



SRV310

Amazon RDS: Deep Dive

Richard Waymire

Principal Solutions Architect, AWS

What is Amazon RDS?


Managed relational database service with a choice of six popular engines

Amazon
Aurora



Microsoft SQL Server

ORACLE




Easy to administer

No need to provision infrastructure, install, and maintain DB software



Available & durable

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



Highly scalable


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



Fast & secure

SSD storage and guaranteed provisioned I/O; data encryption at rest and in transit

Why use Amazon RDS?



Lower TCO because we manage critical administrative functions

- Automated hardware provisioning, database setup, patching, and backups
- Get more leverage from your teams
- Focus on the things that differentiate you



Built-in high availability and disaster recovery across multiple data centers

- Enabled with a single API call or click of a button in the console
- Even a small startup can leverage enterprise-level availability, durability, and scalability

Configuring your database instance in Amazon RDS

Which RDS engine should I use?

Commercial

ORACLE



Amazon EBS-based Storage

Open Source



PostgreSQL



Cloud Native

**Amazon
Aurora**

MySQL Compatible
PostgreSQL Compatible

Aurora Storage System

Which instance type should I choose?

T2 Family

- Burstable instances
- 1 vCPU/1 GB RAM > 8 vCPU 32 GB RAM
- Moderate networking performance
- Good for smaller or variable workloads
- Monitor CPU credit metrics in Amazon CloudWatch
- T2.micro is eligible for free tier

M4 Family

- General-purpose instances
- 2 vCPU/8 GiB RAM > 64 vCPU 256 GiB RAM
- High-performance networking
- Good for running CPU intensive workloads (e.g., WordPress)

R4 Family

- Memory-optimized instances
- 2 vCPU/16 GiB RAM > 64 vCPU 488 GiB RAM
- High-performance networking
- Good for query-intensive workloads or high connection counts



Which storage type should I choose?

General purpose (GP2)

- SSD storage
- Maximum of 16 TB!
- Leverages Amazon EBS Elastic Volumes
- IOPS determined by volume size
- Minimum of 100 IOPS (below 33.33GiB)
- Bursts to 3,000 IOPS (applicable below 1.3 TB)
- Baseline of 10,000 IOPS (at 3.3 TB and above)
- Affordable performance

Provisioned IOPS (IO1)

- SSD storage
- Maximum of 16 TB!
- Leverages Amazon EBS Elastic Volumes
- Maximum of 40 K IOPS (20 K on SQL Server)
- Delivers within 10% of the IOPS performance 99.9% of the time
- High performance and consistency

Magnetic

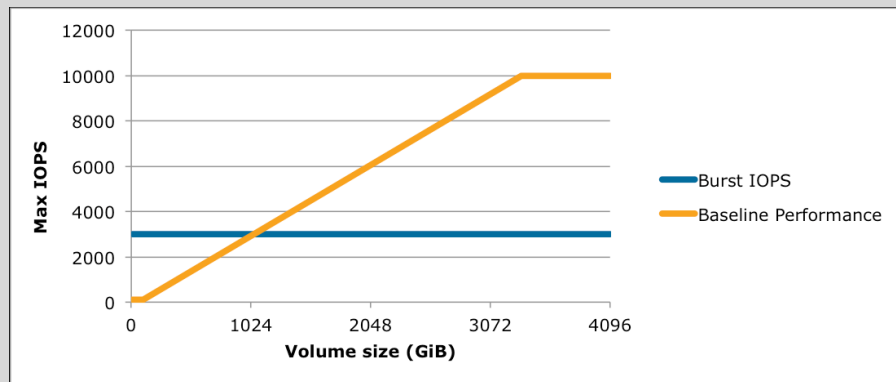
- Magnetic storage
- Maximum of 1 TB
- Supported for legacy databases



How do I decide between GP2 and IO1?


Why am I not seeing 30 K IOPS?

- GP2 is a great choice, but be aware of burst credits on volumes < 1 TB
 - Hitting credit-depletion results in IOPS drop—latency and queue depth metrics will spike until credits are replenished
 - Monitor BurstBalance to see percent of burst-bucket I/O credits available
 - Monitor read/write IOPS to see if average IOPS is greater than the baseline
- Think of GP2 burst rate and PIOPS stated rate as maximum I/O rates



How do I scale my database instance?

Why am I not seeing 30 K IOPS?



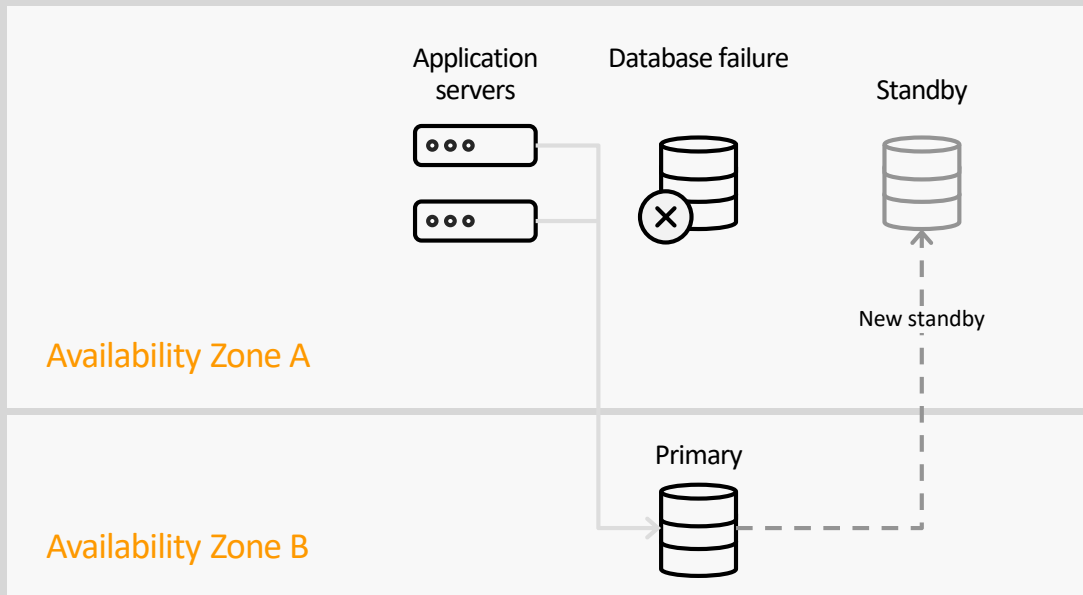
- Scale compute/memory vertically up or down
 - Handle higher load to grow over time
 - Lower usage to control costs
 - New host is attached to existing storage with minimal downtime
- Scale up Amazon EBS storage (now up to 16 TB!)
 - Amazon EBS engines now support Elastic Volumes for fast scaling (now including SQL Server)
 - No downtime for storage scaling
 - Initial scaling operation may take longer, because storage is reconfigured on older instances
 - Can reprovision IOPS on the fly

Managing high availability, read replicas, and backups in Amazon RDS

How do I ensure database high availability?

Multi-AZ provides enterprise-grade fault-tolerance across multiple data centers

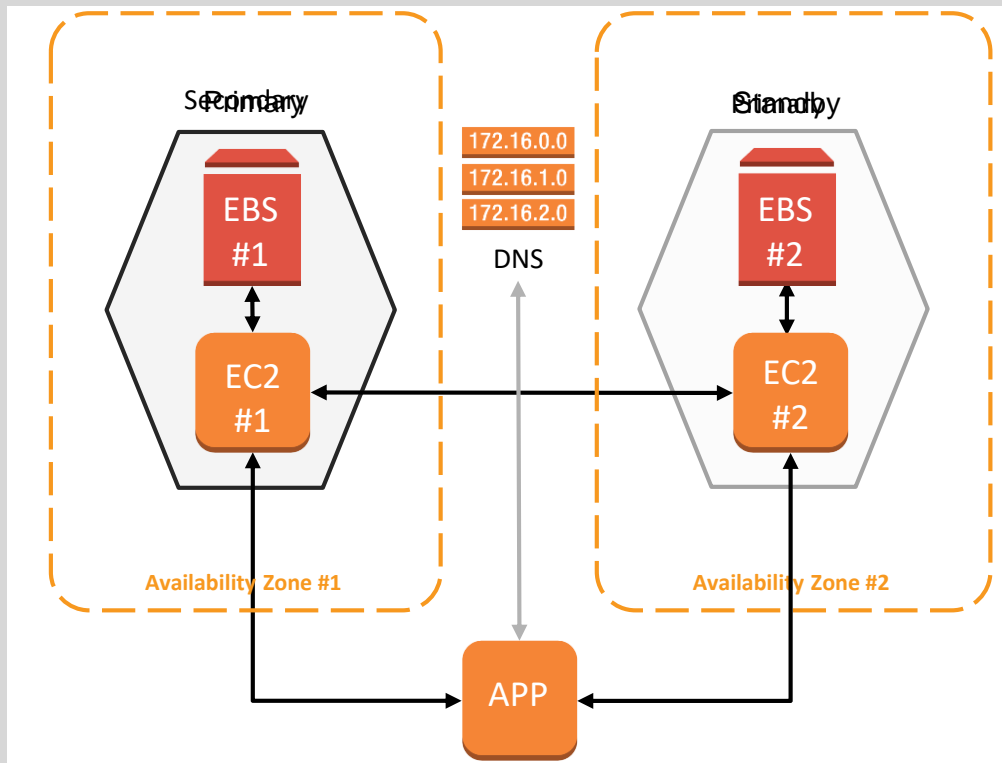
- Automatic failover
- Synchronous replication
- Enabled with one click



What happens during a Multi-AZ failover?

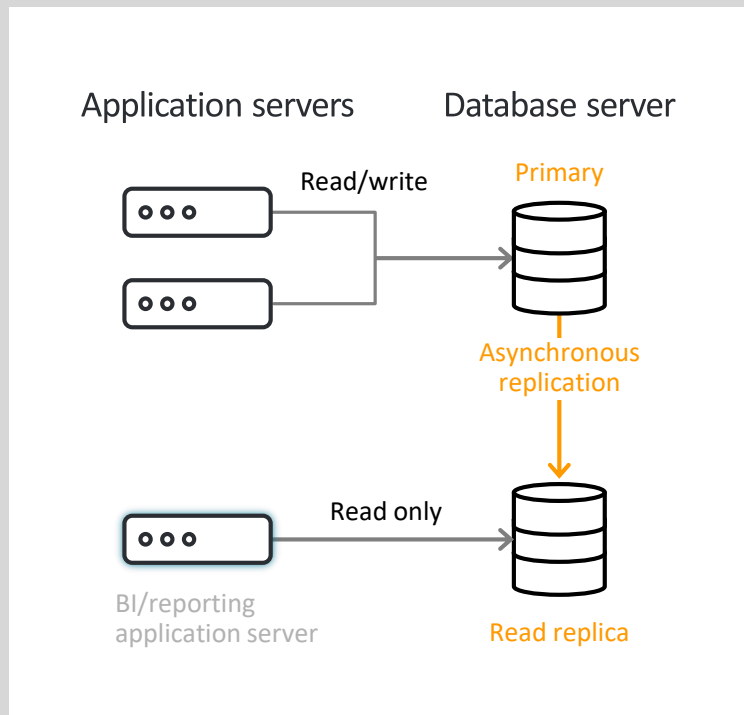
How long does it take?

- Each host manages set of Amazon EBS volumes with a full copy of the data
- Instances are monitored by an external observer to maintain consensus over quorum
- Failover initiated by automation or through the Amazon RDS API
- Redirection to the new primary instance is provided through DNS



Why would I use Read Replicas?

- Relieve pressure on your source database with additional read capacity
- Bring data close to your applications in different regions
- Promote a Read Replica to a master for faster recovery in the event of disaster
- Upgrade a Read Replica to a new engine version
- Supported for MySQL, MariaDB, and PostgreSQL



When should I use Multi-AZ as opposed to Read Replicas?

Multi-AZ

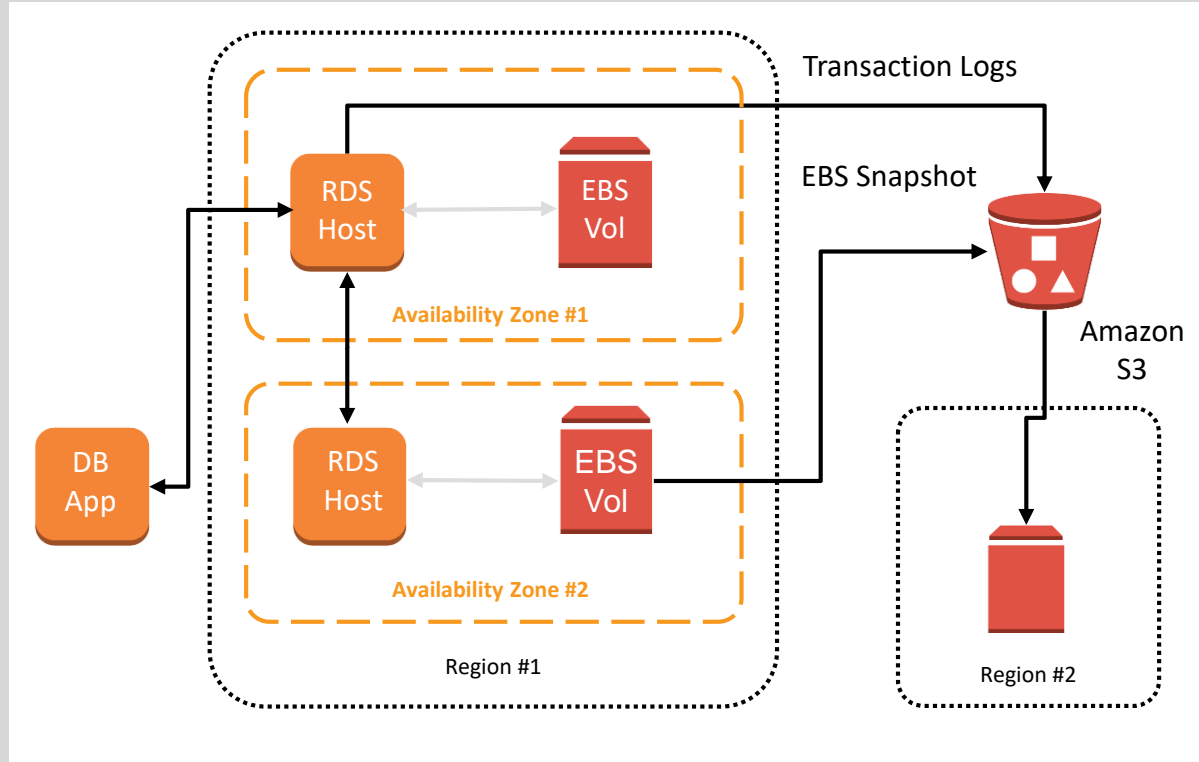
- Synchronous replication – highly durable
- Only primary instance is active at any point in time
- Backups can be taken from secondary
- Always in two Availability Zones within a Region
- Database engine version upgrades happen on primary
- Automatic failover when a problem is detected

Read Replicas

- Asynchronous replication – highly scalable
- All replicas are active and can be used for read scaling
- No backups configured by default
- Can be within an Availability Zone, cross-AZ, or cross-region
- Database engine version upgrades independently from source instance
- Can be manually promoted to a standalone database

How does Amazon RDS manage backups?

- Two options – automated backups and manual snapshots
- Amazon RDS backups leverage Amazon EBS snapshots stored in Amazon S3
- Transaction logs are stored every 5 minutes in Amazon S3 to support point-in-time recovery (PITR)
- No performance penalty for backups
- Snapshots can be copied across regions or shared with other accounts



When should I use automated backups as opposed to snapshots?

Automated backups

- Specify backup retention window per instance (7-day default)
- Kept until outside of window (35-day maximum) or instance is deleted
- Supports PITR
- Good for disaster recovery

Manual snapshots

- Manually created through AWS console, AWS CLI, or Amazon RDS API
- Kept until you delete them
- Restores to saved snapshot
- Use for checkpoint before making large changes, non-production/test environments, final copy before deleting a database

How do I restore a backup?

Why does it take so long?

- Restoring creates an entirely new database instance
 - Define the instance configuration just like a new instance
 - Will get the default parameter, security, and option groups
- New volumes are hydrated from Amazon S3
 - While the volume is usable immediately, full performance requires the volume to warm up until fully instantiated
 - Migrate to a DB instance class with high I/O capacity
 - Maximize I/O during restore process

Launch DB Instance

You are creating a new DB Instance from a source DB Instance at a specified time. This new DB Instance will have the default DB security group and DB parameter groups.

Restore time

Point in time to restore from

Latest restorable time
November 21, 2017 at 7:24:26 PM UTC-8

Custom
Specify a custom date and time to restore from

Custom Date: Custom Time: : : UTC-8

Restore DB Instance

You are creating a new DB Instance from a source DB Instance at a specified time. This new DB Instance will have the default DB Security Group and DB Parameter Groups. This feature is currently **supported for InnoDB storage engine only** if you are using MySQL, refer to details [here](#).

Instance specifications

DB Engine

Name of the Database Engine

License Model

License type associated with the database engine

DB Instance Class

Contains the compute and memory capacity of the DB Instance.

Multi-AZ Deployment

Specifies if the DB Instance should have a standby deployed in another Availability Zone.


Yes

No

Storage type [info](#)

Securing your Amazon RDS database instance

How do I secure my Amazon RDS database?



- Amazon RDS is designed to be secure by default
- Network isolation with Amazon Virtual Private Cloud (Amazon VPC)
- AWS Identity and Access Management (IAM)-based resource-level permission controls
- Encryption at rest using AWS KMS (all engines) or Oracle/Microsoft TDE
- Use SSL protection for data in transit

What does Amazon VPC provide?

- Places your instance in a private subnet, making it secure from public routes on the Internet
- Database instance IP firewall protection lets you securely control network configuration
- Turn off *Public Accessibility* in DB instance settings to restrict access outside Amazon VPC
- Use ClassicLink to network with non-VPC resources



Routing rules



AWS Direct Connect



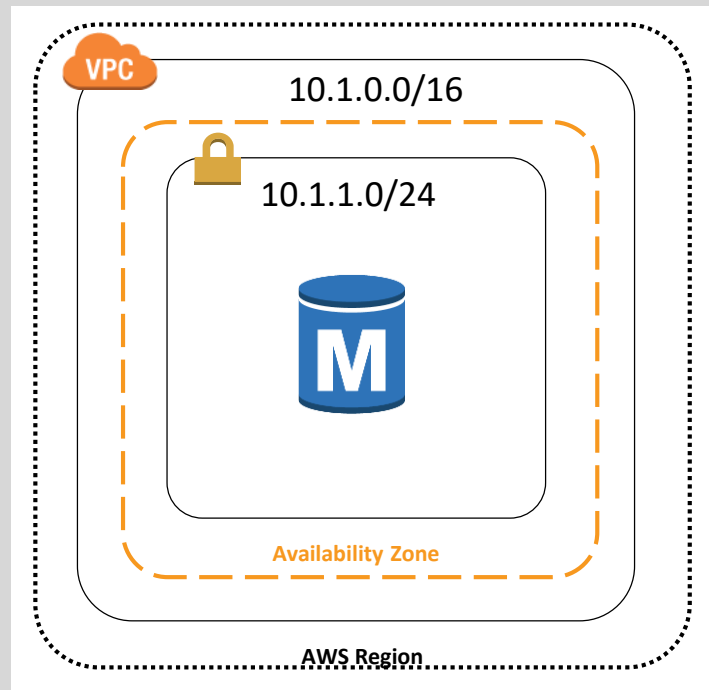
VPN connection



VPC peering




Internet gateway



How do I grant access to my database?

- Use IAM to control who can perform actions on RDS resources
- Do not use AWS root credentials to manage Amazon RDS resources – you should create an IAM user for everyone, including yourself
- Can use AWS Multi-Factor Authentication (MFA) to provide extra level of protection



DBA and Ops

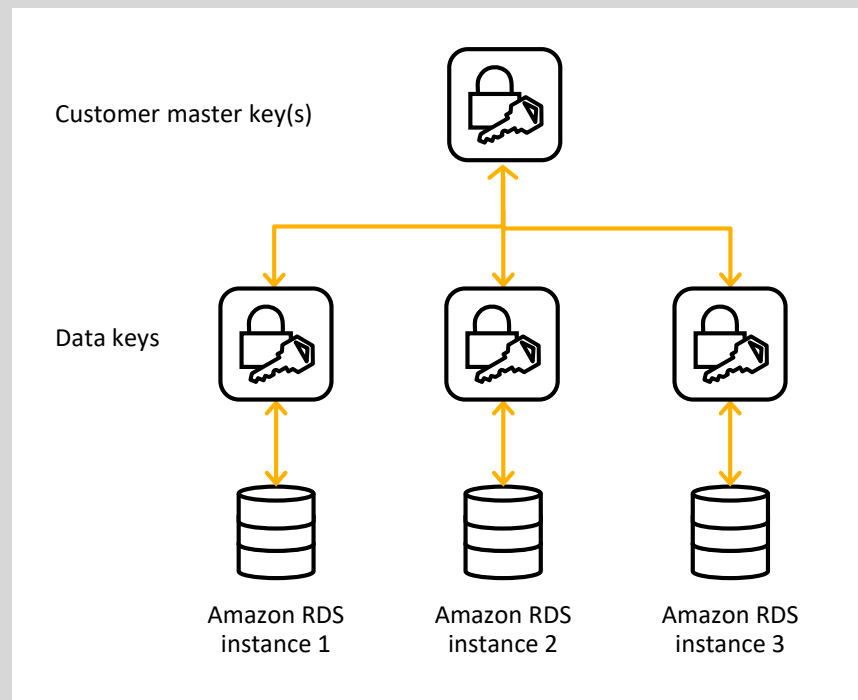


RDS

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowCreateDBInstanceOnly",
      "Effect": "Allow",
      "Action": [
        "rds:CreateDBInstance"
      ],
      "Resource": [
        "arn:aws:rds:*:123456789012:db:test*",
        "arn:aws:rds:*:123456789012:og:default*",
        "arn:aws:rds:*:123456789012:pg:default*",
        "arn:aws:rds:*:123456789012:subgrp:default"
      ],
      "Condition": {
        "StringEquals": {
          "rds:DatabaseEngine": "mysql",
          "rds:DatabaseClass": "db.t2.micro"
        }
      }
    }
  ]
}
```

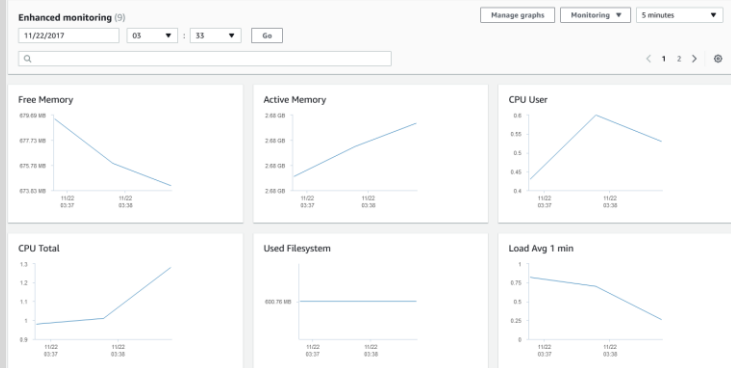
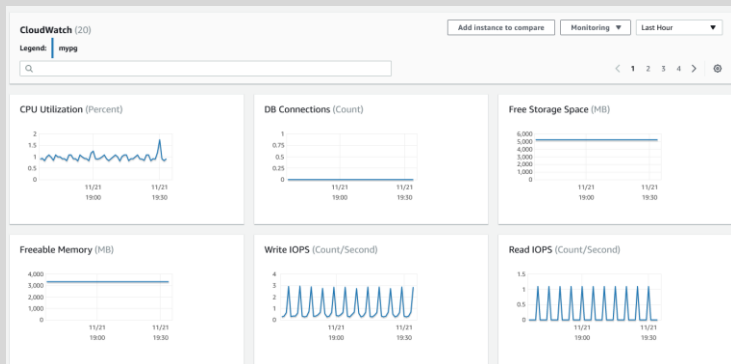
How do I encrypt my database?

- Use AWS KMS-based encryption in the AWS console
- No performance penalty for encrypting data
- Centralized access and audit of key activity
- Best practices
 - Encryption cannot be removed from DB instances
 - If source is encrypted, Read Replicas must be encrypted
 - Add encryption to an unencrypted DB instance by encrypting a snapshot copy



Monitoring your Amazon RDS database instance

How do I monitor my Amazon RDS database?



The screenshot displays the Amazon CloudWatch console for the operating system process list. The legend is set to 'Monitoring'. The table below shows the process list:

NAME	VIRT	RES	CPU%	MEM%
postgres [3213]	1.04 GB	52.74 MB	0	1.33
postgres: rdsadmin rdsadmin localhost(28320) idle [1771]	1.14 GB	8.04 MB	0	0.2
postgres: logger process [3214]	67.42 MB	1.69 MB	0	0.04
postgres: checkpointer process [3216]	1.04 GB	26.22 MB	0	0.66
postgres: writer process [3217]	1.04 GB	9.51 MB	0	0.24

Amazon CloudWatch metrics & alarms

Upload DB logs directly to CloudWatch Logs

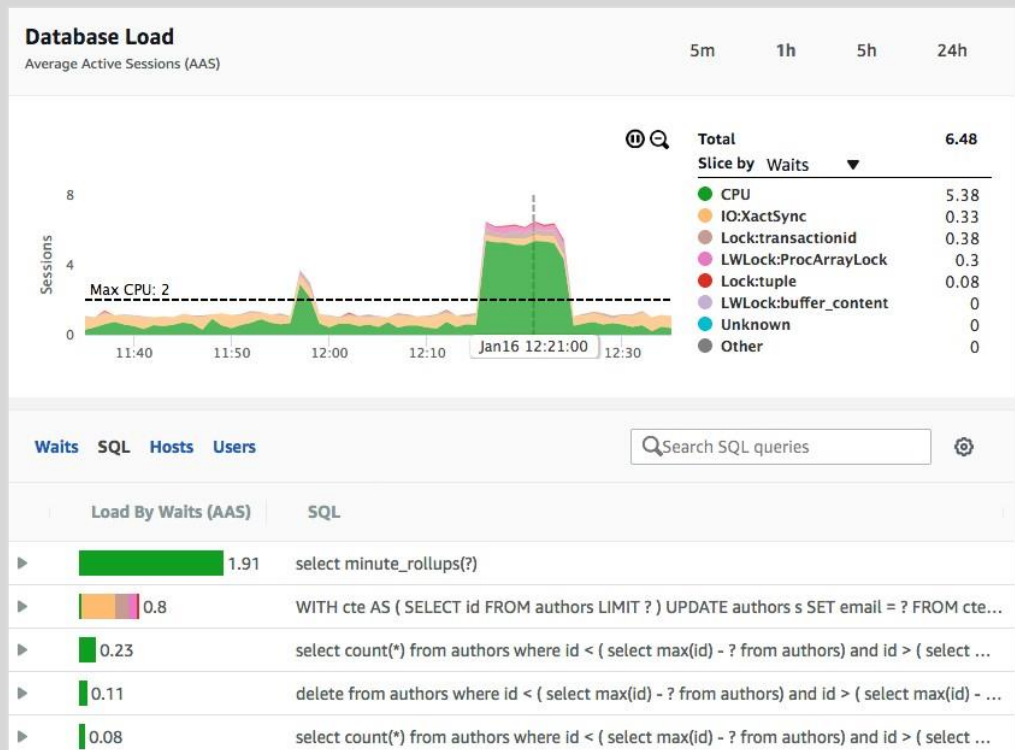
- RDS for MySQL/MariaDB

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 third-party monitoring tools

How do I improve database performance?



- Introducing Amazon RDS Performance Insights
- DB load: Average active sessions
- Identifies database bottlenecks
 - Easy
 - Powerful
 - Top SQL/most intensive queries
- Identifies source of bottlenecks
- Enables problem discovery
- Adjustable timeframe
 - Hour, day, week, and longer
- Available now for Aurora PostgreSQL
- Coming soon for all RDS engines

Can I know when service events happen?

- Amazon RDS uses Amazon SNS to receive notification when an event occurs
- Notifications can be in any form supported by Amazon SNS (email, text message, or call to an HTTP endpoint)
- Six different source types (DB instance, DB parameter group, DB security group, DB snapshot, DB cluster, DB cluster snapshot)
- 17 different event categories (availability, backup, deletion, configuration change, etc.)

The screenshot displays the configuration interface for an Amazon RDS subscription. It is divided into two main sections: 'Target' and 'Source'.

Target Section:

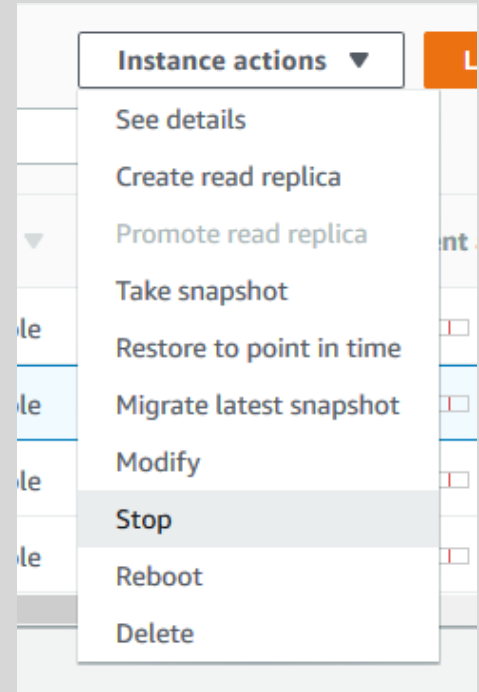
- Send notifications to:** Two radio buttons are present. 'ARN' is selected (indicated by a blue dot), and 'New email topic' is unselected.
- ARN:** A text input field contains 'notifications', followed by a dropdown arrow.

Source Section:

- Source type:** A dropdown menu is set to 'Instances'.
- Instances to include:** Two radio buttons are present. 'Select specific instances' is selected (indicated by a blue dot), and 'All instances' is unselected.
- Specific instances:** A text input field contains 'select instances', followed by a dropdown arrow. Below it, two tags are visible: 'mypg' and 'mysql5619a', each with a close button (X).
- Event categories to include:** Two radio buttons are present. 'Select specific event categories' is selected (indicated by a blue dot), and 'All event categories' is unselected.
- Specific event:** A text input field contains 'select event categories', followed by a dropdown arrow. Below it, two tags are visible: 'backup' and 'maintenance', each with a close button (X).

Can I stop my database when it's not in use?

- Stop and start a running database instance from the console or AWS Command Line Interface (AWS CLI)
- Available for Single-AZ DB instances
- While instance is stopped, you only pay for storage
- Backup retention window is maintained while stopped
- Instances are restarted after 7 days
 - Pending maintenance operations are applied
 - Instances can be stopped again if desired



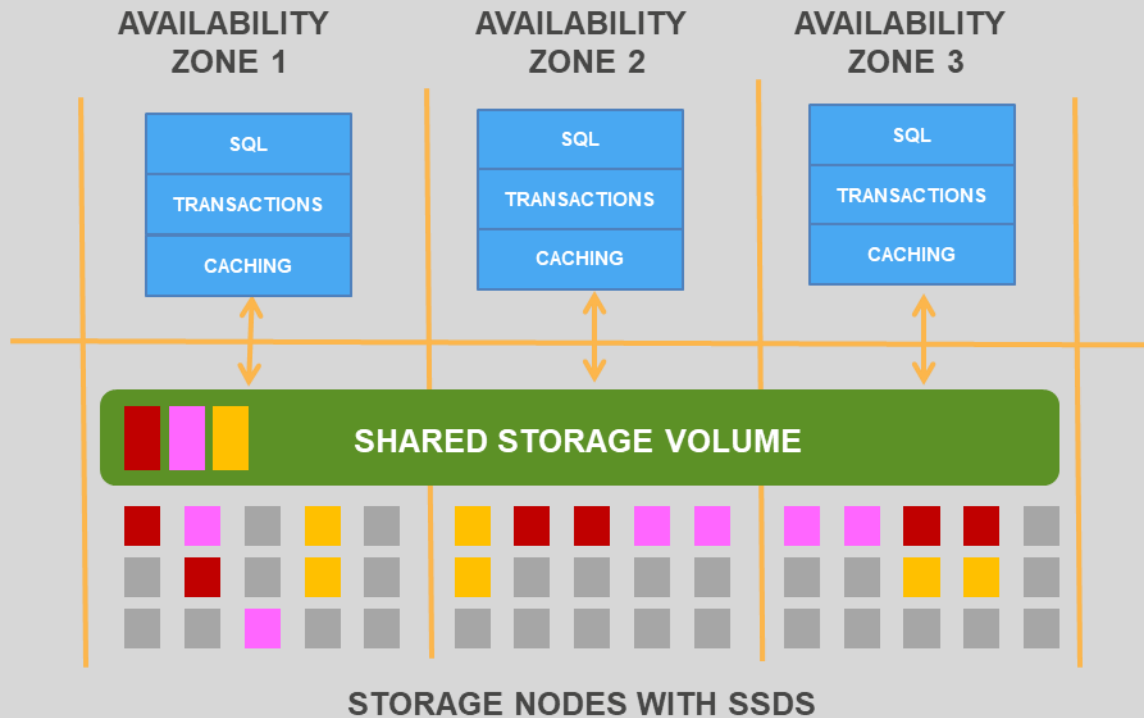
Amazon Aurora

Scale-out, distributed architecture

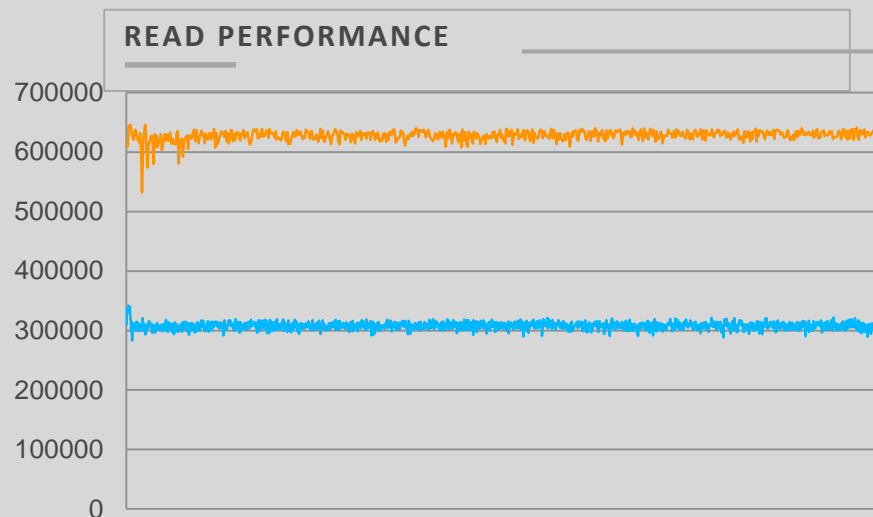
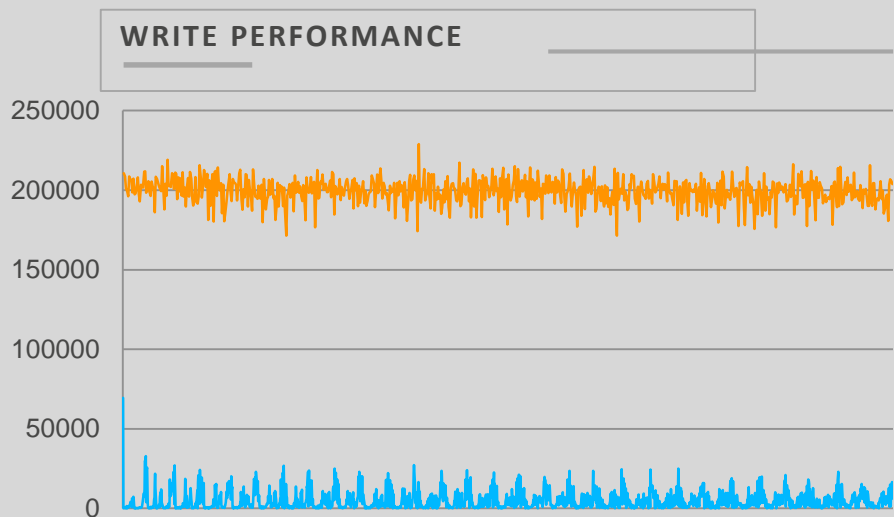
Logging pushed down to a purpose-built log-structured distributed storage system

Storage volume is striped across hundreds of storage nodes distributed across 3 Availability Zones (AZ)

Six copies of data, two copies in each AZ



Aurora performance



MySQL SysBench results; R4.16XL: 64cores / 488 GB RAM

Aurora 

MySQL 5.6 

**Aurora read/write throughput compared to MySQL 5.6
based on industry standard benchmarks.**

How did we achieve this?

DO LESS WORK

Do fewer IOs

Minimize network packets

Cache prior results

Offload the database engine

BE MORE EFFICIENT

Process asynchronously

Reduce latency path

Use lock-free data structures

Batch operations together

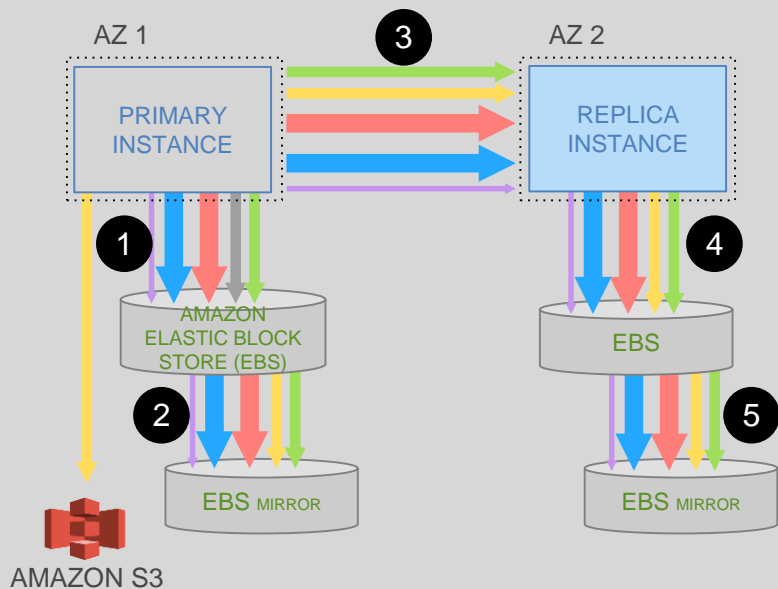
DATABASES ARE ALL ABOUT **I/O**

NETWORK-ATTACHED STORAGE IS ALL ABOUT **PACKETS/SECOND**

HIGH-THROUGHPUT PROCESSING IS ALL ABOUT **CONTEXT SWITCHES**

I/O traffic in MySQL

MYSQL WITH REPLICA



I/O FLOW

Issue write to Amazon EBS – EBS issues to mirror, back when both done

Stage write to standby instance

Issue write to EBS on standby instance

OBSERVATIONS

Steps 1, 3, 4 are sequential and synchronous

This amplifies both latency and jitter

Many types of writes for each user operation

Have to write data blocks twice to avoid torn writes

PERFORMANCE

780K transactions

7,388K I/Os per million txns (excludes mirroring, standby)

Average 7.4 I/Os per transaction

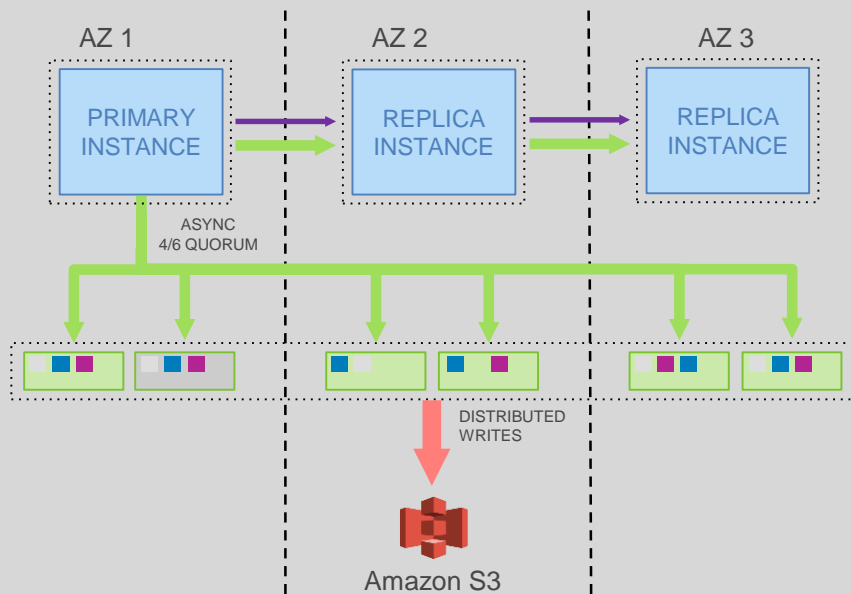
30 minute SysBench writeonly workload, 100GB dataset, **RDS MultiAZ**, 30K PIOPS

TYPE OF WRITE



I/O traffic in Amazon Aurora

AMAZON AURORA



I/O FLOW

Boxcar redo log records – fully ordered by LSN
Shuffle to appropriate segments – partially ordered
Boxcar to storage nodes and issue writes

OBSERVATIONS

Only write redo log records; all steps asynchronous
No data block writes (checkpoint, cache replacement)
6X more log writes, but **9X less** network traffic
Tolerant of network and storage outlier latency

PERFORMANCE

27,378K transactions **35X MORE**
950K I/Os per 1M txns (6X amplification) **7.7X LESS**

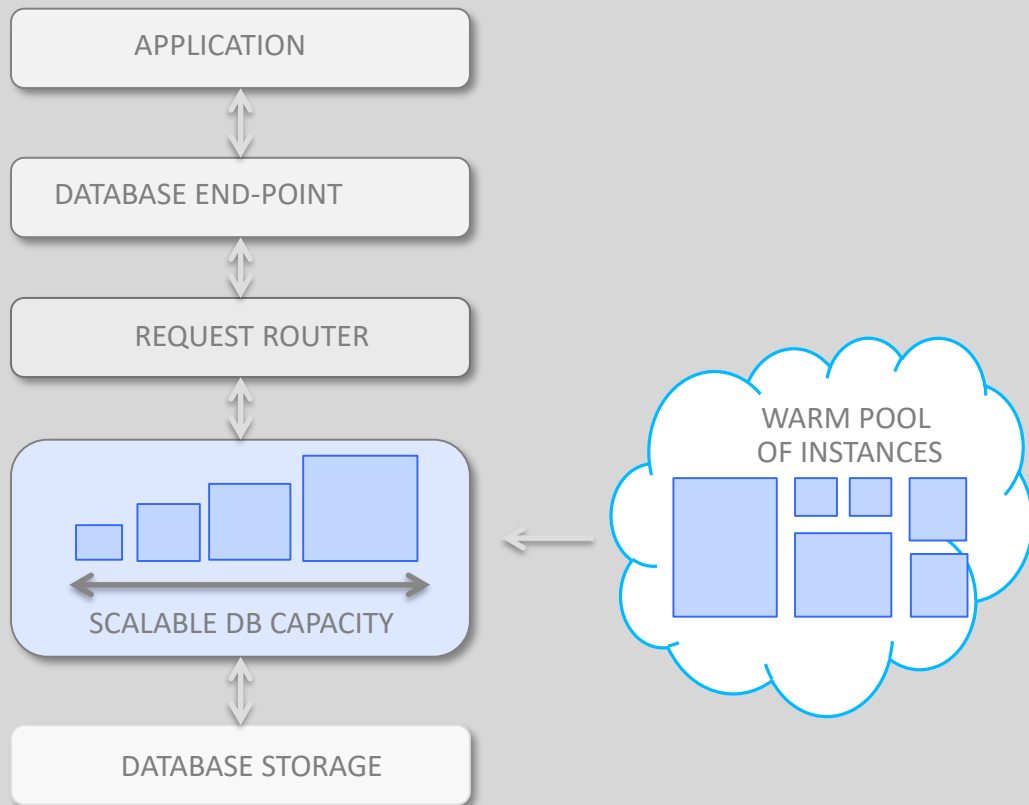
TYPE OF WRITE



Amazon Aurora Serverless

Aurora Serverless

- Starts up on demand, shuts down when not in use
- Scales up/down automatically
- No application impact when scaling
- Pay per second, 1-minute minimum

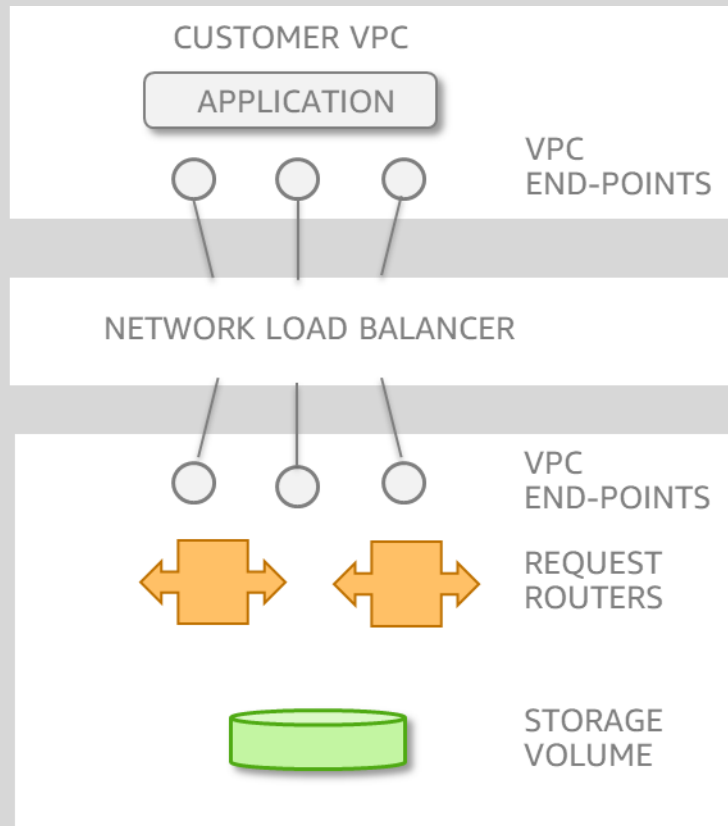


Database end-point provisioning

When you provision a database, Aurora Serverless:

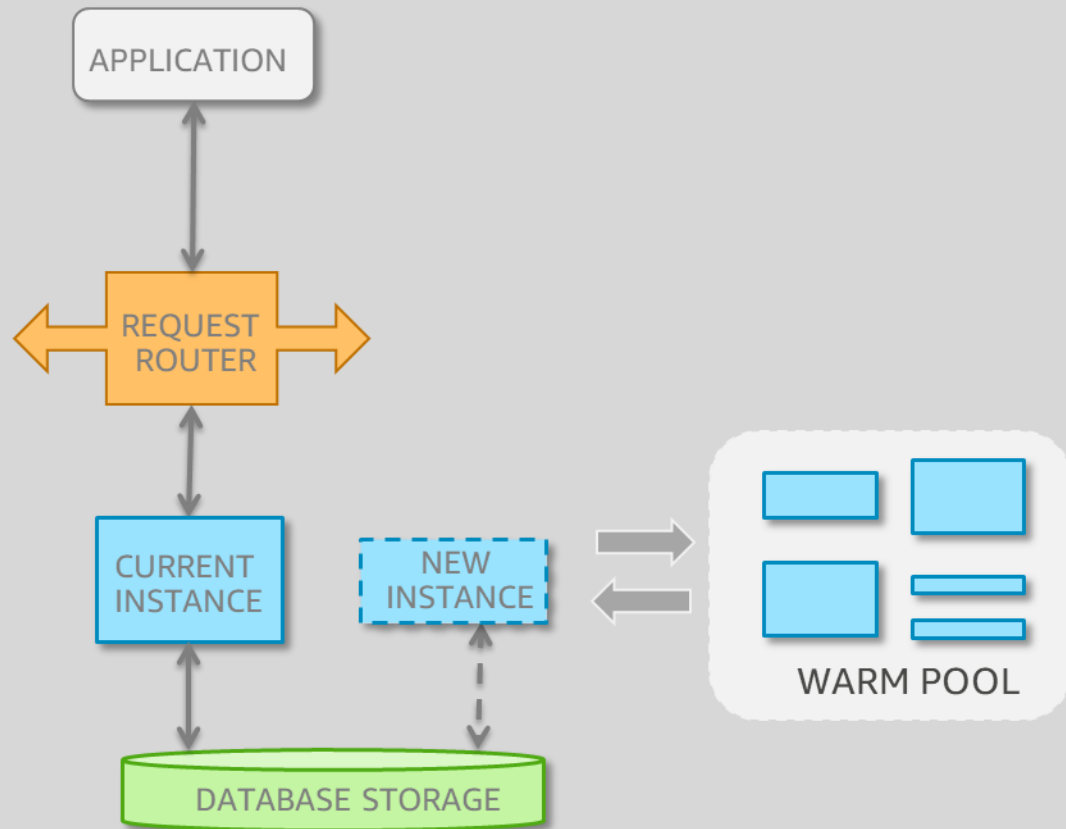
- Provisions VPC end-points for the application connectors
- Initializes request routers to accept connections
- Creates an Aurora storage volume

A database instance is only provisioned when the first request arrives



Instance provisioning and scaling

- First request triggers instance provisioning. Usually **3-5 seconds**
- Instance auto-scales up and down as workload changes. Usually **1-3 seconds**
- Instances hibernate after user-defined period of inactivity
- Scaling operations are transparent to the application – user sessions are not terminated
- Database storage is persisted until explicitly deleted by user



Speaker contact

Richard Waymire
Principal Solutions Architect
waymire@amazon.com

Submit session feedback

1. Tap the **Schedule** icon.
2. Select the session you attended.
3. Tap **Session Evaluation** to submit your feedback.

Thank you!

aws.amazon.com/rds

aws.amazon.com/dms