PUBLIC SECTOR SUMMIT ONLINE

Migrating Microsoft SQL Server to the AWS Cloud

Bill Jacobi
Principal Solutions Architect
Amazon Web Services

aws PUBLIC SECTOR SUMMIT ONLINE

Agenda

Windows Workloads

Microsoft SQL Server on AWS

Best practices for running SQL

Migration methods

AWS Database Migration Service (AWS DMS)

Selecting the migration method

Innovation and experience

12 YEARS OF INNOVATING FOR OUR CUSTOMERS

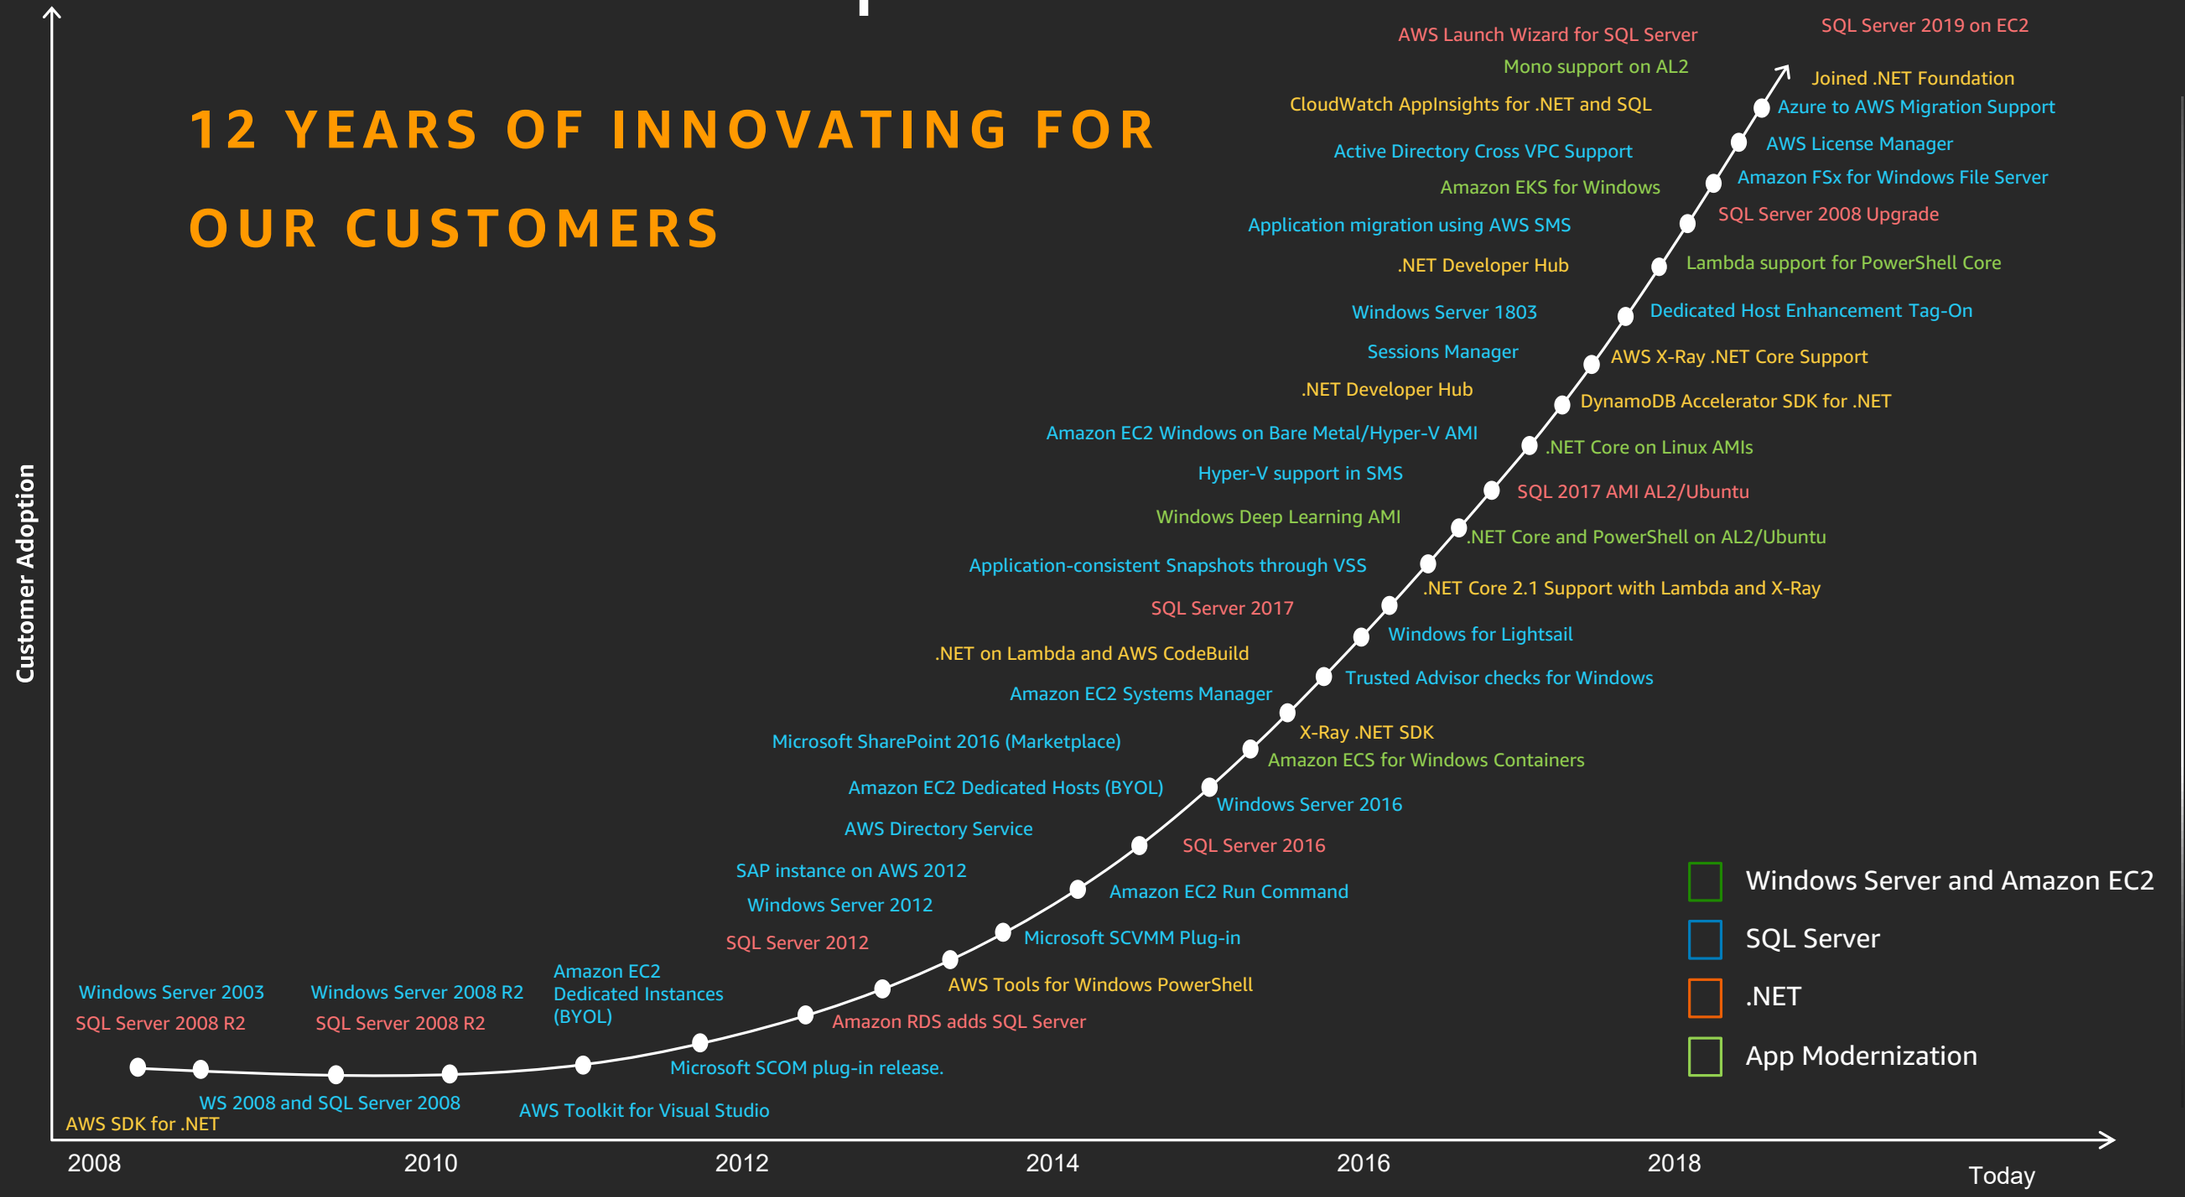

12+

Years of experience

https://zkresearch.com/blog/2018/11/comparing-sql-server-deployments-on-microsoft-azure-and-amazon-web-services

SQL Server on AWS

Which migration strategy is right for you?

Amazon EC2

Amazon RDS

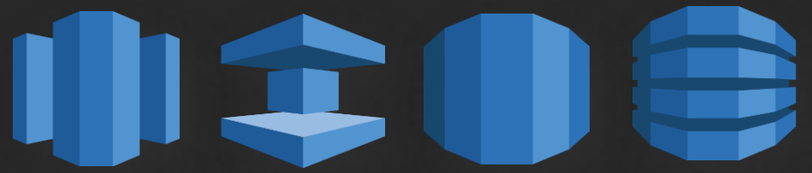

Rehost:

SQL Server on Amazon EC2

- Familiar administration experience
- Full control over the environment
- All SQL Server features available
- All SQL Server versions supported

Replatform:

SQL Server on Amazon RDS

- Optimized architecture
- Automated patching
- Automated backups
- Proven high availability

Refactor:

Adopt Cloud Native Services

- Amazon Aurora – SQL/OLTP
- Amazon Redshift – SQL/OLAP
- Amazon DynamoDB – NoSQL
- Amazon Neptune – Graph
- Eliminate SQL Server licensing costs

Migrating SQL Server databases to AWS

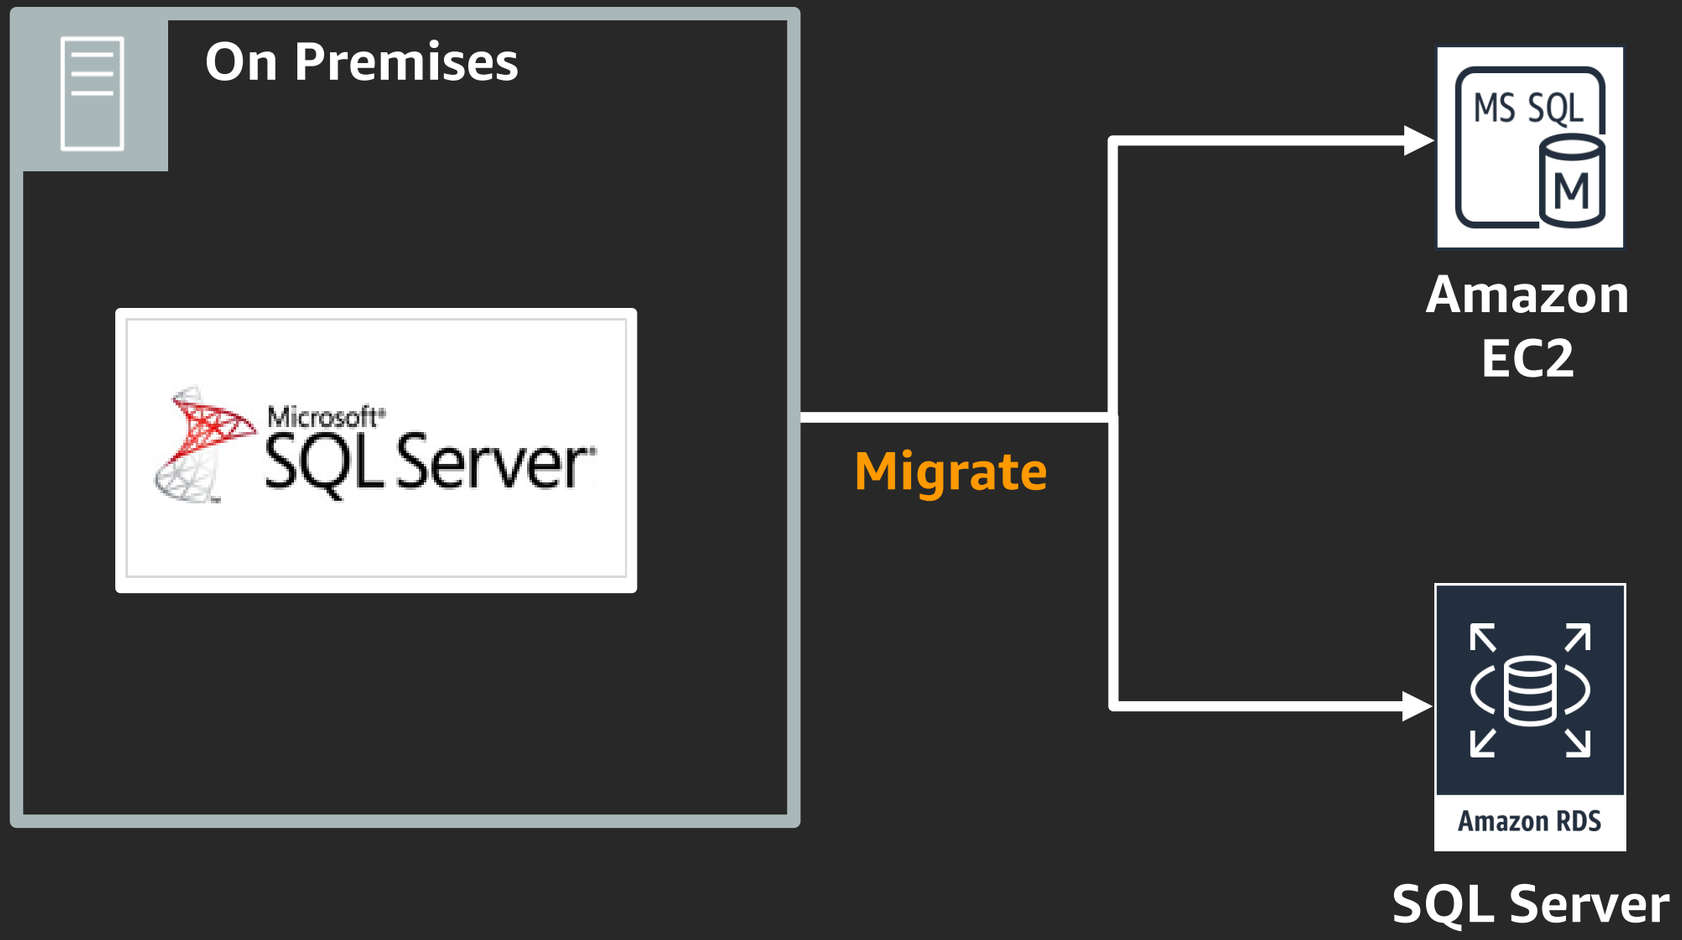

SQL Server on AWS

Amazon RDS for SQL Server

- **Consider Amazon RDS first**
- Focus on business value tasks
- High-level tuning asks
- Schema optimization
- No in-house database expertise

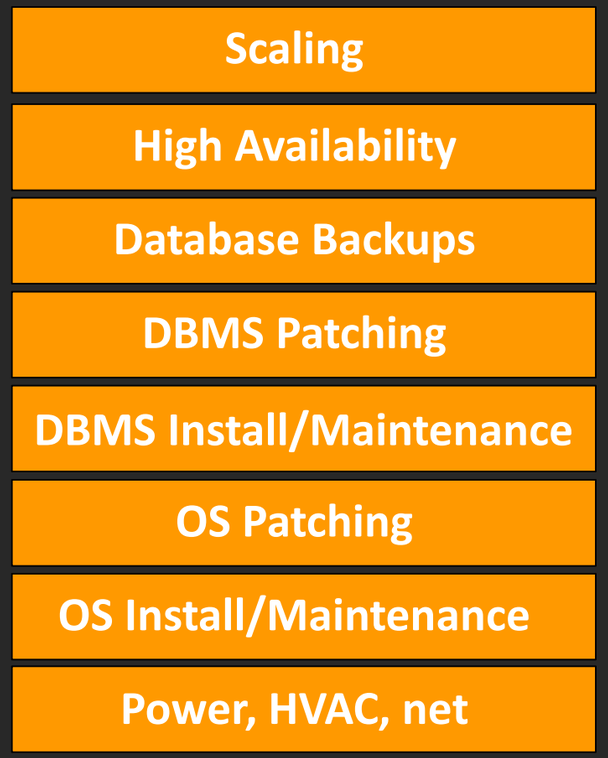

SQL Server on Amazon EC2

- Need full control over DB instance
- Backups
- Replication
- Clustering
- Options that are not available in RDS

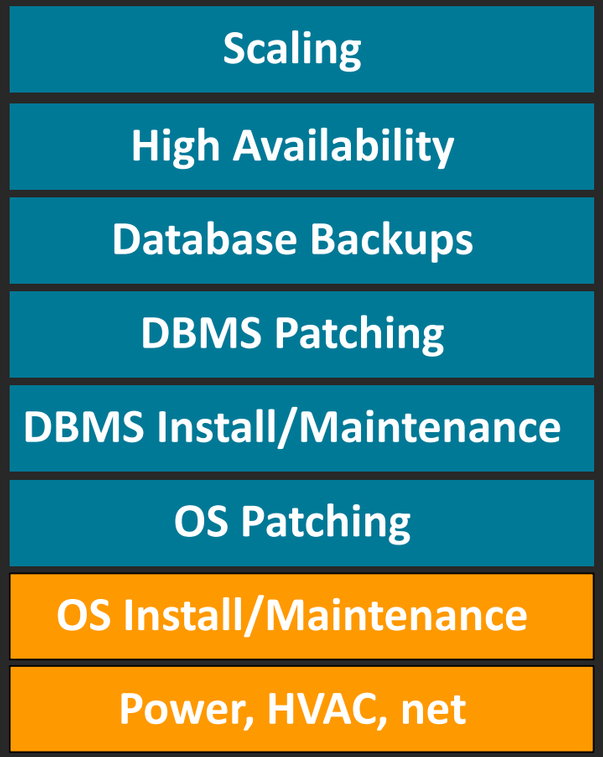

SQL Server features at a glance

Amazon RDS

Amazon EC2

*** Self-installed**

Versions Supported:	~~2008 R2~~, 2012, 2014, 2016, 2017	All
Editions Supported:	Express, web, standard, enterprise	
High Availability:	AWS managed	Self managed; always on, mirror, log ship
Encryption:	Encrypted storage using AWS KMS (all editions); TDE support	
Authentication:	Windows and SQL authentication	
Backups:	Managed automated backups	Maintenance plans and third-party tools
Maintenance:	Automatic software patching	Self managed

SQL Server as a managed service

AMAZON RDS

- Same SQL Server DB engine as with Amazon EC2
- Management, monitoring, and automation layer around the DB engine
- Automated full DB instance backups, with point-in-time restore
- Automated high availability (HA)
- Automated provisioning, patching, monitoring, directory integration
- Support for SSRS and SSAS

LIMITATIONS

- Cannot run SSIS on the DB instance (works as data source) -
- No sysadmin role, server administrator, or direct file system access
  - https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/UsingWithRDS.MasterAccounts.html
- Not supported: MSDTC, maintenance plans, database mail

Storage performance planning

AMAZON RDS STORAGE

- Low latency, persistent, network-attached block storage
- Easy to change after initial selection
- Maximum storage: **16 TB**
- Maximum IOPS: **64,000**
- Maximum throughput: 500 MiB/sec
- Amazon RDS storage throughput depends on DB instance class (see equivalent Amazon EC2, Amazon EBS optimized instance type)
- Keep in mind this includes TempDB

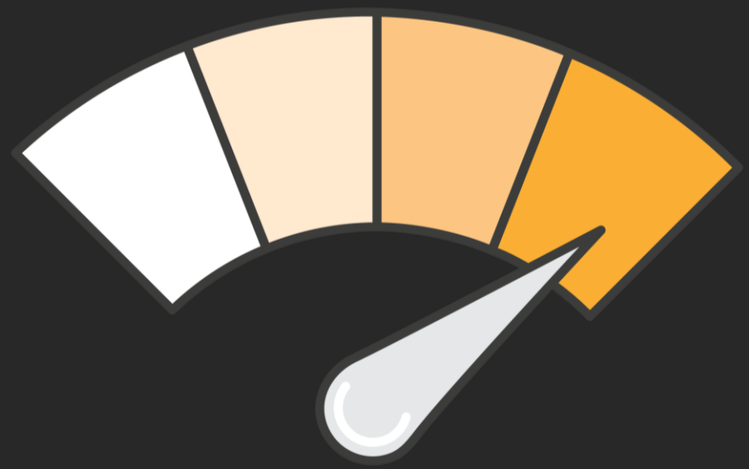

MONITORING I/O EFFICIENCY

- CloudWatch metric a**verage queue depth**
  - I/O requests waiting to be serviced

Amazon RDS SQL Server – performance Insights

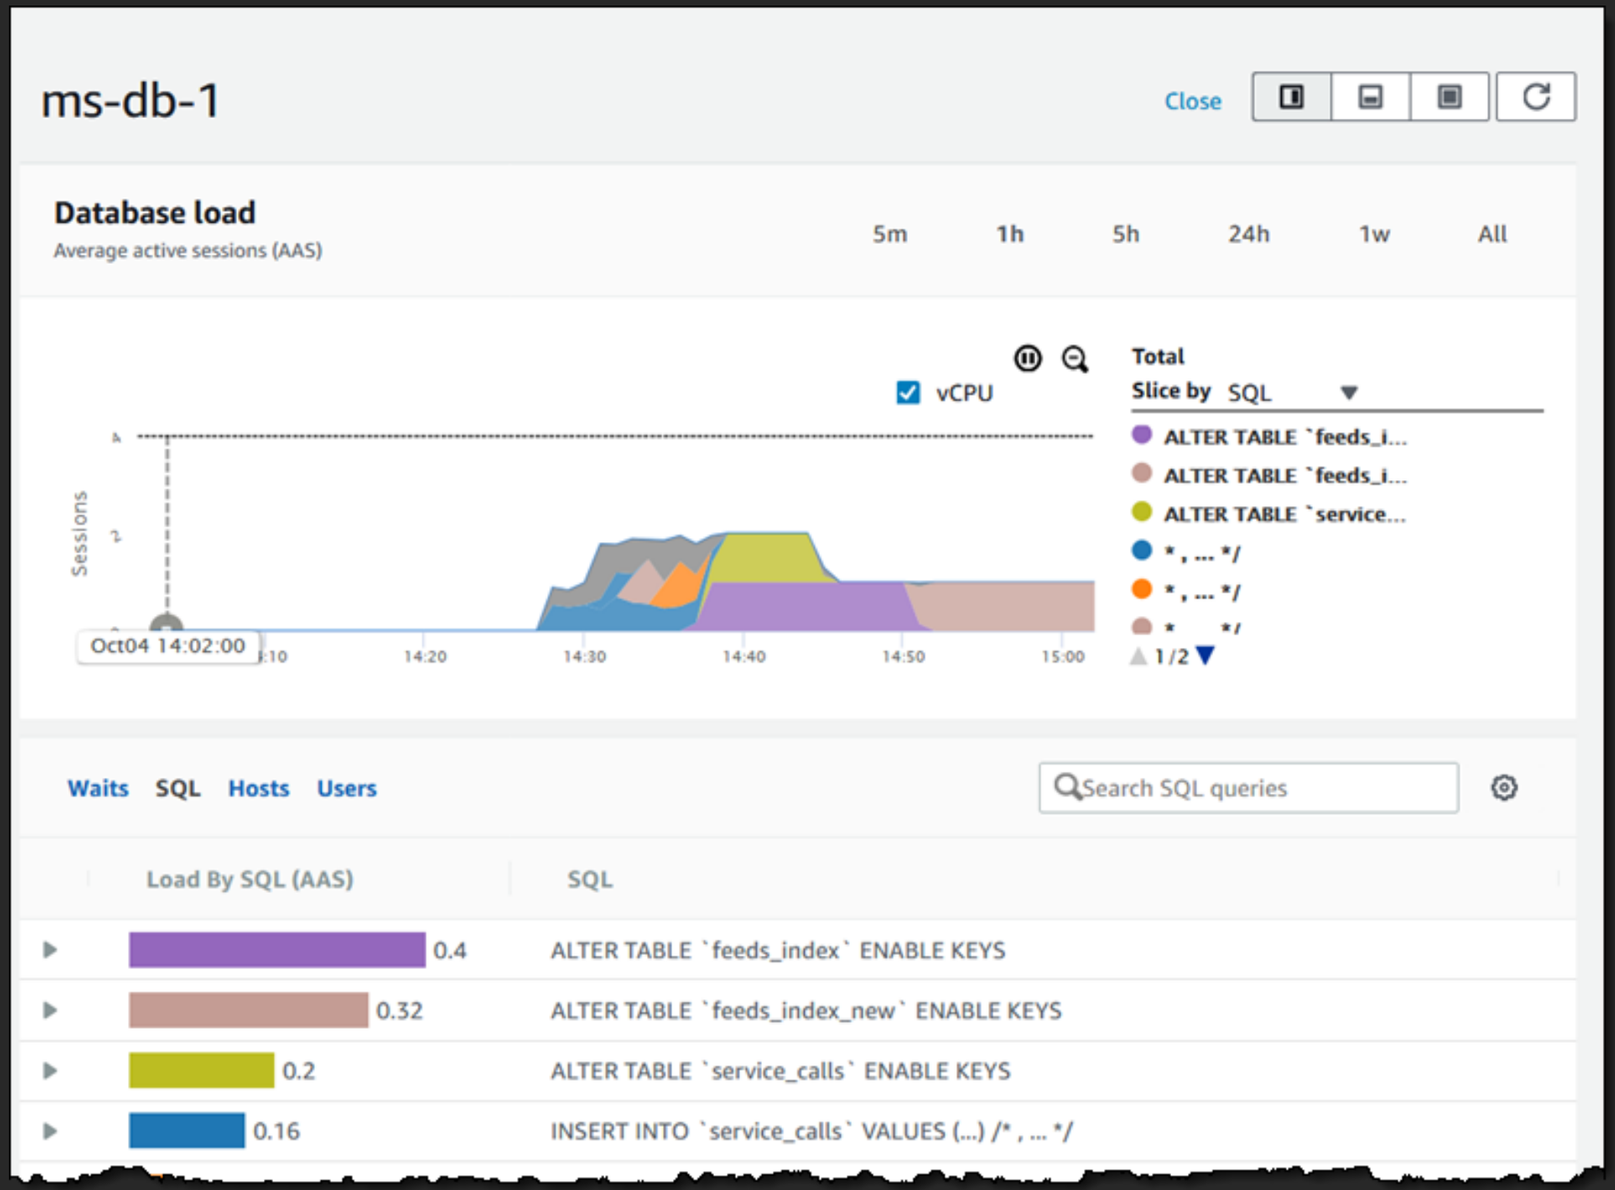

SQL Server on Amazon EC2

Amazon EC2: Purpose-built compute families

Current Instance Families and Generation	Family/Usage
M5, M4	General purpose compute
T2, T3	Burstable performance
C5, C4	Compute optimized
X1, X1E, R5, R5d, R4, R3	Memory optimized
P2, G3, F1	Accelerated computing
I3	Storage optimized (I/O)
D2	Storage optimized (Density)

Example: r5d instance family

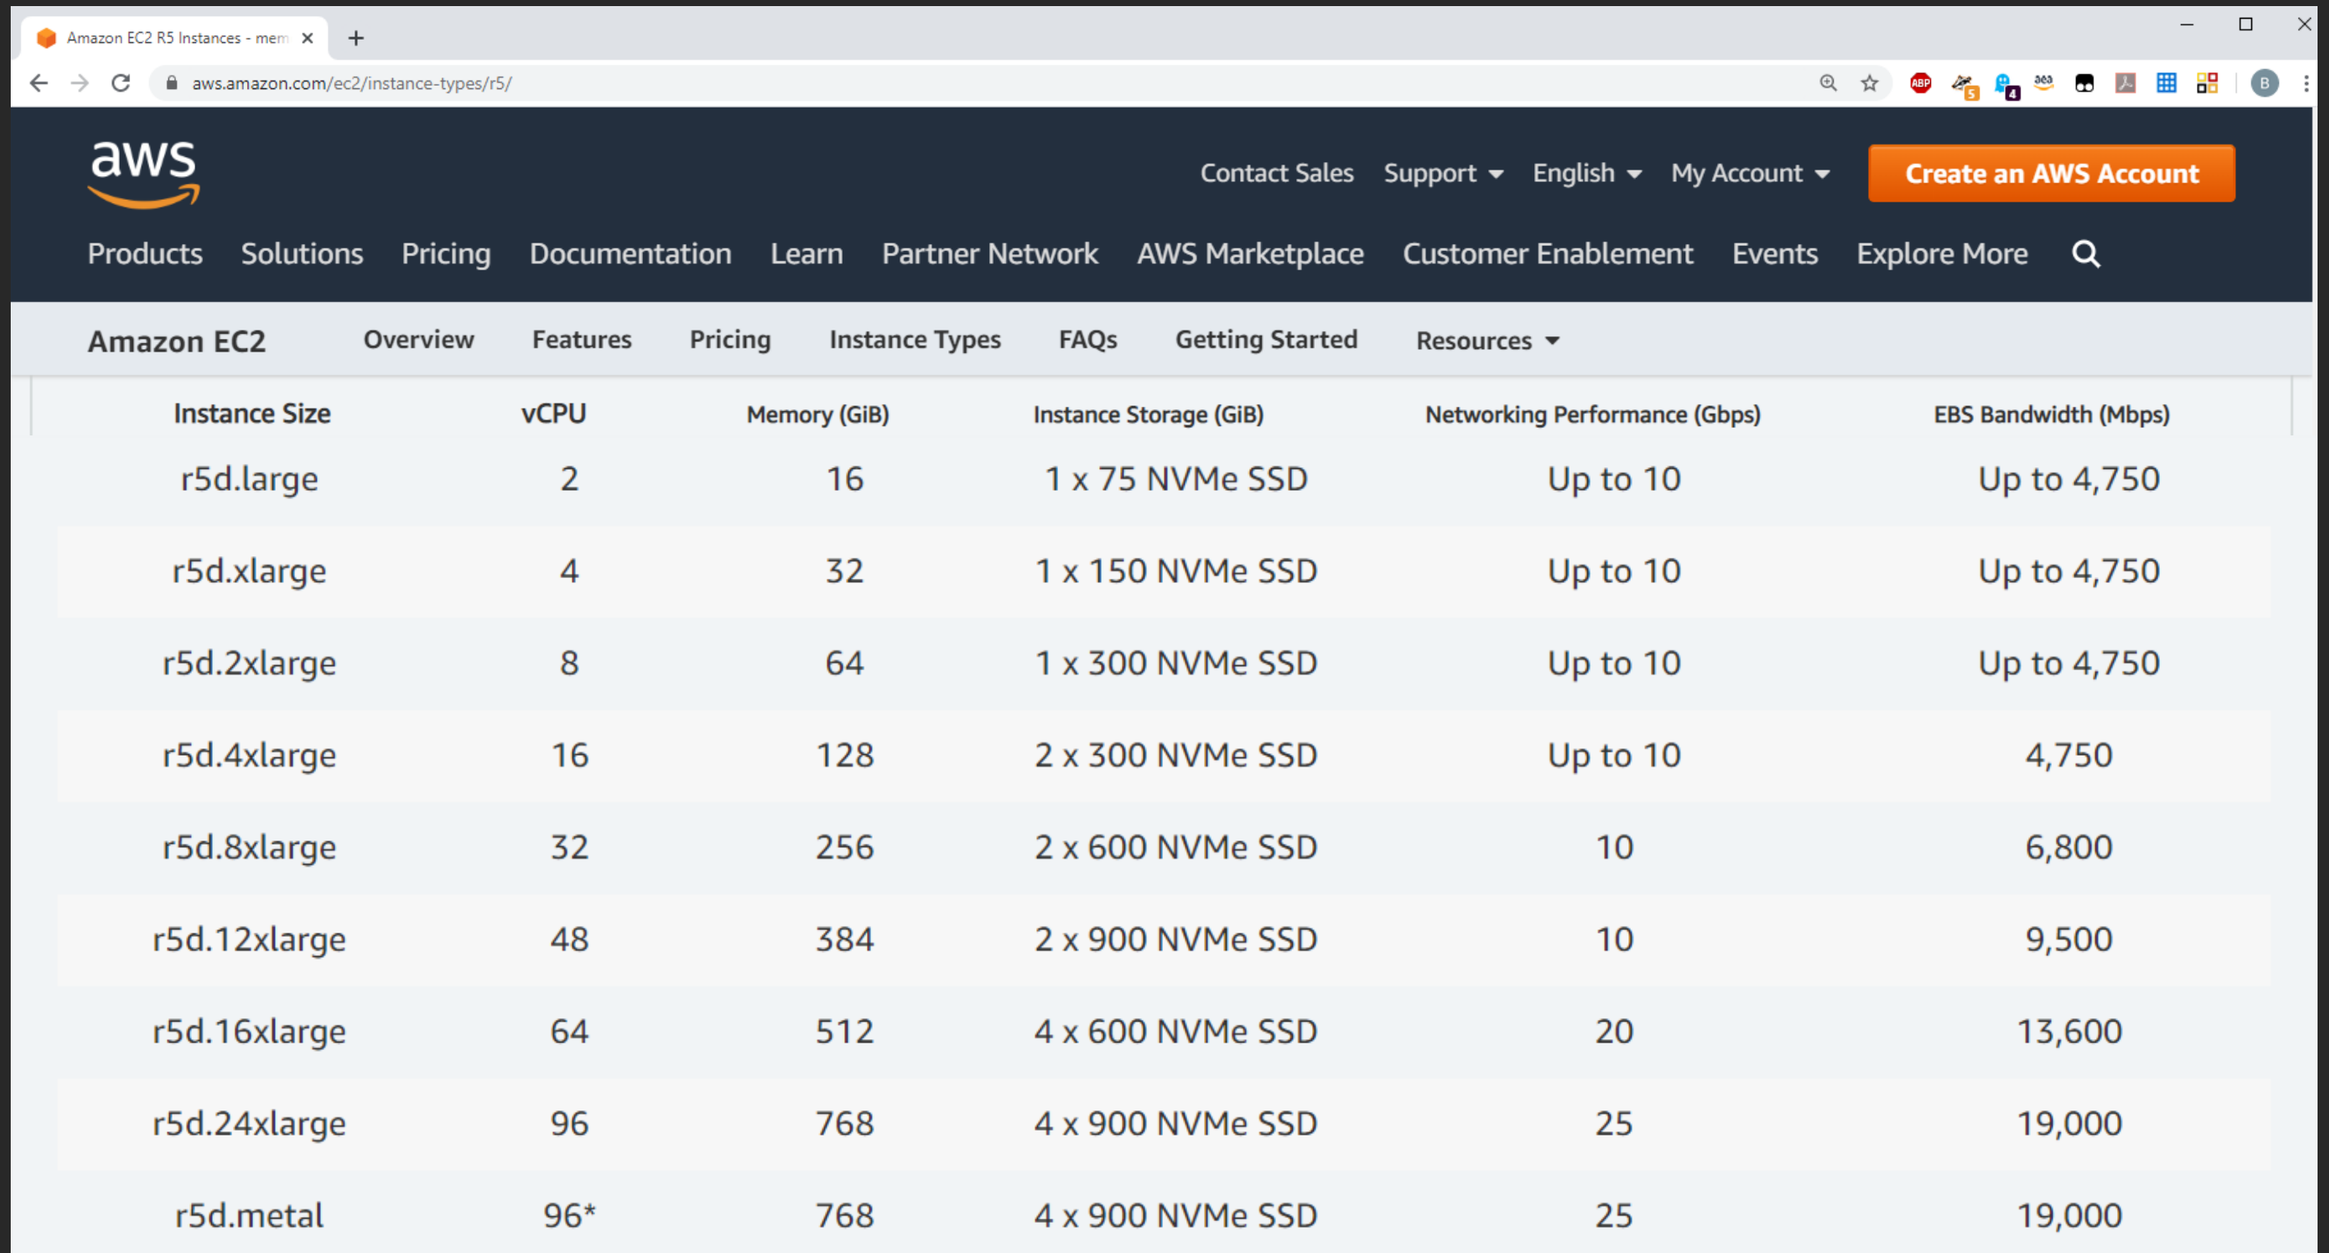

Instance Size	vCPU	Memory (GiB)	Instance Storage (GiB)	Networking Performance (Gbps)	EBS Bandwidth (Mbps)
r5d.large	2	16	1 x 75 NVMe SSD	Up to 10	Up to 4,750
r5d.xlarge	4	32	1 x 150 NVMe SSD	Up to 10	Up to 4,750
r5d.2xlarge	8	64	1 x 300 NVMe SSD	Up to 10	Up to 4,750
r5d.4xlarge	16	128	2 x 300 NVMe SSD	Up to 10	4,750
r5d.8xlarge	32	256	2 x 600 NVMe SSD	10	6,800
r5d.12xlarge	48	384	2 x 900 NVMe SSD	10	9,500
r5d.16xlarge	64	512	4 x 600 NVMe SSD	20	13,600
r5d.24xlarge	96	768	4 x 900 NVMe SSD	25	19,000
r5d.metal	96*	768	4 x 900 NVMe SSD	25	19,000

AWS Instance Types and Intel Technologies

14 years of partnership with over 200 instance types

AWS Instance Type	General Purpose			Compute Optimized		Memory Optimized					Accelerated Compute (GPU)			Storage Optimized (High I/O)	
	T3 (burstable)	M5	M5n	C5	C5n	R5	R5n	X1e / X1	High Memory	Z1d	P3	G4	F1	I3	I3en
Intel® Processor	Intel® Xeon® Scalable Processors	Intel® Xeon® Platinum 8175M Processors	Intel® Xeon® Scalable Processors	Intel® Xeon® Scalable Processors	Intel® Xeon® Platinum 8124M Processors	Intel® Xeon® Platinum 8175 Processors	Intel® Xeon® Scalable Processors	Intel® Xeon® E7 8880 v3 Processors	Intel® Xeon® Platinum 8176M or Scalable Processors	Intel® Xeon® Platinum 8151 Processors	Intel® Xeon® E5-2686 v4 or P-8175M Processors	Intel® Xeon® Scalable Processors	Intel® Xeon® E5-2686 v4 Processors	Intel® Xeon® E5-2686 v4 Processors	Intel® Xeon® Scalable Processors
Intel® Process Technology	Skylake	Skylake	Cascade Lake	Cascade Lake	Skylake	Skylake	Cascade Lake	Haswell	Skylake or Cascade Lake	Skylake	Broadwell or Skylake	Cascade Lake	Broadwell	Broadwell	Skylake
Intel® Advanced Vector Extensions	AVX-512	AVX-512	AVX-512	AVX-512	AVX-512	AVX-512	AVX-512	AVX2	AVX-512	AVX-512	AVX2 or AVX-512	AVX-512	AVX2	AVX2	AVX-512
Intel® AES New Instructions	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Intel® Turbo Boost	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Intel® Deep Learning Boost	-	-	Yes	Yes	-	-	Yes	-	Yes (18 and 24 TiB)	-	-	Yes	-	-	-

Amazon EBS volume types

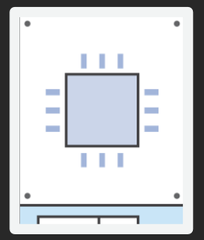

SSD

HDD

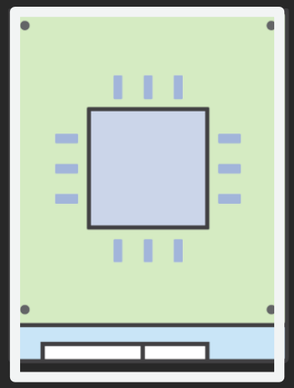

gp2

General purpose

$0.10 per GiB

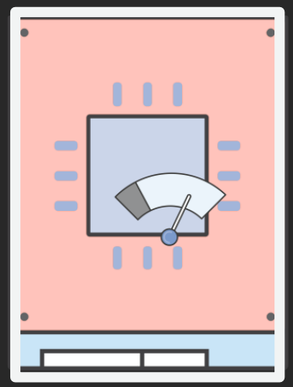

io1

Provisioned IOPS

$0.125 per GiB

$0.065 per PIOPS

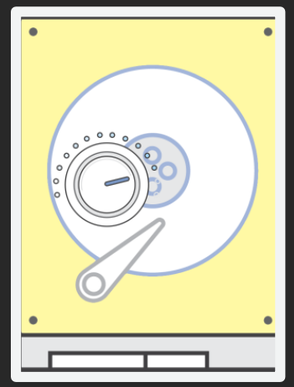

st1

Throughput optimized

$0.045 per GiB

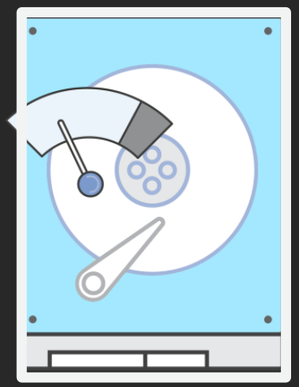

sc1

Cold

$0.025 per GiB

Snapshot storage for all volume types is $0.05 per GiB per month

* All prices are per month, prorated to the second, and from the us-east-1 region as of May, 2020

Migration methods

Assessment and planning

- Inventory SQL Server all dependencies
- Authentication requirements (e.g., Windows Authentication vs. SQL)
- Identify SQL Server version or edition features currently used
- Know you licensing options (e.g. Leverage BYOL)
- Understand high availability and disaster recovery requirements
- Performance requirements (e.g., IOPS) and capacity planning
- Leverage your retention policy
- Understand migration options
- List all database properties (e.g., recovery model/compatibility level)

Hybrid architecture

- Integration of on-premises resources with cloud resources
- Migrate SQL Server data to the AWS Cloud

https://aws.amazon.com/enterprise/hybrid/

SQL Server backups to Amazon S3

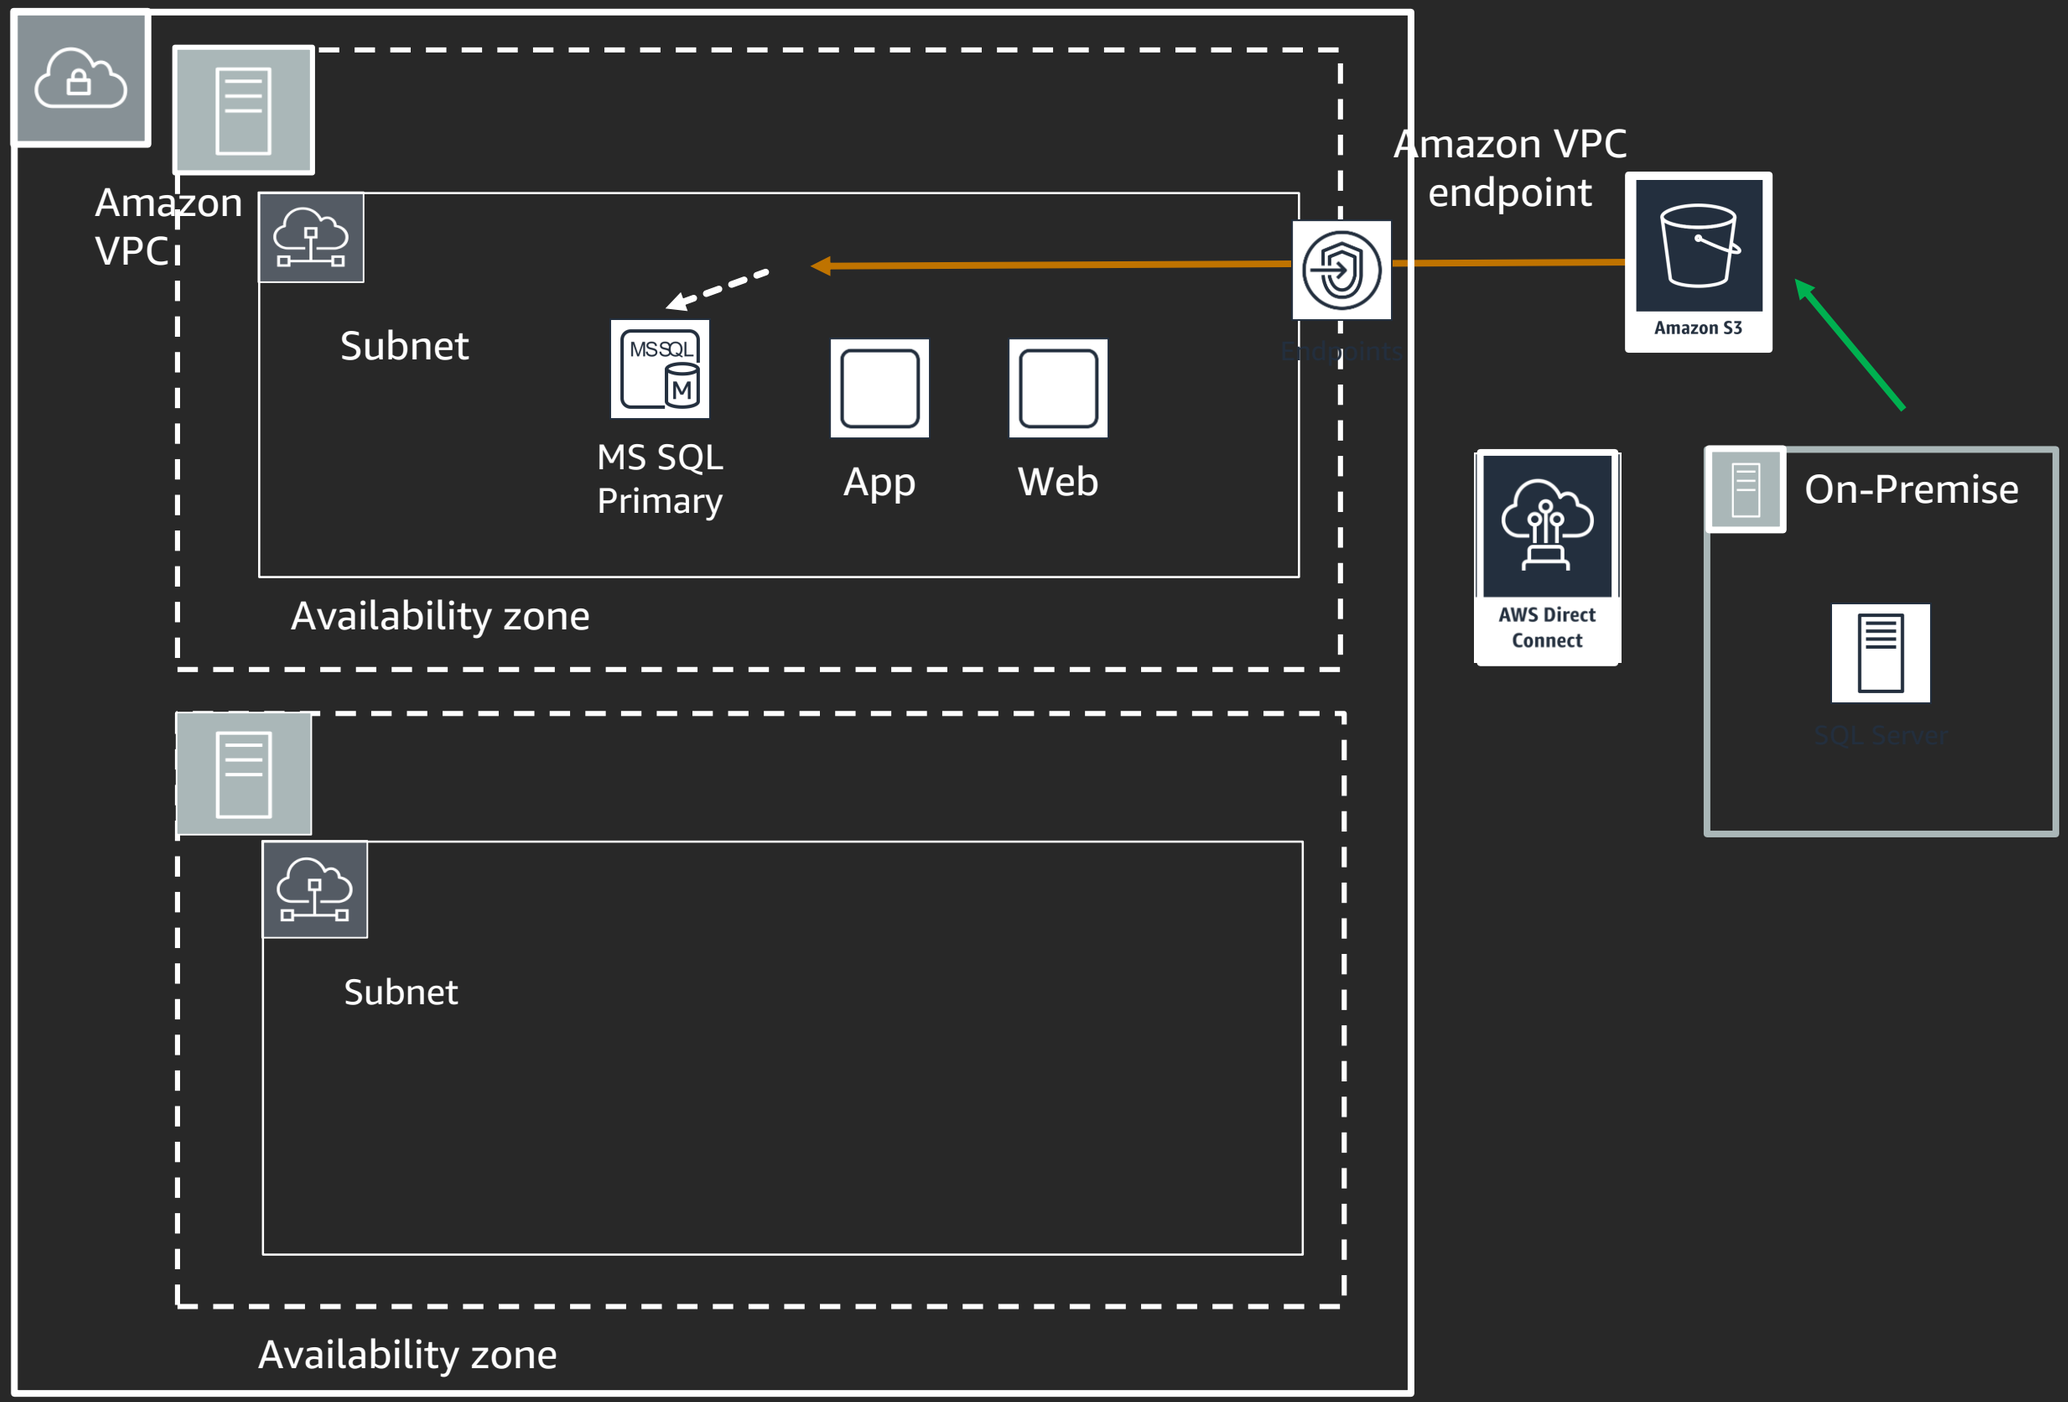

SQL Server backups to Amazon S3

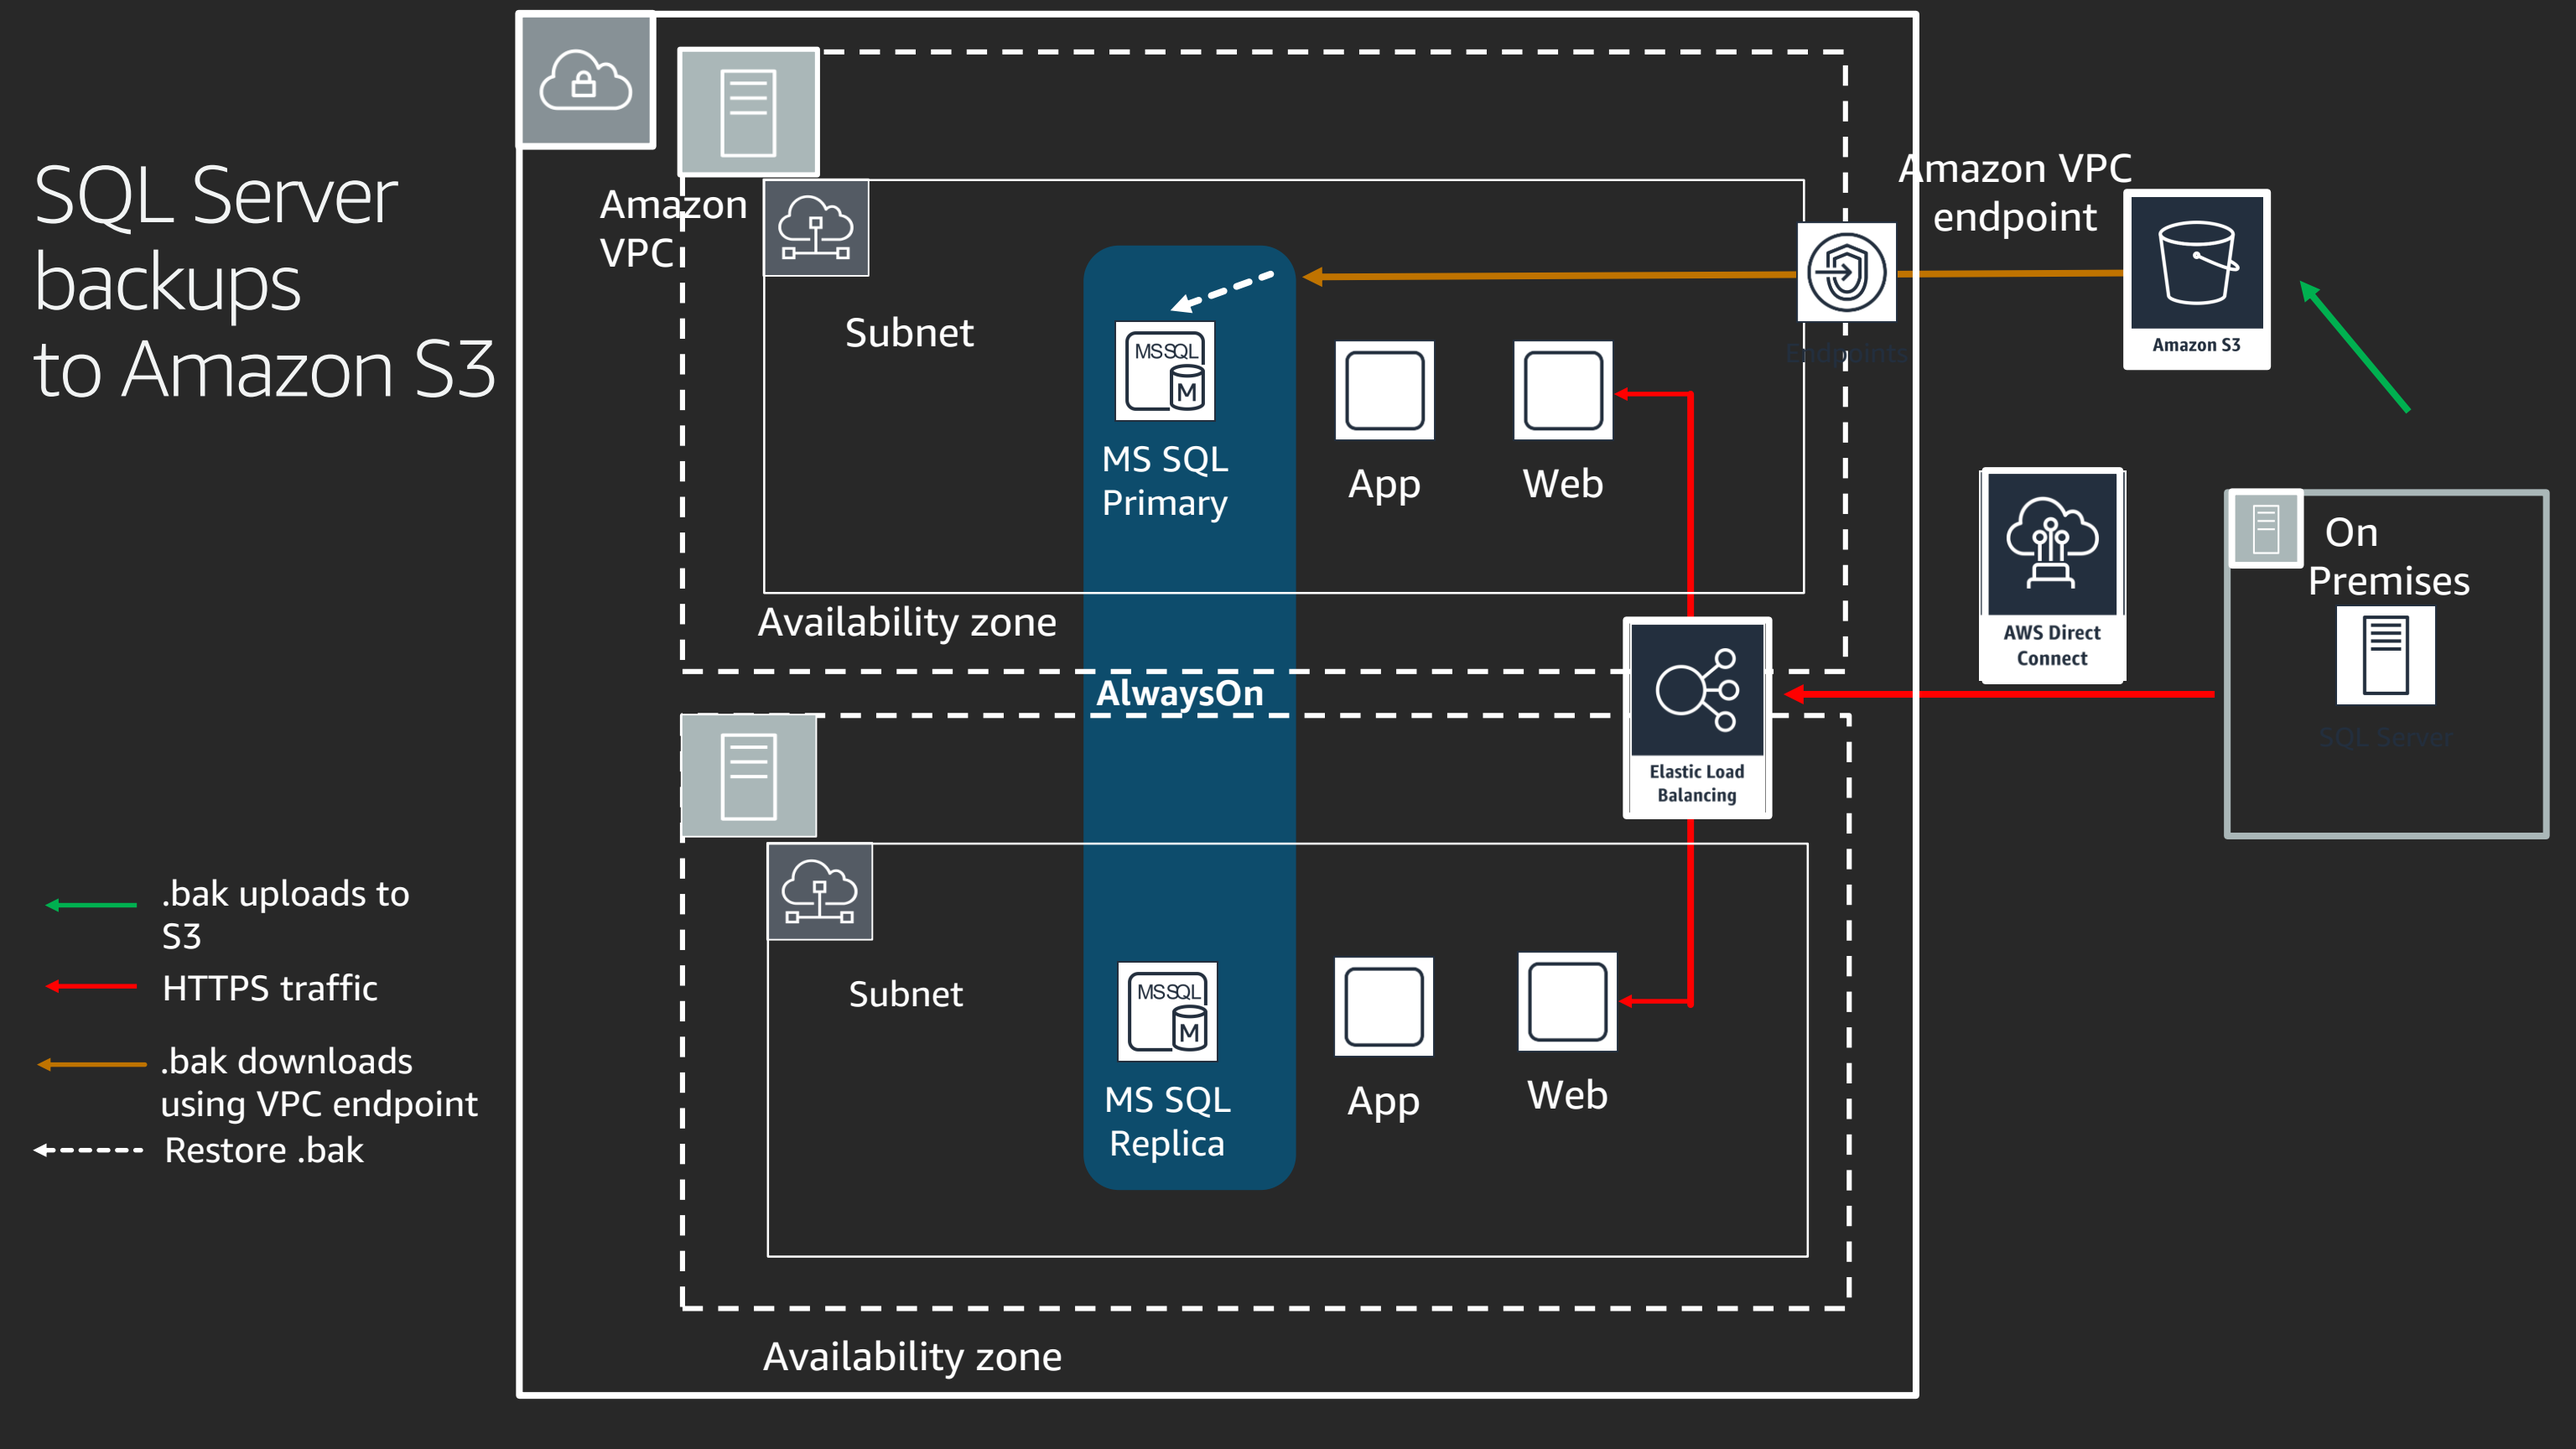

Native SQL backup to Amazon S3 via SMB

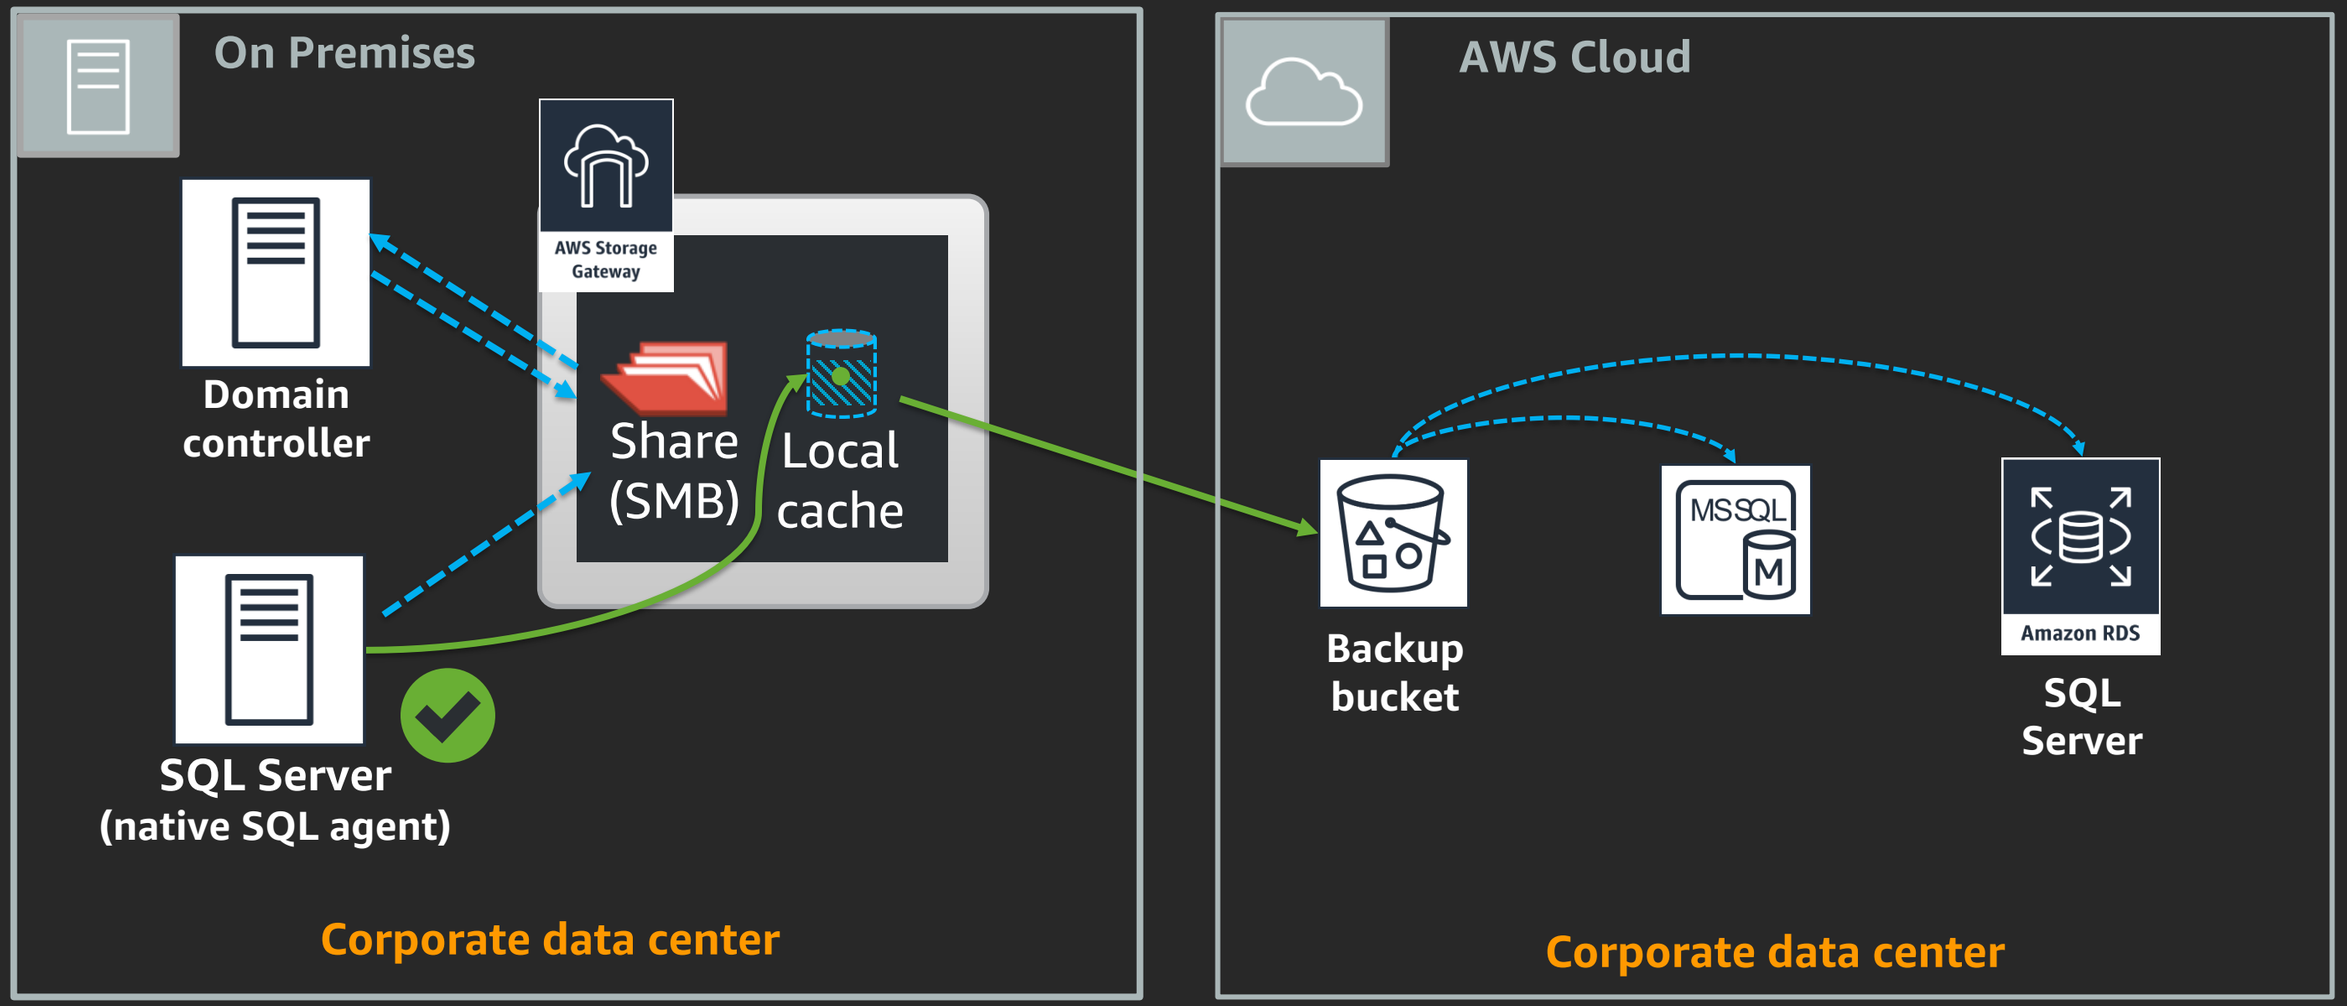

Amazon RDS SQL Server backup/restore

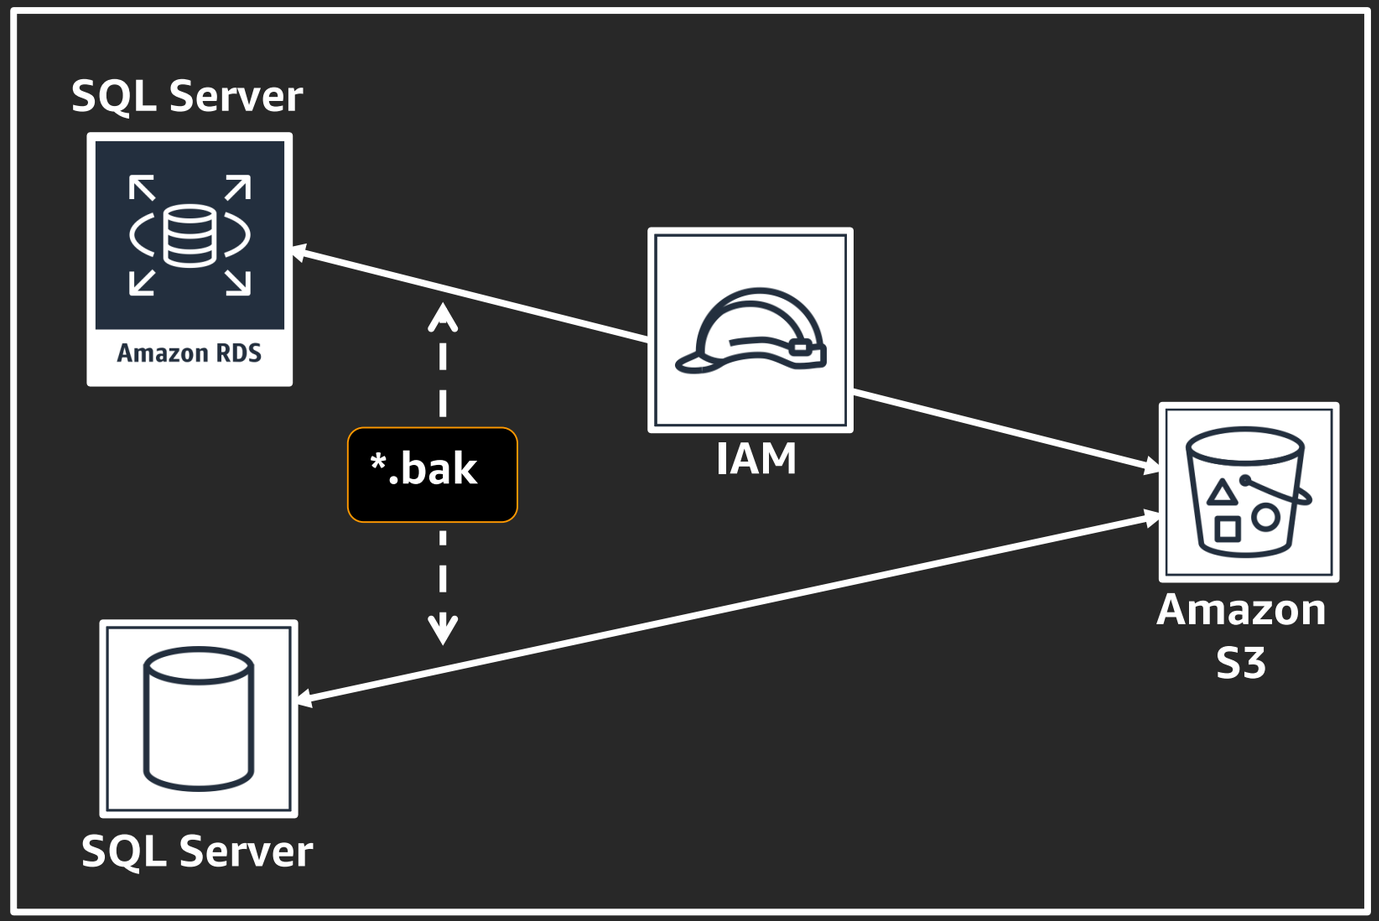

Amazon RDS SQL Server backup/restore

- IAM Role to connect services

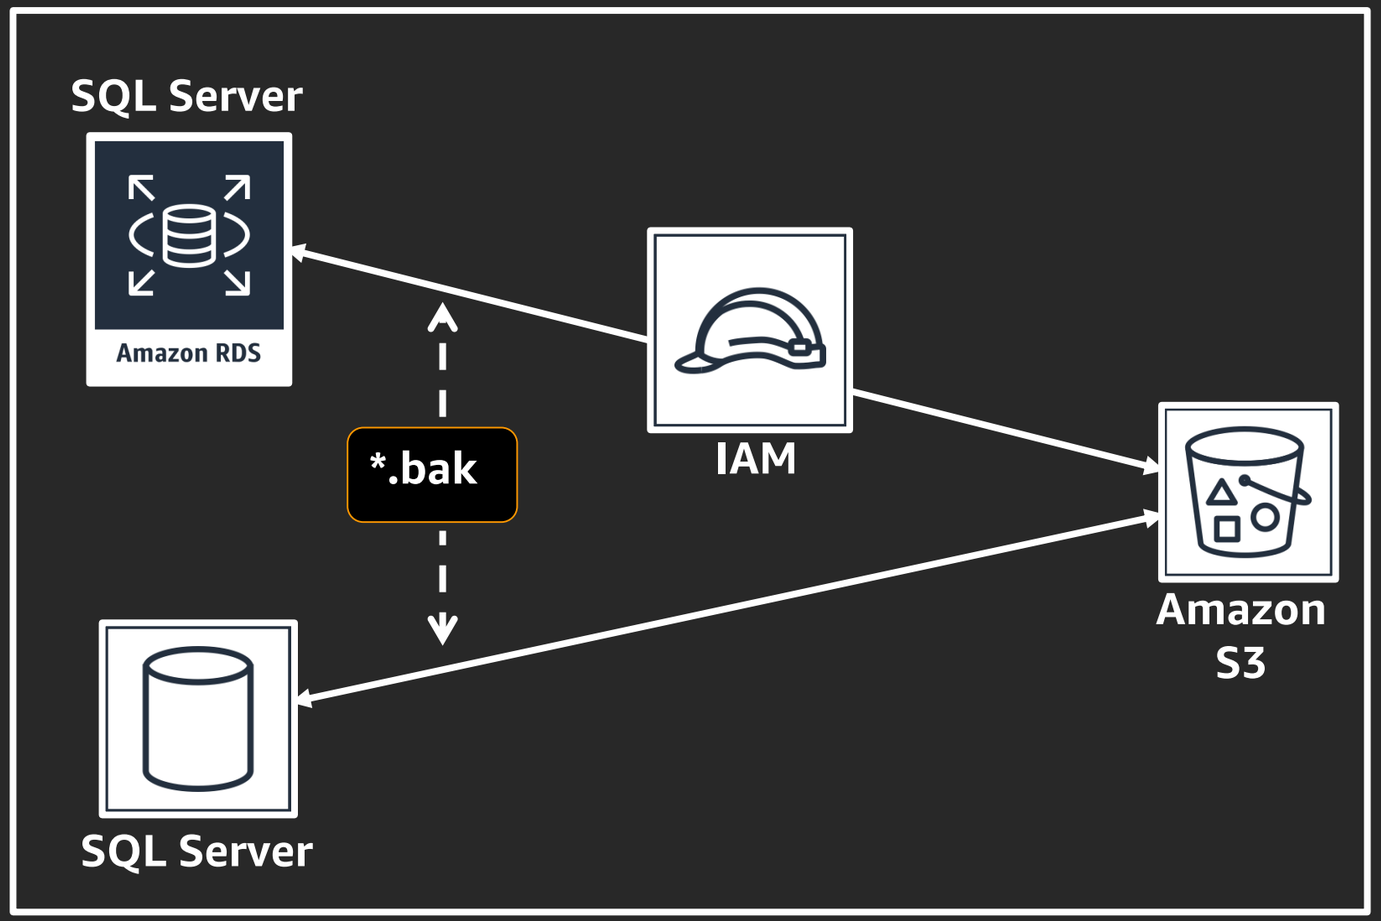

Amazon RDS SQL Server backup/restore

- IAM Role to connect services
- Configure option group to enable functionality

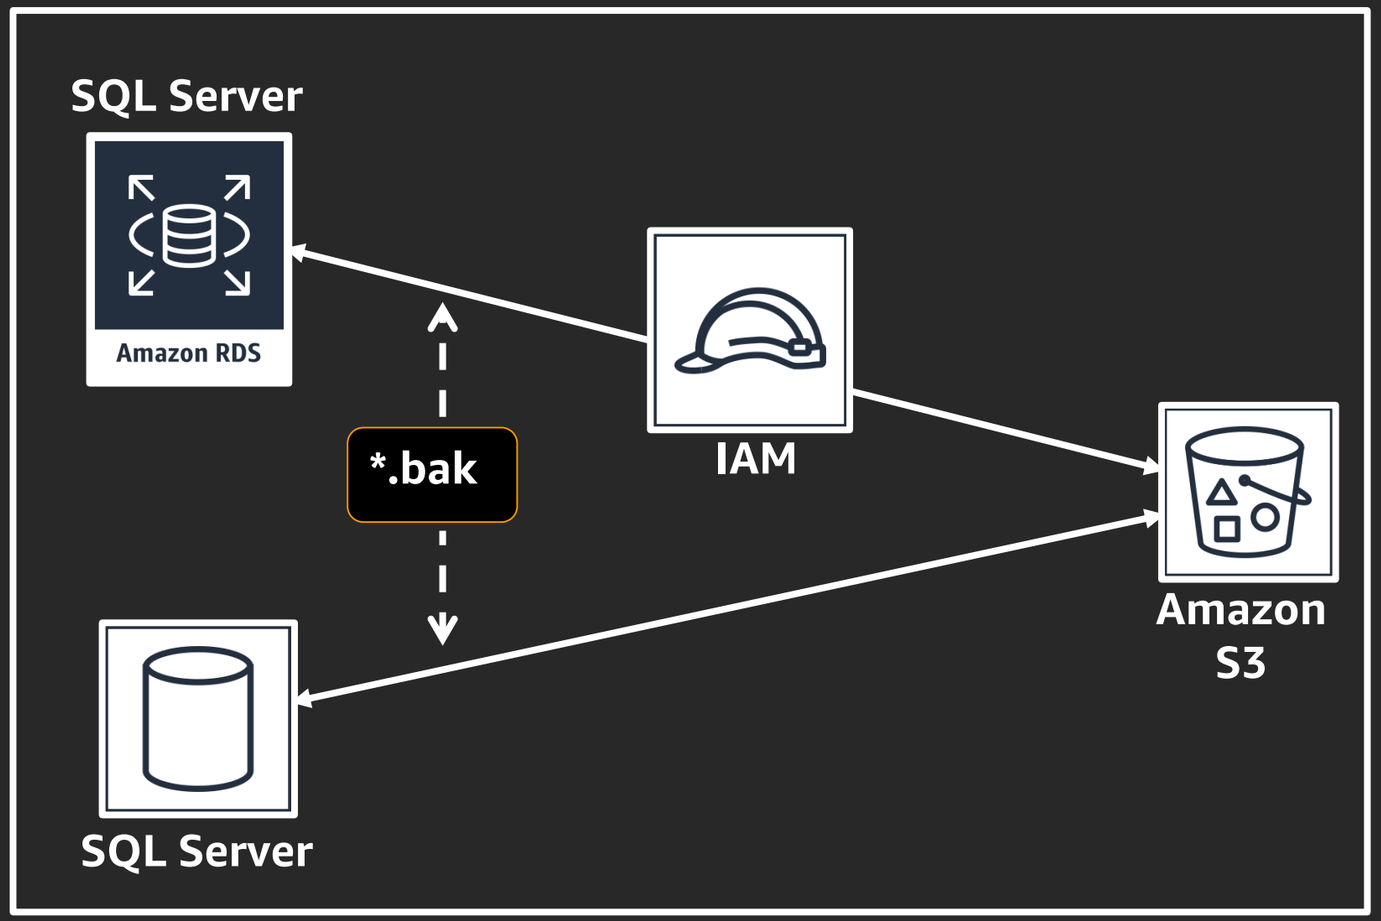

Amazon RDS SQL Server backup/restore

- IAM Role to connect services
- Configure option group to enable functionality
- Specify an Amazon S3 bucket as part of configuration

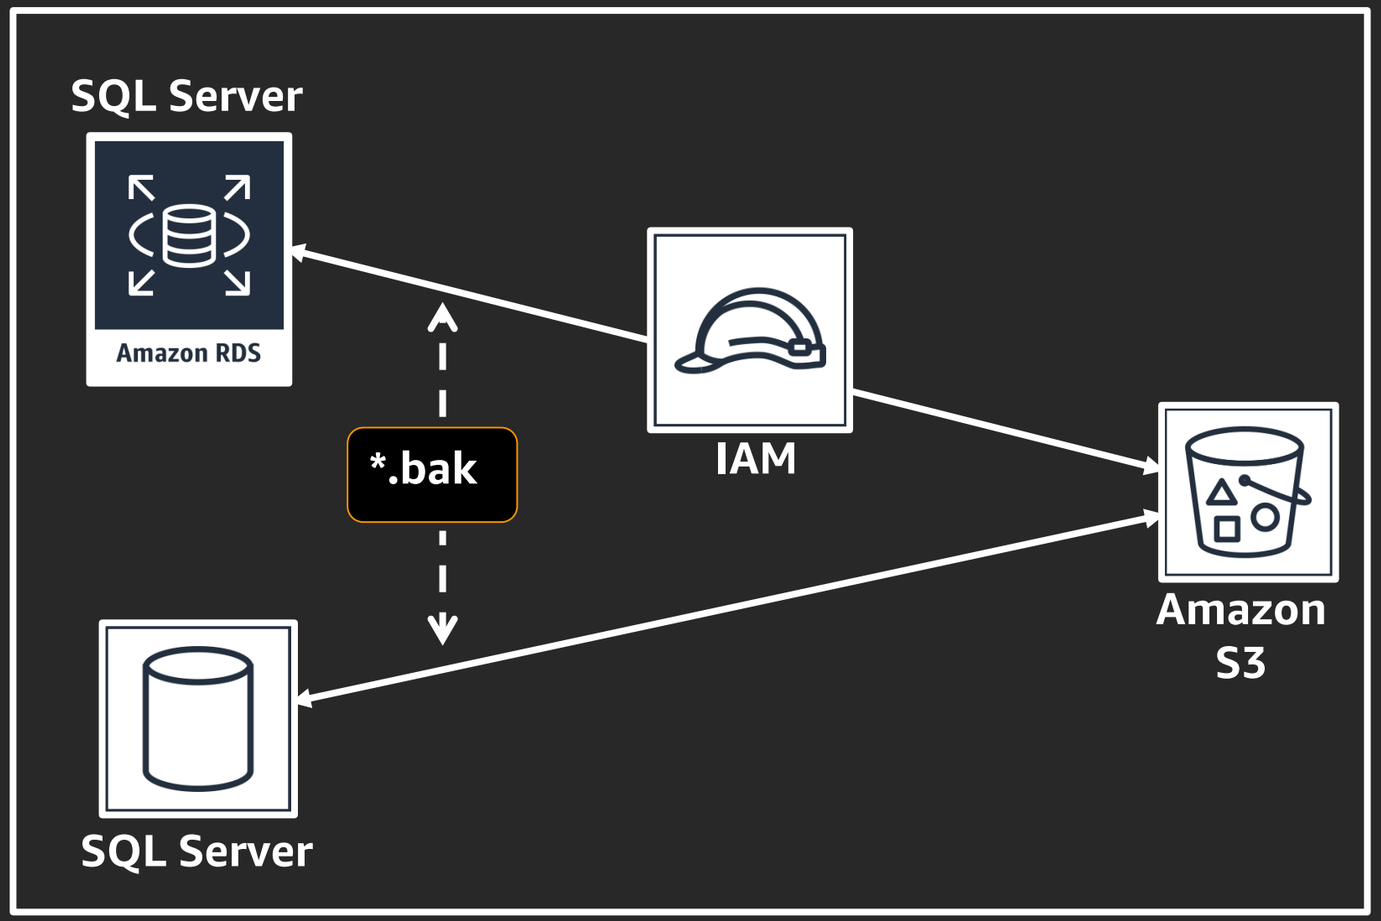

Amazon RDS SQL Server backup/restore

- IAM Role to connect services
- Configure option group to enable functionality
- Specify an Amazon S3 bucket as part of configuration
- Run stored procedure to perform restore

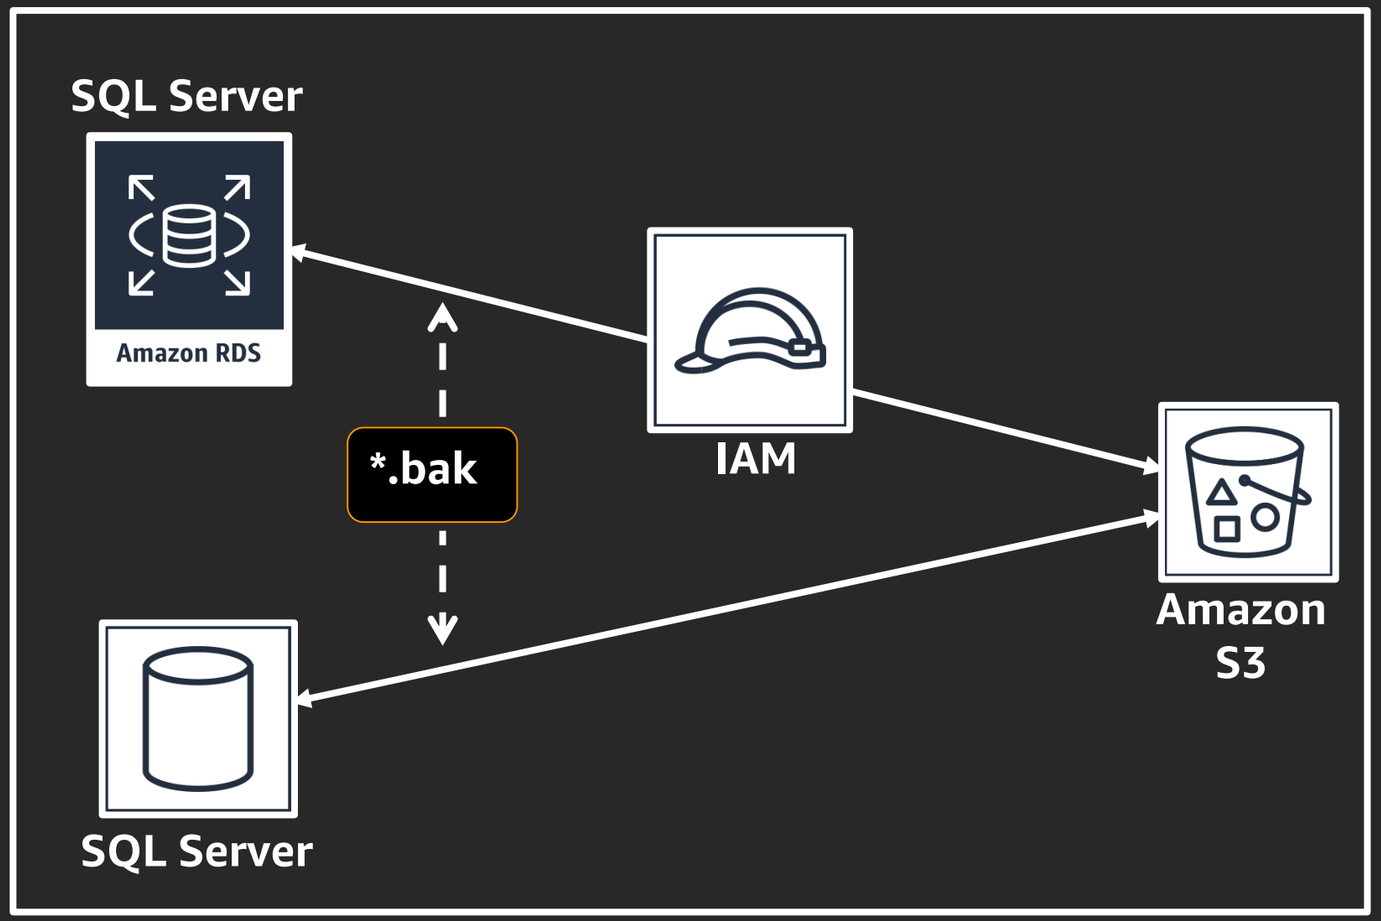

Amazon RDS SQL Server backup/restore

- IAM Role to connect services
- Configure option group to enable functionality
- Specify an Amazon S3 bucket as part of configuration
- Run stored procedure to perform restore
- Heavily optimized

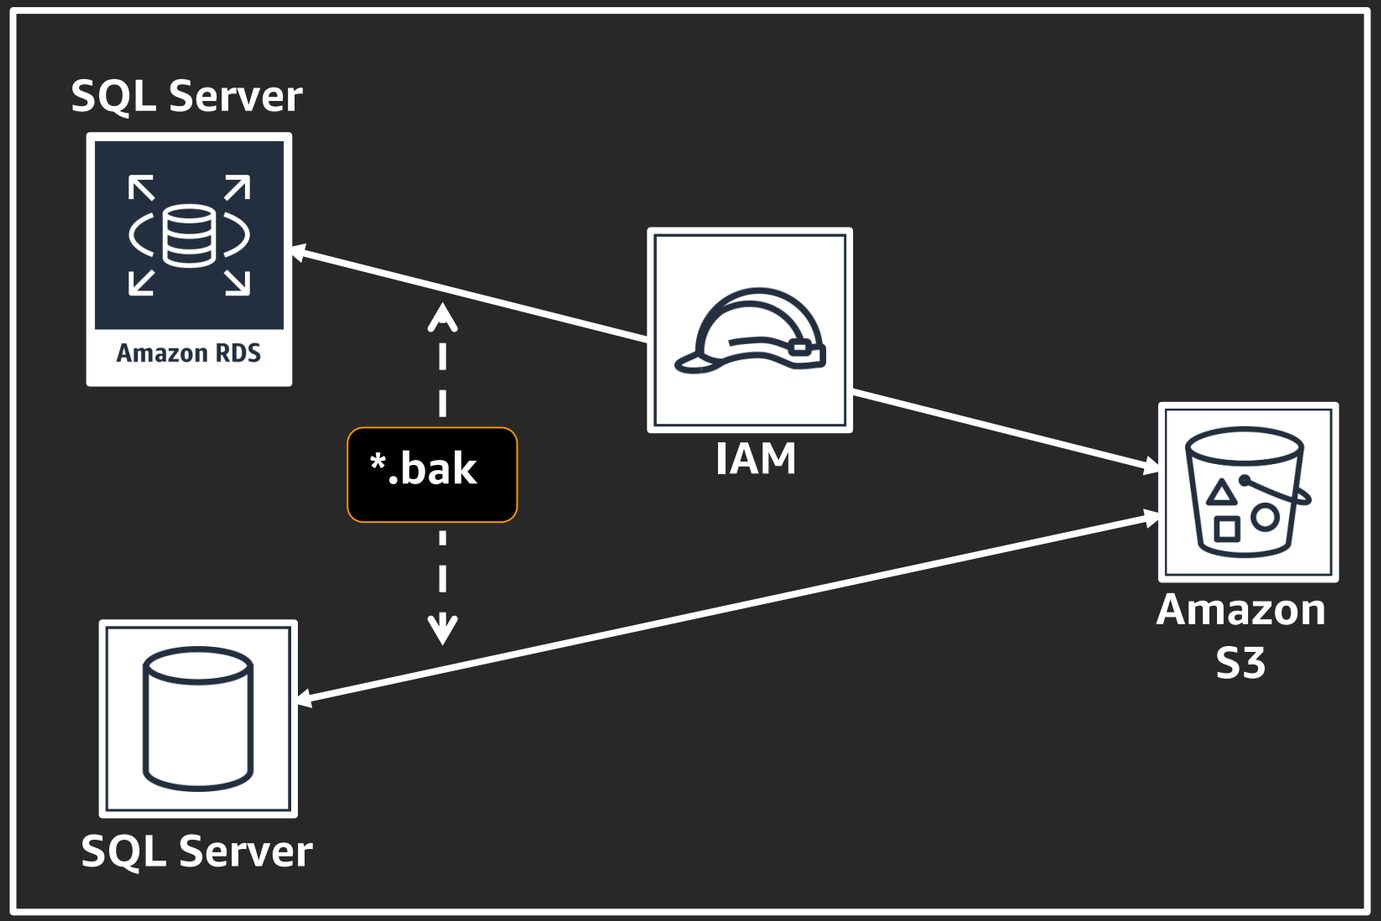

When to use AWS Import/Export Snowball

Cloud Migration

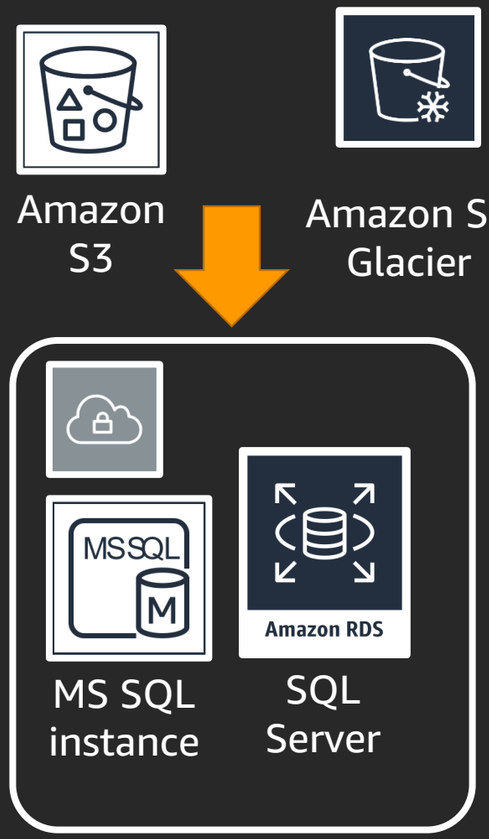

Pearson Large SQL Server Migration

"Our educational and assessment tools, content, products, and services are designed to help learners at every stage open doors to new experiences. Because wherever learning flourishes so do people. MathXL was a proof point that we could migrate a complex, internally-hosted application to the cloud environment. If MathXL would migrate so would other Windows applications that leveraged Windows Server and SQL Server."

"Moving large applications opened the possibility of ending expensive data center costs and leases and reducing costs. Migrating MathXL proved that we could reuse existing .NET skills in the cloud."

– Kathryn Wood, Director of Engineering

- Custom .NET app and migration
- 91 Windows servers
- 25 higher ed front end servers
- 12 school front end servers
- 6 application servers
- **10 SQL Servers**

Native SQL Server migration methods

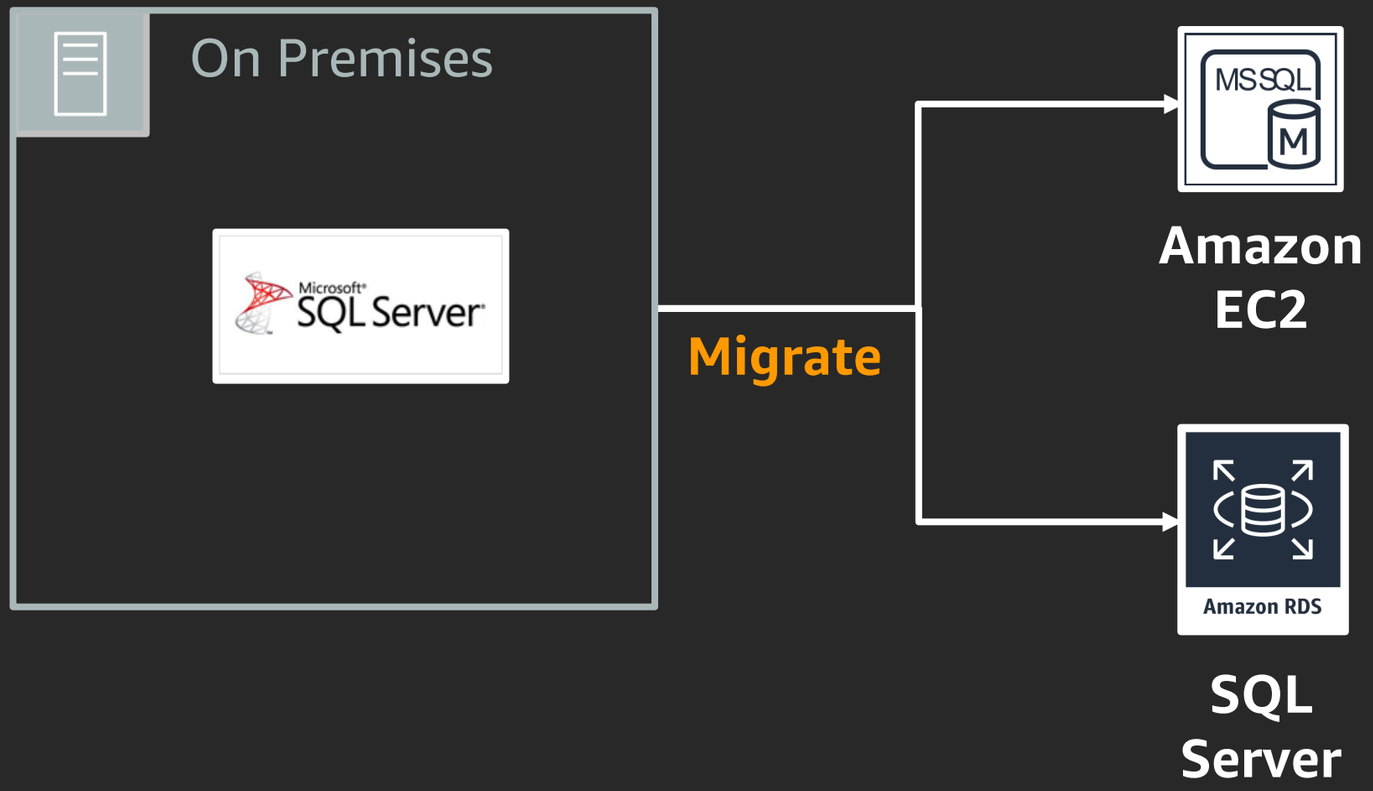

Log shipping

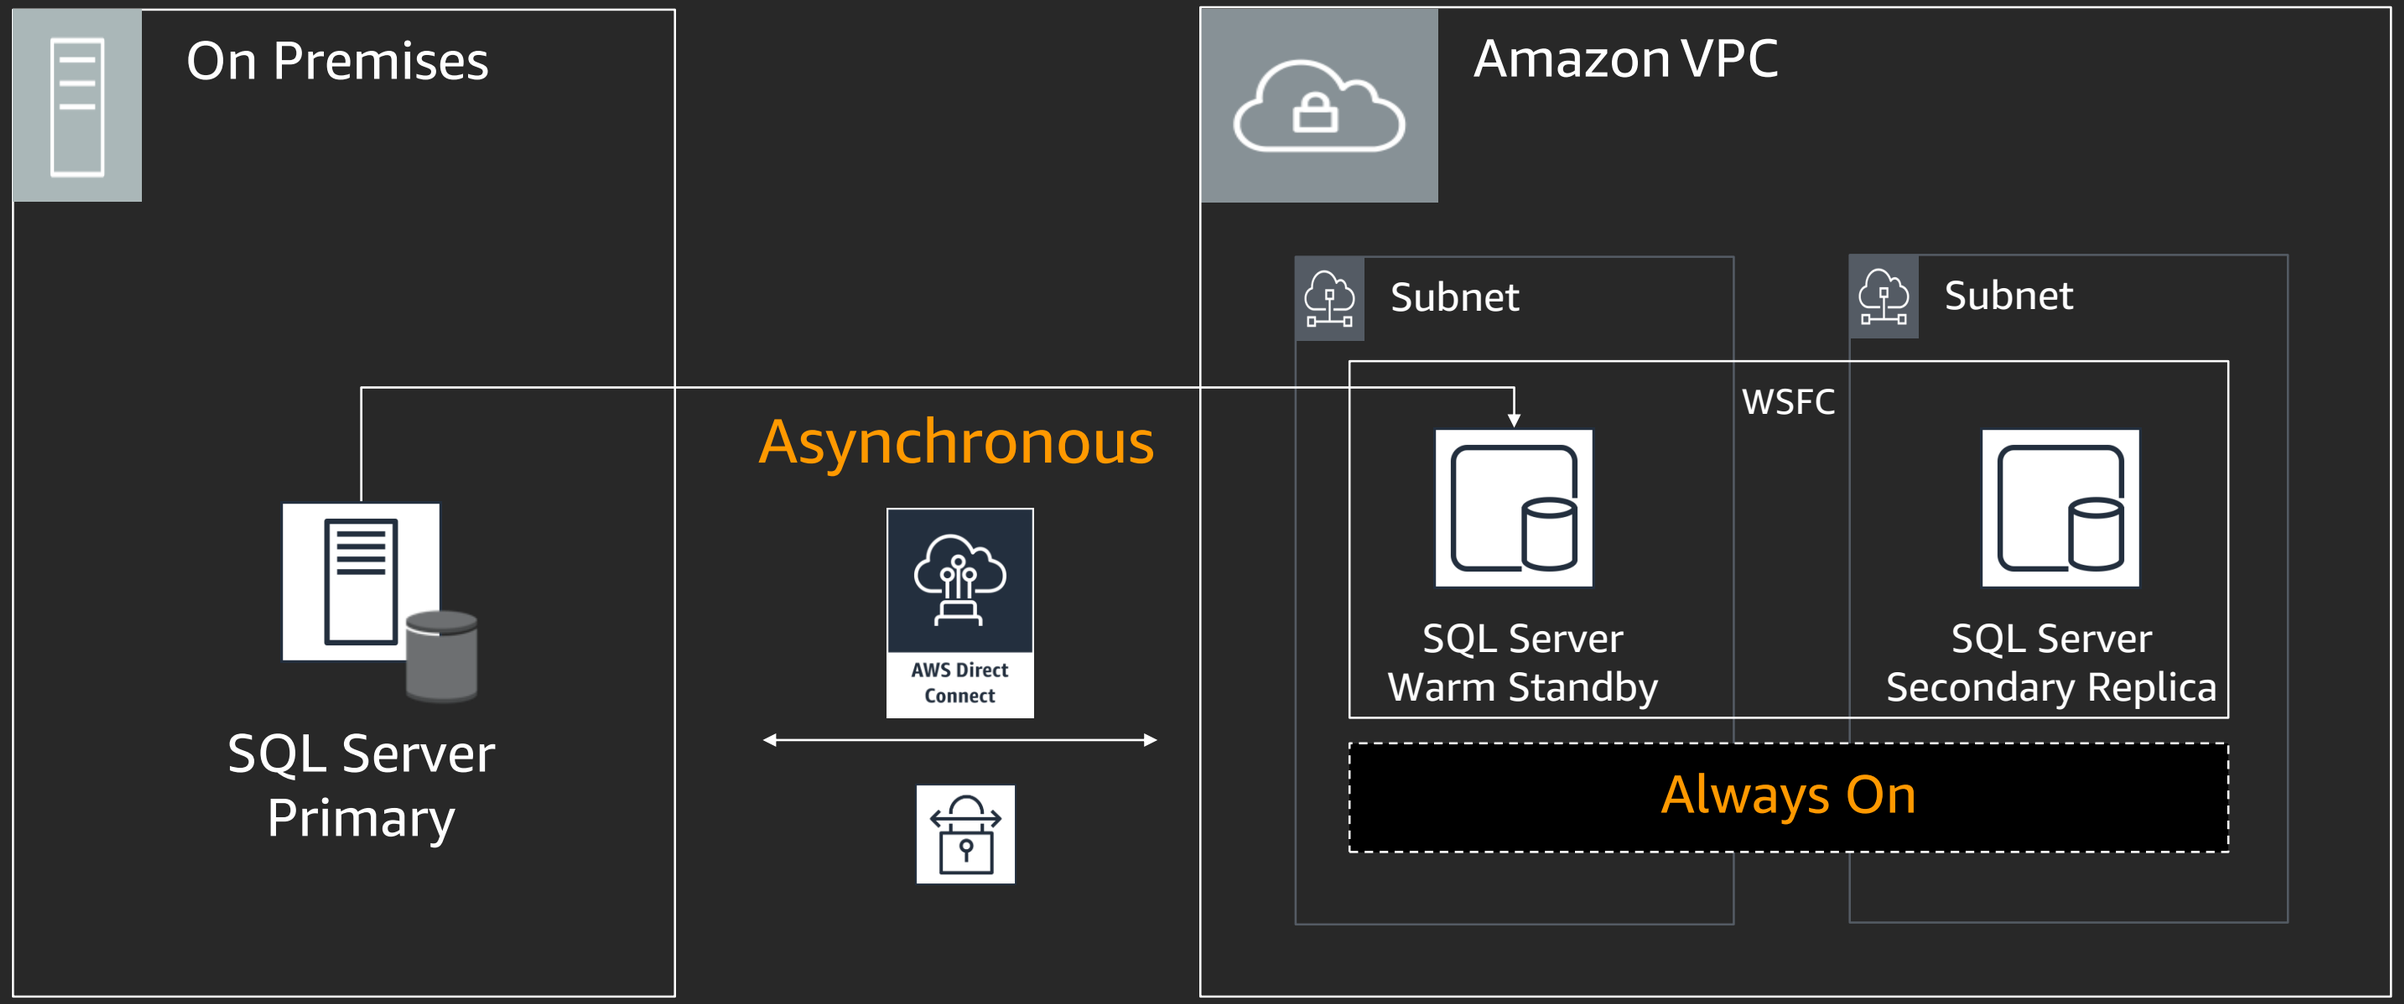

Database mirroring

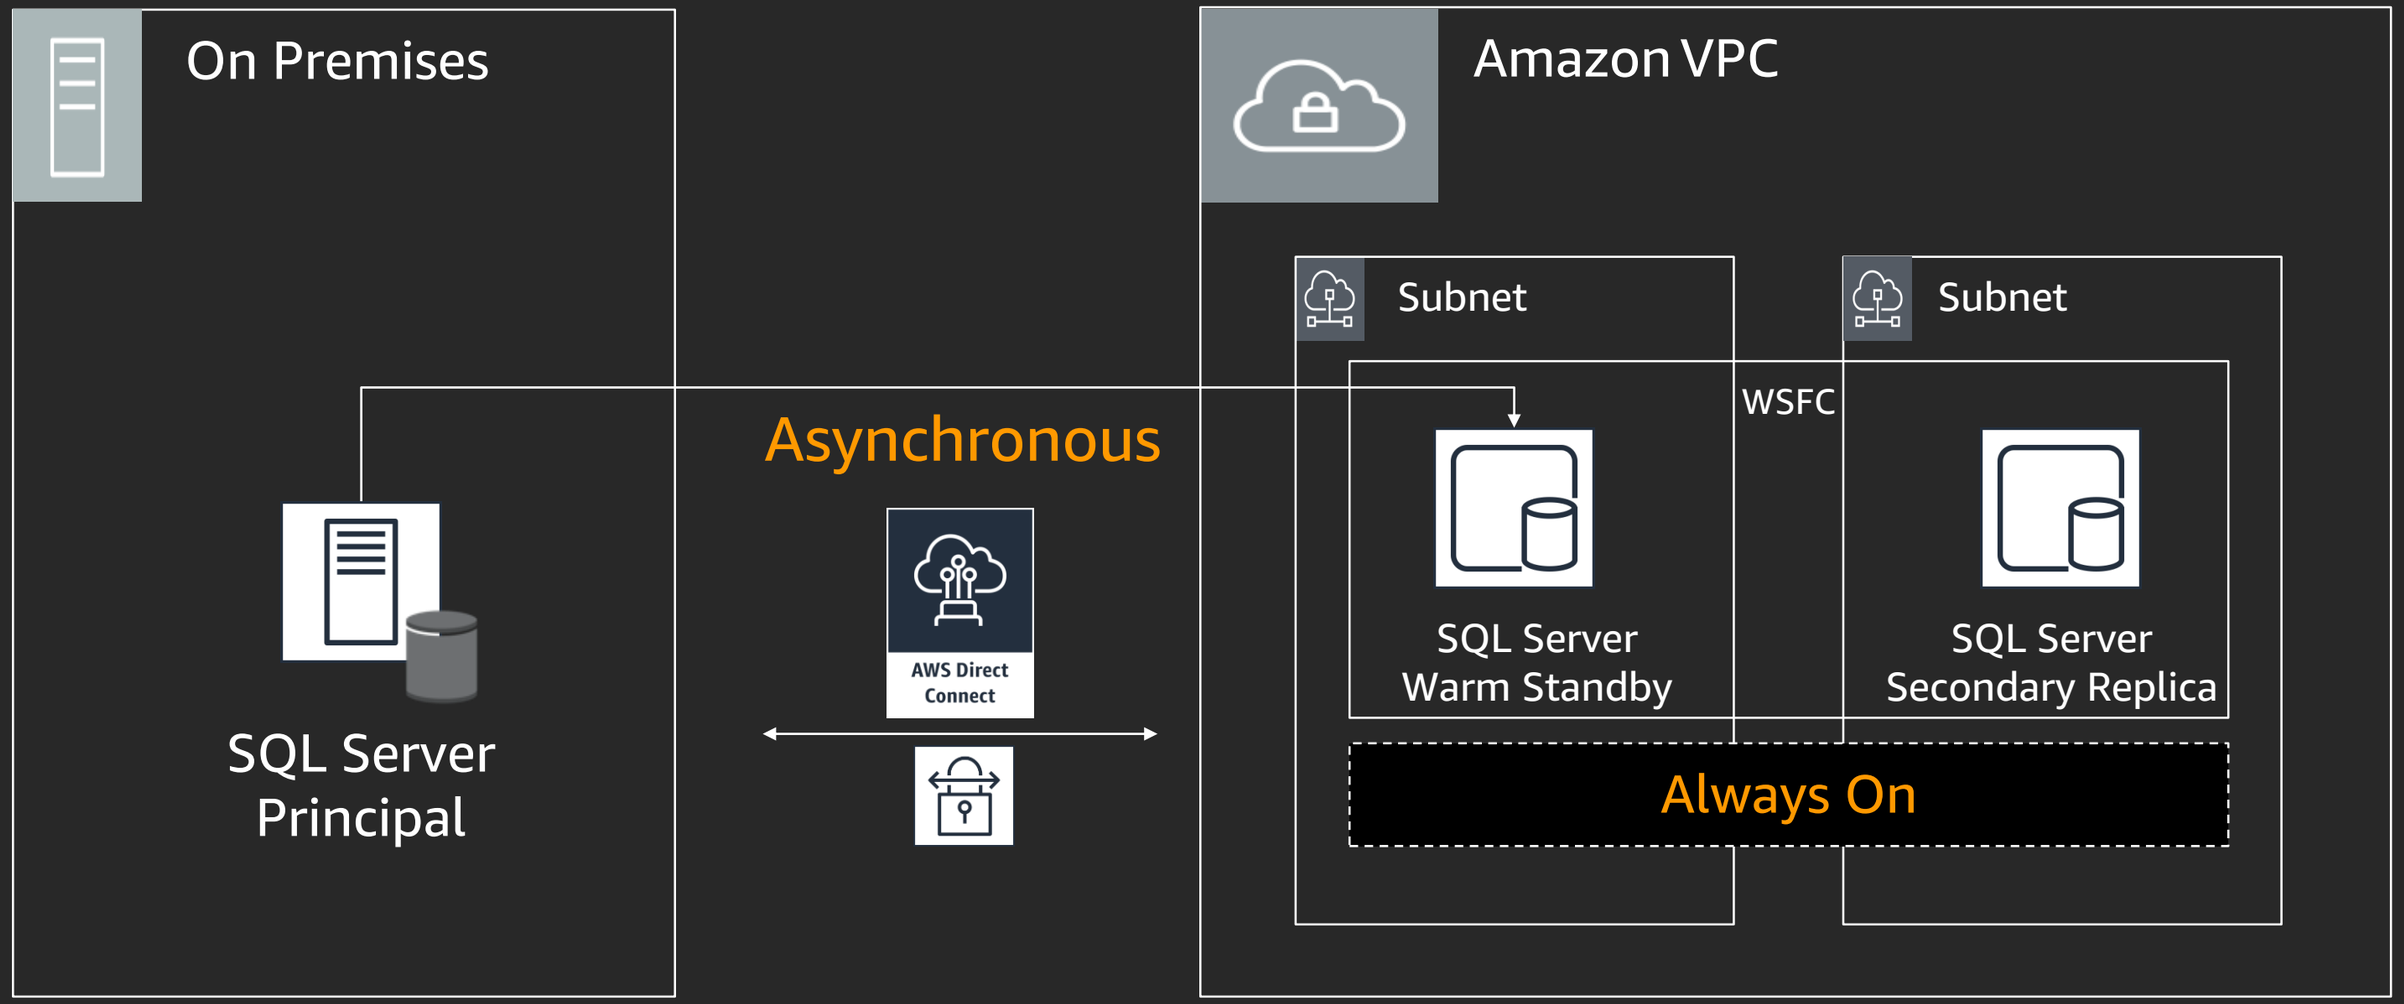

Always on availability groups

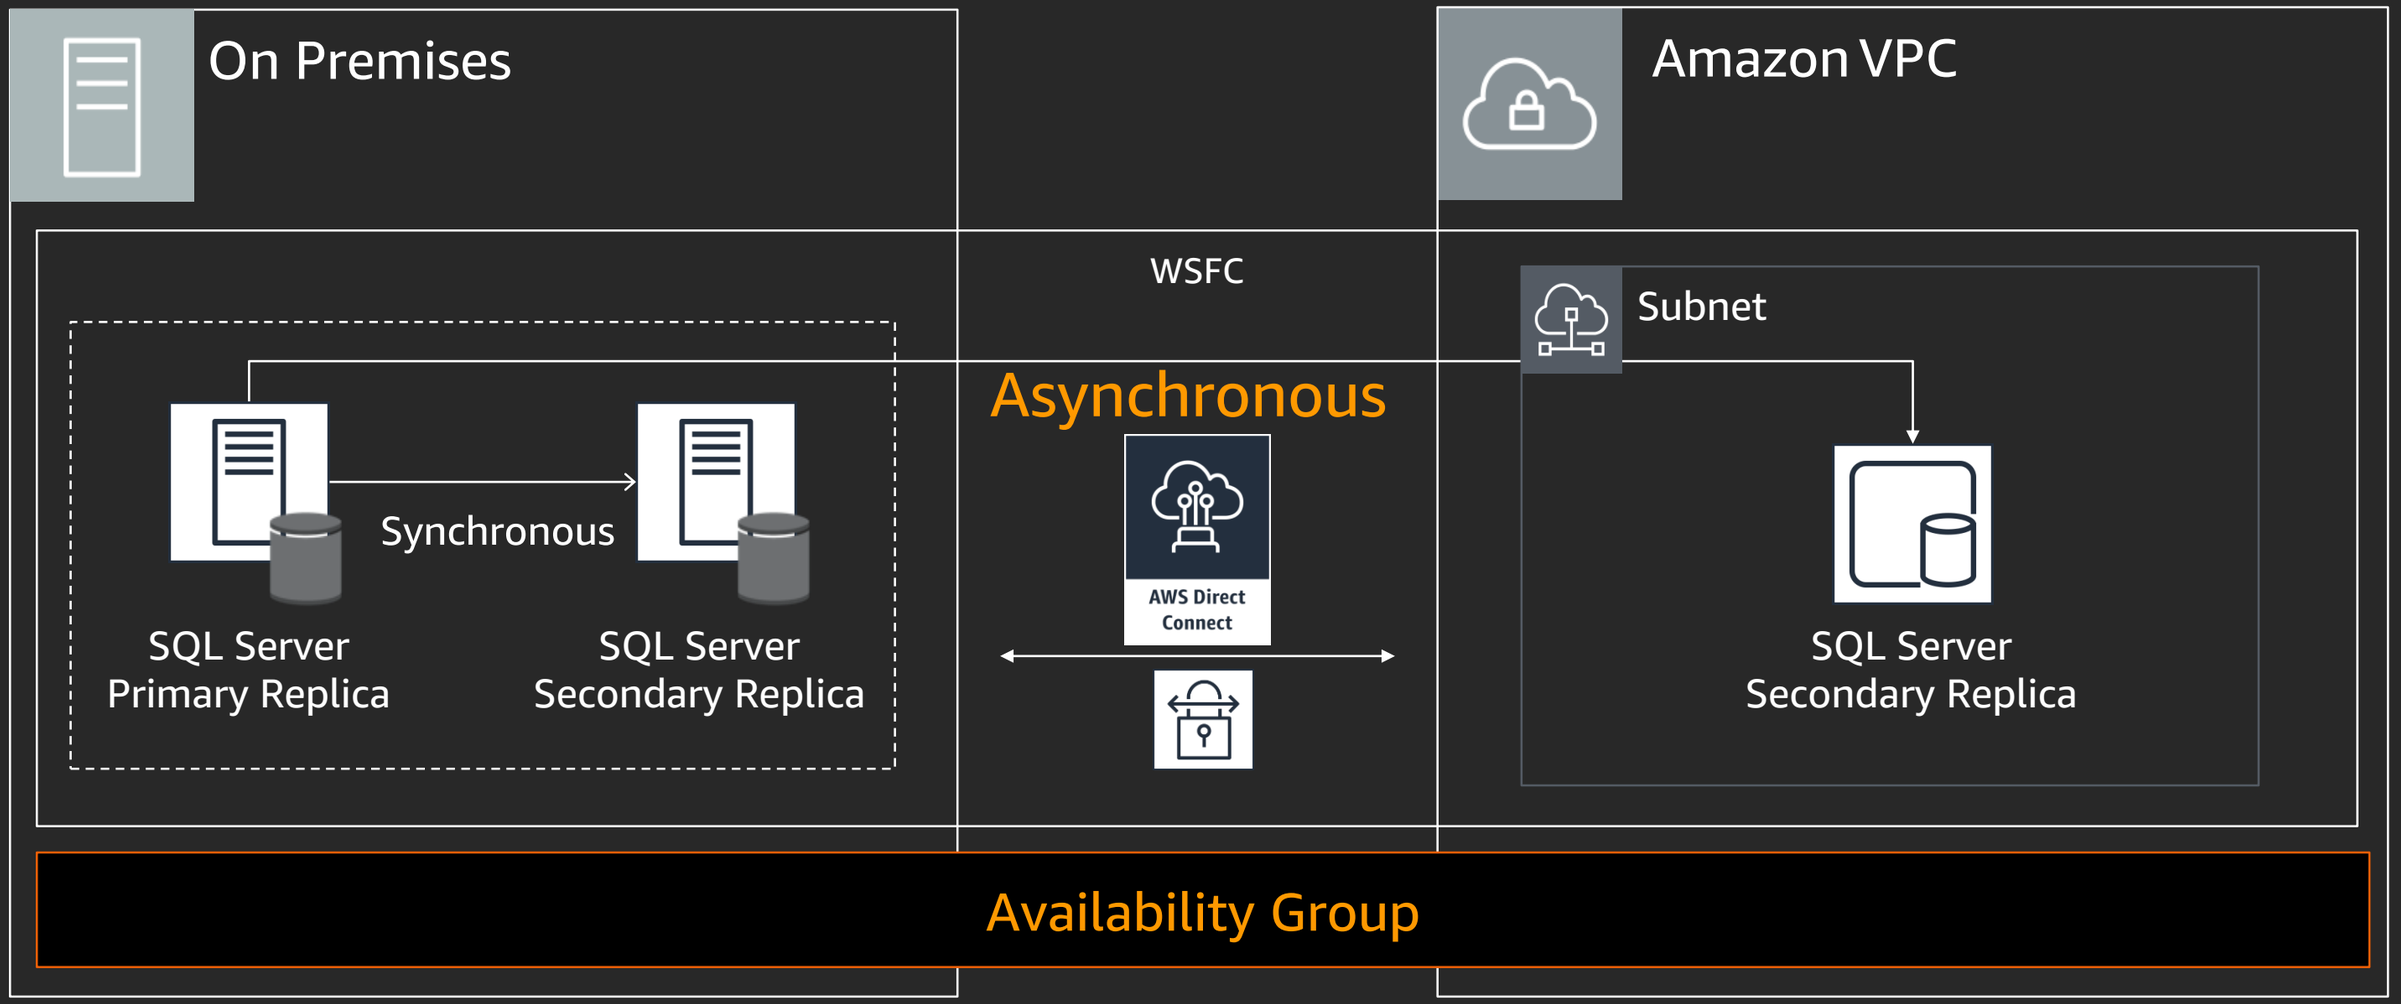

AWS Database Migration Service

AWS Database Migration Service

AWS DMS simply and securely migrate and/or replicate your databases *and* data warehouses to AWS

AWS Schema Conversion Tool (SCT) convert your commercial database and data warehouse schemas to open-source engines or AWS-native services, such as Amazon Aurora and Amazon Redshift

When to use AWS DMS and AWS SCT?

Modernize

Modernize your database tier

- SQL Server to open-source
- SQL Server to Amazon Aurora or PostgeSQL
- SQL Server to Amazon Redshift

Migrate

- Migrate business-critical applications
- Migrate data warehouse to Amazon Redshift
- Consolidate shards into Amazon Aurora

Replicate

- Create cross-regions read replicas
- Run your analytics in the cloud
- Keep your dev/test and production environment sync

Keep your apps running during the migration

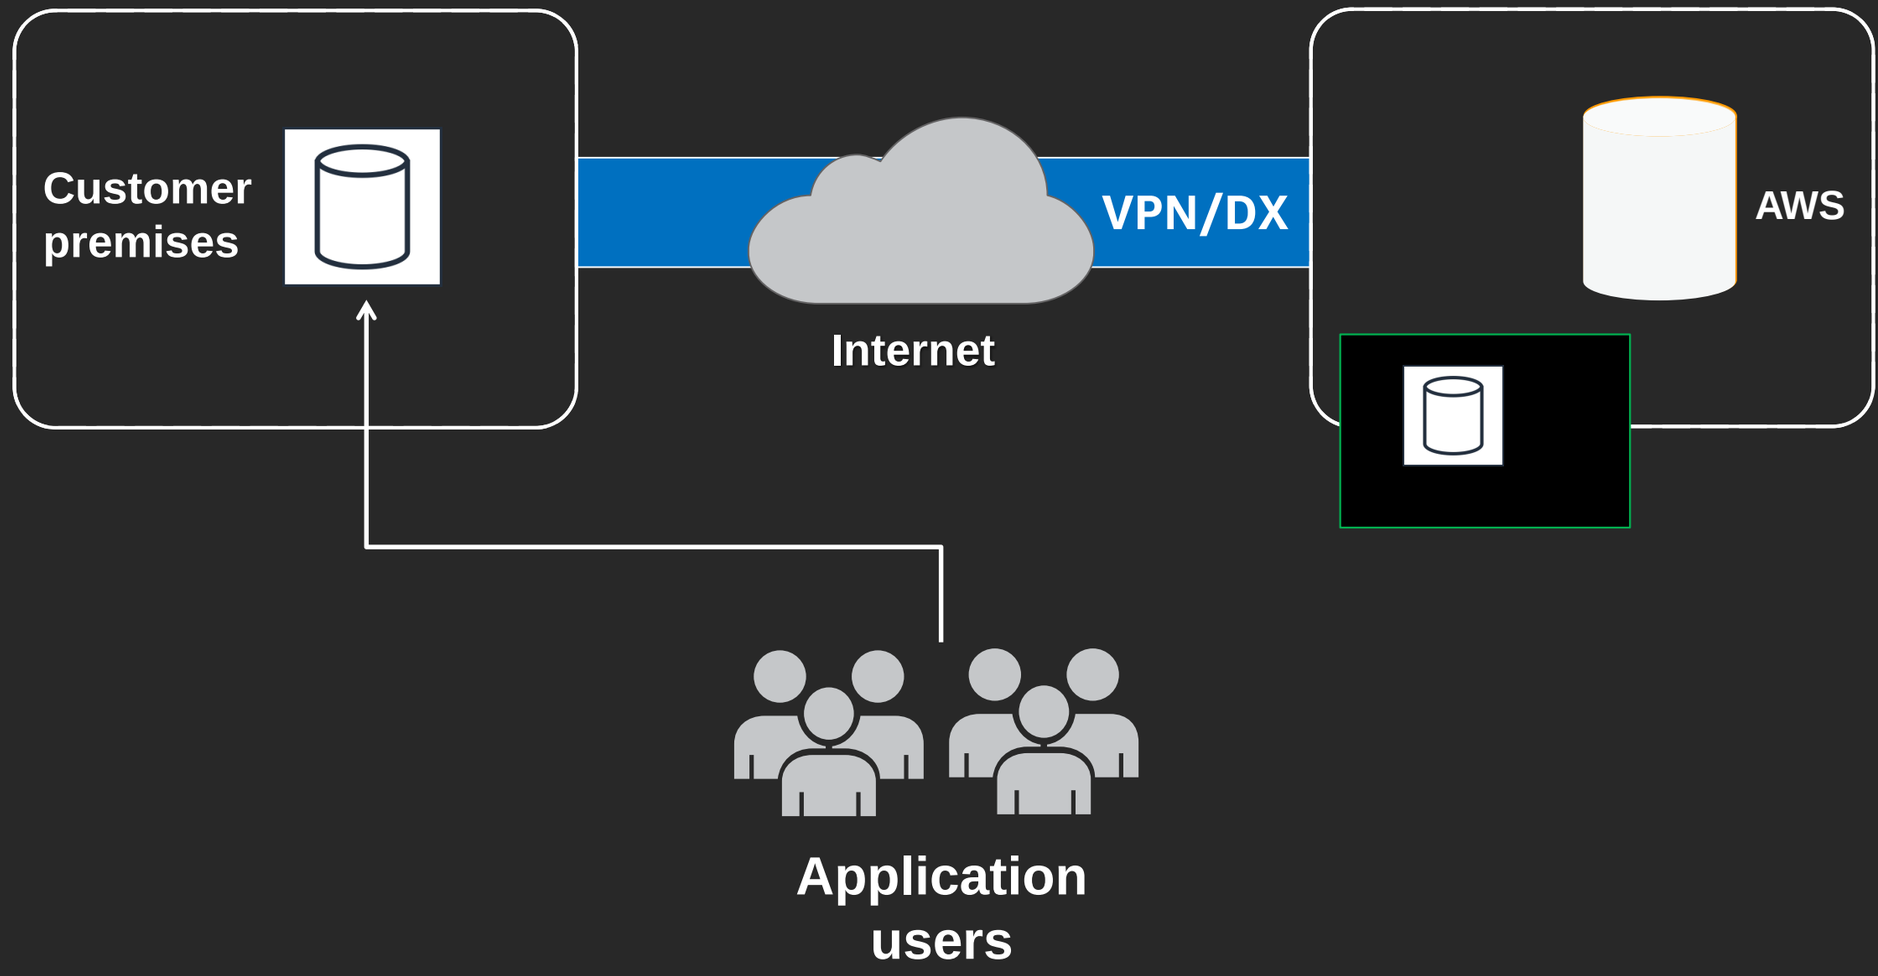

Keep your apps running during the migration

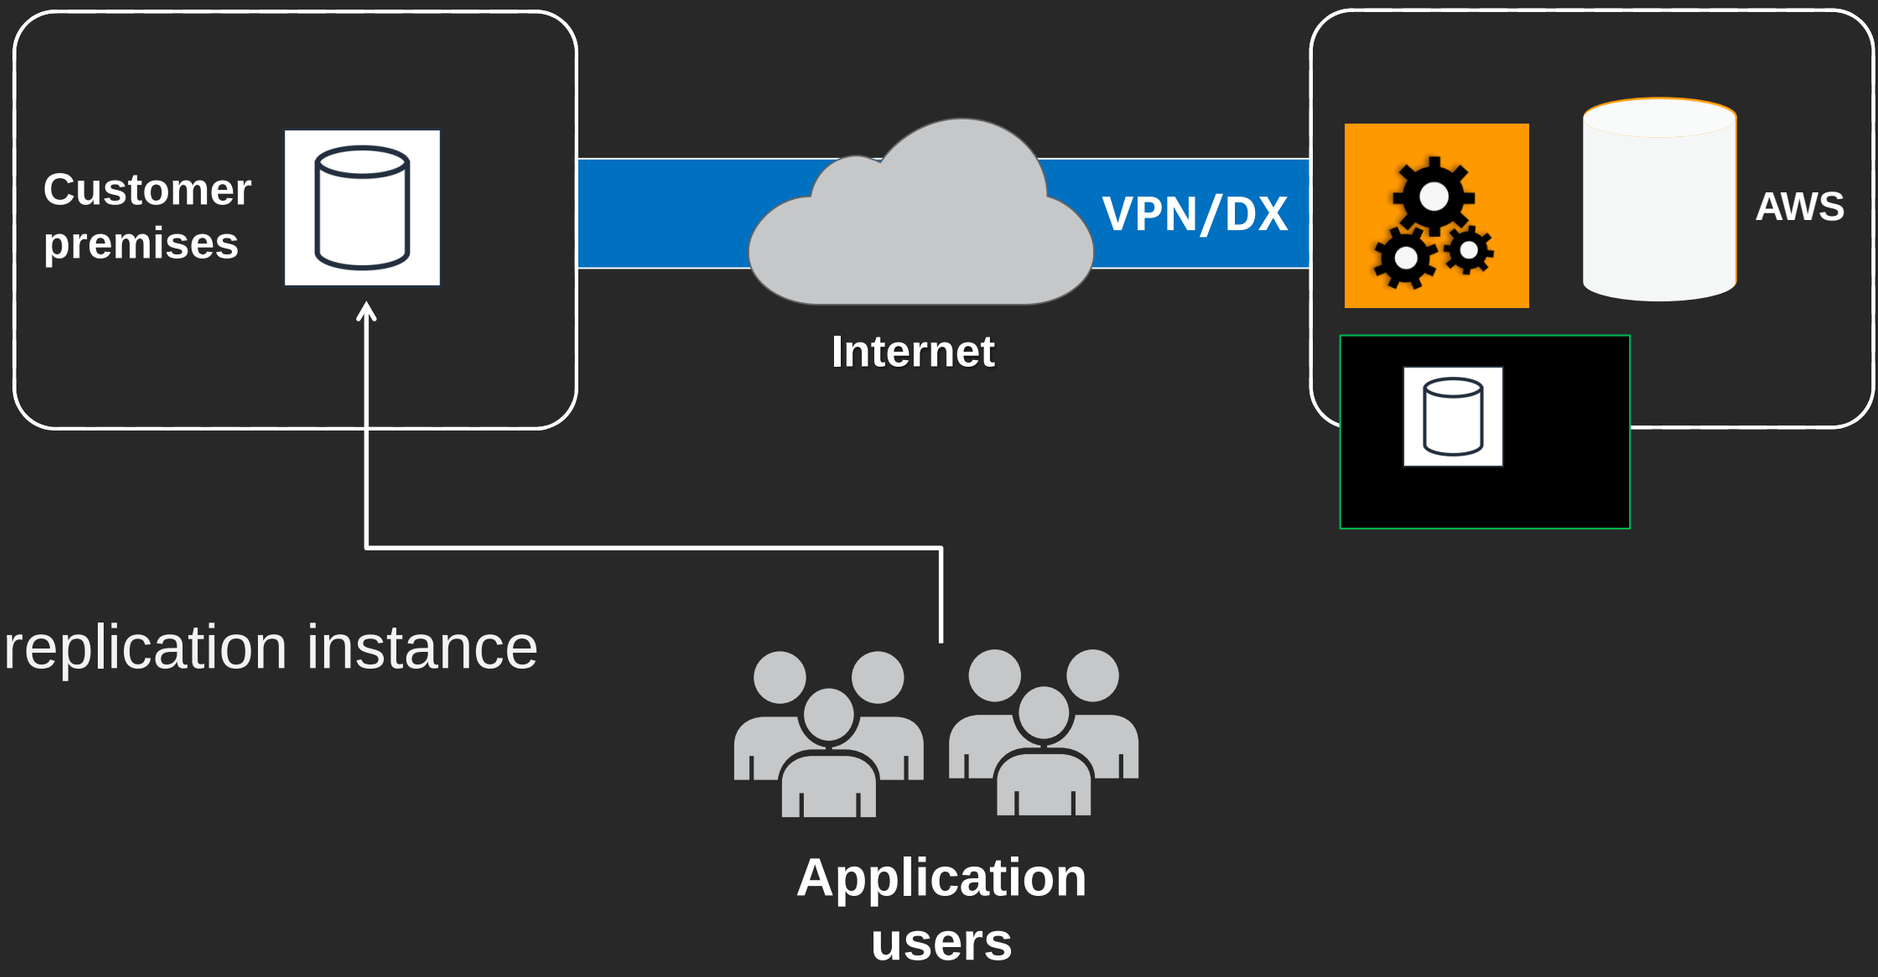

- Start a replication instance

Keep your apps running during the migration

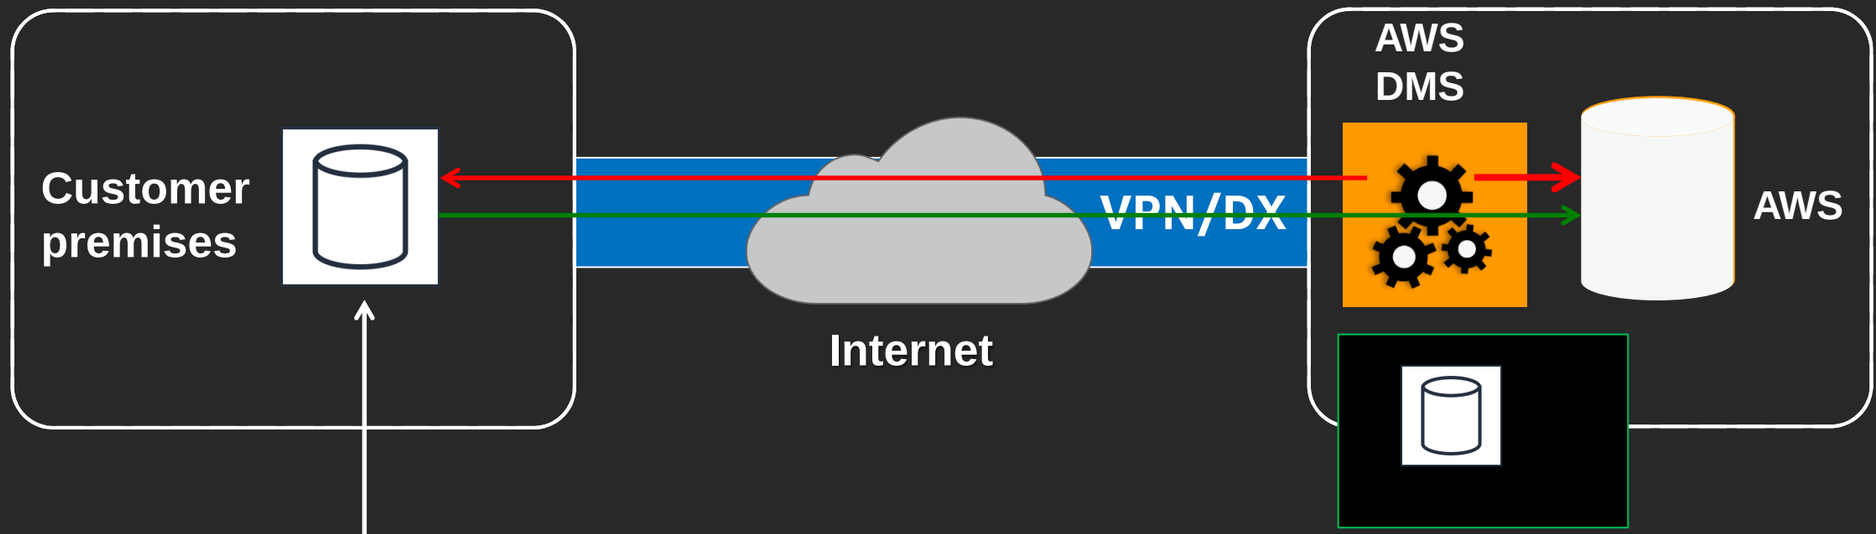

- Start a replication instance
- Connect to source and target databases
- Select tables, schemas, or databases

Application users

Keep your apps running during the migration

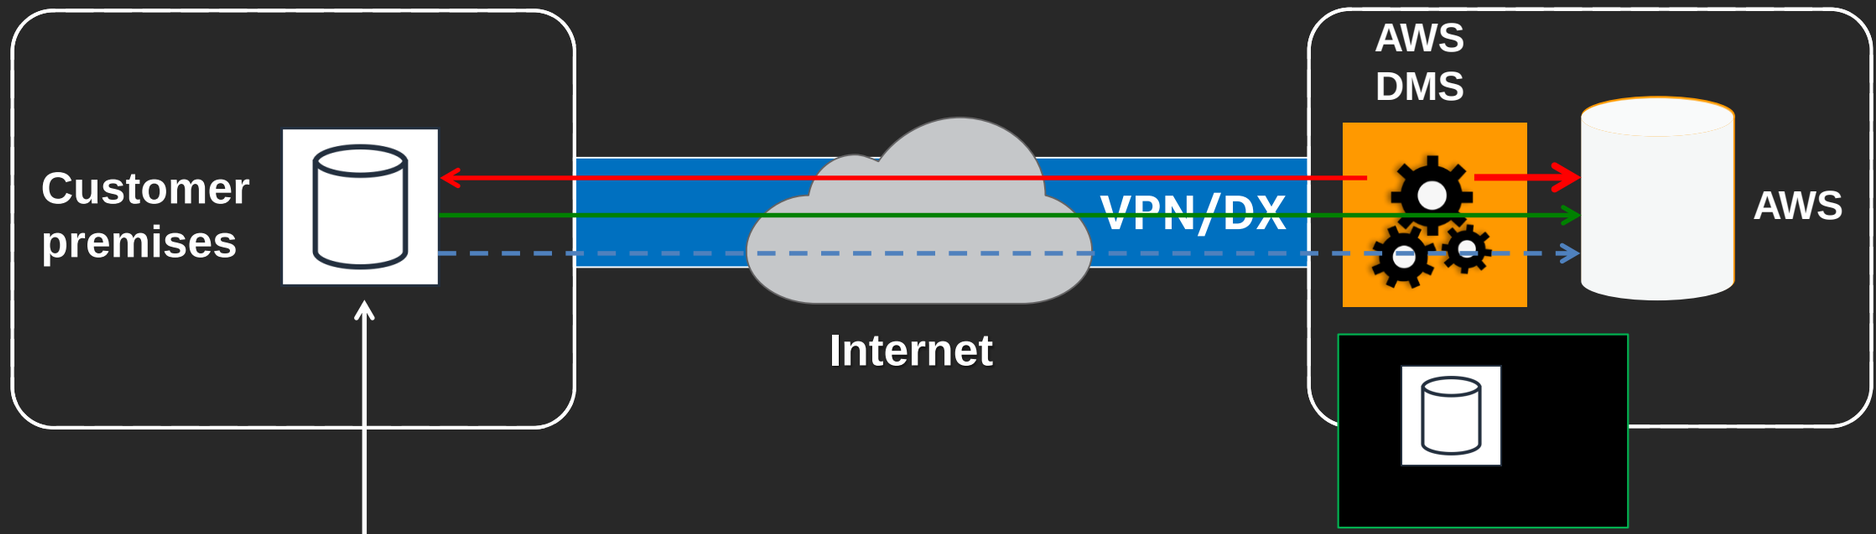

- Start a replication instance
- Connect to source and target databases
- Select tables, schemas, or databases

Application users

- Let AWS DMS load data, and keep them in sync

Keep your apps running during the migration

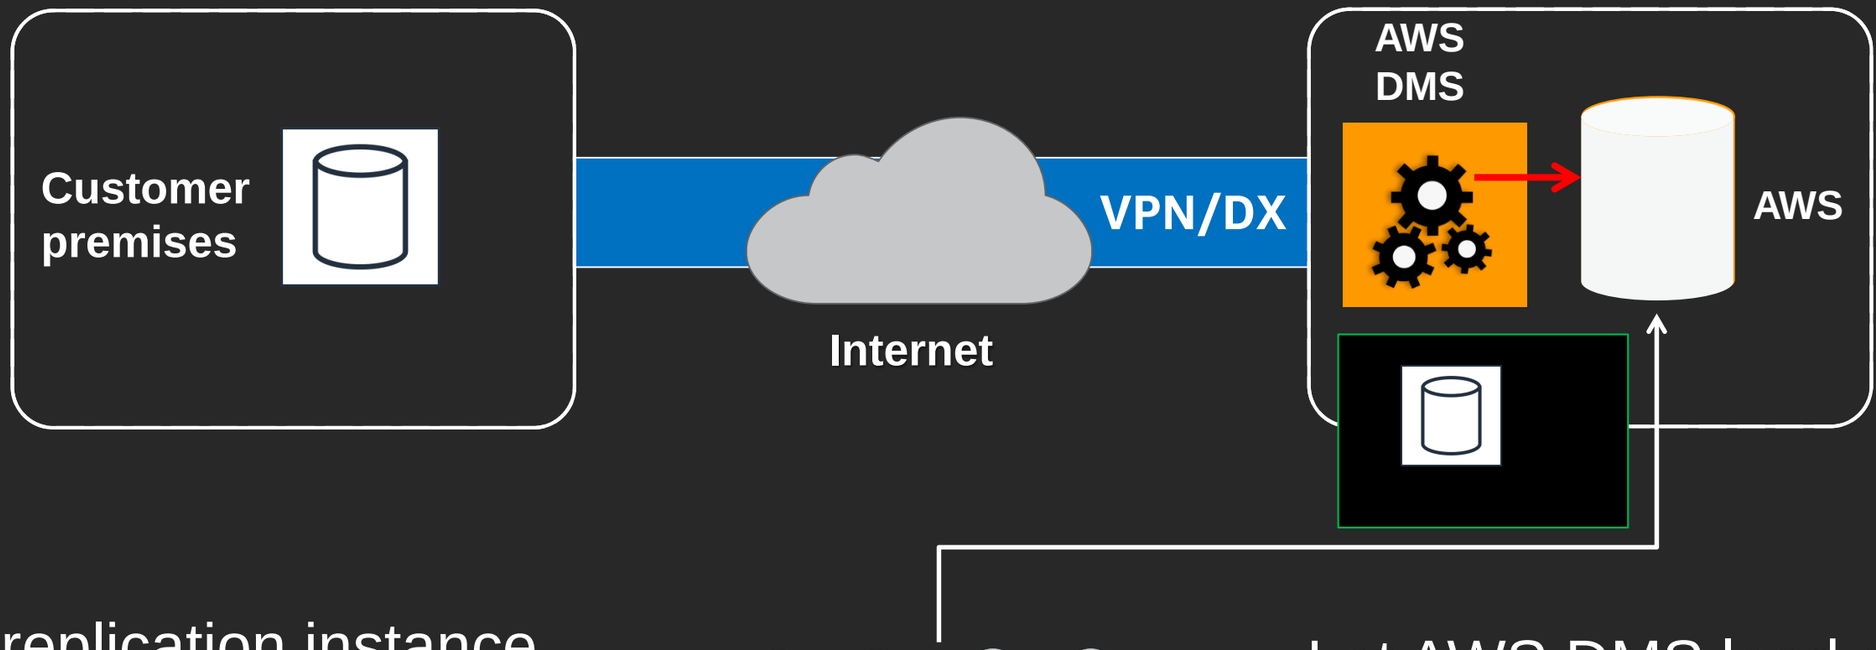

- Start a replication instance
- Connect to source and target databases
- Select tables, schemas, or databases

Application users

- Let AWS DMS load data, and keep them in sync
- Switch applications over to the target at your convenience

Selecting the right migration method

Migration method: Which should I use?

	Backup/ Restore	Transactional Replication	AAGs	Log Shipping	DB Mirroring	DMS/ SCT	AWS Snowball
SQL Server Standard	✓	✓		✓		✓	✓
SQL Server Enterprise	✓	✓	✓	✓	✓	✓	✓
Ongoing Replication		✓	✓	✓	✓	✓	
Migrate Specific DB Objects (e.g. sprocs, tables, indexes, etc.)	✓	✓				✓	
SQL Server 2008/2008R2	✓ (ALL)			✓ (SE,EE)	✓ (EE)	✓ (ALL)	✓ (ALL)
SQL Server 2012+	✓ (ALL)	✓	✓ (EE)	✓ (SE,EE)	✓ (EE)	✓ (ALL)	✓ (ALL)
SQL Server on Amazon EC2	✓	✓	✓	✓	✓	✓	✓
Amazon RDS for SQL Server	✓	✓				✓	✓

Summary of migration considerations

- SQL Server version and edition features
- Authentication requirements
- Amount of data being migrated
- Connectivity to AWS
- Migration method
- New home for the databases
- AWS Professional Services or partner help
- AWS Well Architected Framework
- Optimize after migration

Thank you!

Bill Jacobi

aws PUBLIC SECTOR SUMMIT ONLINE