

AWS Whitepaper

Best Practices for Deploying Microsoft SQL Server on Amazon EC2



Best Practices for Deploying Microsoft SQL Server on Amazon EC2: AWS Whitepaper

Copyright © 2024 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

Abstract and introduction	i
Abstract	1
Introduction	1
Are you Well-Architected?	2
High availability and disaster recovery	4
Availability Zones and multi-AZ deployment	5
Using AWS Launch Wizard to deploy Microsoft SQL Server on Amazon EC2 instances	7
Multi-Region deployments	8
Disaster recovery	10
Performance optimization	12
Using Amazon Elastic Block Store (Amazon EBS)	12
Instance storage	13
Amazon FSx for Windows File Server	14
Bandwidth and latency	15
Read replicas	16
Security optimization	17
Amazon VPC	17
Encryption at rest	17
Encryption in transit	18
Encryption in use	18
AWS Key Management Service (AWS KMS)	18
Security patches	18
Cost optimization	20
Using SQL Server Developer Edition for non-production	20
Amazon EC2 CPU optimization	21
Switch to SQL Server Standard Edition	21
z1d and r5b EC2 instance types	22
Eliminating active replica licenses	23
SQL Server on Linux	25
Operational excellence	26
Observability and root cause analysis	26
Reducing mean time to resolution (MTTR)	27
Patch management	27
Contributors	28

Document history	29
Notices	30
AWS Glossary	31

Best Practices for Deploying Microsoft SQL Server on Amazon EC2

Publication date: July 28, 2021 ([Document history](#))

Abstract

This whitepaper focuses on best practices to attain the most value for the least cost when running [Microsoft SQL Server on AWS](#). Although, for many general-purpose use cases, [Amazon RDS for Microsoft SQL Server](#) (MS SQL) provides an easy and quick solution, this paper focuses on scenarios where you need to push the limits to satisfy your special requirements.

In particular, this paper explains how you can minimize your costs, maximize availability of your SQL Server databases, optimize your infrastructure for maximum performance, and tighten it for security compliance, while enabling operational excellence for ongoing maintenance. The flexibility of AWS services, combined with the power of Microsoft SQL Server, can provide expanded capabilities for those who seek innovative approaches to optimize their applications and transform their businesses.

The main focus of this paper is on the capabilities available in Microsoft SQL Server 2019, which is the most current version at the time of this publication. Existing databases that run on previous versions (2008, 2012, 2014, 2016, and 2017) can be migrated to SQL Server 2019 and run in compatibility mode.

Mainstream and extended support for SQL Server 2000, 2005, and 2008 has been [discontinued by Microsoft](#). Any database running on those versions of SQL Server must be upgraded to a supported version first. Although it is possible to run those versions of SQL Server on AWS, that discussion is outside the scope of this whitepaper.

Introduction

AWS offers the best cloud for SQL Server, and it is the proven, reliable, and secure cloud platform for running Windows-based applications today and in the future. SQL Server on Windows or Linux on Amazon EC2 enables you to increase or decrease capacity within minutes, not hours or days. You can commission one, hundreds, or even thousands of server instances simultaneously.

Deploying self-managed, fully-functioning, and production-ready Microsoft SQL Server instances on Amazon EC2 is now possible within a few minutes for anyone, even those without deep skills on SQL Server and cloud features or configuration nuances, thanks to [AWS Launch Wizard](#) for SQL Server. Using AWS Launch Wizard, you can quickly deploy SQL Server on EC2 Windows or Linux instances, with all the best practices already implemented and included in your deployment.

Independent [benchmarks](#) have proven that SQL Server runs 2X faster with 64% lower costs on AWS when compared with the next biggest cloud provider. AWS continues to be the most preferred option for deploying and running Microsoft SQL Server. This is due to the unique combination of breadth and depth of services and capabilities offered by AWS, providing the optimum platform for MS SQL Server workloads.

Requirements for running SQL Server often fall under following categories:

- High availability and disaster recovery
- Performance
- Security
- Cost
- Monitoring and maintenance

These requirements map directly to the five pillars of the [AWS Well-Architected Framework](#), namely:

- Reliability
- Performance efficiency
- Security
- Cost optimization
- Operational excellence

This paper discusses each of these requirements in further detail, along with best practices using AWS services to address them.

Are you Well-Architected?

The [AWS Well-Architected Framework](#) helps you understand the pros and cons of the decisions you make when building systems in the cloud. The six pillars of the Framework allow you to learn

architectural best practices for designing and operating reliable, secure, efficient, cost-effective, and sustainable systems. Using the [AWS Well-Architected Tool](#), available at no charge in the [AWS Management Console](#), you can review your workloads against these best practices by answering a set of questions for each pillar.

For more expert guidance and best practices for your cloud architecture—reference architecture deployments, diagrams, and whitepapers—refer to the [AWS Architecture Center](#).

High availability and disaster recovery

Every business seeks data solutions that can address their operational requirements. These requirements often translate to specific values of the Recovery Time Objective (RTO), and Recovery Point Objective (RPO). The RTO indicates how long the business can endure database and application outages, and the RPO determines how much data loss is tolerable. For example, an RTO of one hour tells you that, in the event of an application outage, the recovery plans should aim to bring the application back online within one hour. An RPO of zero indicates that, should there be any minor or major issues impacting the application, there should be no data loss after the application is brought back online.

The combination of RTO and RPO requirements dictates what solution should be adopted. Typically, applications with RPO and RTO values close to zero should use a high availability (HA) solution, whereas disaster recovery (DR) solutions can be used for those with higher values. In many cases, HA and DR solutions can be mixed to address more complex requirements.

Microsoft SQL Server offers several HA/DR solutions, each suitable for specific requirements. The following table compares these solutions:

Table 1: HA/DR options in Microsoft SQL Server

Solution	HA	DR	Enterprise edition	Standard edition
Always On availability groups	Yes	Yes	Yes	Yes (2 replicas)*
Always On failover cluster instances	Yes	Yes**	Yes	Yes (2 nodes)
Distributed availability groups	Yes	Yes	Yes	No
Log shipping	No	Yes	Yes	Yes

Solution	HA	DR	Enterprise edition	Standard edition
Mirroring (deprecated)	Yes	Yes	Yes	Yes (Full safety only)

* [Always On basic availability groups](#) in SQL Server 2019 Standard edition support a single passive replicas (in addition to the primary replica) for a single database per availability group. If you need multiple databases in HA mode, a separate availability group needs to be defined for each database.

** MSSQL Failover Cluster Instance is often used as a pure HA solution. However, as discussed later in this document, in AWS the FCI can also serve as a complete HA/DR solution.

These solutions rely on one or more secondary servers, with SQL Server running as active or passive standby. Based on the specific HA/DR requirements, these servers can be located in close proximity to each other or far apart.

In AWS, you can choose between low latency or an extremely low probability of failure. You can also combine these options to create the solution that is most suitable to your use case. This paper looks at these options and how they can be used with SQL Server workloads.

Availability Zones and multi-AZ deployment

[AWS Availability Zones](#) are designed to provide separate failure domains, while keeping workloads in relatively close proximity for low latency communications. Availability Zones are a good solution for synchronous replication of your databases using Mirroring, Always On Availability Groups, Basic Availability Groups, or Failover Cluster Instances. SQL Server provides zero data loss and, when combined with the low-latency infrastructure of Availability Zones, provides high performance.

This is one of the main differences between most on-premises deployments and AWS. For example, Always On Failover Cluster Instance (FCI) is often used inside a single data center, because all nodes in an FCI cluster must have access to the same shared storage. Locating these nodes in different data centers could degrade performance. However, with AWS, FCI nodes can be located in separate Availability Zones and still provide high performance because of the low-latency network link between all Availability Zones within a Region.

This feature enables a higher level of availability and could eliminate the need for a third node, which is often coupled with an FCI cluster for disaster recovery purposes.

SQL Server FCI relies on shared storage being accessible from all nodes participating in FCI. [Amazon FSx for Windows File Server](#) is a fully managed service providing shared storage that automatically replicates the data synchronously across two Availability Zones, provides high availability with automatic failure detection, failover, and failback, and fully supports the [Server Message Block \(SMB\) Continuous Availability \(CA\)](#) feature. This enables you to simplify your SQL Server Always On deployments and use Amazon FSx as storage tier for MS SQL FCI.

Scenarios where Amazon FSx is applicable for performance tuning and cost optimization are discussed in subsequent sections of this document.

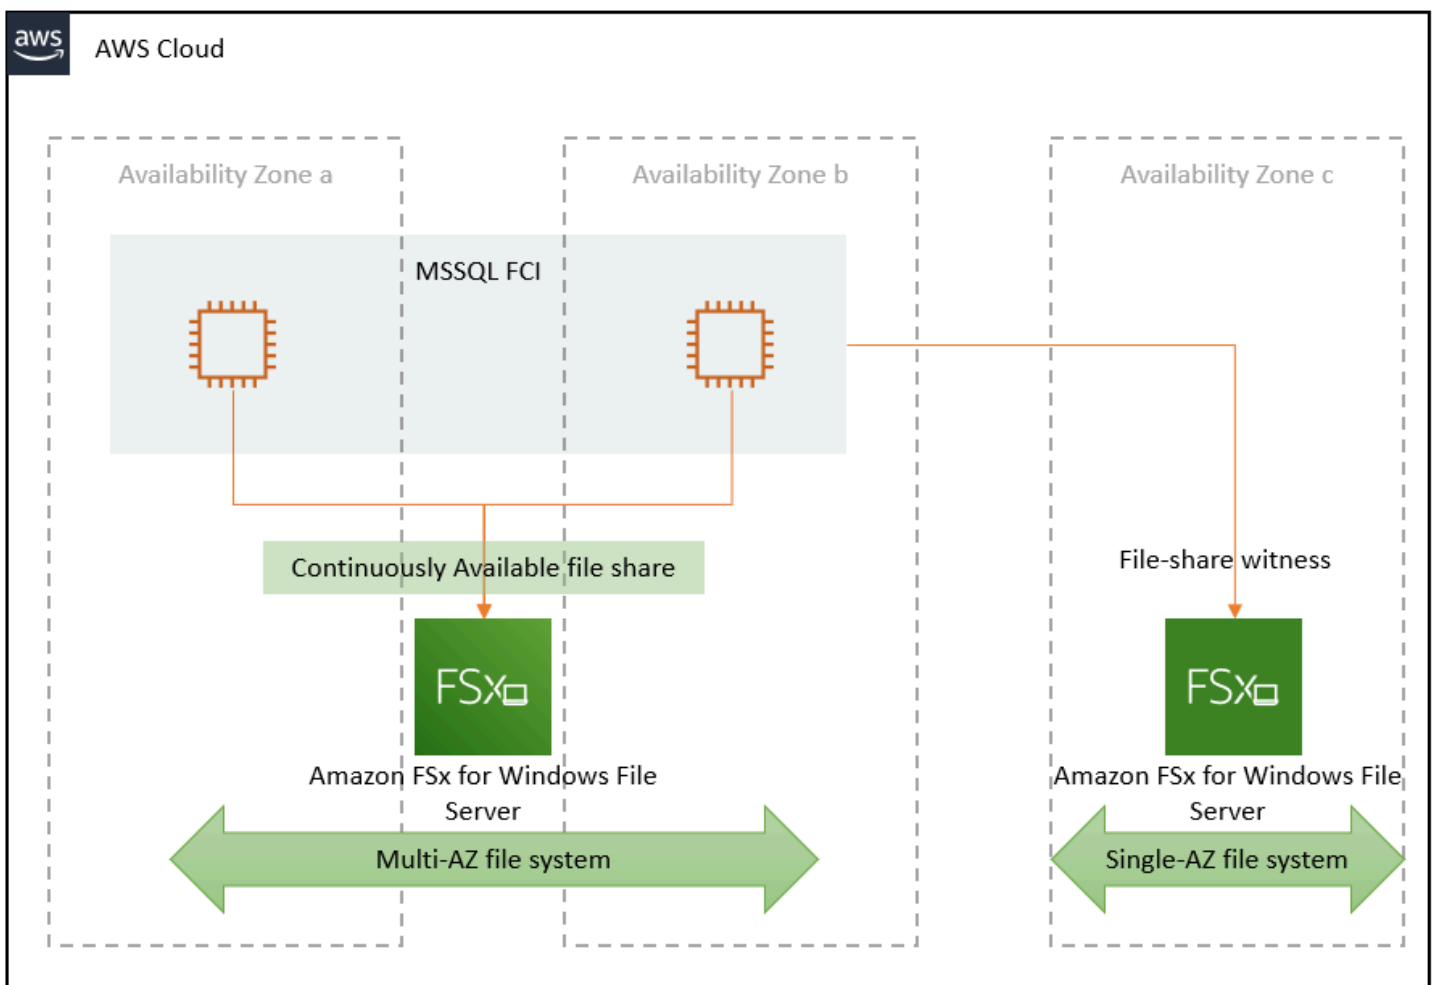


Figure 1: Using Amazon FSx as file-share for Failover Cluster Instance, or as file-share witness in Windows Server Failover Cluster

Using AWS Launch Wizard to deploy Microsoft SQL Server on Amazon EC2 instances

[AWS Launch Wizard](#) is a service that offers a guided way of sizing, configuring, and deploying AWS resources for third party applications, such as Microsoft SQL Server, without the need to manually identify and provision individual AWS resources.

Today, you can use AWS Launch Wizard to deploy Microsoft SQL Server with following configurations:

- SQL Server **single instance** on Windows
- SQL Server **single instance** on Linux
- SQL Server HA using **Always On Availability Groups** on Windows
- SQL Server HA using **Always On Availability Groups** on Linux
- SQL Server HA using **Always On Failover Cluster Instance** on Windows

To start, you input your MS SQL workload requirements, including performance, number of nodes, licensing model, MS SQL edition, and connectivity on the service console. Launch Wizard then identifies the correct AWS resources, such as EC2 instances and EBS volumes, to deploy and run your MS SQL instance. Launch Wizard provides an estimated cost of deployment, and enables you to modify your resources to instantly view an updated cost assessment. After you approve the AWS resources, Launch Wizard automatically provisions and configures the selected resources to create a fully-functioning, production-ready application.

AWS Launch wizard handles all the heavy-lifting, including installation and configuration of Always On Availability Groups or Failover Cluster Instance. This is especially useful with the Linux support, as most MS SQL administrators find Linux configuration non-trivial when done manually.

AWS Launch Wizard also creates CloudFormation templates that can serve as a baseline to accelerate subsequent deployments. For post-deployment management, [AWS Systems Manager \(SSM\) Application Manager](#) automatically imports application resources created by AWS Launch Wizard. From the Application Manager console, you can view operations details, and perform operations tasks. As discussed later in this document, you can also use SSM Automation documents to manage or remediate issues with application components or resources.

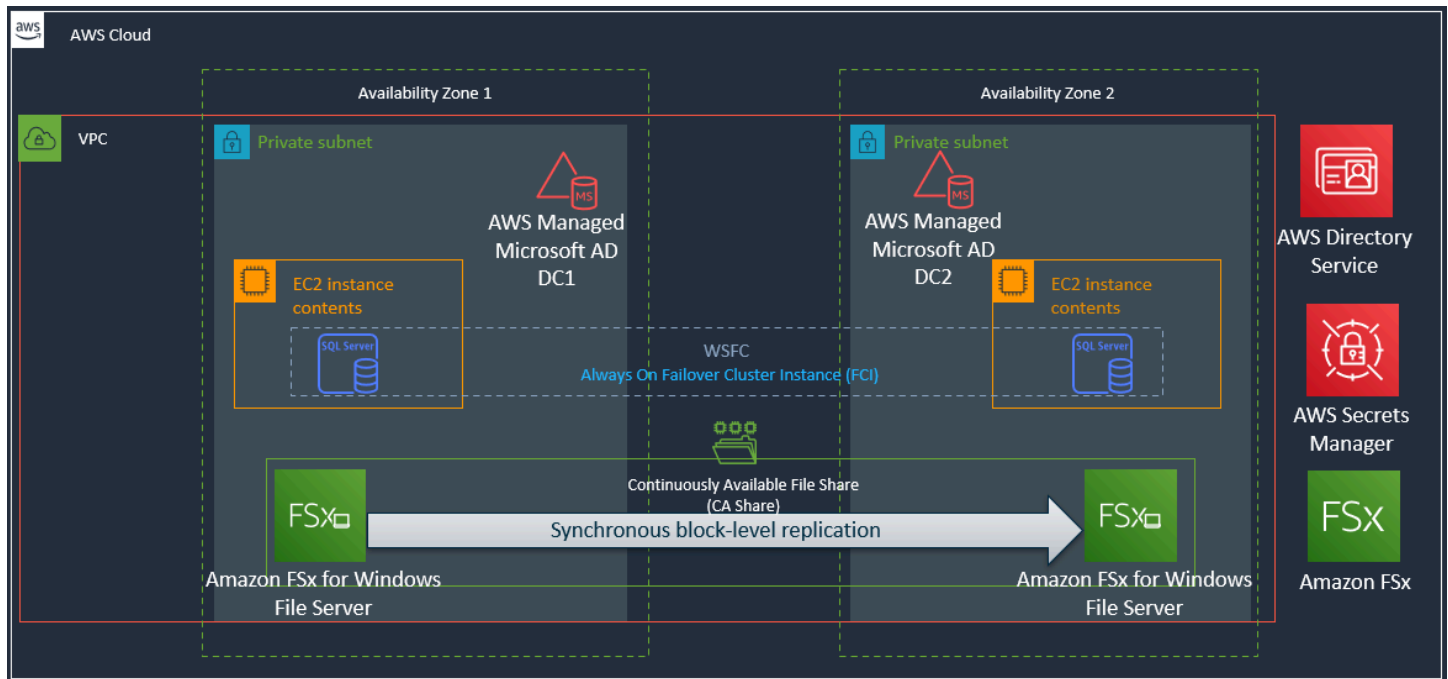


Figure 2: AWS Launch Wizard deploys MS SQL FCI using Amazon FSx for Windows File Server

Multi-Region deployments

For those workloads that require even more resilience against unplanned events, you can leverage the global scale of AWS to ensure availability under almost any circumstances.

By default, Amazon Virtual Private Cloud (Amazon VPC) is confined within a single Region. Therefore, for a multi-region deployment, you need to establish connectivity between your SQL Server instances that are deployed in different Regions. In AWS, there are a number of ways to do this, each suitable for a range of requirements:

- **[VPC peering](#)** — Provides an encrypted network connectivity between two VPCs. The traffic flows through the AWS networking backbone, eliminating latency and other hazards of the internet.
- **[AWS Transit Gateway](#)** — If you need to connect two or more VPCs or on-premises sites, you can use AWS Transit Gateway to simplify management and configuration overhead of establishing network connection between them.
- **[VPN connections](#)** — AWS VPN solutions are especially useful when you need to operate in a hybrid environment, and connect your AWS VPCs to your on-premises sites and clients.
- **[VPC sharing](#)** — If your applications or other clients are spread across multiple AWS accounts, an easy way to make your SQL Server instance available to all of them is using virtual private cloud (VPC) Sharing. A shared VPC can also be connected to other VPCs using [AWS Transit Gateway](#),

[AWS VPN CloudHub](#), VPN connections, or VPC peering. These connections are useful when workloads are spread across multiple accounts and Regions.

If you have applications or users that are deployed in remote Regions which need to connect to your SQL Server instances, you can use the [AWS Direct Connect](#) feature that provides connectivity from any Direct Connect connection to all AWS Regions.

Although it is possible to have synchronous replication in a multi-region SQL Server deployment, the farther apart your selected Regions are, the more severe the performance penalty is for a synchronous replication. Often the best practice for multi-region deployments is to establish an asynchronous replication, especially for Regions that are geographically distant.

For those workloads that come with aggressive RPO requirements, asynchronous multi-Region deployment can be combined with a Multi-AZ or Single-AZ synchronous replication. You can also combine all three methods into a single solution. However, these combinations would impose a significant increase in your SQL Server license costs, which must be considered as part of your planning.

In cases involving several replicas across two or more Regions, [distributed availability groups](#) might be the most suitable option. This feature enables you to combine availability groups deployed in each Region into a larger distributed availability group.

Distributed availability groups can also be used to increase the number of read replicas. A traditional availability group allows up to eight read replicas. This means you can have a total of nine replicas, including the primary. Using a distributed availability group, a second availability group can be added to the first, increasing the total number of replicas to 18. This process can be repeated with a third availability group and a second distributed availability group. The second distributed availability group can be configured to include either the first or second availability groups as its primary. Distributed availability group is the means through which SQL Server Always On can achieve virtually unlimited scale.

Another use case of a distributed availability group is for zero downtime database migrations, when during migration a read-only replica is available at target destination. The independence of SQL Server [Distributed Availability Group](#) from Active Directory and Windows Server Failover Cluster (WSFC) is the main benefactor for these cases. It enables you to keep both sides of the migration synchronized without having to worry about the complexities of Active Directory or WSFC. See [How to architect a hybrid Microsoft SQL Server solution using distributed availability groups](#) for more details.

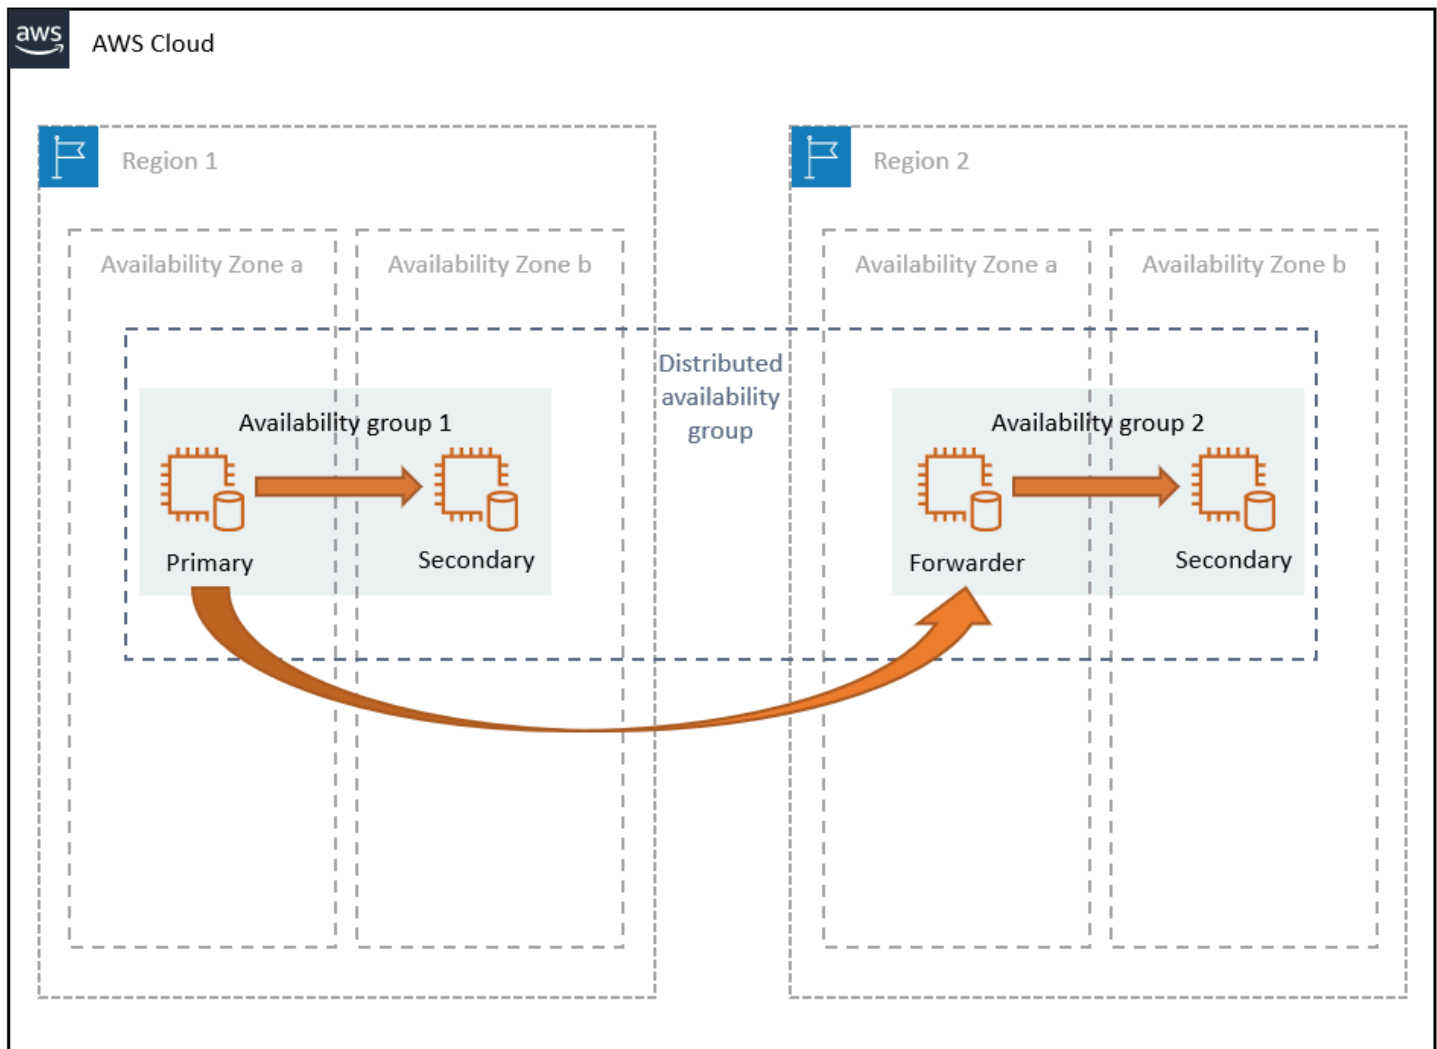


Figure 3: SQL Server distributed availability group in AWS

Disaster recovery

Similar to HA solutions, disaster recovery (DR) solutions require a replica of SQL Server databases in another server. However, for DR, the other server is often in a remote site far away from the primary site. This means higher latency, and therefore, lower performance if you rely on HA solutions that use synchronous replication.

DR solutions often rely on asynchronous replication of data. Similar to HA, DR solutions are based on either block-level or database-level replication. For example, SQL Server Log Shipping replicates data at the database-level, while Windows Storage Replica can be used to implement block-level replication.

DR solutions are selected based on their requirements, such as cost, RPO, RTO, complexity, and the effort to implement each solution.

In addition to common SQL Server DR solutions, such as Log Shipping and Windows Storage Replica, AWS also provides [CloudEndure Disaster Recovery](#). You can use CloudEndure Disaster Recovery to reduce downtime to a few minutes, protect against data loss for sub-second RPO, simplify implementation, increase reliability, and decrease the total cost of ownership.

CloudEndure is an agent-based solution that replicates entire virtual machines, including the operating system, all installed applications, and all databases, into a staging area. The staging area contains low-cost resources automatically provisioned and managed by CloudEndure Disaster Recovery. This greatly reduces the cost of provisioning duplicate resources.

Because the staging area does not run a live version of your workloads, you don't need to pay for duplicate software licenses or high-performance compute. Rather, you pay for low-cost compute and storage. The fully provisioned recovery environment, with the right-sized compute and higher-performance storage required for recovered workloads, is launched only during a disaster or drill.

AWS also makes [CloudEndure](#) available at no additional cost for migration projects.

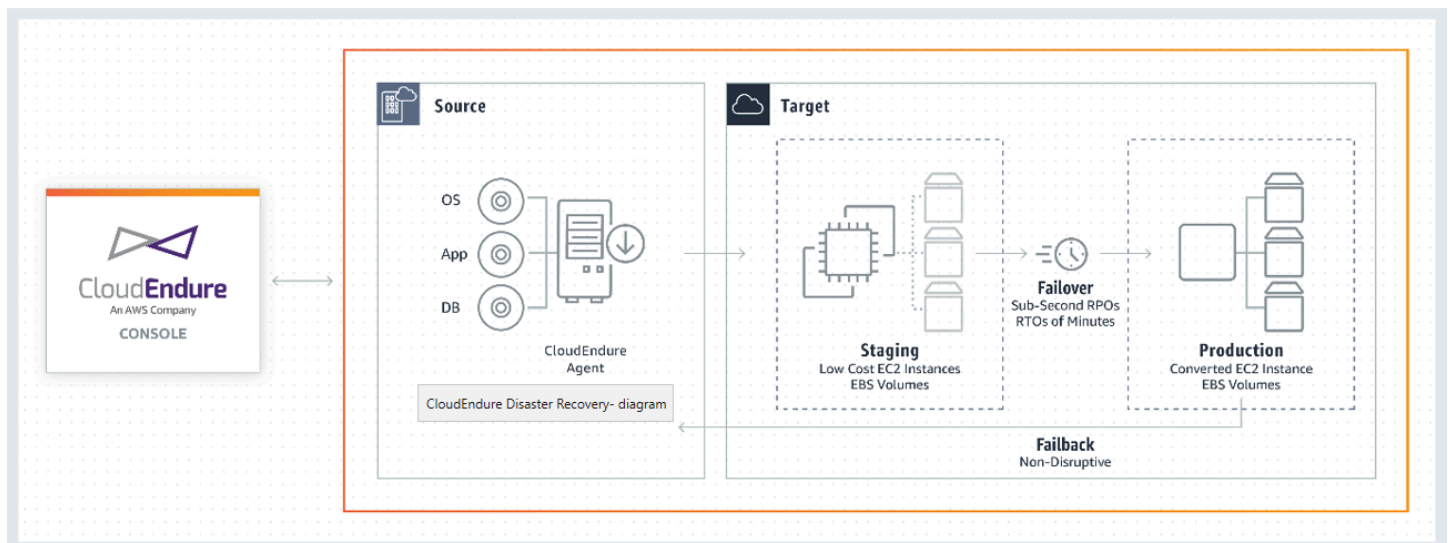


Figure 4: CloudEndure disaster recovery

Performance optimization

In some cases, maximizing performance might be your utmost priority. Both SQL Server and AWS have several options to substantially increase performance of your workloads.

Using Amazon Elastic Block Store (Amazon EBS)

[Amazon EBS](#) is a Single-AZ block storage service with a number of flexible options to cater to diverse requirements. When it comes to maximizing performance with consistent and predictable results on a single volume, using a [Provisioned IOPS](#) Solid State Drive (SSD) volume type (io2 and io2 Block Express) is the easiest choice. You can provision up to 64,000 input/output operations per second (IOPS) per io2 EBS volume (based on 16 KiB I/O size), along with 1000-MiB/s throughput. For more demanding workloads, the io2 [Block Express](#) EBS volumes guarantee 256,000 IOPS and 4,000 MiB/s throughput per volume.

If you need more IOPS and throughput than provided by a single EBS volume, you can create multiple volumes and stripe them in your Windows or Linux instance (Microsoft SQL Server 2017 and later can be installed on both Windows and Linux systems). Striping enables you to further increase the available IOPS per instance up to 260,000, and throughput per instance up to 7,500 MB/s.

Remember to use [EBS-optimized](#) EC2 instance types. This means a dedicated network connection is allocated to serve requests between your EC2 instance and the EBS volumes attached to it.

While you can use a single Provisioned IOPS (io1, io2, or io2 Block Express) volume to meet your IOPS and throughput requirements, General Purpose SSD (gp2 and gp3) volumes offer a better balance of price and performance for SQL Server workloads when configured appropriately.

General Purpose SSD (gp2) volumes deliver single-digit millisecond latencies and the ability to burst to 16,000 IOPS for extended periods. This ability is well suited to SQL Server. The IOPS load generated by a relational database like SQL Server tends to spike frequently. For example, table scan operations require a burst of throughput, while other transactional operations require consistent low latency.

One of the major benefits of using EBS volumes is the ability to create point-in-time and instantaneous EBS snapshots. This feature copies the EBS snapshot to [Amazon Simple Storage Service](#) (Amazon S3) infrastructure, which provides 99.999999999% durability. Despite EBS volumes being confined to a single AZ, EBS snapshots can be restored to any AZ within the same

Region. Note that block-level snapshots are not the same as database backups, and not all features of database backups are attainable this way. Therefore, this method is often combined and complemented with a [regular database backup plan](#).

Although each EBS volume can be as large as 64 TB, and therefore, could take a long time to transfer all its data to Amazon S3, [EBS snapshots](#) are always point-in-time. This means SQL Server and other applications can continue reading and writing to and from the EBS volume while data is being transferred in the background.

When you restore a volume from a snapshot, the volume is immediately available to applications for read and write operations. However, it takes some time until it gets to its full performance capacity. Using [Amazon EBS fast snapshot restore](#), you can eliminate the latency of input/output (I/O) operations on a block when it is accessed for the first time. Volumes created using fast snapshot restore instantly deliver all of their provisioned performance.

You can use [AWS Systems Manager Run Command](#) to take [application-consistent EBS snapshots](#) of your online SQL Server files at any time, with no need to bring your database offline or in read-only mode. The snapshot process uses Windows Volume Shadow Copy Service (VSS) to take image-level backups of VSS-aware applications. Microsoft SQL Server is VSS-aware and is perfectly compatible with this technique. It is also possible to take VSS snapshots of Linux instances, however, that process requires some manual steps, because Linux does not natively support VSS.

You can also take [crash-consistent EBS snapshots](#) across multiple EBS volumes, attached to a Windows or Linux EC2 instance, without using orchestrator applications. Using this method, you only lose uncommitted transactions and writes that are not flushed to the disk. SQL Server is capable of restoring databases to a consistent point before the crash time. This feature is also [supported through AWS Backup](#).

EBS volumes are simple and convenient to use, and in most cases effective, too. However, there might be circumstances where you need even higher IOPS and throughput than what is achievable using Amazon EBS.

Instance storage

[Storage-optimized EC2 instance types](#) use fixed-size local disks and a variety of different storage technologies are available. Among these, Non-Volatile Memory express (NVMe) is the fastest technology with the highest IOPS and throughput. The i3 class of instance types provides NVMe SSD drives. For example, i3.16xlarge, comes with eight disks, each with 1.9 TB of SSD storage.

When selecting storage-optimized EC2 instance types for maximum performance, it is essential to understand that some of the smaller instance types provide instance storage that is shared with other instances. These are virtual disks that reside on a physical disk attached to the physical host. By selecting a bigger instance type, such as i3.2xlarge, you ensure that there is a 1:1 correspondence between your instance store disk and the underlying physical disk. This ensures consistent disk performance and eliminates the noisy-neighbor problem.

Instance disks are ephemeral and live only as long as their associated EC2 instance. If the EC2 instance fails, or is stopped or ended, all of its instance storage disks are wiped out and the data stored on them is irrecoverable. Unlike EBS volumes, instance storage disks cannot be backed up using a snapshot. Therefore, if you choose to use EC2 instance storage for your permanent data, you need to provide a way to increase its durability.

One suitable use for instance storage may be the [tempdb](#) system database files because those files are recreated each time the SQL Server service is restarted. SQL Server drops all tempdb temporary tables and stored procedures during shut down. As a best practice, the tempdb files should be stored on a fast volume, separate from user databases. For the best performance, ensure that the tempdb data files within the same filegroup are the same size and stored on striped volumes.

Another use for EC2 instance storage is the [buffer pool extension](#). Buffer pool extension is available on both the Enterprise and Standard editions of Microsoft SQL Server. This feature uses fast random-access disks (SSD) as a secondary cache between RAM and persistent disk storage, striking a balance between cost and performance when running workloads on SQL Server.

Although instance storage disks are the fastest available to EC2 instances, their performance is capped at the speed of the physical disk. You can go beyond the single disk maximum by striping across several disks.

You could also use instance storage disks as the cache layer in [Storage Spaces](#) (for single Windows instances) and [Storage Spaces Direct](#) (for Windows Server failover clusters) storage pools.

Amazon FSx for Windows File Server

[Amazon FSx for Windows File Server](#) is another storage option for SQL Server on Amazon EC2. This option is suitable for three major use-cases:

- As shared storage used by SQL Server nodes participating in a Failover Cluster Instance

- As file-share witness to be used with any SQL Server cluster on top of Windows Server Failover Cluster
- As an option to attain higher throughput levels than available in dedicated [EBS optimization](#)

The first two cases were discussed in an earlier section of this document. To better understand the third case, notice that EBS throughput depends on EC2 instance size. Smaller EC2 instance sizes provide lower EBS throughput; therefore, to attain EBS higher throughput, you need bigger instance sizes. However, higher instance sizes cost more. If a workload leaves a big portion of its network bandwidth unused, but requires higher throughput to access underlying storage, using a shared file system over SMB may unlock its required performance, while reducing cost by using smaller EC2 instance sizes.

Amazon FSx provides fast [performance](#) with baseline throughput up to 2 GB/second per file system, hundreds of thousands of IOPS, and consistent sub-millisecond latencies. To provide the right performance for your SQL instances, you can choose a throughput level that is independent of your file system size. Higher levels of throughput capacity also come with higher levels of IOPS that the file server can serve to the SQL Server instances accessing it. The storage capacity determines not only how much data you can store, but also how many I/O operations per second (IOPS) you can perform on the storage – each GB of storage provides three IOPS. You can provision each file system to be up to 64 TB in size.

Bandwidth and latency

When tuning for performance, it is important to remember the difference between latency and bandwidth. You should find a balance between network latency and availability. To gain the highest bandwidth on AWS, you can leverage enhanced networking and [Elastic Network Adapter](#) (ENA), or the new [Elastic Fabric Adapter](#) (EFA) which, when combined with new generation of EC2 instances such as C6gn, C5n, R5n, I3en, or G4dn instances, can provide up to 100Gbps bandwidth. But this quite high bandwidth has no effect on latency.

Network latency changes in direct correlation with the distance between interconnecting nodes. Clustering nodes is a way to increase availability, but placing cluster nodes too close to each other increases the probability of simultaneous failure, reducing availability. Putting them too far apart yields the highest availability, but at the expense of higher latency.

AWS Availability Zones within each AWS Region are engineered to provide a balance that fits most practical cases. Each AZ is engineered to be physically separated from other Availability

Zones, while keeping in close geographic proximity to provide low network latency. Therefore, in the overwhelming number of cases, the best practice is to spread cluster nodes across multiple Availability Zone.

Read replicas

You might determine that many of your DB transactions are read-only queries, and that the sheer number of incoming connections is flooding your database. [Read replicas](#) are a known solution for this situation. You can offload your read-only transactions from your primary SQL Server instance to one or more read replica instances. Read replicas can also be used to perform backup operations, relieving primary instance from performance hits during backup windows. When using availability group listeners, if you mark your connection strings as read-only, SQL Server routes incoming connections to any available read replicas and only sends read/write transactions to the primary instance.

[Always On Availability Groups](#), introduced in SQL Server 2012, supports up to four secondary replicas. In more recent versions of SQL Server (2014, 2016, 2017, and 2019), Always On Availability Groups support one set of primary databases and one to eight sets of corresponding secondary databases.

There might be cases where you have users or applications connecting to your databases from geographically dispersed locations. If latency is a concern, you can locate read replicas close to your users and applications.

When you use a secondary database for read-only transactions, you must ensure that the server software is properly licensed.

Security optimization

Cloud security at AWS is the highest priority, and there are many AWS security features available to you. These features can be combined with the built-in security features of Microsoft SQL Server to satisfy even the most stringent requirements and expectations.

Amazon VPC

There are many features in [Amazon VPC](#) that can help you secure your data in transit. You can use [security groups](#) to restrict access to your EC2 instances and allow only certain endpoints and protocols. You can also use network access control lists to deny known sources of threats.

A best practice is to deploy your SQL Server instances in private subnets inside a VPC, and only allow access to the internet through a VPC network address translation (NAT) gateway, or a custom NAT instance.

Encryption at rest

If you are using EBS volumes to store your SQL Server database files, you have the option to enable block-level encryption. Amazon EBS transparently handles encryption and decryption for you. This is available through a simple check box, with no further action necessary. Amazon FSx for Windows File Server also includes built-in encryption at rest. Both EBS and Amazon FSx are integrated with AWS Key Management Service (AWS KMS) for managing encryption keys. This means, through AWS KMS, you can either use keys provided by AWS, or bring your own keys. For more information, see the [AWS KMS documentation](#).

At the database level, you can use [SQL Server Transparent Data Encryption](#) (TDE), a feature available in Microsoft SQL Server that provides transparent encryption of your data at rest. TDE is available on Amazon RDS for SQL Server, and you can also enable it on your SQL Server workloads on EC2 instances.

Previously, TDE was only available on SQL Server Enterprise Edition. However, SQL Server 2019 has also made it available on Standard Edition. If you want to have encryption-at-rest for your database files on Standard Edition on an earlier version of SQL Server, you can use EBS encryption instead.

It's important to understand the tradeoffs and differences between EBS encryption and TDE. EBS encryption is done at the block level, that is, data is encrypted when it is stored and decrypted

when it is retrieved. However, with TDE, the encryption is done at the file level. Database files are encrypted, and can only be decrypted using the corresponding certificate.

For example, this means if you use EBS encryption without TDE and copy your database data or log files from your EC2 instance to an S3 bucket that does not have encryption enabled, the files will not be encrypted. Furthermore, if someone gains access to your EC2 instance, database files will be exposed instantly.

However, there is no performance penalty when using EBS encryption, whereas enabling TDE adds additional drag on your server resources.

Encryption in transit

As a best practice, you can enable encryption in transit for your SQL Server workloads using the SSL/TLS protocol. Microsoft SQL Server supports [encrypted connections](#), and SQL Server workloads in AWS are no exception. When using SMB protocol for SQL Server storage layer, [Amazon FSx automatically encrypts](#) all data in transit using SMB encryption as you access your file system, without the need for you to modify SQL Server or other applications' configurations.

Encryption in use

Microsoft SQL Server offers [Always Encrypted](#) to protect sensitive data using client certificates. This provides a separation between those who own the data and can view it, and those who manage the data but should have no access. This feature is also available on both Amazon RDS for SQL Server, as well as SQL Server workloads on Amazon EC2.

AWS Key Management Service (AWS KMS)

AWS KMS is a fully managed service to create and store encryption keys. You can use KMS-generated keys or bring your own keys. In either case, keys never leave AWS KMS and are protected from any unauthorized access. You can use KMS keys to encrypt your SQL Server backup files when you store them on [Amazon S3](#), [Amazon S3 Glacier](#), or any other storage service. Amazon EBS encryption also integrates with AWS KMS.

Security patches

One of the common security requirements is the regular deployment of security patches and updates. In AWS, you can use [AWS Systems Manager Patch Manager](#) to automate this process. Note

that use cases for Patch Manager are not restricted to security patches. For more details, refer to the [Patch management](#) section of this whitepaper.

Cost optimization

SQL Server can be hosted on AWS through License Included (LI) and Bring Your Own License (BYOL) licensing models. With LI, you run SQL Server on AWS and pay for the licenses as a component of your AWS hourly usage bill. The advantage of this model is that you do not need to have any long-term commitments and can stop using the product at any time and stop paying for its usage.

However, many businesses already have considerable investments in SQL Server licenses and might want to reuse their existing licenses on AWS. This is possible using BYOL:

- If you have Software Assurance (SA), one of its benefits is the [Microsoft License Mobility through Software Assurance program](#). This program enables you to use your licenses on server instances running anywhere, including on Amazon EC2 instances.
- If you don't have SA, you may still be able to use your own licenses on AWS using [Amazon EC2 Dedicated Hosts](#). For more details, consult the licensing section of the [FAQ](#) for Microsoft workloads on AWS to ensure license compliance.

The BYOL option on EC2 Dedicated Hosts can significantly reduce costs, as the number of physical cores on an EC2 host is about half of the total number of vCPU available on that host. However, one of the common challenges with this option is the difficulty of tracking license usage and compliance. [AWS License Manager](#) helps you solve this problem by tracking license usage, and optionally enforcing license compliance, based on your defined license terms and conditions. AWS License Manager is available to AWS customers at no additional cost.

Using SQL Server Developer Edition for non-production

One of the easiest ways to save licensing costs is to use MS SQL Developer Edition for environments that are not going to be used by application end-users. These are typically Dev, Staging, Test, and User-Acceptance Testing (UAT) environments.

For this, you can download SQL Server Developer Edition installation media from the [Microsoft website](#) and install it on your EC2 instances.

SQL Server Developer Edition is equivalent to SQL Server Enterprise Edition, with full features and functionality.

Amazon EC2 CPU optimization

The z1d instance types provide the maximum CPU power, enabling you to reduce the number of CPU cores for compute-intensive SQL Server deployments. However, your SQL Server deployments might not be compute-intensive and require an EC2 instance type that provides intensity on other resources, such as memory or storage.

Because EC2 instance types that provide these resources are also providing a fixed number of cores that might be more than your requirement, the end result can be a higher licensing cost for cores that are not used at all. AWS offers a solution for these situations. You can use [EC2 CPU optimization](#) to reduce the number of cores available to an EC2 instance and avoid unnecessary licensing costs.

Switch to SQL Server Standard Edition

Enterprise-grade features of SQL Server are exclusively available in the Enterprise edition. However, many of these features have also been made available in the Standard edition, enabling you to switch to the Standard edition if you've been using Enterprise edition only for those features. One example is encryption at rest using Transparent Data Encryption (TDE), which is now available in the Standard edition as of SQL Server 2019.

One of the most common reasons for using Enterprise edition has always been its mission-critical HA capabilities. However, there are alternative options that enable switching to Standard edition without degrading availability. You can use these options to cost-optimize your SQL Server deployments.

One option is using [Always On Basic Availability Groups](#). This option is similar to Always On Availability Groups, but comes with a number of limitations. The most important limitation is that you can have only one database in a basic availability group. This limitation rules this option out for applications that rely on multiple databases.

The other option is using Always On Failover Cluster Instance (FCI). Since FCI provides HA at the instance level, it does not matter how many databases are hosted on your SQL Server instance. Traditionally this option was restricted to HA within a single data center. However, as discussed earlier, you can use Amazon FSx for Windows File Server to overcome that limitation. See the [High availability and disaster recovery](#) section of this document.

You can simplify the complexity and cost of running Microsoft SQL FCI deployments using Amazon FSx in the following scenarios:

- Due to the complexity and cost of implementing a shared storage solution for FCI, you might have opted to use availability groups and SQL Server Enterprise Edition. However, you can now switch to Standard edition and significantly reduce your licensing costs, and also simplify the overall complexity of your SQL deployment and ongoing management.
- You might already use SQL Server FCI with shared storage using a third-party storage replication software solution. That implies that you purchased a license for the storage replication solution, and then deployed, administered, and maintained the shared storage solution yourself. You can now switch to using a fully managed shared storage solution with Amazon FSx, simplifying and reducing costs for your SQL Server FCI deployment.
- You ran your SQL Server Always On deployment on-premises, using a combination of FCI and AG-FCI to provide high availability within your primary data center site, and AG provided a DR solution across sites. The combination of Availability Zones and the support in Amazon FSx for highly available shared storage deployed across multiple Availability Zones now makes it possible for you to eliminate the need for separate HA and DR solutions, reducing costs as well as simplify deployment complexities.

For a more detailed discussion of Microsoft SQL Server FCI deployments using Amazon FSx, see the blog post [Simplify your Microsoft SQL Server high availability deployments using Amazon FSx for Windows File Server](#).

z1d and R5b EC2 instance types

The high-performance z1d and R5b instance type is optimized for workloads that carry high licensing costs, such as Microsoft SQL Server and Oracle databases. The z1d instance type is built with a custom Intel Xeon Scalable Processor that delivers a sustained all core Turbo frequency of up to 4.0 GHz, which is significantly faster than other instances. The R5b uses similar technology with 3.1 to 3.5 GHz. For workloads that need faster sequential processing, you can run fewer cores with a z1d instance and get the same or better performance as running other instances with more cores.

For example, moving an SQL Server Enterprise workload running on an r4.4xlarge to a z1d.3xlarge can deliver up to 24% in savings because of licensing fewer cores. The same savings are realized when moving a workload from an r4.16xlarge to a z1d.12xlarge as it is the same 4:3 ratio.

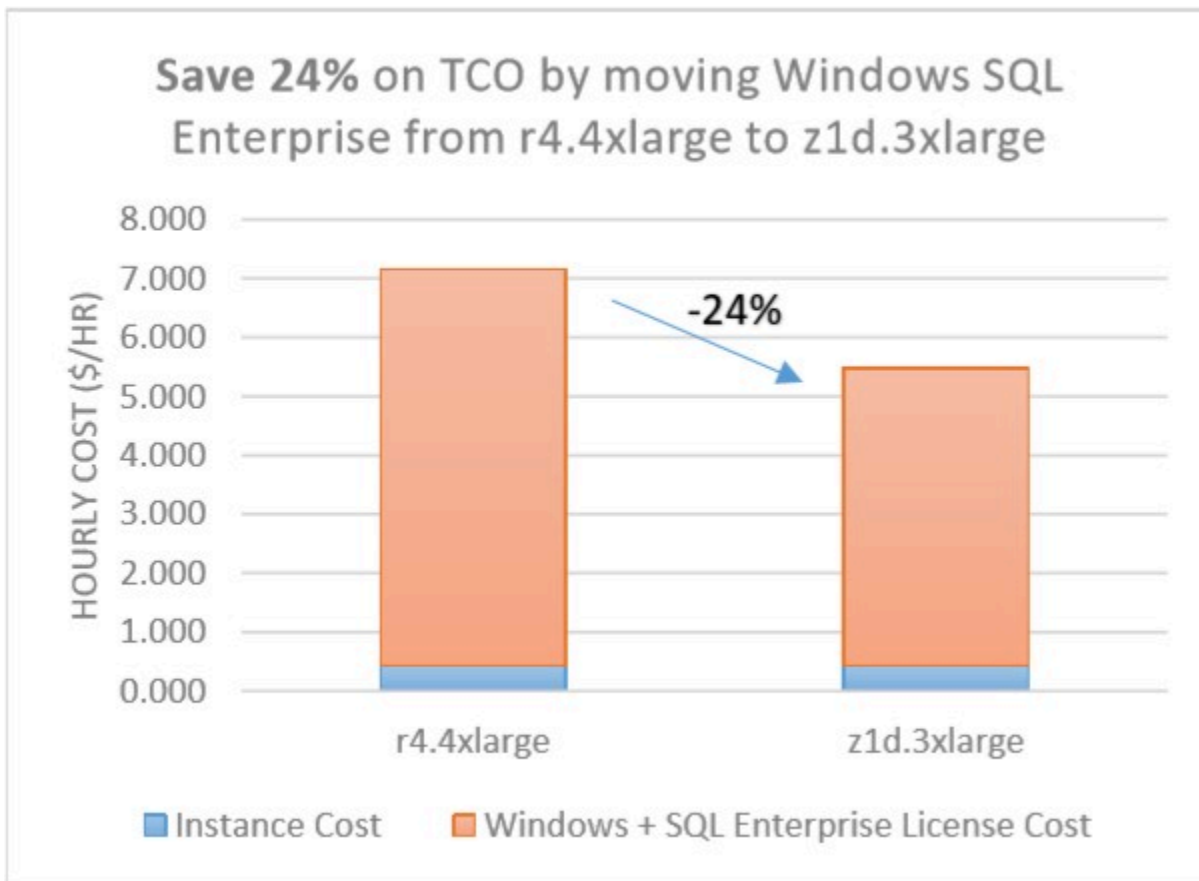


Figure 5: TCO comparison between SQL Server on r4.4xlarge and z1d.3xlarge

Eliminating active replica licenses

Another opportunity for cost optimization in the cloud is through applying a combination of BYOL and LI models. A common use case is SQL Server Always On Availability Groups with active replicas. Active replicas are used primarily for:

- Reporting
- Backup
- OLAP Batch jobs
- HA

Out of the above four operations, the first three are often performed intermittently. This means you would not need an instance continuously up and dedicated to running those operations. In a traditional on-premises environment, you would have to create an active replica that is

continuously synchronized with the primary instance. This means you need to obtain an additional license for the active replica.

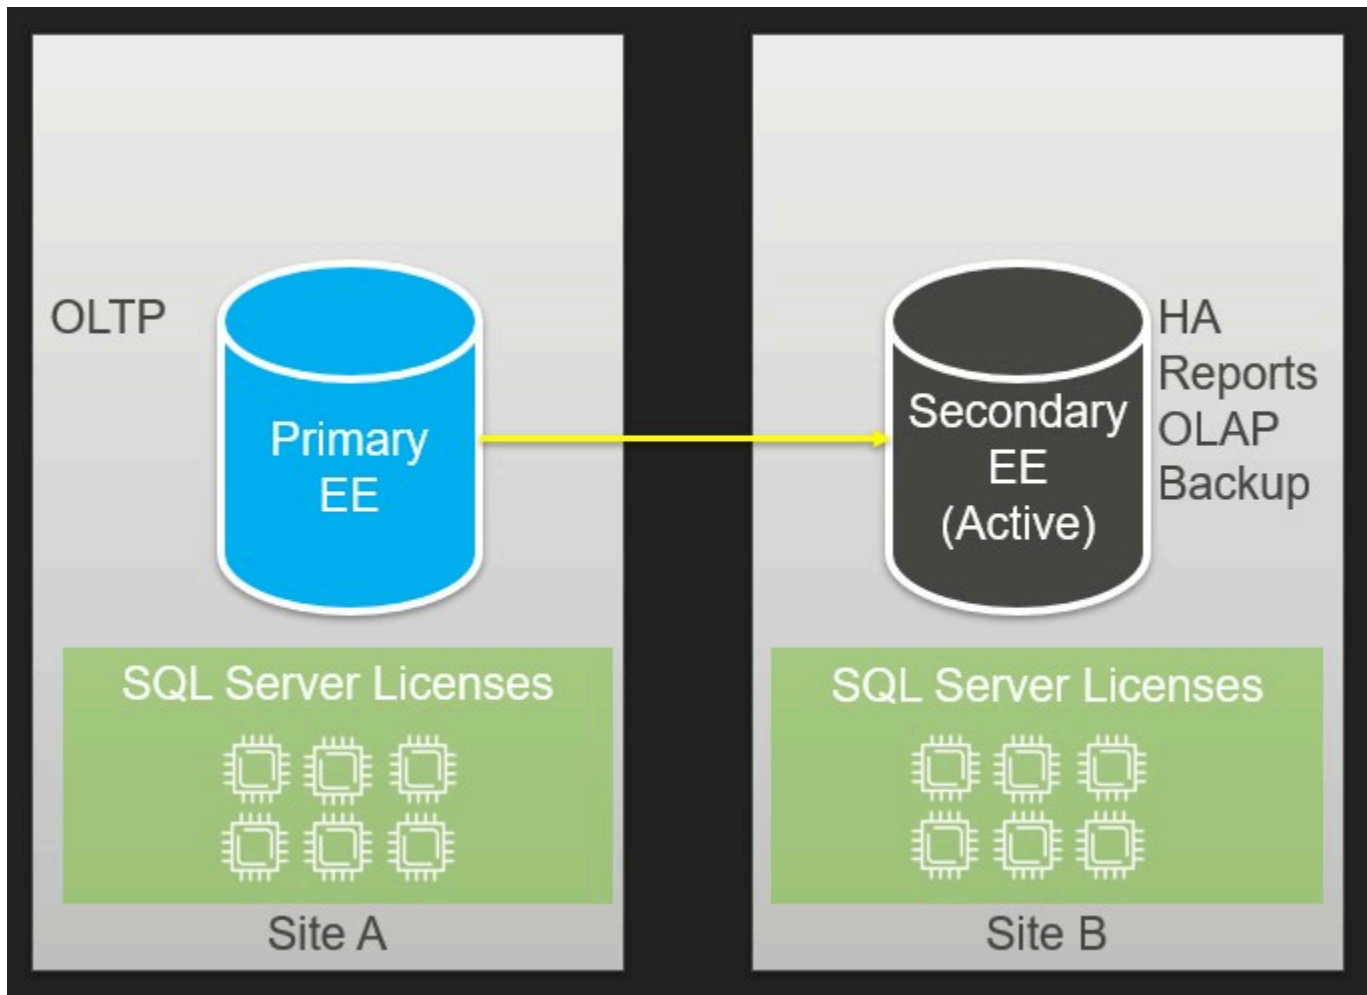


Figure 6: SQL Server active replication on-premises

In AWS, there is an opportunity to optimize this architecture by replacing the active replica with a passive replica, therefore relegating its role solely to the purpose of HA. Other operations can be performed on a separate instance using License Included, which could run for a few hours and then be shut down or ended. The data can be restored through an EBS snapshot of the primary instance. This snapshot can be taken using VSS-enabled EBS snapshots, ensuring no performance impact or downtime on the primary.

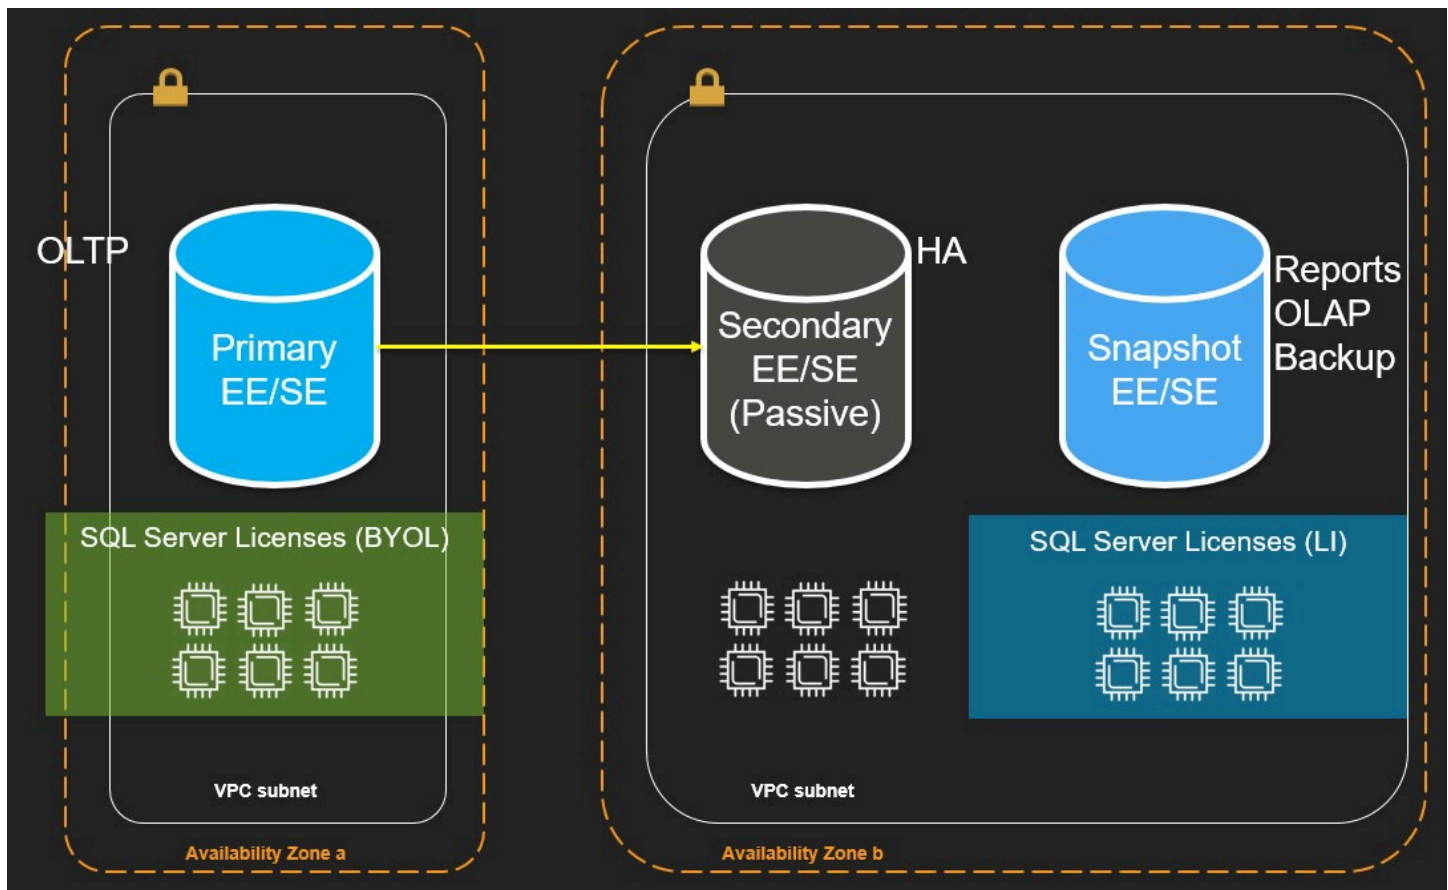


Figure 7: Eliminating active replica licenses in AWS

This solution is applicable when jobs on the active replica run with a low frequency. However, if you need a replica for jobs that run continuously or at a high frequency, consider using [AWS Database Migration Service](#) (AWS DMS) to continuously replicate data from your primary instance into a secondary. The primary benefit of this method is because you can do it using SQL Server Standard edition, it avoids the high cost of SQL Server Enterprise edition licensing.

Refer to the [AWS Microsoft licensing page](#) for more details on ways you can optimize licensing costs on AWS.

SQL Server on Linux

Deploying SQL Server on Linux is a way to eliminate Windows license costs. Installation and configuration of MS SQL on Linux can be non-trivial. However, as discussed earlier in this document, [AWS Launch Wizard](#) helps simplifying that by taking care of most of the heavy lifting for you.

Operational excellence

Most of the discussions explored in this whitepaper pertain to the best practices available for deploying Microsoft SQL Server in AWS. However, another crucial dimension is operating and maintaining these workloads post-deployment.

As a general principle, the best practice is to assume that failures and incidents happen all the time. It's important to be prepared and equipped to respond to these incidents. This objective is composed of three subobjectives:

- Observe and detect anomaly
- Detect the root cause
- Act to resolve the problem

AWS provides tools and services for each of these purposes.

Observability and root cause analysis

[Amazon CloudWatch](#) is a service that enables real-time monitoring of AWS resources and other applications. You can use CloudWatch to collect and track metrics, which are variables you can measure for your resources and applications.

[Amazon CloudWatch Application Insights for .NET and SQL Server](#) is a feature of Amazon CloudWatch that is designed to enable operational excellence for Microsoft SQL Server and .NET applications. Once enabled, it identifies and sets up key metrics and logs across your application resources and technology stack. It continuously monitors the metrics and logs to detect anomalies and errors, while using artificial intelligence and machine learning (AI/ML) to correlate detected errors and anomalies.

When errors and anomalies are detected, Application Insights generates [CloudWatch Events](#). To aid with troubleshooting, it creates automated dashboards for the detected problems, which include correlated metric anomalies and log errors, along with additional insights to point you to the potential root cause.

Using AWS Launch Wizard, you can choose to enable Amazon CloudWatch for Application Insights with a single click. The AWS Launch Wizard handles all the configuration necessary to make your MSSQL instance observable through [Amazon CloudWatch Application Insights](#).

Reducing mean time to resolution (MTTR)

The automated dashboards generated by Amazon CloudWatch Application Insights help you to take swift remedial actions to keep your applications healthy and to prevent impact to the end users of your application. It also creates [OpsItems](#) so you can resolve problems using [AWS Systems Manager OpsCenter](#).

[AWS Systems Manager](#) is a service that enables you to view and control your infrastructure in AWS, on premises, and in other clouds. OpsCenter is a capability of AWS Systems Manager, designed to reduce the mean time to resolution. OpsCenter also provides [Systems Manager Automation](#) documents (runbooks) that you can use to fully or partially automate resolution of issues.

Patch management

[AWS Systems Manager Patch Manager](#) is a comprehensive patch management solution, fully integrated with native Windows APIs, and supporting Windows Server and Linux operating systems, as well as Microsoft applications, including Microsoft SQL Server.

Systems Manager Patch Manager integrates with [AWS Systems Manager Maintenance Windows](#), allowing you to define a predictable schedule to prevent potential disruption of business operations.

You can also use [AWS Systems Manager Configuration Compliance](#) dashboards to quickly see patch compliance state or other configuration inconsistencies across your fleet.

Conclusion

This whitepaper described a number of best practices for deploying Microsoft SQL Server workloads on AWS. It discussed how AWS services can be used to compliment Microsoft SQL Server features to address different requirements.

AWS offers the greatest breadth and depth of services in the cloud, and Amazon EC2 is the most flexible option for deploying Microsoft SQL Server workloads. Each solution and associated trade-offs may be embraced according to particular business requirements.

The five pillars of AWS Well-Architected Framework (reliability, security, performance, cost optimization, and operational excellence) are explored as applicable to SQL Server workloads and AWS services supporting each requirement are introduced.

Contributors

The following individuals and organizations contributed to this document:

- Sepehr Samiei, Solutions Architect, Amazon Web Services

Document history

To be notified about updates to this whitepaper, subscribe to the RSS feed.

Change	Description	Date
Whitepaper updated	Updated for new AWS services and capabilities supporting Microsoft SQL Server workloads.	July 28, 2021
Whitepaper updated	Updated for new AWS services and capabilities supporting Microsoft SQL Server workloads.	May 1, 2020
Whitepaper updated	Updated for Total Cost of Ownership (TCO), using z1d instance types, and Amazon EC2 CPU Optimization.	March 1, 2019
Initial publication	Whitepaper first published.	September 1, 2018

Note

To subscribe to RSS updates, you must have an RSS plug-in enabled for the browser that you are using.

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

© 2021 Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS Glossary Reference*.