Relational Database Services (RDS) for Db2

Hans Siebrand

Senior Specialist Solution Architect

hanssieb@amazon.com

Overview of Db2 RDS

Introducing Amazon Relational Database Service (RDS)

RUN FULLY-MANAGED IBM Db2 DATABASES ON AWS

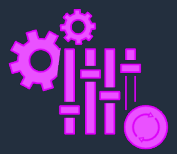

Increase efficiency

Focus on innovation

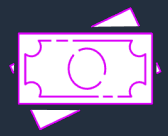

Reduce costs

- Automates time-consuming Db2 undifferentiated tasks, such as provisioning, backups, patching, monitoring and more
- Easily migrate existing IBM Db2 for LUW (Linux, UNIX, and Windows) databases using AWS DMS and/or Db2 native tools
- Launch Db2 v11.5+ databases in minutes and enable high availability with RDS Multi-AZ
- Bring your own IBM software licenses with support for Standard and Advanced Editions or pay hourly for on-demand licenses through AWS Marketplace
- AWS/IBM Partnership to bring best of IBM Db2 on an AWS fully managed platform to meet security, durability, availability and scalability requirements

Amazon Relational Database Service (RDS)

Set up, operate, and scale a relational database in the cloud with just a few clicks

Easy to administer

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

Secure and compliant

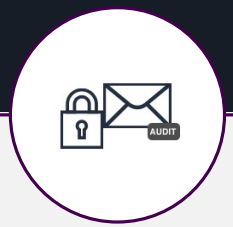

Data encryption at rest and in transit, with industry compliance and assurance programs

Available and durable

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

Performant and scalable

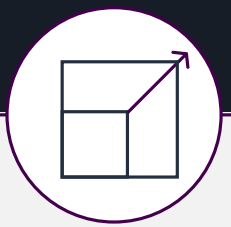

Scale compute and storage with a few clicks, plus minimal downtime for your application

Accelerate path to innovation with managed databases

Spend time innovating and building apps, not managing infrastructure

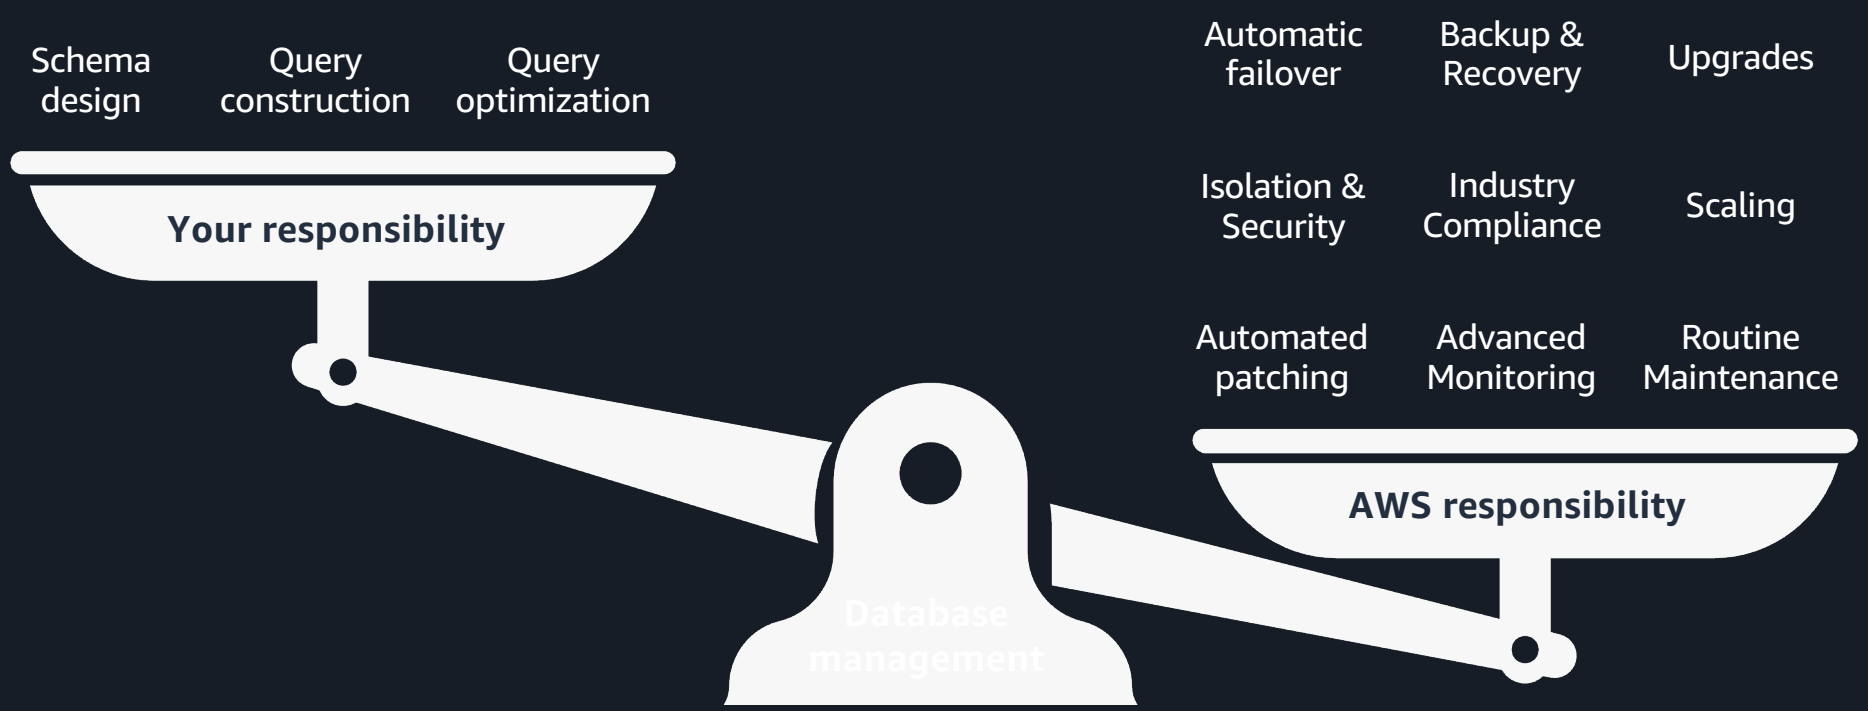

Amazon RDS – fully managed

Spend time innovating and building new apps, not managing infrastructure

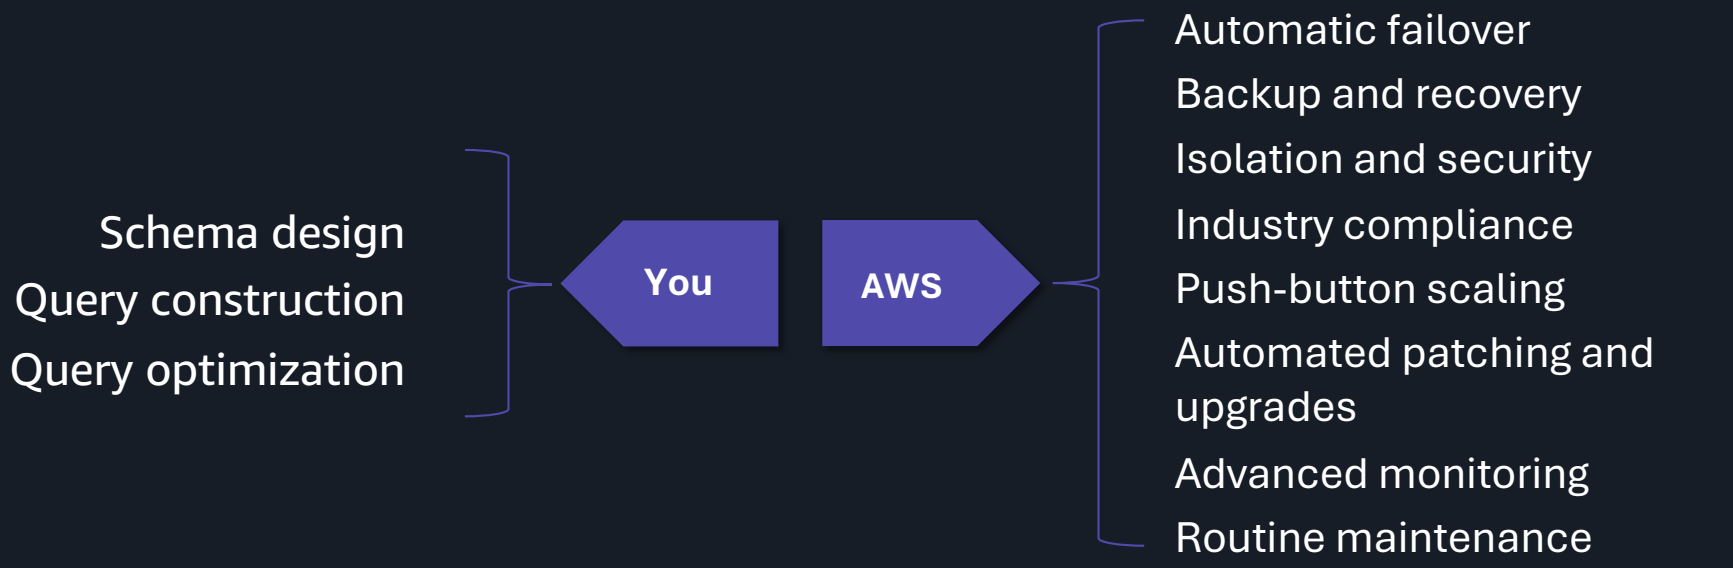

Scale compute and storage with ease

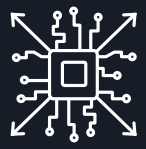

Scale compute to handle increased load

- Up to 192 vCPU and 1536 GiB of RAM per instance
- Scale out with read replicas

Scale storage for larger datasets

- Quickly scale EBS storage up to 64 TiB
- No downtime for storage scaling
- Storage Autoscaling

Scale down to control costs

- As little as 1 vCPU/ 1 GiB of RAM

Database server instance types

Burstable Instances
T family

- Moderate networking performance
- Good for smaller or variable workloads
- 1 vCPU/1 GB RAM > 8 vCPU 32 GB RAM
- T2.micro is eligible for the AWS Free Tier
- T3 will enable unlimited mode—can burst above baseline for extra charge

General Purpose
M family

- High-performance networking
- Good for running CPU-intensive workloads
- 2 vCPU/8 GiB RAM > 128 vCPU 512 GiB RAM

Memory Optimized
R family

- High-performance networking
- Good for query-intensive workloads or high connection counts
- 2 vCPU/16 GiB RAM > 192 vCPU 1536 GiB RAM

High-performance database storage

General purpose (GP2 / GP3)

- SSD storage
- Auto scale up to 64 TiB
- Latency in milliseconds
- IOPS determined by volume size
- Minimum of 100 IOPS (below 33.33 GiB)
- Bursts to 3,000 IOPS (applicable below 1.3 TiB)
- Baseline of 16,000 IOPS per volume (at 5.34 TiB and above)
- Affordable performance

Provisioned IOPS (IO1 / IO2)

- SSD storage
- Auto scale up to 64 TiB
- Single digit millisecond latencies
- Allocate the desired IOPS to volume
- Maximum of 256K IOPS
- Delivers within 10% of the IOPS performance, 99.9% of the time
- High performance and consistency

Amazon RDS Multi-AZ deployments

Enterprise-grade high availability

Fault tolerance across multiple data centers

- Automatic failover
- Synchronous replication
- Enabled with one click
- Failover in 1-2 minutes
- 99.95% monthly uptime SLA

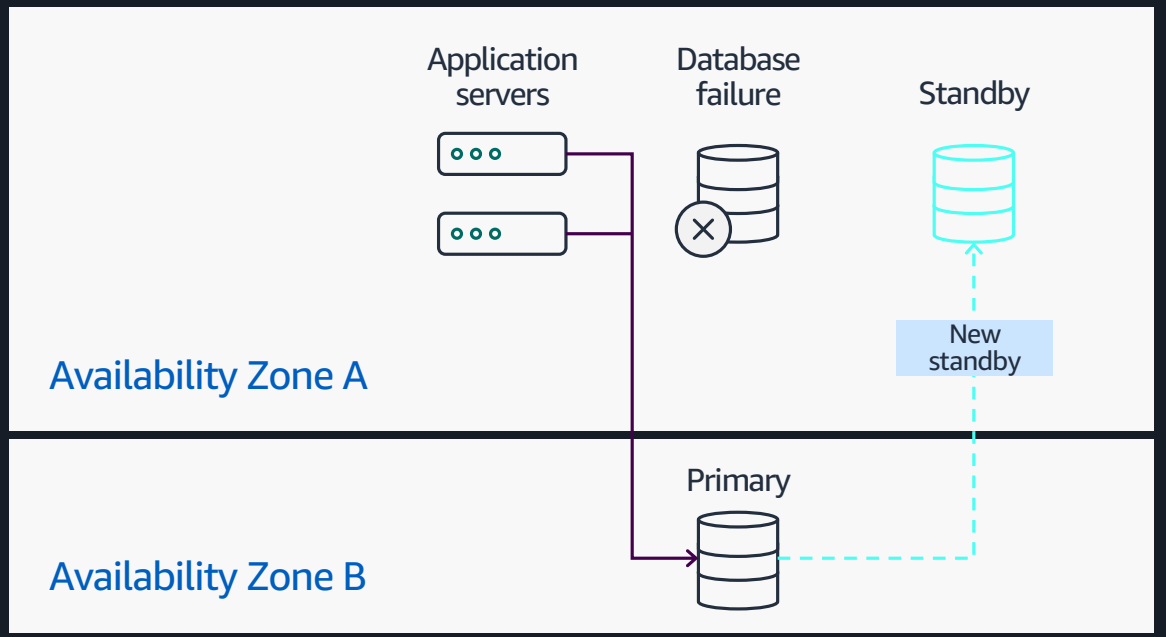

Amazon RDS backups

Automated backups

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

Manual backups

- Take a snapshot any time
- Kept until you delete

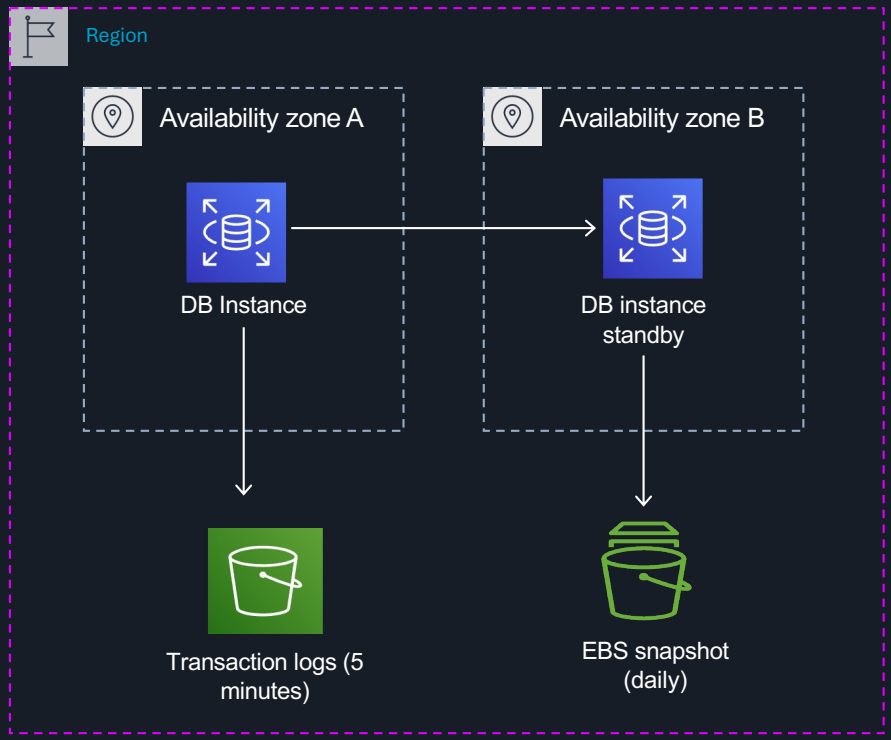

Database snapshots

Backups of your entire DB instance on Amazon S3

- Always incremental
- Amazon S3 → 99.999999999% durability
- Supports encryption
- Copy across accounts and Regions
- Export snapshot in Apache Parquet format

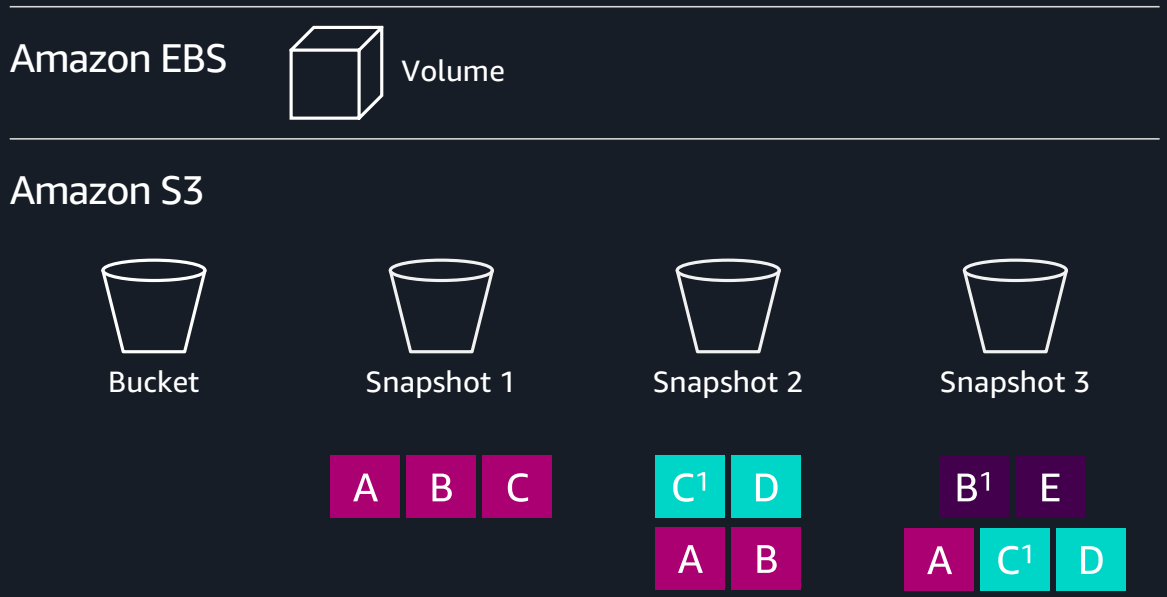

Security and compliance

Network security

- Amazon VPC security groups act as a virtual firewall to control inbound and outbound traffic

Resource access permissions

- AWS Identity and Access Management (IAM) provides resource-level role permission controls

Data encryption

- Encryption at rest using AWS KMS
- SSL protection for data in transit

Compliance and assurance programs for finance, healthcare, government, and more

- HIPAA eligibility under a Business Associate Agreement (BAA) with AWS

RDS for Db2 Integration with Amazon Services

AWS Database Migration Service (AWS DMS)

Amazon CloudWatch

AWS Key Management Service (AWS KMS)

AWS License Manager

Amazon Elastic Block Store (Amazon EBS)

AWS Identity and Access Management (IAM)

AWS CloudFormation

Amazon Virtual Private Cloud (Amazon VPC)

AWS Secrets Manager

AWS Organizations

AWS CloudTrail

New! : Launches in the last 3 months

Recent launches for RDS for Db2

Migrations

- EBCDIC collation support for migrations from z/OS

Manageability

- AWS License Manager integration
- Audit logging
- Up to 5K database users
- Local time zone support
- Native Backups *New!*
- Streaming backups to S3 *New!*
- Reverse log shipping from RDS Db2 *New!*

Availability

- Cross-region automated backups with PITR for DR
- Support managed-AD in opt-in regions
- Cross-region standby replica for DR
- Cross-region read replica for DR

Security

- Group authorization using on-prem AD *New!*
- Encryption support for cross-region backup *New!*

Flexibility

- io2 block express
- M6i, r6i, and x2iedn instances in Asia Pacific (Jakarta) and Asia Pacific (Hyderabad)
- Hourly Db2 license for RDS through AWS Marketplace
- Support for multiple databases
- Support for M7i, R7i instances
- Support for Reserved instances *New!*

New! : Launches in the last 3 months

Upcoming launches for RDS for Db2

Features and new versions

- Db2 v12.1
- Support for Community edition

Availability

- Launch in additional Commercial regions (Auckland, NZ)
- Launch in GovCloud
- FedRAMP High certification

Security

- Self-managed Active Directory integration

Future Direction

- Support for 64TB+ storage
- Optimize CPU for BYOL
- Database Insights integration
- FedRAMP High certification
- Support Db2 v12.1 with Starter edition

Licensing & Pricing with RDS for Db2

Use what you need, pay for what you use

RDS PRICING MODEL

Only pay for what you use

- Flexibility to provision exactly what you need
- Storage can scale automatically
- Scale up and down compute as needed
- RDS decouples storage and compute, and they are charged individually

AWS MARKETPLACE LICENSE

Pay-as-you-go licensing

- Pay an hourly rate for on-demand licenses to help with dev, test, prod or migration needs
- Use this option for seasonal or bursty workloads to avoid the cost of overprovisioning
- Explore disaster recovery testing and database validation exercises
- Get started on RDS for Db2 instantly without an existing license

BRING YOUR OWN LICENSE (BYOL)

Use existing Db2 database licenses

- Db2 VPC based Standard and Advanced edition
- Make use of the Db2's sub-capacity licensing policy in RDS
- Continue to use your active IBM support account, and you contact IBM directly for Db2 service requests
- If you have an AWS Support account with case support, contact AWS for RDS service requests

RDS Db2 now supports Db2's PVU licenses

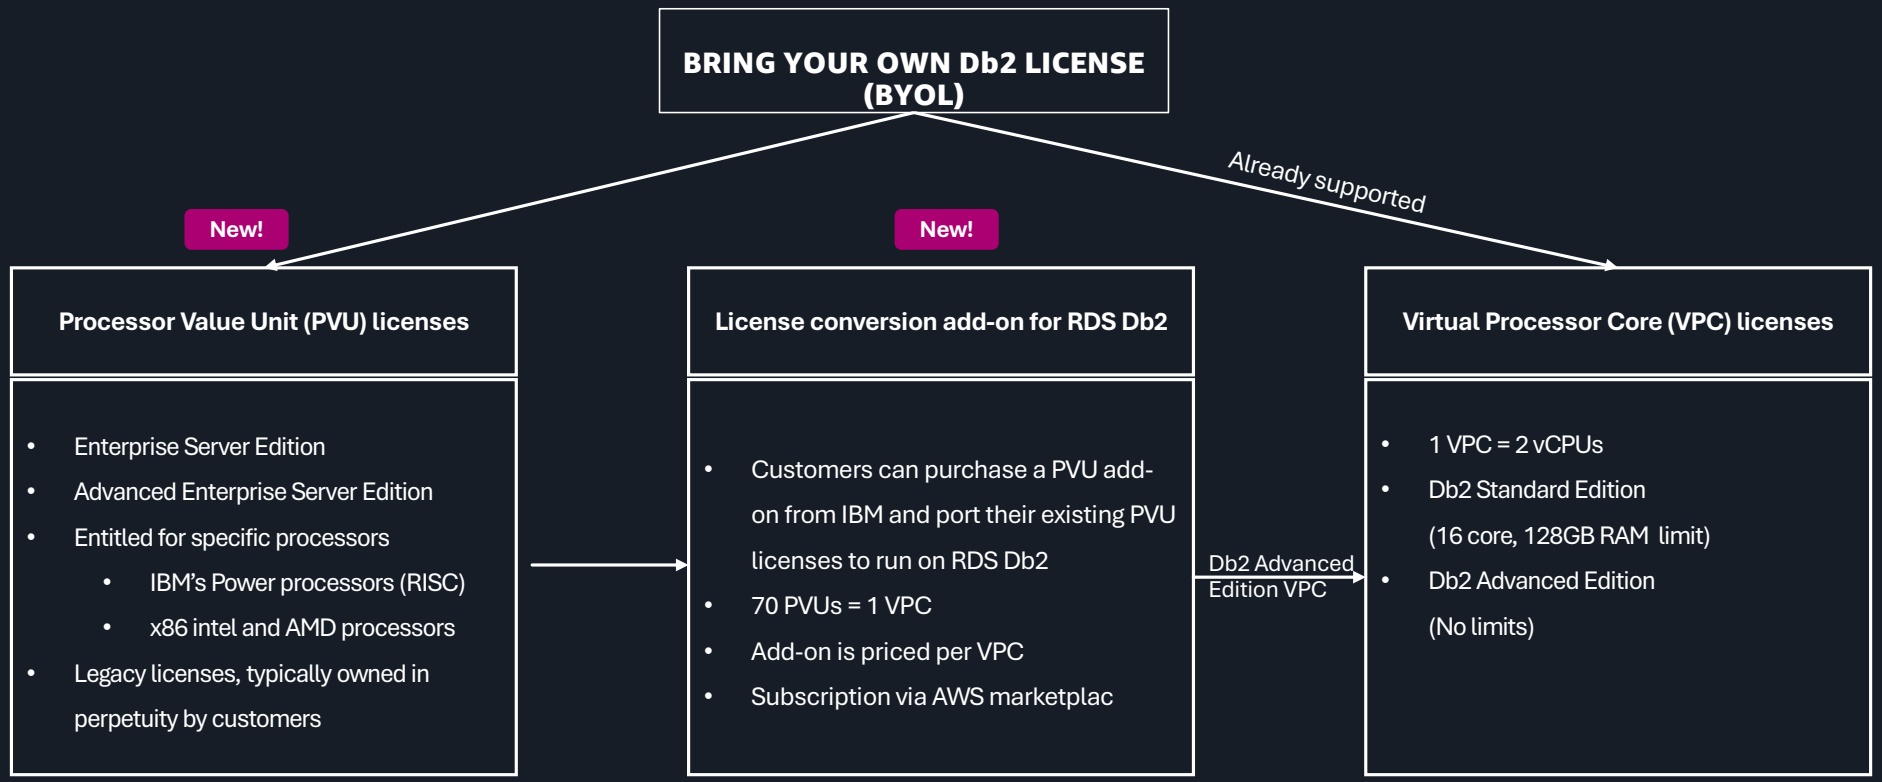

Pricing for Db2 licenses and add-ons are set and managed by IBM. Please contact your IBM Sales rep.

Scaling

Amazon RDS for Db2

RDS Performance Factors

RDS DB Instance Class

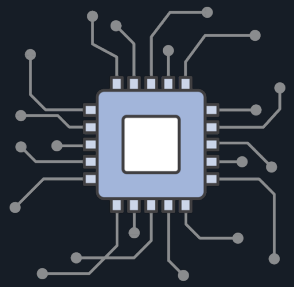

Compute Capabilities

vCPUs

Memory Capabilities

GB of RAM

Network Performance

MB/s (Throughput)

Network attached storage

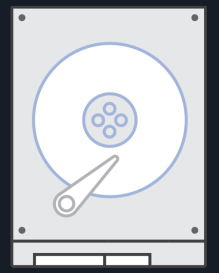

Storage Performance

I/O Throughput

RDS Storage Type

3/20/26

Scale Compute and Storage with Ease & independence

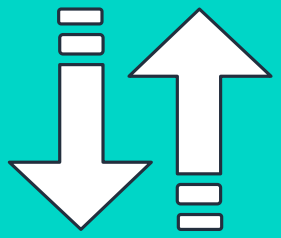

Scale Compute to Handle Increased Load

Up to 192 vCPUs and
Up to 4 TiB of RAM

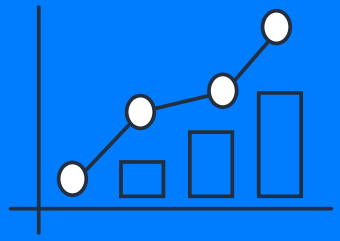

Scale Storage for Larger Data Sets

Scalable EBS storage up to 128 TiB

Scale Down to Control Costs

As little as 2 vCPU and
2 GiB of RAM

Amazon RDS Db2 Storage

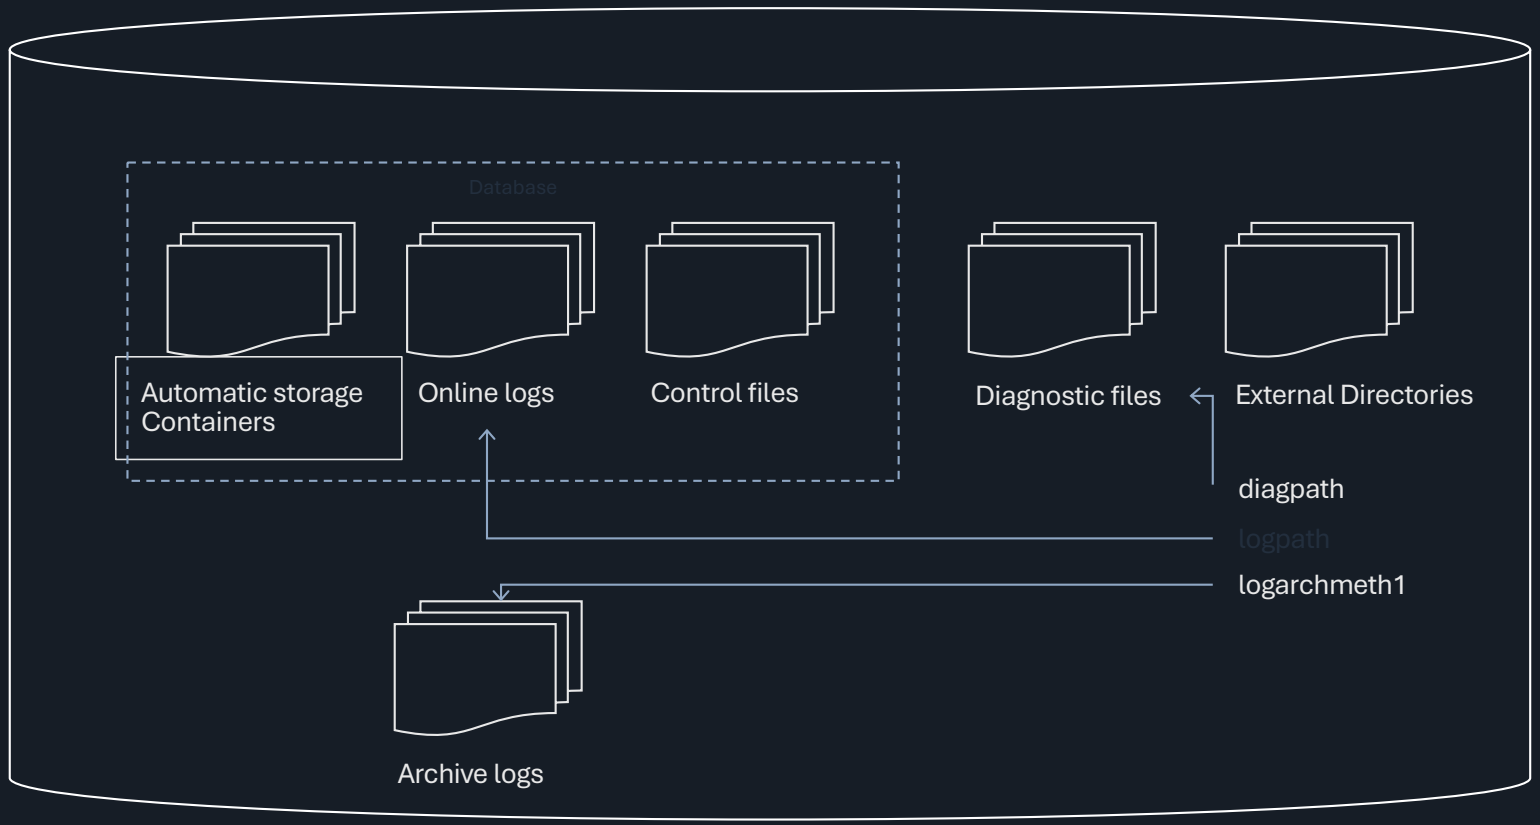

Multi-AZ Secondary

Amazon EBS Snapshot

RDS Storage Scaling

Type	Size	Performance	Maximum throughput
gp3 (General Purpose SSD)	20 GiB - 64 TiB	3,000 - 64,000 IOPS	4,000 MiB/s
io1 (Provisioned IOPS SSD)	100 GiB to 64 TiB	1,000 - 256,000 IOPS*	4,000 MiB/s
io2 Block Express storage (recommended)	100 GiB to 64 TiB	1,000 - 256,000 IOPS*	4,000 MiB/s

- DB instance can be modified to change storage
  - Can modify size (increase size), type, and PIOPs
  - No downtime, performance may degrade during change
- Storage Autoscaling supported

https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/CHAP_Storage.html

RDS Storage Scaling

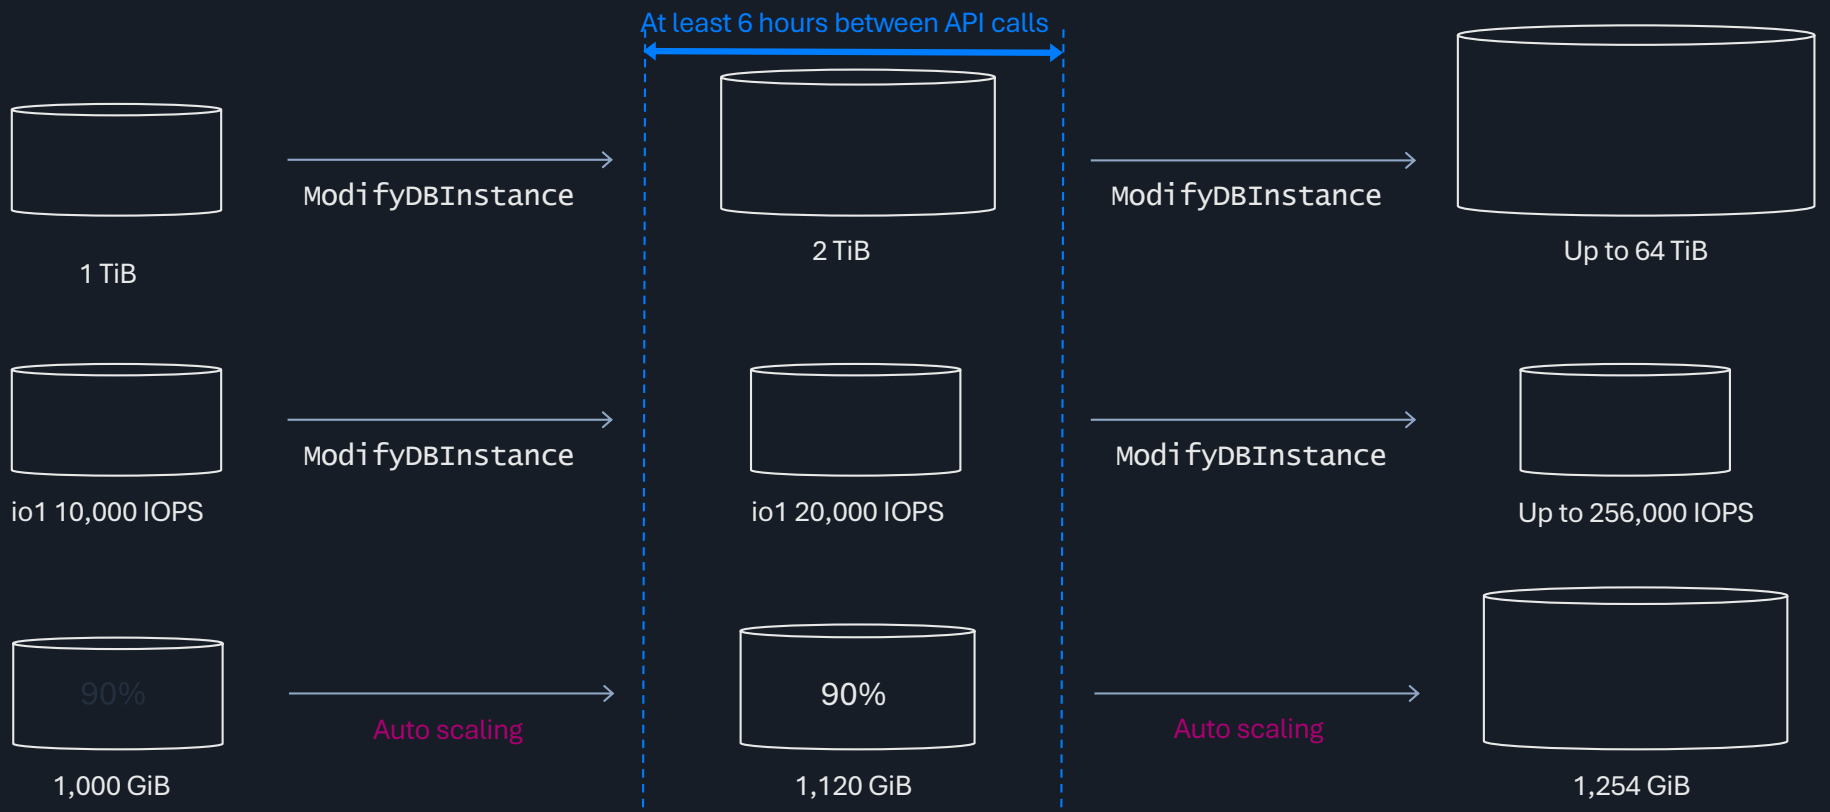

Amazon RDS for Db2 – Building blocks

Compute

Burstable Performance

- T3: 2 vCPU / 2 GB RAM – 8 vCPU 32 GB RAM
- Moderate networking performance
- Built on the AWS Nitro System
- Ideal for database workloads with moderate CPU usage that experience temporary spikes in use
- EBS-optimized by default

General Purpose

- 2 vCPU / 8 GiB RAM - 192 vCPU 768 GiB RAM
- High performance networking
- Built on the AWS Nitro System
- Offers a balance of compute, memory, and networking resources for a broad range of workloads.

Memory Optimized

- 8 vCPU / 256 GiB RAM - 192 vCPU 4096 GiB RAM
- High performance networking
- Built on the AWS Nitro System
- Ideal fit for memory-intensive workloads.

- **Burstable instance types** → Typically steady, with periodic short spikes
- **Scale up and down** → Predictable, steadily increasing, periodic peaks

https://aws.amazon.com/rds/instance-types/

RDS Instance Vertical Scaling

Scale up and down

- Single-AZ total downtime during scaling
- Minimal downtime with Multi-AZ

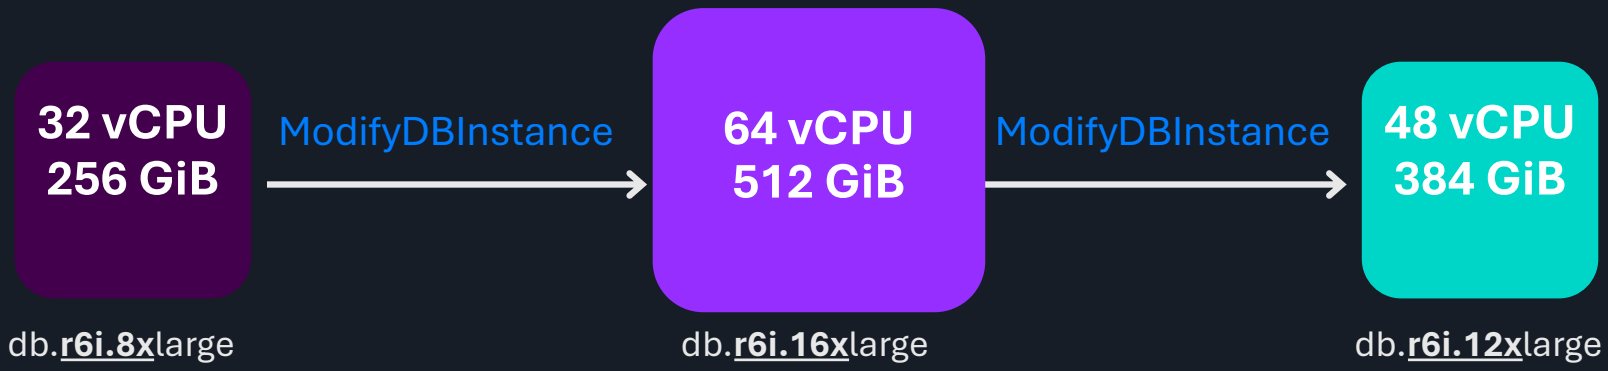

Performance Planning

- All databases including Db2 workloads typically benefit from large amounts of memory (caching)
    - Consider db.r6, db.x2 - Memory Optimized instances for production
- DB instances can be modified to change the DB instance class
    - Requires a reboot (or failover in Multi-AZ)
    - Can scale compute capacity with the workload, if practical
- Use DB Parameter Groups to adjust DB instance configuration
    - Default parameter groups follow AWS & Db2 recommendations
- Amazon RDS **storage throughput** depends on DB instance class

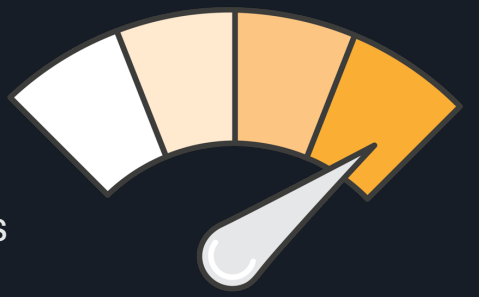

High Availability on RDS for Db2

RDS for Db2 – High Availability (MAZ)

- Amazon RDS Multi-AZ deployments provide enhanced availability and durability for Database (DB) Instances
- When enabled, Amazon RDS
  - automatically creates a primary DB Instance
  - synchronously replicates the data to a standby instance
- In a Multi-AZ setup, secondary instance is launched in a different Availability Zone
- Each AZ runs on its own physically distinct, independent infrastructure, and is engineered to be highly reliable

RDS for Db2 – High Availability (MAZ)

- Amazon RDS automatically performs a failover in the event of any of the following:
  - Loss of availability in primary Availability Zone
  - Loss of network connectivity to primary
  - Compute unit failure on primary
  - Storage failure on primary
  - Detects infrastructure issues, not database engine problems
- 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
- Since the endpoint for your DB Instance remains the same after a failover, your application can resume database operation without the need for manual intervention

High Availability with RDS Multi-AZ (MAZ)

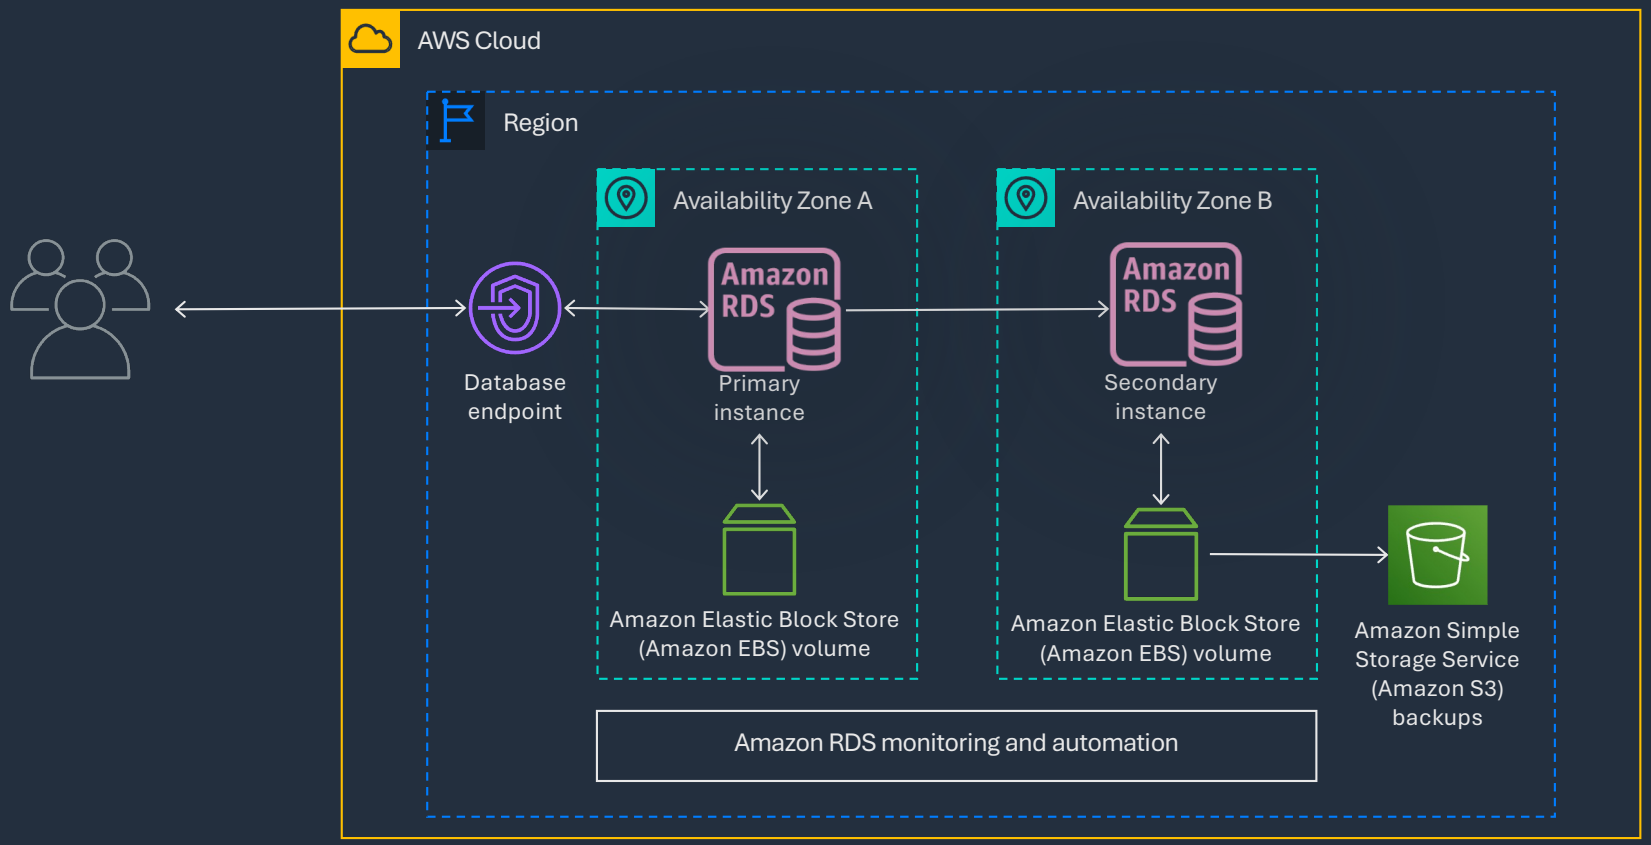

RDS Single-AZ Architecture

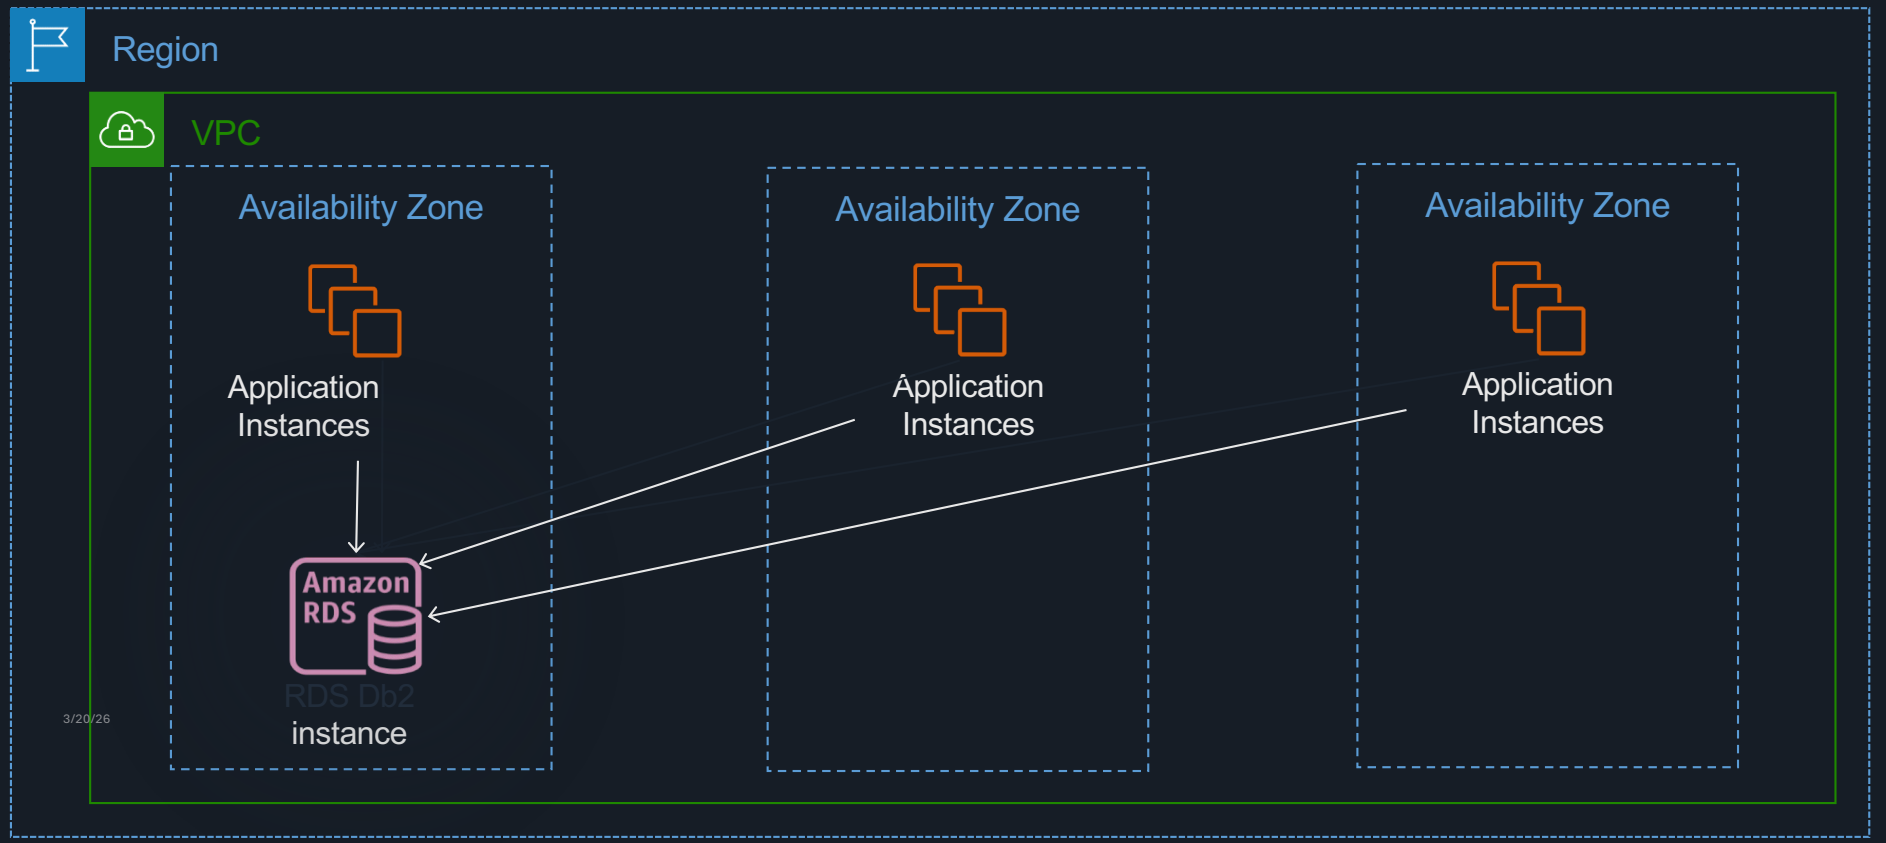

3/20/26

RDS Multi-AZ Architecture

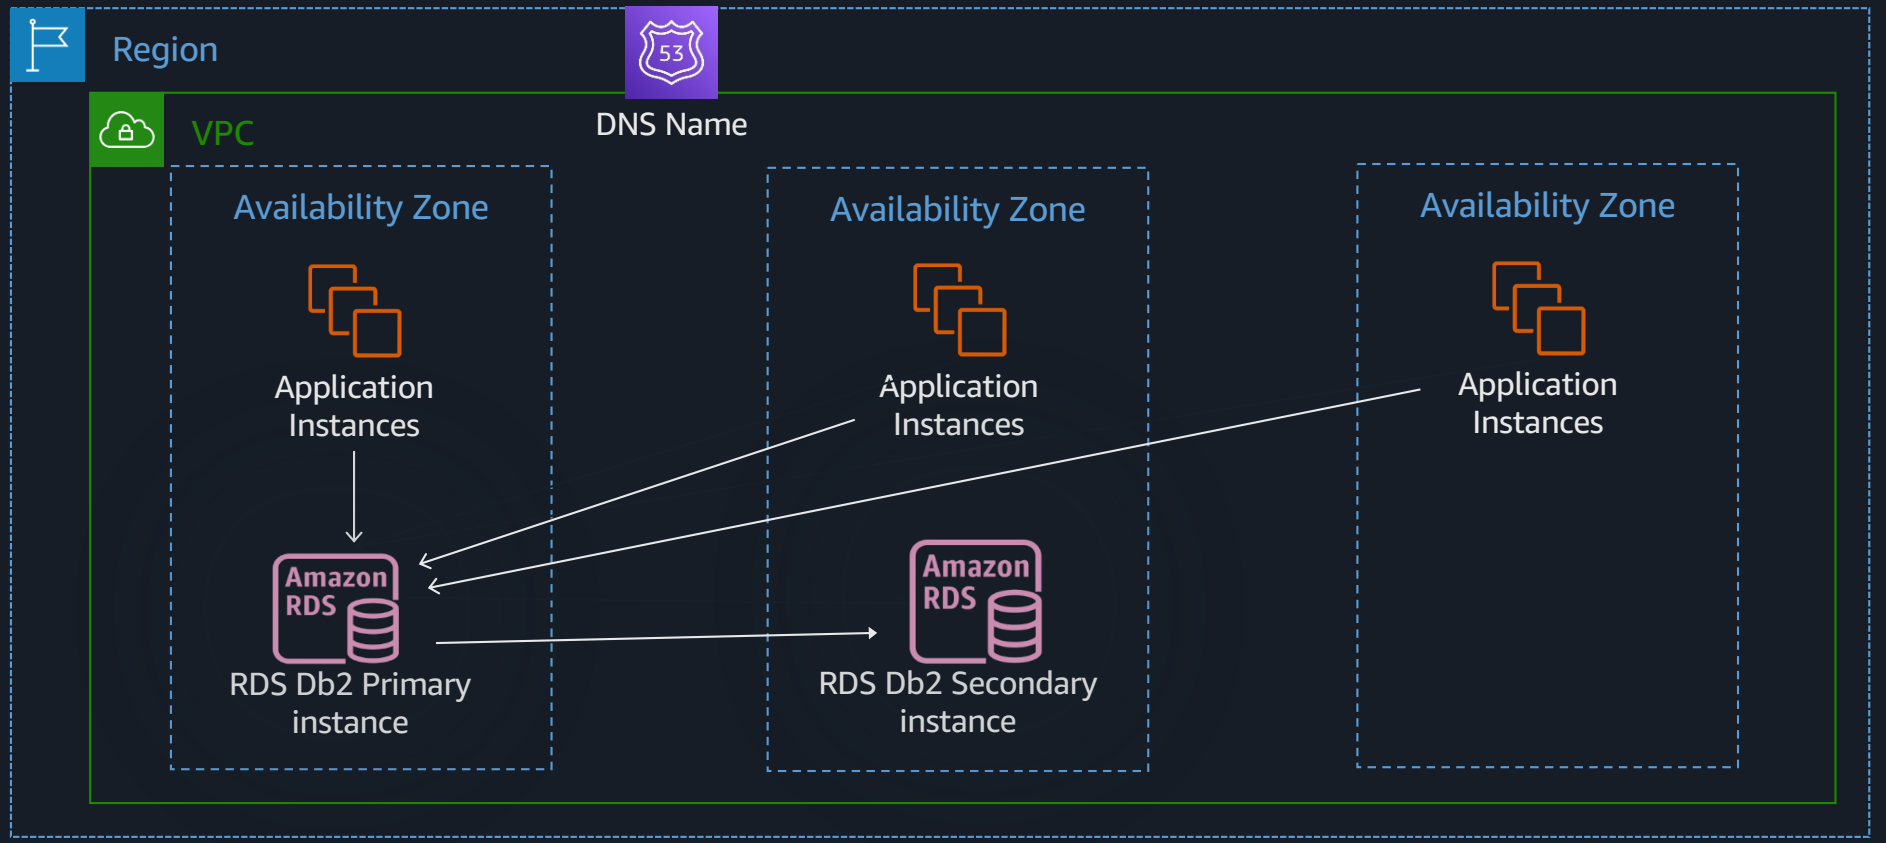

RDS Multi-AZ Failover

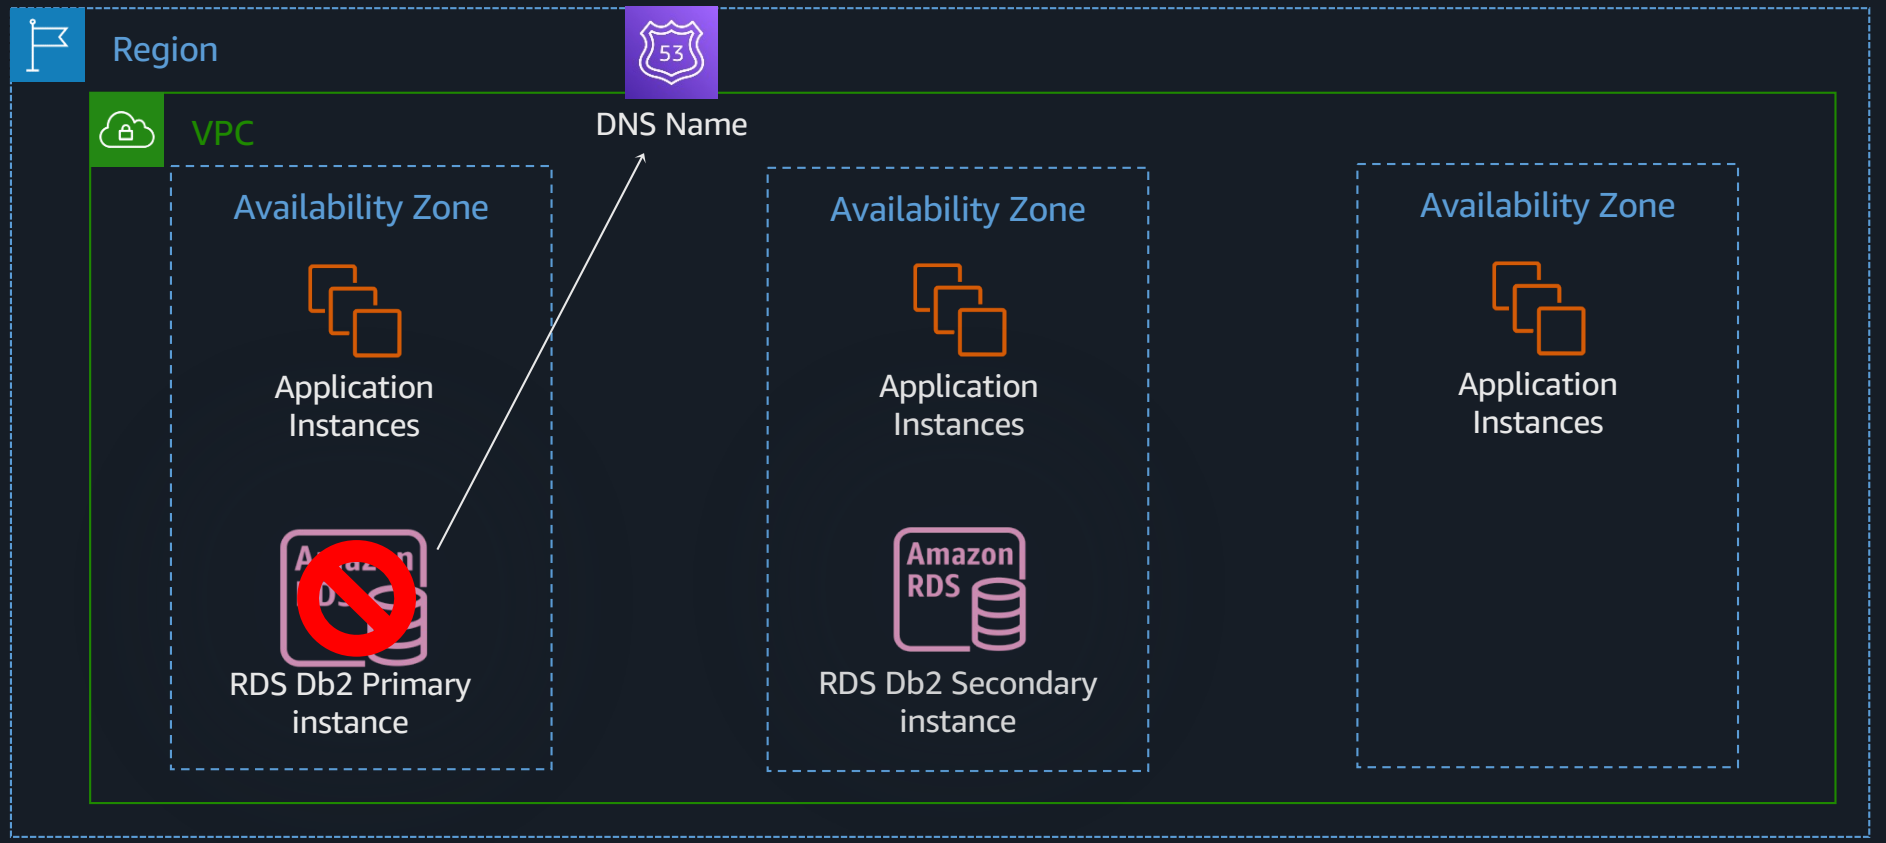

RDS Multi-AZ Failover

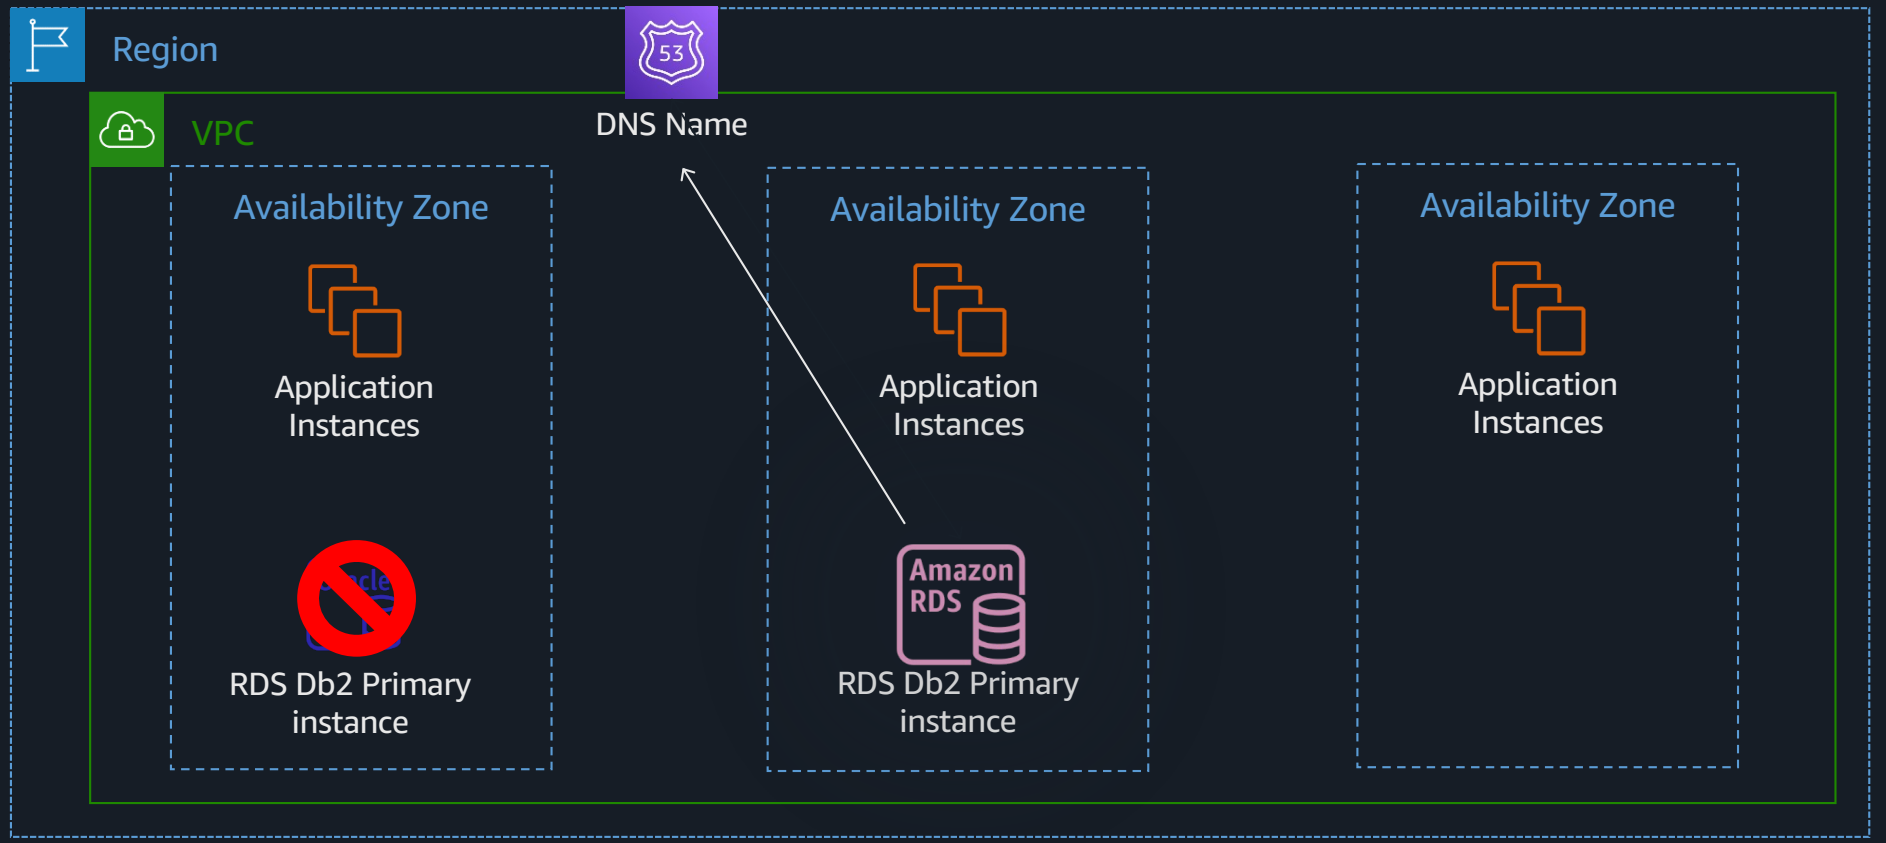

RDS Multi-AZ Failover

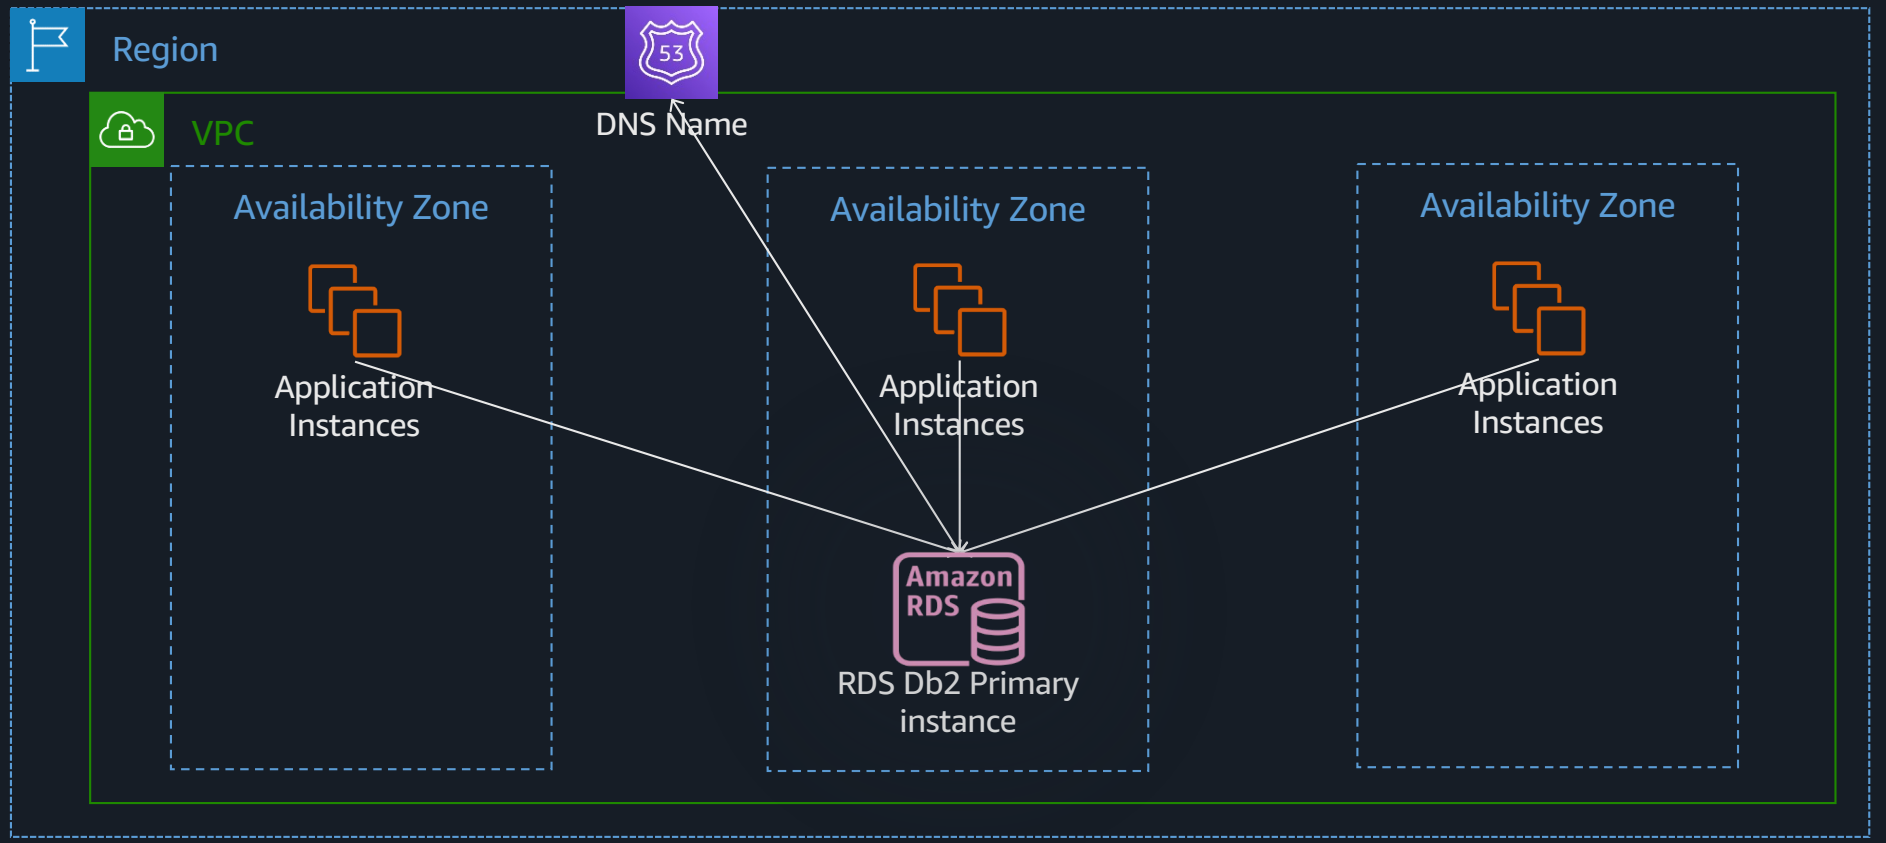

RDS Multi-AZ Failover

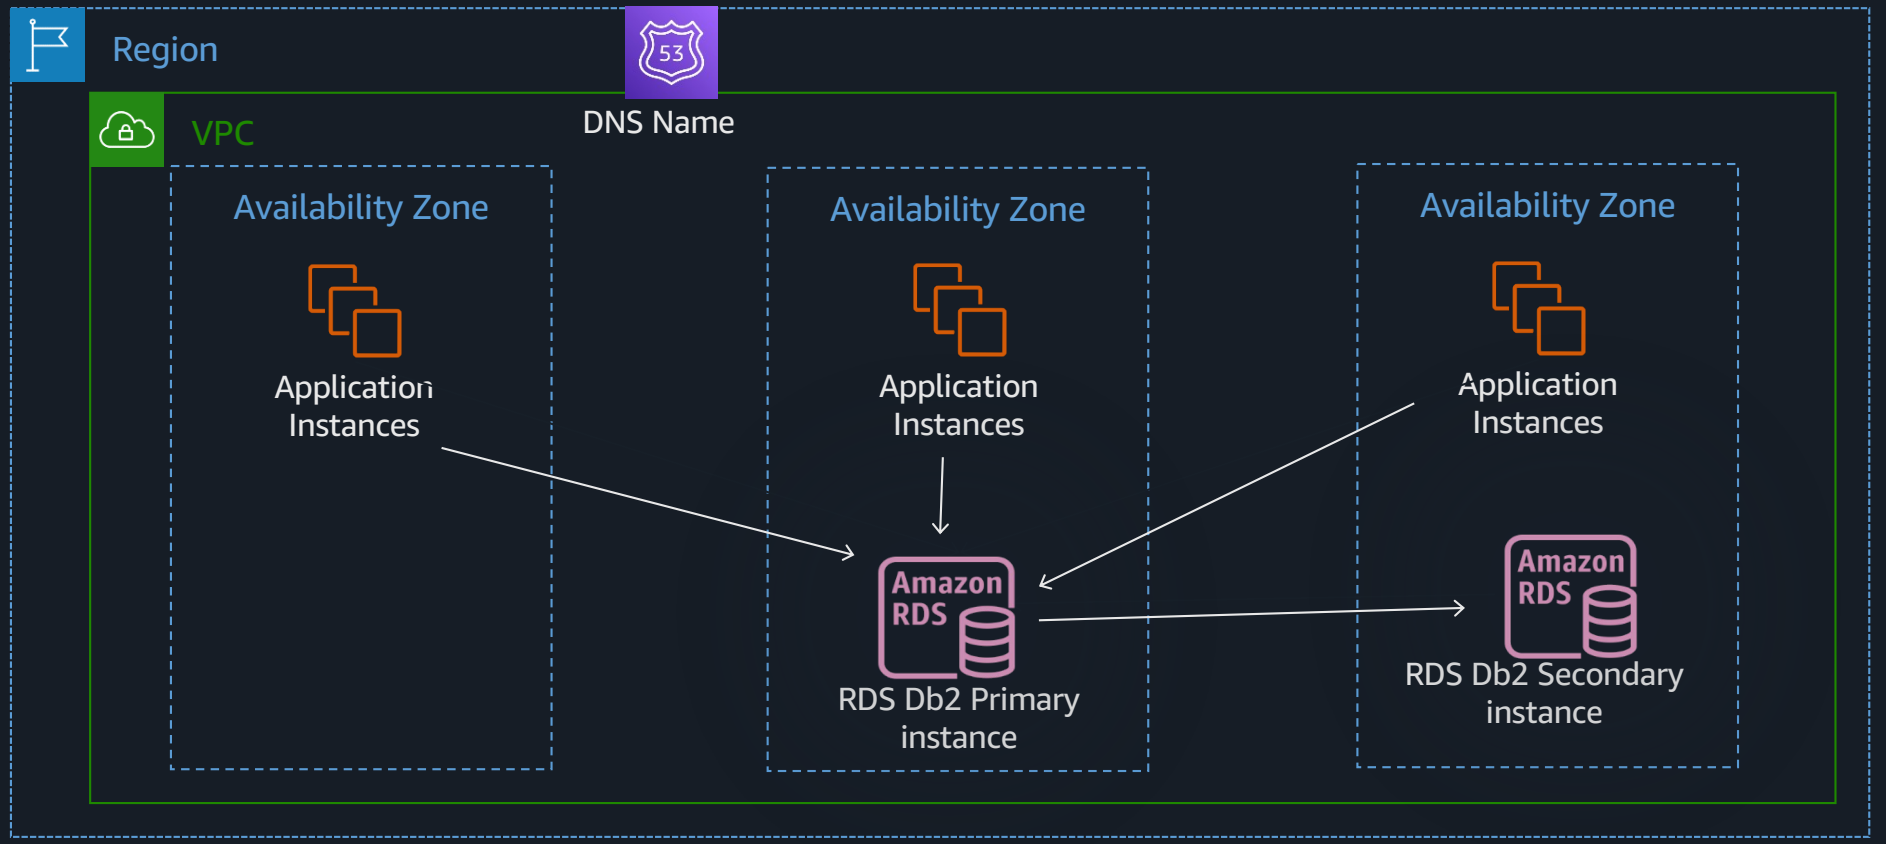

Multi-AZ failover timeline

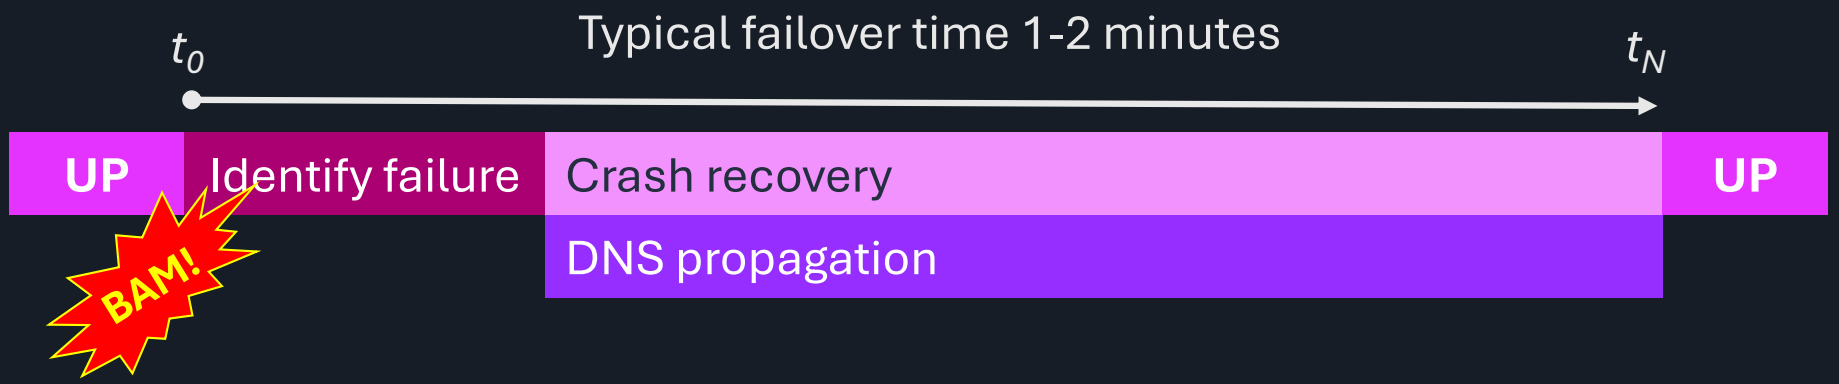

Backup and Restore

Amazon RDS for Db2 – Working with Backups

- Automated backups - Enabled with backup retention period to a positive non-zero value
- Automated backups occur daily during the preferred backup window
- If you don't specify a preferred backup window when you create the DB instance, Amazon RDS a ssigns a default 30-minute backup window
- This window is selected at random from an 8-hour block of time for each AWS Region

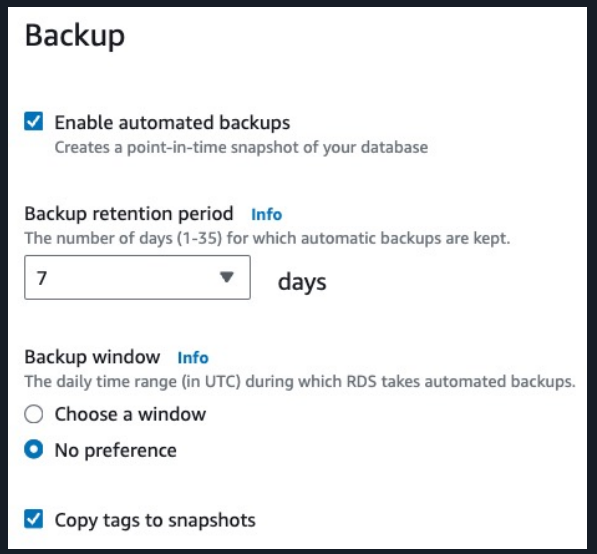

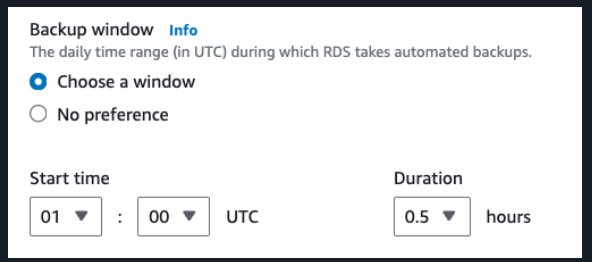

Amazon RDS automated backups

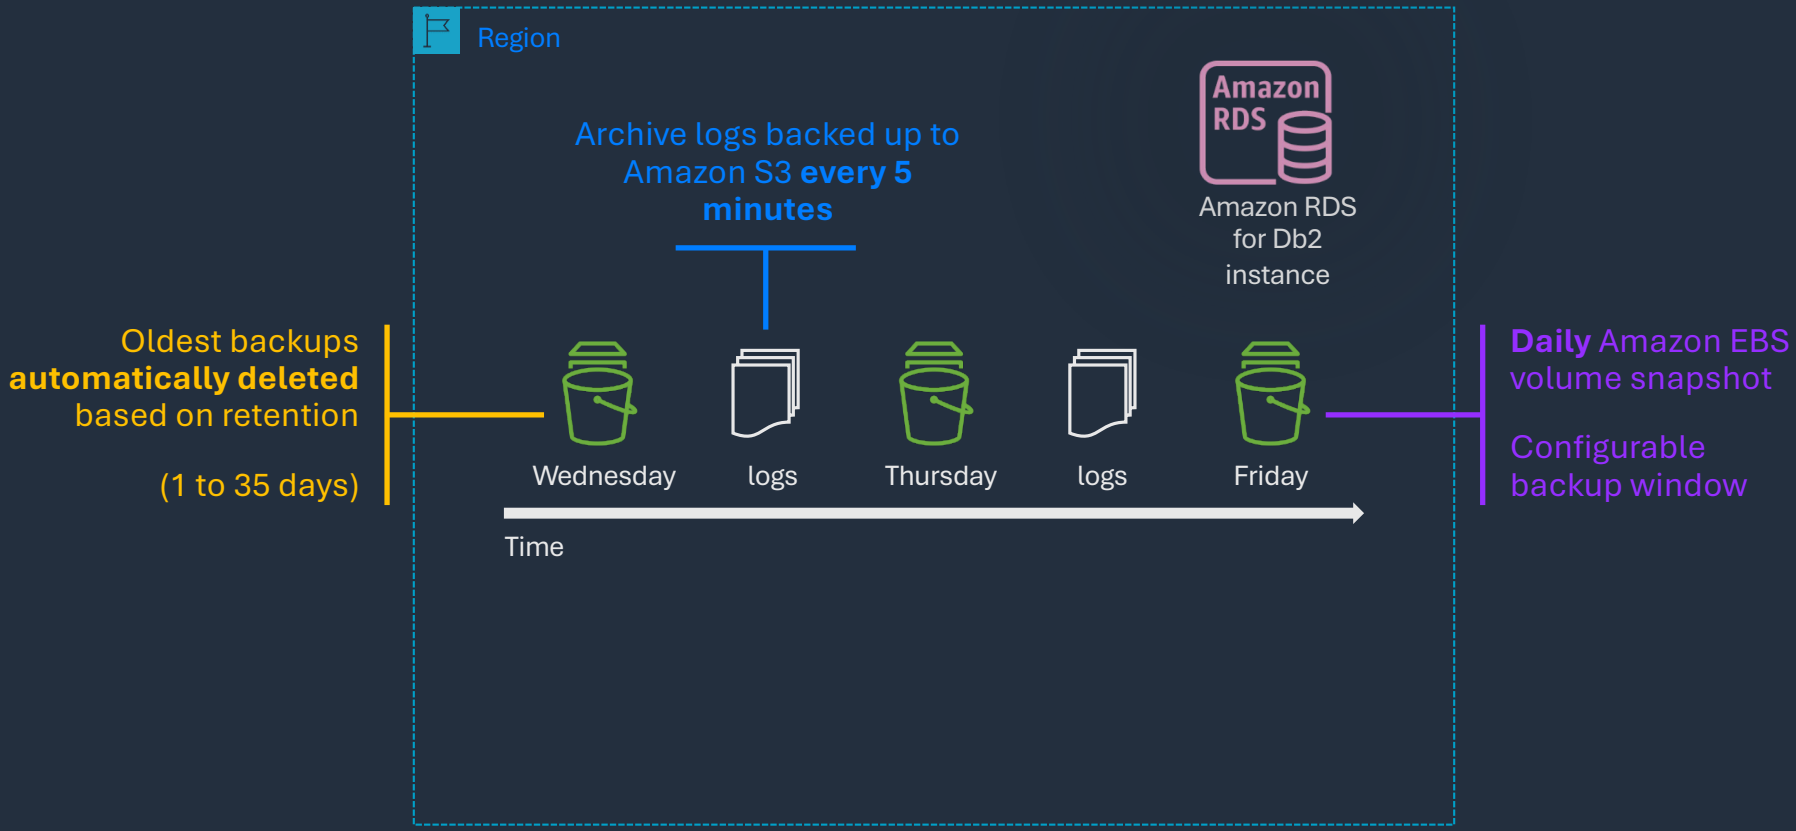

Automated backups

Point-in-time recovery for your DB instance

- Scheduled daily volume backup of entire instance
- Archive database change logs
- 35–days maximum retention
- Minimal impact on database performance
- Taken from standby when running Multi-AZ

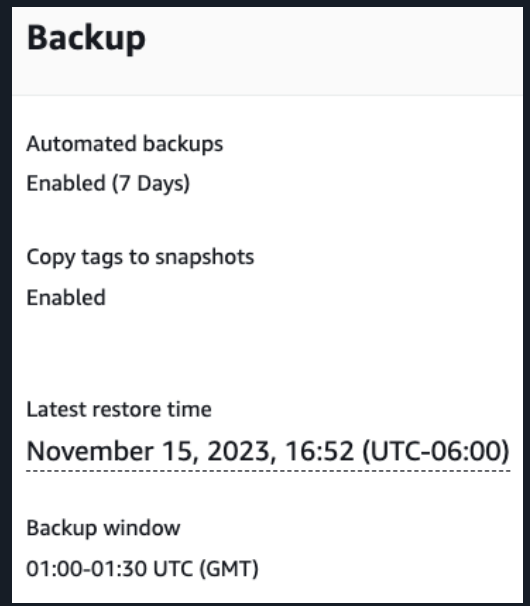

Every day during your backup window, RDS creates a storage volume snapshot of your instance

Every five minutes, RDS backs up the archive logs of your database

Database snapshots

Backups of your entire DB instance in Amazon S3

- User initiated
- Always incremental
- Amazon S3 → 99.999999999% durability
- Supports encryption
- Copy across accounts, across regions

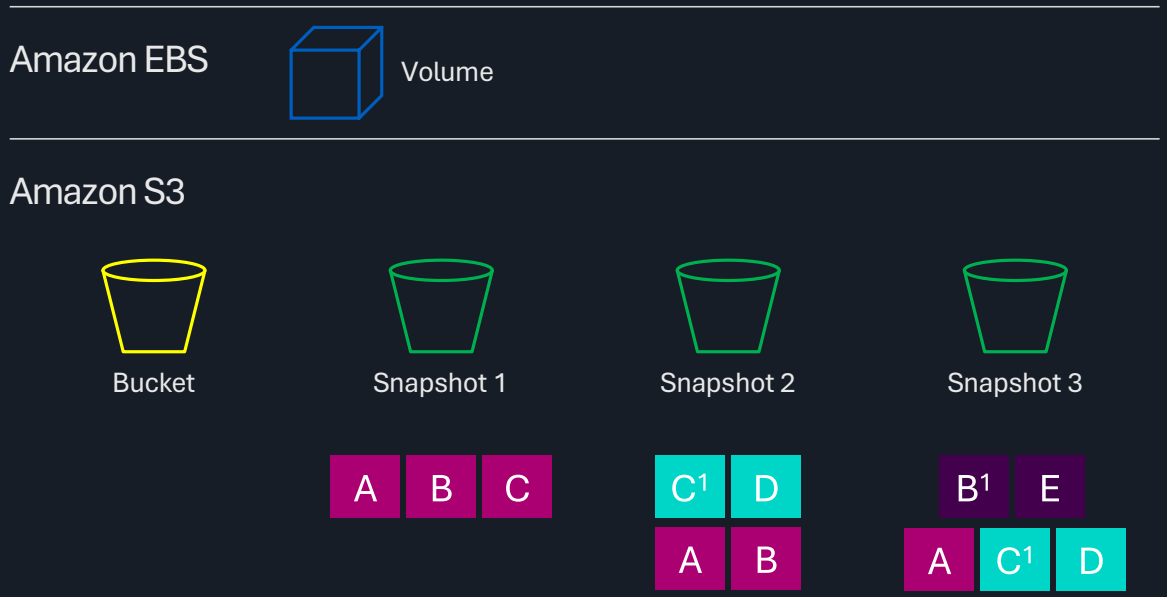

Automated backups vs snapshots

Automated backups	Manual Snapshots
• 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 • You can also retain automated backups upon instance deletion now	• 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 • Used for “sharing” snapshots.

RDS for Db2 – backup configuration and transaction logs

- Backup retention = 0 days → Circular logging mode. No automated backups.
- Backup retention >= 1 → Archive logging mode. Automated backups enabled.
- Default values: LOGFILSIZ = 25000, LOGPRIMARY=50, LOGSECOND=100
- db2 ARCHIVE LOG for database <dbName> – This is run every 5 minutes through RDS Db2 automation
- Default log size is logfilsiz * pagesize i.e. 25000*4 KB = 100000KB

Amazon RDS for Db2 – performance during backup

- During the automatic backup window, storage I/O might be suspended briefly while the backup process initializes (typically under a few seconds)
- You might experience elevated latencies for a few seconds during backups for Multi-AZ deployments
- Single-AZ DB instance results in a brief I/O suspension that can last from a few seconds to a few minutes

Restoring backup

- Restoring creates an entirely new database instance
- Define the instance configuration just like a new instance
- 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.

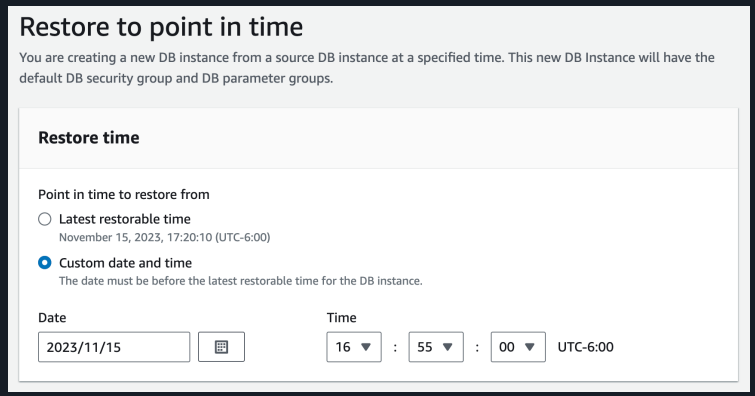

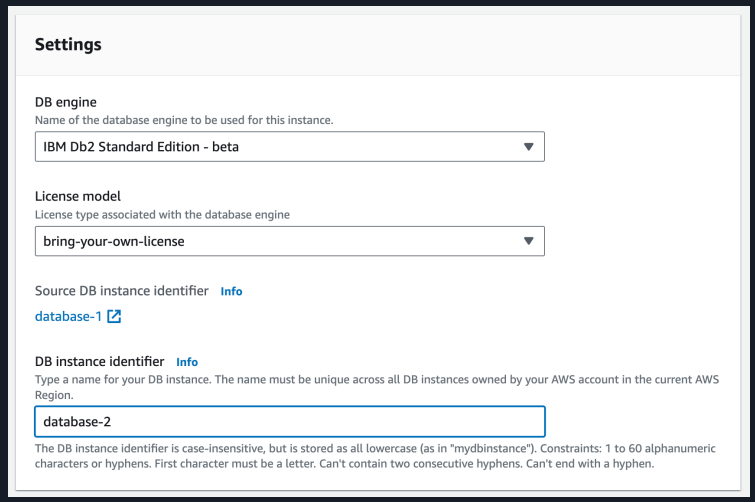

Snapshot/Point-in-time restore

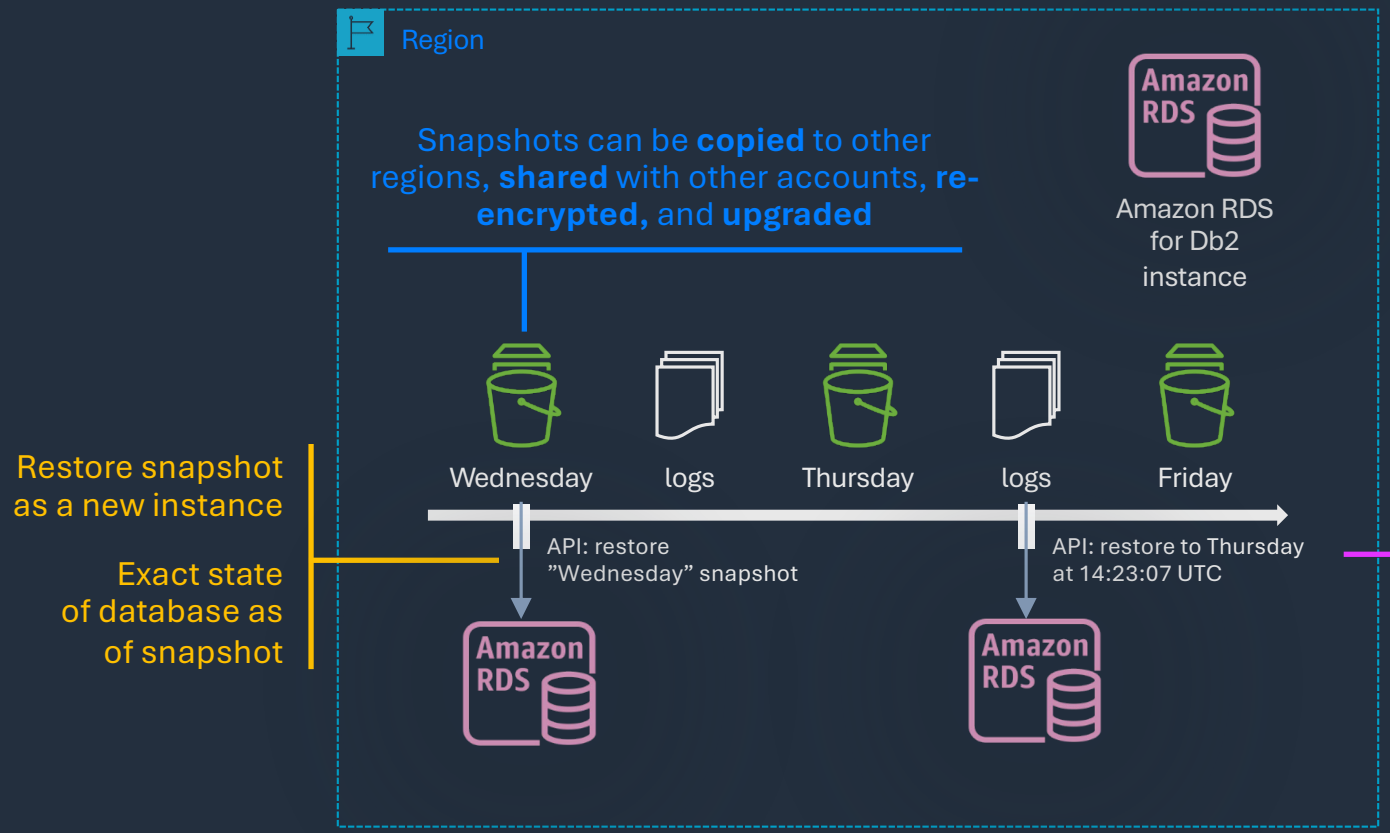

Amazon RDS backups and restores

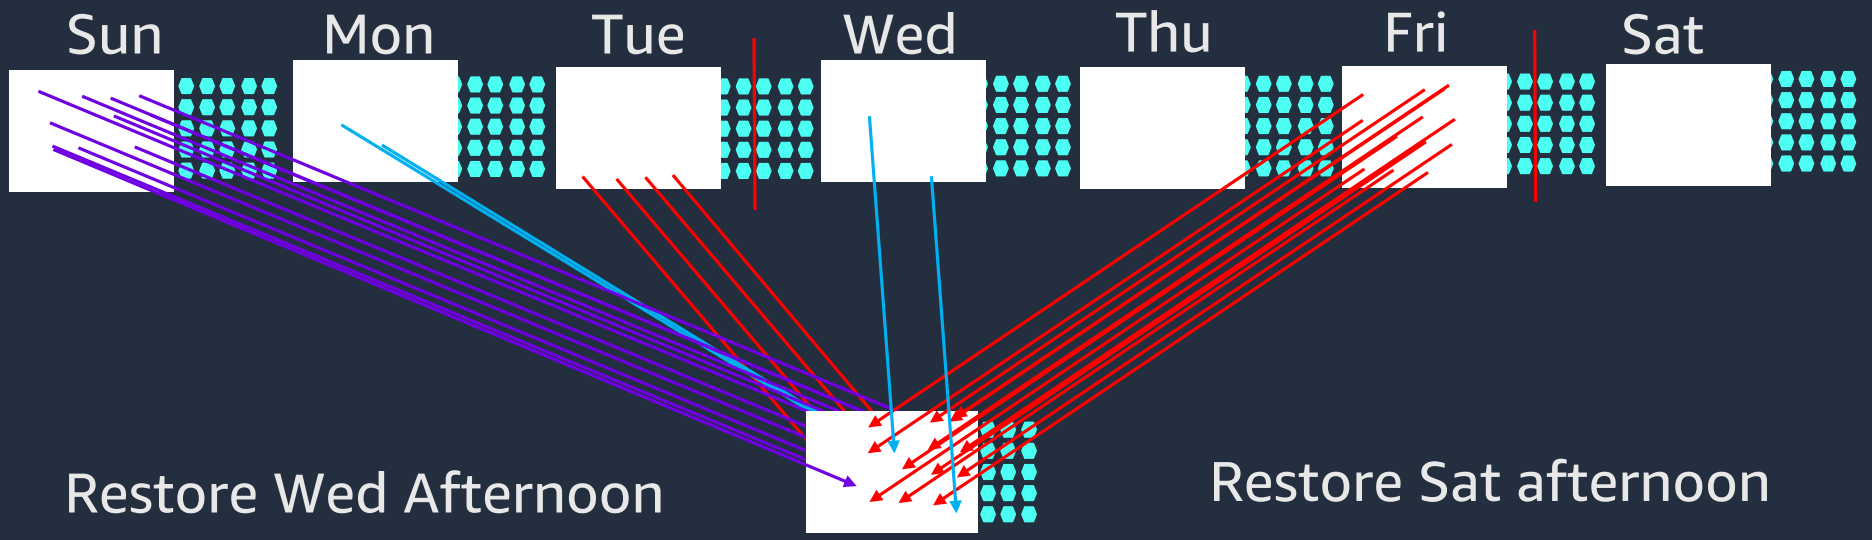

Fixed time for EBS storage
Variable for Archive logs

Amazon RDS for Db2 – Copy Snapshots

- You can only copy manual snapshots.
- First create a manual snapshot by copying an automatic snapshot or by taking a manual snapshot explicitly
- Typical use cases:
    - Copy to another region / Another account
    - Encrypt / change encryption keys

Disaster Recovery with Cross-Region Automated Backups

Cross-Region Key Aspects

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

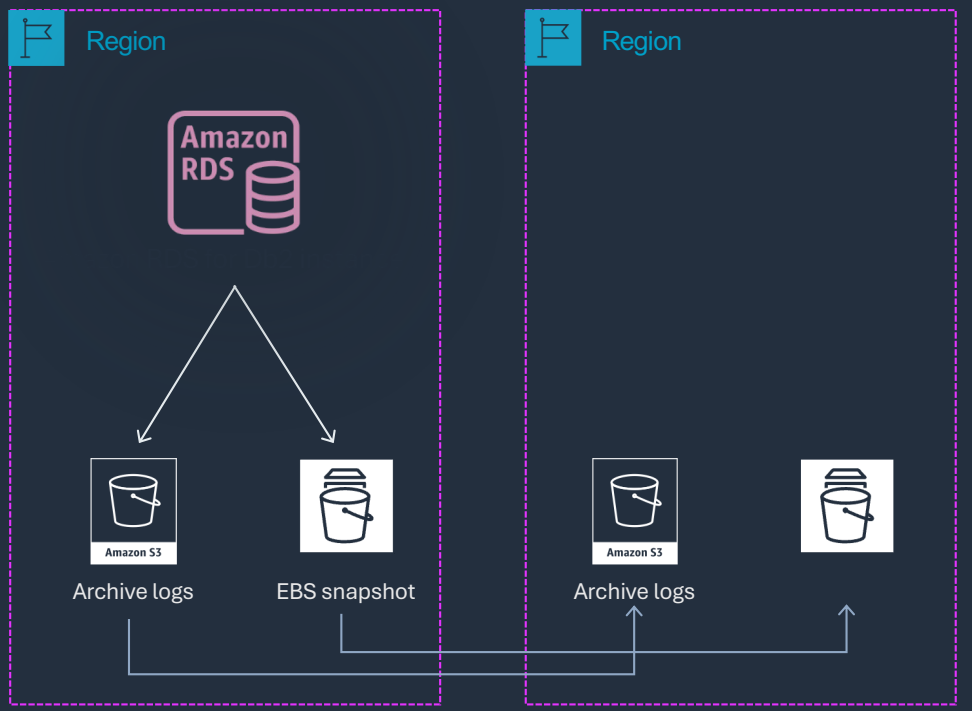

Amazon RDS cross-region automated backups

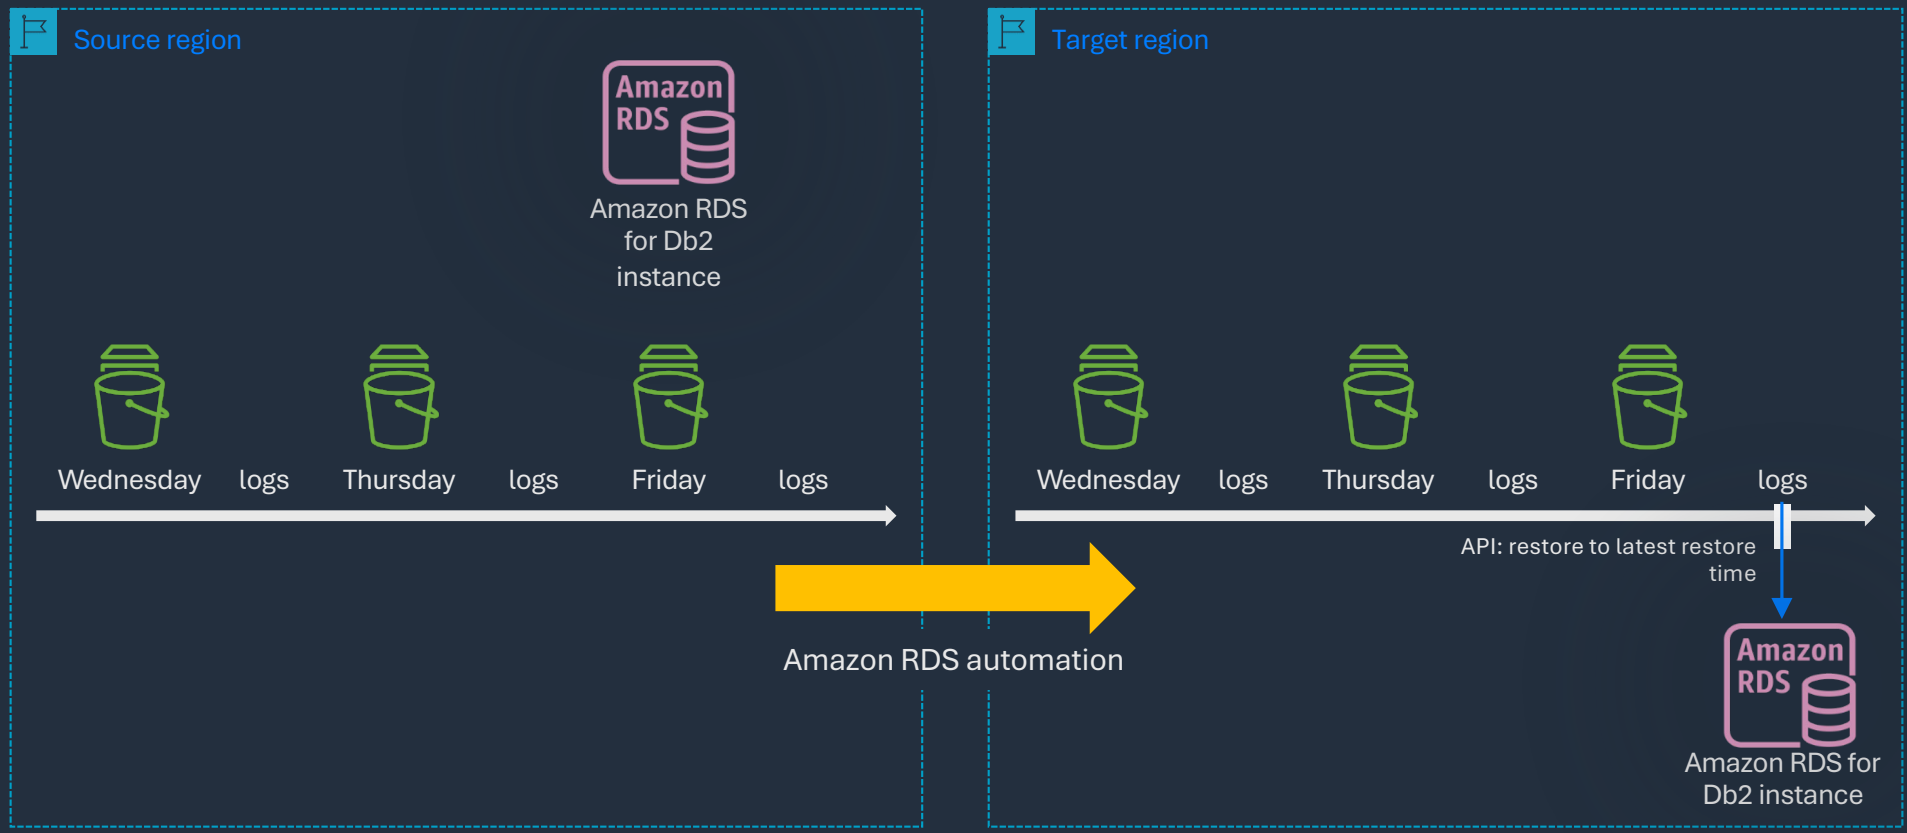

Backup and restore best practices

- Disable backups for data load (No archive logging mode) - WARNING: deletes existing automated backups
- Enable backups for critical workloads (Archive logging mode)
- Set backup window to low-usage time
- Take manual snapshots to reduce PiTR replay window
- Use restores to test upgrades/parameters/app changes
- Copy snapshots to other accounts/regions
- Manually hydrate using INSPECT – “db2 inspect check database results results.out”


aws
Thank You