Under the hood: How to run your Oracle databases on AWS

Nathan Fuzi, Sr. Database Solutions Architect

Agenda

- Running Oracle workloads on AWS
- Scaling RDS for Oracle performance
- High Availability and Disaster Recovery
- Migrating to RDS for Oracle
- RDS for Oracle best practices

aws

Meet Your Presenter

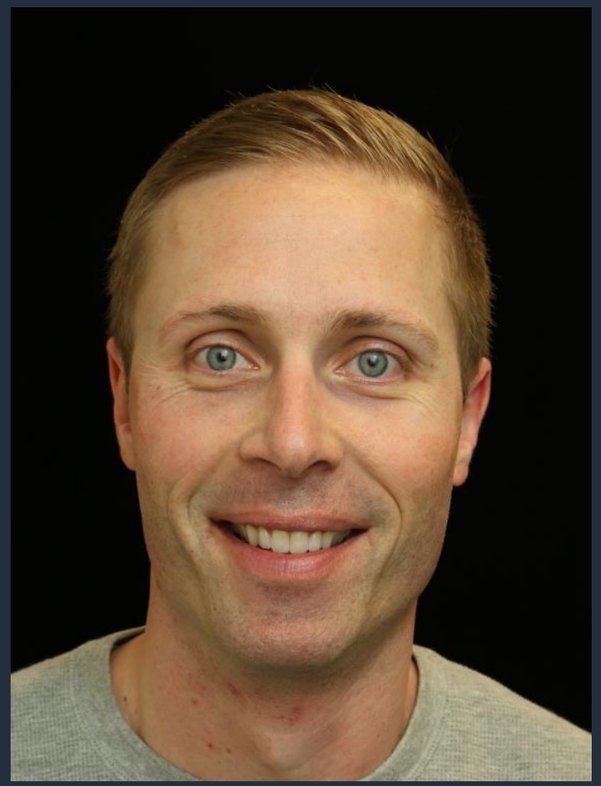

Nathan Fuzi
Sr. Oracle Database Specialist SA

- 4 years with AWS
- 20+ years working with Oracle Database and related technologies

aws

Running Oracle workloads on AWS

Relational databases are complex

Our experience running Amazon.com taught us that relational databases are challenging to manage and operate with high availability.

It's expensive and complex to manage administrative functions including regular patching cycles, performance optimization, and backup and disaster recovery—all for constantly changing applications.

aws

Amazon Relational Database Service

Managed relational database service with a choice of six popular database engines

Amazon Aurora | MySQL | PostgreSQL | MariaDB | MicroSoft SQLServer | Oracle

Easy to administer

Easily deploy and maintain hardware, OS and DB software; built-in monitoring

Secure & compliant

Data encryption at rest and in transit; industry compliance and assurance programs

Available & durable

Automatic Multi-AZ data replication; automated backup, snapshots, failover

Performant & scalable

Scale compute and storage with a few clicks; minimal downtime for your application

aws

Amazon Relational Database Service

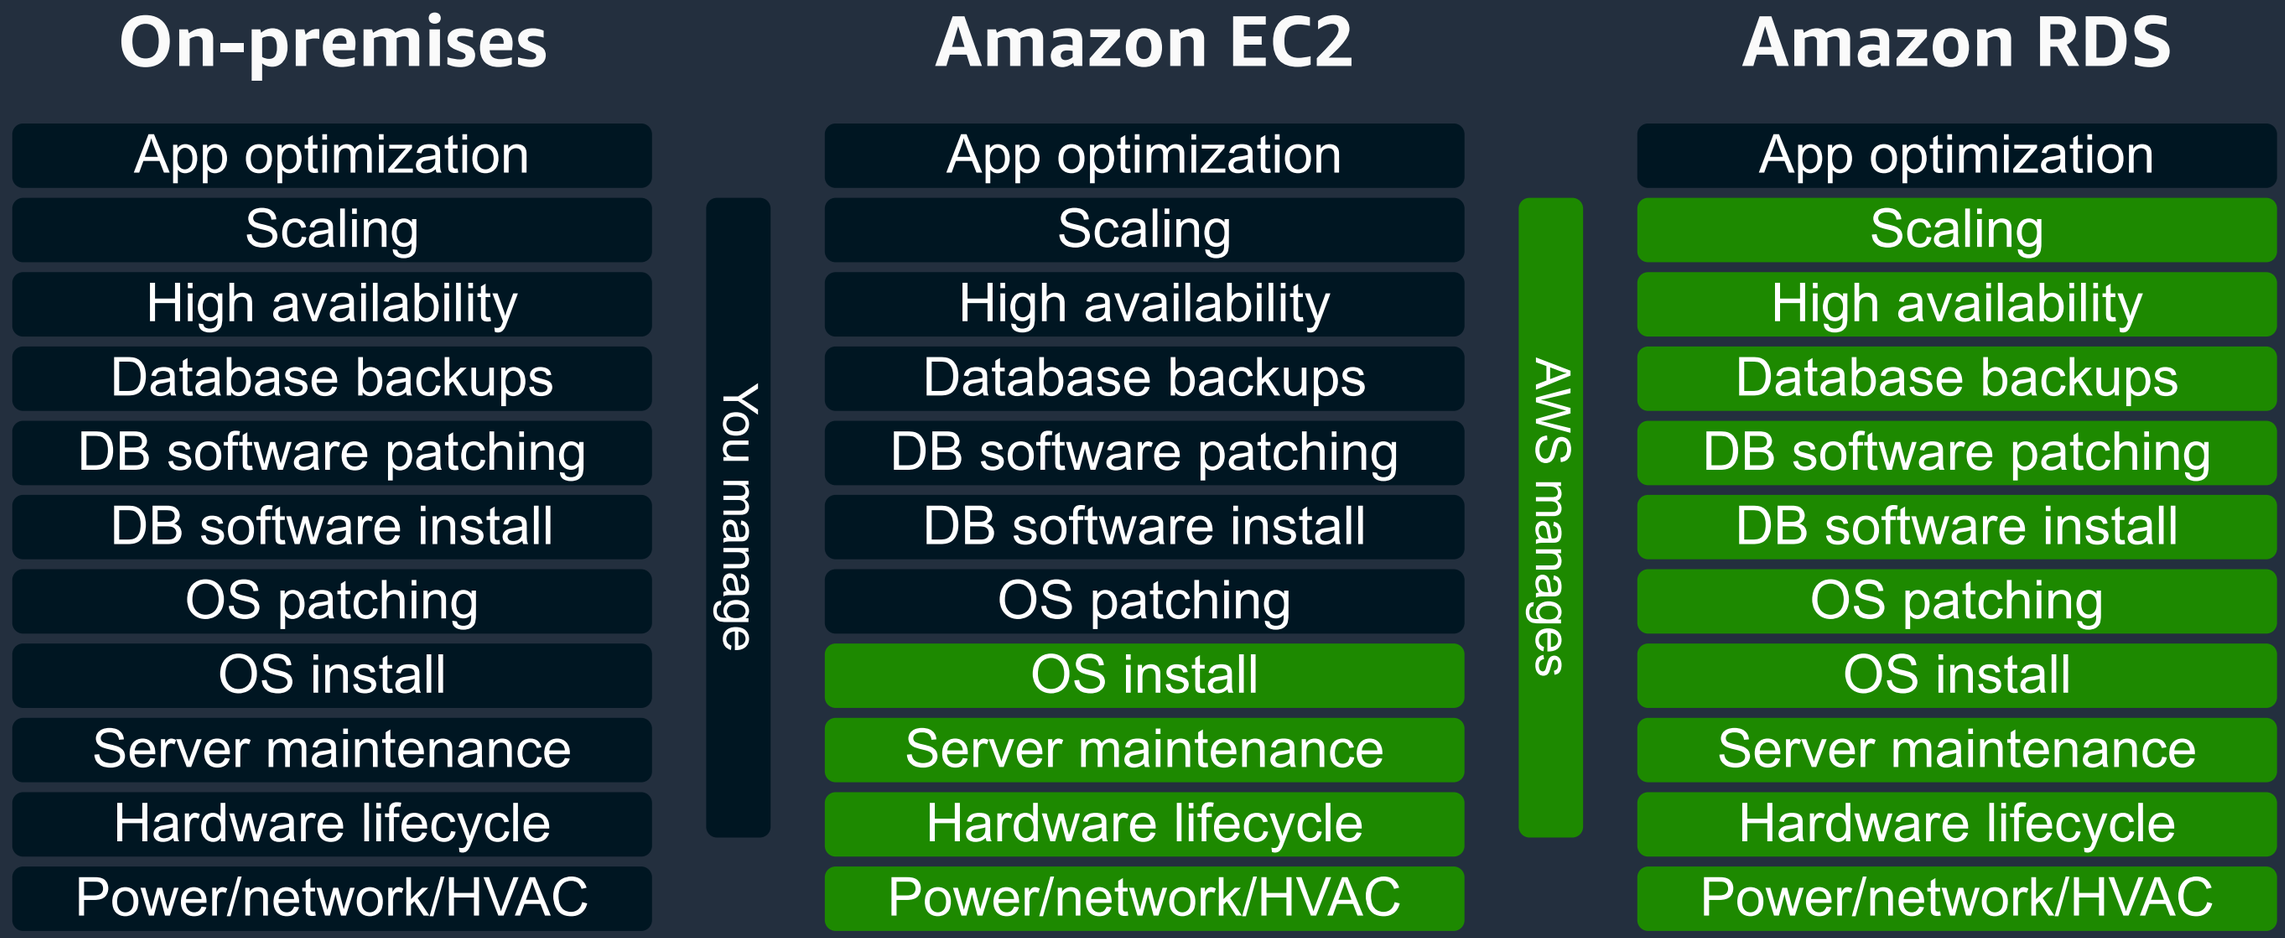

aws

When to choose to run Oracle on Amazon EC2

- You run a configuration outside what RDS can support
  - Storage configuration
  - DB size or performance
  - Third party backup or replication solution

- You need a feature that RDS does not support
  - Operating System access
  - Specific database features

- You run a specific version outside what is available in RDS
  - Specific version or patches
  - Avoiding software patches

aws

Choosing the best RDS solution for your needs

Amazon RDS for Oracle License Included (LI)

- Oracle SE1 and SE2
- No need to purchase Oracle licenses from Oracle
- Support is handled through AWS

Amazon RDS for Oracle Bring Your Own License (BYOL)

- SE1, SE2 and EE
- Bring your own licenses and support from Oracle
- Call Oracle for Oracle database support

aws

Scaling RDS for Oracle Performance

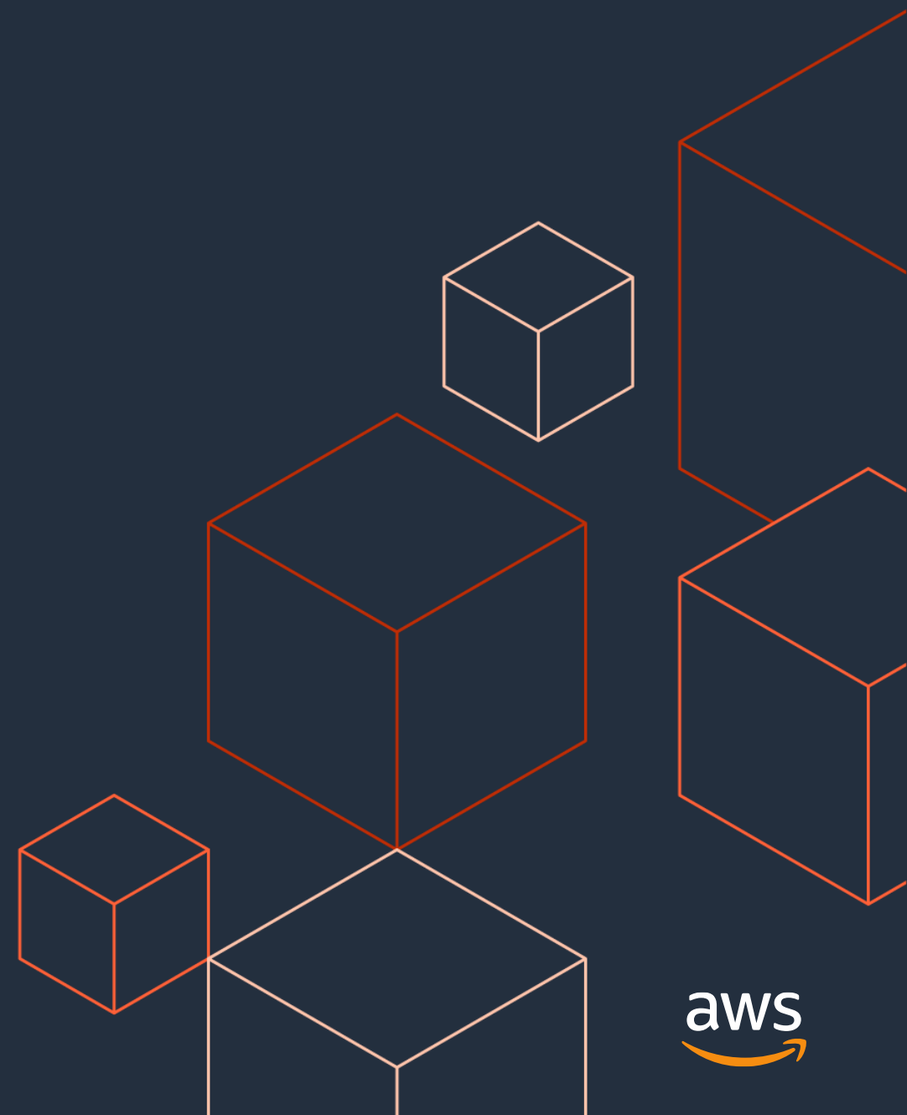

aws

RDS performance factors

RDS DB Instance Class

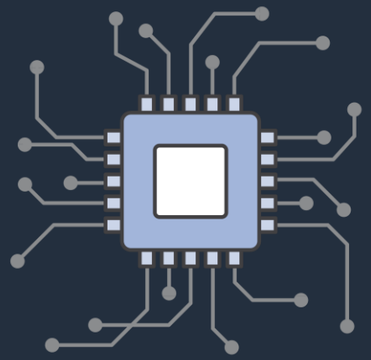

Compute Capabilities

vCPUs

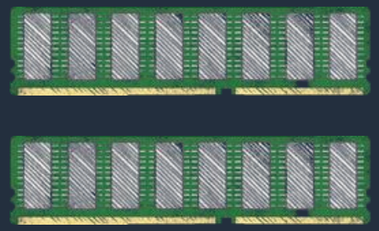

Memory Capabilities

GB of RAM

Network Performance

MB/s (Throughput)

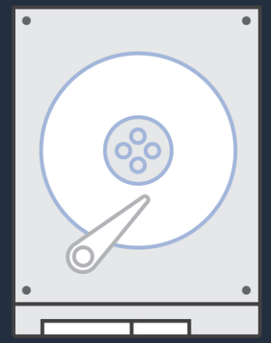

Storage Performance

I/O Performance

RDS Storage Type

aws

Scaling compute and storage

- ✓ Nearly 100 RDS instance class and size options
- ✓ Processing power: 1 vCPU up to 128 vCPUs
- ✓ Memory size: 1 GiB up to 3,904 GiB
- ✓ Storage from 20 GB to 64 TB and baseline IOPS from 60 to 80,000
- ✓ Dedicated storage bandwidth up to 19,000 Mbps
- ✓ Network performance up to 25 Gbps

aws

Monitoring Amazon RDS for Oracle databases

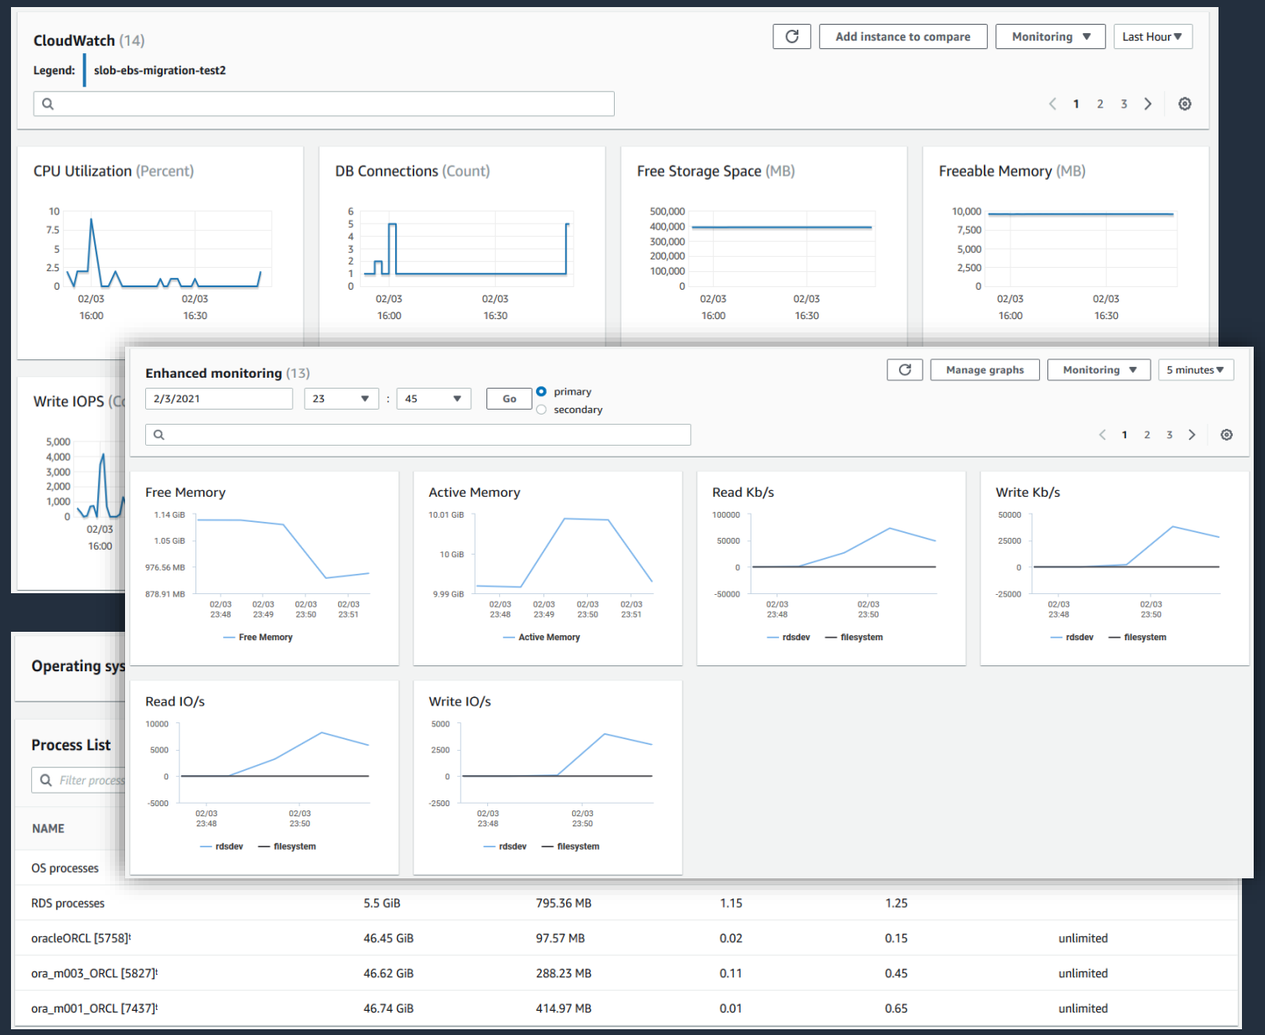

Amazon CloudWatch metrics

- CPU / Storage / Memory
- Swap usage
- I/O (read and write)
- Latency (read and write)
- Throughput (read and write)
- Replica lag

Amazon CloudWatch Alarms

- Similar to on-premises monitoring tools

Enhanced Monitoring for Amazon RDS

- Access to over 50 CPU, memory, file system, and disk I/O metrics
- As low as 1 second intervals

Integration with 3rd Party Monitoring Tools

aws

Performance Insights for real-time analysis

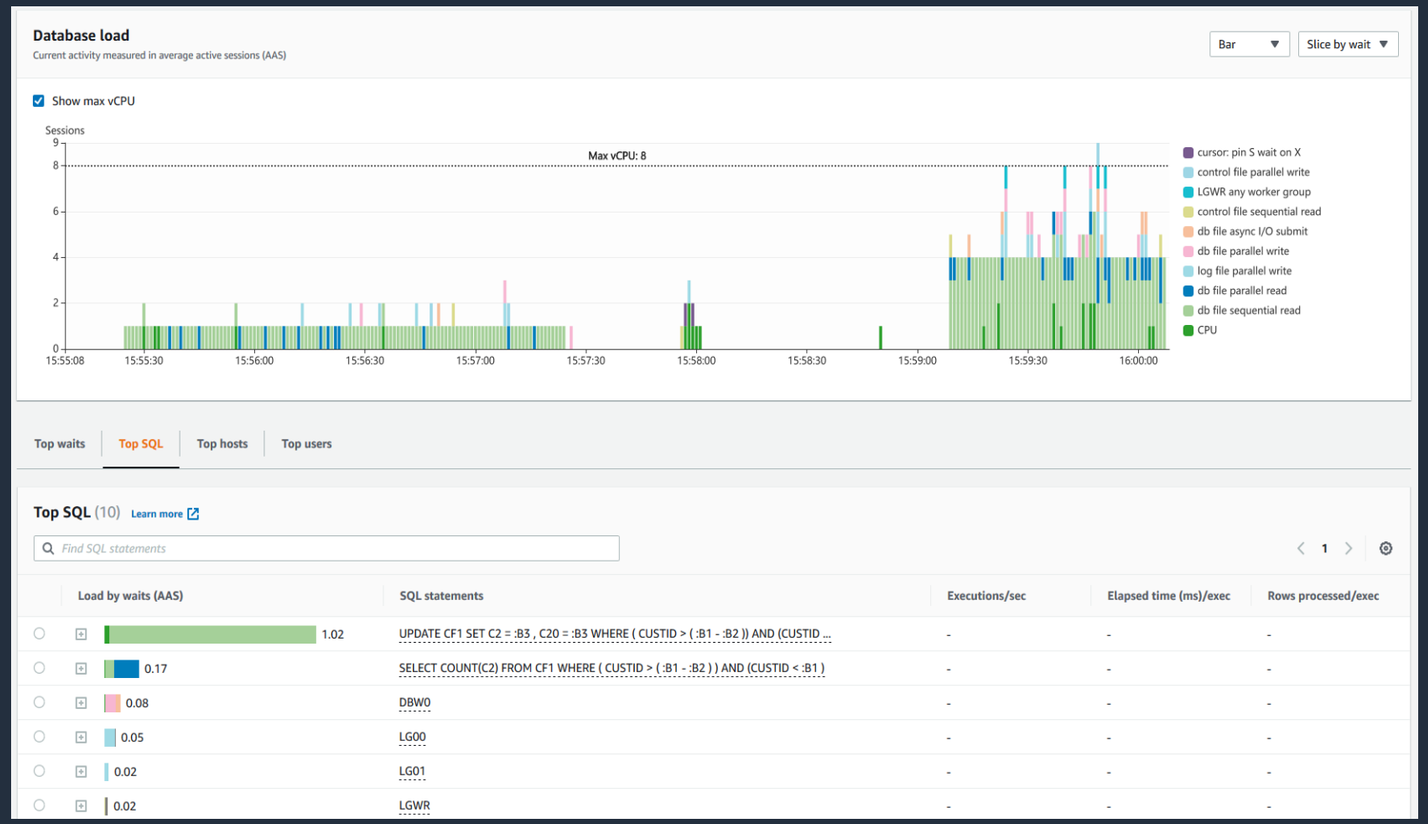

Amazon RDS Performance Insights measures database load over time

Easy to identify database bottlenecks

- Top SQL/most intensive queries

Enables problem discovery for real time or historical analysis

Adjustable timeframe

- Hour, day, week, and longer

Available for all Amazon RDS database engines

Scaling for read-heavy workloads

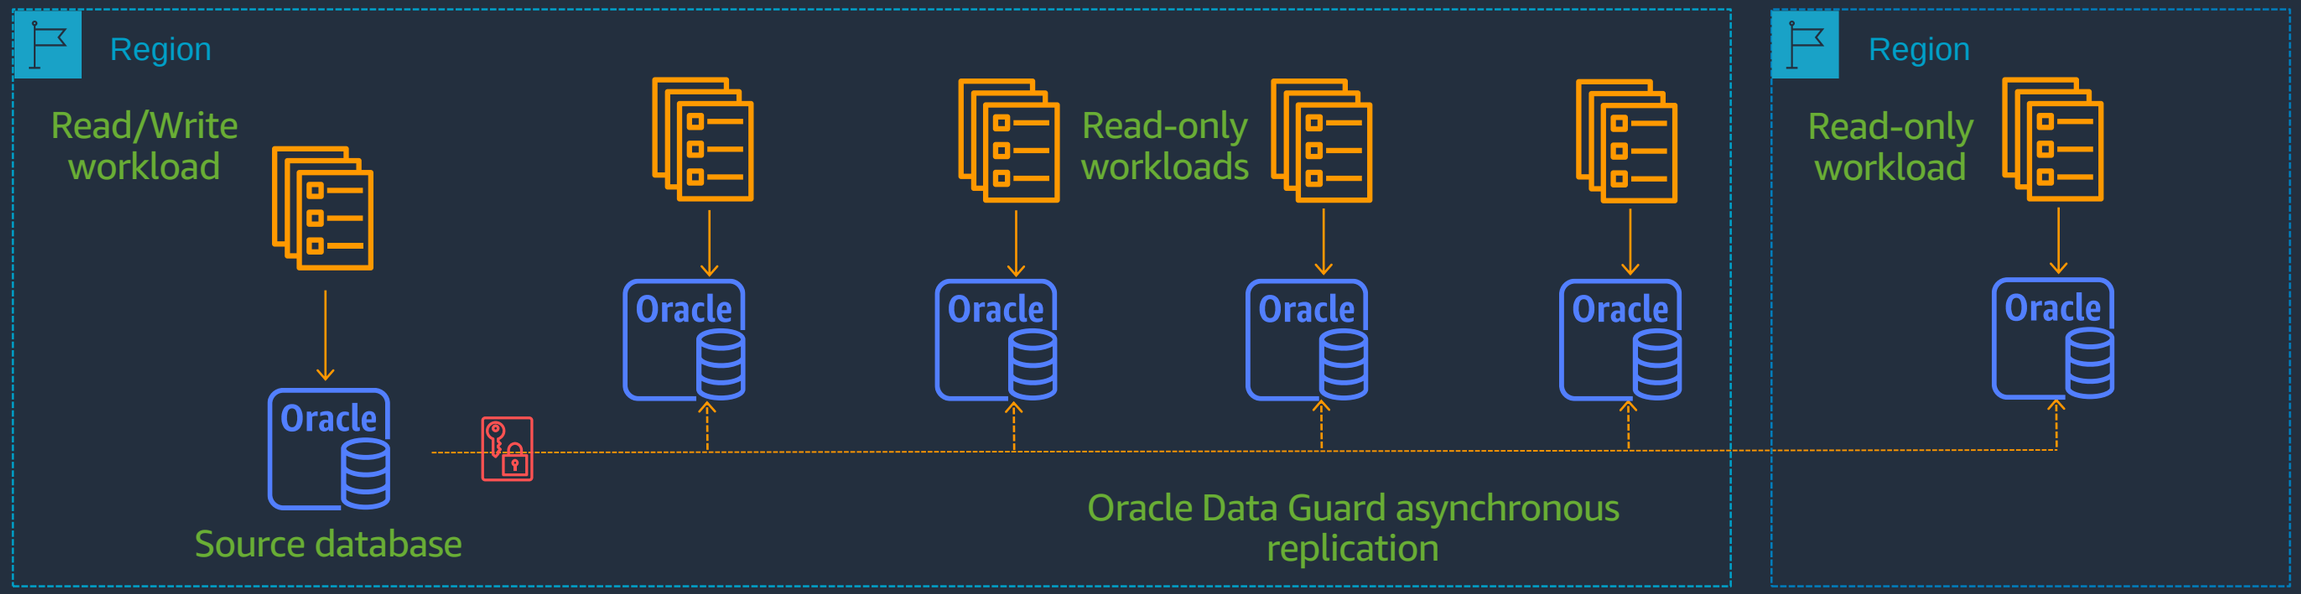

- ✓ Up to five replicas per source RDS instance, in same region or cross-region
- ✓ Scale out or back, and size each replica instance independently
- ✓ Replicas may be mounted or available for reads (Active Data Guard)
- ✓ Promote replica as new standalone database
- ✓ Source has no downtime during creation or promotion

aws

High Availability and Disaster Recovery

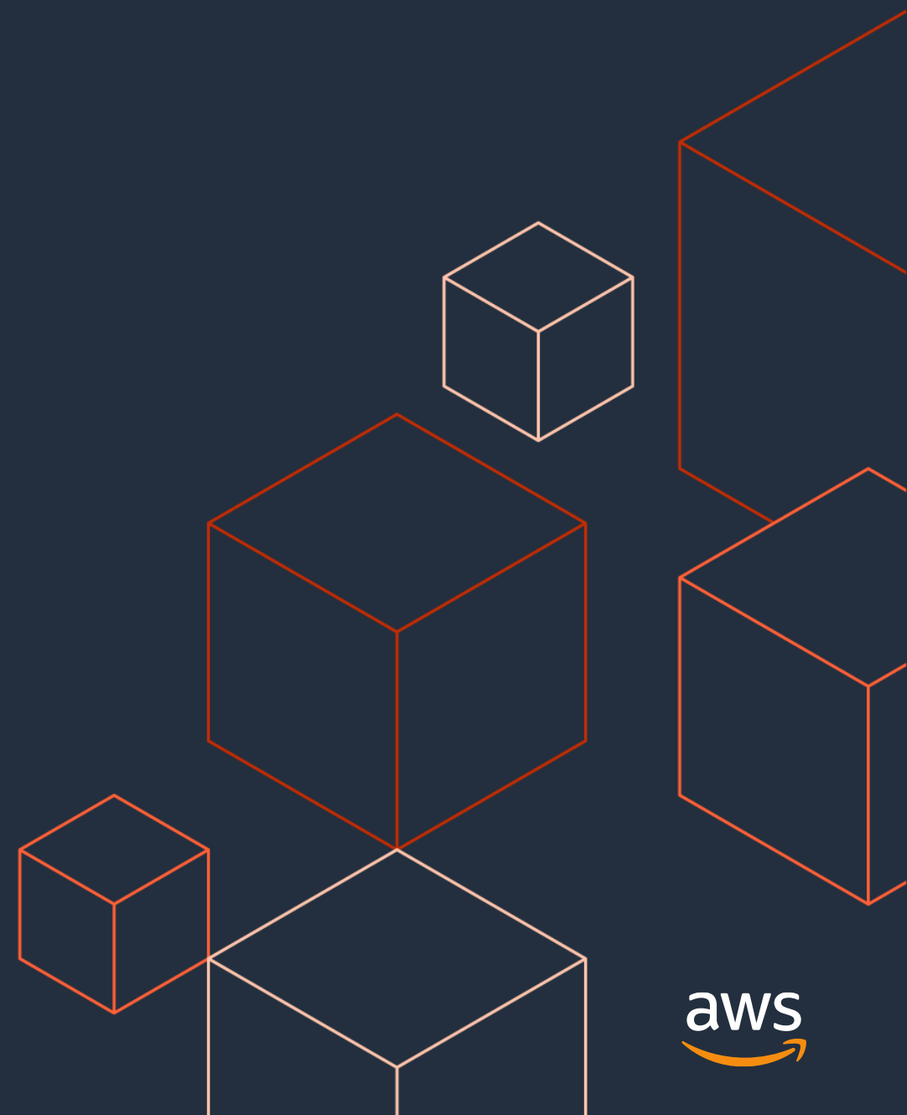

Availability and recoverability

Recovery Point Objective (RPO) and Recovery Time Objective (RTO):

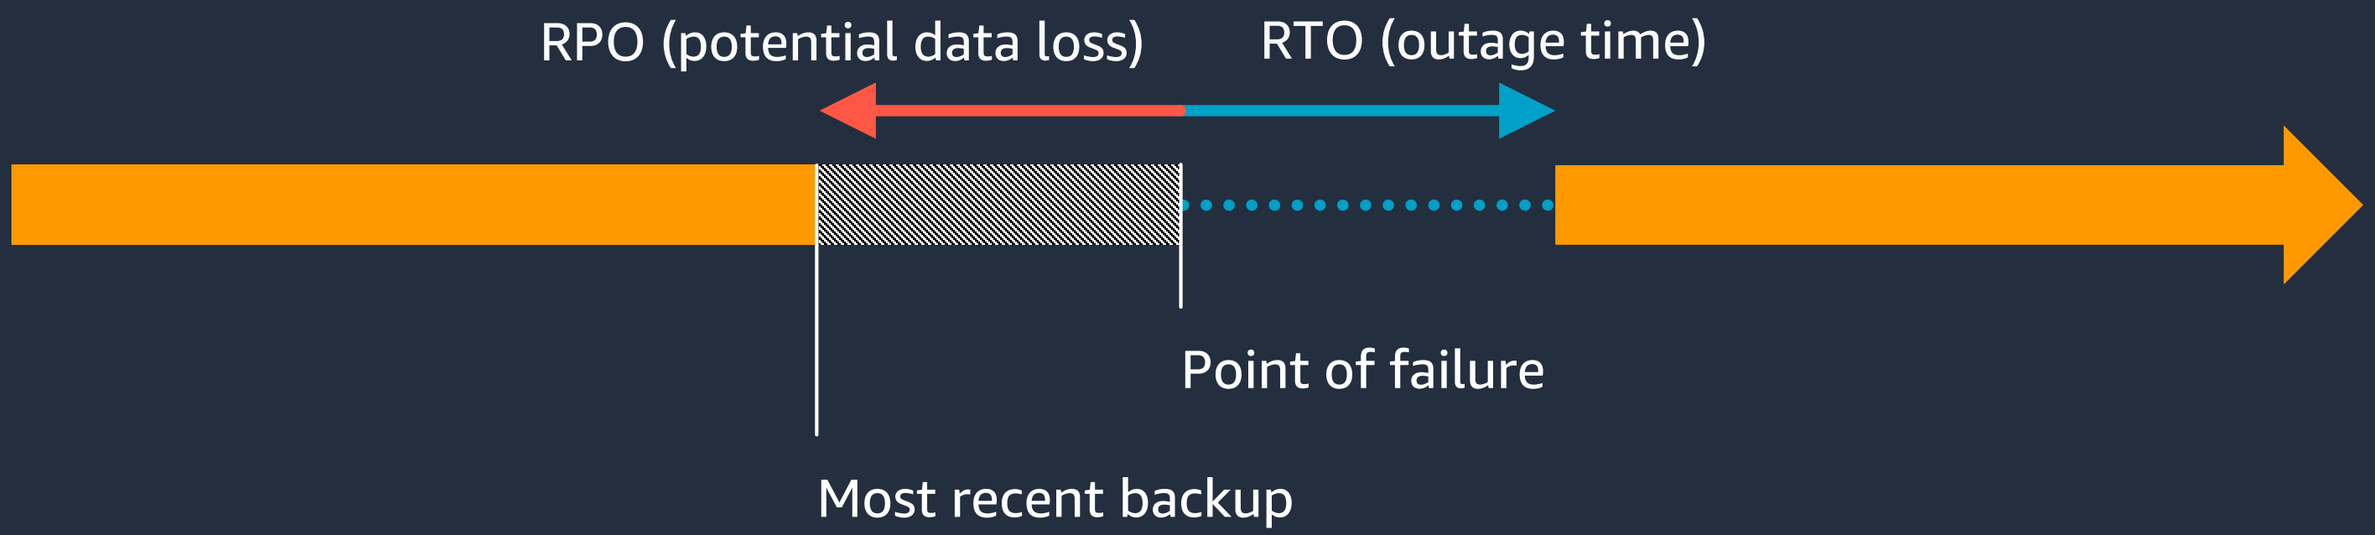

aws

RDS HA / DR comparison

Feature	RPO (approx)	RTO (approx)	Licensing
Multi-AZ for high availability	0	1 to 2 minutes	All editions
Snapshot restore	Hours	< 1 hour	
Point-in-time restore (cross-region)	25 minutes	Hours	
Point-in-time restore (in-region)	5 minutes	Hours	
Mounted replica promotion (in-region)	Minutes	Minutes	Enterprise edition
Mounted replica promotion (cross-region)	Minutes	Minutes	
Read replica promotion (in-region)	Minutes	Minutes	EE + Active DataGuard
Read replica promotion (cross-region)	Minutes	Minutes	

aws

RDS Multi-AZ architecture

Multi-AZ Key Aspects

- Compute instances manage synchronous storage replication
- Writes occur on both instances and storage; reads only from Primary
- Failover results in changing roles (Primary, Standby) and re-establishing Standby
- Snapshots taken from Standby

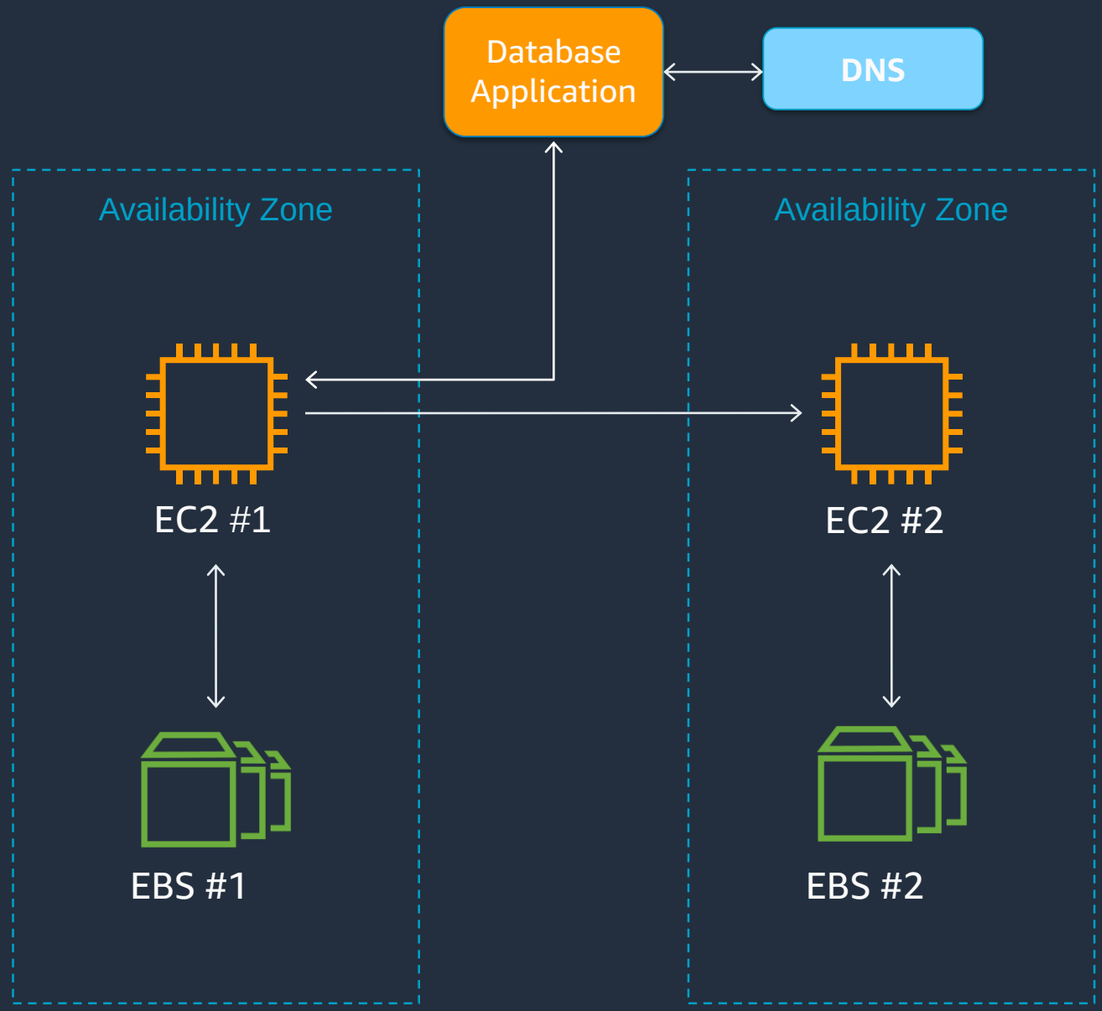

aws

RDS for Oracle - Backups

Automated backups

- Daily snapshot during backup window
- Redo logs to S3 every 5 minutes
- Retained 1-35 days

Manual backups

- Take a snapshot any time
- Kept until you delete

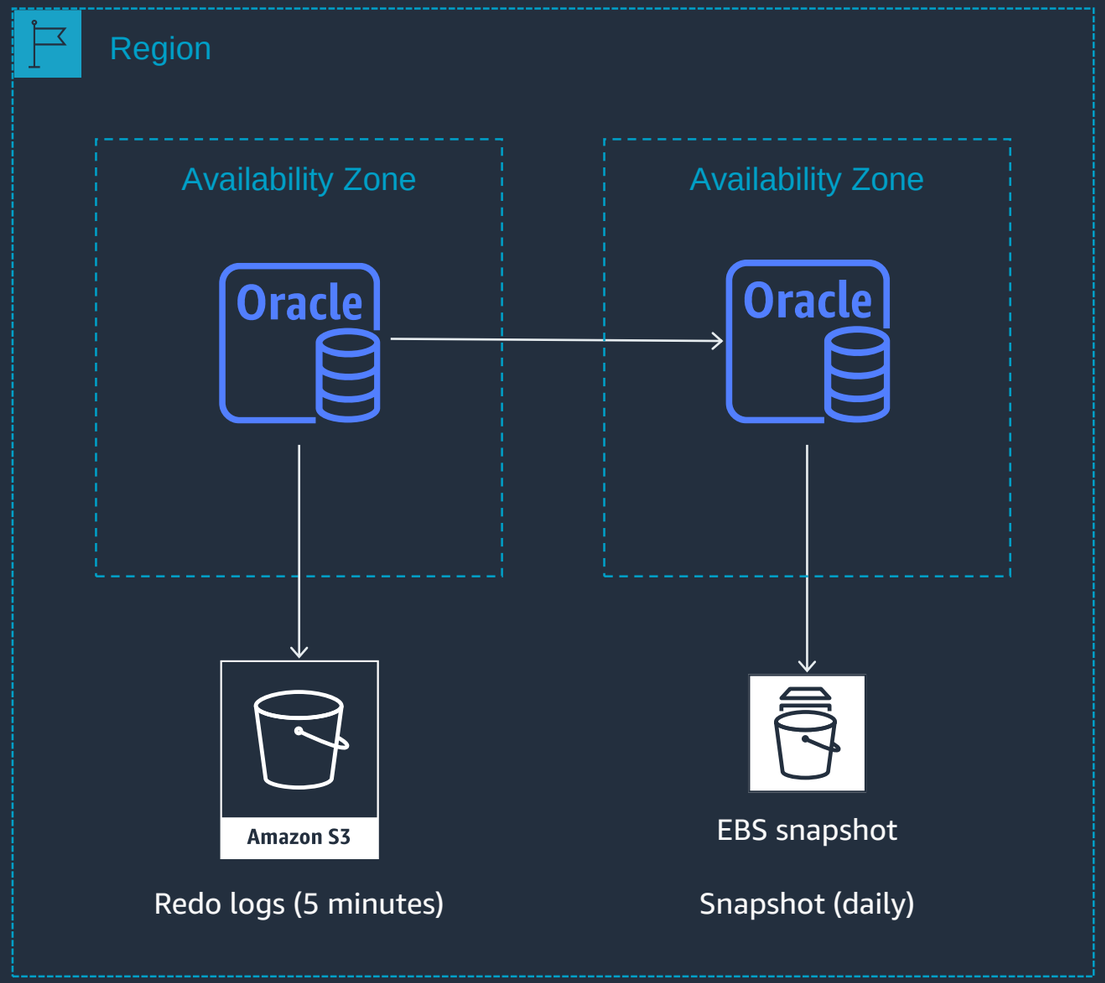

RDS for Oracle – Restore from snapshot

Functionality

- Restore from any snapshot
- Copy snapshots to other regions or accounts

Use cases

- Refresh test environments
- Test upgrades and changes
- Instantiate logical replicas

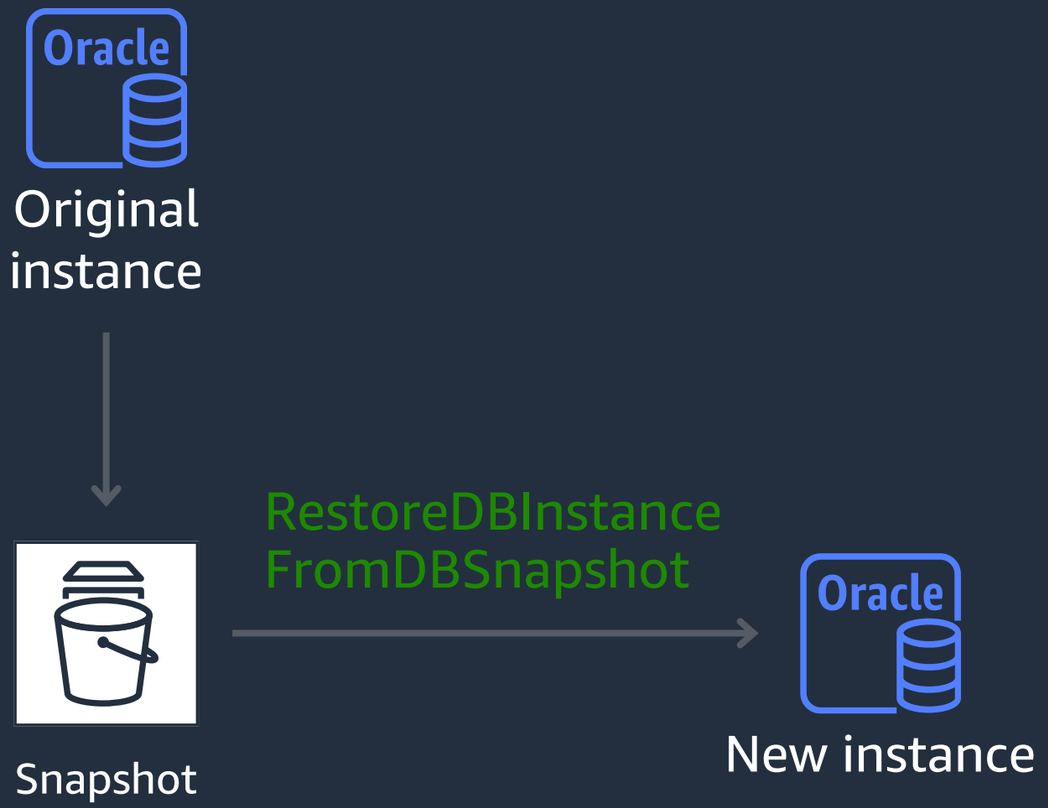

aws

RDS for Oracle – Restore to a point in time

Functionality

- Restore to any second in backup retention
- Available in-region/account
- Latest restorable time typically < 5 minutes

Use cases

- Oops… I dropped a table
- Recover from application errors or logical corruption

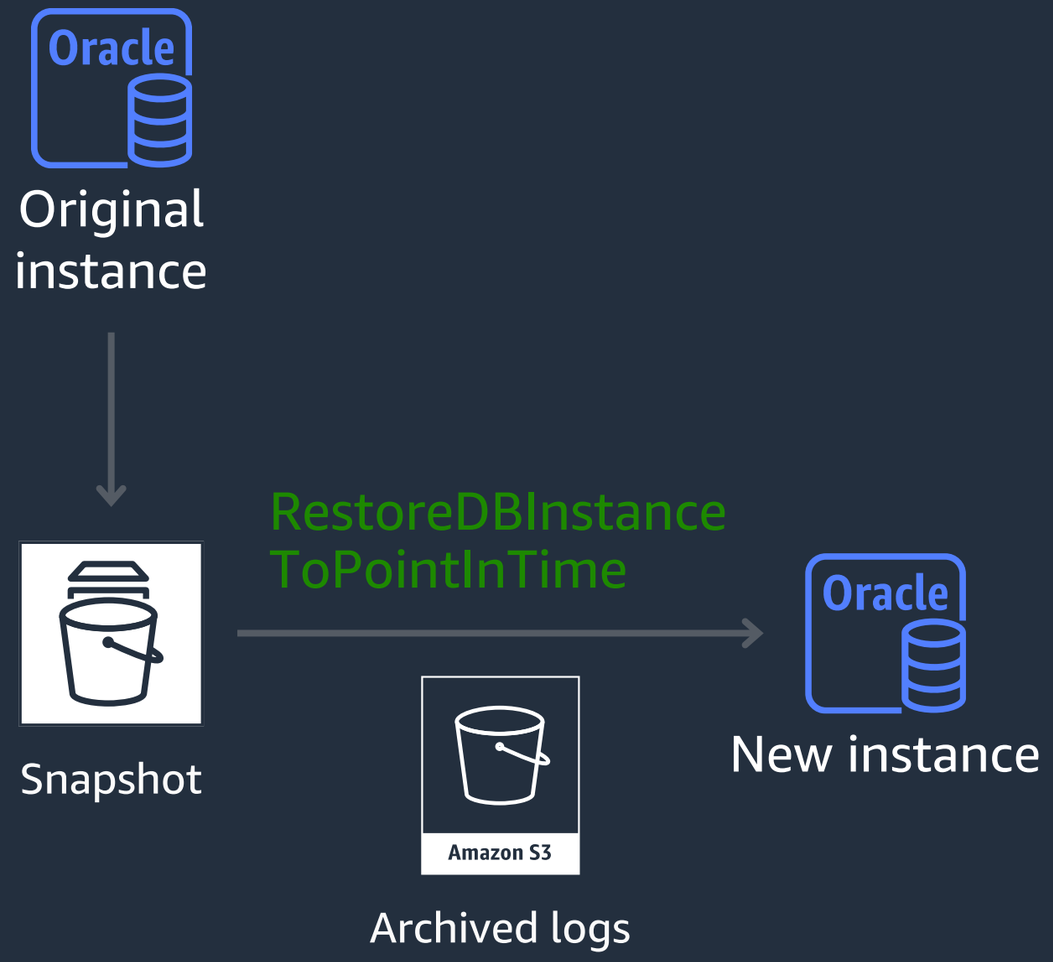

aws

RDS for Oracle – Cross region automated backups

Cross-Region Key Aspects

- Automated snapshots and archived redo logs replicated to target region as soon as available in source region
- Specify independent recovery window for replicated backup region
- Enables Point In Time Restore (PiTR) in second region for mission-critical databases

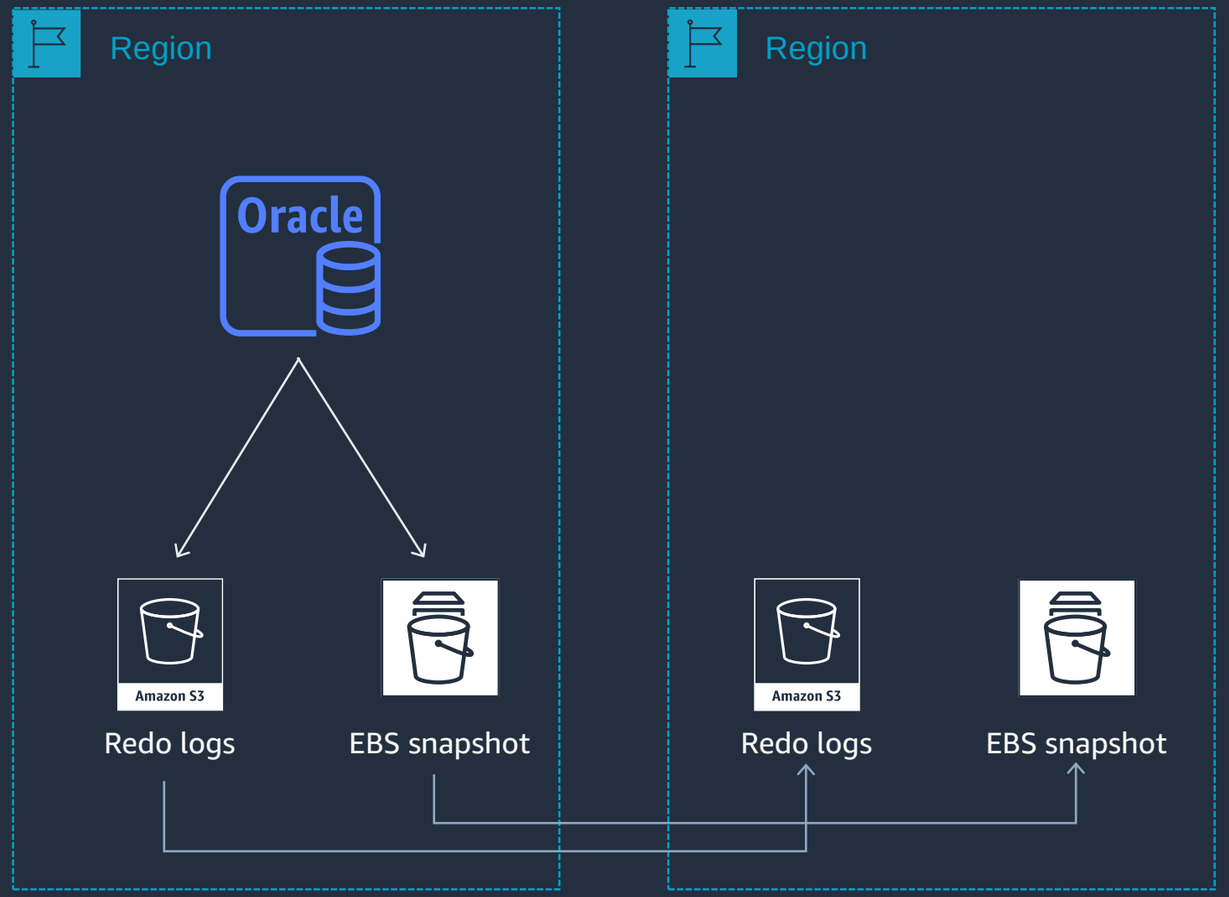

aws

RDS for Oracle – Backup and restore best practices

- ✓ Disable backups for data load (NOARCHIVELOG mode)
  - WARNING: deletes existing automated backups
- ✓ Enable backups for critical workloads (ARCHIVELOG mode)
- ✓ Set backup window to low-usage time
- ✓ Take manual snapshots to reduce PiTR replay duration
- ✓ Use restores to test upgrades/parameters/app changes
- ✓ Copy snapshots to other accounts/regions

aws

RDS for Oracle – Replicas for resiliency

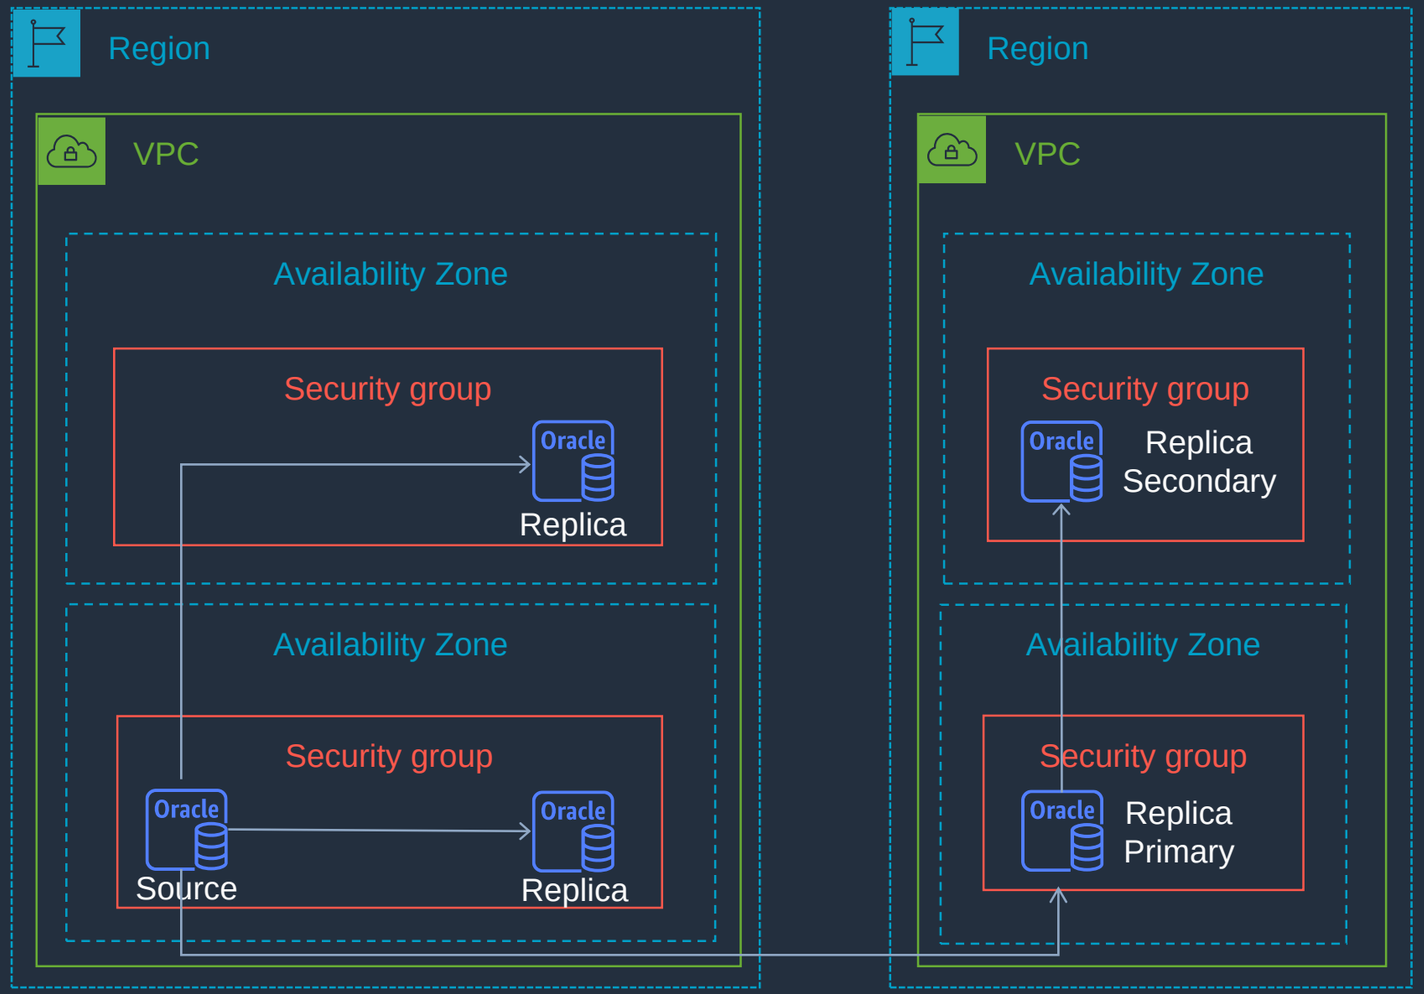

Any replica may also be deployed in a Multi-AZ configuration

aws

Migrating to RDS for Oracle

aws

Migrating Oracle databases to AWS

- Tools available for both Oracle on EC2 and RDS for Oracle:
  - Oracle Data Pump
  - Oracle Export/Import Utility
  - Oracle Materialized Views
  - Oracle GoldenGate
  - SQL*Loader
  - Amazon Database Migration Service (DMS)

- Tools available for Oracle on EC2 only:
  - RMAN restore and duplicate database
  - File system sync or other datafile file transfer method

aws

Amazon Database Migration Service

- Performs logical replication of source to target database
- Many source-target scenarios supported
  - On-premises ←→ Oracle on EC2 ←→ RDS ←→ On-premises
  - Etc.
- Homogenous and heterogeneous migrations
  - Oracle ←→ Oracle
  - Oracle ←→ PostgreSQL
  - Microsoft SQL Server ←→ MySQL
  - Etc.
- Heterogeneous migrations enabled using AWS Schema Conversion Tool

aws

Migrating with Oracle Data Pump and DMS

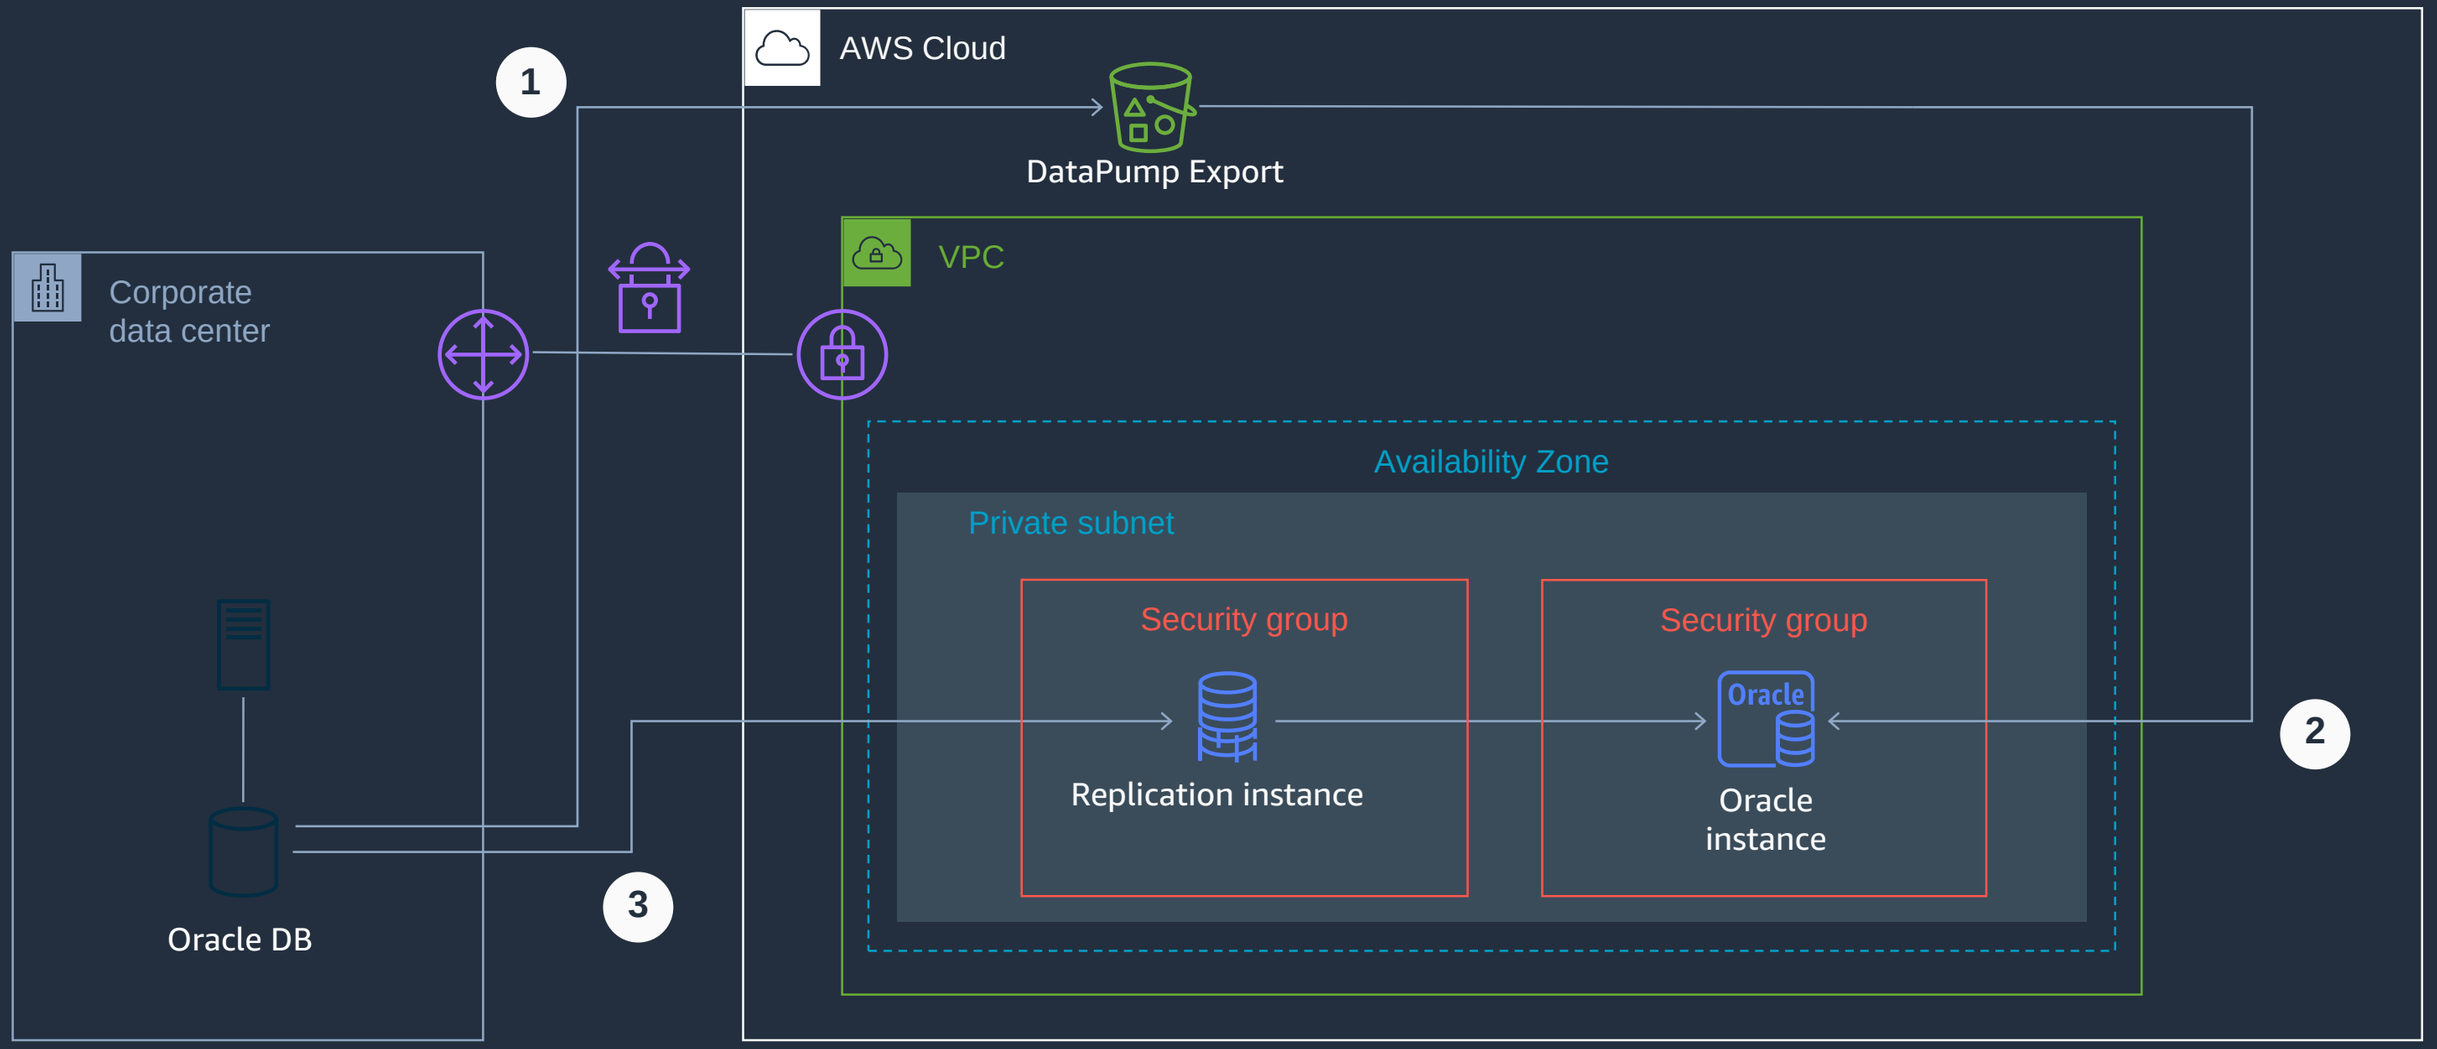

RDS for Oracle Best Practices

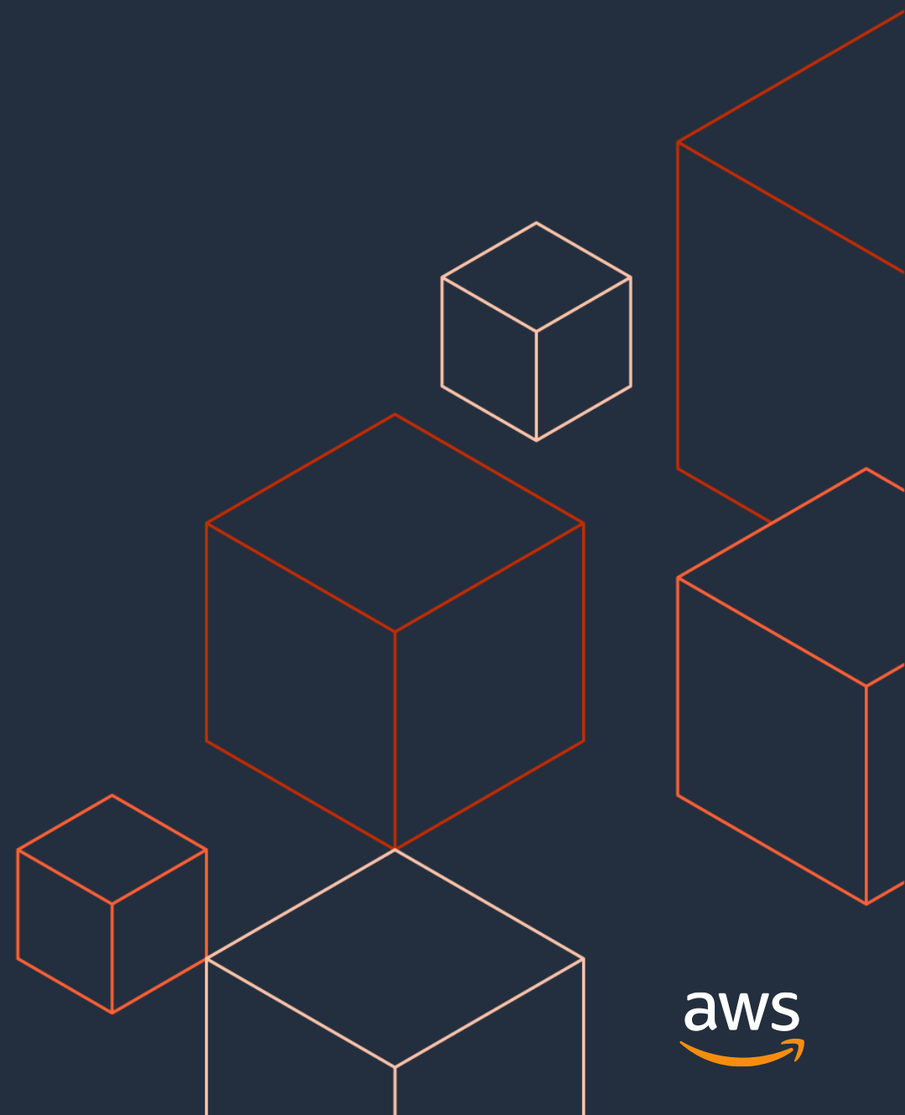

RDS for Oracle best practices

- ✓ Run only as big a configuration as you need
  - vCPU, memory, storage type, IOPS, network bandwidth
  - Consider burstable instance types (t2, t3)
  - Don’t forget you can scale down too
  - Consider Standard Edition vs Enterprise Edition
  - Remember: You license compute, not storage or memory
- ✓ Size instance storage appropriately for initial load
  - Storage autoscaling is great for steady-state operations
- ✓ Storage modification throttling considerations
  - One storage change every 6 hours – includes autoscaling

aws

RDS for Oracle best practices

- ✓ Troubleshooting performance issues, identifying bottlenecks
  - Use available tools: CloudWatch, Enhanced Monitoring, Performance Insights, AWR / Statspack reports
- ✓ Instance restore performance considerations
  - Restore is optimized for availability and loads lazily
- ✓ DBA role changes and expectations
  - Offload tedious tasks, spend more time at application layer

aws

Rethinking Standard Edition

- ✓ High availability with RDS Multi-AZ
  - Synchronous replication and automated failover – 99.95% uptime SLA
  - Independent infrastructure
- ✓ Disaster Recovery with RDS Cross-region Automated Backups
  - Automatically copies snapshots and transactions logs to other region
  - Issue a point-in-time restore from target region in the event of disaster
- ✓ RDS Storage encryption with Amazon KMS
  - AES-256 encryption at rest
  - Bring your own keys
- ✓ Tuning with RDS Enhanced Monitoring and Performance Insights
  - Per-process host metrics, granularity down to 1 second
  - Analyze database load and active sessions: real-time and historical

aws

Q&A

Nathan Fuzi

aws