



# Best Practices for Running Microsoft SQL Server on AWS

Domagoj Kovac – AWS Database Sales Specialist  
Baris Furtinalar – AWS Specialist Solutions Architect

# Table of contents

- Where can I run SQL Server on AWS
- Amazon RDS for SQL Server
- Optimizing Performance for SQL Server on EC2
- Highly Availability Best Practices for SQL Servers on AWS
- Migrating SQL Server to EC2
- Q&A

# Where can I run my SQL Server workloads on AWS?



## Amazon Relational Database Service (Amazon RDS)

Managed service with up to 128 vCPU, 4TB RAM, and 16 TB storage



## Amazon Elastic Compute Cloud (Amazon EC2)

Self-managed virtual machine with up to 448 vCPU, 24TB RAM, and 400-TB storage

# 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		✓

# 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

# Amazon RDS with SQL Server





# Amazon RDS: Hassle-free flexibility

Instance type selection:

- t3 instances for lower cost

- m5/m5d instances

- r5/z1d/x1e instances for more RAM


EBS volume type selection:

- GP2 volumes for lower cost

- IO1 volumes for I/O intensive, critical workloads



# Compute and Storage Scaling



**Scale Compute to  
Handle Increased Load**  
Up to 128 vCPUs  
4TiB of RAM



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



**Scale Down to  
Control Costs**  
As little as 4 vCPU and  
16 GiB of RAM

# Automated backups, Snapshots and Native Backups

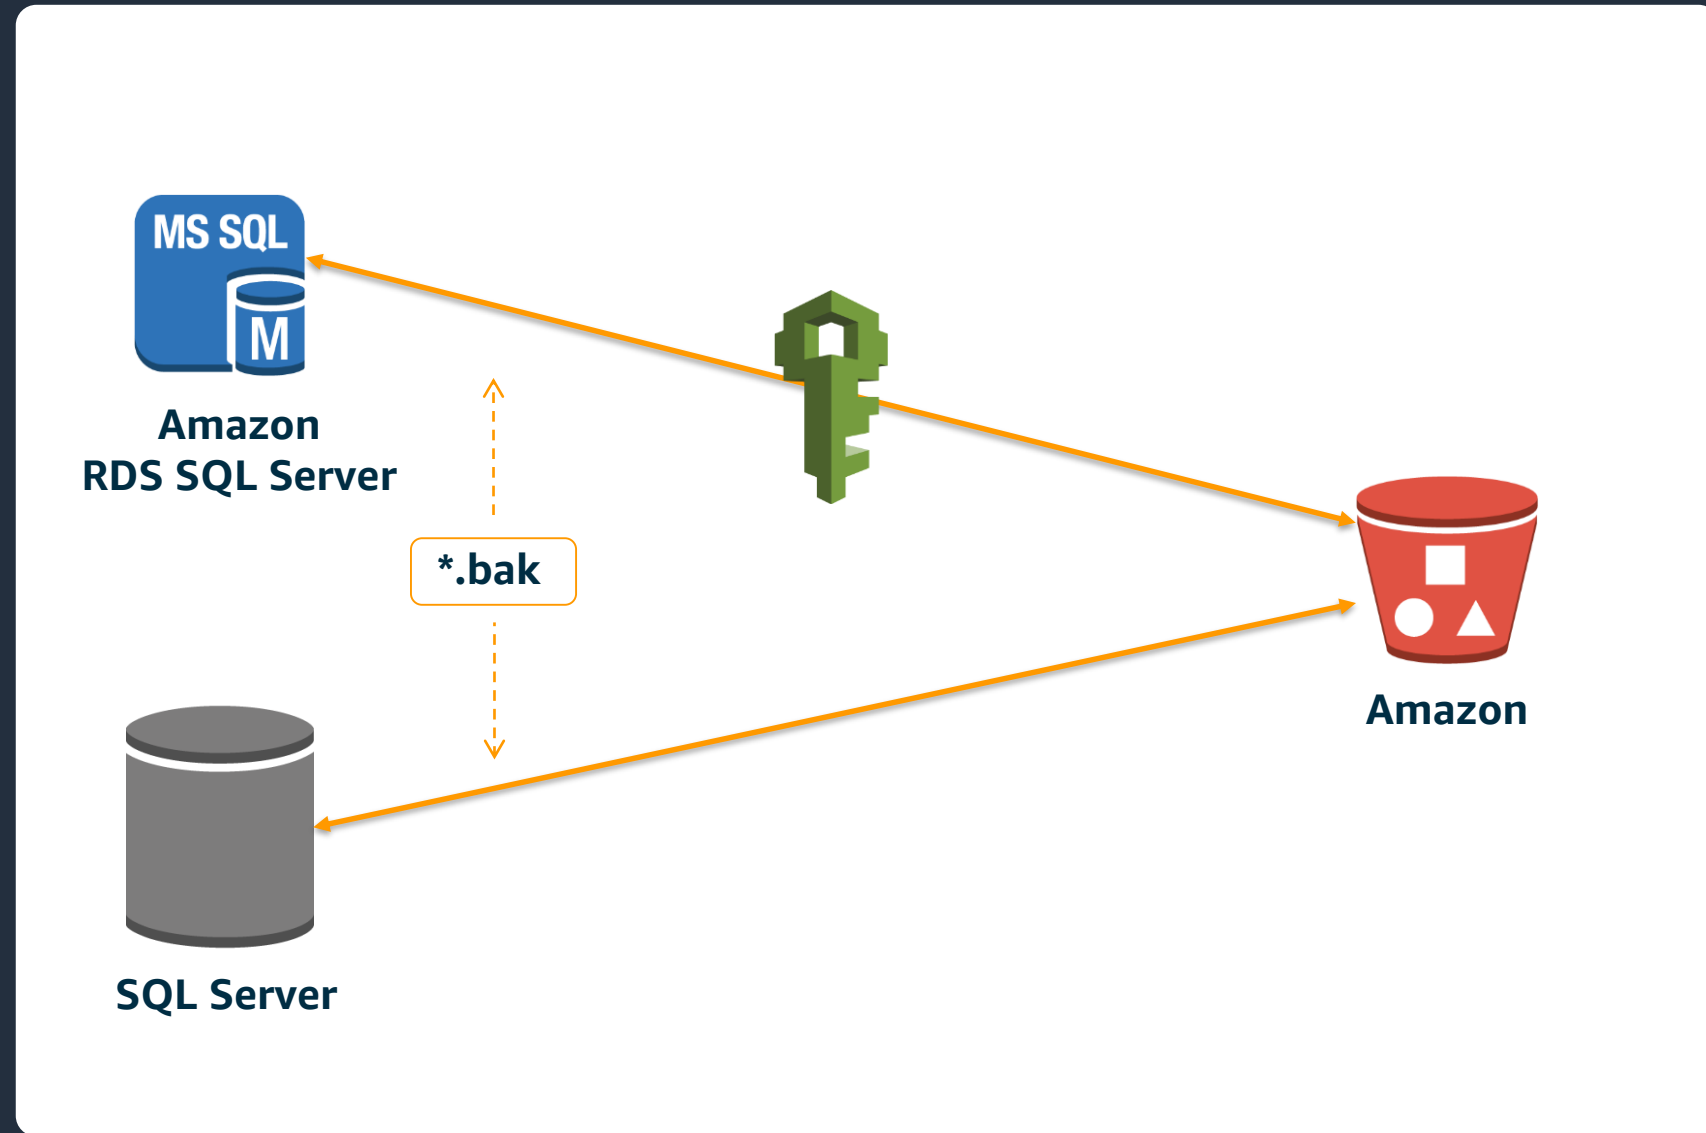
- Backup and restore directly to and from S3 bucket
- Supports Compression
- Only full & diff backups (no t-log)
- Full, diff & t-log restores
- Multi-file backup/restore

## 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=0/1],
    [@type='DIFFERENTIAL/FULL'],
    [@number_of_files=n];
```

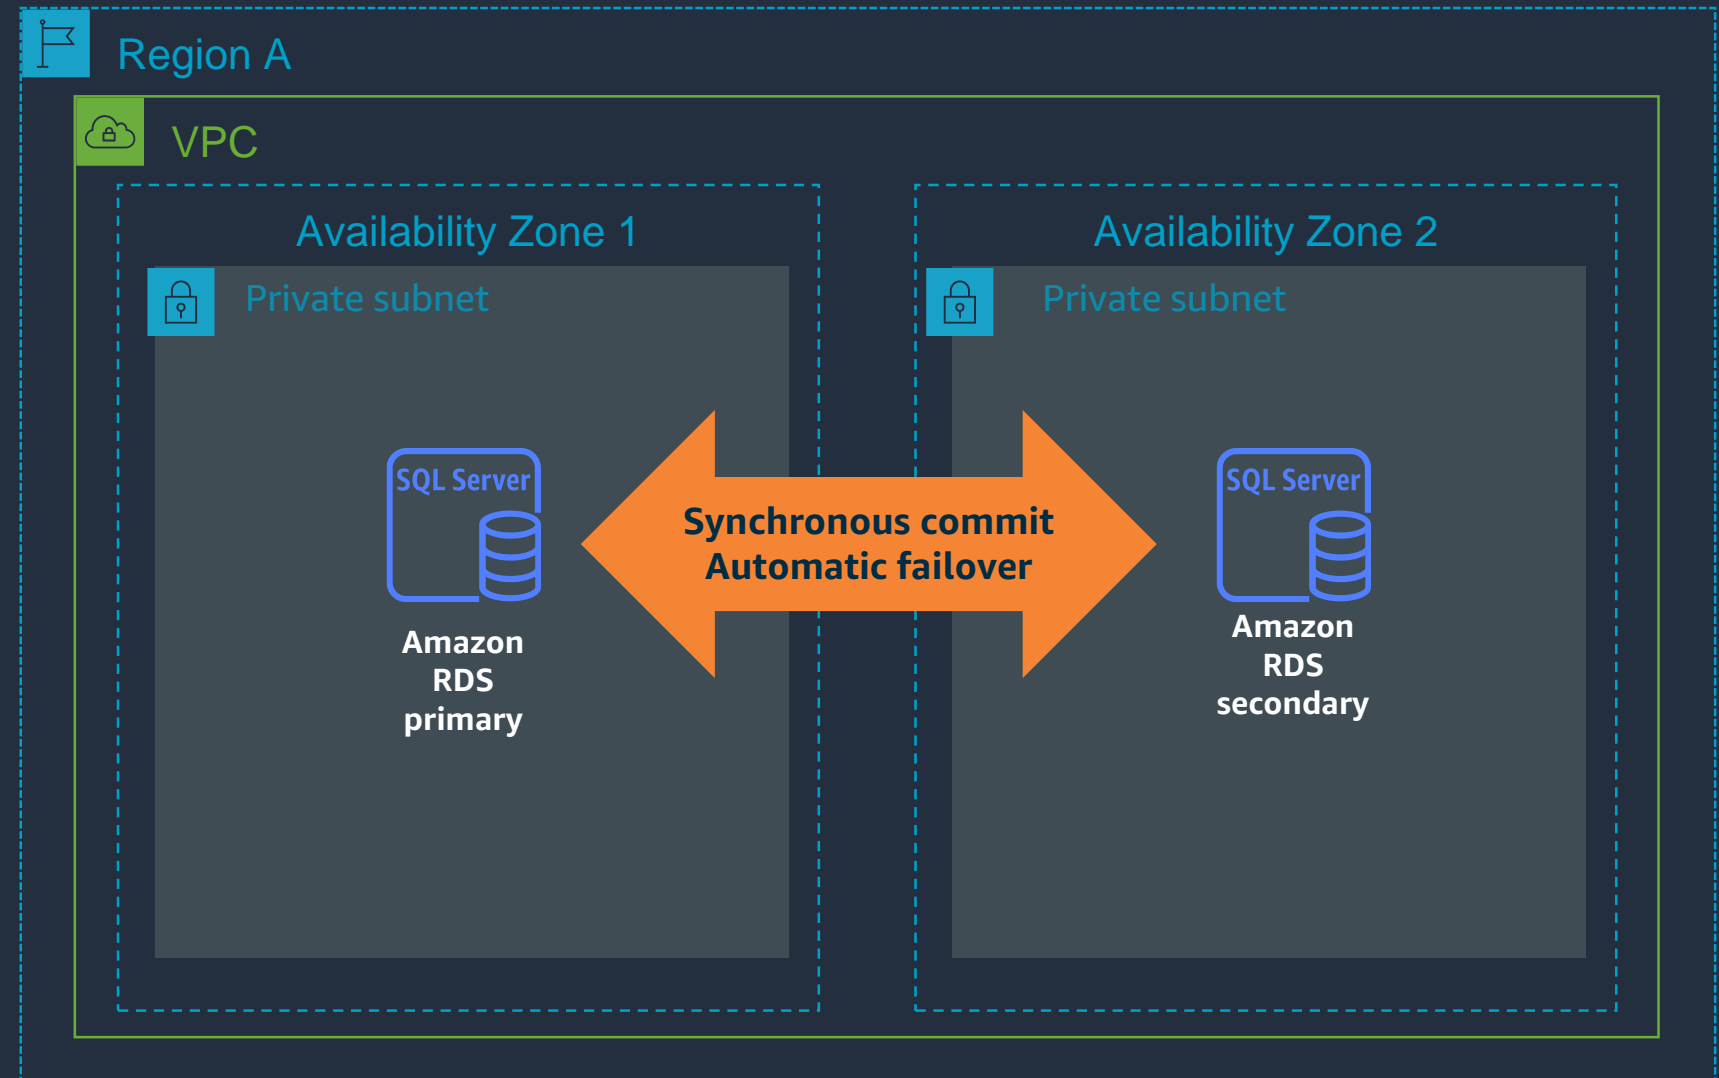
The following parameters are required:

- @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.



# Amazon RDS Multi-AZ SQL Server

- 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



# Read scale-out with Amazon RDS

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

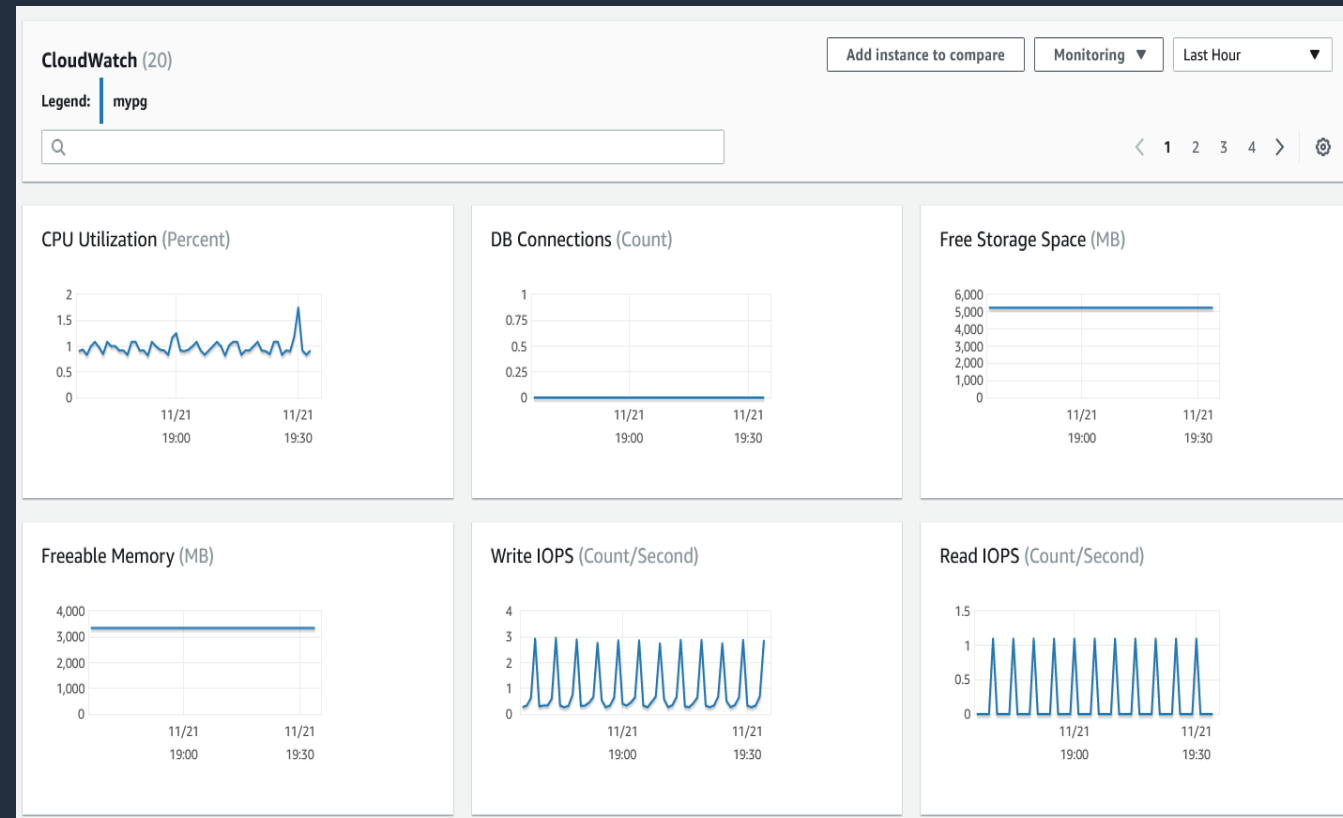
The screenshot displays the AWS Management Console interface for Amazon RDS. The left sidebar shows navigation options like Dashboard, Databases, Query Editor, and Performance Insights. The main content area shows the 'Databases' page with a table of database instances. The 'Actions' menu is open for 'database-3', showing options like Stop, Reboot, Delete, and 'Create read replica' (which is highlighted). The table below shows the following data:

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

# 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



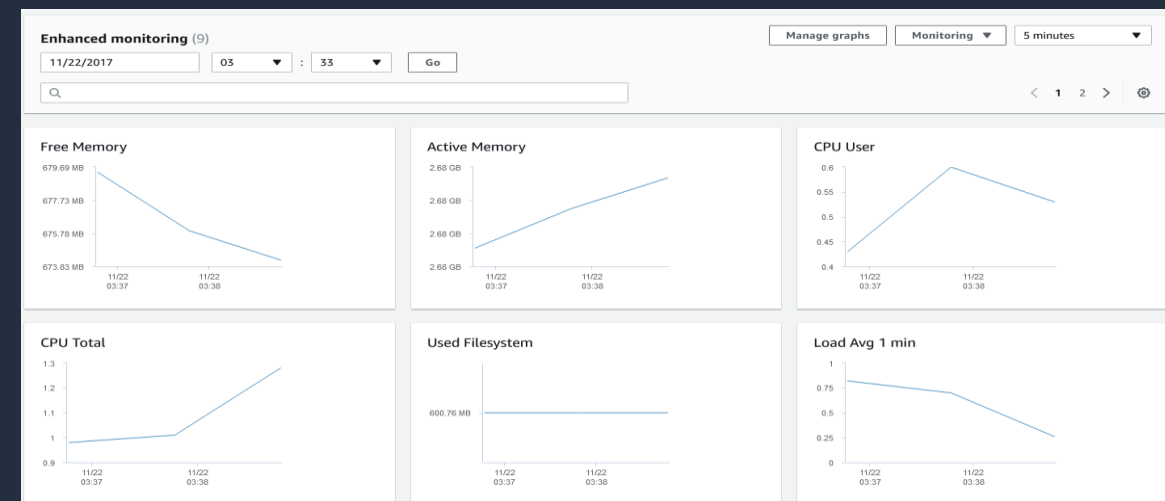
# Amazon RDS Enhanced Monitoring

## Overview:

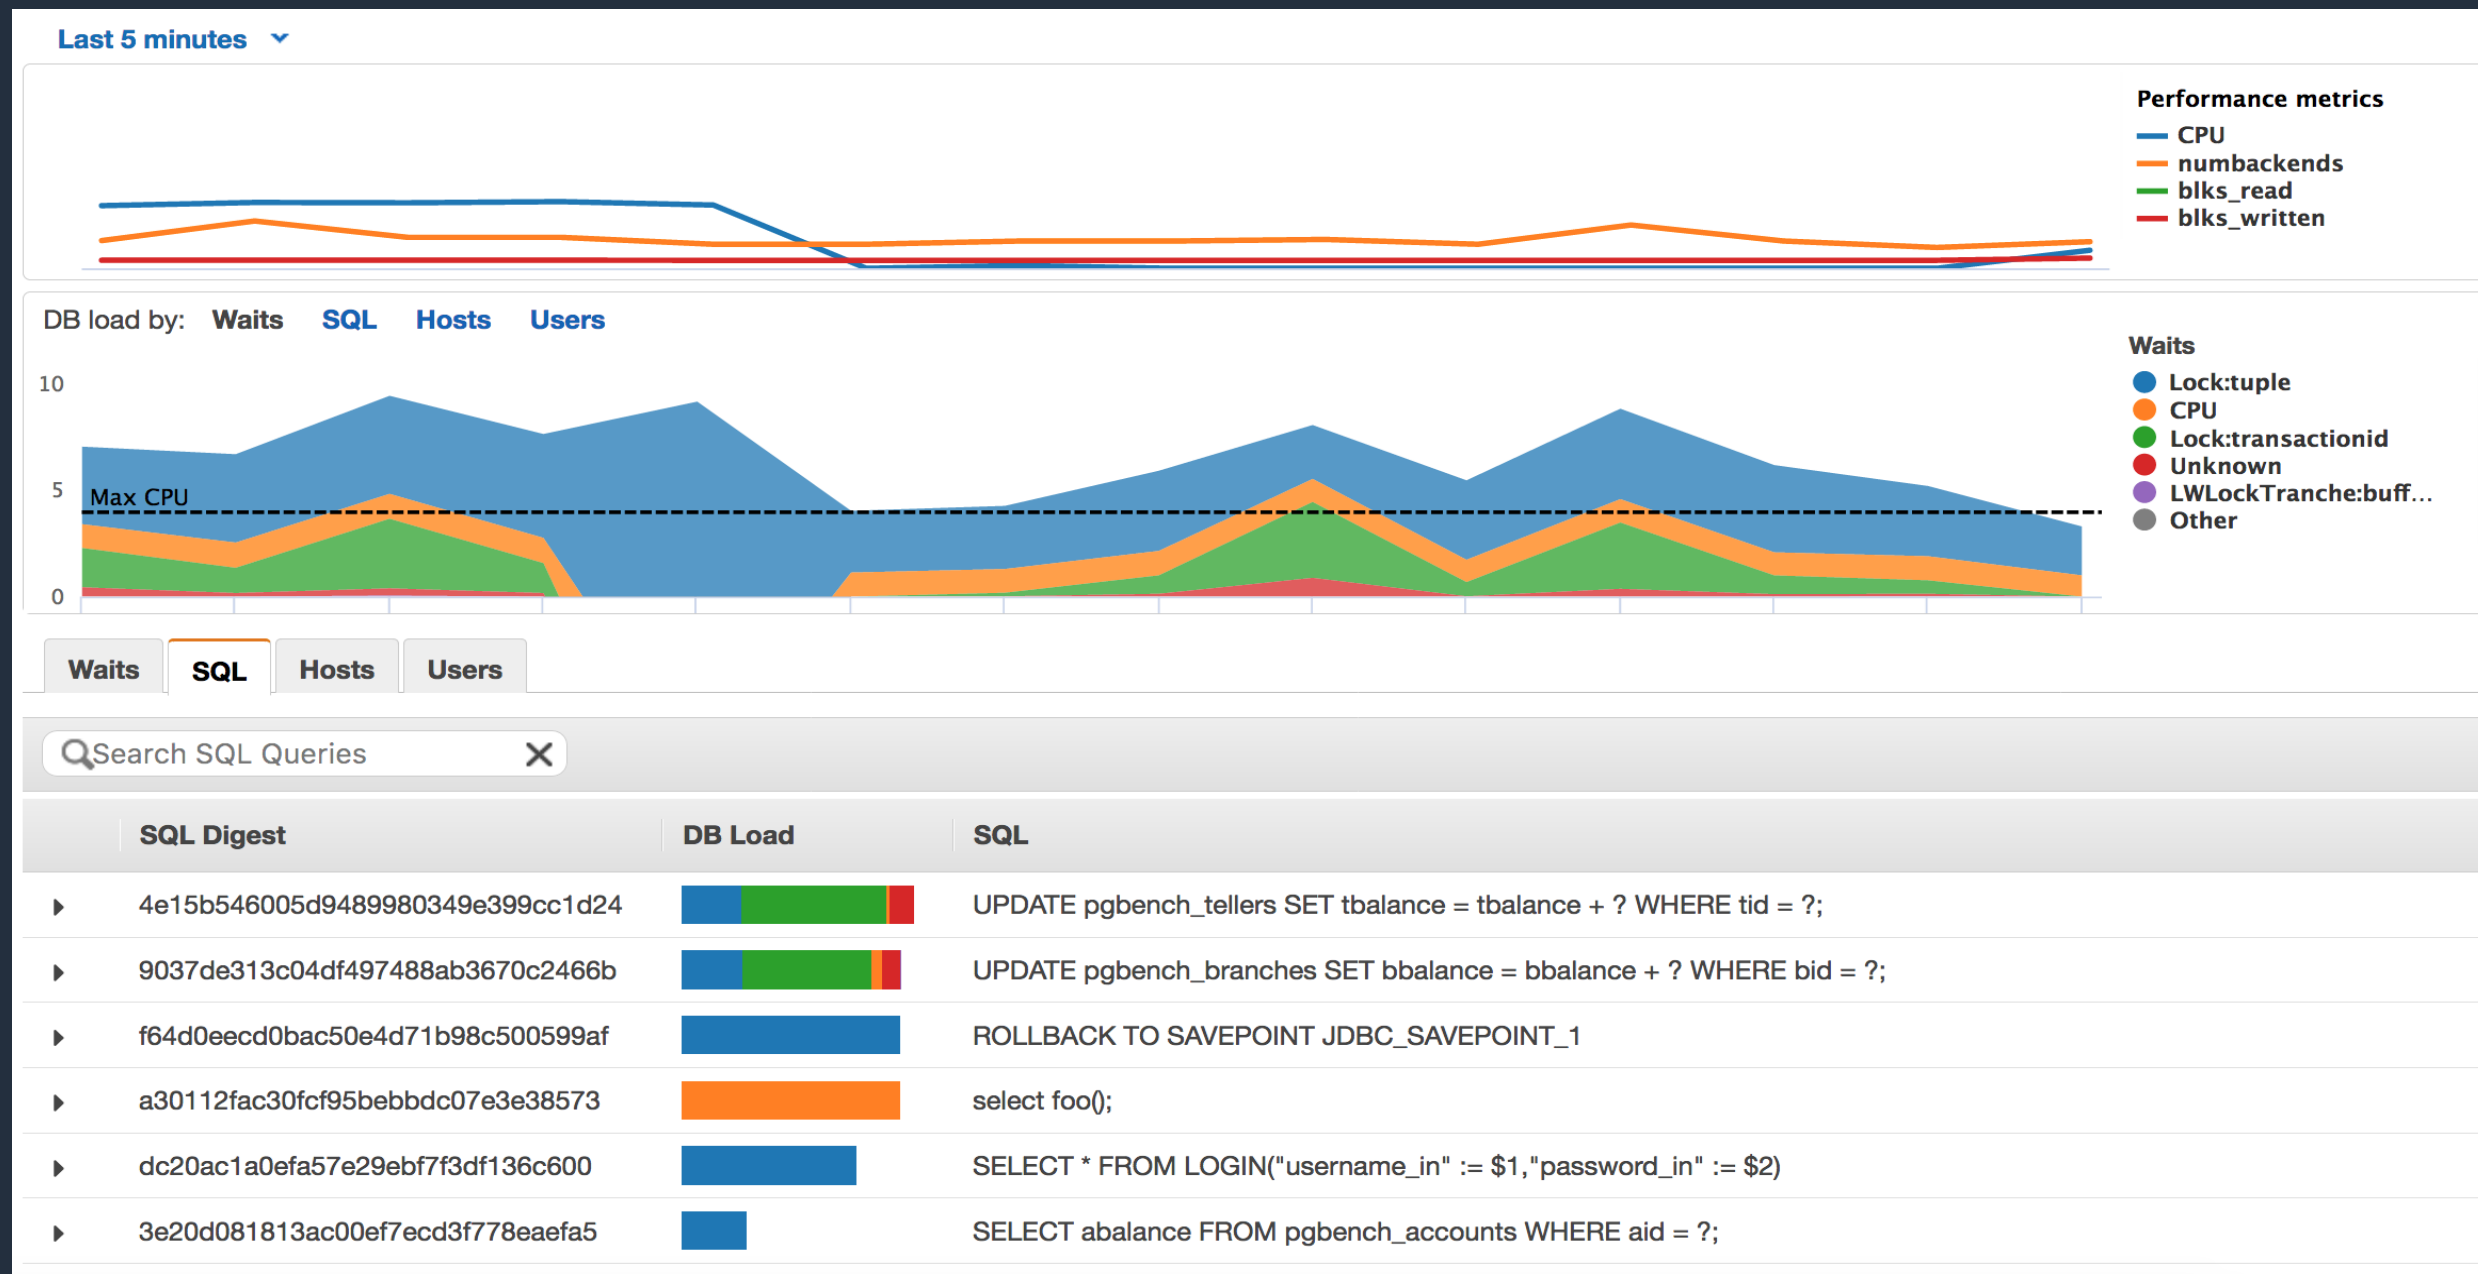
- OS Level Monitoring Metrics – 26 system and per process metrics
- Metrics delivered to CloudWatch Logs
- Up to 1 second granularity

## Compared to CloudWatch Metrics:

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



# Performance Insights for RDS SQL Server







# 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)**

# 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 EC2 with SQL Server



# SQL Server on Amazon EC2



- How to start
  - Spin up an EC2 instance and install from customer media (BYOL)
  - Spin up an EC2 instance that includes SQL in the Amazon Machine Instance (AMI)
- Self-Managed SQL Server
  - Security, patching & updates
  - Backups
  - High Availability
- Full SQL Server sysadmin privileges

# Optimizing Performance for SQL Server on EC2



# Amazon EC2: Purpose-built computing

Current instance families and generation	Family/usage
<b>M5, M5a, M5zn (4.5GHz)</b>	<b>General purpose compute</b>
T2, T3	Burstable performance
C5, C5n, C5a, C4	Compute optimized
<b>X1, X1E, R5, R5d, R5a, Z1d (4.0GHz)</b>	<b>Memory optimized</b>
P2, G3, F1	Accelerated computing
<b>I3, I3en</b>	<b>Storage optimized (I/O)</b>
H1, D2	Storage optimized (Density)

# Get the maximum EBS performance on Amazon EC2

## General Purpose SSD

GP3



99.8% - 99.9% durability

1 GiB - 16 TiB

Max IOPS  
per volume  
(16 KiB I/O)

16,000

16,000




32,000

Max throughput  
per volume

1,000 MiB/s

## Amazon EC2



R5b.4xlarge

Max EBS Bandwidth 10,000Mbps

Maximum IOPS 43,333

# Get the maximum EBS performance on Amazon EC2

## Provisioned IOPS SSD

IO2



99.999% durability

4 GiB - 16 TiB

Max IOPS  
per volume  
(16 KiB I/O)

64,000

64,000




128,000

Max throughput  
per volume

1,000 MiB/s

## Amazon EC2



R5b.12xlarge

Max EBS Bandwidth 30,000Mbps

Maximum IOPS 130,000



# Place tempdb in EC2 Instance Store

## Amazon EC2




## What is instance storage?

- Some instance types come with direct attached disk-based storage
  - SSD
  - NVMe SSD
- Included in the hourly cost
- Data on instance storage does not persist a user- initiated instance stop/start or hardware failure
- Fast disk I/O without going over the network
- Traffic does not count towards EBS IOPS

# Place tempdb in EC2 Instance Store

## Amazon EC2



## What is instance storage?

- Some instance types come with direct attached disk-based storage
  - SSD
  - NVMe SSD
- Included in the hourly cost
- Data on instance storage does not persist a user- initiated instance stop/start or hardware failure
- Fast disk I/O without going over the network
- Traffic does not count towards EBS IOPS

# Amazon FSx for Windows File Server

FSx

Amazon FSx for Windows  
File Server



Region

FSx Amazon FSx  
for Windows File Server

\\fs-0123456789.example.com



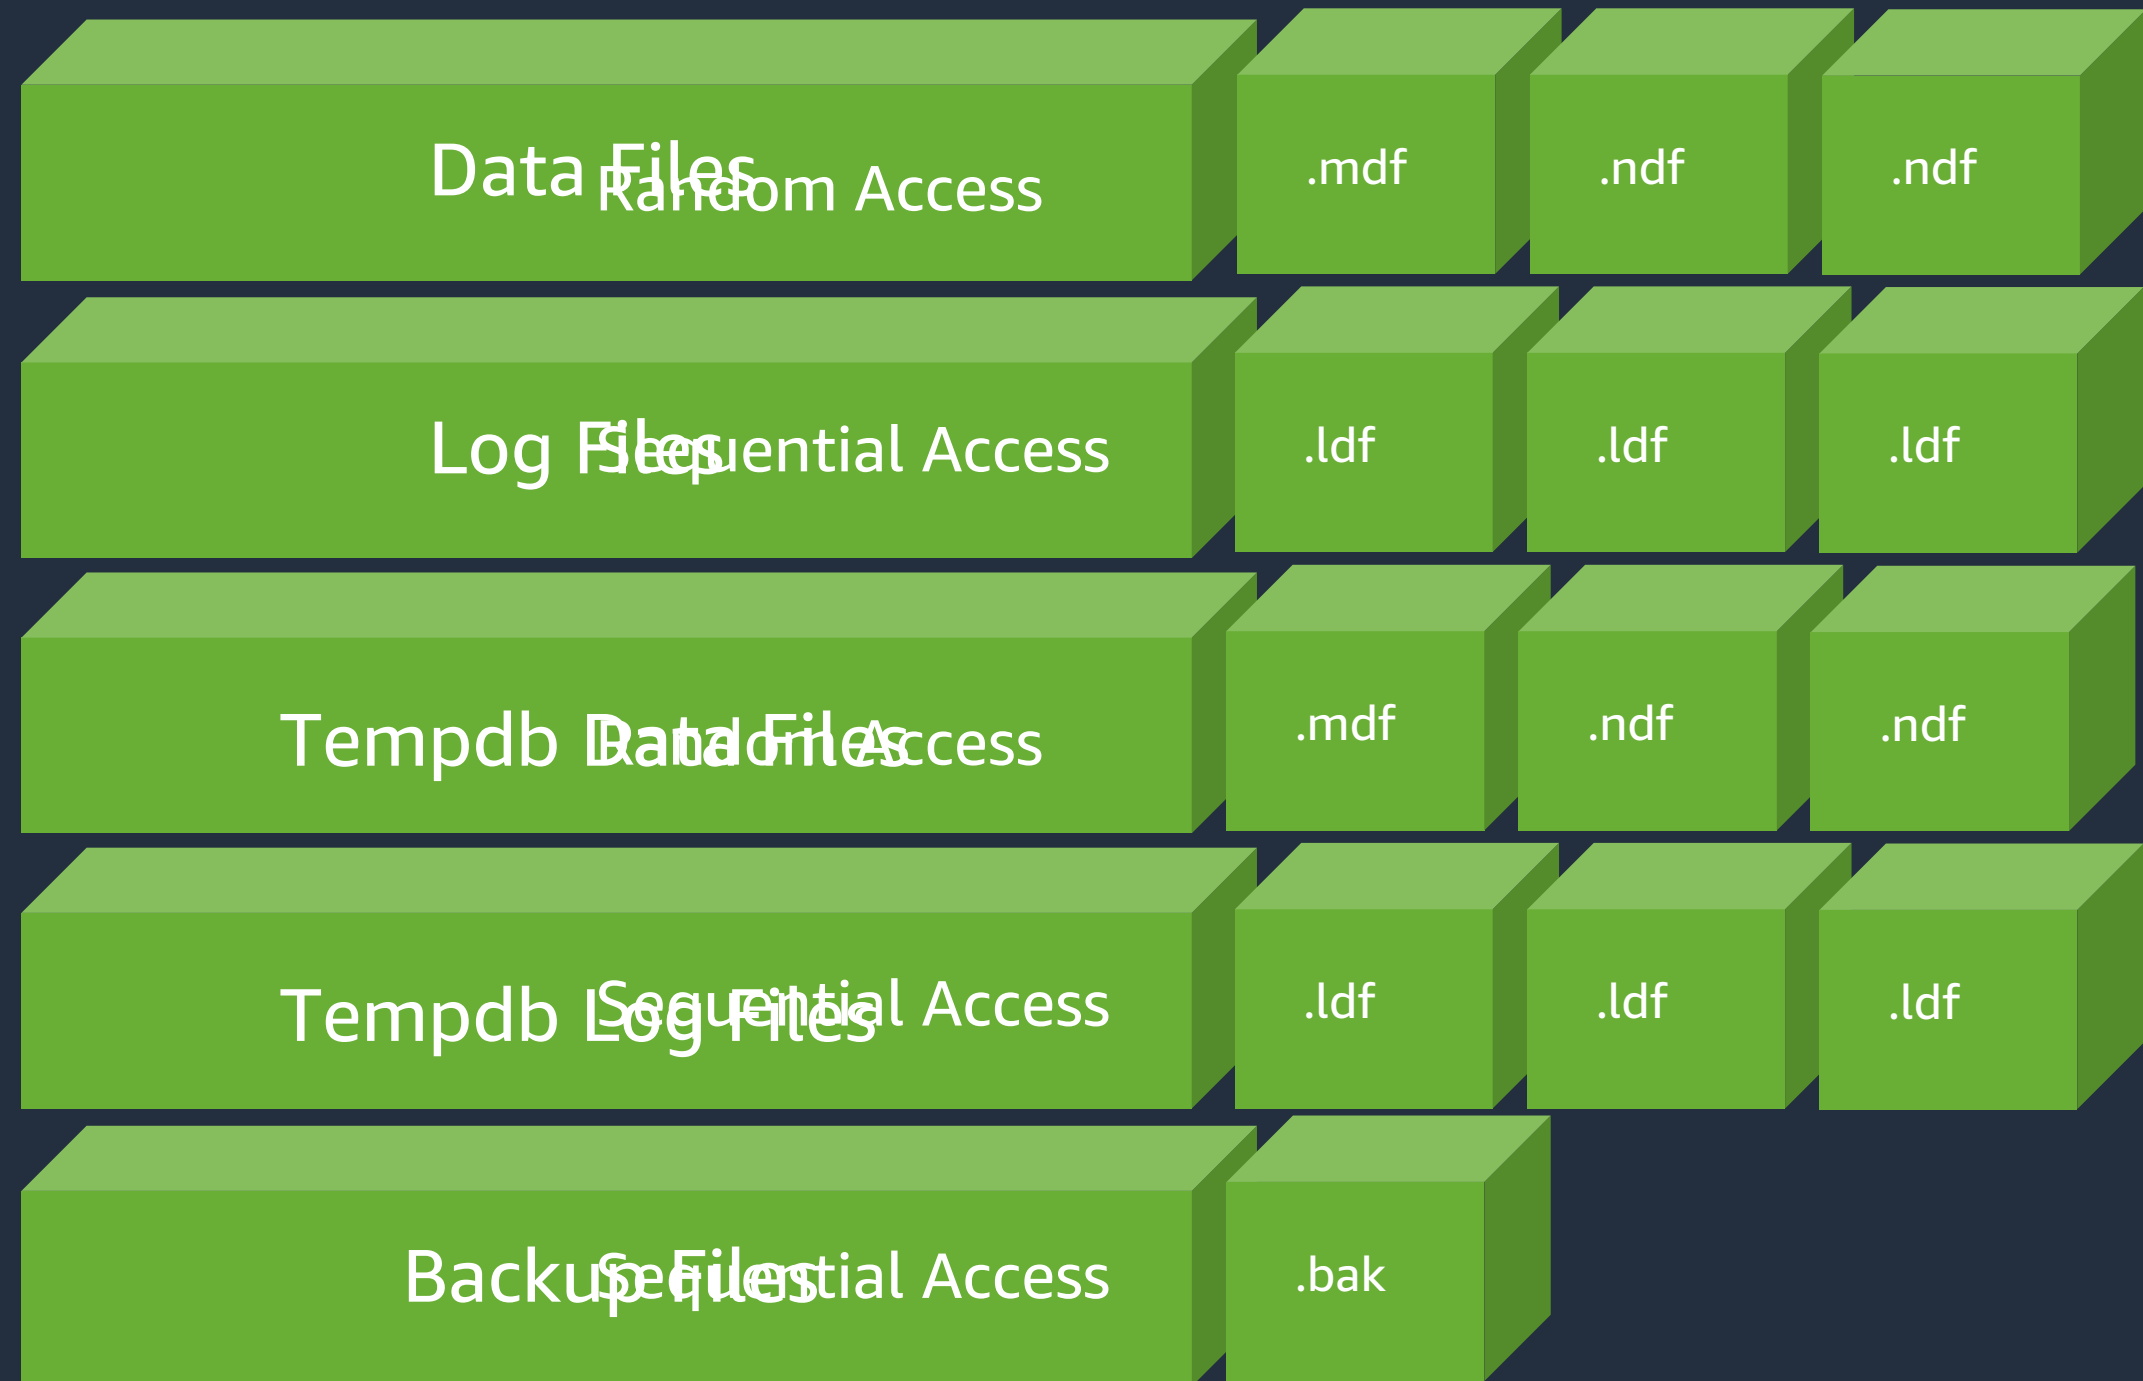
mdf

mdf


ldf

```
CREATE DATABASE [DBonFSx] ON PRIMARY ( name = N'DBonFSx_data',  
filename = N'\\Fsx_UNC_Path\DB\DBonFSx_data.mdf',  
size = 8192kb, maxsize = unlimited,filegrowth = 1024kb )  
log ON ( name = N'DBonFSx_log',  
filename =N'\\ Fsx_UNC_Path \DB\DBonFSx_log.ldf',  
size = 1024kb, maxsize = 2048gb, filegrowth = 10%)  
GO
```

# Optimize your disk layout or file distribution



# Optimize your disk layout or file distribution



# High Availability Best Practices for SQL Server on EC2

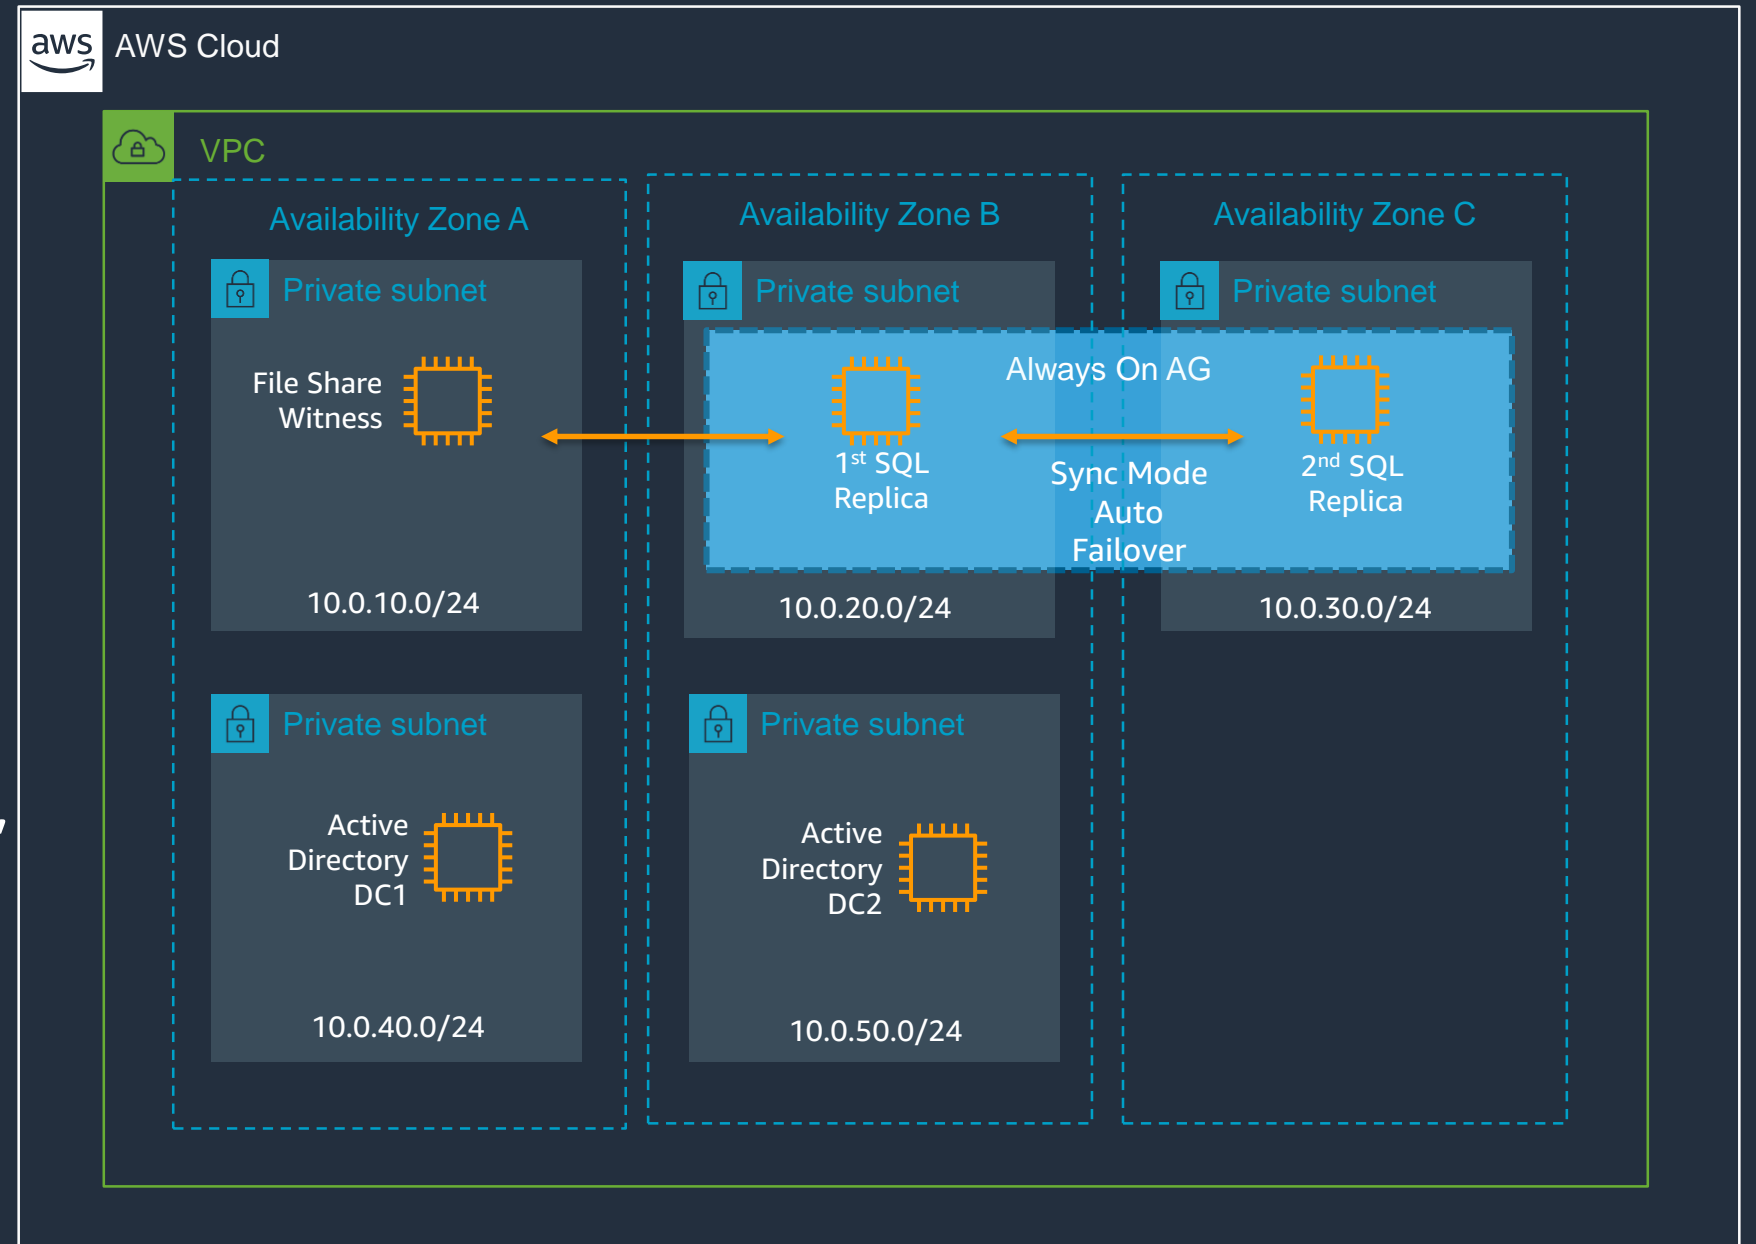


# SQL Server HA/DR on EC2

- Windows clusters can span Availability Zones or regions
- Mirroring
- AlwaysOn Availability Groups
- Transaction Log Shipping
- SQL Failover Cluster Instance

# Multi-AZ Always On Availability Group

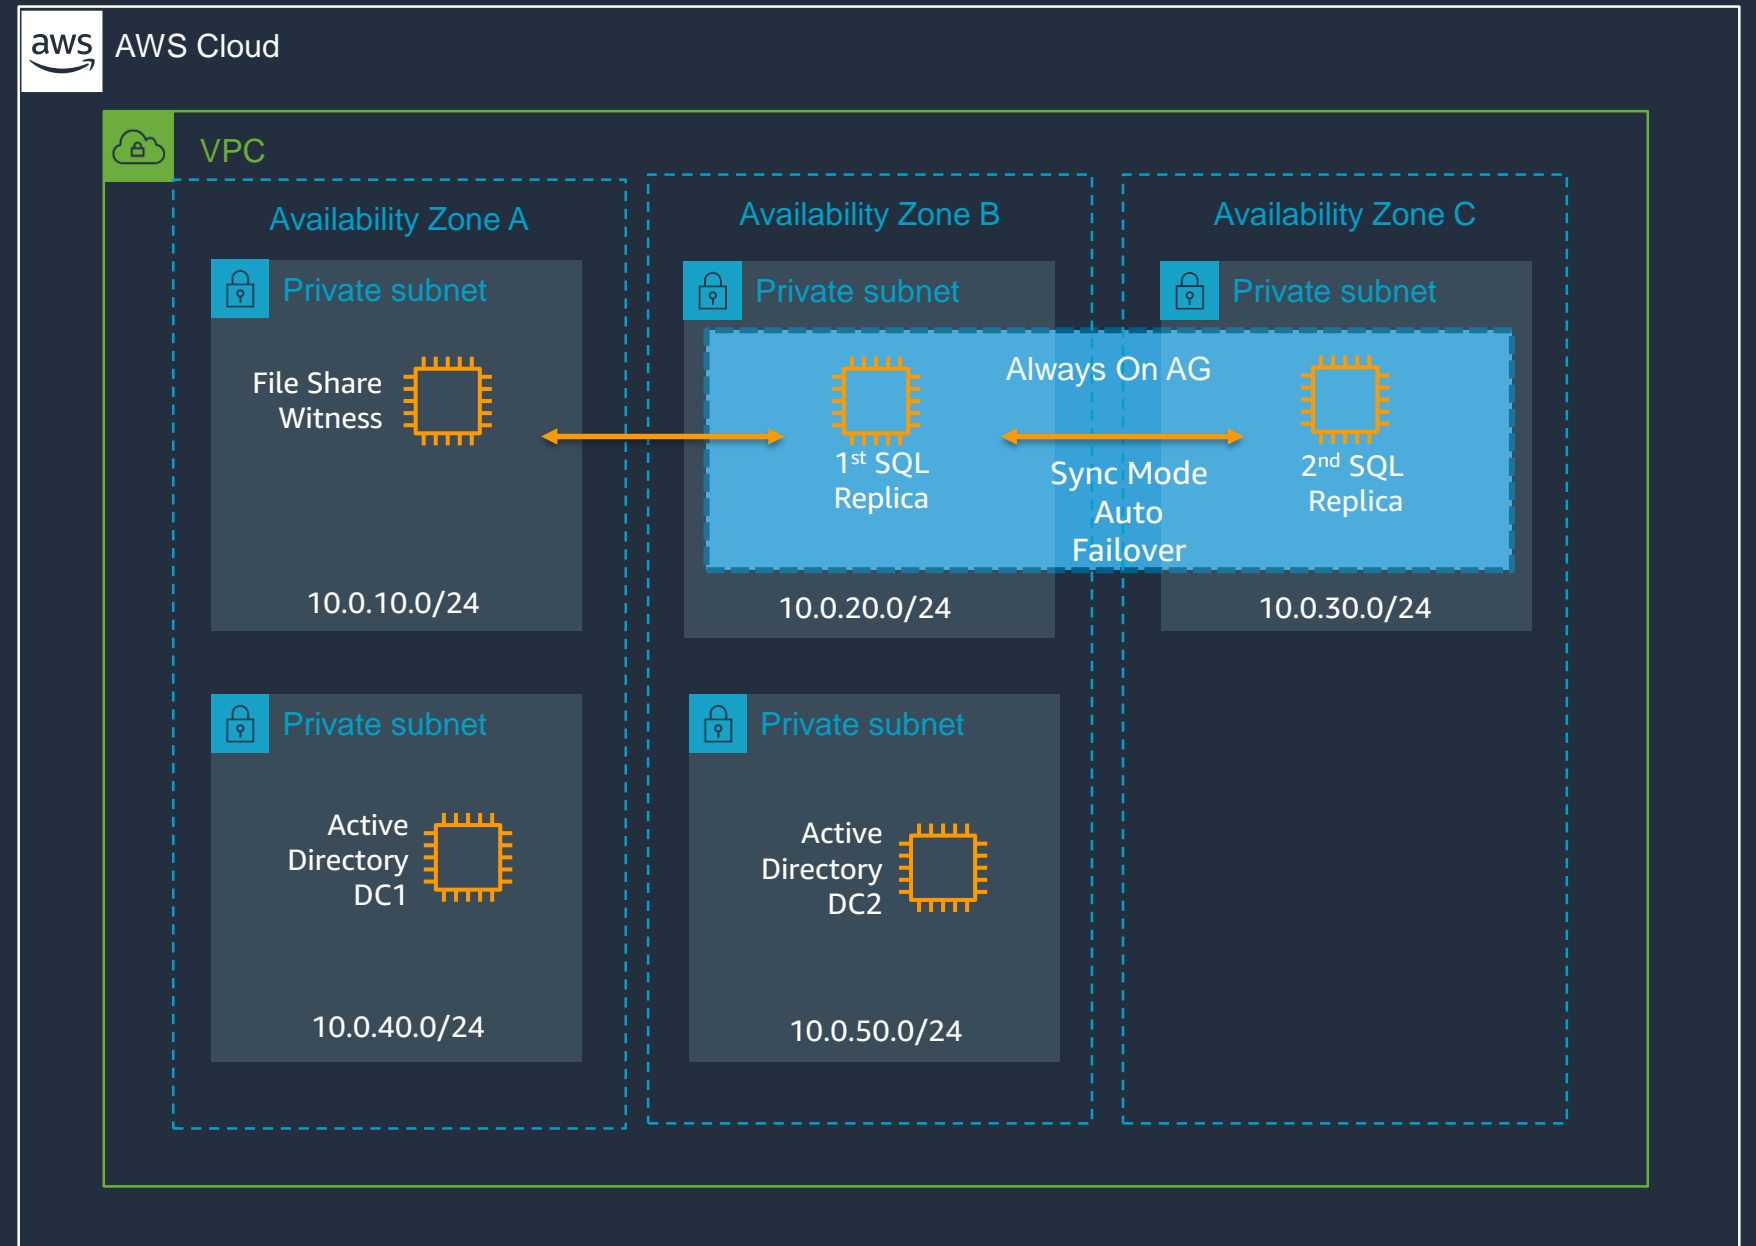
- Infrastructure:
  - Shared-nothing configuration
  - Storage configured and managed outside of WSFC
  - Flexibility and granularity of storage configuration and maintenance
  - Limited SQL resources managed by WSFC
- SQL Server:
  - Database-level HA solution
  - Manual adding logins, SQL Agent jobs, certificates and other SQL Server instance level objects to all secondary
  - Data/Transactions replicated through SQL Server



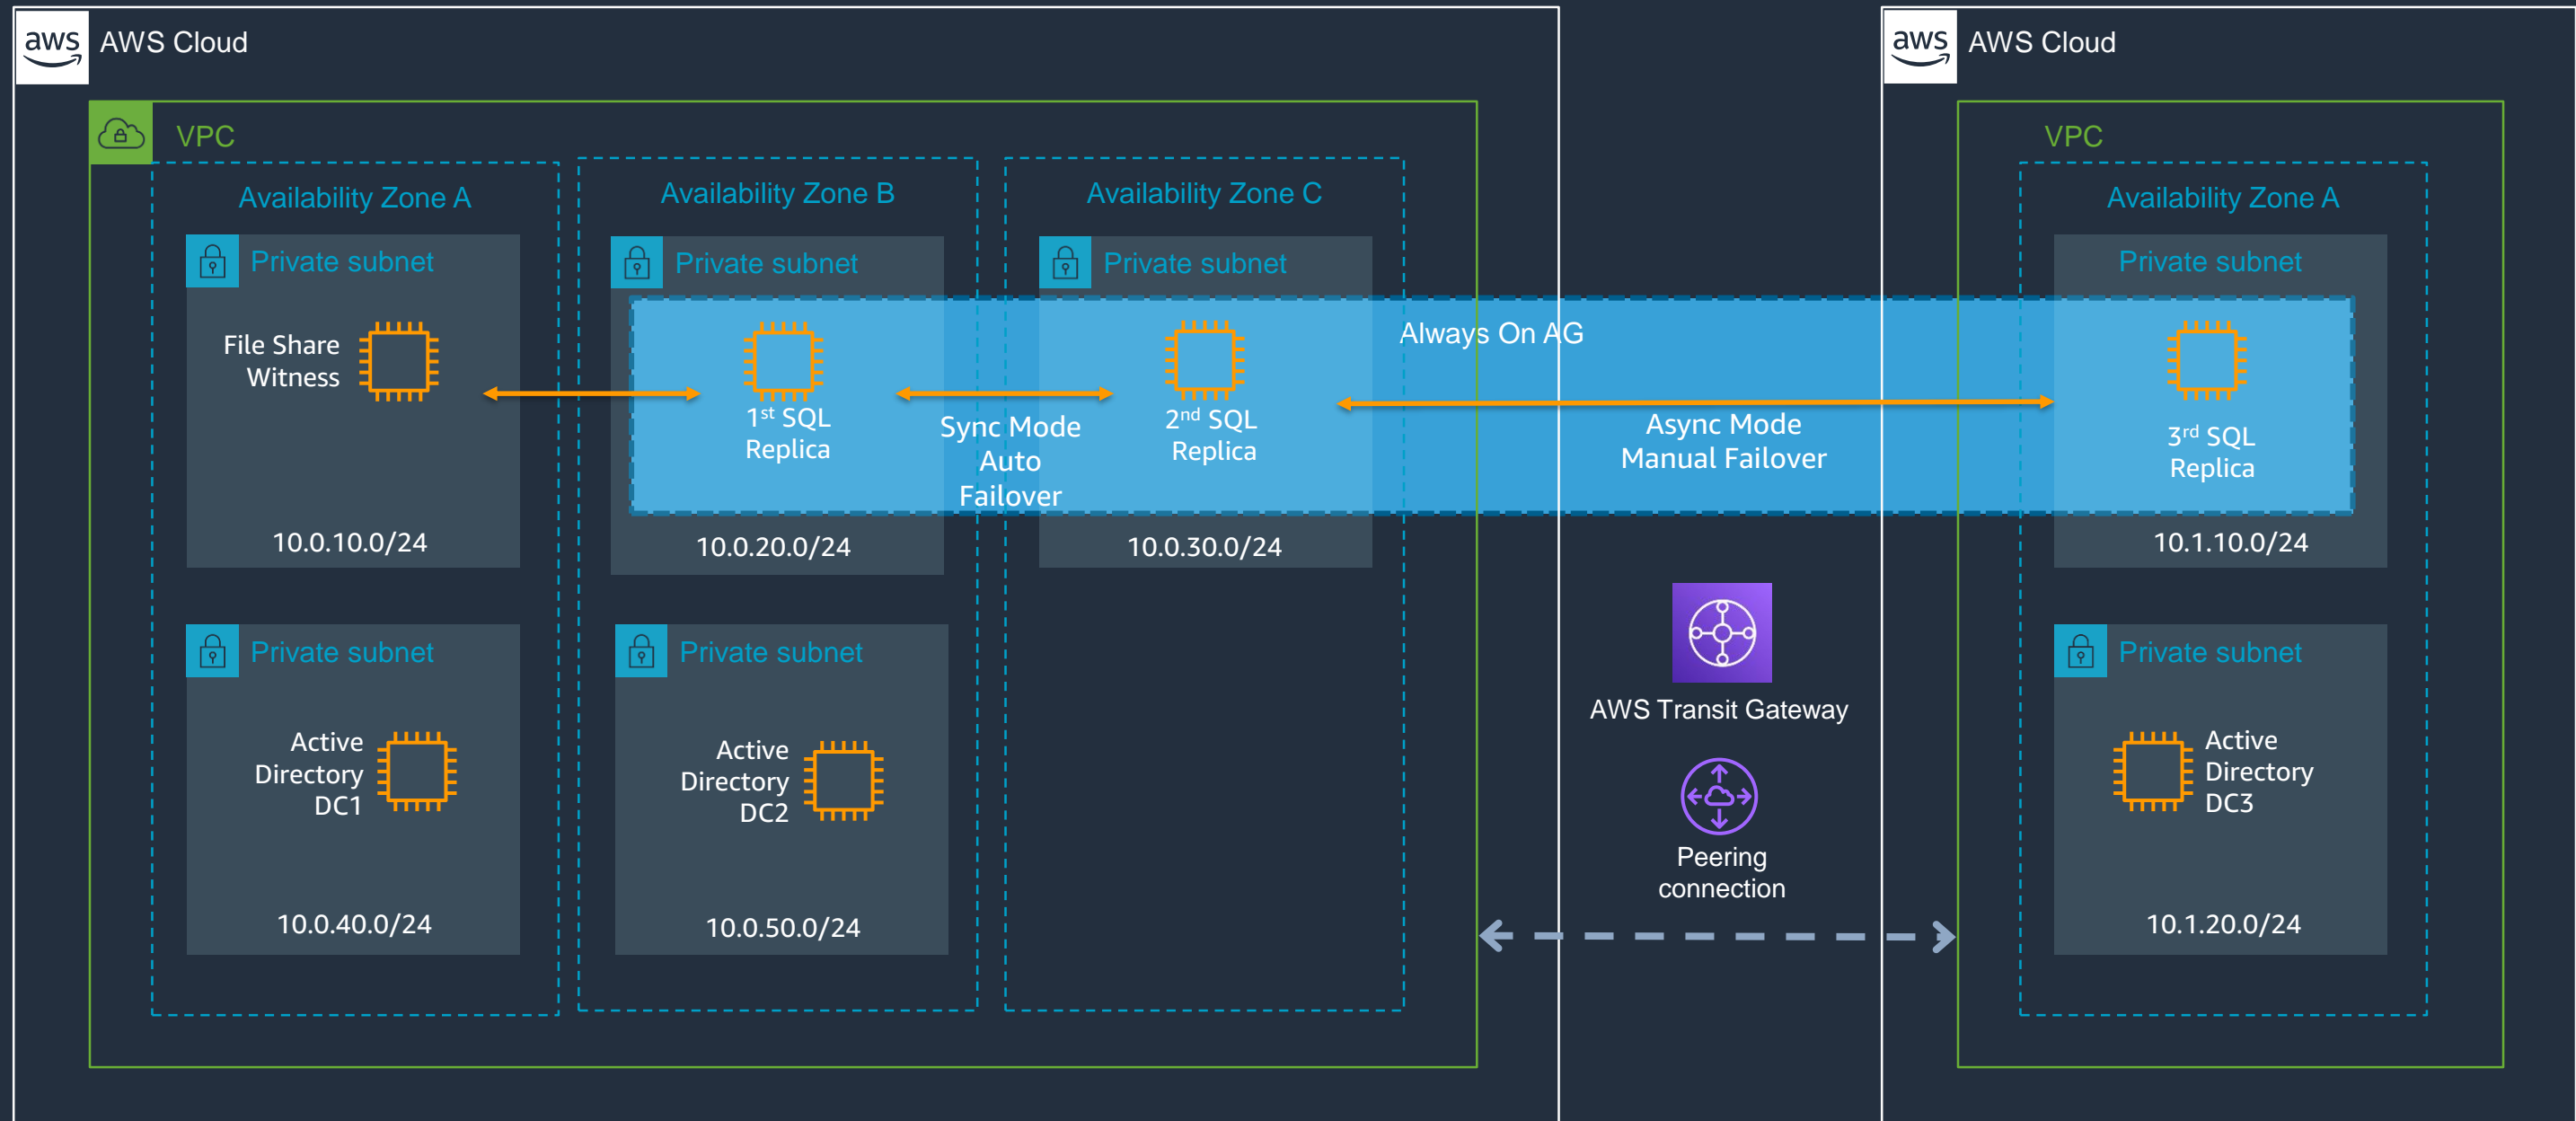


# Multi-AZ Always On Availability Group

- Best Practices:
  - Multi AZ setup
  - Use File Share Witness
  - MultiSubnetFailover = Yes
  - Connect to Listener
  - RegisterAllProvidersIP = Yes
  - Change default HostRecordTTL
  - SQL Native Client Providers

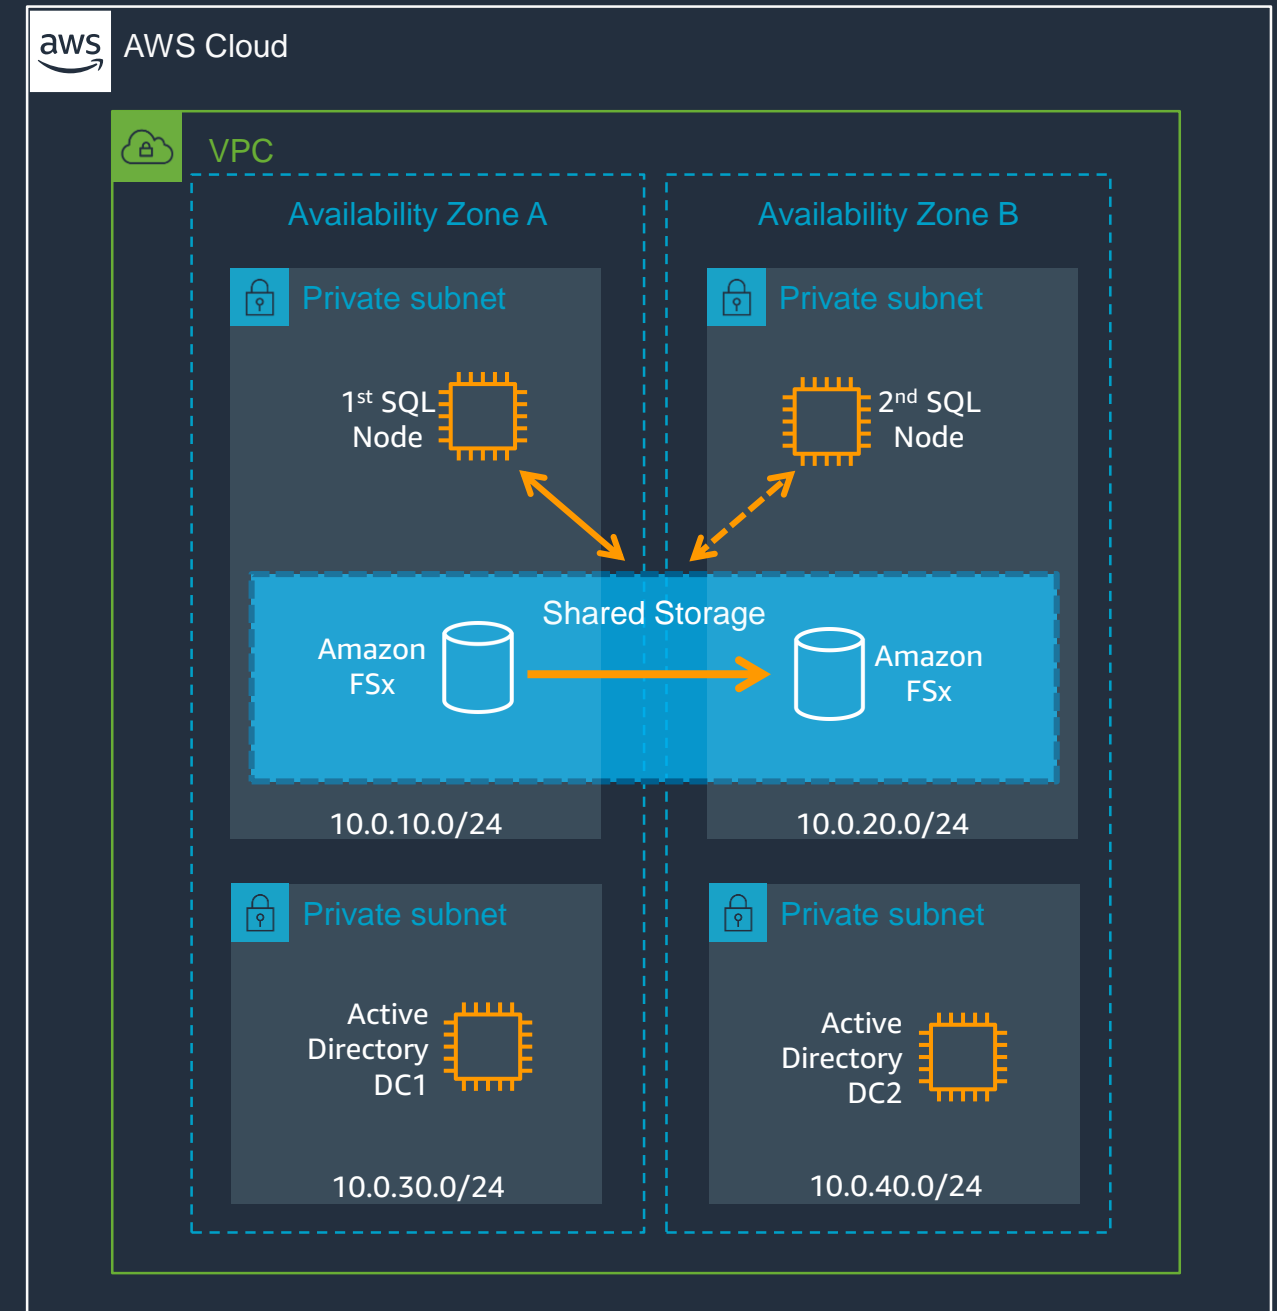


# Multi-Region Always On Availability Group



# SQL Failover Cluster Instances and Amazon FSx

- Infrastructure:
  - Clustered disk resource via SMB and continuously available shares (CAFS)
  - AWS managed SMB file system
  - Storage and Network traffic using the same network pipe
  - Managed storage configuration and maintenance
  - Very limited control over file system configuration (e.g. 4k vs. 64k allocation unit)
- SQL Server:
  - Instance-level HA solution
  - Logins, SQL Agent jobs, certificates and other SQL Server instance level objects are in-tact after failover




# Migrating SQL Server to EC2



# Migrating Data to & from Amazon EC2

- 1 Backup & Restore**  
SQL Server's native backup functionality
- 2 SQL Always On Availability Group**  
Asynchronous-commit mode, distributed AGs (SQL Server 2016+)
- 3 AWS Database Migration Service**  
Minimize downtime during migrations, migrate between different DB platforms



# Thank you

[migrations@amazon.com](mailto:migrations@amazon.com)

