Choosing a DR capability for standard editions of Amazon RDS for Oracle and SQL Server

AWS Prescriptive Guidance

Copyright © 2026 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

AWS Prescriptive Guidance: Choosing a DR capability for standard editions of Amazon RDS for Oracle and SQL Server

Copyright © 2026 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

Introduction **1**
The scope of standard edition database engines 1
Objectives **2**
DR sites and common Amazon RDS DR strategies **3**
Synchronous DR 3
Asynchronous DR options 3
Designing cross-Region DR solutions for Amazon RDS standard editions **5**
Amazon RDS PITR snapshot replication 6
Automated snapshots using AWS Backup 8
Native transaction log backup replication 9
AWS DMS 9
DR solution decision matrix **11**
FAQ **13**
Do the DR solutions explained in this guide apply only to standard editions? 13
After a primary site Multi-AZ failover, what will happen to the cross-Region automated backup replication? 13
Does the native transaction logs solution work for Amazon RDS for Oracle as well? 13
What are the options for automating the failover process in a cross-Region DR setup with standard editions of Amazon RDS commercial DB engines? 13
How can I achieve near real-time replication for cross-Region DR with standard editions of Amazon RDS commercial DB engines? 14
Are there any architectural differences or restrictions when using third-party tools or services for cross-Region DR with standard editions of Amazon RDS commercial DB engines? 14
What are the considerations for network connectivity when implementing cross-Region DR for standard editions of Amazon RDS commercial DB engines? 14
Next steps **15**
Resources **16**
Document history **17**

Choosing a DR capability for standard editions of Amazon RDS for Oracle and SQL Server

Gokul Gunasekaran and Sawan Saxena, Amazon Web Services

In data-driven industries, it's essential to have high availability and swift recovery for the data that is stored in databases. Therefore, you must understand the importance of disaster recovery (DR) strategies, and build a robust solution to recover the databases rapidly in the event of a crisis. A wide range of DR options are available for Amazon Relational Database Service (Amazon RDS) on the Amazon Web Services (AWS) Cloud.

This guide focuses on the DR options for standard editions of Oracle and Microsoft SQL Server on Amazon RDS. It covers both active-active and active-passive DR scenarios, and it discusses the benefits and limitations of each option. This guide also provides examples and case studies to help you understand how DR plans can be tailored to your organization's specific needs.

The scope of standard edition database engines

Standard editions of Oracle and SQL Server database engines are widely adopted by customers for their known cost effectiveness and basic use cases. Compared with enterprise editions, standard editions offer the minimal features that are essential to run basic database workloads. If you are managing your workloads with the available limited features in standard editions, you might prefer this for its affordable license costs. If you are migrating your database workloads to the cloud, this is a good time to evaluate whether you can fit your use cases to standard edition versions and reduce costs.

These factors have significantly increased the scope of standard edition database workloads in the cloud and the need for disaster recovery solutions across AWS Regions to increase resiliency.

Objectives

By implementing disaster recovery (DR) solutions to recover your data, your business can achieve the following benefits:

- **Increased resiliency and durability** – Having a strong DR solution for your databases helps ensure that multiple copies of data are available 24/7. This results in maximizing the data availability, and it boosts the durability of the data and the system.
- **Defined RPO and RTO** – Every system in the business will have a different recovery point objective (RPO) and recovery time objective (RTO) based on the criticality of the data or the ease of reproducing the lost data. With a range of DR solutions available, you can implement the one that best suits your system's defined RPO and RTO.
- **Business continuity** – Your DR plan should be a subset of your organization's business continuity plan (BCP). With the right DR strategy employed for your database and the other systems that are part of your application, your business can continue to operate across Availability Zones or even across AWS Regions when disaster strikes.
- **Cost optimization and efficiency** – Implementing these curated DR solutions for standard editions of Oracle and SQL Server helps you to achieve cost efficiency by carefully evaluating and using cost-effective replication tools, cloud-native services, and optimized storage options. This results in minimizing the associated infrastructure and operational expenses while fulfilling your DR requirements. You can implement reliable and optimized DR operations while maximizing resource utilization and minimizing costs.

DR sites and common Amazon RDS DR strategies

A disaster recovery (DR) site is a secondary location used by an organization for restoring its business-critical IT infrastructure and applications when a primary site has been impacted by disaster. DR sites are often built in a remote location to help ensure that the disaster affecting the main site will not also affect the secondary site. DR sites available to you within AWS can be broadly categorized as near DR and far DR:

- Near DR is often referred to as in-Region disaster recovery. The DR solution is configured among the Availability Zones to protect the system when there is an Availability Zone outage or planned maintenance.
- Far DR is often referred to as cross-Region disaster recovery. The solution is configured across AWS Regions to protect the system from Region-wide outages or for conducting planned switchovers to meet compliance policies.

Synchronous DR

In synchronous DR, data is replicated to the DR site at the same time as new data is being created or updated on the primary site. To achieve a synchronous DR solution for standard edition databases, you can use Multi-AZ deployment.

Multi-AZ deployment (near DR) is an AWS managed near-DR solution that provides synchronous replication of your RDS database to a standby instance in a different Availability Zone within the same Region. If the primary database becomes unavailable because of an outage at the Availability Zone level, AWS automatically promotes the standby database to the primary role. This ensures minimal data loss and downtime. However, it's important to note that Multi-AZ deployment does not protect against Region-level outages. It also doesn't provide any DR capabilities outside of the AWS Cloud.

Asynchronous DR options

In asynchronous DR, replication is not performed at the same time as changes are made in the primary. Data is replicated only at the intervals defined by the recovery point objective. To implement an asynchronous DR solution, you can use the following options:

- Automated snapshots managed using AWS Backup – AWS Backup is a fully managed backup service that automates the backup and recovery of your Amazon RDS database based on your point-in-time recovery (PITR) settings. With AWS Backup, you can take snapshots, create backup schedules, retention policies, and backup plans to protect your data against accidental deletion, corruption, or hardware failures. You can restore your database to a new RDS instance in the same AWS Region or in a different Region.

  AWS Backup is an active-passive DR option because you have to manually initiate the restore process in the case of a primary database failure. You can use AWS Backup in conjunction with a Multi-AZ deployment, combining synchronous and asynchronous solutions to provide an additional layer of protection against data loss.
- Amazon RDS PITR snapshot replication – RDS snapshots are a manual DR option. You define a point-in-time snapshot configuration for your RDS database instance, and the snapshots are stored in Amazon Simple Storage Service (Amazon S3). You can then enable cross-Region replication to replicate changes from the primary database to the standby database in a different Region.

  Amazon RDS PITR snapshot replication is an active-passive DR option because you have to manually initiate the restore process and promote the standby database to the primary role if the original primary database fails. However, this option provides more flexibility than Multi-AZ deployment, and you can use it to protect against Region-level outages.

Designing cross-Region DR solutions for Amazon RDS standard editions

When it comes to disaster recovery capabilities, standard editions offer the following cross-Region options:

- Automated snapshots using AWS Backup
- Amazon RDS point-in-time recovery (PITR) snapshot replication
- Native transaction log backup replication
- AWS Database Migration Service (AWS DMS)

It's essential to understand the recovery time objective (RTO) and recovery point objective (RPO) requirements along with the budget while designing cross-Region DR solutions. You can minimize RPO by configuring a combination of synchronous and asynchronous DR strategies and options. However, this will increase infrastructure cost.

It's hard to achieve near zero RPO with Amazon RDS standard editions because there are no synchronous replication options. However, AWS can provide competitive RPO by using various cost-effective options.

The following table summarizes the DR solutions and their availability in each license model, with enterprise editions included to give you a complete view:

- Amazon RDS for Microsoft SQL Server Standard Edition (SE)
- Amazon RDS for Oracle Standard Edition 2 (SE2)
- Enterprise Edition (EE) of Amazon RDS for Oracle and Amazon RDS for SQL Server

Solution	DR site scope	RDS for SQL Server		RDS for Oracle	
		SE	**EE**	**SE2**	**EE**
Multi-AZ deployment	In-Region	Yes	Yes	Yes	Yes

Mounted or read-only replica promotion	In-Region, cross-Region	No	Yes	No	Yes
Automated snapshots using AWS Backup	In-Region, cross-Region	Yes	Yes	Yes	Yes
Amazon RDS PITR snapshot replication	In-Region, cross-Region	Yes	Yes	Yes	Yes
Native transaction log backup replication	In-Region, cross-Region	Yes	Yes	Not Applicable	Not Applicable
AWS DMS	In-Region, cross-Region	Yes	Yes	Yes	Yes

The following sections discuss the options available for cross-Region disaster recovery for standard editions.

Amazon RDS PITR snapshot replication

An Amazon RDS database instance can be configured to replicate snapshots and transaction logs to a destination AWS Region of your choice. After you configure backup replication for a DB instance, Amazon RDS initiates a cross-Region copy of all snapshots and transaction logs as soon as they are ready on the DB instance.

The following diagram shows how Amazon RDS facilitates the automated transfer of snapshots and transaction logs across AWS Regions to maintain the point-in-time recovery (PITR) requirements as configured in the settings of the primary Amazon RDS instance. From the primary

Region, snapshots are copied to the secondary Region. Logs are stored in an S3 bucket in the primary Region and copied to an S3 bucket in the secondary Region.

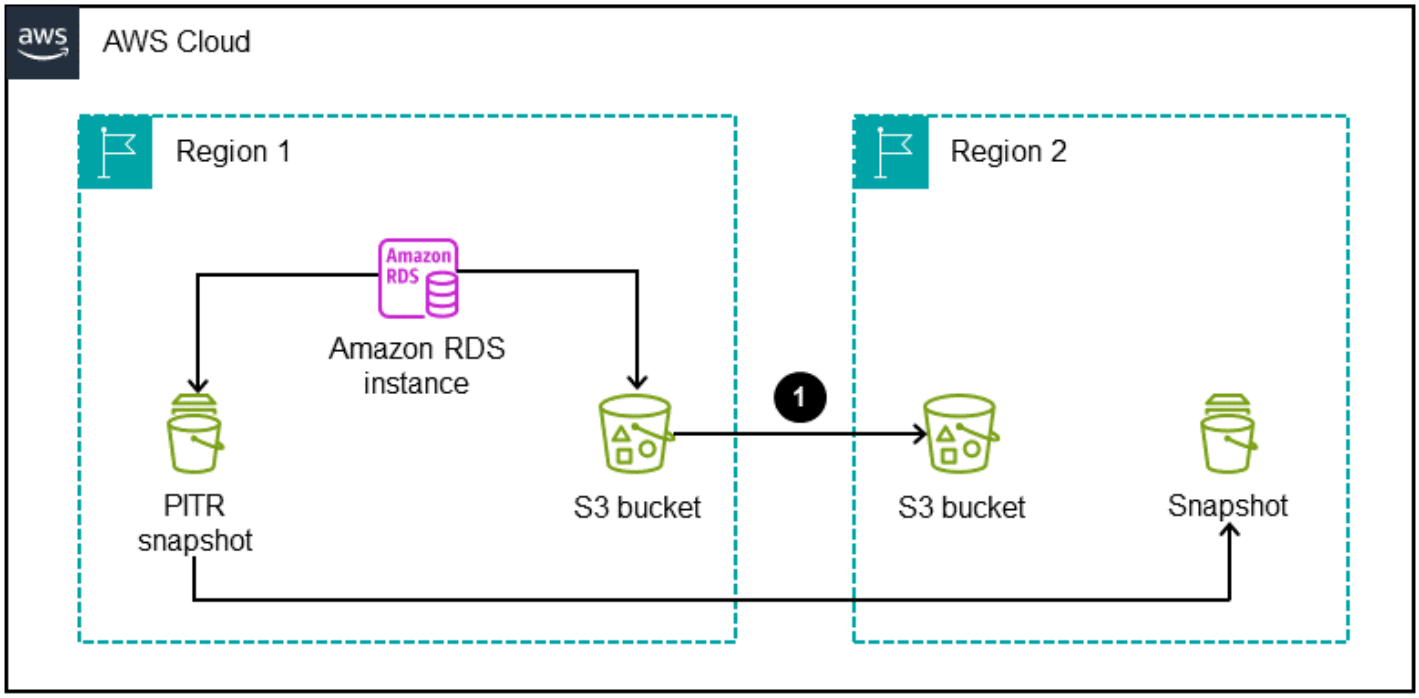

1. Transaction or redo logs

Amazon RDS PITR snapshot replication is based on asynchronous replication, so the underlying RPO is on the slightly higher side, ranging 5–30 minutes. The time depends on redo or transaction log generation and network transfer time.

Recovery time objective (RTO) for the cross-Region Amazon RDS PITR snapshot replication backup strategy is based on the snapshot restore time and recovering the database to a point in time by applying the archived redo or transaction logs. The process can take up to a few hours.

One advantage of this strategy is that it's managed. It's also cost effective because the setup requires no additional infrastructure.

Cross-Region Amazon RDS PITR snapshot replication has limited availability across AWS Regions. Be sure to check availability when you choose your DR Region.

Automated snapshots using AWS Backup

AWS Backup centralizes and automates data protection across AWS services. AWS Backup is a fully managed, policy-based service for data protection at scale. The service is ideal for use cases such as regulatory compliance obligations, business policies for data protection, and business continuity goals.

The following diagram shows using an AWS Backup plan and backup vault to take snapshots of the Amazon RDS instance at scheduled intervals and copy the snapshots to the DR Region.

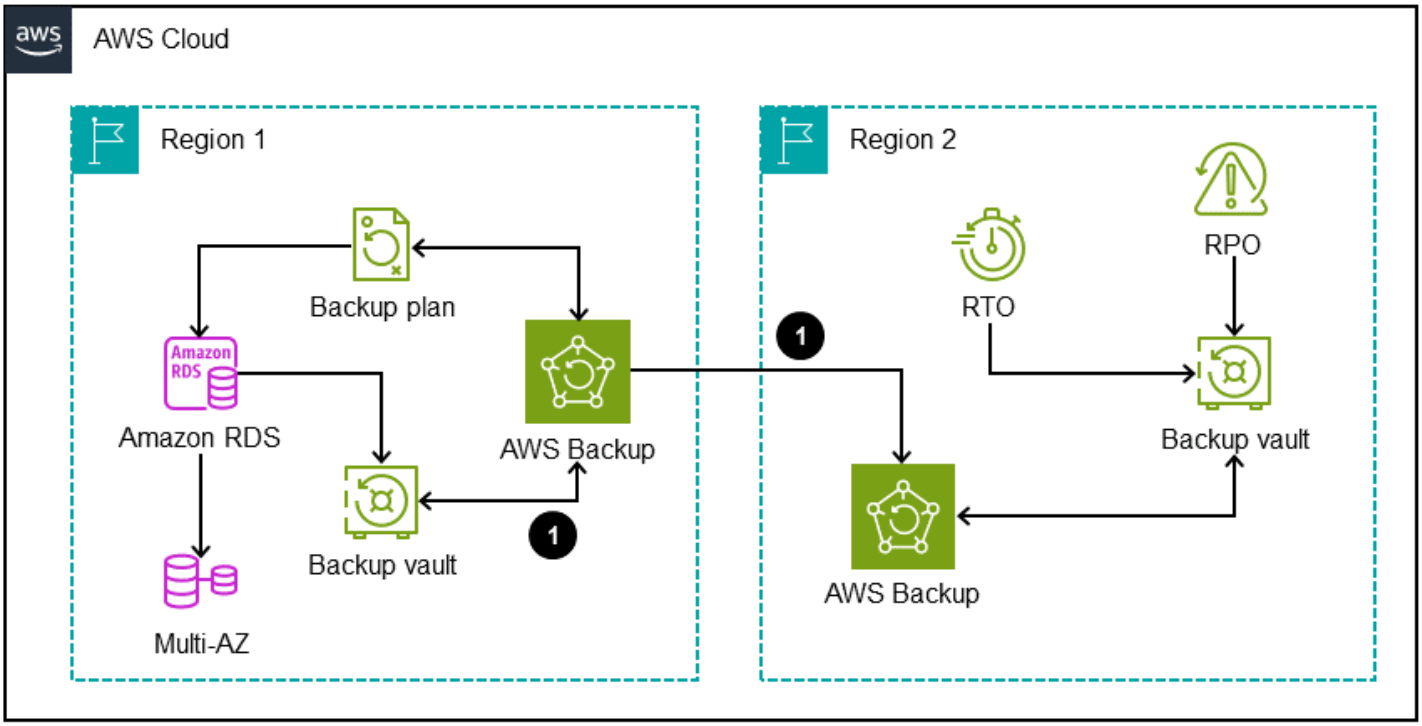

1. Snapshot restore points

AWS Backup doesn't have continuous replication enabled across AWS Regions, which means you will have extended recovery point objective (RPO). The RPO will depend on how frequently you are taking snapshots (defined in backup plan) on the primary database and copying to the standby Region. Because this is a snapshot-based solution, the RPO and recovery time objective (RTO) will depend on the snapshot frequency and restore time. This can vary, and it can be up to a couple hours depending on the database size.

The major benefits of using AWS Backup are the automated backup schedule and retention management. Specific requirements, such as taking a backup once a month or on a particular cadence, can be implemented.

The failover by restoring the snapshot from an AWS Backup service backup vault is manual and not transparent to the application. After you restore the snapshot as a new RDS instance in the standby Region, you must modify the application connection settings.

Native transaction log backup replication

Amazon RDS for SQL Server takes periodic transaction log backups. It provides access to transaction log backups in an AWS managed Amazon Simple Storage Service (Amazon S3) bucket. With this feature, you can access these transaction log backup files and copy them to an S3 bucket in your own account.

This approach provides the cross-Region disaster recovery solution with the lowest recovery point objective (RPO) for Amazon RDS for SQL Server standard editions. An overview of the solution and detailed instructions are explained in the blog post Implement Amazon RDS for SQL Server Standard edition cross-Region disaster recovery using access to transaction log backups feature.

This solution uses logs copy to provide continuous replication across Regions, so you have minimal RPO. The RPO depends on transaction log generation and network transfer time.

The cost associated with this solution is higher because the Amazon RDS for SQL Server instance in the DR must be up and running all the time. Additionally, you will need a tracking instance of Amazon RDS for SQL Server.

The failover is manual and not transparent to the application. After failover, you must modify the application connection settings or DNS values.

AWS DMS

AWS Database Migration Service (AWS DMS) provides ongoing replication of data, keeping the source and target databases in sync. By using AWS DMS as a DR option, you can establish synchronous replication with minimal recovery point objective (RPO). The RPO depends on redo or transaction log generation and network transfer time.

The following diagram shows the ongoing replication of data from the primary Region to the standby Region by using AWS DMS.

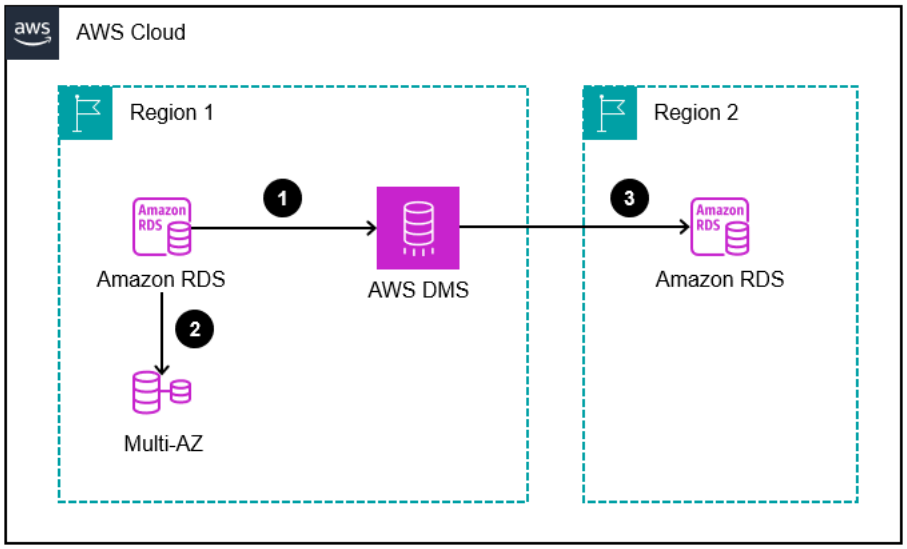

1. In the primary Region, data is copied from the Amazon RDS DB instance to AWS DMS.
2. The Amazon RDS database in the primary Region also uses a Multi-AZ deployment.
3. From AWS DMS, data is loaded onto the standby Amazon RDS database in the secondary Region.

When choosing AWS DMS as a DR option, take note of the following considerations:

- This is a cost-intensive-solution. The cost associated mainly depends on the type of replication instance that you choose and the storage allocated. There is also a small cost for data transfer.
- AWS DMS replicates only a limited amount of data definition language (DDL) statements. AWS DMS doesn't propagate items such as indexes, users, privileges, stored procedures, and other database changes not directly related to table data.
- Before you deploy AWS DMS as your DR solution, consider the data type involved. AWS DMS has limitations for some data types. Also consider the types of transactions happening on the primary RDS DB instance. For more information, see the AWS DMS documentation:
  - Limitations on using Oracle as a source for AWS DMS
  - Limitations on using SQL Server as a source for AWS DMS
  - Limitations on Oracle as a target for AWS DMS
  - Limitations on using SQL Server as a target for AWS DMS
- Ongoing maintenance is required for schema changes.
- The failover using AWS DMS is manual and not transparent. You must modify the application connection settings after failover.

DR solution decision matrix

Use the following table to review RPO, RTO, licensing and other features of disaster recovery options for Amazon RDS for Oracle and SQL Server and find the solutions that best fit your business requirements.

Feature	RPO (approximate)	RTO (approximate)	Licensing	Readable standby	Automatic failover	Endpoint change in failover	Setup and maintenance complexity
Multi-AZ deployment (in-Region)	0	1–2 minutes	SE, SE2, and EE	No	Yes	No	Low
Amazon RDS PITR snapshot replication (cross-Region)	25 minutes	Hours	SE, SE2, and EE	No	No	Yes	Low
Automated snapshots using AWS Backup (cross-Region)	Hours	Hours	SE, SE2, and EE	No	No	Yes	Medium

Native transaction log backup replication (SQL Server only)	5–15 minutes	Minutes	SE and EE	Yes	No	Yes	High
AWS DMS	Secondsor minutes	Minutes	SE, SE2, and EE	Yes	No	Yes	High
Replica promotion (in-Region)	0–5 Minutes	Minutes	EE only	Yes	No	Yes	Low
Automated snapshots using AWS Backup (in-Region)	5 minutes	Hours	SE, SE2, and EE	No	No	Yes	Medium
Replica promotion (cross-Region)	0–15 Minutes	Minutes	EE only	Yes	No	Yes	Medium

FAQ

Find answers to frequently raised questions about disaster recovery (DR) solutions for standard editions of Amazon RDS for Oracle and Amazon RDS for SQL Server.

Do the DR solutions explained in this guide apply only to standard editions?

The options given for standard editions can be used for enterprise editions as well if they meet your RTO and RPO requirements.

After a primary site Multi-AZ failover, what will happen to the cross-Region automated backup replication?

Cross-Region automated backup replication works fine after a Multi-AZ failover on the primary site. No manual intervention is needed.

Does the native transaction logs solution work for Amazon RDS for Oracle as well?

Amazon RDS for Oracle supports accessing the archived redo logs and copying to and from an S3 bucket, but RDS for Oracle can't perform manual recovery using these logs. So this solution applies only to Amazon RDS for SQL Server.

What are the options for automating the failover process in a cross-Region DR setup with standard editions of Amazon RDS commercial DB engines?

Automating the failover process often involves combining multiple AWS services, such as AWS Lambda, Amazon CloudWatch, Amazon EventBridge, and Amazon RDS APIs. You can use these services to initiate failover actions based on predefined criteria, and you can automate the process for minimal downtime.

How can I achieve near real-time replication for cross-Region DR with standard editions of Amazon RDS commercial DB engines?

Achieving near real-time replication requires careful consideration of factors such as network latency, workload complexity, and the chosen replication mechanism. Using services such as AWS DMS can help achieve this objective.

Are there any architectural differences or restrictions when using third-party tools or services for cross-Region DR with standard editions of Amazon RDS commercial DB engines?

Third-party tools or services might introduce specific architectural requirements or restrictions. It's crucial to review their documentation and ensure compatibility with standard editions of Amazon RDS commercial DB engines before implementation.

What are the considerations for network connectivity when implementing cross-Region DR for standard editions of Amazon RDS commercial DB engines?

Network connectivity between AWS Regions plays a critical role in cross-Region DR. To help ensure efficient and reliable data replication, it's important to evaluate network bandwidth, latency, and any limitations imposed by your network infrastructure.

Next steps

If you are ready to implement a cross-Region disaster recovery (DR) solution for your Amazon RDS for SQL Server or Amazon RDS for Oracle instance, we recommend the following next steps:

1. Define the recovery point objective (RPO) and recovery time objective (RTO) based on your business criticality and continuity plan.
2. Choose the DR Region.
3. Choose the DR solution that suits your RTO, RPO, and Region requirements.
4. Determine your cost requirements by using the AWS Pricing Calculator.
5. Implement the solution and perform failover drills.

Resources

- Implement Amazon RDS for SQL Server Standard edition cross-Region disaster recovery using access to transaction log backups feature (blog post)
- Managed disaster recovery with Amazon RDS for Oracle cross-Region automated backups – Part 1 (blog post)
- Rethink Oracle Standard Edition Two on Amazon RDS for Oracle (blog post)
- Cross-Region automated backups
- Disaster recovery on Amazon RDS for Oracle using AWS DMS (blog post)
- Choosing the right database for your RTO and RPO requirements (guide)
- Architect a Managed Disaster Recovery on Amazon RDS for SQL Server: Part 2 (blog post)
- Cross-Region disaster recovery of Amazon RDS for SQL Server (blog post)

Document history

The following table describes significant changes to this guide.

Change	Description	Date
Initial publication	—	August 25, 2023