard to continue with an application deployment when the Backint installation fails.

- **Verify that you have attached the required policy for Backint operations to the following role.** Select this check box after you have attached the required policy to the `AmazonEC2RoleForLaunchWizard`. This policy allows Launch Wizard to perform Backint Agent operations on your behalf. The policy and instructions to attach the policy to the role are provided by Launch Wizard during deployment. This information can also be found in [Step 2 of the Backint Agent IAM documentation](#).
  - If you select **No**, only the AWS infrastructure is provisioned so you can manually deploy an SAP HANA database post deployment .
3. After you enter your application settings, choose **Next**.

(Use the tab for **Single instance deployment**, **Multiple instance deployment**, or **High availability deployment**, depending on your configuration)

#### Single instance deployment

On the **Configure deployment model** page, enter the deployment details for the SAP HANA database deployment.

1. **Deployment model.** Launch Wizard supports single instance deployments, multiple instance deployments, and high availability deployments. Select **Single instance deployment**.
2. **Settings for SAP HANA database on one instance**
  - **Instance details.**
    - Under **Image type**, choose whether to use **AWS/Marketplace/Community images** or **Bring your own images (BYOI)**.
      - **Operating System.** Select a supported operating system version for the ERS instance.
      - **AMI ID.** For BYOI, select the AMI that you want to use from the dropdown.
      - **Host name.** Enter the host name for the EC2 instance.
      - **Auto Recovery.** Auto recovery is an Amazon EC2 feature to increase instance availability. Select the check box to enable EC2 automatic recovery for the instance. For more information, see [Recover Your Instance](#) in the Amazon EC2 User Guide.
  - Under **Instance sizing**, choose **Use AWS recommended resources** or **Choose your instance**.
  - **Use AWS recommended resources.**
    - **Define requirements.** Choose the requirements for your recommended resources from the dropdown list.
      - **Based on CPU/Memory.** If you select this option, enter the required number of vCPU **Cores** and **Memory**. Amazon EC2 supports up to 448 logical processors. If the amount of memory required exceeds 4TB, [dedicated hosts](#) are required.
      - **SAPS (SAP Application Performance Standard).** If you select this option, enter the **SAPS** rating for the SAP certified instance types.
  - **Choose your instance.**
    - **Instance type.** Choose the instance type from the dropdown list.
    - **Recommended Resources.** Launch Wizard displays the **Estimated monthly cost of operation** based on your instance sizing selections. This is an estimate of AWS costs to deploy additional resources and does not include applicable taxes or discounts.
3. **Pre- and post-deployment configuration scripts — optional**

You can run pre- and post-deployment configuration scripts during application provisioning. For more information about how Launch Wizard accesses and deploys these scripts, see [Custom deployment configuration scripts \(p. 46\)](#).

### Pre-deployment configuration script — optional

- **Deployment settings.** In the event of a configuration script failure or time out, select whether to ignore all failures and proceed with the deployment. If you do not select this option, when the configuration scripts fail or time out, Launch Wizard will roll back the workload and delete all of the AWS resources created by Launch Wizard. Pre-deployment configuration scripts must finish running in 45 minutes or less.
- **Configuration script.** Choose to use a script located in Amazon S3 and enter the URL path of the script, or enter a script manually by uploading a script file. To remove the configuration script, choose **Remove script**.

### Post-deployment configuration script — optional

- **Deployment settings.** In the event of a configuration script failure or time out, select whether to ignore all failures and proceed with the deployment. If you do not select this option, when the configuration scripts fail or time out, Launch Wizard will roll back the workload and delete all of the AWS resources created by Launch Wizard. Post-deployment configuration scripts must finish running in 2 hours or less.
  - **Configuration script.** Choose to use a script located in Amazon S3 and enter the URL path of the script, or enter a script manually by uploading a script file. To remove the configuration script, choose **Remove script**.
4. After you enter your deployment settings, choose Next.

(See the **Review** tab)

#### Multiple instance deployment

On the **Configure deployment model** page, enter the deployment details for the SAP HANA database deployment.

1. **Deployment model.** Launch Wizard supports single instance deployments, multiple instance deployments, and high availability deployments. Select **Multiple instance deployment**.
2. **SAP HANA on multiple EC2 instances**
  - **Instance details.**
    - Under **Instance sizing**, choose whether to use **AWS/Marketplace/Community images** or **Bring your own images (BYOI)**.
    - **Operating System.** Select a supported operating system version for the SAP HANA servers.
    - **AMI ID.** For BYOI, select the AMI that you want to use from the dropdown.
  - Under **Instance sizing**, choose **Use AWS recommended resources** or **Choose your instance**.
  - **Use AWS recommended resources.**
    - **Infrastructure requirements.** Choose the requirements for your recommended resources from the dropdown list.
      - **Based on CPU/Memory.** If you select this option, enter the required number of vCPU **Cores** and **Memory**. Amazon EC2 supports up to 448 logical processors. If the amount of memory required exceeds 4TB, [dedicated hosts](#) are required.
      - **SAPS (SAP Application Performance Standard).** If you select this option, enter the **SAPS** rating for the SAP certified instance types.
  - **Choose your instance.**
    - **Instance type.** Choose the instance type from the dropdown list.

- **Host Name for SAP system.** Enter the host name for the EC2 instance.
- **Number of worker nodes.** Enter the number of EC2 instances to be configured as worker nodes for this SAP HANA system.
- **Worker node hostname prefix.** Enter the hostname prefix for the worker nodes.
- **Auto Recovery.** Auto recovery is an Amazon EC2 feature to increase instance availability. Select the check box to enable EC2 automatic recovery for the instance. For more information, see Recover Your Instance in the Amazon EC2 User Guide.
- **Recommended Resources.** Launch Wizard displays the **Estimated monthly cost of operation** based on your instance sizing selections. This is an estimate of AWS costs to deploy additional resources and does not include applicable taxes or discounts.

### 3. Pre- and post-deployment configuration scripts — optional

You can run pre- and post-deployment configuration scripts during application provisioning. For more information about how Launch Wizard accesses and deploys these scripts, see [Custom deployment configuration scripts \(p. 46\)](#).

#### Pre-deployment configuration script — optional

- **Deployment settings.** In the event of a configuration script failure or time out, select whether to ignore all failures and proceed with the deployment. If you do not select this option, when the configuration scripts fail or time out, Launch Wizard will roll back the workload and delete all of the AWS resources created by Launch Wizard. Pre-deployment configuration scripts must finish running in 45 minutes or less.
- **Configuration script.** Choose to use a script located in Amazon S3 and enter the URL path of the script, or enter a script manually by uploading a script file. To remove the configuration script, choose **Remove script**.

#### Post-deployment configuration script — optional

- **Deployment settings.** In the event of a configuration script failure or time out, select whether to ignore all failures and proceed with the deployment. If you do not select this option, when the configuration scripts fail or time out, Launch Wizard will roll back the workload and delete all of the AWS resources created by Launch Wizard. Post-deployment configuration scripts must finish running in 2 hours or less.
- **Configuration script.** Choose to use a script located in Amazon S3 and enter the URL path of the script, or enter a script manually by uploading a script file. To remove the configuration script, choose **Remove script**.

### 4. After you enter your deployment settings, choose **Next**.

(See **Review** tab)

#### High availability deployment

On the **Configure deployment model** page, enter the deployment details for the SAP HANA database deployment.

1. **Deployment model.** Launch Wizard supports single instance deployments, multiple instance deployments, and high availability deployments. Select **High availability deployment**.
2. **Instance details.**
  - Under **Instance details**, choose whether to use **AWS/Marketplace/Community images** or **Bring your own images (BYOI)**.
  - **Operating System.** Select a supported operating system version for the SAP HANA servers.
  - **AMI ID.** For BYOI, select the AMI that you want to use from the dropdown.

- **Primary and secondary instance details.** Enter details for both the primary and secondary instances.
    - **SAP HANA host name.** Enter the host name for the SAP HANA primary and secondary instances.
    - **Server site name.** Enter the primary and secondary site name for the SAP HANA system replication.
  - **Overlay IP address.** Enter the overlay IP address to assign to the active node. The IP address should be outside of the VPC CIDR and must not be used by any other HA cluster. It is configured to always point to the active SAP HANA node.
  - **Pacemaker tag name.** Enter the tag to assign to each EC2 instance. This tag is used by the pacemaker component of SLES HAE and RHEL for SAP high availability solutions and must not be used by any other EC2 instance in your account.
  - Under **Instance sizing**, choose **Use AWS recommended resources** or **Choose your instance**.
    - **Use AWS recommended resources.**
      - **Infrastructure requirements.** Choose the requirements for your recommended resources from the dropdown list.
        - **Based on CPU/Memory.** If you select this option, enter the required number of vCPU **Cores** and **Memory**. Amazon EC2 supports up to 448 logical processors. If the amount of memory required exceeds 4TB, [dedicated hosts](#) are required.
        - **SAPS (SAP Application Performance Standard).** If you select this option, enter the **SAPS** rating for the SAP certified instance types.
      - **Choose your instance.**
        - **Instance type.** Choose the instance type from the dropdown list.
        - **Recommended Resources.** Launch Wizard displays the **Estimated monthly cost of operation** based on your instance sizing selections. This is an estimate of AWS costs to deploy additional resources and does not include applicable taxes or discounts.
3. **Pre- and post-deployment configuration scripts — optional**

You can run pre- and post-deployment configuration scripts during application provisioning. For more information about how Launch Wizard accesses and deploys these scripts, see [Custom deployment configuration scripts \(p. 46\)](#).

#### **Pre-deployment configuration script — optional**

- **Deployment settings.** In the event of a configuration script failure or time out, select whether to ignore all failures and proceed with the deployment. If you do not select this option, when the configuration scripts fail or time out, Launch Wizard will roll back the workload and delete all of the AWS resources created by Launch Wizard. Pre-deployment configuration scripts must finish running in 45 minutes or less.
- **Configuration script.** Choose to use a script located in Amazon S3 and enter the URL path of the script, or enter a script manually by uploading a script file. To remove the configuration script, choose **Remove script**.

#### **Post-deployment configuration script — optional**

- **Deployment settings.** In the event of a configuration script failure or time out, select whether to ignore all failures and proceed with the deployment. If you do not select this option, when the configuration scripts fail or time out, Launch Wizard will roll back the workload and delete all of the AWS resources created by Launch Wizard. Post-deployment configuration scripts must finish running in 2 hours or less.



- **Configuration script.** Choose to use a script located in Amazon S3 and enter the URL path of the script, or enter a script manually by uploading a script file. To remove the configuration script, choose **Remove script**.
4. After you enter your deployment settings, choose **Next**.

(See **Review** tab)

#### Review

- On the **Review** page, review your infrastructure, application, and deployment model settings. If you are satisfied with your selections, choose **Deploy** . If you want to change settings, choose **Previous**.
- When you choose **Deploy** , you are redirected to the **Deployments page**, where you can view the status of your deployment, and also the deployment details.

## Manage application resources with AWS Launch Wizard for SAP

After you have deployed an SAP application, you can manage and update it by following these steps.

### Manage deployments

1. From the left navigation pane, choose **SAP**.
2. Under the **Deployments** tab, select the check box next to the application that you want to manage, and then choose **Actions**. You can do the following:
  1. **Manage resources on the EC2 console.** You are redirected to the Amazon EC2 console, where you can view and manage your SAP application resources, such as Amazon EC2, Amazon EBS, Amazon VPC, Subnets, NAT Gateways, and Elastic IPs.
  2. **View resource group with Systems Manager.** In the Systems Manager console, you can manage your application with built-in integrations through resource groups. Launch Wizard automatically tags your deployment with resource groups. When you access Systems Manager through Launch Wizard, the resources are automatically filtered for you based on your resource group. You can manage, patch, and maintain your applications in Systems Manager.
  3. **View CloudWatch application logs.**
  4. **View CloudFormation template.**
3. To delete a deployment, select the application that you want to delete, and select **Delete**. You are prompted to confirm the deletion.

#### Important

When you delete a deployment, Launch Wizard attempts to delete only the AWS resources it created in your account as part of the deployment. Launch Wizard considers certain resources, such as security groups, infrastructure configuration templates created during a deployment, and EFS file systems created for a transport directory, as shared resources between multiple deployments. Shared resources are not deleted when you delete a deployment.

4. For more information about your application resources, choose the **Application name**. You can then view the **Deployment events** and **Summary** details for your application using the tabs at the top of the page.

### Delete infrastructure configuration

1. From the left navigation pane, choose **SAP**.



2. Under the **Saved infrastructure configurations** tab, select the configuration name you want to delete, and then choose **Delete**. You are prompted to confirm the deletion.

**Important**

When you delete an infrastructure configuration, it will not be available for future deployments. Resources created from the configuration, such as VPCs, availability groups, subnets, and key pair names are not deleted.

3. For more information about an infrastructure configuration, choose the **Configuration name**.

## Troubleshoot AWS Launch Wizard for SAP

Each application in your account in the same AWS Region can be uniquely identified by the application name specified at the time of a deployment. The application name can be used to view the details related to the application launch.

**Contents**

- [Launch Wizard provisioning events](#) (p. 69)
- [CloudWatch Logs](#) (p. 69)
- [AWS CloudFormation stack](#) (p. 69)
- [Pre- and post-deployment configuration scripts](#) (p. 70)
- [Application launch quotas](#) (p. 70)
- [Errors](#) (p. 70)

## Launch Wizard provisioning events

Launch Wizard captures events from SSM Automation and AWS CloudFormation to track the status of an ongoing application deployment. If an application deployment fails, you can view the deployment events for this application by selecting **Deployments** from the navigation pane. A failed event shows a status of **Failed** along with a failure message.

## CloudWatch Logs

Launch Wizard streams provisioning logs from all of the AWS log sources, such as AWS CloudFormation, SSM, and CloudWatch Logs. CloudWatch Logs for a given application name can be viewed on the CloudWatch console for the log group name `LaunchWizard-APPLICATION_NAME` and log stream `ApplicationLaunchLog`.

## AWS CloudFormation stack

Launch Wizard uses AWS CloudFormation to provision the infrastructure resources of an application. AWS CloudFormation stacks can be found in your account using the AWS CloudFormation [describe-stacks](#) API. Launch Wizard launches various stacks in your account for validation and application resource creation. The following are the relevant filters for the [describe-stacks](#) API.

- **Application resources**

`LaunchWizard-APPLICATION_NAME`.

You can view the status of these AWS CloudFormation stacks. If any of them fail, you can view the cause of the failure.

## Pre- and post-deployment configuration scripts

### Can't find the output of my scripts

- **Cause:** Customizations are key scripts that you want to run on the EC2 instances and the logs from script deployments are not included with the provisioning logs.
- **Solution:** The logs for scripts that run on EC2 instances are included in the CloudWatch log group that Launch Wizard creates in your account for the workload. The CloudWatch log group can be identified as `LaunchWizard-APPLICATION_NAME`. You can find the following logs in this log group.
  - `lw-customization/<instance-id>/preDeploymentConfiguration` — For pre-deployment configuration scripts that run on the specified EC2 instance.
  - `lw-customization/<instance-id>/postDeploymentConfiguration` — For post-deployment configuration scripts that run on the specified EC2 instance.

## Application launch quotas

Launch Wizard allows for a maximum of 50 active applications (with status in `progress` or `completed`) for any given application type. If you want to increase this limit, contact [AWS Support](#).

## Errors

### Your requested instance type is not supported in your requested Availability Zone

- **Cause:** This failure might occur during the launch of your instance, or during the validation of the instances that Launch Wizard launches in your selected subnets.
- **Solution:** For this scenario, you must choose a different Availability Zone and retry the deployment from the initial page of the Launch Wizard console.

### Infrastructure template already exists

- **Cause:** This failure occurs when you choose to create a new infrastructure configuration and then navigate back to the first step in the wizard to review or adjust any settings. Launch Wizard has already registered the configuration template, so choosing **Next** results in the error "Template name already exists. Select a new template name."
- **Solution:**

Perform one of the following actions to continue with your deployment.

  - Change the name of the configuration template and continue.
  - Choose another template and continue.
  - Delete the template causing the error by navigating to the **Saved Infrastructure Setting** tab under **Deployments – SAP**, and then continue with your configuration using the same configuration name.

## Make SAP HANA software available for AWS Launch Wizard to deploy a HANA database

This section describes steps to download the SAP HANA software and upload it to Amazon S3 to make it available for Launch Wizard to deploy a HANA database.

### Topics

- [Download SAP HANA software \(p. 71\)](#)
- [Upload SAP HANA software to Amazon S3 \(p. 71\)](#)

## Download SAP HANA software

To download the SAP HANA software, go to the **SAP Software Downloads** page and download the installation files directly to your local drive.

1. Navigate to the [SAP Software Downloads](#) page and log in to your account.
2. Under **Installation and Upgrades**, choose **Access Downloads>A-Z index**.
3. Choose **H** in the **Installations and Upgrades** window, and select **SAP HANA Platform Edition** from the list.
4. Choose **SAP HANA Platform Edition>Installation**.
5. In the **Downloads** window, find the revision you want to download and download each file to your local drive.

### Note

If you do not have access to the software and believe you should, contact the [SAP Global Support Customer Interaction Center](#).

### Important

Do not extract the downloaded HANA software. Instead, stage the files in your Amazon S3 bucket as is. Launch Wizard will extract the media and install the software for you.

## Upload SAP HANA software to Amazon S3

To upload the SAP HANA software to your Amazon S3 bucket, you must create and set up your destination bucket.

### Set up destination bucket

1. Navigate to the Amazon S3 console at <https://console.aws.amazon.com/s3>.
2. Choose **Create Bucket**.
3. In the **Create Bucket** dialog box, provide a name for your new S3 bucket with the prefix `launchwizard`. Choose the **AWS Region** where you want to create the S3 bucket, which should be a Region that is close to your location, and then choose **Create Bucket**. For detailed information about bucket names and Region selection, see [Create a Bucket](#) in the **Amazon S3 Getting Started Guide**.
4. Choose the bucket that you created and, from the **Overview** tab, **Create folders** to organize your SAP HANA downloads. We recommend that you create a folder for each version of SAP HANA.
5. To add the unextracted SAP HANA files to the appropriate folder, choose **Upload** from the **Overview** tab.

If the path for the specific version of SAP HANA software is `s3://launchwizardhanamedia/SP23` or `s3://launchwizardhanamedia/SP24`, then use this path in the Amazon S3 URL for SAP HANA software (`HANAInstallMedia`) parameter.

### Note

We recommend that you place only the main SAP HANA installation files in the S3 bucket. Do not place multiple SAP HANA versions in the same folder. SAP provides the software as a single `.zip` file or as multiple files depending on the SAP HANA version (one `.exe` file and multiple `.rar` files). Upload them to the version-specific folder that you created.

## Supported operating system versions for SAP deployments

The following table provides the details for the operating systems supported by Launch Wizard for SAP deployments .

Operating system version	Single-node deployment	ASCS	ERS	PAS	SAP HANA database	SAP HANA database in HA cluster
Red-Hat-Enterprise-Linux-7.6-For-SAP-HA-US-HVM	✓	✓	✓	✓	✓	✓
Red-Hat-Enterprise-Linux-8.1-For-SAP-HA-US-HVM	✓	✓	✓	✓	✓	✓
SuSE-Linux-15-HVM	✓			✓	✓	
SuSE-Linux-15-For-SAP-HVM	✓	✓	✓	✓	✓	✓
SuSE-Linux-15-For-SAP-BYOS-HVM	✓	✓	✓	✓	✓	✓
SuSE-Linux-12-SP4-HVM	✓			✓	✓	
SuSE-Linux-12-SP4-For-SAP-HVM	✓	✓	✓	✓	✓	✓
SuSE-Linux-12-SP4-For-SAP-BYOS-HVM	✓	✓	✓	✓	✓	✓
SuSE-Linux-12-SP5-HVM	✓			✓	✓	

Operating system version	Single-node deployment	ASCS	ERS	PAS	SAP HANA database	SAP HANA database in HA cluster
SuSE-Linux-12-SP5-For-SAP-HVM	✓	✓	✓	✓	✓	✓
SuSE-Linux-12-SP5-For-SAP-BYOS-HVM	✓	✓	✓	✓	✓	✓
SuSE-Linux-15-SP1-HVM	✓			✓	✓	
SuSE-Linux-15-SP1-For-SAP-HVM	✓	✓	✓	✓	✓	✓
SuSE-Linux-15-SP1-For-SAP-BYOS-HVM	✓	✓	✓	✓	✓	✓

For Additional Application Server (AAS), the operating system is inherited from the operating system selected for the PAS server.

## Security groups in AWS Launch Wizard for SAP

This section describes the security groups that Launch Wizard for SAP creates and assigns to the database and application instances. It also describes how the entries in the outbound and inbound communication rules for database and application security groups are updated.

### Topics

- [Security groups \(p. 73\)](#)
- [Connectivity to external systems and users \(p. 75\)](#)

## Security groups

A security group acts as a virtual firewall that controls the traffic for one or more instances. When you allow Launch Wizard to create security groups, it creates a set of security groups and assigns them to the SAP database and application instances to allow for inbound traffic. Security groups use the following naming conventions:

- `<Infrastructure_Configuration_Name>_App_SecurityGroup`
- `<Infrastructure_Configuration_Name>_DB_SecurityGroup`

### <Infrastructure\_Configuration\_Name>\_App\_SecurityGroup

<Infrastructure\_Configuration\_Name>\_App\_SecurityGroup is configured as follows to allow inbound access to the database servers.

Source	Protocol	Port Range
All instances attached to this security group	all	
All instances attached to the DB security group	TCP	1-65535

This configuration allows:

- inbound communication on all TCP ports from all of the SAP application servers deployed using the same configuration name
- inbound communication on all TCP ports from all of the database servers deployed using the same configuration name.

### <Infrastructure\_Configuration\_Name >\_DB\_SecurityGroup

<Infrastructure\_Configuration\_Name>\_DB\_SecurityGroup is configured as follows to allow inbound access to the database servers.

Source	Protocol	Port Range
All instances attached to this security group	all	
All instances attached to the App security group	TCP	1-65535
All instances attached to the App security group	UDP	111
All instances attached to the App security group	UDP	2049
All instances attached to the App security group	UDP	4000-4002

This configuration allows:

- inbound communication on all TCP ports from all of the SAP database servers deployed using the same configuration name.
- inbound communication on all TCP ports from all of the SAP application servers deployed using the same configuration name.
- inbound communication on UDP 111,2049 and 4000 to 4002 from all the SAP application servers deployed using the same configuration name.

## Connectivity to external systems and users

CIDR/IP address and security group entries are entered in the infrastructure configuration. This allows access to SAP systems by front end users and upstream/downstream systems that are running in that CIDR block, or by end users (IP address) or systems assigned to those security groups. Port ranges are included in the rule definition that allow inbound access so that you can reuse the infrastructure configuration and deploy SAP systems with an instance number 00 to 99. Each entry in the outbound and inbound communication rules for a database security group, created either by the service or provided by the user, are updated as follows.

Source	Protocol	Port Range
Input	TCP	22
Input	TCP	1128 - 1129
Input	TCP	4300 - 4399
Input	TCP	8000 - 8099
Input	TCP	8443
Input	TCP	30013 - 39913
Input	TCP	30015 - 39915
Input	TCP	30017 - 39917
Input	TCP	30041 - 39941
Input	TCP	30044 - 39944
Input	TCP	50013 - 59914

Each entry in the outbound and inbound communication rules for the application security group, created either by the service or by the user, are updated as follows.

Source	Protocol	Port Range
Input	TCP	22
Input	TCP	3200 - 3399
Input	TCP	8080
Input	TCP	8443
Input	TCP	3600-3699
Input	TCP	4237

**Note**

When the deployment is complete, you can update the security group information by adjusting the port range and source information.

**Note**

Launch Wizard considers a security group that it created as a shared resource. It does not delete the security group if you delete a deployment or if a deployment is rolled back.

# AWS Launch Wizard security

Cloud security at AWS is the highest priority. As an AWS customer, you benefit from a data center and network architecture that is built to meet the requirements of the most security-sensitive organizations.

Security is a shared responsibility between AWS and you. The [shared responsibility model](#) describes this as security *of* the cloud and security *in* the cloud:

- **Security of the cloud** – AWS is responsible for protecting the infrastructure that runs AWS services in the AWS Cloud. AWS also provides you with services that you can use securely. Third-party auditors regularly test and verify the effectiveness of our security as part of the [AWS Compliance Programs](#). To learn about the compliance programs that apply to AWS Launch Wizard, see [AWS Services in Scope by Compliance Program](#).
- **Security in the cloud** – Your responsibility is determined by the AWS service that you use. You are also responsible for other factors including the sensitivity of your data, your company's requirements, and applicable laws and regulations.

This documentation helps you understand how to apply the shared responsibility model when using AWS Launch Wizard. The following topics show you how to configure Launch Wizard to meet your security and compliance objectives. You also learn how to use other AWS services that help you to monitor and secure your Launch Wizard resources.

AWS Launch Wizard deploys Amazon EC2 instances into Amazon VPCs. For security information for Amazon EC2 and Amazon VPC, see the security sections in the [Amazon EC2 Getting Started Guide](#) and the [Amazon VPC User Guide](#).

## Contents

- [Infrastructure security in Launch Wizard \(p. 76\)](#)
- [Resilience in Launch Wizard \(p. 76\)](#)
- [Data protection in Launch Wizard \(p. 77\)](#)
- [Identity and Access Management for AWS Launch Wizard \(p. 77\)](#)
- [Update management in Launch Wizard \(p. 78\)](#)

## Infrastructure security in Launch Wizard

As a managed service, Launch Wizard is protected by the AWS global network security procedures that are described in the [Amazon Web Services: Overview of Security Processes](#) whitepaper.

## Resilience in Launch Wizard

The AWS global infrastructure is built around AWS Regions and Availability Zones. Regions provide multiple physically separated and isolated Availability Zones, which are connected through low-latency, high-throughput, and highly redundant networking. With Availability Zones, you can design and operate applications and databases that automatically fail over between Availability Zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures.

For more information about AWS Regions and Availability Zones, see [AWS Global Infrastructure](#).



AWS Launch Wizard sets up an application across multiple Availability Zones to ensure automatic failover between Availability Zones without interruption. Availability Zones are more highly available, fault tolerant, and scalable than traditional single or multiple datacenter infrastructures.

## Data protection in Launch Wizard

The AWS [shared responsibility model](#) applies to data protection in AWS Launch Wizard. As described in this model, AWS is responsible for protecting the global infrastructure that runs all of the AWS Cloud. You are responsible for maintaining control over your content that is hosted on this infrastructure. This content includes the security configuration and management tasks for the AWS services that you use. For more information about data privacy, see the [Data Privacy FAQ](#). For information about data protection in Europe, see the [AWS Shared Responsibility Model and GDPR](#) blog post on the *AWS Security Blog*.

For data protection purposes, we recommend that you protect AWS account credentials and set up individual user accounts with AWS Identity and Access Management (IAM). That way each user is given only the permissions necessary to fulfill their job duties. We also recommend that you secure your data in the following ways:

- Use multi-factor authentication (MFA) with each account.
- Use SSL/TLS to communicate with AWS resources. We recommend TLS 1.2 or later.
- Set up API and user activity logging with AWS CloudTrail.
- Use AWS encryption solutions, along with all default security controls within AWS services.
- Use advanced managed security services such as Amazon Macie, which assists in discovering and securing personal data that is stored in Amazon S3.
- If you require FIPS 140-2 validated cryptographic modules when accessing AWS through a command line interface or an API, use a FIPS endpoint. For more information about the available FIPS endpoints, see [Federal Information Processing Standard \(FIPS\) 140-2](#).

We strongly recommend that you never put sensitive identifying information, such as your customers' account numbers, into free-form fields such as a **Name** field. This includes when you work with Launch Wizard or other AWS services using the console, API, AWS CLI, or AWS SDKs. Any data that you enter into Launch Wizard or other services might get picked up for inclusion in diagnostic logs. When you provide a URL to an external server, don't include credentials information in the URL to validate your request to that server.

## Identity and Access Management for AWS Launch Wizard

AWS Launch Wizard uses the following AWS managed policies to grant permissions to users and services.

- **AmazonEC2RolePolicyForLaunchWizard**

AWS Launch Wizard creates an IAM role with the name **AmazonEC2RoleForLaunchWizard** in your account if the role already does not already exist in your account. If the role exists, the role is attached to the instance profile for the Amazon EC2 instances that Launch Wizard will launch into your account. This role is comprised of two IAM managed policies: **AmazonSSMManagedInstanceCore** and **AmazonEC2RolePolicyForLaunchWizard**.

When you choose to deploy your SAP application with AWS Backint Agent for SAP HANA, you must attach the IAM inline policy provided in [Step 2 of the AWS Identity and Access Management](#)

[documentation for AWS Backint Agent for SAP HANA](#). This policy and instructions to attach the policy to the role are provided by Launch Wizard.

- **AmazonSSMManagedInstanceCore**

This policy enables AWS Systems Manager service core functionality on Amazon EC2. For information, see [Create an IAM Instance Profile for Systems Manager](#).

- **AmazonLaunchWizard\_Fullaccess**

This policy provides full access to AWS Launch Wizard and other required services.

- **AWSLambdaVPCLambdaAccessExecutionRole**

This policy provides minimum permissions for a Lambda function to execute while accessing a resource within a VPC. These permissions include create, describe, delete network interfaces, and write permissions to CloudWatch Logs.

- **AmazonLambdaRolePolicyForLaunchWizardSAP**

This policy provides minimum permissions to enable SAP provisioning scenarios on Launch Wizard. It allows invocation of Lambda functions to be able to perform certain actions, such as validation of route tables and perform pre-configuration and configuration tasks for HA mode enabling.

- To run custom pre- and post-configuration deployment scripts, you must manually add the permissions provided in [Add permissions to run custom pre- and post-deployment configuration scripts \(p. 49\)](#) to the `AmazonEC2RoleForLaunchWizard` role.

## Update management in Launch Wizard

We recommend that you regularly patch, update, and secure the operating system and applications on your EC2 instances. You can use [AWS Systems Manager Patch Manager](#) to automate the process of installing security-related updates for both the operating system and applications. Alternatively, you can use any automatic update services or recommended processes for installing updates that are provided by the application vendor.

# AWS Launch Wizard documentation history

The following table describes the documentation for this release of AWS Launch Wizard.

- **Latest documentation update:** November 17, 2020

update-history-change	update-history-description	update-history-date
<a href="#">AWS Launch Wizard for SAP support for custom pre-deployment and post-deployment scripts. (p. 38)</a>	You can run custom pre- and post-deployment configuration scripts using AWS Launch Wizard for SAP.	November 17, 2020
<a href="#">AWS Launch Wizard support for SQL Server application single-node deployments. (p. 1)</a>	You can deploy your SQL Server application on a single instance.	October 28, 2020
<a href="#">AWS Launch Wizard for SAP support for application single-node deployments. (p. 38)</a>	You can deploy your SAP application on a single instance.	October 15, 2020
<a href="#">Route 53/DNS association support (p. 38)</a>	You can provide your DNS domain name or Route53 hosted zone to enable DNS association for your deployed EC2 instances.	June 18, 2020
<a href="#">AWS Launch Wizard for SQL Server integration with CloudWatch Application Insights (p. 1)</a>	You can set up monitoring for your application with CloudWatch Application Insights.	June 18, 2020
<a href="#">SQL Server witness node support. (p. 1)</a>	You can add a witness node to your SQL Server Always On configuration.	May 11, 2020
<a href="#">Proxy server support (p. 38)</a>	You can route outbound internet traffic for deployed EC2 instances through a proxy server.	May 11, 2020
<a href="#">Initial release (p. 38)</a>	Initial release of AWS Launch Wizard for SAP User Guide.	April 8, 2020
<a href="#">Initial release (p. 79)</a>	Initial release of the AWS Launch Wizard for SQL Server User Guide.	November 14, 2019