
Microsoft SQL Server on Amazon EC2

User Guide



Microsoft SQL Server on Amazon EC2: User Guide

Copyright © 2023 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

What is Microsoft SQL Server on Amazon EC2?	1
Options	1
SQL Server on Amazon EC2	2
RDS for SQL Server	2
Amazon RDS Custom	3
Decision matrix	3
Concepts	13
Features	15
Pricing	16
Set up SQL Server on Amazon EC2	17
Prerequisites	17
Sign up for an AWS account	17
Create an administrative user	17
Create a key pair	18
Create a security group	19
Permissions	21
Get started	22
Licensing SQL Server	22
Licensing options	22
Licensing considerations	23
SQL Server license-included AMIs	26
SQL Server license-included AMI discovery options	26
AMI versions	28
Deploy a SQL Server on Amazon EC2	28
Considerations	28
Deployment options	28
Connect to a SQL Server on Amazon EC2	33
SSMS	33
Configuration Manager	33
Best practices	34
Assign IP addresses	34
Cluster properties	35
Cluster quorum votes and 50/50 splits in a multi-site cluster	35
DNS registration	35
Elastic Network Adapters (ENAs)	36
Multi-site clusters and EC2 instance placement	36
Instance type selection	36
Assign elastic network interfaces and IPs to the instance	37
Heartbeat network	37
Configure the network adapter in the OS	37
IPv6	37
Host record TTL for SQL Availability Group Listeners	37
Logging	38
NetBIOS over TCP	38
NetFT Virtual Adapter	38
Set possible owners	38
Tune the failover thresholds	39
Witness importance and Dynamic Quorum Architecture	40
Troubleshoot	40
Security	41
Document history	42

What is Microsoft SQL Server on Amazon EC2?

You can run Microsoft SQL Server on Amazon Elastic Compute Cloud (Amazon EC2). Microsoft SQL Server is a relational database management system (RDBMS) whose primary purpose is to store and retrieve data. SQL Server includes additional services, such as Analysis Services (SSAS), Reporting Services (SSRS), Integration Services (SSIS), and Machine Learning (ML). AWS provides a comprehensive set of services and tools to deploy Microsoft SQL Server on the reliable and secure AWS Cloud infrastructure. The benefits of running SQL Server on AWS include cost savings, scalability, high availability and disaster recovery, improved performance, and ease of management. For more information, see [Learn why AWS is the best cloud to run Microsoft Windows Server and SQL Server workloads](#) on the AWS Compute blog.

Amazon Elastic Compute Cloud (Amazon EC2) supports a self-managed SQL Server. That is, it gives you full control over the setup of the infrastructure and the database environment. Running SQL Server on Amazon EC2 is very similar to running SQL Server on your own server. You have full control of the database and operating system-level access, so you can use your choice of tools to manage the operating system, database software, patches, data replication, backup, and restoration. You are responsible for data replication and recovery across your instances in the same or different AWS Regions. For more information, refer to the [AWS Shared Responsibility Model](#).

Overview topics

- [SQL Server on the AWS Cloud \(p. 1\)](#)
- [Microsoft SQL Server on Amazon EC2 concepts \(p. 13\)](#)
- [Microsoft SQL Server on Amazon EC2 features \(p. 15\)](#)
- [Microsoft SQL Server on Amazon EC2 pricing \(p. 16\)](#)

SQL Server on the AWS Cloud

AWS provides the option to run Microsoft SQL Server in a cloud environment. For developers and database administrators, running SQL Server in the AWS Cloud is similar to running SQL Server databases in a data center. There are three primary options to run SQL Server on AWS:

- Microsoft SQL Server on Amazon Elastic Compute Cloud (Amazon EC2)
- Amazon RDS for Microsoft SQL Server
- Amazon RDS Custom for Microsoft SQL Server

Your application requirements, database features, functionality, growth capacity, and overall architecture complexity will determine which option to choose. If you are migrating multiple SQL Server databases to AWS, some of them might be a great fit for Amazon RDS, whereas others might be better suited to run directly on Amazon EC2. You might have databases that are running on SQL Server Enterprise edition but are a good fit for SQL Server Standard edition. You may also want to modernize your SQL Server database running on Windows to run on a Linux operating system to save on cost and licenses. Many AWS customers run multiple SQL Server database workloads across Amazon RDS and Amazon EC2.

When to choose:

- [Microsoft SQL Server on Amazon EC2 \(p. 2\)](#)
- [Amazon RDS for Microsoft SQL Server \(p. 2\)](#)
- [Amazon RDS Custom for SQL Server \(p. 3\)](#)

- [Decision matrix \(p. 3\)](#)

Microsoft SQL Server on Amazon EC2

When to choose Microsoft SQL Server on Amazon Elastic Compute Cloud (Amazon EC2):

- You want full control over the database and access to its underlying operating system, database installation, and configuration.
- You want to administer your database, including backups and recovery, patching the operating system and the database, tuning the operating system and database parameters, managing security, and configuring high availability or replication.
- You want to use features and options that aren't currently supported by Amazon RDS. For more information, see [Features not supported and features with limited support](#) in the Amazon RDS documentation.
- You require a specific SQL Server version that isn't supported by Amazon RDS. For a list of supported versions and editions, see [SQL Server versions on Amazon RDS](#) in the *RDS for Microsoft SQL Server User Guide*.
- Your database size and performance requirements exceed the current RDS for Microsoft SQL Server offerings. For more information, see [Amazon RDS DB instance storage](#) in the *Amazon RDS User Guide*.
- You want to avoid automatic software patches that might not be compliant with your applications.
- You want to bring your own license instead of using the RDS for Microsoft SQL Server license-included model.
- You want to achieve higher IOPS and storage capacity than the current limits. For more information, see [Amazon RDS DB instance storage](#) in the *Amazon RDS User Guide*.

Amazon RDS for Microsoft SQL Server

RDS for Microsoft SQL Server is a managed database service that simplifies the provisioning and management of SQL Server on AWS. With Amazon RDS, you can quickly deploy multiple versions and editions of SQL Server, with cost-efficient and resizable compute capacity. You can provision Amazon RDS for SQL Server DB instances with either General Purpose SSD or Provisioned IOPS SSD storage. Provisioned IOPS SSD is optimized for I/O-intensive, transactional (OLTP) database workloads.

Amazon RDS manages database administration tasks, including provisioning, backups, software patching, monitoring, and hardware scaling. Amazon RDS also offers Multi-AZ deployments and read replicas (for SQL Server Enterprise edition) to provide high availability, performance, scalability, and reliability for production workloads. For more information, see [Amazon RDS for Microsoft SQL Server](#).

When to choose RDS for Microsoft SQL Server:

- You want to focus on your business and applications, and you want AWS to take care of undifferentiated heavy lifting tasks, such as the provisioning of the database, management of backup and recovery tasks, management of security patches, minor SQL Server version upgrades, and storage management.
- You want a highly available database solution, and you want to take advantage of the push-button, synchronous Multi-AZ replication offered by Amazon RDS, without having to manually set up and maintain database mirroring, failover clusters, or Always On availability groups.
- You want to pay for the SQL Server license as part of the instance cost on an hourly basis, instead of making a large, up front investment.
- Your database size and IOPS requirements are supported by Amazon RDS for SQL Server. See [Amazon RDS DB Instance Storage](#) in the AWS documentation for the current maximum limits.
- You don't want to manage backups or point-in-time recoveries of your database.

- You want to focus on high-level tasks, such as performance tuning and schema optimization, instead of the daily administration of the database.
- You want to scale the instance type up or down based on your workload patterns without being concerned about licensing complexities.

Amazon RDS Custom for SQL Server

Amazon RDS Custom for SQL Server is a managed database service for legacy, custom, and packaged applications that require access to the underlying operating system and database environment. Amazon RDS Custom for SQL Server automates setup, operation, and scaling of databases in the AWS Cloud while granting you access to the database and underlying operating system on Amazon EC2 to configure settings, install patches, and enable native features to meet the dependent application's requirements. For more information, see [Working with RDS Custom for SQL Server](#) in the *Amazon Relational Database Service User Guide*.

When to choose Amazon RDS Custom for SQL Server:

- You want the benefits of Amazon RDS, but your requirements include the need to customize the underlying operating system and database environment for legacy, custom, and packaged applications.
- You need administrative rights to the database and underlying operating system.
- You need to install custom database and OS patches and packages.
- You need to configure file systems to share files directly with their applications.

Decision matrix

The following table provides a side-by-side comparison of SQL Server features supported on Amazon RDS, Amazon RDS Custom, and Amazon EC2. Use this information to understand the differences between these services, and to choose the best approach for your use case. For the most current information for Amazon RDS, see [Microsoft SQL Server on Amazon RDS](#) and [Working with RDS Custom for SQL Server](#) in the *Amazon RDS User Guide*.

Development

Development feature	Amazon RDS	Amazon RDS Custom	Amazon EC2	Notes
Buffer pool extensions	⊗ No	✔ Yes	✔ Yes	
BULK INSERT	✔ Yes	✔ Yes	✔ Yes	See Integrating an Amazon RDS for SQL Server DB instance with Amazon S3 in the Amazon RDS documentation.
Change data capture (CDC)	✔ Yes (Enterprise Edition: all versions; Standard Edition: 2016 SP1 and later)	✔ Yes	✔ Yes	See Using change data capture in the Amazon RDS documentation.

Microsoft SQL Server on Amazon EC2 User Guide
Decision matrix

Development feature	Amazon RDS	Amazon RDS Custom	Amazon EC2	Notes
Change tracking	✔ Yes	✔ Yes	✔ Yes	
Columnstore indexes	✔ Yes	✔ Yes (Enterprise Edition: 2019)	✔ Yes (Enterprise Edition)	
Data Quality Services	✘ No	✔ Yes	✔ Yes	
Database Mail	✔ Yes	✔ Yes	✔ Yes	<p>See the blog post Using Database Mail on Amazon RDS for SQL Server.</p> <p>We encourage you to use the Amazon Simple Email Service (Amazon SES) to send outbound email originating from AWS resources, to ensure a high degree of deliverability.</p>
Database Engine Tuning Advisor	✔ Yes	✔ Yes	✔ Yes	
DB event notifications	✔ Yes	✔ Yes	✘ No (manually track and manage DB events)	See Using Amazon RDS event notification in the Amazon RDS documentation.
DDL event notifications	✘ No	✔ Yes	✔ Yes	
Delayed transaction durability (lazy commit)	✔ Yes	✔ Yes (SQL Server 2019)	✔ Yes	
Distributed queries	✔ Yes (SQL Server targets)	✔ Yes (SQL Server targets)	✔ Yes (SQL Server targets)	See the Implementing linked servers with Amazon RDS for SQL Server blog post.

Microsoft SQL Server on Amazon EC2 User Guide
Decision matrix

Development feature	Amazon RDS	Amazon RDS Custom	Amazon EC2	Notes
Extended stored procedures, including xp_cmdshell	✗ No	✓ Yes	✓ Yes	
File tables	✗ No	✓ Yes	✓ Yes	
FILESTREAM	✗ No	✓ Yes	✓ Yes	FILESTREAM isn't compatible with Amazon RDS. However, you can configure the in-memory database.
Full-text search	✓ Yes (except semantic search)	✓ Yes	✓ Yes	
In-memory database	✓ Yes	✓ Yes (SQL Server 2019)	✓ Yes	
Linked servers	✓ Yes (SQL Server targets)	✓ Yes	✓ Yes	See the Implementing linked servers with Amazon RDS for SQL Server blog post.
Machine Learning Services (with R scripts)	✗ No	✓ Yes	✓ Yes	Machine Learning Services must be installed separately on a Windows or Linux machine. It's supported on an Always On Failover Cluster Instance (FCI) only in SQL Server 2019 and later. Although R isn't supported on Amazon RDS, you can use it on AWS (see the blog post Getting started with R on AWS).

Microsoft SQL Server on Amazon EC2 User Guide
Decision matrix

Development feature	Amazon RDS	Amazon RDS Custom	Amazon EC2	Notes
Maintenance plans	✘ No	✔ Yes	✔ Yes	Amazon RDS provides a separate set of features to facilitate backup and recovery of databases. For backup, you can configure automated backup.
Master Data Services	✘ No	✔ Yes	✔ Yes	
Microsoft Distributed Transaction Coordinator (MSDTC)	✔ Yes	✔ Yes	✔ Yes	See the blog post Enabling distributed transaction support for domain-joined Amazon RDS for SQL Server instances.
OPENROWSET	✔ Yes	✔ Yes	✔ Yes	
Partially contained databases	✔ Yes	✔ Yes (SQL Server 2019)	✔ Yes	
Performance Data Collector	✘ No	✔ Yes	✔ Yes	On Amazon RDS, you can use Amazon CloudWatch, AWS CloudTrail, and Performance Insights to monitor your SQL Server performance (see Overview of monitoring Amazon RDS in the Amazon RDS documentation).
Policy-Based Management	✘ No	✔ Yes	✔ Yes	
PolyBase	✘ No	✔ Yes	✔ Yes	
Resource Governor	✘ No	✔ Yes	✔ Yes	

Microsoft SQL Server on Amazon EC2 User Guide
Decision matrix

Development feature	Amazon RDS	Amazon RDS Custom	Amazon EC2	Notes
Safe CLR	✔ Yes	✔ Yes	✔ Yes	
Server-level triggers	✘ No	✔ Yes	✔ Yes	
Service Broker	✔ Yes (except endpoints)	✔ Yes	✔ Yes	
Spatial and location features	✔ Yes	✔ Yes	✔ Yes	
SQL Server Agent	✔ Yes	✔ Yes	✔ Yes	
SQL Server Analysis Services (SSAS)	✔ Yes (SQL Server 2016 and later)	✔ Yes	✔ Yes	See Support for SSAS in Amazon RDS for SQL Server in the Amazon RDS documentation.
SQL Server Integration Services (SSIS)	✔ Yes (SQL Server 2016 and later)	✔ Yes	✔ Yes	See Support for SSIS in Amazon RDS for SQL Server in the Amazon RDS documentation.
SQL Server Management Studio (SSMS)	✔ Yes	✔ Yes	✔ Yes	
SQL Server Profiler	✔ Yes (server-side and client-side traces)	✔ Yes	✔ Yes	
SQL Server Reporting Services (SSRS)	✔ Yes (SQL Server 2016 and later)	✔ Yes	✔ Yes	See Support for SSRS in Amazon RDS for SQL Server in the Amazon RDS documentation.
sqlcmd	✔ Yes	✔ Yes	✔ Yes	
Stretch Database	✘ No	✔ Yes	✔ Yes	
THROW statement	✔ Yes	✔ Yes (SQL Server 2019)	✔ Yes	

Development feature	Amazon RDS	Amazon RDS Custom	Amazon EC2	Notes
Transact-SQL endpoints	✗ No	✓ Yes	✓ Yes	All operations that use CREATE ENDPOINT are unavailable on Amazon RDS. We recommend that you install SQL Server on an EC2 instance for these operations.
UTF-16 support	✓ Yes	✓ Yes	✓ Yes	
WCF Data Service	✗ No	✓ Yes	✓ Yes	

HA/DR

Development feature	Amazon RDS	Amazon RDS Custom	Amazon EC2	Notes
Always On availability groups	✓ Yes	✓ Yes (both synchronous and asynchronous)	✓ Yes	If you need a self-managed Always On availability group, we recommend that you use AWS Launch Wizard to simplify SQL Server HA deployment on an EC2 instance. See AWS Launch Wizard for SQL Server in the AWS documentation.
Always On Failover Cluster Instances (FCIs)	✗ No	✓ Yes	✓ Yes	You can use AWS Launch Wizard to simplify your SQL Server FCI deployment on Amazon EC2. See AWS Launch Wizard for SQL Server in the AWS documentation.
Backing up to Amazon S3	✓ Yes	✓ Yes	✓ Yes	Amazon RDS supports native backup and restore for SQL Server databases by using full

Microsoft SQL Server on Amazon EC2 User Guide
Decision matrix

Development feature	Amazon RDS	Amazon RDS Custom	Amazon EC2	Notes
				backup files (.bak files) and Amazon S3 as a repository. See Importing and Exporting SQL Server databases in the Amazon RDS documentation.
BACKUP command	✘ No	✔ Yes	✔ Yes	See How do I perform native backups of an Amazon RDS DB instance that's running SQL Server? in AWS Knowledge Center.
Database mirroring	✔ Yes (Multi-AZ)	✔ Yes	✔ Yes	
Database replication	✘ No (limited push subscription)	✔ Yes	✔ Yes	If you want to replicate a single table on Amazon RDS, you can also use AWS DMS or set up read replicas.
Distributed availability groups	✘ No	✔ Yes	✔ Yes	
Log shipping	✘ No	✔ Yes	✔ Yes	For disaster recovery purposes, you can use read replicas or AWS DMS .
Managed automated backups	✔ Yes	✔ Yes	✘ No (requires configuring and managing maintenance plans, or using third-party solutions)	See Working with backups in the Amazon RDS documentation.

Development feature	Amazon RDS	Amazon RDS Custom	Amazon EC2	Notes
Multi-AZ with automated failover	✔ Yes	✔ Yes (with manual configuration of Always On availability groups)	✔ Yes (Enterprise Edition only, with manual configuration of Always On availability groups)	See Multi-AZ deployments for Amazon RDS for SQL Server in the Amazon RDS documentation.
Read replicas	✔ Yes (SQL Server 2016 and later)	✔ Yes (with manual configuration of Always On availability groups)	✔ Yes (with manual configuration of Always On availability groups)	
RESTORE command	✔ Yes	✔ Yes	✔ Yes	See AWS Knowledge Center .

Scalability

Development feature	Amazon RDS	Amazon RDS Custom	Amazon EC2	Notes
Built-in instance and database monitoring and metrics	✔ Yes	✘ No	✘ No (export your own metrics to CloudWatch or use a third-party solution)	See the blog post Monitor your SQL Server database by using custom metrics with Amazon CloudWatch and AWS Systems Manager .
Configurable storage size	✔ Yes	✘ No	✔ Yes	
Maximum number of databases per instance	Depends on the instance size and Multi-AZ configuration	5,000	32,767	See Maximum capacity specifications for SQL Server in the Microsoft SQL Server documentation.
Maximum storage size of a DB instance	16 TiB	16 TiB	✔ No limitation	Amazon RDS also supports tempdb databases on local disks by using Non-Volatile Memory Express (NVMe) instance storage.

Development feature	Amazon RDS	Amazon RDS Custom	Amazon EC2	Notes
				See Instance store support for the tempdb database on Amazon RDS for SQL Server in the Amazon RDS documentation.
Minimum storage size of a DB instance	20 GiB (Enterprise, Standard, Web, and Express Editions)	20 GiB (Enterprise, Standard, Web, and Express Editions)	✔ No limitation	
New Query Optimizer	✔ Yes	✔ Yes	✔ Yes	
Read replicas	✔ Yes (SQL Server 2016 and later)	✔ Yes (with manual configuration of Always On availability groups)	✔ Yes (with manual configuration of Always On availability groups)	

Security

Development feature	Amazon RDS	Amazon RDS Custom	Amazon EC2	Notes
Automatic software patching	✔ Yes	✘ No	✘ No	
Encrypted storage using AWS KMS	✔ Yes (all SQL Server editions except Express)	✔ Yes	✔ Yes	See the blog post Securing data in Amazon RDS using AWS KMS encryption .
Flexible server roles	✔ Yes	✔ Yes (SQL Server 2019)	✔ Yes	
SQL authentication	✔ Yes	✔ Yes	✔ Yes	
SQL Server audit	✔ Yes	✔ Yes	✔ Yes	
SSL (encryption in transit)	✔ Yes	✔ Yes	✔ Yes	See Using SSL with a Microsoft SQL Server DB instance in the

Development feature	Amazon RDS	Amazon RDS Custom	Amazon EC2	Notes
				Amazon RDS documentation.
sysadmin role	✗ No	✓ Yes	✓ Yes	<p>For unsupported server-level roles, see Microsoft SQL Server security in the Amazon RDS documentation.</p> <p>When you create a new RDS DB instance, the default IAM user with administrator-level credentials that you use gets certain privileges for that DB instance (see Account privileges in the Amazon RDS documentation).</p>
TDE (encryption at rest)	✓ Yes	✓ Yes	✓ Yes	See Support for transparent data encryption in SQL Server in the Amazon RDS documentation.
Windows Authentication	✓ Yes	✓ Yes	✓ Yes	

Other features

Development feature	Amazon RDS	Amazon RDS Custom	Amazon EC2	Notes
Ability to install a third-party agent	✗ No	✓ Yes	✓ Yes	
Ability to rename existing databases	✓ Yes (Single-AZ only)	✓ Yes (not available for databases in availability groups or enabled for mirroring)	✓ Yes (not available for databases in availability groups or enabled for mirroring)	For Multi-AZ deployments on Amazon RDS, see Renaming a Microsoft SQL Server database in a Multi-AZ deployment in the Amazon RDS documentation.

Development feature	Amazon RDS	Amazon RDS Custom	Amazon EC2	Notes
Control over DB instance and operating system	⊗ No	✔ Yes	✔ Yes	
Custom set time zones	✔ Yes	⊗ No	✔ Yes	
Distributed Replay	⊗ No	✔ Yes	✔ Yes	The SQL Server Distributed Replay client service requires sysadmin permissions , which is why it isn't supported in Amazon RDS.
Import data into the msdb database	⊗ No	✔ Yes	✔ Yes	
Installation methods	N/A	N/A	Amazon Machine Image (AMI) or manual installation	
SQL Server editions	Enterprise, Standard, Web, Express	Enterprise, Standard, Web	Enterprise, Standard, Web, Developer, Express	
SQL Server versions	2014, 2016, 2017, 2019	2019	2014, 2016, 2017, 2019	

Microsoft SQL Server on Amazon EC2 concepts

The following concepts introduce you to the fundamental terminology used when working with Microsoft SQL Server on Amazon Elastic Compute Cloud (Amazon EC2) instances:

- [Amazon Machine Images \(AMIs\)](#)
- [Backup](#)
- [Billing](#)
- [High availability and disaster recovery \(HADR\)](#)
- [Instance](#)
- [Instance types](#)
- [Launching](#)
- [Security](#)
- [Storage](#)

Amazon Machine Images (AMIs)

SQL Server on Amazon EC2 instances are created from Amazon Machine Images (AMIs). AMIs are similar to templates. SQL Server on Amazon EC2 AMIs are pre-installed with an operating system, typically Microsoft Windows Server, and other software. Together, these determine the operating environment. You can select an AMI provided by AWS, create your own AMI, or select an AMI from the AWS Marketplace. To find a SQL Server on Amazon EC2 AMI, see the options under [Find a Windows AMI](#) in the *Amazon EC2 User Guide*.

Backup

Your backup and recovery design for SQL Server on Amazon EC2 is flexible, depending on your RTO and RPO requirements. AWS provides the ability to perform server-level backups using Windows Volume Shadow Copy Service (VSS)-enabled Amazon Elastic Block Store (Amazon EBS) snapshots and with AWS Backup. You can also perform database-level backups using [native backup and restore procedures](#) for SQL Server databases. Database-level backups can be stored on Amazon EBS, FSx for Windows File Server, or Amazon Simple Storage Service using AWS Storage Gateway. For more information about backing up SQL Server on Amazon EC2, see [Backup and restore options for SQL Server on Amazon EC2](#) in the *AWS Prescriptive Guidance*.

Billing

A SQL Server on Amazon EC2 instance is charged by the second, with a minimum of 1 minute. Applied rates are based on the type and size of the selected instance, the edition of SQL Server when using a license-included instance, along with the cost of any additional services, such as storage or networking. AWS provides a variety of instance families that are favorable to the performance requirements of SQL Server workloads.

You can rent an instance based on your unique CPU, memory, and storage throughput requirements. You can also stop or terminate an instance at any time to pause or stop billing for the instance. The main advantage of the On-Demand model is the ability to save on CAPEX when an instance is no longer required.

Warning

Any data on [Amazon EC2 instance store](#) volumes are lost if your instance is stopped or terminated. You'll still incur costs for EBS volumes when your instance is stopped. For more information, see [Stop and start your instance](#) in the *Amazon EC2 User Guide for Windows Instances*.

High availability and disaster recovery (HADR)

You can take advantage of Windows Server Failover Cluster for high availability and disaster recovery (HADR) with SQL Server on Amazon EC2. SQL Server on Amazon EC2 supports both failover cluster instances (SQL FCIs) and Always On availability groups (AG).

For more information see [How do I create a SQL Server Always On availability group cluster in the AWS Cloud?](#) in the AWS knowledge center.

Instance

A SQL Server on Amazon EC2 instance is a virtual (or bare metal) server that runs in the AWS Cloud.

A SQL Server on Amazon EC2 instance is provisioned on demand. The subscriber rents the virtual server by the hour/minute/second, and can use it to deploy specific configurations of SQL Server. For more information about On-Demand instances, see [On-Demand instances](#) in the *Amazon EC2 User Guide*.

An Amazon EC2 Dedicated Hosts is a physical server with EC2 instance capacity that is fully dedicated to your use. Dedicated Hosts allow you to use your existing per-socket, per-core, or per-VM Microsoft SQL Server software licenses. For more information about Dedicated Hosts, see [Dedicated Hosts](#) in the *Amazon EC2 User Guide*.

Instance types

AWS provides various types of instances with different CPU, memory, storage, and networking configurations to support your application requirements. Each instance type is available in various sizes to address specific workload requirements. Instance types are grouped into families according to target application profiles, such as general purpose, compute-optimized, memory-optimized, and storage-optimized. The memory-optimized family of instances is a popular choice for SQL Server on Amazon EC2 because instances in this family have a high memory to CPU ratio for optimal performance. You can choose bare metal instances to support capabilities such as [Always Encrypted with secure enclaves on Amazon EC2 bare metal instances](#). For more information about individual and families of instance types, see [Amazon EC2 Instance Types](#) in the AWS product pages.

Launching SQL Server on Amazon EC2

SQL Server on Amazon EC2 instances can be launched directly from the [Amazon EC2 console](#), with AWS CloudFormation, by using [AWS Tools for PowerShell](#), or by using the [AWS CLI](#). For a guided deployment of Microsoft SQL Server, use [AWS Launch Wizard](#).

Security

AWS supports all security standards and compliance certifications, such as PCI-DSS, HIPAA/HITECH, FedRAMP, GDPR, FIPS, FIPS 140-2, and more. These standards enable you to build a fully compliant application on Amazon EC2. AWS also supports all SQL Server security features such as [Transparent Data Encryption](#) (TDE) and [Always Encrypted with Secure Enclaves](#) (when using bare metal instances).

Security and compliance is a shared responsibility between you and AWS. This shared model helps to relieve your operational burden because AWS operates, manages, and controls the components from the host operating system and virtualization layer to the physical security of the facilities in which the service operates.

For SQL Server on Amazon EC2, you assume responsibility and management of the guest operating system, including updates and security patches, other associated application software, and the configuration of AWS provided security group firewalls.

For more information about the shared responsibility model, see [Shared Responsibility Model](#).

Storage

AWS provides many storage options to host your database files. In addition to EBS volume types, you can attach volumes to SQL Server on Amazon EC2 instances using an Amazon FSx managed file system service, such as FSx for Windows File Server and Amazon FSx for NetApp ONTAP. Some instance types provide an Amazon EC2 instance store which provides temporary block level storage on NVMe solid state drive (SSD) disks that are physically attached to the host computer. For more information, see [Best practices for deploying Microsoft SQL Server on Amazon EC2](#) in the *AWS Prescriptive Guidance*.

Microsoft SQL Server on Amazon EC2 features

SQL Server on Amazon EC2 provides the following features:

- **Flexible licensing options** — When you use Amazon EC2 instances with the license included, you are using instances with fully-compliant Windows Server and SQL Server that are licensed through AWS. Flexible BYOL options include default tenant EC2 for products that are eligible for [Microsoft License Mobility through Software Assurance](#), as well as [Amazon EC2 Dedicated Hosts](#) and [Amazon EC2 Dedicated Instances](#). You can use [AWS License Manager](#) to track the usage of software licenses and reduce the risk of non-compliance. For more information, see [Licensing](#) in the *Amazon Web Services and Microsoft Frequently Asked Questions*.
- **High performance block storage** — [Amazon Elastic Block Store](#) provides multiple options for high-performance block storage for Microsoft SQL Server. EC2 Instances using [io2 Block Express](#) give you the highest block storage performance with a single storage volume. Other SSD-backed Amazon EBS

options include io2 volumes for business-critical applications and gp3 volumes for general purpose applications. Amazon EBS also offers crash-consistent snapshots, and enables application-consistent snapshots through Windows VSS (Volume Shadow Copy Services) to help protect your SQL Server deployments.

- **Fully-managed shared storage** — [Amazon FSx for Windows File Server](#) and Amazon FSx for NetApp ONTAP offer fully-managed shared storage for high-availability SQL Server failover cluster instances (FCI) workloads.
- **Windows-based services** — [AWS Directory Service](#) offers managed Microsoft Active Directory with identity and access management.
- **Scalable processors** — [Intel Xeon Scalable Processors on AWS](#) provide you with better data protection, faster processing of more data volumes, and increased service flexibility for Amazon EC2.
- **Migration programs** — AWS offers programs for migration for customers looking to migrate SQL Server workloads to AWS. AWS [Migration Acceleration Program \(MAP\) for Windows](#) provides services, best practices, and tools to help you save costs and accelerate your migration on AWS.
- **Windows workload optimization** — After you move your SQL Server workloads to AWS, you can continue to optimize costs, usage, and licenses to suit your business requirements. With [Cost Explorer Service](#), you can visualize, understand, and manage your AWS costs and usage over time. [AWS Compute Optimizer](#) recommends optimal AWS compute resources for your workloads so that you can reduce costs up to 25% by analyzing historical utilization data. [AWS Trusted Advisor](#) can check that your EC2 instances have the required amount of SQL Server licenses and that the EC2 instance vCPU count doesn't exceed what is permitted for the SQL Server edition. [AWS Managed Services](#) can help operate your cloud environment post-migration by analyzing alerts and responding to incidents, reducing operational overhead and risk. You can use [AWS Systems Manager](#) to automate operational tasks across your AWS resources and better manage your infrastructure at scale.

AWS can help you to modernize you Windows-based applications with AWS open source services if you want to reduce the high cost of commercial licensing. Options include running SQL Server database applications on Linux, moving workloads to [Amazon Aurora](#), containerizing your Windows applications with [Amazon EKS](#), going serverless with [AWS Lambda](#), or taking advantage of micro-services based architecture.

For more features specific to Amazon EC2, see [Features of Amazon EC2](#).

Microsoft SQL Server on Amazon EC2 pricing

For information about pricing for Amazon EC2, see the [Amazon EC2 pricing](#) page.

For information about creating a price estimate for Microsoft Windows Server and Microsoft SQL Server, see [Tutorial: Using Windows Server and SQL Server on Amazon EC2 calculator](#) in the *AWS Pricing Calculator User Guide*.

Set up Microsoft SQL Server on Amazon EC2

Describes the prerequisites, permissions, and configurations that you should consider when preparing to use Microsoft SQL Server on Amazon Elastic Compute Cloud (Amazon EC2) instances for your SQL Server workloads.

Topics for setting up SQL Server on Amazon EC2

- [Prerequisites for using SQL Server on Amazon EC2 \(p. 17\)](#)
- [Permissions required to use SQL Server on Amazon EC2 \(p. 21\)](#)

Prerequisites for using SQL Server on Amazon EC2

Complete the tasks in this section to start using SQL Server on Amazon EC2 instances for the first time:

1. [Sign up for an AWS account \(p. 17\)](#)
2. [Create an administrative user \(p. 17\)](#)
3. [Create a key pair \(p. 18\)](#)
4. [Create a security group \(p. 19\)](#)

Sign up for an AWS account

If you do not have an AWS account, complete the following steps to create one.

To sign up for an AWS account

1. Open <https://portal.aws.amazon.com/billing/signup>.
2. Follow the online instructions.

Part of the sign-up procedure involves receiving a phone call and entering a verification code on the phone keypad.

When you sign up for an AWS account, an *AWS account root user* is created. The root user has access to all AWS services and resources in the account. As a security best practice, [assign administrative access to an administrative user](#), and use only the root user to perform [tasks that require root user access](#).

AWS sends you a confirmation email after the sign-up process is complete. At any time, you can view your current account activity and manage your account by going to <https://aws.amazon.com/> and choosing **My Account**.

Create an administrative user

After you sign up for an AWS account, create an administrative user so that you don't use the root user for everyday tasks.

Secure your AWS account root user

1. Sign in to the [AWS Management Console](#) as the account owner by choosing **Root user** and entering your AWS account email address. On the next page, enter your password.

For help signing in by using root user, see [Signing in as the root user](#) in the *AWS Sign-In User Guide*.

2. Turn on multi-factor authentication (MFA) for your root user.

For instructions, see [Enable a virtual MFA device for your AWS account root user \(console\)](#) in the *IAM User Guide*.

Create an administrative user

- For your daily administrative tasks, grant administrative access to an administrative user in AWS IAM Identity Center (successor to AWS Single Sign-On).

For instructions, see [Getting started](#) in the *AWS IAM Identity Center (successor to AWS Single Sign-On) User Guide*.

Sign in as the administrative user

- To sign in with your IAM Identity Center user, use the sign-in URL that was sent to your email address when you created the IAM Identity Center user.

For help signing in using an IAM Identity Center user, see [Signing in to the AWS access portal](#) in the *AWS Sign-In User Guide*.

Create a key pair

AWS uses public-key cryptography to secure the login information for your instance. You specify the name of the key pair when you launch your instance, then provide the private key to obtain the administrator password for your Windows instance so you can log in using RDP.

If you haven't created a key pair already, you can create one by using the Amazon EC2 console. Note that if you plan to launch instances in multiple Regions, you'll need to create a key pair in each Region. For more information about Regions, see [Regions and Zones](#) in the *User Guide for Windows Instances*.

To create your key pair

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the navigation pane, choose **Key Pairs**.
3. Choose **Create key pair**.
4. For **Name**, enter a descriptive name for the key pair. Amazon EC2 associates the public key with the name that you specify as the key name. A key name can include up to 255 ASCII characters. It can't include leading or trailing spaces.
5. For **Key pair type**, choose either **RSA** or **ED25519**. Note that **ED25519** keys are not supported for Windows instances.
6. For **Private key file format**, choose the format in which to save the private key. To save the private key in a format that can be used with OpenSSH, choose **pem**. To save the private key in a format that can be used with PuTTY, choose **ppk**.

If you chose **ED25519** in the previous step, the **Private key file format** options do not appear, and the private key format defaults to **pem**.

7. Choose **Create key pair**.

- The private key file is automatically downloaded by your browser. The base file name is the name you specified as the name of your key pair, and the file name extension is determined by the file format you chose. Save the private key file in a safe place.

Important

This is the only chance for you to save the private key file.

For more information, see [Amazon EC2 key pairs and Windows instances](#) in the *User Guide for Windows Instances*.

Create a security group

Security groups act as a firewall for associated instances, controlling both inbound and outbound traffic at the instance level. You must add rules to a security group that enable you to connect to your instance from your IP address using RDP. You can also add rules that allow inbound and outbound HTTP and HTTPS access from anywhere.

Note that if you plan to launch instances in multiple Regions, you'll need to create a security group in each Region. For more information about Regions, see [Regions and Zones](#) in the *User Guide for Windows Instances*.

Prerequisites

You'll need the public IPv4 address of your local computer. The security group editor in the Amazon EC2 console can automatically detect the public IPv4 address for you. Alternatively, you can use the search phrase "what is my IP address" in an Internet browser, or use the following service: [Check IP](#). If you are connecting through an Internet service provider (ISP) or from behind a firewall without a static IP address, you need to find out the range of IP addresses used by client computers.

You can create a custom security group using one of the following methods.

New Amazon EC2 console

To create a security group with least privilege

- Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
- From the top navigation bar, select a Region for the security group. Security groups are specific to a Region, so you should select the same Region in which you created your key pair.
- In the left navigation pane, choose **Security Groups**.
- Choose **Create security group**.
- For **Basic details**, do the following:
 - Enter a name for the new security group and a description. Use a name that is easy for you to remember, such as your user name, followed by `_SG_`, plus the Region name. For example, `me_SG_uswest2`.
 - In the **VPC** list, select your default VPC for the Region.
- For **Inbound rules**, create rules that allow specific traffic to reach your instance. For example, use the following rules for a web server that accepts HTTP and HTTPS traffic. For more examples, see [Security group rules for different use cases](#) in the *User Guide for Windows Instances*.
 - Choose **Add rule**. For **Type**, choose **HTTP**. For **Source**, choose **Anywhere**.
 - Choose **Add rule**. For **Type**, choose **HTTPS**. For **Source**, choose **Anywhere**.
 - Choose **Add rule**. For **Type**, choose **RDP**. For **Source**, do one of the following:
 - Choose **My IP** to automatically add the public IPv4 address of your local computer.

- Choose **Custom** and specify the public IPv4 address of your computer or network in CIDR notation. To specify an individual IP address in CIDR notation, add the routing suffix /32, for example, 203.0.113.25/32. If your company or your router allocates addresses from a range, specify the entire range, such as 203.0.113.0/24.

Warning

For security reasons, do not choose **Anywhere** for **Source** with a rule for RDP. This would allow access to your instance from all IP addresses on the internet. This is acceptable for a short time in a test environment, but it is unsafe for production environments.

7. For **Outbound rules**, keep the default rule, which allows all outbound traffic.
8. Choose **Create security group**.

Old Amazon EC2 console

To create a security group with least privilege

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. In the left navigation pane, choose **Security Groups**.
3. Choose **Create Security Group**.
4. Enter a name for the new security group and a description. Use a name that is easy for you to remember, such as your user name, followed by _SG_, plus the Region name. For example, *me_SG_uswest2*.
5. In the **VPC** list, select your default VPC for the Region.
6. On the **Inbound rules** tab, create the following rules (choose **Add rule** for each new rule):
 - Choose **HTTP** from the **Type** list, and make sure that **Source** is set to **Anywhere** (0.0.0.0/0).
 - Choose **HTTPS** from the **Type** list, and make sure that **Source** is set to **Anywhere** (0.0.0.0/0).
 - Choose **RDP** from the **Type** list. In the **Source** box, choose **My IP** to automatically populate the field with the public IPv4 address of your local computer. Alternatively, choose **Custom** and specify the public IPv4 address of your computer or network in CIDR notation. To specify an individual IP address in CIDR notation, add the routing suffix /32, for example, 203.0.113.25/32. If your company allocates addresses from a range, specify the entire range, such as 203.0.113.0/24.

Warning

For security reasons, do not allow RDP access from all IP addresses to your instance. This is acceptable for a short time in a test environment, but it is unsafe for production environments.

7. On the **Outbound rules** tab, keep the default rule, which allows all outbound traffic.
8. Choose **Create security group**.

Command line

To create a security group with least privilege

Use one of the following commands:

- [create-security-group](#) (AWS CLI)
- [New-EC2SecurityGroup](#) (AWS Tools for Windows PowerShell)

For more information, see [Amazon EC2 security groups for Windows instances](#) in the *Amazon EC2 User Guide for Windows Instances*.

Permissions required to use SQL Server on Amazon EC2

For information about the permissions required to create or modify Amazon EC2 resources, or to perform tasks using the Amazon EC2 API, see [IAM policies for Amazon EC2](#) in the *User Guide for Windows Instances*.

Get started with Microsoft SQL Server on Amazon EC2

This section contains information to help you get started with Microsoft SQL Server on Amazon Elastic Compute Cloud (Amazon EC2), including licensing options, how to create a SQL Server on Amazon EC2 instance, and how to connect to a SQL Server on Amazon EC2 instance.

Getting started topics

- [Licensing Microsoft SQL Server on Amazon EC2 \(p. 22\)](#)
- [Find a SQL Server license-included AMI \(p. 26\)](#)
- [Deploy Microsoft SQL Server on Amazon EC2 \(p. 28\)](#)
- [Connect to a Microsoft SQL Server on Amazon EC2 \(p. 33\)](#)

Licensing Microsoft SQL Server on Amazon EC2

There are two ways in which you can license Microsoft SQL Server on Amazon Elastic Compute Cloud (Amazon EC2) on the AWS Cloud. You own existing SQL Server licenses, or those which are provided by AWS. The most cost-effective license strategy for your workload will depend on multiple factors.

Topics

- [Licensing options \(p. 22\)](#)
- [Licensing considerations \(p. 23\)](#)

Licensing options

You can launch Amazon Elastic Compute Cloud (Amazon EC2) instances with Microsoft SQL Server licenses included from AWS, or you can bring your own SQL Server licenses for use on AWS. You can perform a license type conversion for SQL Server in certain configurations if your needs change. For the most license flexibility, you can import your VM into AWS. For more information, see [Eligible license types for license type conversion](#) in the *AWS License Manager User Guide*.

Licensing options topics

- [License-included \(p. 22\)](#)
- [BYOL \(p. 23\)](#)

License-included

Windows Server with currently supported versions of Microsoft SQL Server AMIs are available from AWS in a variety of combinations. AWS provides these AMIs with SQL Server software and operating system updates already installed. When you purchase an Amazon EC2 instance with a Windows Server AMI, licensing costs and compliance are handled for you. For more information, see [the section called "SQL Server license-included AMIs" \(p. 26\)](#).

Amazon EC2 offers a variety of instance types and sizes that you can configure for your target workload. Amazon EC2 AMIs with Windows Server require no Client Access Licenses (CALs). They also include two Microsoft Remote Desktop Services licenses for administrative purposes.

For SQL Server license-included AMIs, use the installation and setup media included in C:\SQLServerSetup to make changes to the default installation, add new features, or install additional named instances.

BYOL

When you launch a SQL Server instance from an imported AMI, you can bring your existing licenses with the Bring Your Own License model (BYOL), and let AWS manage them to ensure compliance with licensing rules that you set. After you import your licensed image, and it is available as a private AMI in your AWS account on the Amazon EC2 console, you can use the AWS License Manager service to create a license configuration.

After you create the license configuration, you must associate the AMI that contains your licensed operating system image with the configuration. Then, you must create a host resource group and associate it with the license configuration. After you associate your host resource group with the configuration, License Manager automatically manages your hosts when you launch instances into a host resource group, and ensures that you do not exceed your configured license count limits. For more information, see the [Getting started](#) section of the *License Manager User Guide*.

You can also bring your own SQL Server licenses with Active Software Assurance to default (shared) tenant Amazon EC2 through Microsoft License Mobility through Software Assurance. For information about how to sign up for Microsoft License Mobility, see [License Mobility](#).

Licensing considerations

There are many considerations for cost effectively licensing your Microsoft SQL Server on Amazon Elastic Compute Cloud (Amazon EC2) workload. Your use case, and existing license agreements, will determine whether to bring your own license to AWS with the Bring Your Own License model (BYOL) or to use license included AMIs from AWS. The following topics should help determine which approach you might take. For more information, see [Licensing – SQL Server](#) on the *Amazon Web Services and Microsoft Frequently Asked Questions* page.

Licensing considerations topics

- [Choose a SQL Server edition \(p. 23\)](#)
- [Purchase SQL Server from AWS \(p. 24\)](#)
- [Use BYOL for SQL Server on AWS \(p. 24\)](#)
- [Quantify the required SQL Server licenses for BYOL \(p. 24\)](#)
- [License Mobility with SQL Server \(p. 24\)](#)
- [Track BYOL license consumption \(p. 25\)](#)
- [SQL Server client access licenses \(CALs\) \(p. 25\)](#)
- [Licensing for passive failover \(p. 25\)](#)

Choose a SQL Server edition

The edition of SQL Server that is used will determine the supported features your implementation will have available. For example, the edition determines the maximum compute capacity used by a single instance of the SQL Server Database Engine, and the high availability options you might implement. For a comparison of SQL Server editions and supported features, see [Editions and supported features of SQL Server 2022](#) in the Microsoft documentation.

Purchase SQL Server from AWS

You can utilize Microsoft SQL Server licenses included from AWS. You can choose any of the following editions for your use on Amazon EC2 instances.

- SQL Server Express
- SQL Server Web
- SQL Server Standard
- SQL Server Enterprise

Note

SQL Server Developer edition is eligible for use in non-production, development, and test workloads. Once downloaded from Microsoft, you can bring and install SQL Server Developer edition on Amazon EC2 instances in the AWS Cloud. Dedicated infrastructure is not required for SQL Server Developer edition. For more information, see <https://www.microsoft.com/en-us/sql-server/sql-server-downloads>.

Use BYOL for SQL Server on AWS

You can use BYOL licenses for SQL Server on AWS. The requirements differ depending on if the licenses have active Software Assurance.

SQL Server licenses with active Software Assurance

You can bring your SQL Server licenses with active Software Assurance to default (shared) tenant Amazon EC2 through License Mobility benefits. Microsoft requires that you complete and send a License Mobility verification form which can be downloaded [here](#). For more information, see [License Mobility](#).

SQL Server licenses without active Software Assurance

SQL Server licenses without Software Assurance can be deployed on Amazon Elastic Compute Cloud Dedicated Hosts if the licenses are purchased prior to 10/1/2019 or added as a true-up under an active Enterprise Enrollment that was effective prior to 10/1/2019. In these specific BYOL scenarios, the licenses can only be upgraded to versions that were available prior to 10/1/2019. For more information, see [Dedicated Hosts](#) in the *Amazon EC2 User Guide*, and the [Amazon EC2 Dedicated Hosts FAQs](#).

Quantify the required SQL Server licenses for BYOL

If you are licensing SQL Server under Microsoft License Mobility through Software Assurance, the number of licenses required varies based on the instance type, version of SQL Server, and the Microsoft licensing model you choose. For assistance with virtual core licensing calculations under the Microsoft Product Terms based on the instance type, see [SQL License Mobility](#).

If you are using Dedicated Hosts, Amazon EC2 provides you with the number of physical cores installed on the Dedicated Host. Using this information, you can calculate the number of SQL Server licenses that you need to bring in. For more information, see [Amazon EC2 Dedicated Hosts Pricing](#) and the [SQL Server 2022 licensing guide](#).

License Mobility with SQL Server

SQL Server licenses with active Software Assurance are eligible for Microsoft License Mobility and can be deployed on default or dedicated tenant Amazon EC2. For more information on bringing SQL Server licenses with active Software Assurance to default tenant EC2, see [Microsoft License Mobility](#).

It is also possible to bring SQL Server licenses without active Software Assurance to EC2 Dedicated Hosts. To be eligible, the licenses must be purchased prior to October 1, 2019 or added as a true-up under

an active Enterprise Enrollment that was effective prior to October 1, 2019. For additional FAQs about Dedicated Hosts, see the [Dedicated Hosts](#) section of the *Amazon Web Services and Microsoft FAQ*.

Track BYOL license consumption

You can use AWS License Manager to manage your software licenses for SQL Server. With License Manager, you can create license configurations, take inventory of your license-consuming resources, associate licenses with resources, and track inventory and compliance. For more information, see [What is AWS License Manager?](#) in the *AWS License Manager User Guide*.

SQL Server client access licenses (CALs)

When you are using SQL Server on Amazon EC2, license included instances do not require client access licenses (CALs) for SQL Server. An unlimited number of end users can access SQL Server on a license-included instance.

When you bring your own SQL Server licenses to Amazon EC2 through Microsoft License Mobility or BYOL, you must continue to follow the licensing rules in place on-premises. If you purchased SQL Server under the Server/CAL model, you still require CALs to meet Microsoft licensing requirements, but these CALs would remain on-premises and enable end user access SQL Server running on AWS.

Licensing for passive failover

There are various factors to consider when licensing passive failover for SQL Server. The information in this section pertains only to the SQL Server licenses and not the Windows Server licenses. In all cases, you must license Windows Server.

Using instances that include the license for SQL Server

When you purchase SQL Server license included instances on EC2, you must license passive failover instances.

Bringing SQL Server licenses with active Software Assurance to default tenant Amazon EC2

When you bring SQL Server 2014 and later versions with Software Assurance to default tenant EC2, you must license the virtual cores (vCPUs) on the active instance. In return, Software Assurance permits one passive instance (equal or lesser size) where SQL Server licensing is not Amazon EC2 Dedicated Hosts required.

Bringing SQL Server to Amazon EC2 Dedicated Instances

SQL Server 2014 and later versions require Software Assurance for SQL Server passive failover benefits on dedicated infrastructure. When you bring SQL Server with Software Assurance, you must license the cores on the active instance/host and are permitted one passive instance/host (equal or lesser size) where SQL Server licensing is not required.

SQL Server 2008 - SQL Server 2012R2 are eligible for passive failover on an Amazon EC2 Dedicated Hosts infrastructure without active Software Assurance. In these scenarios, you will license the active instance/host, and it will be permitted one passive instance/host of equal or lesser size where SQL Server licensing is not required.

There are specific BYOL scenarios that do not require Microsoft License Mobility through Software Assurance. An Amazon EC2 Dedicated Hosts infrastructure is always required in these scenarios. To be eligible, the licenses must be purchased prior to October 1, 2019 or added as a true-up under an active Enterprise Enrollment that was effective prior to October 1, 2019. In these specific BYOL scenarios, the licenses can only be upgraded to versions that were available prior to October 1, 2019.

Find a SQL Server license-included AMI

This topic describes how you can find SQL Server license-included AMIs using the Amazon EC2 console, the AWS Tools for PowerShell, the AWS CLI, or by searching the AWS Marketplace. For SQL Server license-included AMIs, use the installation and setup media included in C:\SQLServerSetup to make changes to the default installation, add new features, or install additional named instances.

As you select a SQL Server license-included AMI, consider the following requirements you might have for the instances that you'll launch:

- The AWS Region
- The operating system
- The architecture: 64-bit (x86_64)
- The [root device](#) type: Amazon EBS-backed (EBS)
- The provider (for example, Amazon Web Services)
- Additional software (for example, SQL Server)

SQL Server license-included AMI discovery options

AWS Marketplace

To view a list of SQL Server AMIs available from AWS in AWS Marketplace, see [Windows AMIs](#).

Console

You can find SQL Server license-included AMIs using the Amazon EC2 console. You can select from the list of AMIs when you use the launch instance wizard to launch an instance, or you can search through all available AMIs using the **Images** page. AMI IDs are unique to each AWS Region.

To find a SQL Server license-included AMI using the launch instance wizard

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. From the navigation bar, select the Region in which to launch your instances. You can select any Region that's available to you, regardless of your location.
3. From the console dashboard, choose **Launch instances**.
4. Under **Application and OS Images (Amazon Machine Image)**, enter SQL in the search bar and choose **Enter**. You will be taken to the **AMIs** page, where you can browse and choose from AMIs with SQL Server included. You can choose from AMIs under the **Quickstart AMIs**, **My AMIs**, **AWS Marketplace AMIs**, and the **Community AMIs** tabs. You can filter by cost, operating system, and architecture.
5. To launch an instance from this AMI, select it and then choose **Launch instance**. For more information about launching an instance using the console, see [Launch an instance using the new launch instance wizard](#). If you're not ready to launch the instance now, take note of the AMI ID for later.

To find a SQL Server AMI using the AMIs page

1. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
2. From the navigation bar, select the Region in which to launch your instances. You can select any Region that's available to you, regardless of your location.
3. In the navigation pane, choose **AMI Catalog**.

4. Enter SQL in the search bar and choose **Enter**. You can choose from SQL Server license-included AMIs under the **Quickstart AMIs**, **My AMIs**, **AWS Marketplace AMIs**, and the **Community AMIs** tabs. You can filter by cost, operating system, and architecture.
5. To launch an instance from this AMI, select it and then choose **Launch instance**. For more information about launching an instance using the console, see [Launching your instance from an AMI](#). If you're not ready to launch the instance now, take note of the AMI ID for later.

PowerShell

You can use cmdlets for Amazon EC2 to list only the Windows AMIs that match your requirements. After locating an AMI that matches your requirements, take note of its ID so that you can use it to launch instances. For more information, see [Launch an Instance Using Windows PowerShell](#) in the *AWS Tools for Windows PowerShell User Guide*.

To list the latest SQL Server license-included AMIs provided by Amazon, you can use the `Get-SSMLatestEC2Image` cmdlet. The following command lists the latest Windows AMIs with SQL in their image name:

```
Get-SSMLatestEC2Image -Path ami-windows-latest -ImageName *SQL*
```

To list SQL Server license-included AMIs using commands that match specific criteria, you can use the `Get-EC2Image` cmdlet in addition to filters. The following commands filter for AMIs owned by you, or Amazon, with SQL in their name:

```
$name_values = New-Object 'collections.generic.list[string]'  
$name_values.add("*SQL*")  
$filter_name = New-Object Amazon.EC2.Model.Filter -Property @{Name = "name"; Values =  
    $name_values}  
Get-EC2Image -Owner amazon, self -Filter $filter_name
```

For more information and examples, see [Find an AMI Using Windows PowerShell](#) in the *AWS Tools for Windows PowerShell User Guide*.

AWS CLI

You can use AWS CLI commands for Amazon EC2 to list only the SQL Server license-included AMIs that match your requirements. After locating an AMI that matches your requirements, take note of its ID so that you can use it to launch instances. For more information, see [Launching an Instance Using the AWS CLI](#) in the *AWS Command Line Interface User Guide*.

The `describe-images` command supports filtering parameters. For example, use the `--owners` parameter with `amazon` to display public AMIs owned by Amazon or `self` to list AMIs you own. You can specify multiple values for the `--owners` parameter as in the following example:

```
aws ec2 describe-images --owners self amazon
```

You can add the following filter to the previous command to display only SQL Server license-included AMIs:

```
--filters "Name=name,Values=*SQL*"
```

You can use the following filter with the command to display only AMIs backed by Amazon EBS:

```
--filters "Name=root-device-type,Values=ebs"
```

You can combine multiple filters together. For example, this command will list all AMIs owned by you or Amazon with *SQL* in the AMI name and the `--root-device-type` parameter as *ebs*:

```
aws ec2 describe-images --owners self amazon --filters "Name=name,Values=*SQL*"
"Name=root-device-type,Values=ebs"
```

Note

Omitting the `--owners` flag from the `describe-images` command will return all images for which you have launch permissions, regardless of ownership.

SQL Server license-included AMI version history

To view changes to each release of the AWS Windows AMIs, including SQL Server updates, see the [AWS Windows AMI version history](#) in the *Amazon EC2 User Guide*.

Deploy Microsoft SQL Server on Amazon EC2

To launch Microsoft SQL Server using Amazon Elastic Compute Cloud (Amazon EC2) instances with Windows Server, perform the following steps according to your use case.

New SQL environment deployments are classified under three categories:

- SQL Server standalone
- SQL Server Failover Cluster Instances (FCI)
- SQL Server Always On availability groups (AG)

Considerations

Before you launch SQL Server on your instance, consider the following:

- If you use an AWS provided AMI, you must initially manage SQL Server as the local administrator. For more information, see [the section called "Connect to a SQL Server on Amazon EC2" \(p. 33\)](#).
- The built-in availability form of clustering in Windows Server is activated by a feature named Failover Clustering. This feature allows you to build a Windows Server Failover Cluster (WSFC) to use with an availability group or failover cluster instances (FCI).
- Always On is an umbrella term for the availability features in SQL Server, and the term covers both availability groups and FCIs. Always On isn't the name of the Always On availability group (AG) feature.
- The major difference between FCI and AG is that all FCIs require some sort of shared storage, even if it's provided through networking. The FCI's resources can be run and owned by one node at any given time. AG doesn't require that shared storage is also highly available. It's a best practice to have replicas that are local in one data center for high availability, and remote ones in other data centers for disaster recovery, each with separate storage.
- An availability group also has another component called the listener. The listener allows applications and end users to connect without needing to know which SQL Server instance is hosting the primary replica. Each availability group has its own listener.

Deployment options

Use one of the following options to deploy SQL Server on Amazon EC2.

Deploy SQL Server on Amazon EC2 with AWS Launch Wizard

AWS Launch Wizard is a service that guides you through the sizing, configuration, and deployment of enterprise applications following AWS Cloud best practices. AWS Launch Wizard for SQL Server supports both high availability and single instance deployments according to AWS and SQL Server best practices. For more information, see the [AWS Launch Wizard for SQL Server User Guide](#).

Always On availability groups (AG)

Deploy your SQL Server Always On availability groups with primary and secondary replicas for database level protection. Each replica is hosted by a SQL Server instance with its own local storage.

Always On Failover Cluster Instances (SQL FCI)

Deploy SQL Server Always On using Failover Cluster Instances (FCI) for instance-level protection. A single SQL Server instance is installed across Windows Server Failover Clustering (WSFC) nodes to ensure high availability and storage sharing.

Launch Wizard uses Amazon FSx to provide the following shared storage options required for SQL FCI deployments:

- Amazon FSx for NetApp ONTAP using Microsoft iSCSI endpoint
- Amazon FSx for Windows File Server using SMB 3.0 continuously available Windows file share

For more information on how to deploy SQL Server with Launch Wizard, see [Deploy an application with AWS Launch Wizard for SQL Server on Windows](#) in the *AWS Launch Wizard User Guide*.

Deploy SQL Server standalone

For a SQL Server standalone deployment, you can use one of the license-include AMIs provided by AWS or by using your own licensed media. For a list of SQL Server AMIs provided by AWS, see [Windows AMIs](#). For more information on licensing options, see [Licensing Microsoft SQL Server on Amazon EC2 \(p. 22\)](#).

Deploy SQL Server failover cluster instances (FCIs)

Failover cluster instances (FCIs) provide availability for the entire installation of SQL Server known as an instance. Everything that is included in the instance, such as databases, SQL Server Agent jobs, and linked servers, move to a different server when the underlying server fails.

You can use AWS Launch Wizard to deploy SQL Server FCIs in the AWS Cloud. Launch Wizard identifies the AWS resources to automatically provision the SQL Server databases based on your use case. For more information, see [Get started with AWS Launch Wizard for SQL Server](#).

You can reference the following AWS blogs to manually deploy SQL Server FCIs using Amazon FSx:

- [Deploy a SQL Server FCI using SMB 3.0 Continuously Available File Shares \(CAFS\) as shared storage](#)
- [Deploy a SQL Server FCI using Microsoft iSCSI Initiator as shared storage](#)

Deploy SQL Server Always On availability groups (AG)

Always on availability groups provide high availability and disaster recovery of user databases through data replication. Availability groups can also distribute read operations amongst member nodes.

You can use [AWS Launch Wizard \(p. 29\)](#) to deploy a SQL Server Always On availability group in the AWS Cloud. Launch Wizard identifies the AWS resources to automatically provision the SQL Server databases based on your use case. For more information, see [Get started with AWS Launch Wizard for SQL Server](#).

To manually deploy a SQL Server Always On availability group, perform the following steps:

Prerequisites

Before you manually deploy a SQL Server Always On availability group, you must perform the following prerequisites.

- Launch two Amazon EC2 Windows Server instances (Windows Server 2012 or later) across two Availability Zones within an Amazon VPC.
- Install SQL Server 2014 or later 64-bit Enterprise edition. For testing, use SQL Server 2014 or later 64-bit Evaluation edition.
- Configure secondary Amazon EBS volumes to host SQL Server Master Data File (MDF), Log Data File (LDF), and SQL Backup files (.bak).
- Deploy the cluster nodes in private subnets. You can then use Remote Desktop Protocol (RDP) to connect from a jump server to the cluster node instances.
- Configure inbound security group rules and [Windows firewall exceptions](#) to allow the nodes to communicate in a restrictive environment.
- Open all necessary ports for Active Directory domain controllers so that the SQL nodes and witness can join the domain and authenticate against Active Directory.
- Join the nodes to the domain before you create the Windows failover cluster. Ensure that you are logged in with domain credentials before you create and configure the cluster.
- Run the SQL Database instances with an Active Directory service account.
- Create a SQL login with sysadmin permissions using Windows domain authentication. Consult with your database administrator for details. For more information, see [Create a login using SSMS for SQL Server](#) in the Microsoft documentation.
- Properly configure the SQL browser for SQL Server named instances.

Configure the secondary IPs for each cluster node elastic network interface

Two secondary IP addresses are required for each cluster node elastic network interface.

Note

If you do not plan to deploy a SQL Group Listener, add only one secondary IP address for each cluster node elastic network interface.

1. Navigate to the [Amazon EC2 console](#) and choose the AWS Region where you want to host your Always On cluster.
2. Choose **Instances** from the left navigation pane, and then select your Amazon EC2 cluster instance.
3. Choose the **Networking** tab.
4. Under **Network interfaces**, select the network interface and then choose **Actions > Manage IP addresses**.
5. Choose the network interface Id to open the expandable section, and then choose **Assign new IP address**. You can enter a specific IP address or keep the default entry as Auto-assign. Repeat this step to add a second new IP address.
6. Choose **Save > Confirm**.
7. Repeat steps 1 through 7 for the other Amazon EC2 instance that will be included in the cluster.

Create a two-node Windows cluster

Perform the following steps to create a two-node Windows cluster.

1. [Connect to your Amazon EC2 instance](#) with RDP, using a domain account with local administrator permissions on both nodes.
2. On the Windows **Start** menu, open **Control Panel**, and then choose **Network and Internet > Network and Sharing Center**.

3. Choose **Change adapter settings** from the left navigation pane.
4. Choose your network connection, and then choose **Change settings of this connection**.
5. Choose **Internet Protocol Version 4 TCP/IPv4**, and then choose **Properties**.
6. Choose **Advanced**.
7. Under the **DNS** tab, choose **Append primary and connection specific DNS suffixes**.
8. Choose **OK > OK > Close**.
9. Repeat steps 1 through 8 for the other Amazon EC2 instance to include in the cluster.
10. On each instance, install the cluster feature on the nodes from the Server Manager, or run the following Windows PowerShell command:

```
Install-WindowsFeature -Name Failover-Clustering -IncludeManagementTools
```

11. Open the command line as an administrator and enter `cluadmin.msc` to open the Cluster Manager.
12. Open the context menu (right-click) for **Failover Cluster Manager**, and then choose **Create Cluster**.
13. Choose **Next > Browse**.
14. For **Enter the object names to select**, enter the cluster node hostnames, and then choose **OK**.
15. Choose **Next**. You can now choose whether to validate the cluster. We recommend that you perform a cluster validation. If the cluster does not pass validation Microsoft may be unable to provide technical support for your SQL cluster. Choose **Yes** or **No**, and then choose **Next**.
16. Enter a **Cluster Name**, and then choose **Next**.
17. Clear **Add all eligible storage to the cluster**, and then choose **Next**.
18. When the cluster creation is complete, choose **Finish**.

Note

Cluster logs and reports are located at `%systemroot%\cluster\reports`.

19. In the **Cluster Core Resources** section of Cluster Manager, expand the entry for your new cluster.
20. Open the context menu (right-click) for the first IP address entry, and then choose **Properties**. For **IP Address**, choose **Static IP Address**, and then enter one of the secondary IP addresses associated with the `eth0` elastic network interface. Choose **OK**. Repeat this step for the second IP address entry.
21. Open the context menu (right-click) for the cluster name, and then choose **Bring Online**.

Note

We recommend that you configure a [File Share Witness \(FSW\)](#) in addition to your cluster to act as a tie-breaker. You can also use [Amazon FSx for Windows File Server with Microsoft SQL Server](#).

Create Always On availability groups

Perform the following steps to create Always On availability groups.

1. Open SQL Server Configuration Manager.
2. Open the context menu (right-click) for the SQL instance, and then choose **Properties**.
3. On the **AlwaysOn High Availability** tab, select **Enable AlwaysOn Availability Groups**, and then choose **Apply**.
4. Open the context menu (right-click) for the SQL instance, and then choose **Restart**.
5. Repeat steps 1 through 4 on the other cluster node to include in the cluster.
6. Open Microsoft SQL Server Management Studio (SSMS).
7. Log in to one of the SQL instances with your Windows authenticated login that has access to the SQL instance.

Note

We recommend that you use the same MDF and LDF directory file paths across the SQL instances.

8. Create a test database. Open the context menu (right-click) for **Databases**, and then choose **New Database**.

Note

Make sure that you use the **Full** [recovery model](#) on the **Options** page.

9. Enter a **Database name**, and then choose **OK**.
10. Open the context menu (right-click) for the new database name, choose **Tasks**, and then choose **Back Up For Backup type**, choose **Full**.
11. Choose **OK > OK**.
12. Open the context menu (right-click) for **Always On High Availability** and then choose **New Availability Group Wizard**.
13. Choose **Next**.
14. Enter an **Availability group name**, and then choose **Next**.
15. Select your database, and then choose **Next**.
16. A primary replica is already present in the Availability Replicas window. Choose **Add Replica** to create a secondary replica.
17. Enter a **Server name** for the secondary replica and then choose **Connect**.
18. [Decide which Availability Mode you want to use](#) for each replica, and then choose either **Synchronous commit** or **Asynchronous commit**.
19. Choose **Next**.
20. Choose your [data synchronization preference](#), and then choose **Next**.
21. When the validation is successful, choose **Next**.

Note

You can safely ignore **Checking the listener configuration** because you will add it later.

22. Choose **Finish > Close**.

Add a SQL Group Listener

Perform the following steps to add a SQL Group Listener.

1. Open SQL Server Management Studio (SSMS) and expand **Always On High Availability, Availability Groups, <primary replica name>**.
2. Open the context menu (right-click) for **Availability Group Listeners** and then choose **Add Listener**. Enter a **DNS Name**.
3. Enter **Port 1443**.
4. Choose **Static IP for Network Mode**.
5. Choose **Add**.

For the **IPv4 Address**, enter the second secondary IP address from one of the cluster node instances, and then choose **OK**. Repeat this step using the second secondary IP address from the other cluster node instance.

6. Choose **OK**.

Note

If you receive errors when you add a SQL Group Listener, you may be missing permissions. For troubleshooting see:

- [Troubleshooting AlwaysOn availability group listener creation in SQL Server 2012](#)

- [Create Availability Group Listener Fails with Message 19471, 'The WSFC cluster could not bring the Network Name resource online'](#)

Test failover

1. From SSMS, open the context menu (right-click) for the primary replica on the navigation menu, and then choose **Failover**.
2. Choose **Next > Next**.
3. Choose **Connect > Connect**.
4. Choose **Next**, and then choose **Finish**. The primary replica will become the secondary replica after failover.

Connect to a Microsoft SQL Server on Amazon EC2

You can connect to your Microsoft SQL Server instance using one of the following tools:

- [SQL Server Management Studio \(SSMS\) \(p. 33\)](#)
- [SQL Server Configuration Manager \(p. 33\)](#)

SQL Server Management Studio (SSMS)

By default, only the built-in local administrator account can access a SQL Server instance launched from an AWS Windows AMI. You can use SQL Server Management Studio (SSMS) to add domain users so that they can access and manage SQL Server.

Perform the following steps to access a SQL Server instance on Amazon EC2 as a domain user.

1. [Connect to your instance](#) as a local administrator using Remote Desktop Protocol (RDP).
2. Open SQL Server Management Studio (SSMS). For **Authentication**, choose **Windows Authentication** to log in with the built-in local administrator.
3. Choose **Connect**.
4. In Object Explorer, expand **Security**.
5. Open the context menu (right-click) for **Logins** then select **New Login**.
6. For **Login name**, select **Windows authentication**. Enter **Domain\username**, replacing **DomainName** with your domain NetBIOS name and **username** with your Active Directory user name.
7. On the **Server roles** page, select the [server roles](#) that you want to grant to the Active Directory user.
8. Select the **General** page, and then choose **OK**.
9. Log out from the instance and then log in again as a domain user.
10. Open SSMS. For **Authentication**, choose **Windows authentication** to log in with your domain user account.
11. Choose **Connect**.

SQL Server Configuration Manager

To connect to SQL Server using SQL Server Configuration Manager, see [SQL Server Configuration Manager](#) in the Microsoft documentation.

Best practices and recommendations for SQL Server clustering on Amazon EC2

You can configure Microsoft SQL Server on Amazon Elastic Compute Cloud (Amazon EC2) instances for high availability. SQL Server Always On availability groups offer high availability without the requirement for shared storage. The list of best practices in this topic, in addition to the prerequisites listed at [Prerequisites, Restrictions, and Recommendations for Always On availability groups](#), can help you optimize operating a SQL Server Always On availability groups on AWS. The practices listed in this topic also offer a method to gather logs.

Note

When nodes are deployed in different Availability Zones, or in different subnets within the same Availability Zone, they should be treated as a multi-subnet cluster. Keep this in mind as you apply these best practices and when you address possible failure scenarios.

Contents

- [Assign IP addresses \(p. 34\)](#)
- [Cluster properties \(p. 35\)](#)
- [Cluster quorum votes and 50/50 splits in a multi-site cluster \(p. 35\)](#)
- [DNS registration \(p. 35\)](#)
- [Elastic Network Adapters \(ENAs\) \(p. 36\)](#)
- [Multi-site clusters and EC2 instance placement \(p. 36\)](#)
- [Instance type selection \(p. 36\)](#)
- [Assign elastic network interfaces and IPs to the instance \(p. 37\)](#)
- [Heartbeat network \(p. 37\)](#)
- [Configure the network adapter in the OS \(p. 37\)](#)
- [IPv6 \(p. 37\)](#)
- [Host record TTL for SQL Availability Group Listeners \(p. 37\)](#)
- [Logging \(p. 38\)](#)
- [NetBIOS over TCP \(p. 38\)](#)
- [NetFT Virtual Adapter \(p. 38\)](#)
- [Set possible owners \(p. 38\)](#)
- [Tune the failover thresholds \(p. 39\)](#)
- [Witness importance and Dynamic Quorum Architecture \(p. 40\)](#)
- [Troubleshoot \(p. 40\)](#)

Assign IP addresses

Each cluster node should have one elastic network interface assigned that includes three private IP addresses on the subnet: a primary IP address, a cluster IP address, and an Availability Group IP address. The operating system (OS) should have the NIC configured for DHCP. It should not be set for a static IP address because the IP addresses for the cluster IP and Availability Group will be handled virtually in the Failover Cluster Manager. The NIC can be configured for a static IP as long as it is configured to only

use the primary IP of **eth0**. If the other IPs are assigned to the NIC, it can cause network drops for the instance during failover events.

When the network drops because the IPs are incorrectly assigned, or when there is a failover event or network failure, it is not uncommon to see the following event log entries at the time of failure.

```
Isatap interface isatap.{9468661C-0AEB-41BD-BB8C-1F85981D5482} is no longer active.
```

```
Isatap interface isatap.{9468661C-0AEB-41BD-BB8C-1F85981D5482} with address fe80::5efe:169.254.1.105 has been brought up.
```

Because these messages seem to describe network issues, you could potentially mistake the cause of the outage or failure as a network error. However, these errors describe a symptom, rather than cause, of the failure. ISATAP is a tunneling technology that uses IPv6 over IPv4. When the IPv4 connection fails, the ISATAP adapter also fails. When the network issues are resolved, these entries should no longer appear in the event logs. Alternately, you can reduce network errors by safely disabling ISATAP with the following command.

```
netsh int ipv6 isatap set state disabled
```

When you run this command, the adapter is removed from Device Manager. This command should be run on all nodes. It does not impact the ability of the cluster to function. Instead, when the command has been run, ISATAP is no longer used. However, because this command might cause unknown impacts on other applications that use ISATAP, you should test it.

Cluster properties

To see the complete cluster configuration, run the following PowerShell command.

```
Get-Cluster | Format-List -Property *
```

Cluster quorum votes and 50/50 splits in a multi-site cluster

To learn how the cluster quorum works and what to expect if a failure occurs, see [Understanding Cluster and Pool Quorum](#).

DNS registration

In Windows Server 2012, Failover Clustering, by default, attempts to register each DNS node under the cluster name. This is acceptable for applications that are aware the SQL target is configured for multi-site. However, when the client is not configured this way, it can result in timeouts, delays, and application errors due to attempts to connect to each individual node and failing on the inactive ones. To prevent these problems, the Cluster Resource parameter `RegisterAllProvidersIp` must be changed to `0`. For more information, see [RegisterAllProvidersIP Setting](#) and [Multi-subnet Clustered SQL + RegisterAllProvidersIP + SharePoint 2013](#).

The `RegisterAllProvidersIp` can be modified with the following PowerShell script.

```
Import-Module FailoverClusters
```

```
$cluster = (Get-ClusterResource | where {($_.ResourceType -eq "Network Name") -and  
($_.OwnerGroup -ne "Cluster Group")}).Name  
Get-ClusterResource $cluster | Set-ClusterParameter RegisterAllProvidersIP 0  
Get-ClusterResource $cluster | Set-ClusterParameter HostRecordTTL 300  
Stop-ClusterResource $cluster  
Start-ClusterResource $cluster
```

In addition to setting the Cluster Resource parameter to **0**, you must ensure that the cluster has permissions to modify the DNS entry for your cluster name.

1. Log in to the Domain Controller (DC) for the domain, or a server that hosts the forward lookup zone for the domain.
2. Launch the DNS Management Console and locate the A record for the cluster.
3. Choose or right-click the A record, and choose **Properties**.
4. Choose **Security**.
5. Choose **Add**.
6. Choose **Object Types...**, select the box for **Computers**, and choose **OK**.
7. Enter the name of the cluster resource object and choose **Check name** and **OK if resolve**.
8. Select the check box for **Full Control**.
9. Choose **OK**.

Elastic Network Adapters (ENAs)

AWS has identified known issues with some clustering workloads running on ENA driver version 1.2.3. We recommend upgrading to the latest version, and adjusting settings on the NIC in the operating system. For the latest versions, see [Amazon ENA Driver Versions](#). The first setting, which applies to all systems, increases Receive Buffers, which you can do with the following example PowerShell command.

```
Set-NetAdapterAdvancedProperty -Name (Get-NetAdapter | Where-Object  
{$_ .InterfaceDescription -like '*Elastic*'}).Name -DisplayName "Receive Buffers" -  
DisplayValue 8192
```

For instances with more than 16 vCPUs, we recommend preventing RSS from running on CPU 0.

Run the following command.

```
Set-NetAdapterRss -name (Get-NetAdapter | Where-Object {$_ .InterfaceDescription -like  
'*Elastic*'}).Name -Baseprocessorgroup 0 -BaseProcessorNumber 1
```

Multi-site clusters and EC2 instance placement

Each cluster is considered a [multi-site cluster](#). The EC2 service does not share IP addresses virtually. Each node must be in a unique [subnet](#). Though not required, we recommend that each node also be in a unique Availability Zone.

Instance type selection

The type of instance recommended for Windows Server Failover Clustering depends on the workload. For production workloads, we recommend instances that support Amazon Elastic Block Store (Amazon EBS) optimization and enhanced networking. For more information, see [EBS optimization](#) and [Enhanced networking](#) in the *Amazon EC2 User Guide for Windows Instances*.

Assign elastic network interfaces and IPs to the instance

Each node in an EC2 cluster should have only one attached elastic network interface. The network interface should have a minimum of two assigned private IP addresses. However, for workloads that use Availability Groups, such as SQL Always On, you must include an additional IP address for each Availability Group. The primary IP address is used for accessing and managing the server, the secondary IP address is used as the cluster IP address, and each additional IP address is assigned to Availability Groups, as needed.

Heartbeat network

Some Microsoft documentation recommends using a dedicated [heartbeat network](#). However, this recommendation is not applicable to EC2. With EC2, while you can assign and use a second elastic network interface for the heartbeat network, it uses the same infrastructure and shares bandwidth with the primary network interface. Therefore, traffic within the infrastructure cannot be prioritized, and cannot benefit from a dedicated network interface.

Configure the network adapter in the OS

The NIC in the OS can keep using DHCP as long as the DNS servers that are being retrieved from the DHCP Options Set allow for the nodes to resolve each other. You can set the NIC to be configured statically. When completed, you then manually configure only the primary IP address for the elastic network interface. Failover Clustering manages and assigns additional IP addresses, as needed.

For all instance types, you can increase the maximum transmission unit (MTU) on the network adapter to 9001 to support [Jumbo Frames](#). This configuration reduces fragmentation of packets wherever Jumbo Frames are supported. The following example shows how to use PowerShell to configure Jumbo Frames for an Elastic Network Adapter.

```
Get-NetAdapter | Set-NetAdapterAdvancedProperty -DisplayName "MTU" -DisplayValue 9001
```

IPv6

Microsoft does not recommend disabling IPv6 in a Windows Cluster. While Failover Clustering works in an IPv4-only environment, Microsoft tests clusters with IPv6 enabled. See [Failover Clustering and IPv6 in Windows Server 2012 R2](#) for details.

Host record TTL for SQL Availability Group Listeners

Set the host record TTL to **300** seconds instead of the default 20 minutes (1200 seconds). For legacy client comparability, set `RegisterAllProvidersIP` to **0** for SQL Availability Group Listeners. This is not required in all environments. These settings are important because some legacy client applications cannot use `MultiSubnetFailover` in their connection strings. See [HostRecordTTL Setting](#) for more information. When you change these settings, the Cluster Resource must be restarted. The Cluster Group

for the listener stops when the Cluster Resource is restarted, so it must be started. If you do not start the Cluster Group, the Availability Group remains offline in a RESOLVING state. The following are example PowerShell scripts for changing the TTL and RegisterAllProvidersIP settings.

```
Get-ClusterResource yourListenerName | Set-ClusterParameter RegisterAllProvidersIP 0
```

```
Get-ClusterResource yourListenerName | Set-ClusterParameter HostRecordTTL 300
```

```
Stop-ClusterResource yourListenerName
```

```
Start-ClusterResource yourListenerName
```

```
Start-ClusterGroup yourListenerGroupName
```

Logging

The default logging level for the cluster log is **3**. To increase the detail of log information, set the logging level to **5**. See [Set-ClusterLog](#) for more information about the PowerShell cmdlet.

```
Set-ClusterLog -Level 5
```

NetBIOS over TCP

In Windows Server 2012 R2, you can increase the speed of the failover process by disabling NetBIOS over TCP. This feature was removed from Windows Server 2016. You should test this procedure if you are using earlier operating systems in your environment. For more information, see [Speeding Up Failover Tips-n-Tricks](#). The following is an example PowerShell command to disable NetBIOS over TCP.

```
Get-ClusterResource "Cluster IP Address" | Set-ClusterParameter EnableNetBIOS 0
```

NetFT Virtual Adapter

For Windows Server versions earlier than 2016 and non-Hyper-V workloads, Microsoft recommends you enable the NetFT Virtual Adapter Performance Filter on the adapter in the OS. When you enable the NetFT Virtual Adapter, internal cluster traffic is routed directly to the NetFT Virtual Adapter. For more information, see [NetFT Virtual Adapter Performance Filter](#). You can enable NetFT Virtual Adapter by selecting the check box in the NIC properties, or by using the following PowerShell command.

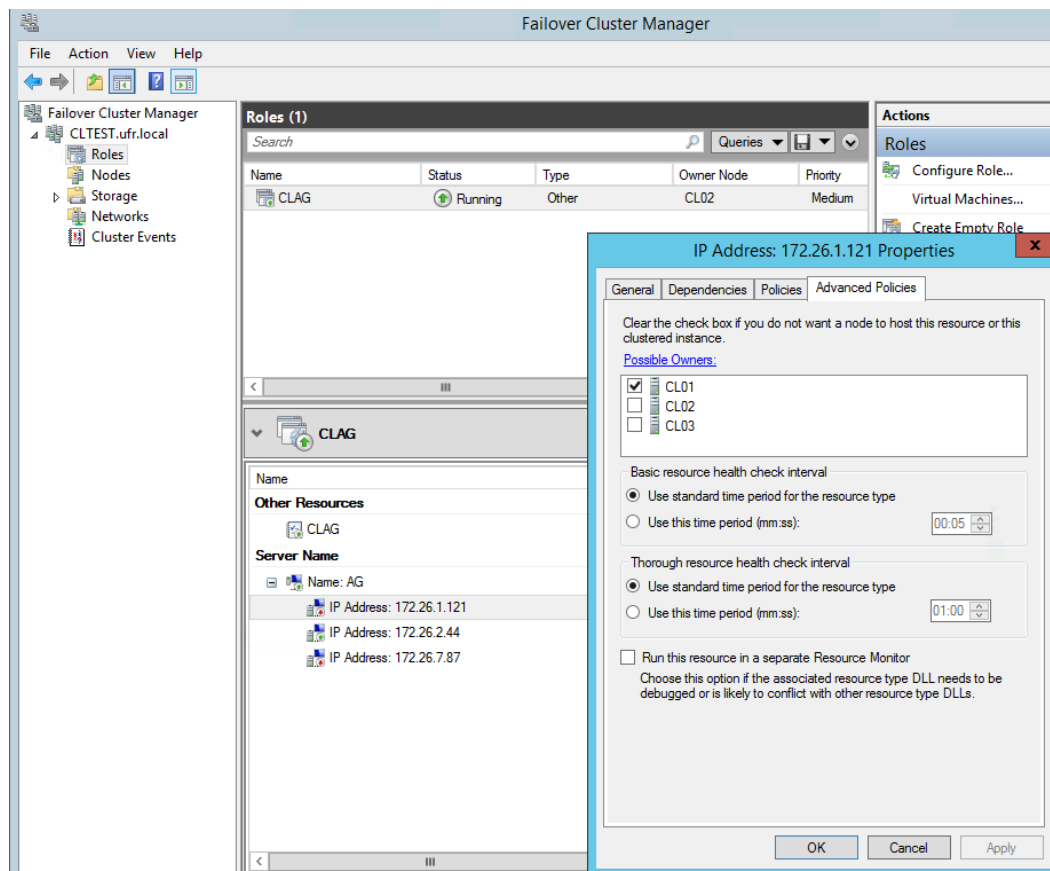
```
Get-NetAdapter | Set-NetAdapterBinding -ComponentID ms_netftflt -Enable $true
```

Set possible owners

The Failover Cluster Manager can be configured so that each IP address specified on the Cluster Core Resources and Availability Group resources can be brought online only on the node to which the IP

belongs. When the Failover Cluster Manager is not configured for this and a failure occurs, there will be some delay in failover as the cluster attempts to bring up the IPs on nodes that do not recognize the address. For more information, see [SQL Server Manages Preferred and Possible Owner Properties for AlwaysOn Availability Group/Role](#).

Each resource in a cluster has a setting for Possible Owners. This setting tells the cluster which nodes are permitted to “online” a resource. Each node is running on a unique subnet in a VPC. Because EC2 cannot share IPs between instances, the IP resources in the cluster can be brought online only by specific nodes. By default, each IP address that is added to the cluster as a resource has every node listed as a Possible Owner. This does not result in failures. However, during expected and unexpected failures, you can see errors in the logs about conflicting IPs and failures to bring IPs online. These errors can be ignored. If you set the Possible Owner property, you can eliminate these errors entirely, and also prevent down time while the services are moved to another node.



Tune the failover thresholds

In Windows Server 2012 R2, the network thresholds for the failover heartbeat network default to high values. See [Tuning Failover Cluster Network Thresholds](#) for details. This potentially unreliable configuration, which applies to clusters with some distance between them, was addressed in Server 2016 with an increase in the number of heartbeats. It was discovered that clusters would fail over due to very brief transient network issues. The heartbeat network is maintained with UDP 3343, which is traditionally far less reliable than TCP and more prone to incomplete conversations. Although there are low-latency connections between AWS Availability Zones, there are still geographic separations with a number of “hops” separating resources. Within an Availability Zone, there may be some distance between clusters unless the customer is using Placement Groups or Dedicated Hosts. As a result, there is a higher possibility for heartbeat failure with UDP than with TCP-based heartbeats.

The only time a cluster should fail over is when there is a legitimate outage, such as a service or node that experiences a hard failover, as opposed to a few UDP packets lost in transit. To ensure legitimate outages, we recommend that you adjust the thresholds to match, or even exceed, the settings for Server 2016 listed in [Tuning Failover Cluster Network Thresholds](#). You can change the settings with the following PowerShell commands.

```
(get-cluster).SameSubnetThreshold = 10
```

```
(get-cluster).CrossSubnetThreshold = 20
```

When you set these values, unexpected failovers should be dramatically reduced. You can fine tune these settings by increasing the delays between heartbeats. However, we recommend that you send the heartbeats more frequently with greater thresholds. Setting these thresholds even higher ensures that failovers occur only for hard failover scenarios, with longer delays before failing over. You must decide how much down time is acceptable for your applications.

After increasing the SameSubnetThreshold or CrossSubnetThreshold, we recommend that you increase the RouteHistoryLength to double the higher of the two values. This ensures that there is sufficient logging for troubleshooting. You can set the RouteHistoryLength with the following PowerShell command.

```
(Get-Cluster).RouteHistoryLength = 20
```

Witness importance and Dynamic Quorum Architecture

There is a difference between Disk Witness and File Share Witness. Disk Witness keeps a backup of the cluster database while File Share Witness does not. Both add a [vote to the cluster \(p. 35\)](#). You can use Disk Witness if you use iSCSI-based storage. For more about witness options, see [File Share witness vs Disk witness for local clusters](#).

Troubleshoot

If you experience unexpected failovers, first make sure that you are not experiencing networking, service, or infrastructure issues.

1. Check that your nodes are not experiencing network-related issues.
2. Check driver updates. If you are using outdated drivers on your instance, you should update them. Updating your drivers might address bugs and stability issues that might be present in your currently installed version.
3. Check for any possible resource bottlenecks that could cause an instance to become unresponsive, such as CPU and disk I/O. If the node cannot service requests, it might appear to be down by the cluster service.

Security in Microsoft SQL Server on Amazon EC2

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 SQL Server on EC2, 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

For detailed information about how to configure Amazon EC2 to meet your security and compliance objectives, see [Security in Amazon EC2](#) in the *User Guide for Windows Instances*.

Document history for the Microsoft SQL Server on Amazon EC2 User Guide

The following table describes the documentation releases for Microsoft SQL Server on Amazon EC2.

Change	Description	Date
Initial release (p. 42)	Initial release of the Microsoft SQL Server on Amazon EC2 User Guide	August 18, 2022