RDS for Oracle High Availability And Performance

Dallas Willett, Senior Database Engineer

Agenda

- **Introduction**
- High Availability
- Performance

Enterprise Apps on AWS – RDS Advantages

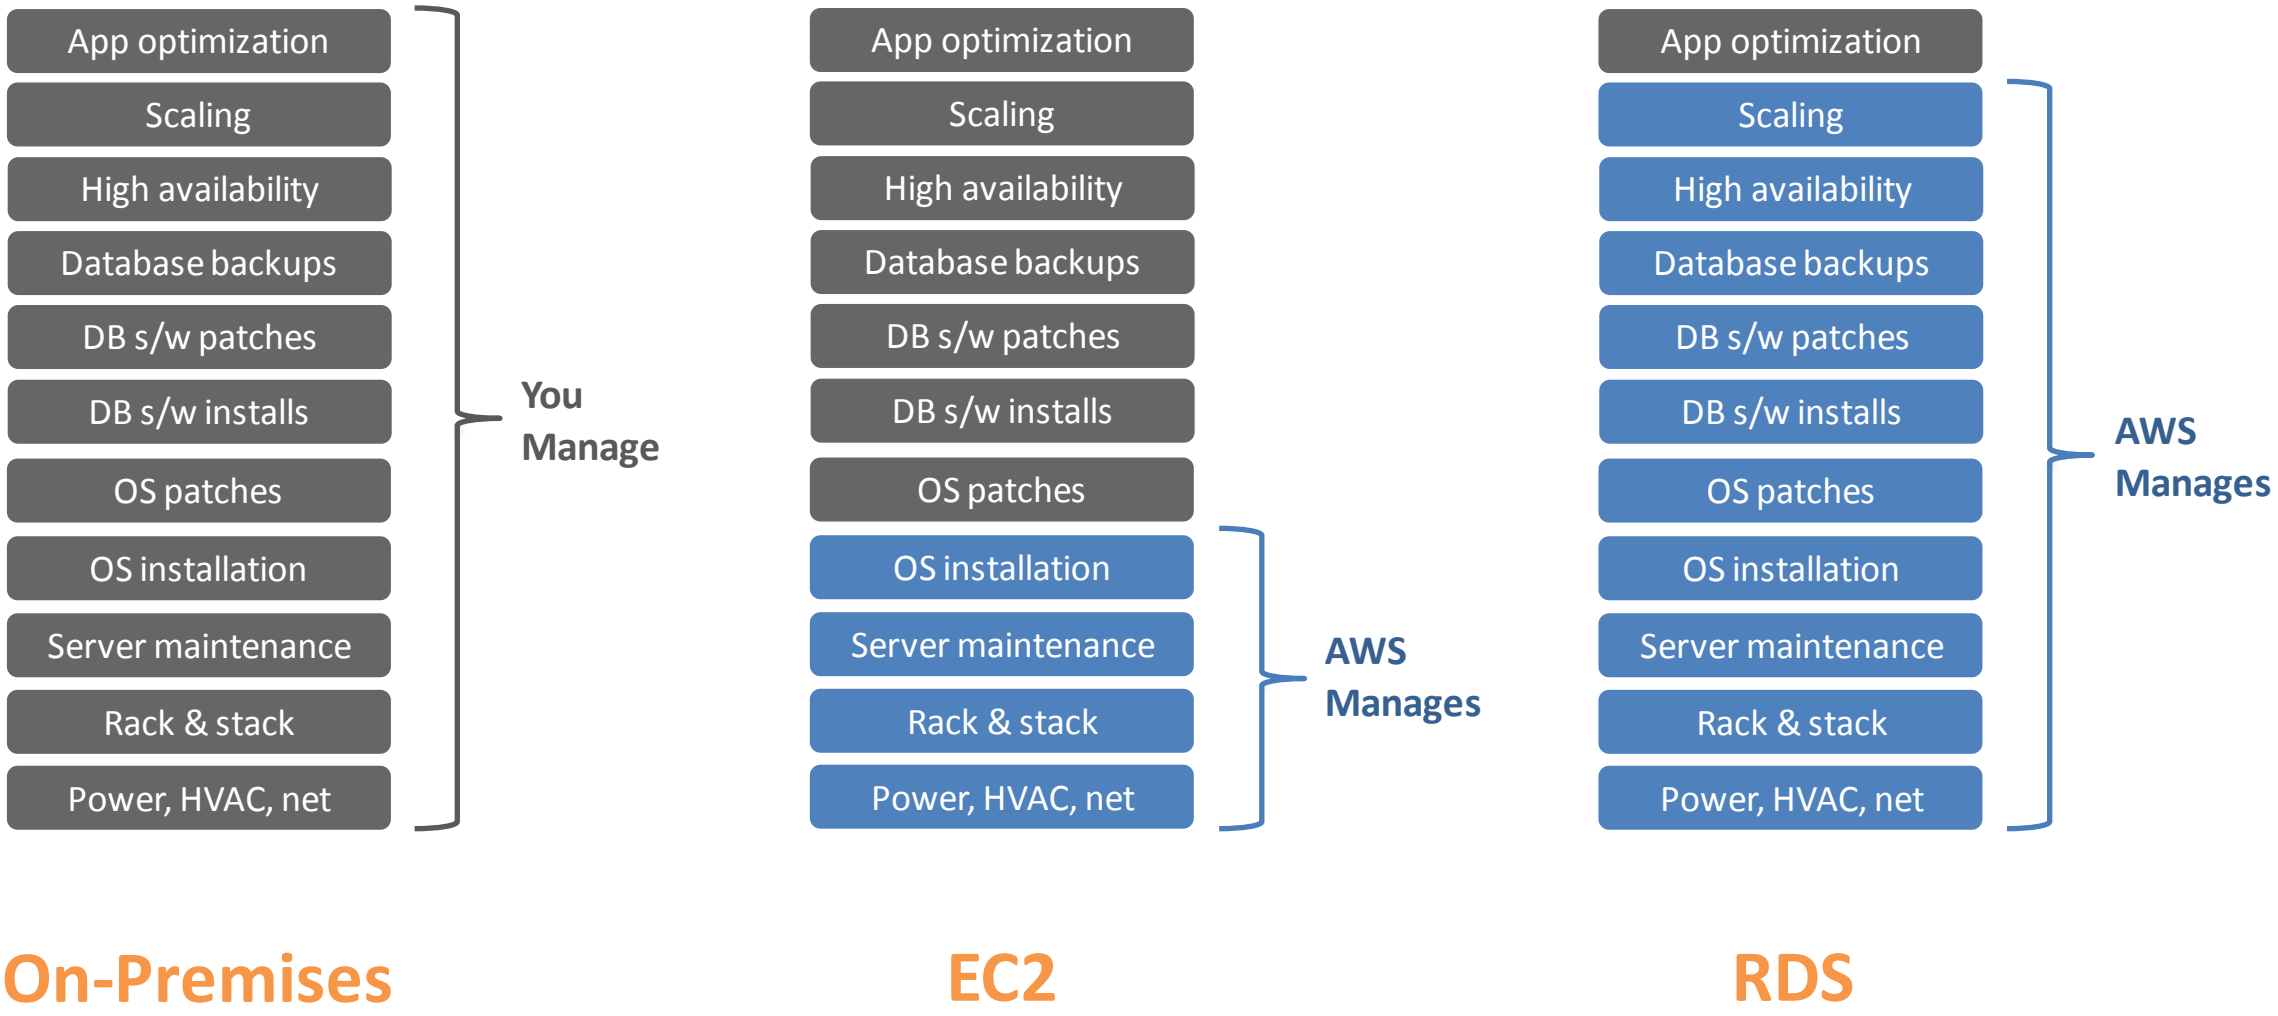

On-Premises EC2 RDS

Enterprise Apps on AWS – RDS Advantages

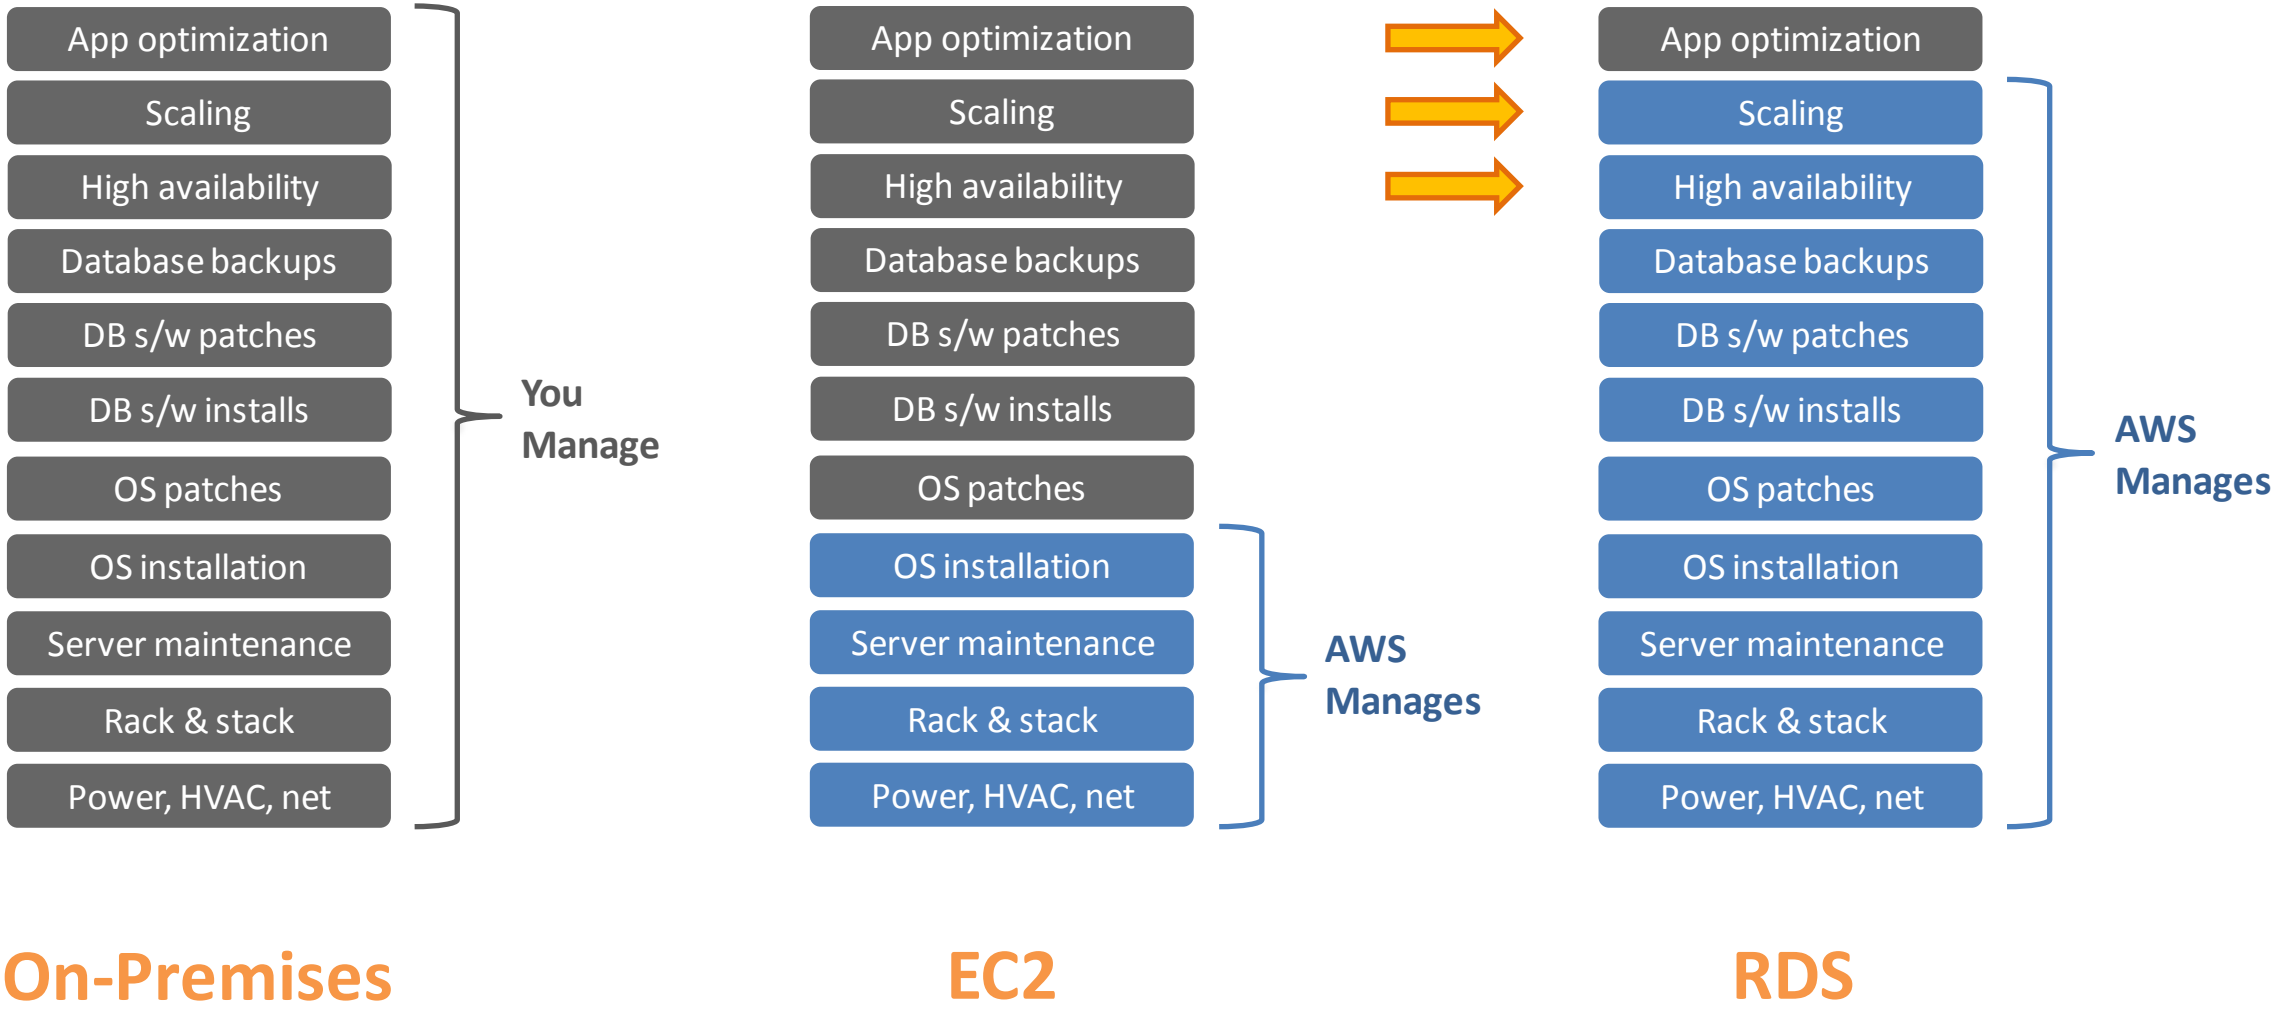

Agenda

- Introduction
- **High Availability**
  - **RDS Mutli-AZ**
  - **Roll-your-own Multi-Master with Logical Replication**
- Performance

RDS Multi-AZ

- Primary in one AWS Availability Zone (AZ)
- Secondary in a different AZ
- Synchronous replication between Primary and Secondary
- AWS monitors database and will automatically fail-over
- Zero data loss

- Used for all RDS Linux based engines
- Multi-AZ can be added or removed at any time (online – no outage)
- Can be used with Oracle Standard Editions

AWS Global Infrastructure (2019)

Regions and Availability Zones (AZ)

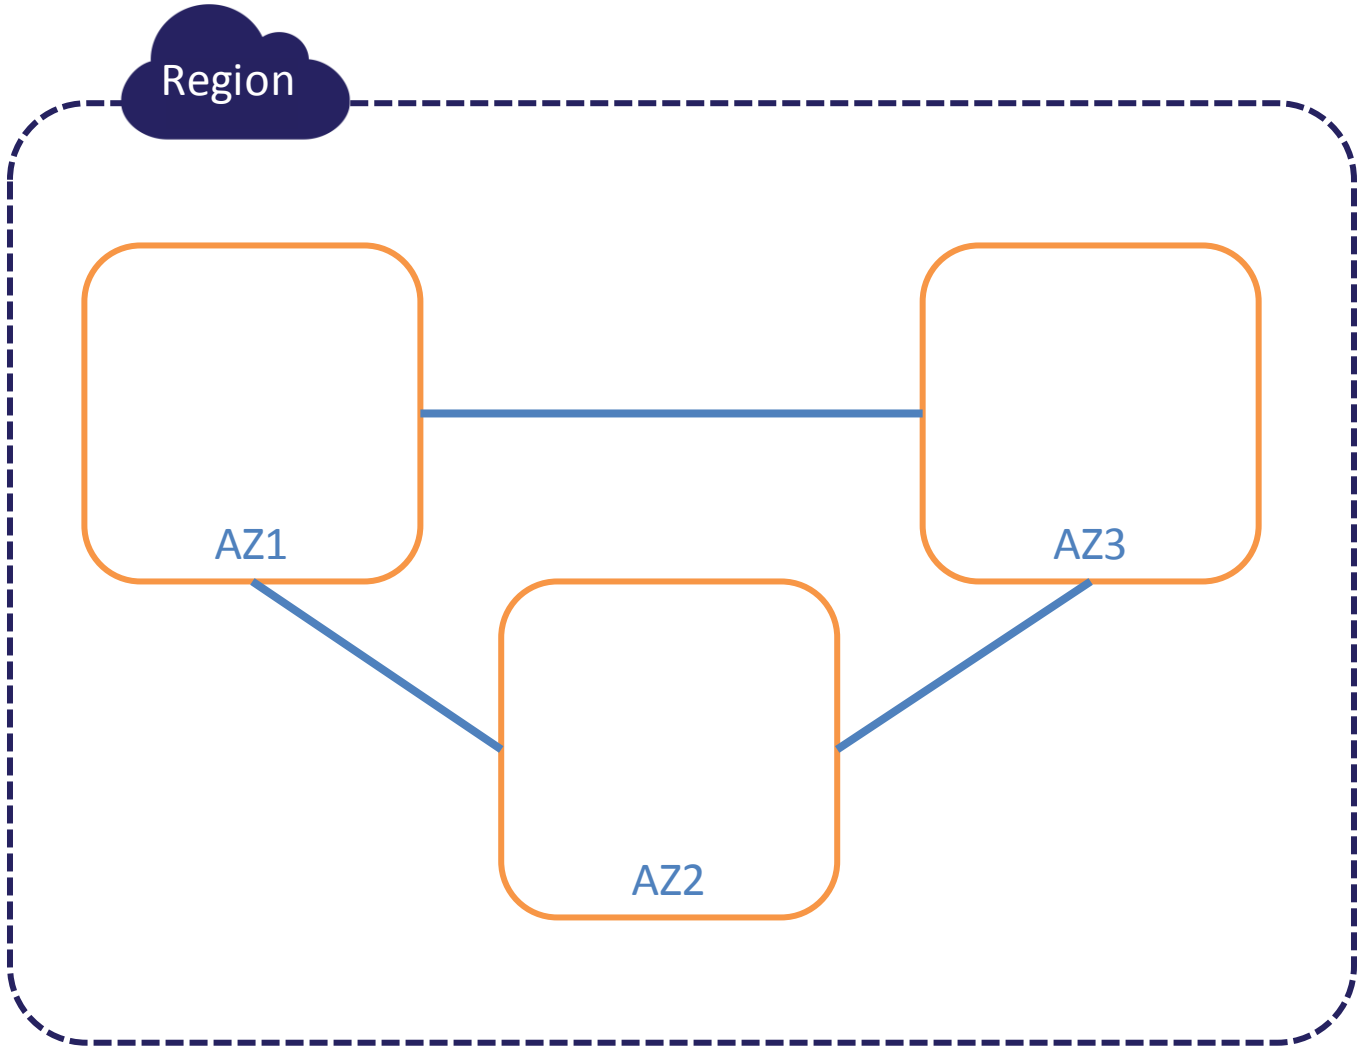

Single Availability Zones (Single-AZ)

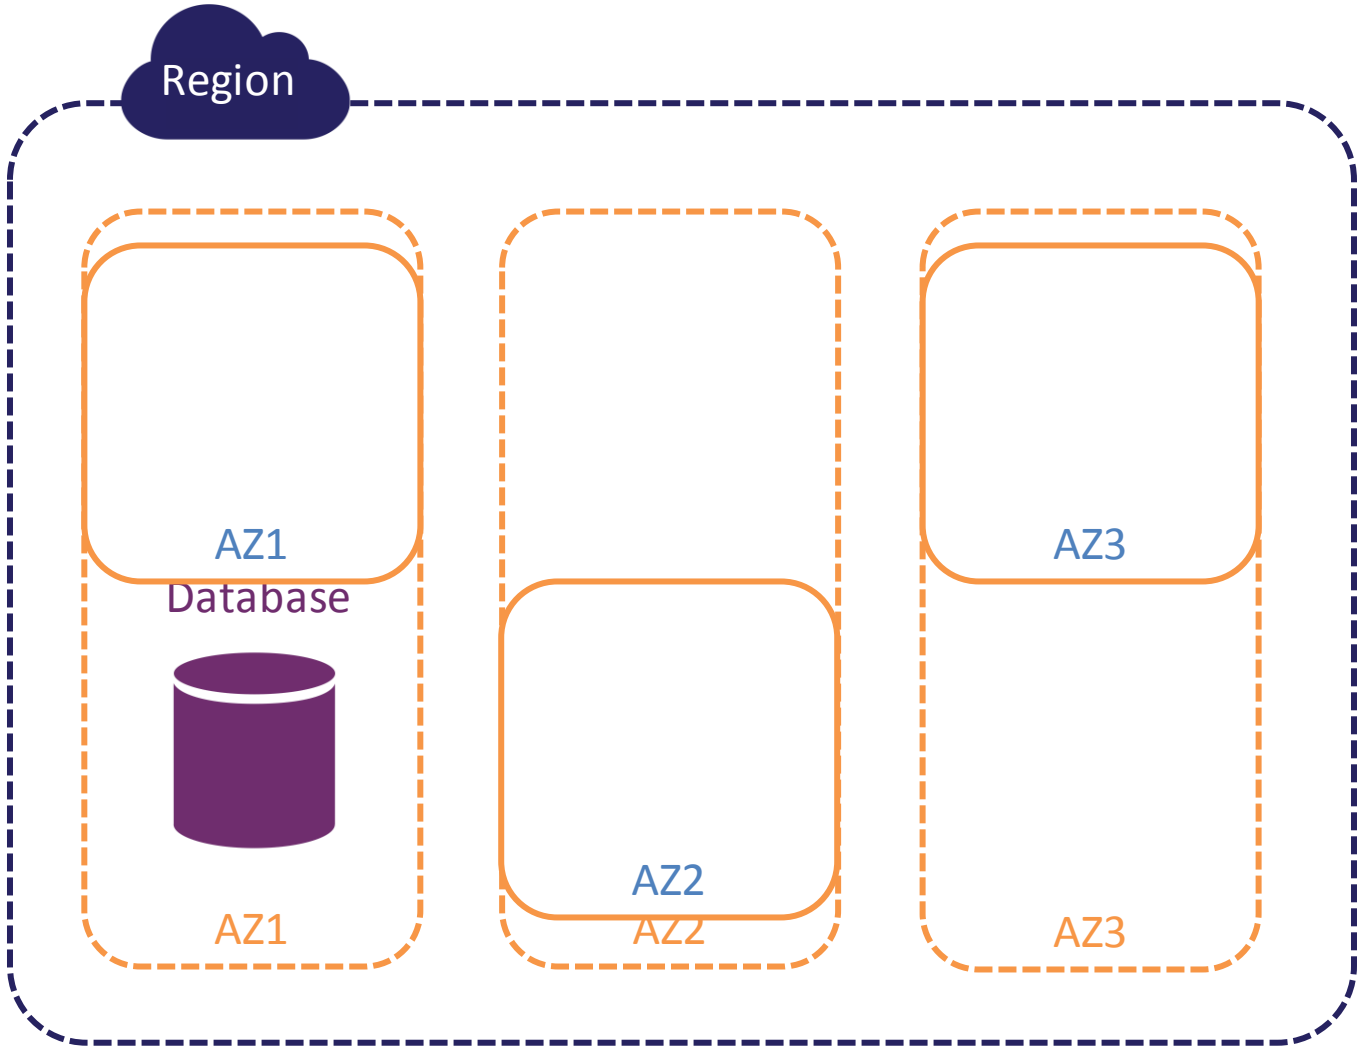

Multiple Availability Zone (Multi-AZ)

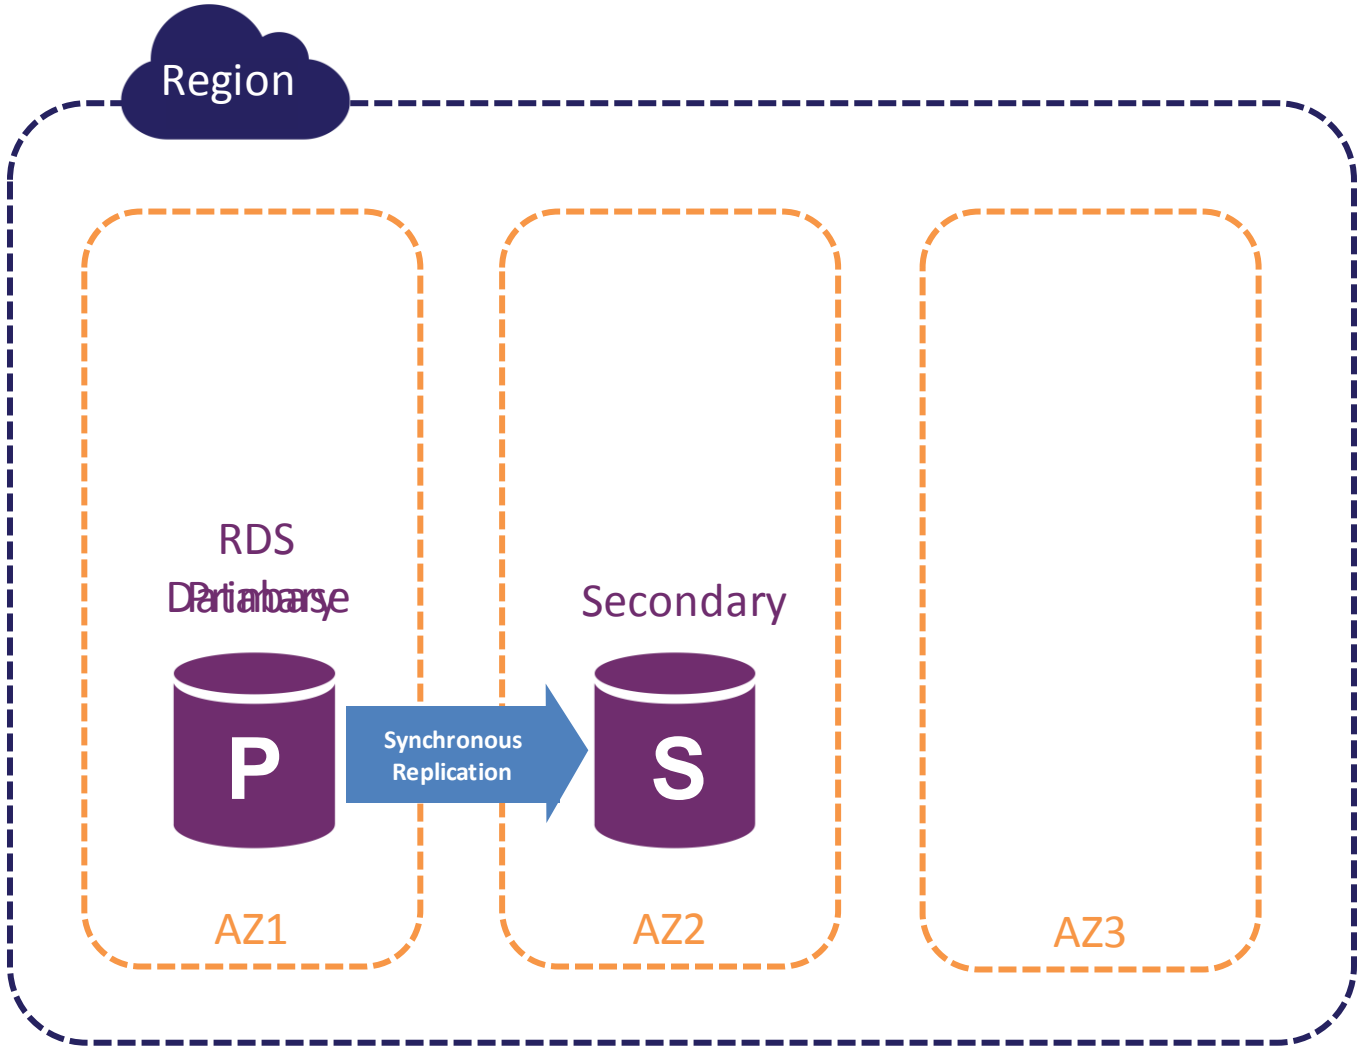

Multi-AZ – Instance Failure Scenario

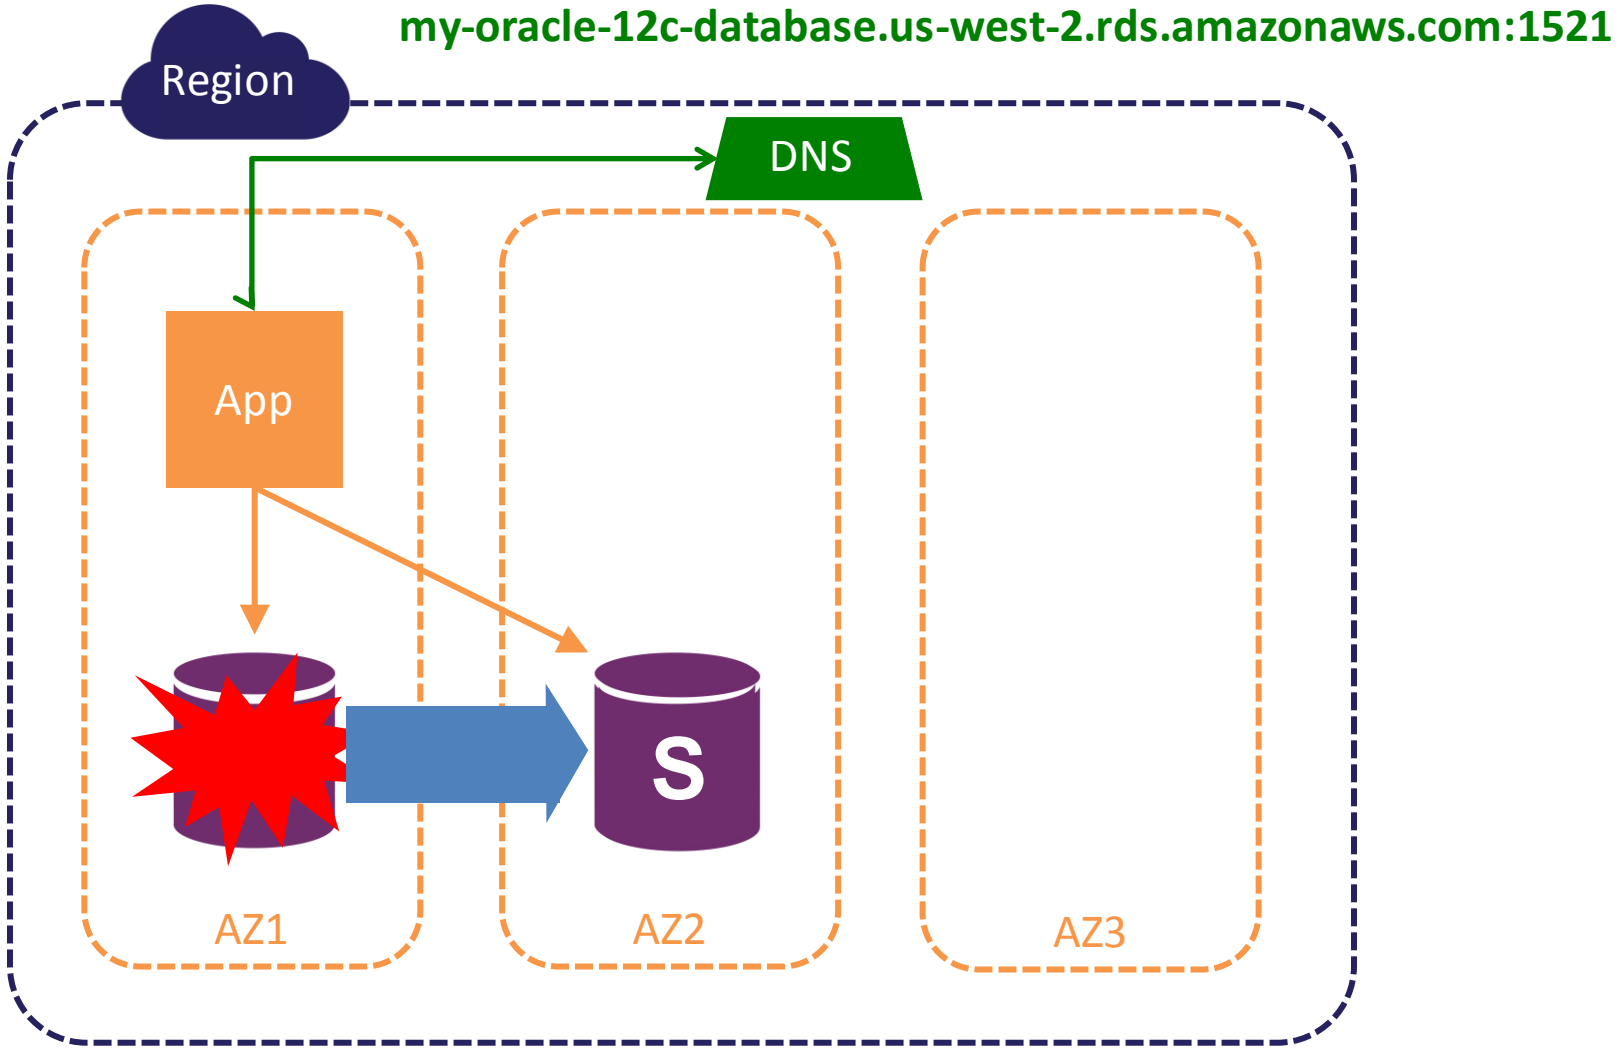

Multi-AZ – Instance Failure Scenario

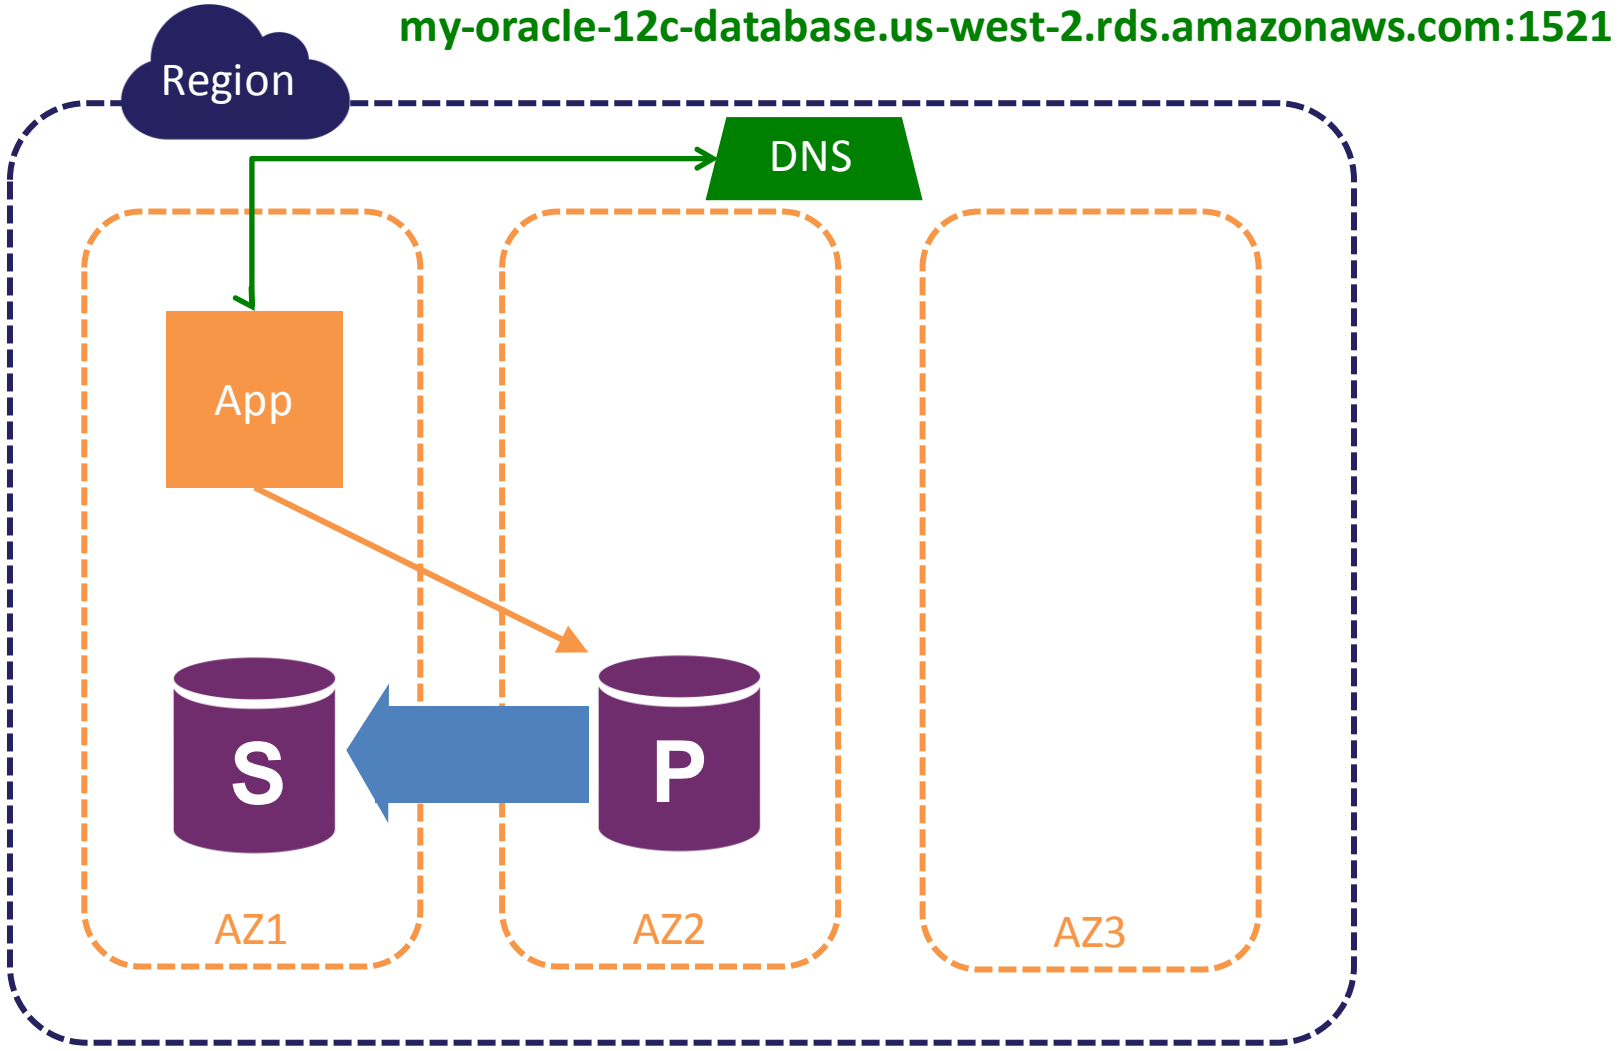

Multi-AZ – AZ Failure and Applications

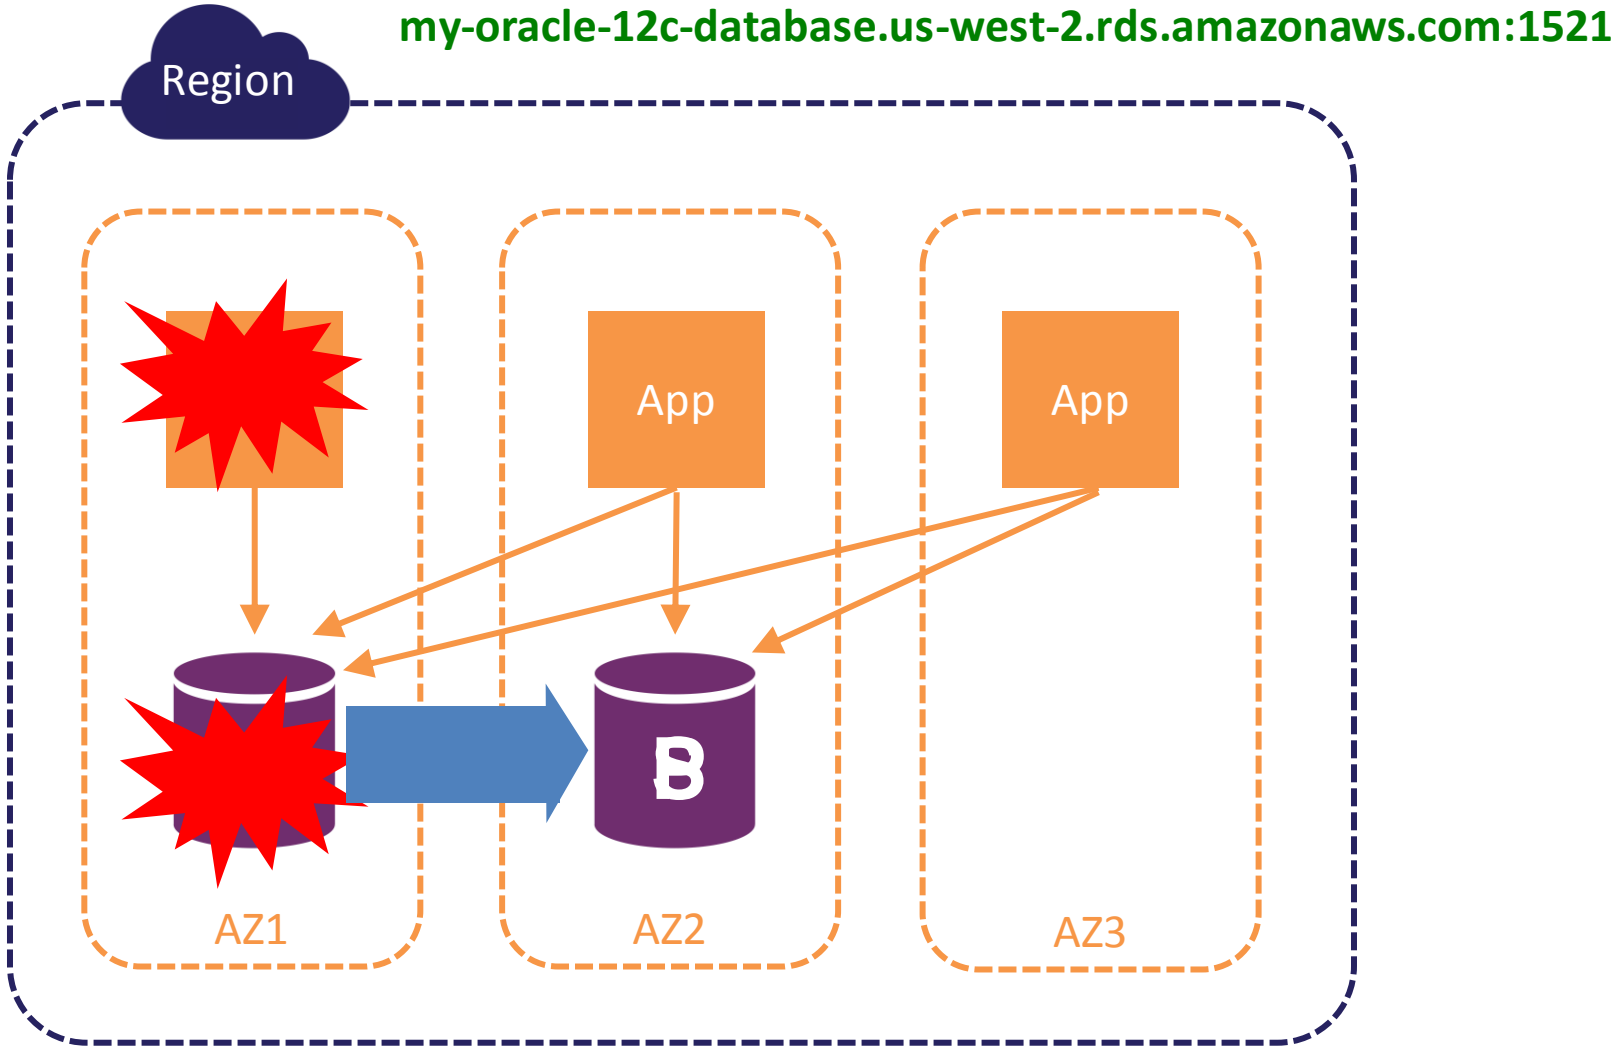

Multi-AZ Overview

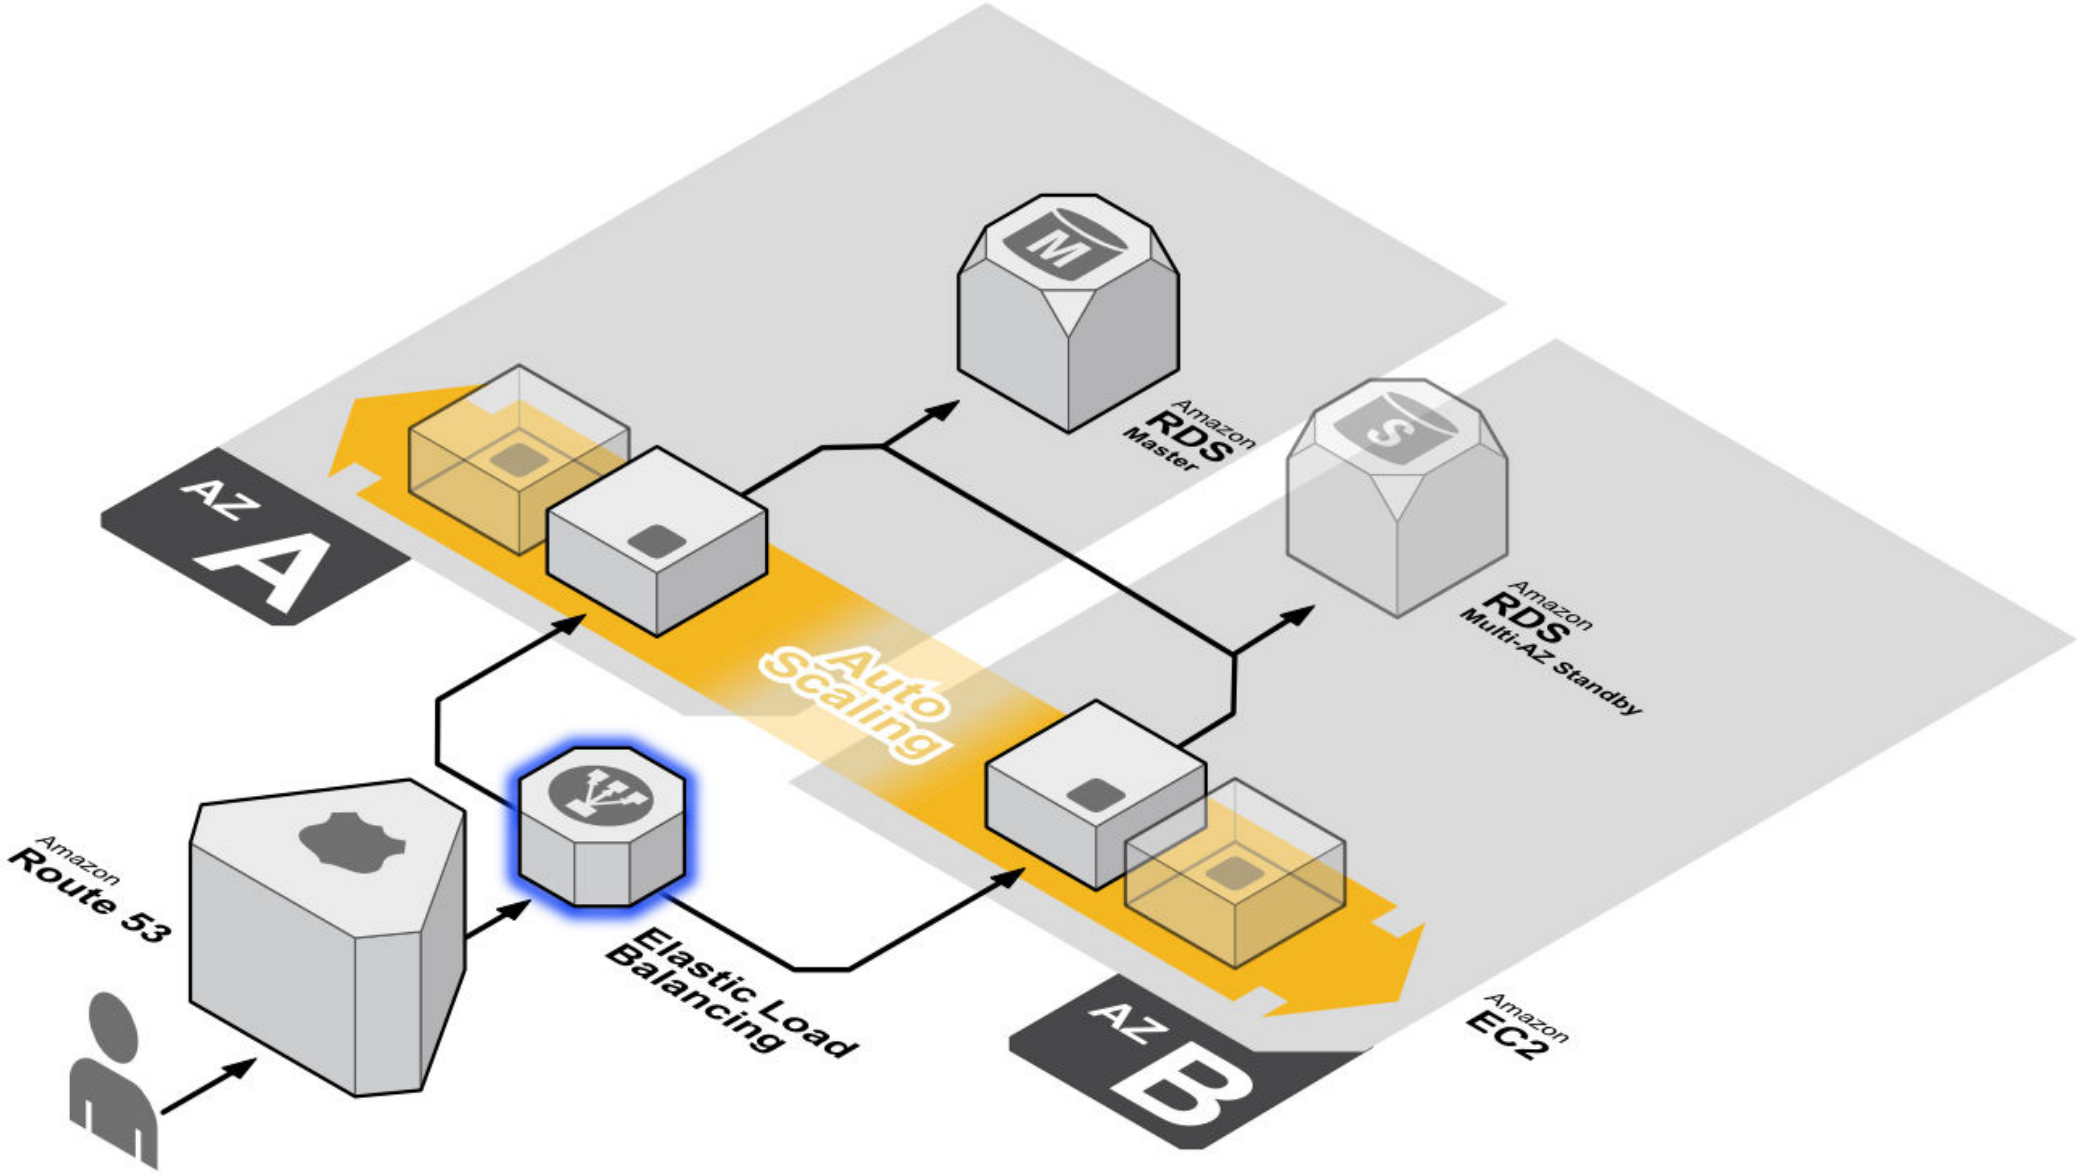

Multi-Master with RDS for Oracle

Less common configuration, but very good availability (can be better than RAC)

Configuration:

- Two separate databases
- Logical replication between databases
  - AWS Database Migration Service
  - Oracle Goldengate
  - Other logical replication technologies
- Application knows about both databases and can fail-over between them
- Avoid split-brain, have a data conflict resolution plan

Agenda

- Introduction
- High Availability
- **Performance**
  - **Performance Tools**
  - **RDS Instance Class**
  - **RDS Storage**
  - **RDS Restores**
  - **RDS Multi-AZ**
  - **Database Parameters**
  - **Migrating to RDS**

Performance Tools

Most of your existing tools to analyze performance work with RDS Oracle:

- V$ views
- OEM_AGENT and OEM
  - Installed via RDS Option Groups
- AWR / Statspack
  - AWR requires Diagnostic and Tuning pack
  - STATSPACK installed via RDS Option Group

We have also developed our own:

- RDS Performance Insights
  - Performance Graphs and SQL Analysis (similar to OEM)
  - Can be used with Standard Editions

Performance Tools – Performance Insights

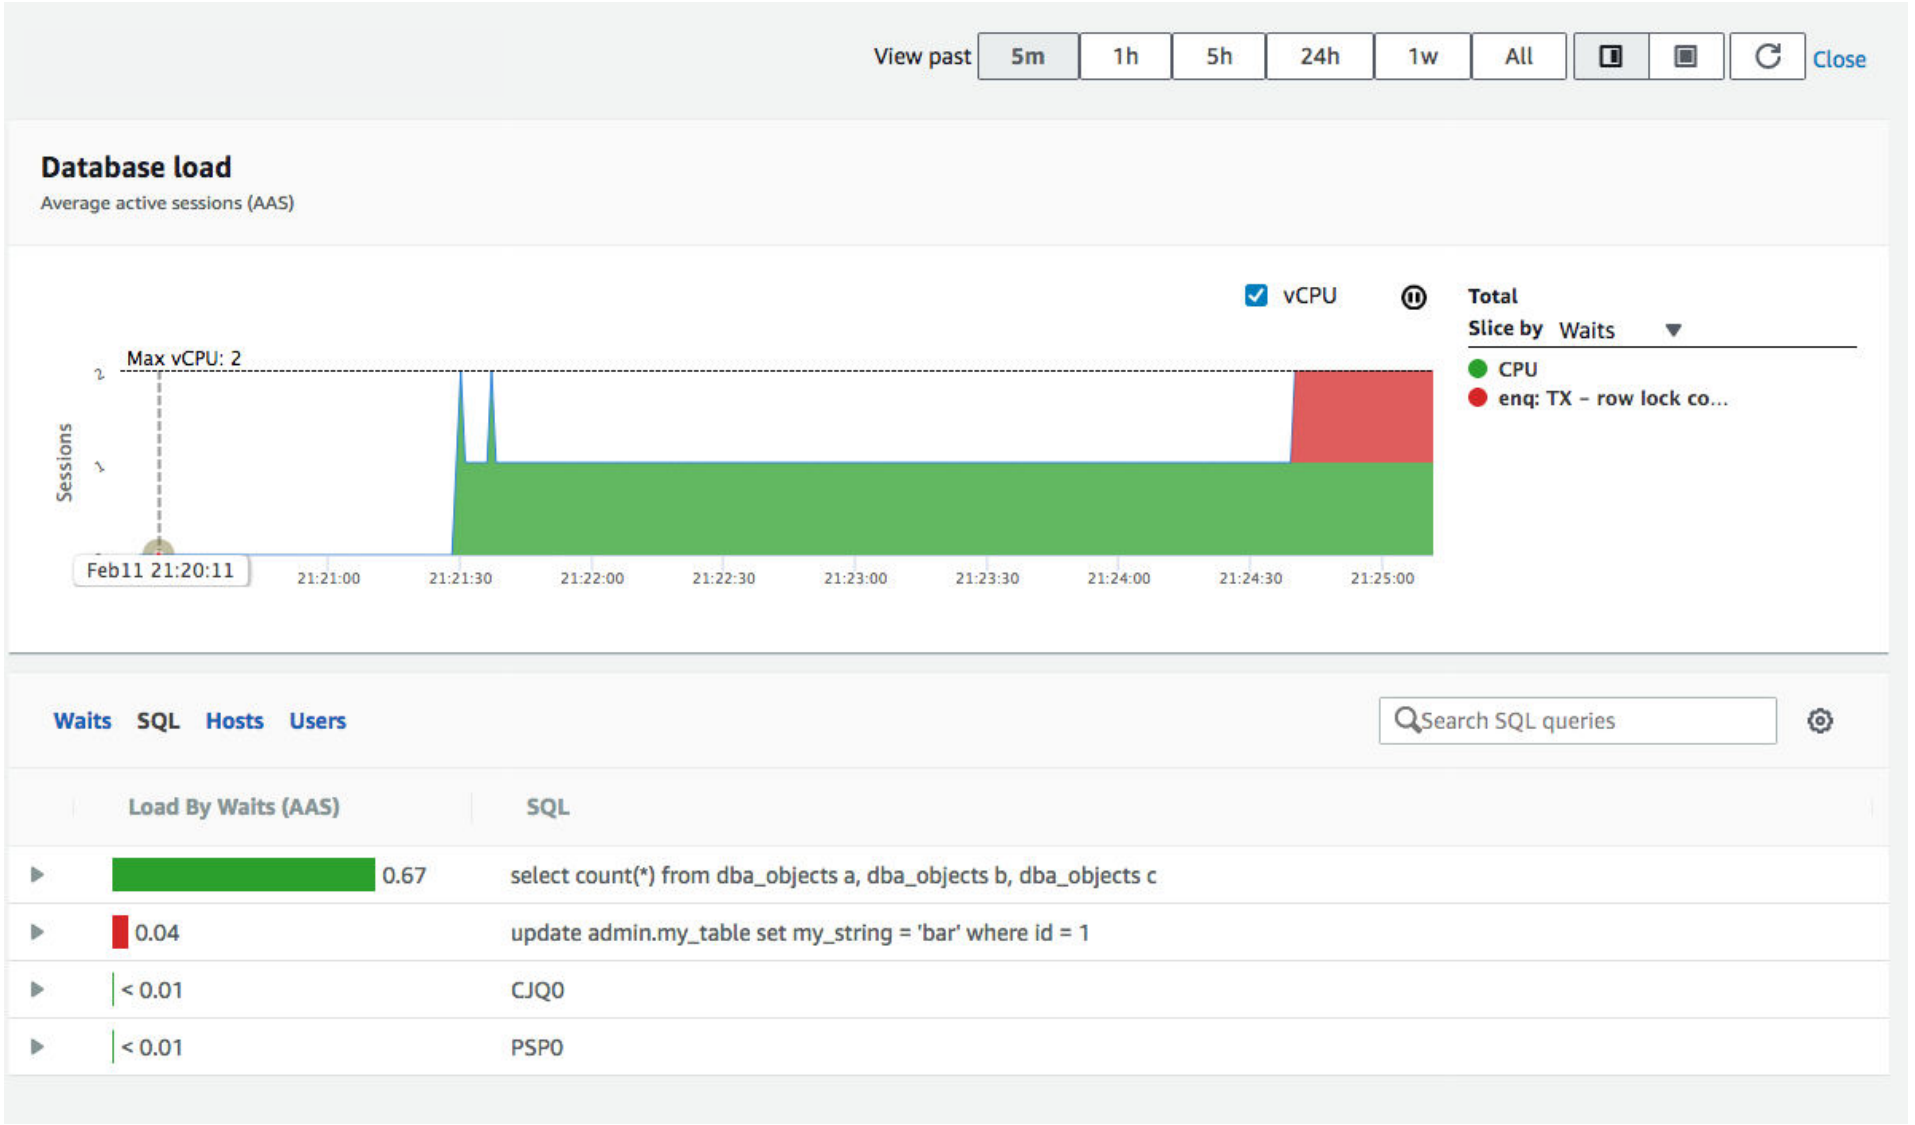

RDS Instance Class

Determines

- Number of vCPUs
- Memory
- Network Throughput
- EBS max IOPS and Throughput

Can be changed at any time (with an outage) - scale-up or scale-down

RDS Instance Class

RDS DB Instance Class

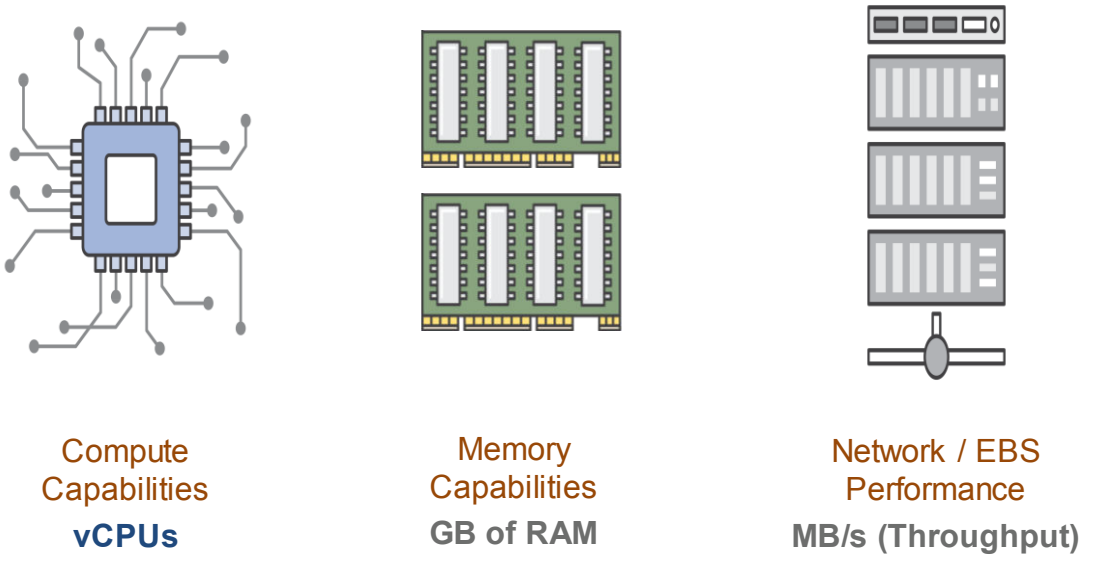

Compute Capabilities
vCPUs

Memory Capabilities
GB of RAM

Network / EBS Performance
MB/s (Throughput)

Instance class families:

General Purpose (**M5, M4**) – 1 vCPU : 4 GB RAM ratio
Memory Optimized (**R5, R4**) – 1 vCPU : 8 GB RAM ratio
Burstable Capacity (**T3, T2**) – lower cost – save/spend CPU credits
High Memory (**X1e, X1**) - 1 vCPU : 32 GB RAM or 16 GB RAM ratio

Range of DB instance classes:

Very small: 1 vCPU, 1 GB RAM (t2.micro)
Largest burstable: 8 vCPU, 32 GB RAM (t3.2xlarge)

CPU Bound workloads:

From: 2 vCPU, 8 GB RAM (m5.large)
To: 96 vCPU and 384 GB RAM (m5.24xlarge)
Or: 128 vCPU and 1952 GB RAM (x1.32xlarge)

Memory Bound workloads:

From: 2 vCPU, 16 GB RAM (r5.large)
To: 96 vCPU and 768 GB RAM (m5.24xlarge)
Or: 128 vCPU and 3904 GB RAM (x1e.32xlarge)

IO Bound workloads:

- Typically choose largest of M or R class
- Optimize CPU feature may keep licensing costs low

RDS Instance Class – CPU

RDS vCPU typically corresponds to a hyper-thread on a CPU core

- **R5** - Intel Xeon Platinum 8000 series (Skylake-SP) processors with a sustained all core Turbo CPU clock speed of up to 3.1 GHz
- **R4** - High Frequency Intel Xeon E5-2686 v4 (Broadwell) processors (2.3 GHz)

- **M5** - 2.5 GHz Intel Xeon® Platinum 8175 processors with new Intel Advanced Vector Extension (AXV-512) instruction set
- **M4** - 2.3 GHz Intel Xeon® E5-2686 v4 (Broadwell) processors or 2.4 GHz Intel Xeon® E5-2676 v3 (Haswell) processors

- **T3** - 2.5 GHz Intel Scalable Processor
- **T2** - High frequency Intel Xeon processors

- **X1e/X1** - High frequency Intel Xeon E7-8880 v3 (Haswell) processors

RDS Instance Class – Memory

- **R5** – 1:8 vCPU to Memory ratio
- **R4** – 1:8 vCPU to Memory ratio

- **M5** - 1:4 vCPU to Memory ratio
- **M4** - 1:4 vCPU to Memory ratio

- **T3** - 1:4 vCPU to Memory ratio typically
- **T2** - 1:4 vCPU to Memory ratio typically

- **X1e** - 1:32 vCPU to Memory ratio (EBS throughput may be an issue)
- **X1** - 1:16 vCPU to Memory ratio (EBS throughput may be an issue)

RDS Instance Class – EBS Bandwidth

EBS Bandwidth

- Max EBS Bandwidth today is 14,000 Mbps
- Max EBS Throughput is 1,750 MB/s
- Max RDS Oracle IOPS is 40,000 (EC2 max IOPS is 80,000)
- https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSOptimized.html

14 Gbps EBS-Optimized Instances

- m5.24xlarge (96 vCPU)
- r5.24xlarge (96 vCPU)
- r4.16xlarge (64 vCPU)
- x1e.32xlarge (128 vCPU)
- x1.32xlarge (128 vCPU)

RDS Instance Class – Nitro Hypervisor

Nitro Instances – Hardware card in EC2 instance

- More of the physical host resources are available for the VM
- Provides CPU and memory isolation for EC2 instances
- VPC networking and EBS storage resources are implemented by dedicated hardware components

RDS Nitro Instances:

- R5, M5, T3

RDS Instance Class – Optimize CPU

Optimize CPU

- Reduce number of vCPUs on a given instance to:
  - Lower licensing costs
  - Scale up to larger instance for more memory or EBS bandwidth
- Minimum 1 core (2 vCPU)
  - On multi-socket instances, minimum is 1 core per socket (to access memory)
  - Can disable Intel Hyper-Threading Technology if you want

NOTE:

- Infrastructure cost remains the same
  - Typically license cost is the majority of the TCO
- Changes require reboot

RDS Storage

RDS Storage includes:

- Data Files
- Temp Files
- Online Redo Logs
- Archive Redo Logs
  - Before they are backed up to S3
  - Optionally if ”retention hours” is configured – will remain on RDS storage
- Other Files
  - Log files
  - Anything in Oracle Directory Objects. E.g., Data Pump files

RDS Storage

https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSVolumeTypes.html

gp2 – General Purpose SSD (burst)

- Lower cost – $0.115 per GB-month in us-west-2
  - E.g., $115.00 for 1000 GiB per month
- Max Throughput – 250 MB/s (per volume)

io1 – Provisioned IOPS SSD

- Higher cost – $0.125 per GB-month in us-west-2 plus $0.10 per IOPS-month
  - E.g., $425.00 for 1000 GiB and 3000 IOPS per month
- Max throughput – 1000 MB/s (per volume)

NOTE: RDS Oracle uses 1 EBS volume for < 200 GB and 4 EBS volumes for >= 200 GB.

Storage type can be changed at any time (online – no outage)

Storage size can be increased at any time (online – no outage)

RDS Storage – gp2 EBS Volumes

Provides a combination of baseline + burst IOPS

- 20 GiB – 32,768 GiB Storage
- Baseline: 3 IOPS / GiB – you specify size, but not IOPS
- Burst: 3000 IOPS (only useful for volumes less than 1 TiB)
- Initial burst balance of 5.4 million I/O credits – enough for 3000 IOPS for 30 minutes

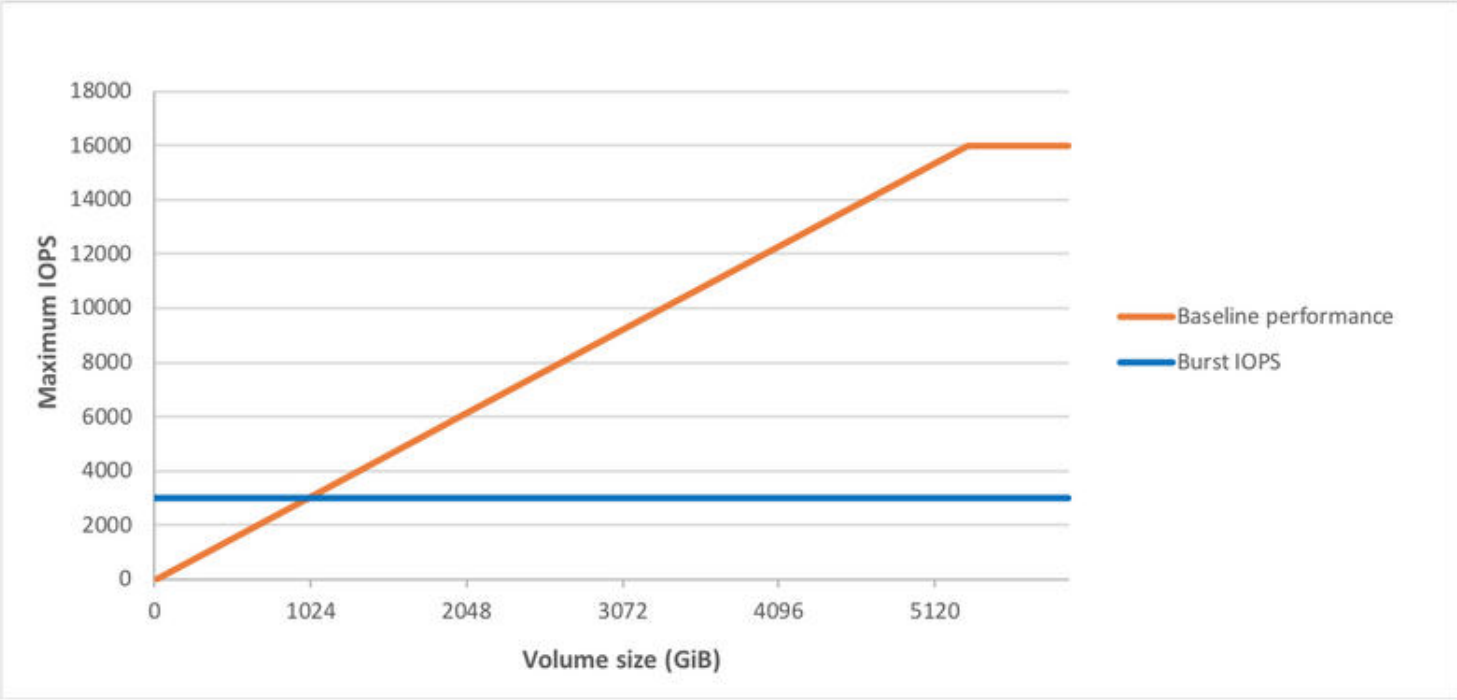

RDS Storage – gp2 EBS Volumes

Example of gp2 burst and credit exhaustion

- All except 2 TiB volume start with 3000 IOPS
- 1 TiB and 2 TiB volumes have consistent performance

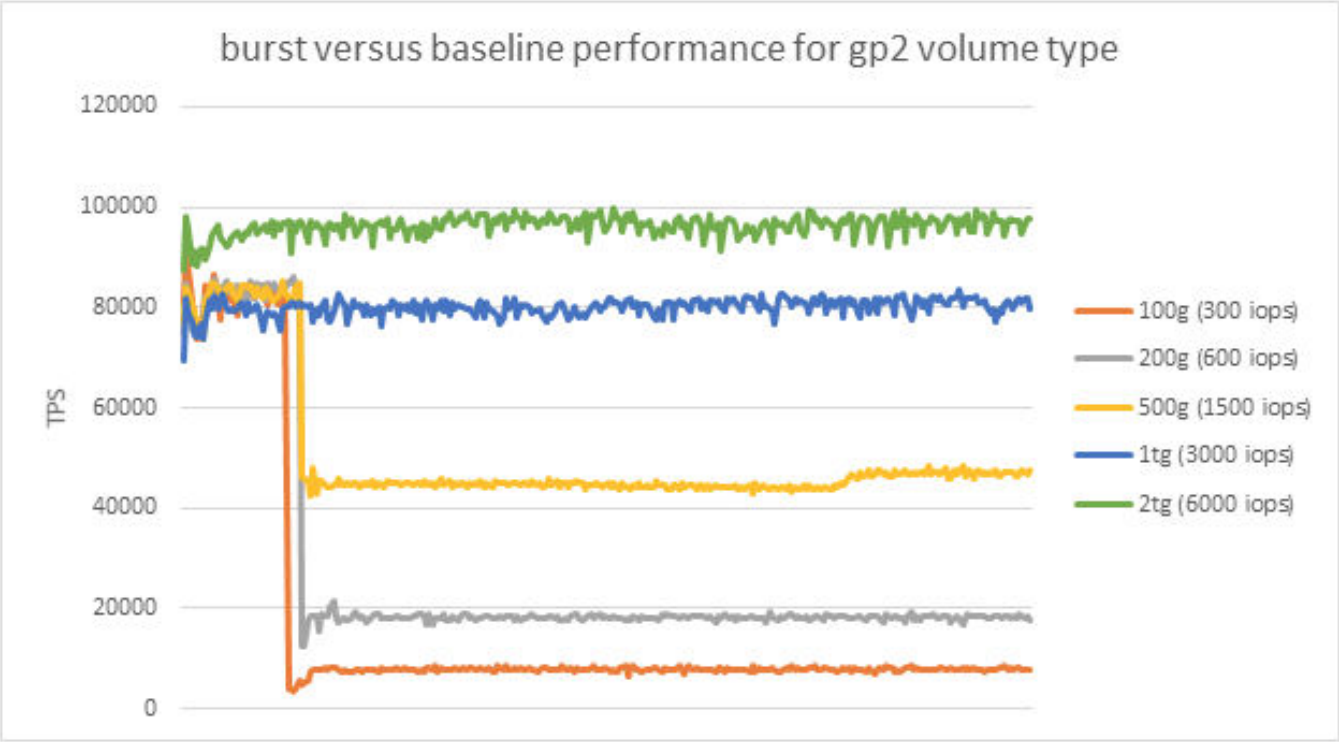

RDS Storage – io1 EBS Volumes

Ideal for workloads sensitive to performance and consistency

- 100 GiB – 32,768 GiB Storage
- 1,000 - 40,000 IOPS with RDS for Oracle
- Max ratio of IOPS:GiB is 50:1 – e.g., 100 GiB volume with 5000 IOPS
- Delivers within 10% of provisioned IOPS 99.9% of the time in a given year

RDS Storage – io1 EBS Volumes

- Up to 32k IOPS, I/O size can be up to 256 KiB, above that 16 KiB size is used
- Single 1,024 KiB operation counts as 4 I/O operations (at 256 KiB)
- Contiguous IO operations can be merged into a single 256 KiB operation

RDS Storage – Throughput Limits

gp2

- 250 MiB/s per volume if 334 GiB or larger
  - 1000 MiB/s with 1,336+ GiB RDS Storage
- 250 MiB/s per volume if between 170-334 GiB and burst credits available
  - 1000 MiB/s with 680-1335 GiB RDS Storage and burst credits available
- 125 MiB/s per volume less than 170 GiB
  - 500 MiB/s with 200-679 GiB RDS Storage
  - 125 MiB/s with 20-199 GiB RDS Storage

io1

- 1000 MiB/s per volume on Nitro systems (R5, M5, T3)
  - 4000 MiB/s with 200+ GiB RDS Storage (largest EC2 instance can do 1,750 MiB/s)
- 500 MiB/s per volume on older systems

RDS Storage

Performance Implication of 'Scale Storage' operation

- Uses EBS Elastic Volumes (unless transitioning across 200 GB threshold)
  - New storage is immediately available
  - Cannot reduce amount of storage
- Limit of one scale storage operation every 6 hours
  - RDS Instance will be in "storage optimizing" state

RDS Performance – Restores

Hydration of EBS volume and ”first-touch penalty”

- EBS Backups are to S3 (Simple Storage Service)
  - S3 provides 11-9’s of durability (99.999999999%) - that’s a lot, however…
  - S3 latency is not great
- EBS volume restores optimize for making the volume available quickly
- EBS blocks are filled-in / hydrated from S3 in the background, or when first accessed
- To force hydration to happen quicker, you can scan all the blocks

```
BEGIN rdsadmin.rdsadmin_rman_util.validate_database (
  p_validation_type => 'PHYSICAL+LOGICAL’,
  p_parallel => 4,
  p_section_size_mb  => 10,
  p_rman_to_dbms_output  => FALSE);
END;
/
```

RDS Performance – Multi-AZ

Performance implication of RDS Multi-AZ

- Read latencies are the same as Single-AZ
- Write latencies are increased compared to Single-AZ
  - Synchronous writes to second AZ
  - Round trip to second AZ (1-2ms)
  - Write to secondary EBS volume from secondary RDS host (sub-ms)
- May see a latency impact when enabling Multi-AZ
  - Multi-AZ secondary created via restore
  - Secondary will have first-touch penalty until hydrated

RDS Performance – Parameters

Linux HugePages

- HugePages is enabled by default on newer instance types or anything over 100 GiB Memory
- Can manually enable HugePages on older instances
- RDS automates OS configuration when ‘use_large_pages=ONLY’

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Oracle.html#Oracle.Concepts.HugePages

memory_target = IF({DBInstanceClassHugePagesDefault}, **0**, {DBInstanceClassMemory*3/4})
memory_max_target = IF({DBInstanceClassHugePagesDefault}, **0**, {DBInstanceClassMemory*3/4})
pga_aggregate_target = IF({DBInstanceClassHugePagesDefault}, **{DBInstanceClassMemory*1/8}**, 0)
sga_target = IF({DBInstanceClassHugePagesDefault}, **{DBInstanceClassMemory*3/4}**, 0)
sga_max_size = IF({DBInstanceClassHugePagesDefault}, **{DBInstanceClassMemory*3/4}**, 0)
use_large_pages = {DBInstanceClassHugePagesDefault}

RDS Performance – Migrating to RDS

Optimize your Data Load

- Use a larger instance for your data load and then scale down
- Use PIOPS (io1) and then scale down IOPS or switch to gp2
- Disable backup retention – NOARCHIVELOG mode until after data is loaded
- Create indexes last

Agenda

- Introduction
- High Availability
- Performance

RDS for Oracle High Availability And Performance

Questions?