



Running MS SQL Server workload on Amazon RDS

Eugene Stepanov

Sr. Database Specialist Solutions Architect

Who am I

Eugene Stepanov

- 20+ years of development experience (web, desktop, middle tier and back-end)
- Focusing on RDBMS since 2005
- Worked with SQL Server since version 7.0
- Worked at number of companies big and small
- Joined AWS one and half years ago
- Prior to AWS worked at Microsoft
- Based in Seattle, WA



Agenda

- Amazon RDS
- EC2 vs. RDS
- Managed experience of RDS
 - Automated Backup, Manual snapshot, Native Backups
 - High Availability
 - Read Scale-out
 - Compute & Storage scaling
 - Performance Insights
- Compute & Storage subsystem
- Active Directory Integration
- Performance Monitoring
- Moving data to and from RDS SQL Server
- New Features of RDS SQL Server
- How to deploy & manage RDS SQL Server

Amazon RDS

Choice of open source and commercial databases

Cloud Native Engine



Automatic fail-over
Backup & recovery
X-region replication

Open Source Engines



RDS Platform

Isolation & security
Industry compliance
Automated patching

Commercial Engines



Advanced monitoring
Routine maintenance
Push-button scaling

Options for Deploying SQL Server on AWS



Amazon RDS for SQL Server

- **Consider RDS first**
- Focus on business value tasks
- High-level tuning asks
- Schema optimization
- No in-house database expertise

| |
|--------------------------|
| Scaling |
| High Availability |
| Database Backups |
| DBMS Patching |
| DBMS Install/Maintenance |
| OS Patching |
| OS Install/Maintenance |
| Power, HVAC, net |

 AWS managed



SQL Server on Amazon EC2

- Need full control over DB instance
- Backups
- Replication
- Clustering
- Options that are not available in RDS

| |
|--------------------------|
| Scaling |
| High Availability |
| Database Backups |
| DBMS Patching |
| DBMS Install/Maintenance |
| OS Patching |
| OS Install/Maintenance |
| Power, HVAC, net |

 Customer managed

SQL Server Features at a Glance



Amazon RDS



Amazon EC2

Versions Supported:

2012 – 2019

All

Editions Supported:

Express, Web, Standard, Enterprise

All

High Availability:

AWS-managed

Self-managed

Encryption:

TDE, Column-level, Always Encrypted, TLS, EBS encryption

Authentication:

Windows & SQL Authentication

BI Stack:

SSIS, SSRS, SSAS (tabular)

All

Backups:

AWS- managed

Self-managed

Maintenance:

AWS- managed

Self-managed

SQL Server EC2 vs. RDS: Which should I use?

| | EC2 | RDS |
|--|-----|-----|
| License included | ✓ | ✓ |
| BYOL | ✓ | |
| Full control over the instance | ✓ | |
| Automated backups | | ✓ |
| AWS-managed Multi-AZ deployment | | ✓ |
| AWS-managed Read scale-out | | ✓ |
| AWS-automated ability to scale Compute & Storage | | ✓ |

Managed Experience

Automated Backups

Point-in-time recovery for your DB instance

- Scheduled daily volume backup of entire instance
- Archive database change logs
- 35-day maximum retention
- Minimal impact on database performance
- In-region and x-region PiTR

DB instance status

available

Multi AZ

Yes

Secondary zone

us-east-1d

Automated backups

Enabled (7 Days)

Latest restore time

March 22, 2018 at 10:25:00 AM
UTC-7

Manual Snapshot

- Triggered by customer
- Instance level
- Keep as long as needed

The screenshot displays the Amazon RDS console interface for managing snapshots. On the left is a navigation sidebar with options like Dashboard, Databases, Query Editor, Performance Insights, Snapshots (highlighted), Automated backups, Reserved instances, Proxies, Subnet groups, Parameter groups, Option groups, Events, Event subscriptions, and Recommendations (2). The main content area is titled 'Snapshots' and has tabs for Manual, System, Shared with me, Public, Backup service, and Exports in Amazon S3. The 'Manual' tab is active, showing 'Manual snapshots (1)'. A search bar 'Filter manual snapshots' is present. Below it is a table with columns for Snapshot name, DB instance or cluster, and Snapshot creation time. One snapshot is listed: 'mysnap' for 'database-5', created on 'Thu Apr 02 2021'. An 'Actions' dropdown menu is open over the snapshot, listing 'Restore Snapshot', 'Copy Snapshot', 'Share Snapshot', 'Migrate snapshot', and 'Delete Snapshot'. A 'Take snapshot' button is located in the top right corner of the table area.

| <input checked="" type="checkbox"/> | Snapshot name | DB instance or cluster | Snapshot creation time |
|-------------------------------------|---------------|------------------------|------------------------|
| <input checked="" type="checkbox"/> | mysnap | database-5 | Thu Apr 02 2021 |

Native Backups

- Backup and restore directly to and from S3 bucket
- Supports Compression
- Only full & diff backups (no t-log)
- Full, diff & t-log restores
- Multi-file backup/restore

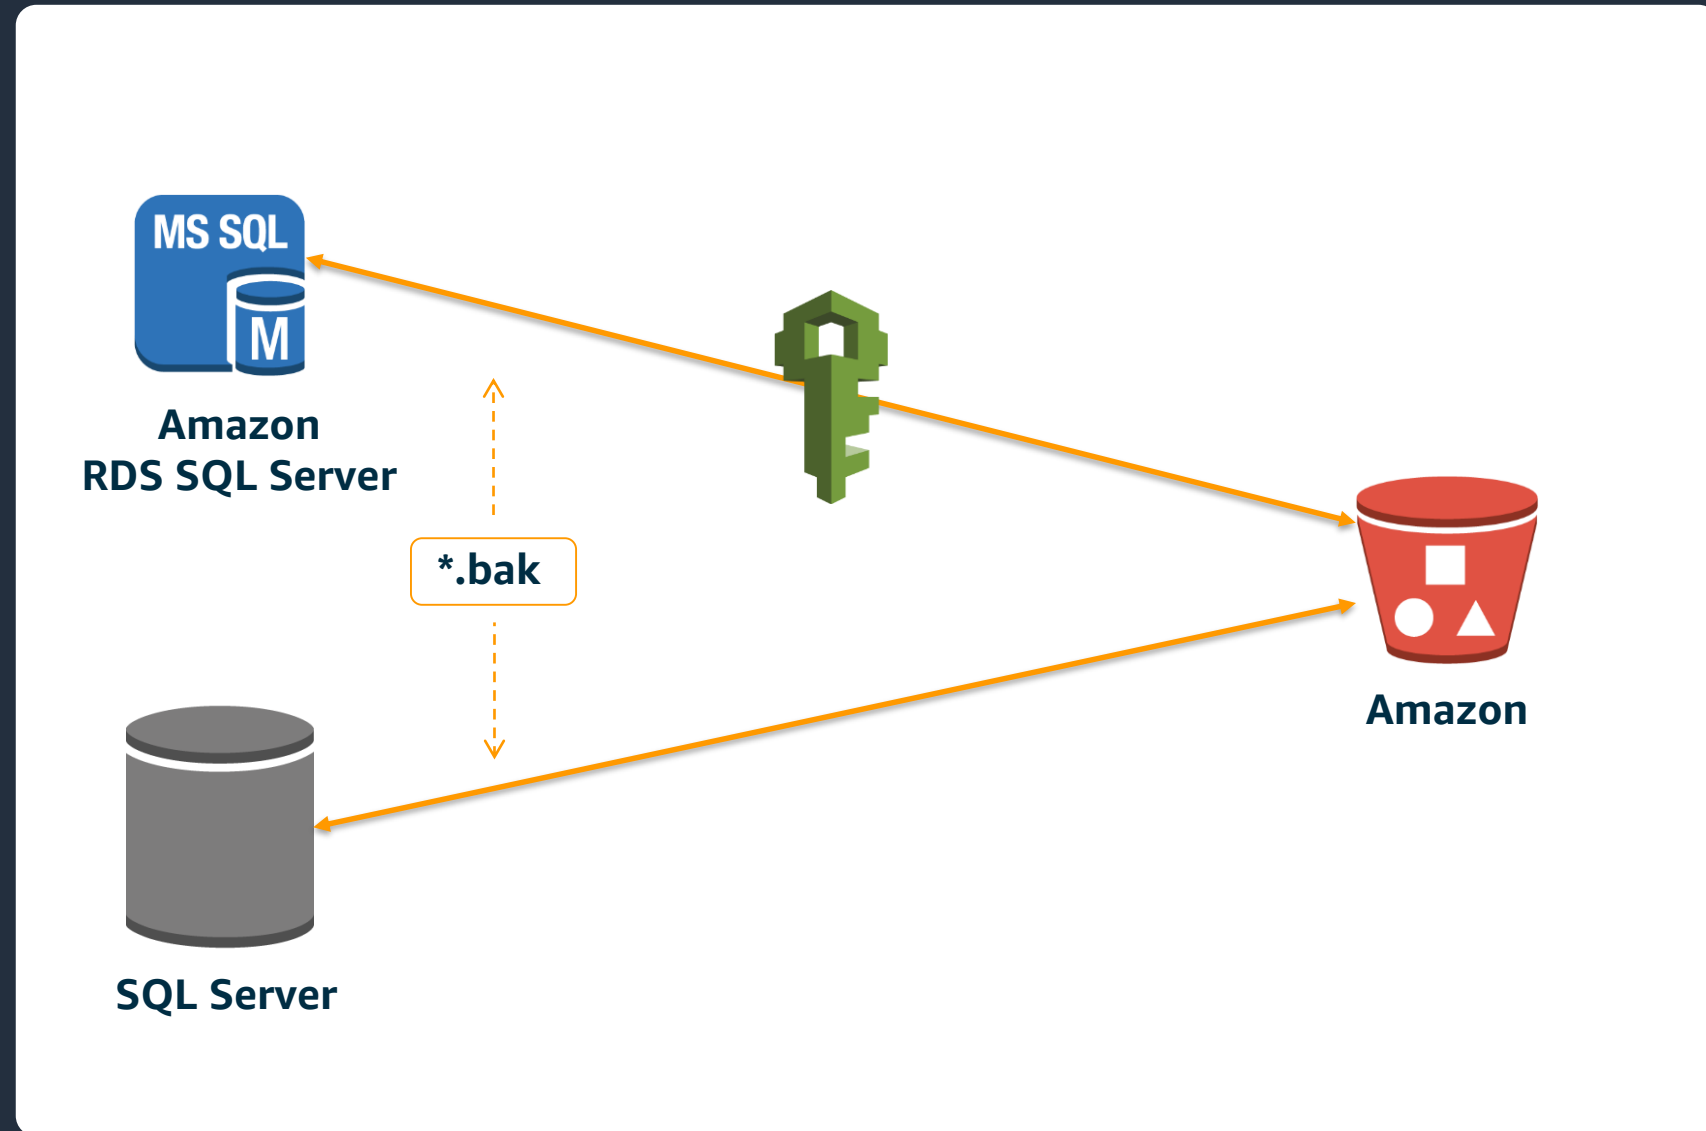
Usage

```
exec msdb.dbo.rds_backup_database
    @source_db_name='database_name',
    @s3_arn_to_backup_to='arn:aws:s3:::bucket_name/file_name.extension',
    [@kms_master_key_arn='arn:aws:kms:region:account-id:key/key-id'],
    [@overwrite_s3_backup_file=@1],
    [@type='DIFFERENTIAL/FULL'],
    [@number_of_files=n];
```

The following parameters are required:

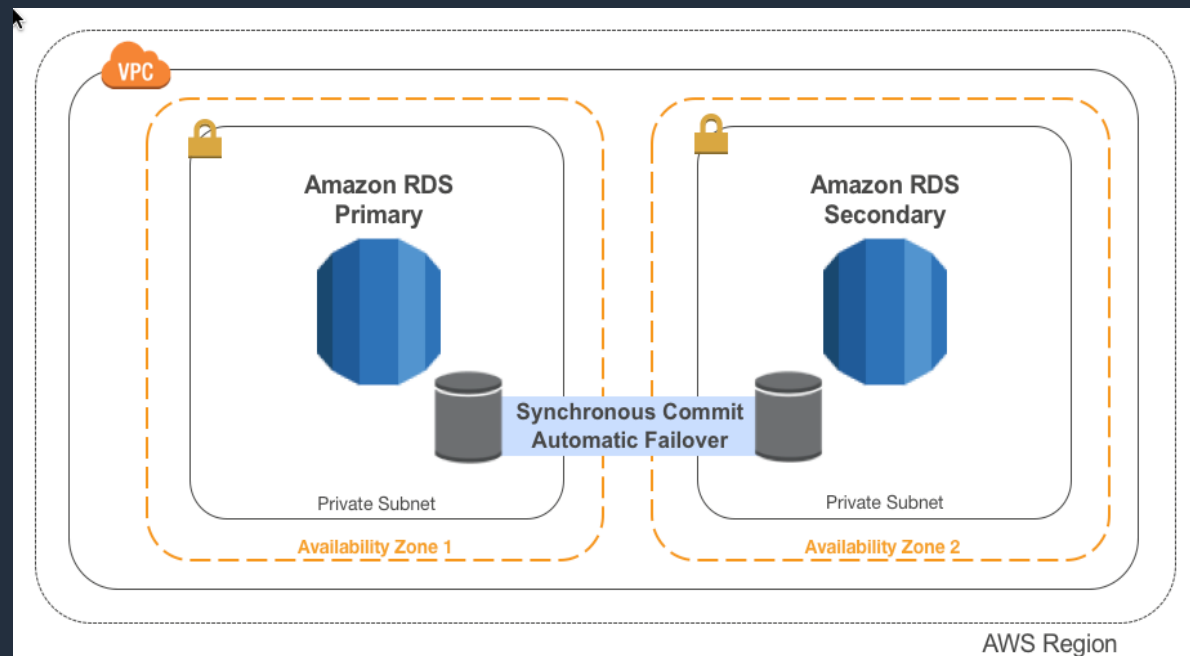
- `@source_db_name` – The name of the database to back up.
- `@s3_arn_to_backup_to` – The ARN indicating the Amazon S3 bucket to use for the backup, plus the name of the backup file.

The file can have any extension, but `.bak` is usually used.



Multi-AZ SQL Server on Amazon RDS

- AlwaysOn AG (Basic AG for Standard) for 2016, 2017 & 2019
- DB Mirroring for 2012 and 2014
- Synchronous Secondary hot StandBy
- Automatic & Manual Failover
- No read traffic



Availability: Failover Times

| Failover reason | Time to reconnect* (seconds) | | |
|-----------------------------------|------------------------------|-----------------------|--------------------|
| | Mirroring Primary DNS | Always On Primary DNS | Always On Listener |
| Manual failover | 60-70 | 30-40 | 5-7 |
| Primary SQL Server crash | 30-40 | 30-40 | 6-9 |
| Primary server network disconnect | 80-110 | 80-110 | 30-35 |

*Crash recovery times not included

**MultiSubnetFailover = True


Read scale-out

- 2016+ Enterprise Edition feature
- Up to 5 asynchronous read replicas
- In-region
- Promotion
- Separate end-point

The screenshot displays the AWS Management Console interface for Amazon RDS. The left sidebar shows navigation options like Dashboard, Databases, Query Editor, and Performance Insights. The main content area is titled 'Databases' and shows a table of database instances. The 'Actions' menu is open for the selected instance 'database-3', with 'Create read replica' highlighted. The table columns include DB identifier, Role, Engine, Region & AZ, and Size.

| DB identifier | Role | Engine | Region & AZ | Size |
|-----------------------|----------|-------------------------------|-------------|--------------|
| database-1 | Regional | Aurora PostgreSQL | us-west-2 | 2 instances |
| database-1-instance-1 | Writer | Aurora PostgreSQL | us-west-2a | db.r5.xlarge |
| reader-1 | Reader | Aurora PostgreSQL | us-west-2b | db.r5.xlarge |
| database-2 | Instance | PostgreSQL | us-west-2a | db.m5.xlarge |
| database-3 | Instance | SQL Server Enterprise Edition | us-west-2a | db.m5.xlarge |

Compute and Storage Scaling



**Scale Compute to
Handle Increased Load**
Up to 96 vCPUs (R5.24XL)
976 GiB of RAM (x1e.8XL)

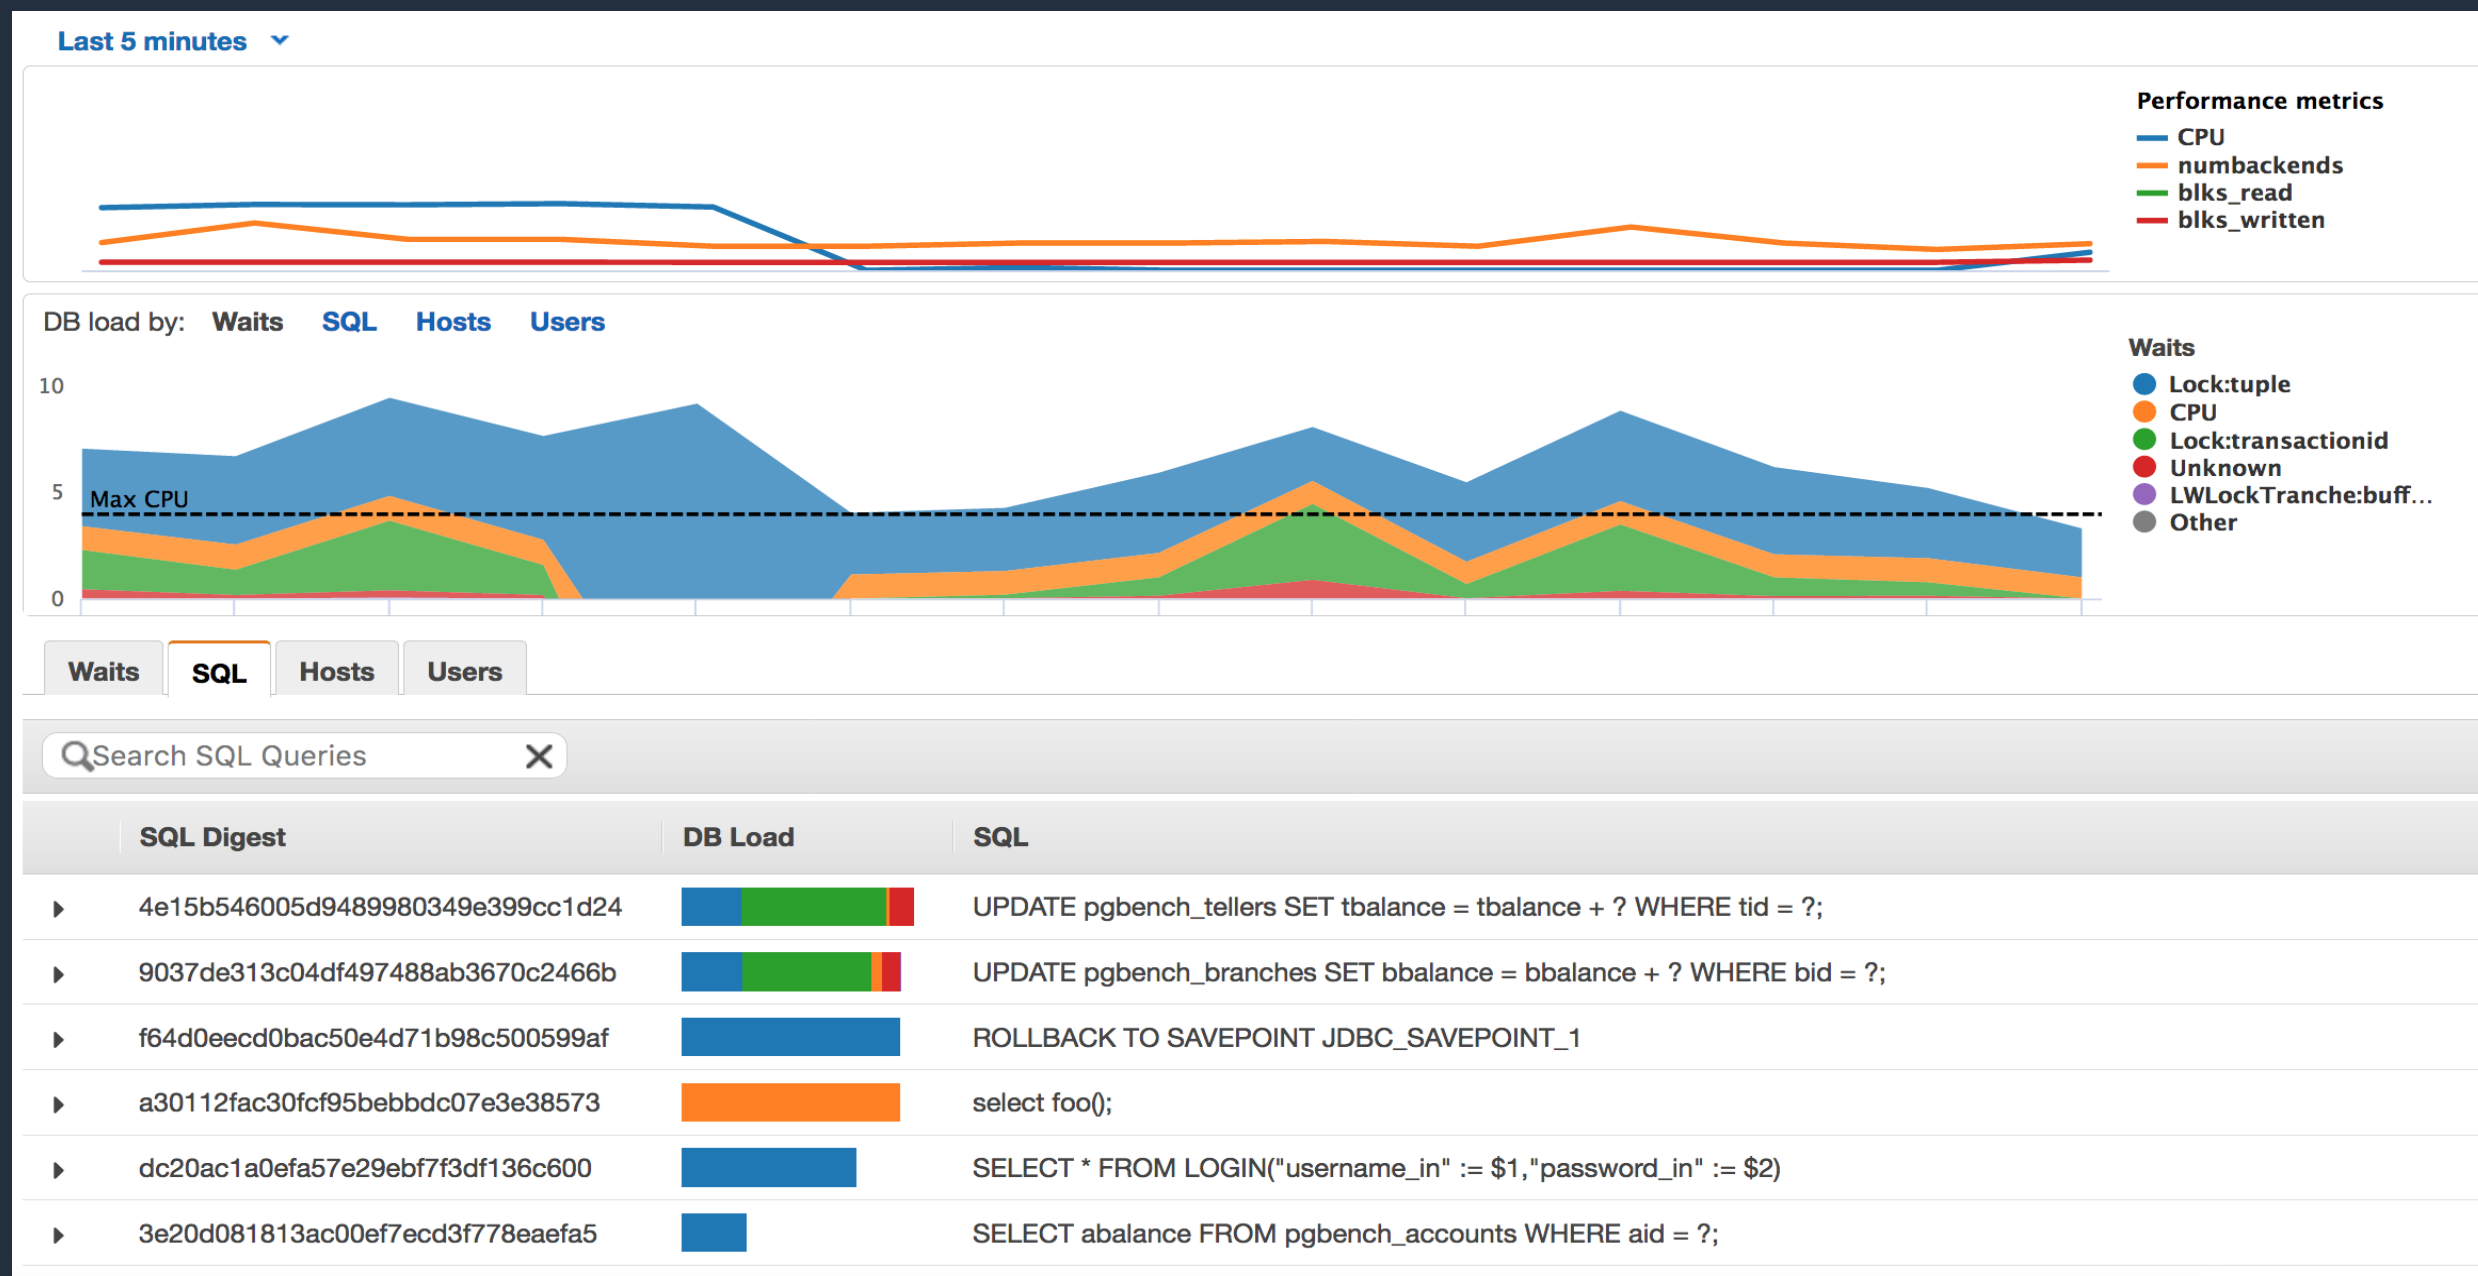


**Scale Storage for
Larger Data Sets**
Scalable EBS storage up to 16TiB



**Scale Down to
Control Costs**
As little as 1 vCPU and
1 GiB of RAM

Performance Insights for RDS SQL Server



Available Instance Types

R5 Family

- Memory Optimized
- R5.Xlarge (4 vCPU/32 GiB)
- R5.24Xlarge (96 vCPU/768 GiB)
- High performance networking

R5d Family (New!)

- Memory Optimized
- Local NVMe for TempDB
- R5.Xlarge (4 vCPU/32 GiB)
- R5.24Xlarge (96 vCPU/768 GiB)
- High performance networking

R5b Family (New!)

- Memory Optimized
- 3x EBS Performance
- R5.Xlarge (4 vCPU/32 GiB)
- R5.24Xlarge (96 vCPU/768 GiB)
- High performance networking

M5 Family

- General Purpose Instances
- M5.Xlarge (4 vCPU/16 GiB)
- M5.24Xlarge (96 vCPU/384 GiB)
- High performance networking

X1E

- Optimized for large-scale, enterprise-class and in-memory applications, and offer one of the lowest price per GiB of RAM among Amazon EC2 instance type
- x1e.Xlarge (4 vCPU/122 GiB)
- x1e.8Xlarge (32 vCPU/976 GiB)

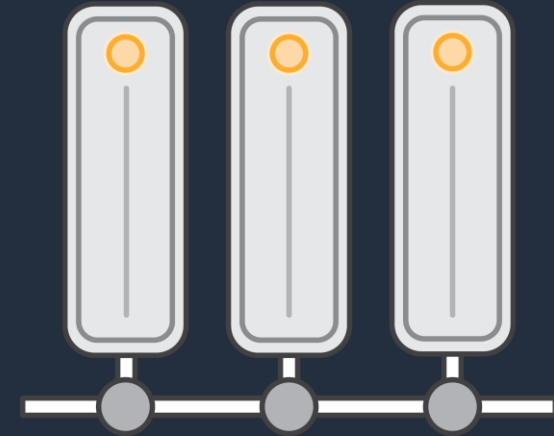
z1d Family

- CPU Optimized, 4GHz
- Memory Optimized
- z1d.Xlarge (4 vCPU/32 GiB)
- z1d.12Xlarge (48 vCPU/384 GiB)

Amazon Elastic Block Storage

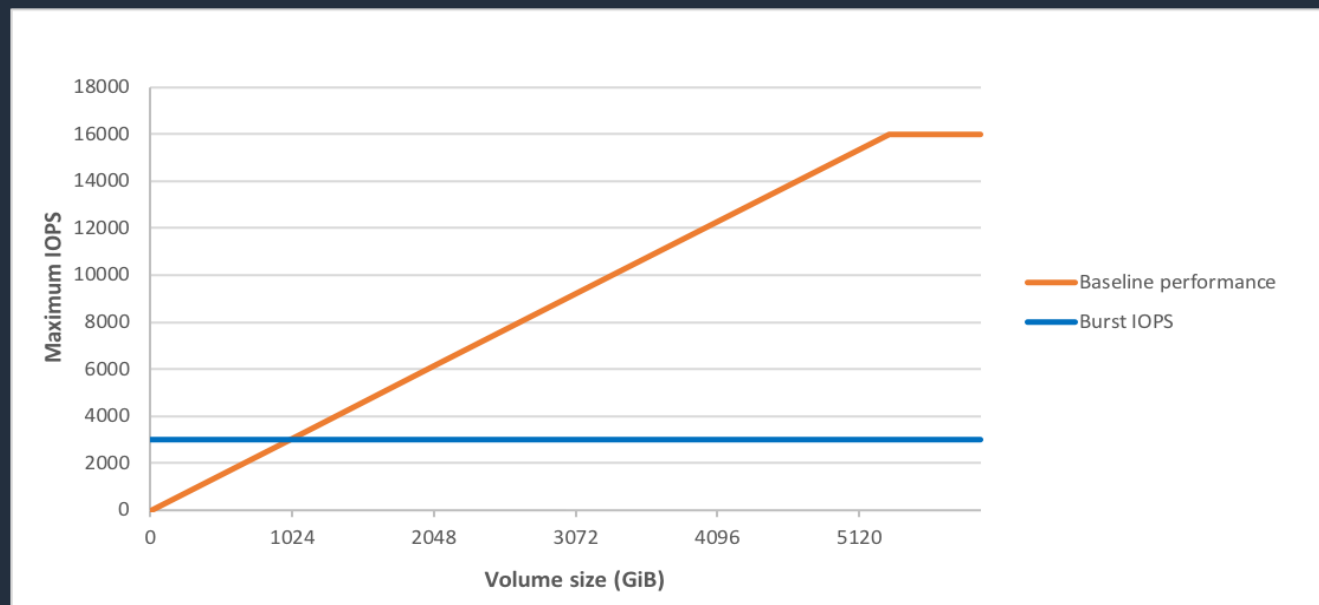
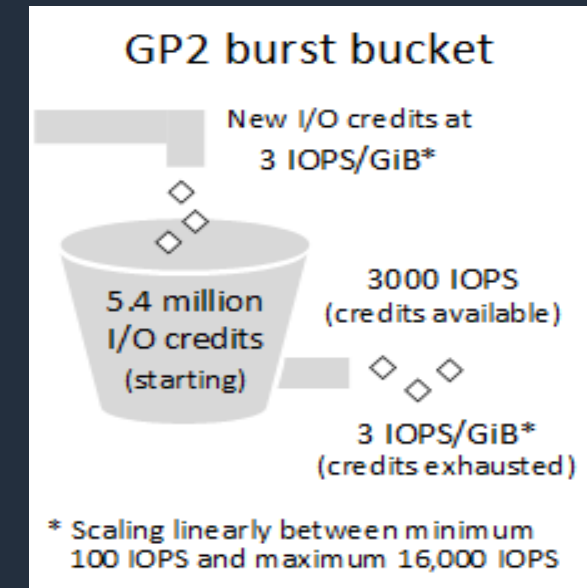
What is Amazon Elastic Block Storage (EBS)?

- Network-attached block storage
- Available for all instance types
- Many instance types support EBS optimization – dedicated channel for network storage I/O, eliminating contention with regular I/O
- Some instance types are EBS optimized, others offer it as an option



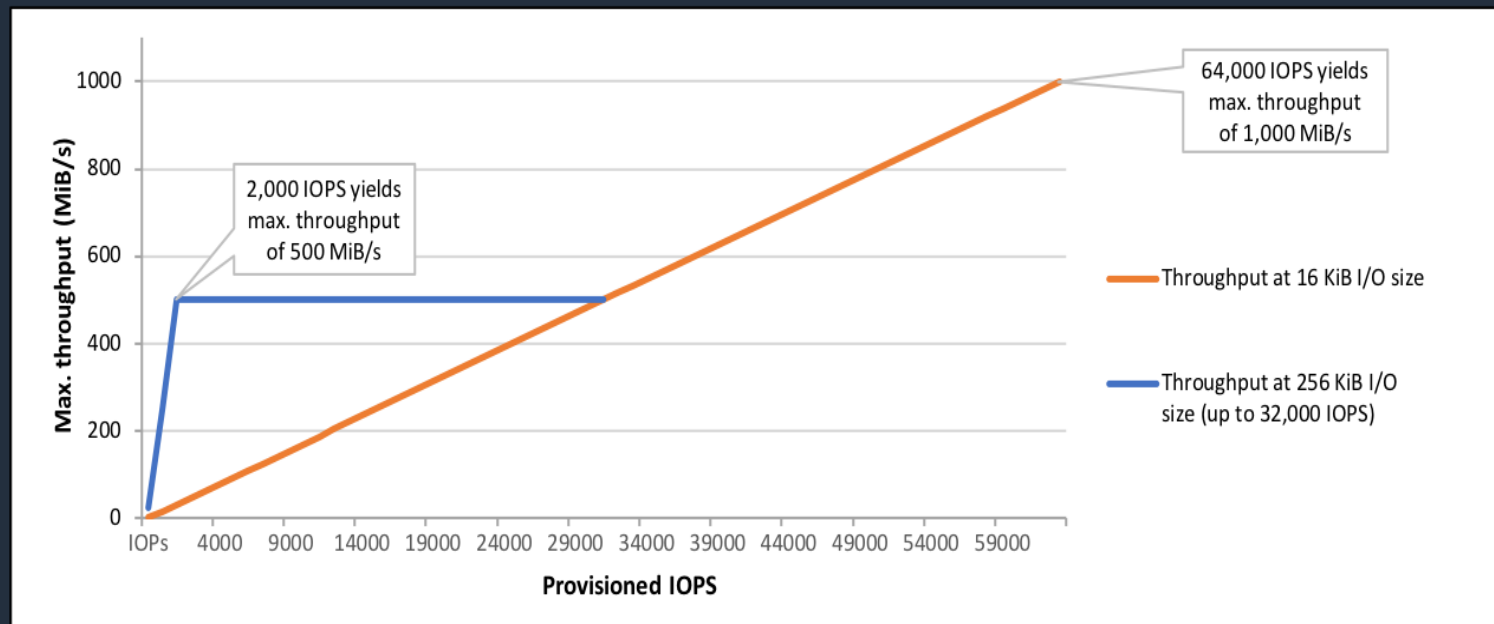
General Purpose (GP2)

- Cost-effective storage
- Wide range of workloads
- Single-digit ms latencies
- Ability to burst to 3,000 IOPS for extended periods of time
- Min of 100 IOPS (at 33.33 GiB and below)
- Max of 16,000 IOPS (at 5,334 GiB and above)
- Baseline performance scales linearly at 3 IOPS per GiB of volume size
- Provisioned performance 99% of the time
- Size from 1 GiB to 16 TiB

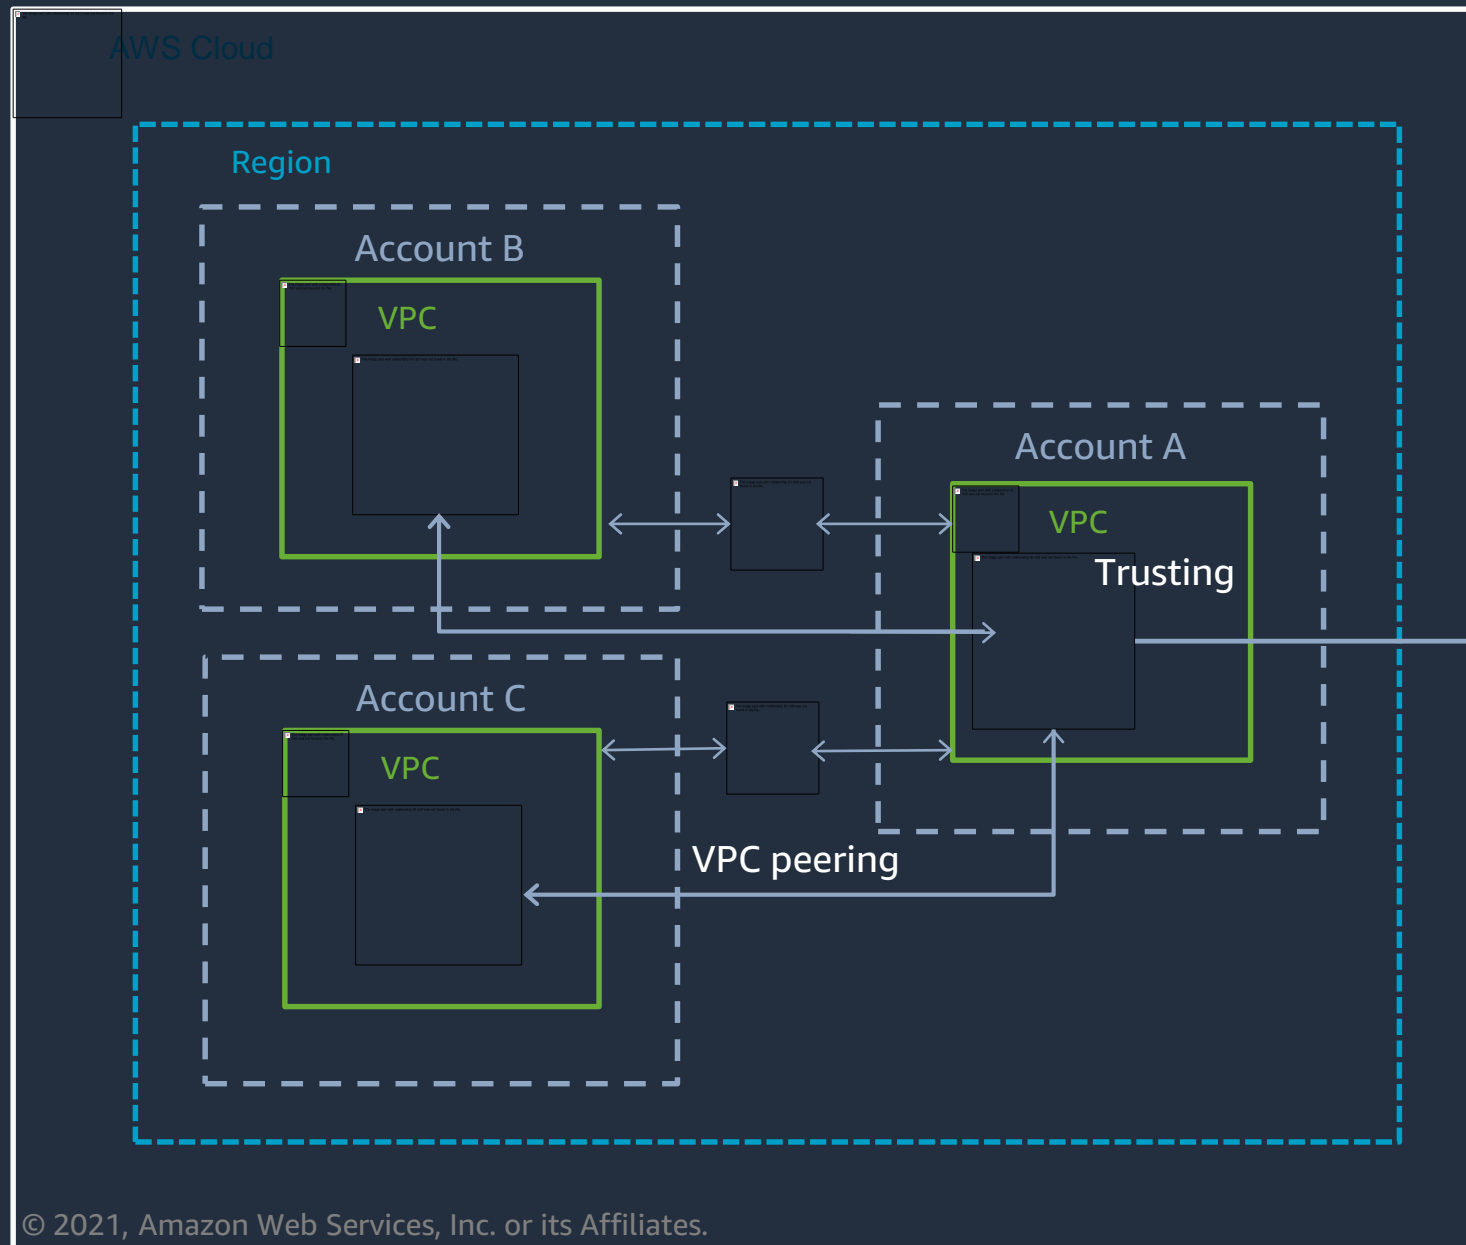


Provisioned IOPS (IO1)

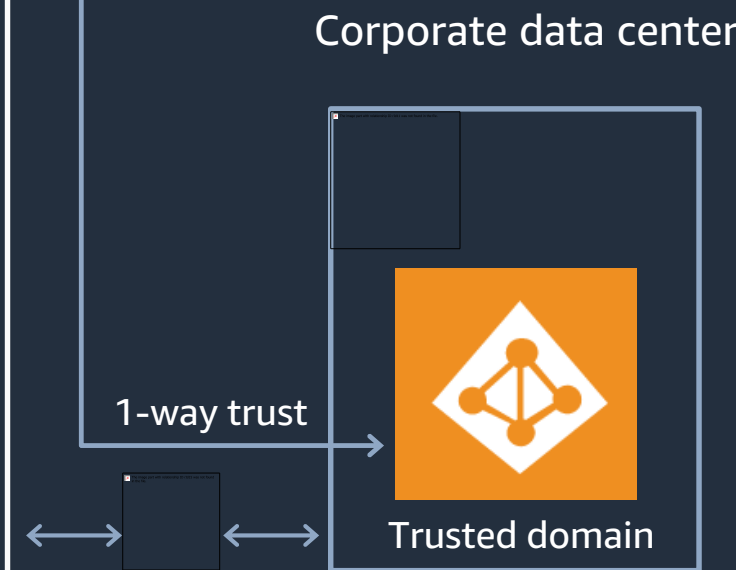
- I/O-intensive workloads
- Allows you to specify a consistent IOPS rate
- Delivers provisioned performance 99.9 percent of the time
- Size ranges from 4 GiB to 16 TiB
- IOPS range from 100 IOPS to 64,000 IOPS (Nitro instances)
- Up to 32,000 (prior to Nitro instances)
- The maximum ratio of provisioned IOPS to requested volume size (in GiB) is 50:1.



Windows Authentication Using On-Prem AD



- Setup SQL Server RDS
- Setup Managed AD
- Enable Windows Integrated Authentication to use Managed AD directory
- Create a Trust with On-Premises Domain
- Assign privileges to On-Premises for access to RDS



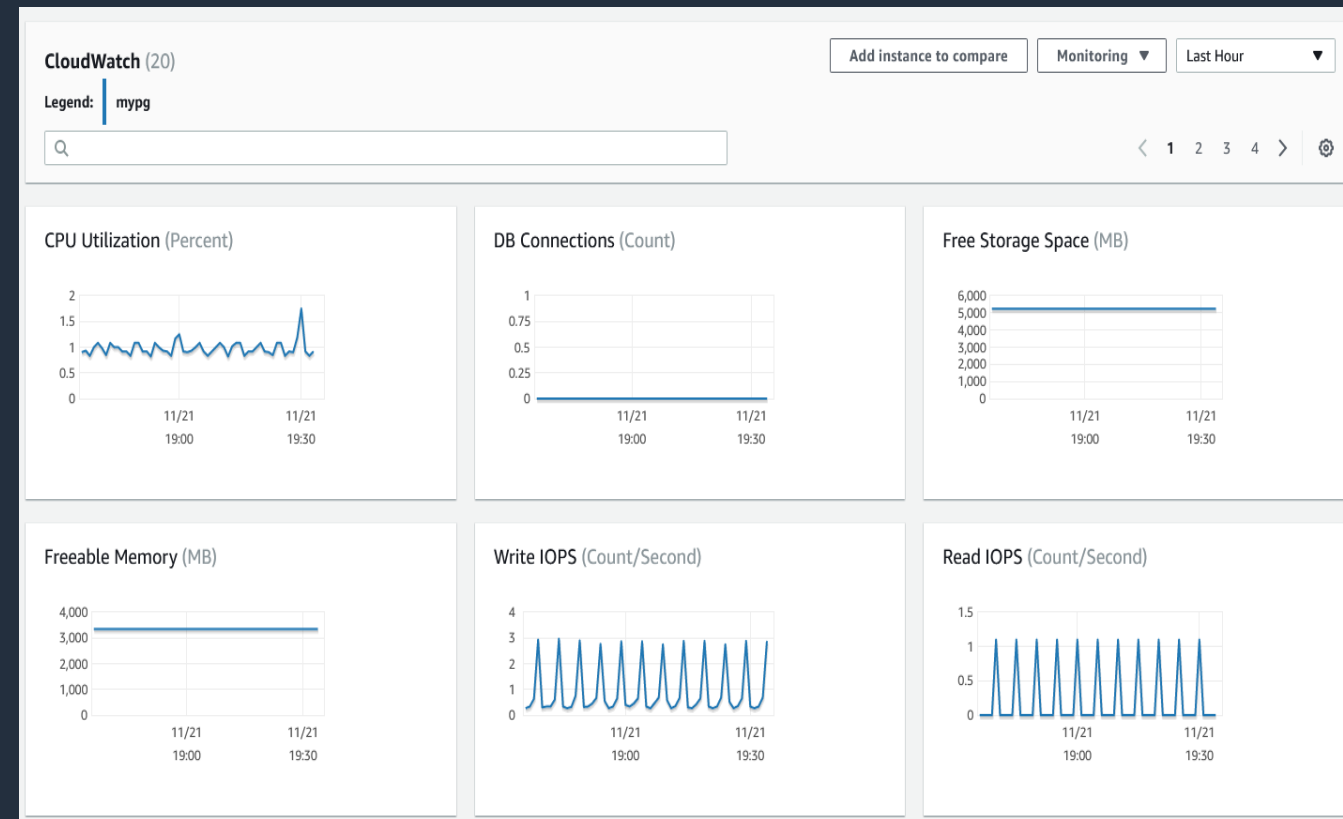
Monitoring RDS SQL Server performance

- 1 Amazon CloudWatch
- 2 Enhanced Monitoring
- 3 Performance Insights
- 4 SQL Server Native (DMVs/DMFs, Profiler, etc)
- 5 3rd Party (Ola Hallengren, SentryOne, etc)

Cloud Watch metrics

Amazon CloudWatch metrics

- CPU Utilization
- DB Connections
- Free Storage Space
- Freeable memory
- Write / Read IOPS
- Queue Depth
- Write / Read Throughput
- Swap usage
- Write / Read Latency
- Network Receive Throughput
- Network Transmit Throughput




Amazon RDS Enhanced Monitoring

Overview:

- OS Level Monitoring Metrics – 26 system and per process metrics
- Metrics delivered to CloudWatch Logs
- Up to 1 second granularity


Compared to CloudWatch Metrics:

- Agent based metrics collections
- There can be differences with CloudWatch metrics due to collection source (hypervisor vs. agent) – eg. CPU



Migrating Data to & from Amazon RDS

- 1 .BAK File Save & Restore**
Leverages SQL Server's native backup functionality
- AWS Database Migration Service**
Minimize downtime during migrations, migrate between different DB platforms, Schema Conversion Tool
- SQL Server Replication**
Push subscriptions to transactional replication
- Microsoft SQL Server Database Publishing Wizard, Import/Export**
Export to T-SQL files, load using `sqlcmd`
- AWS Marketplace**
Third-party data import and export tools and solutions



Amazon RDS hybrid use cases



Latency

Applications and processes sensitive to network and disk latency

Financial services applications such as trading and brokerage

Security and fraud applications that require a quick response time



Residency

Regulations dictate that data and infrastructure reside locally

Contracts specify where applications are deployed

Information security or other reasons prevent adoption of AWS regions

Amazon RDS on VMware

Overview

- Database managed services for the hybrid cloud (AWS regions and on-premises datacenters)
- Available on VMware vSphere

Benefits

- Unified Amazon RDS interface to manage databases on-premises and AWS regions
- Comprehensive database management services for a hybrid cloud environment

Feature

RDS on VMware

DB engines

- SQL Server 2016 SP2 Enterprise edition
- MySQL 5.7, PostgreSQL 10.9

Performance and scalability

- Instances from M to 24XL
- Read replica (MySQL, PostgreSQL)

Availability and durability

- Event integration
- Amazon Cloudwatch events and metrics
- On-premises storage for data residency

Infrastructure

- VMware vSphere 6.5+
- VMFS, NFS, or vSAN

New features of RDS SQL Server

1. Enterprise Edition Multi-AZ Price Reduction!

<https://aws.amazon.com/about-aws/whats-new/2020/07/amazon-rds-for-sql-server-lowers-cost-for-high-availability-db-instances/>

2. Cross-Account & cross-VPC Domain Joins

<https://aws.amazon.com/blogs/database/joining-your-amazon-rds-instances-across-accounts-to-a-single-shared-domain/>

3. Disable older versions of TLS and Ciphers

<https://aws.amazon.com/blogs/database/customizing-security-parameters-on-amazon-rds-for-sql-server/>

4. SQL Server Integration Services

<https://www.youtube.com/watch?v=4-6Jo7RDLoc>

5. SQL Server Reporting Services

<https://www.youtube.com/watch?v=2JyiHnjdRiQ>

New features of RDS SQL Server (Continued)

6. SQL Server Analysis Services (tabular)

https://www.youtube.com/watch?v=G2kaz_G7vNA

7. In-Region Read Replicas

<https://www.youtube.com/watch?v=2NnS55pto4I>

8. MSDTC

<https://aws.amazon.com/blogs/database/enabling-distributed-transaction-support-for-domain-joined-amazon-rds-for-sql-server-instances/>

9. Replicating the Service Master Key

<https://aws.amazon.com/about-aws/whats-new/2020/07/amazon-rds-for-sql-server-now-supports-service-master-key-retention/>

10. Multi-file Backups


<https://aws.amazon.com/about-aws/whats-new/2019/06/amazon-rds-for-sql-server-now-supports-multi-file-native-restores/>

How to deploy and manage RDS SQL Server

Multiple ways to start and manage your SQL Server resources using AWS



AWS Console



AWS CLI



AWS SDKs



CloudFormation

Questions **Answers**

