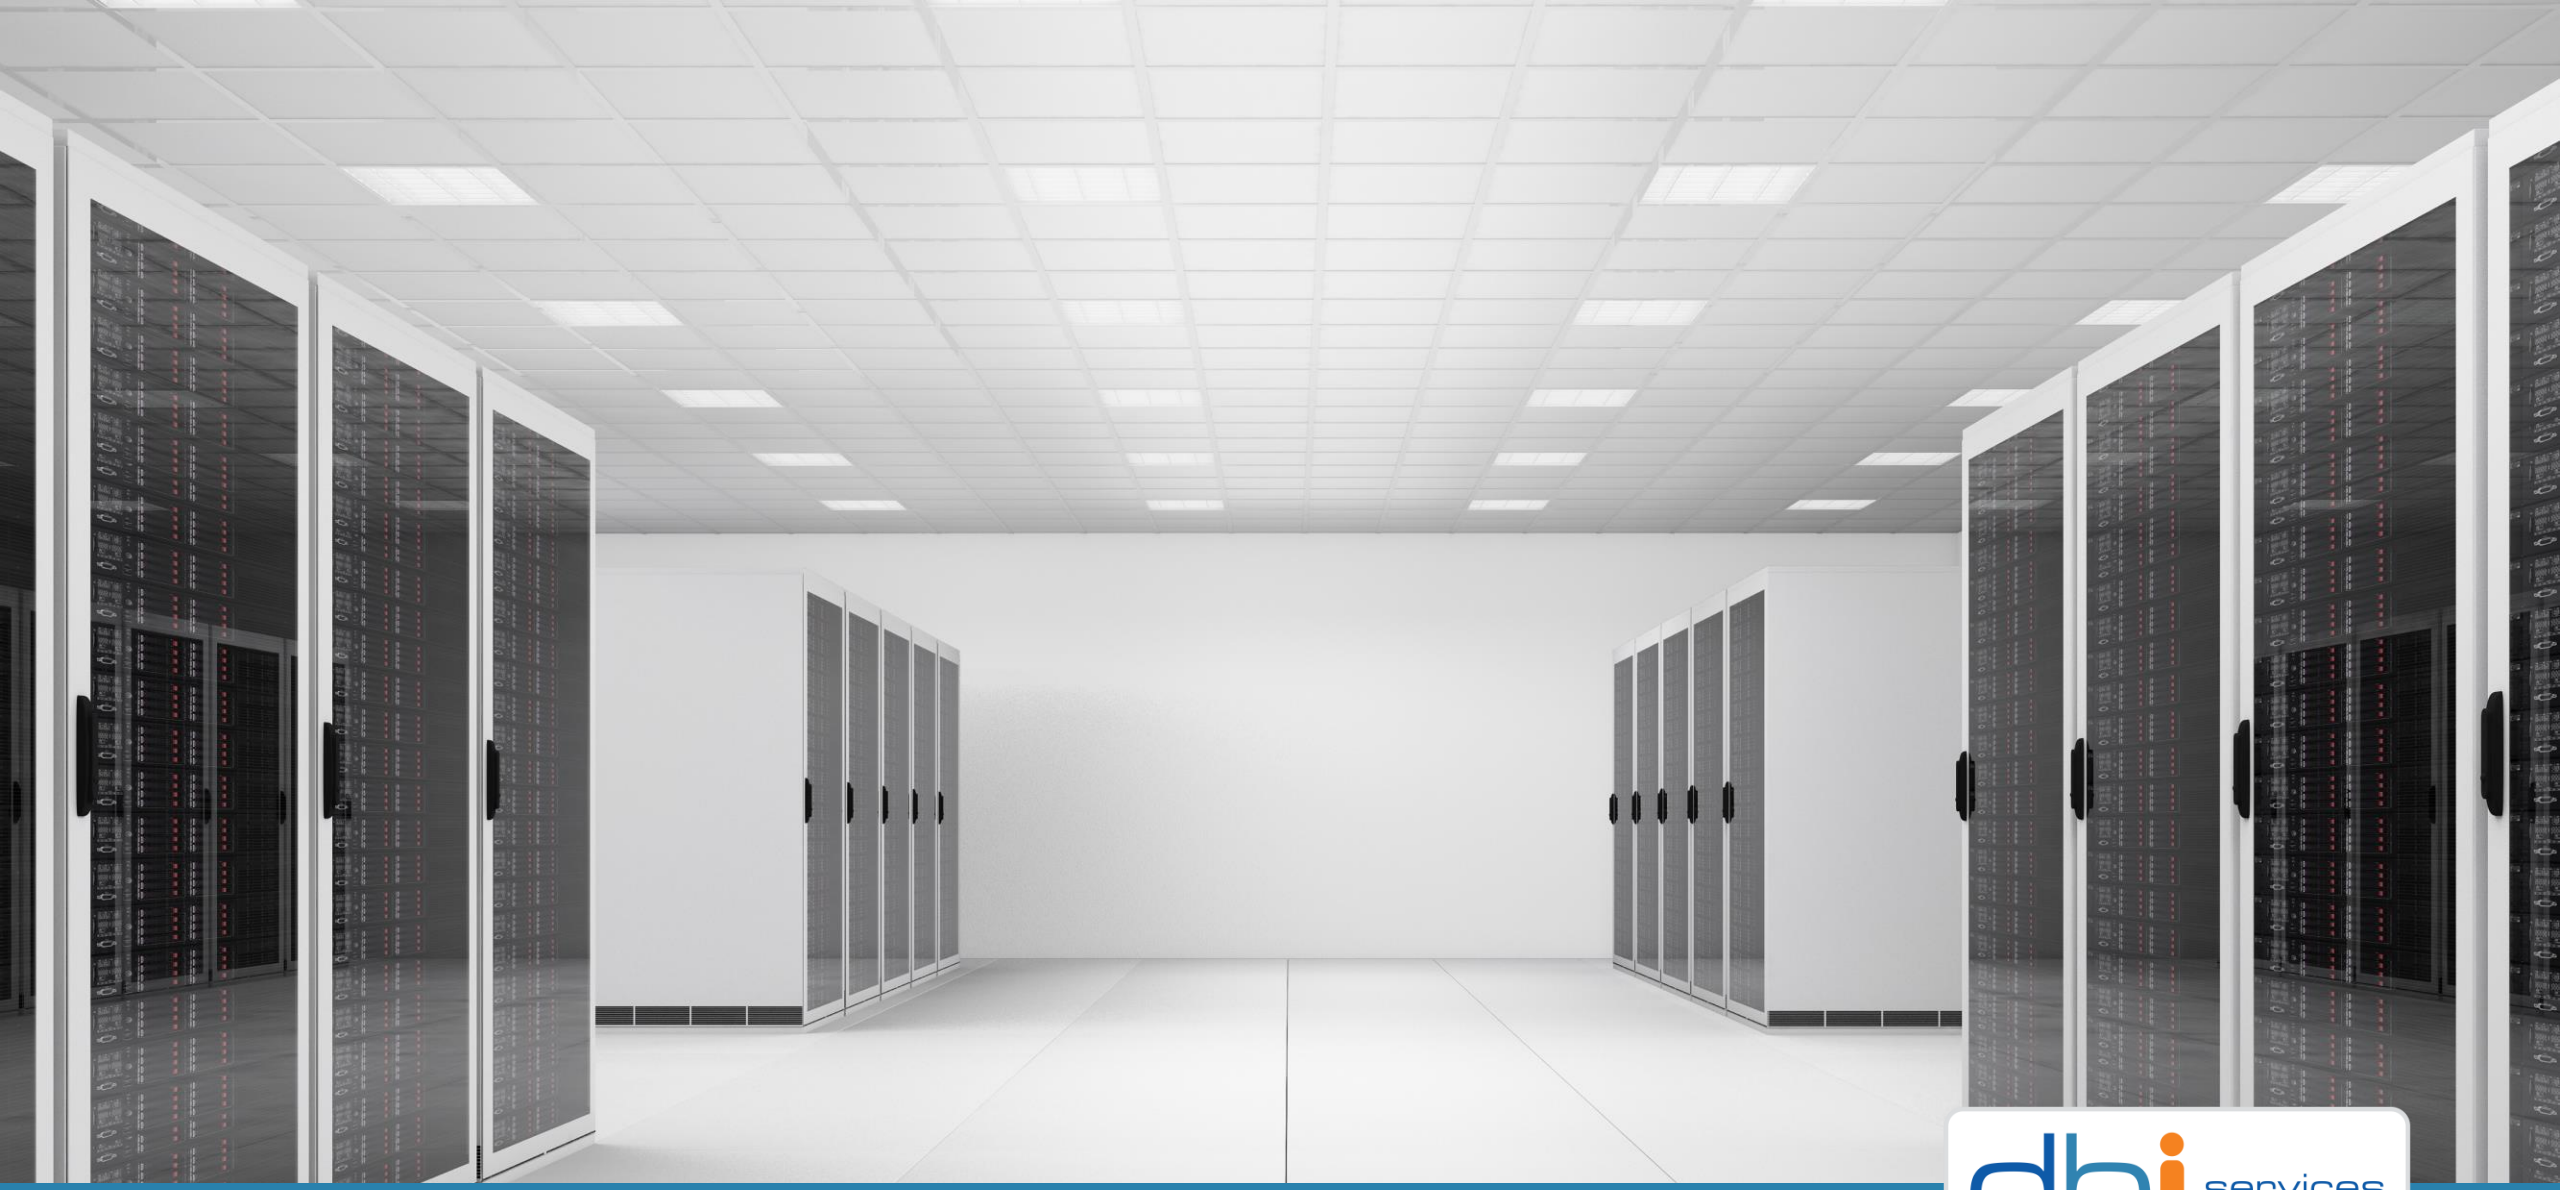

dbi services by Sequotech

Useful AWS services to operate SQL server in the cloud

Who we are

The Company

- Founded in 2010
- More than 100 employees
- Specialized in the Middleware Infrastructure
  - The invisible part of IT
- Customers in Switzerland and all over Europe

Our Offer

- Consulting
- Service Level Agreements (SLA)
- Trainings
- License Management

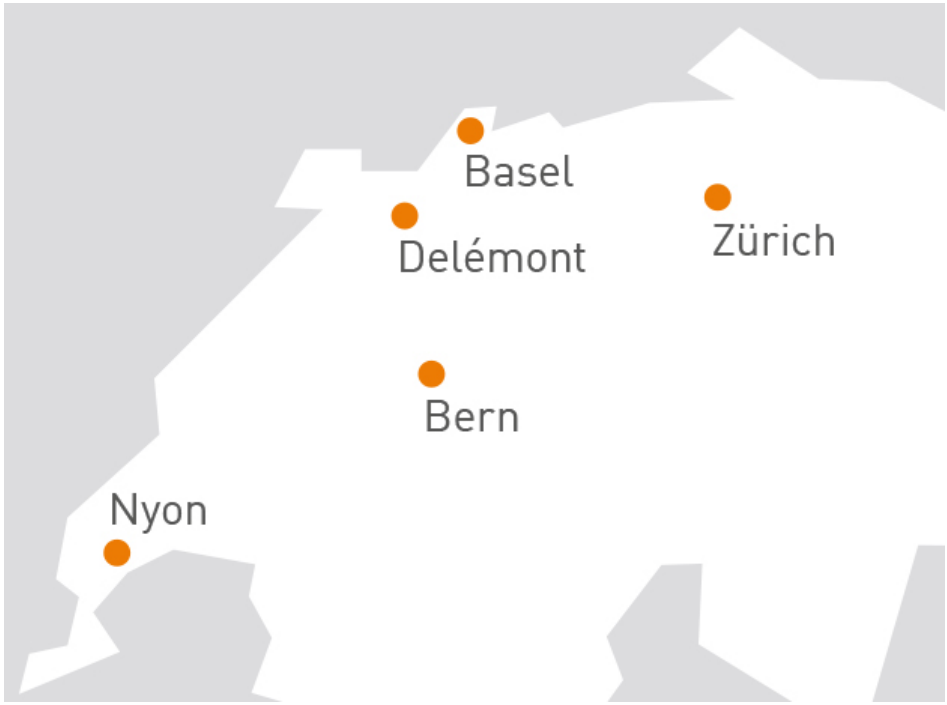

swiss safety center ISO 20000 certified system

swiss safety center ISO 27001 certified system

About me

Nicolas Jardot

Principal Consultant

+41 79 835 02 95

nicolas.jardot[at]dbi-services.com

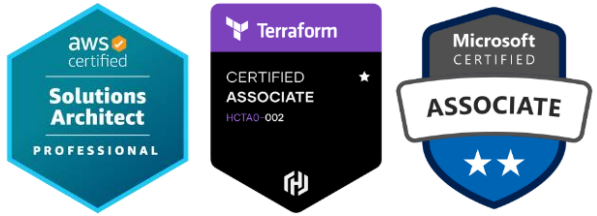

Agenda

1. SQL Server on AWS
2. Storage management
3. Backup
4. High Availability

SQL Server on AWS

- IaaS vs. PaaS
- Amazon EC2
- Amazon RDS
- Amazon RDS Custom

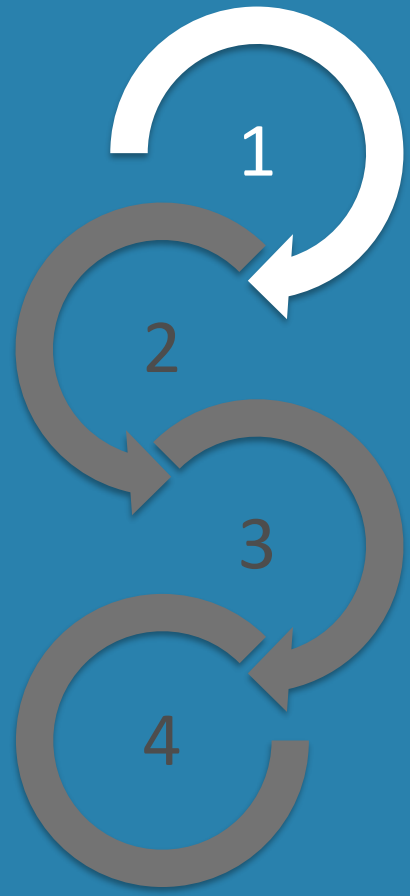

SQL Server on AWS
IaaS vs. PaaS

There are 2 main options to run SQL Server on AWS
> IaaS - On Amazon Elastic Compute Cloud instances (EC2)
> PaaS - On Amazon Relational Database Service (RDS)

Choose based on required control and flexibility needs

Amazon EC2
> AWS maintains infrastructure up to VM provisioning
> OS and SQL Server configuration by customer

Amazon RDS
> AWS maintains up to the database engine
> Customer responsible for databases and data

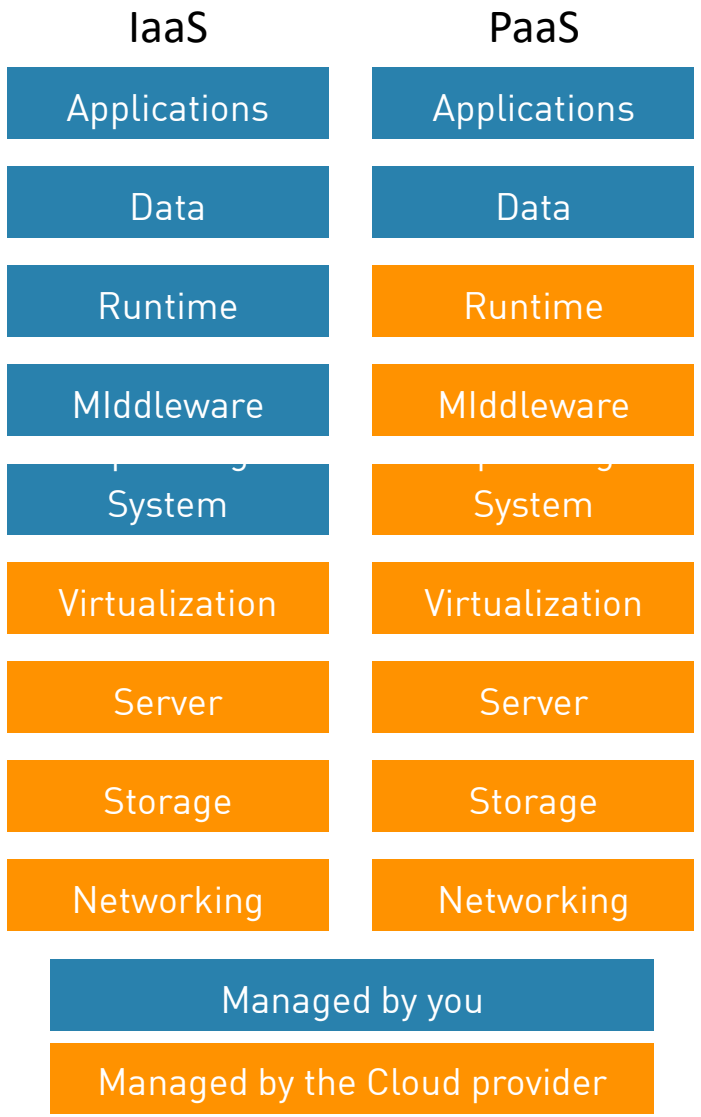

SQL Server on AWS
Amazon EC2

Full control over OS and SQL Server configuration

> Customer deploys a VM in his own VPC

> No deployment automation unless done by customer side (YaK, AWS Launch Wizard, Terraform/Ansible, …)

Bring Your Own License (BYOL)

> Customer is responsible to install and configure SQL Server

> Require Software Assurance in most cases for License Mobility

> https://docs.aws.amazon.com/sql-server-ec2/latest/userguide/sql-server-on-ec2-licensing-considerations.html

Get SQL Server license from AWS

> Preconfigured Amazon Machine Image (AMI) with SQL Server already installed

> Additional software installed : AWS Systems Manager, Storage & Network drivers, etc.

> Pay as you Go, license costs included in running costs, e.g. Dev instances that can be stopped

SQL Server on AWS
Amazon RDS

Amazon RDS is fully managed service

> Supports multiple commercial and open-source relational database engines
> Easily scale compute and storage
> No OS access / No sysadmin privileges
> License Included only: Express, Web, Standard and Enterprise

Main features

> High-availability in one click with Multi-AZ
> Automated backups
> Managed OS and database patching
> Supports Read replicas
> Supports options : SQL Server Audit, Integration Services (SSIS), Analysis Services (SSAS), …
> Supports SQL Server 2016 / 2017 / 2019 / 2022

SQL Server on AWS
Amazon RDS - Administration

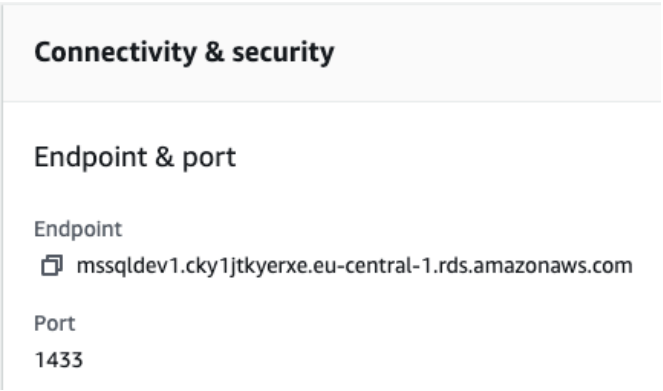

Connect to database endpoint

> Allow traffic in VPC security group

> Can be publicly accessible but not recommended

Use parameter group and option group to configure

Common DBA tasks

> Use Amazon RDS procedures: e.g. take manual SQL backup (require *SQLSERVER_BACKUP_RESTORE* option)

> https://docs.aws.amazon.com/AmazonRDS/latest/UserGuide/Appendix.SQLServer.CommonDBATasks.html

```
exec msdb.dbo.rds_backup_database
      @source_db_name='database_name',
      @s3_arn_to_backup_to='arn:aws:s3:::bucket_name/file_name_and_.extension',
      @kms_master_key_arn='arn:aws:kms:region:account-id:key/key-id',
      @overwrite_S3_backup_file=1,
      @type='FULL',
      @number_of_files=n;
```

SQL Server on AWS
Amazon RDS - Limitations

Maximum number of databases depends on instance class and availability

Instance Class	Single-AZ	Multi-AZ with DBM	Multi-AZ (Always On AGs)
db.*.micro to db.*.medium	30	N/A	N/A
db.*.large	30	30	30
db.*.xlarge to db.*.16xlarge	100	50	75
db.*.xlarge to db.*.16xlarge	100	50	100

Maximum storage size

> SSD – 16 TiB

> Magnetic – 1 TiB

Not supported

> Data import in msdb, Data Quality Services, Master Data Services

SQL Server on AWS
AWS RDS Custom

Built for applications that require access to underlying OS and database environment

> Customer needs to install custom drivers
> Customer needs features requiring elevated privileges, e.g. CLR, xp_cmdshell
> Applications need to install packages on same server, e.g. Microsoft SharePoint
> Use more than 100 databases

Use Amazon RDS as Disaster Recovery

> Replicate on-premise environment to RDS with Always On Availability Groups for DR

Still benefit from AWS Automation

> Automated Backups
> Monitoring with CloudWatch
> OS patching if using RDS provided AMI
> Point-in-Time Restore (if less than 1000 DBs)

SQL Server on AWS
AWS RDS Custom - New concepts

Automation mode

- Default to full automation mode, AWS monitors and performs automatic instance recovery
- Can be paused
  - Ensure customer customization done on SQL DB instance is not interacting with RDS Custom automation
  - From 60 (default and minimum) and up to 1440 minutes

Supported perimeter

- New monitoring capability
- Customer has access to OS and may break AWS automations
- Support perimeter checks that database instance can still be managed by AWS
  - Send event notification if not the case
  - DB stays available but AWS features like automated backup are disabled
  - Once issue is fixed DB comes back to available state

SQL Server on AWS
AWS RDS Custom – When to choose

SQL Server on EC2	RDS Custom for SQL Server	RDS for SQL Server
Self managed	Shared management	AWS managed
Full control	Full control	No sysadmin / OS access
All DB engine features	Access all SQL Server configurations	Optimized architecture
Self managed backup, restore, monitoring, PITR	Automated provisioning, backup, restore, monitoring, PITR	Automated provisioning, backup, restore, monitoring, PITR
Self-managed patching	Automated patching	Automated patching
Self-managed high-availability	Managed high availability	Managed high availability
Allows 3rd party apps on DB host	Allows 3rd party apps on DB host	No 3rd party apps on DB host
BYOL, License Included	BYOM, License Included	License Included
All	Web, Standard, Enterprise	Express, Web, Standard, Enterprise
All	2019, 2022	2016, 2017, 2019, 2022

Storage management

- Block Storage
- File Storage

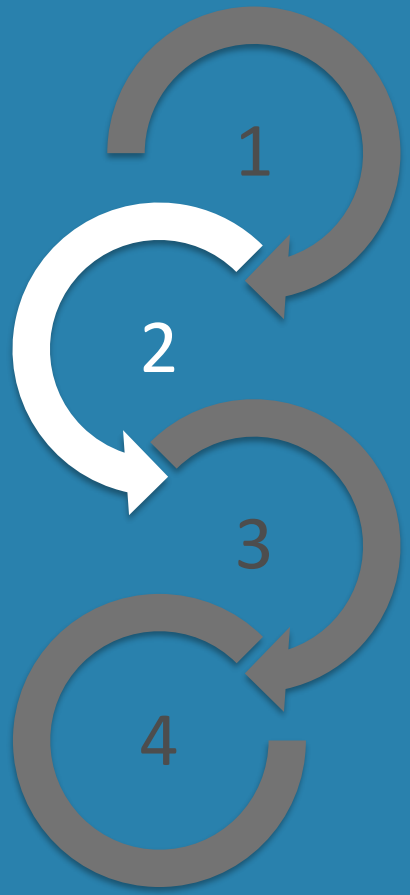

Storage management
Block Storage – EBS

Block storage simulate physical disks
> Raw block-level storage attached to EC2 or RDS instances
> Data written in fixed block size and usually require to be formatting before usage (e.g. NTFS, XFS, …)

Amazon Elastic Block Store (EBS)
> Stored within an availability zone
> Multiple volume types
> Supports snapshots

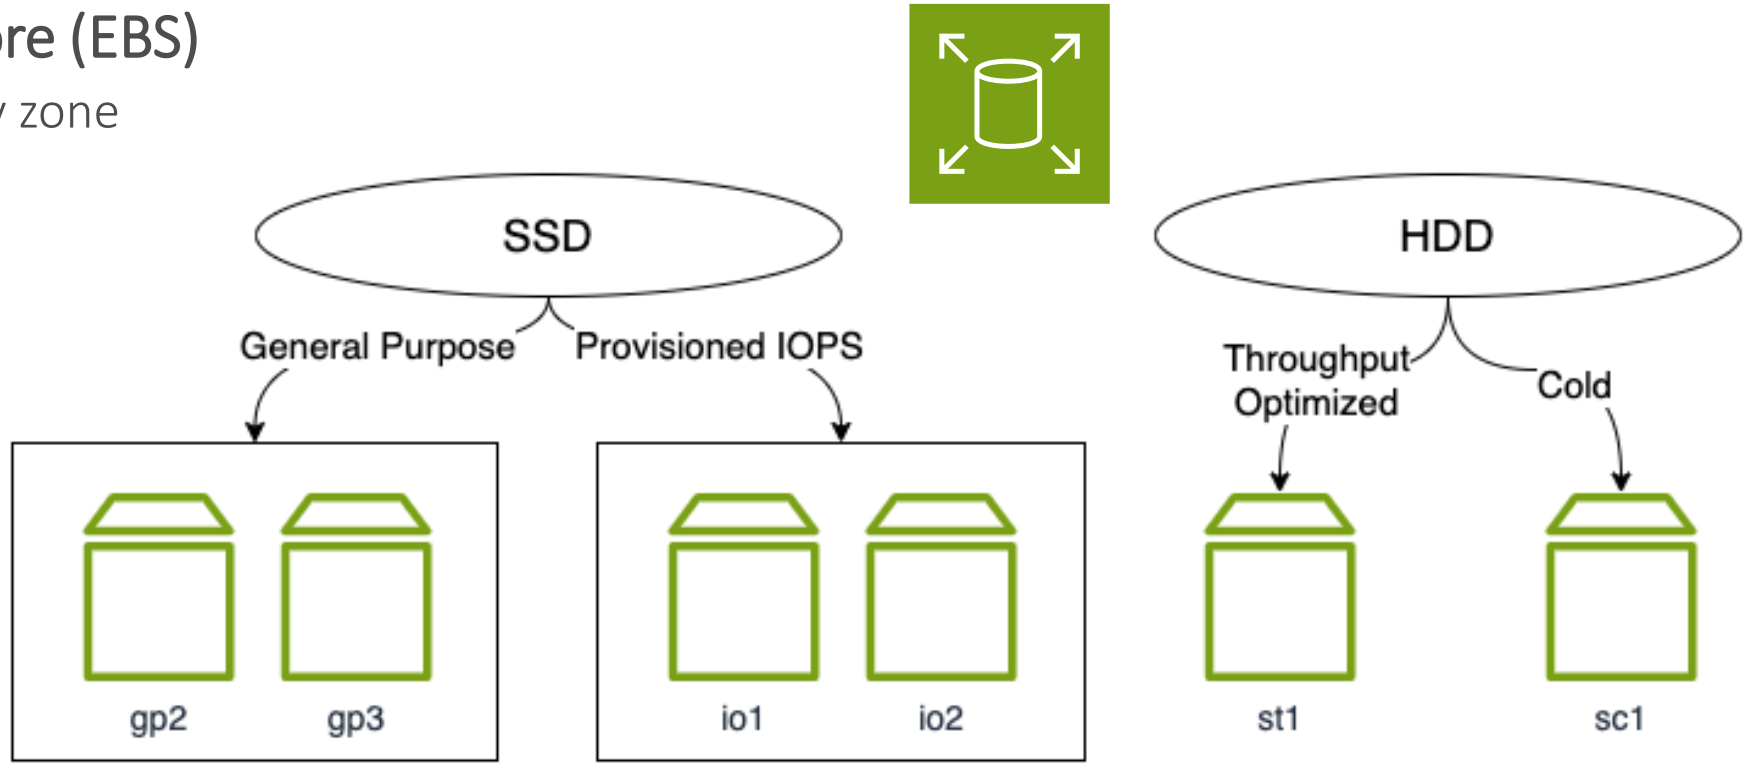

Storage management
Block Storage – EBS characteristics

	General Purpose SSD		Provisioned IOPS SSD		HDD volumes	
Type	gp2	gp3	io1	Io2 Block Express	st1	sc1
Durability	99.8% - 99.9% durability			99.999% durability	99.8% - 99.9% durability	
Size	1 GiB – 16 TiB		4 GiB – 16 TiB	4 GiB – 64 TiB	125 GiB – 16 TiB	
Max IOPS	16 000 (16 KiB) 3 IOPS per GB	16 000 (64 KiB) Min: 3000 IOPS	64 000 (16 KiB)	256 000 (16 KiB)	500	250
Max Throughput	250 MiB/s 128 if < 334 GiB	1000 MiB/s Min: 125 MiB/s	1000 MiB/s	4000 MiB/s	500 MiB/s	250 MiB/s
Multi-attach	Not supported		Supported (Linux)	Supported	Not supported	
Boot	Supported				Not supported	

Storage management
Block Storage – EBS Optimization

Performance depends on volume size for gp2

> Max need to increase volume size to meet performance requirements
> Do not store tempdb on very small dedicated volume

Scenario

> 800 GB database
> 6000 IOPS
> 250 MiB/s throughput

	Option 1	Option 2	Option 3
Volume type	gp2	gp3	io1
Volume size	**2 TB**	1 TB	1 TB
IOPS	6000	6000 (21 $)	6000 (564 $)
Throughput (MiB/s)	250	250 (7.14 $)	500 (256 KiB)
Total (USD)	286 $	142.34 $	**743 $**

Storage management
Block Storage – EBS Optimization

Use striping to bypass IOPS/throughput limitations

Scenario

> 2 TB database
> 32 000 IOPS
> 1000 MiB/S
> (*) io2 not available in all regions

	Option 1	Option 2	Option 3	Option 4 (*)
Volume type	gp2	gp3	io1	io2
Volume size	**2* 5334 GB**	2* 1 TB	2 TB	2 TB
IOPS	2* 16000	2* 16 000 (182 $)	32 000 (3008 $)	32 000 (2496 $)
Throughput (MiB/s)	2*250	2* 500 (42.84 $)	500	500 (16K) – 4000 (256K)
Total (USD)	1525.52 $	453.24 $	**3366 $**	**2794 $**

Storage management
Block Storage – Instance storage

Some instances come with local NVMe SSDs named instance storage

> Highest IOPS and throughput
> Ephemeral and non persistent storage
> May be shared with other instances for smaller instance types

m5d.large	2	8	1 x 75 NVMe SSD	Up to 10	Up to 4,750
m5d.xlarge	4	16	1 x 150 NVMe SSD	Up to 10	Up to 4,750
m5d.2xlarge	8	32	1 x 300 NVMe SSD	Up to 10	Up to 4,750

Use cases

> tempdb storage
> Buffer pool extension (both Standard and Enterprise Edition)
> Cache layer in Storage Spaces / Storage Spaces Direct
> Prepare script to initialize volumes at instance startup

Storage management
File Storage

File storage provides shared access to file data

> Directly available for use
> Use of network protocol (i.e. SMB, NFS)

Amazon FSx for Windows File Server

> Fully managed service
> Provides native Windows compatibility (DFS support, Windows ACLs, end-user file restore, …)
> Replicates data within availability zone and supports multi-AZ for high availability
> Accessible over Server Message Block (SMB) protocol from v2.0 up to 3.1.1

Storage management
File Storage

Advantages
> Pay as you go
> Performance de-coupled from storage size
> Up to 64 TB and 3 GBps per file system

Use cases
> Home directories
> Web serving / Content management
> Witness for Failover clusters
> Shared storage for Failover clusters

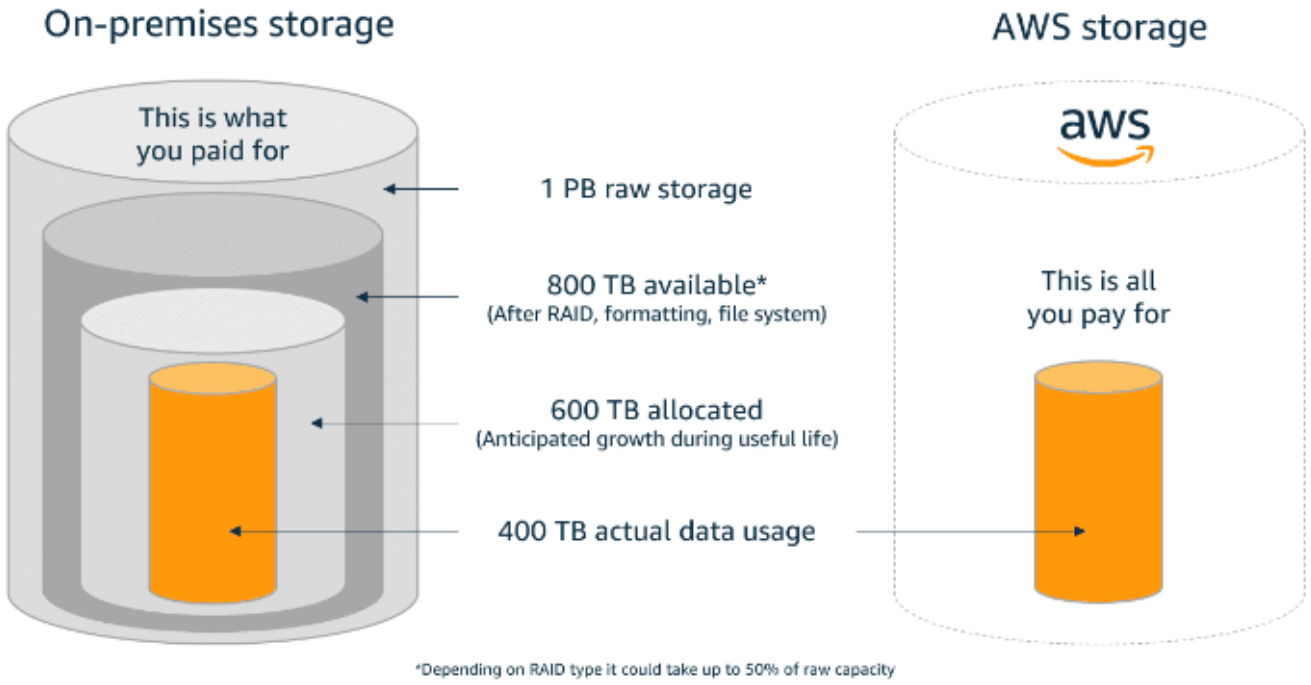

Backup

- Amazon S3
- AWS Storage Gateway
- Archiving and restore

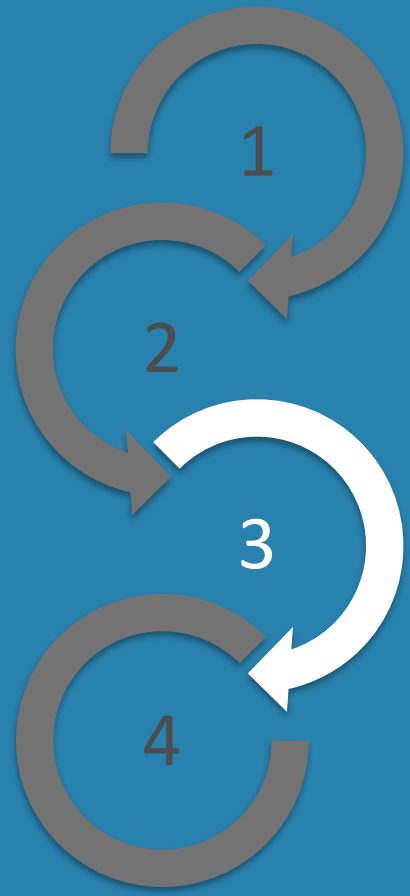

Backup
Amazon S3

Amazon Simple Storage Service (S3) is an object storage

> Access through REST API with HTTP interface
> Store objects within buckets
> No provisioning

Keys features

> Highly durable (99.99999999999%) and available (99.99%)
> Multiple storage classes
> Versioning, Lifecycle, Encryption, …

Amazon S3 costs

> Data transfer OUT – Data IN and transfer within same region, e.g. EC2 instance, is free
> Storage (GB per month)
> Requests, e.g. PUT, COPY, GET (per 1000 requests) – DELETE and CANCEL requests are free

Backup
Amazon S3 – Use cases

Data lakes
> Data can be queries from Amazon Redshift or Amazon Athena

Cheap storage for static content
> Images, Office documents
> Web pages

EBS (gp3)	FSx	S3 Standard	S3 Infrequent Access
0.1142	0.093 – 0.186	0.02695	0.01485

Backup/Archive
> Amazon S3 Glacier storage classes
> Cut prices for long-term storage

Instant Retrieval	Flexible Retrieval	Deep Archive
0.0055	0.004455	0.00198

Backup
Amazon S3 – SQL Server 2022

SQL Server 2022 natively supports S3 endpoint for backup/restore

> Require IAM use to authenticate against Amazon S3 (credential stored in SQL instance)
> Uses S3 multipart upload for each file with maximum part size = MAXTRANSFERSIZE (default 10 MiB)
> Recommended to enable compression to allow bigger MAXTRANSFERSIZE (up to 20 MiB)

```
BACKUP DATABASE db1
TO URL = 's3://sql-backups-2022.s3.eu-central-2.amazonaws.com/backups/db1/db1-part1.bak',
   URL = 's3://sql-backups-2022.s3.eu-central-2.amazonaws.com/backups/db1/db1-part2.bak'
WITH FORMAT, COMPRESSION, MAXTRANSFERSIZE = 20971520;
```

> Backup maximum size is 12,2 TiB (20 MiB * 10000 * 64)
> Supports mirroring but only to another S3 (S3 object replication could be used instead)
> Adding mirrors reduce the maximum number of parts per mirror

Useful resources

> https://aws.amazon.com/blogs/modernizing-with-aws/backup-sql-server-to-amazon-s3/

Backup
AWS Storage Gateway

Mainly designed to provide on-premise accesses to cloud storage

> Amazon S3 File Gateway
> Tape Gateway
> Volume Gateway

S3 File Gateway

> Local virtual machine appliance
> Present S3 buckets as SMB or NFS file shares
> Benefit from local caching
> Optimize data transfers to/from Amazon S3

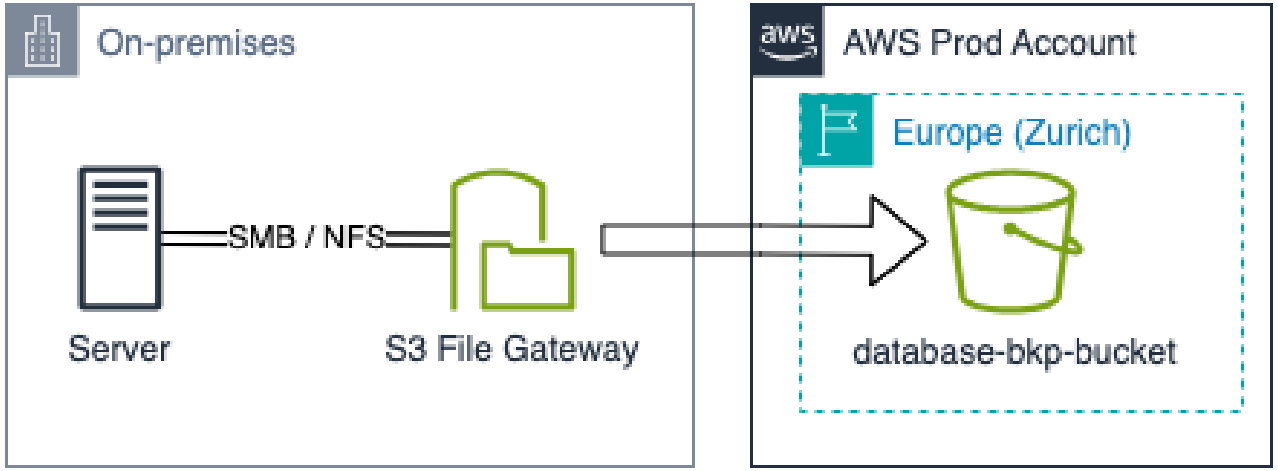

Backup
AWS Storage Gateway – S3 File Gateway

Backup on-premise database to the cloud
> Get off-site backup copy
> Ease Disaster Recovery in the cloud

Use EC2 instance as S3 File Gateway
> Backup any version of SQL Server to SMB compatible file share
> Reduce cost for backup storage by eliminating local EBS volume

Tips and tricks


> Bucket dedicated for database backup with bucket policy for access control
> Use lifecycle policy or S3 Intelligent Tiering (> 30 days)
> File upload to S3 outside of the storage gateway are not visible by default
> **Objects stored in Glacier Flexible Retrieval and Glacier Deep Archive must be restored before access**
> Prepare restore scripts and procedures ahead of time

Backup
Archiving and restore

List Objects

```
PS /Users/nij> Get-S3ObjectV2 -BucketName dbi-sgw-database-backups
 -Prefix SQLServer/DB1/ | Select-Object BucketName, Key, Size, StorageClass


BucketName              Key                                Size StorageClass
----------              ---                                ---- ------------
dbi-sgw-database-backups SQLServer/DB1/                        0 STANDARD
dbi-sgw-database-backups SQLServer/DB1/20240825/               0 GLACIER
dbi-sgw-database-backups SQLServer/DB1/20240825/db1.bkp 104857600 GLACIER
```

Restore Objects

```
PS /Users/nij> Restore-S3Object -BucketName dbi-sgw-database-backups
 -Key SQLServer/DB1/20240825/db1.bkp -Tier Standard -CopyLifetimeInDays 1
```

> https://www.dbi-services.com/blog/restore-s3-object-with-awspowershell

High availability

- Amazon RDS Multi-AZ
- Amazon FSx

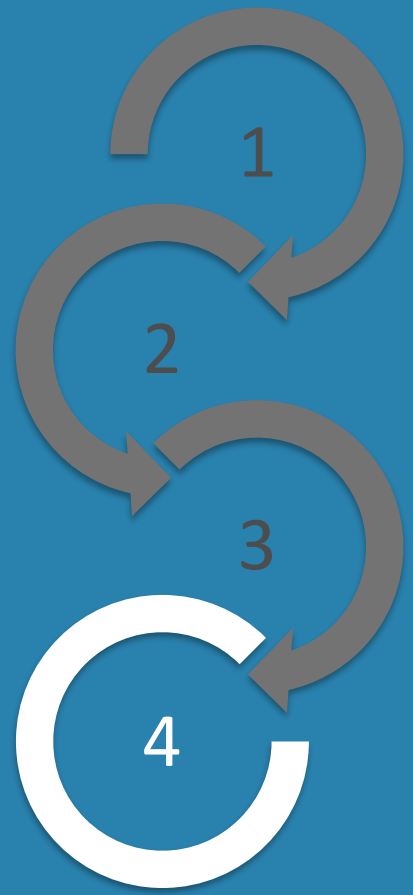

High Availability
Amazon RDS Multi-AZ

Managed replication across 2 Availability zones

> SQL Server Database Mirroring (DBM) or Always On Availability Groups (AGs)

> Automated failovers

> No access to standby database

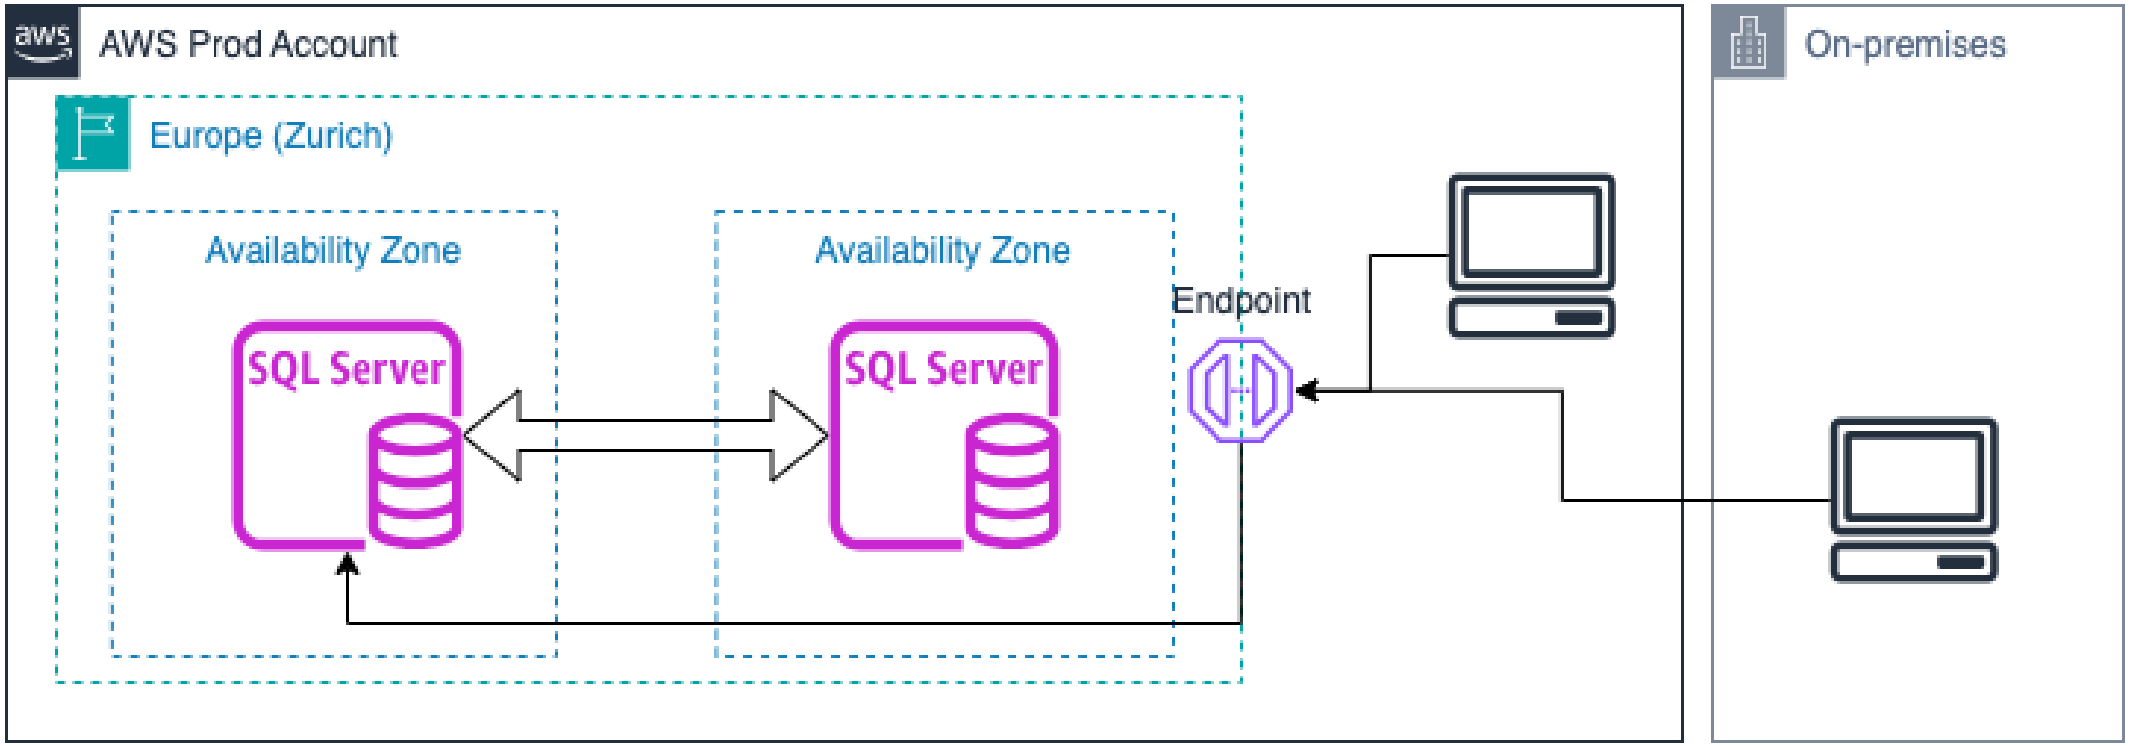

High Availability
Amazon RDS Multi-AZ

Endpoint just points to standby in case of failover

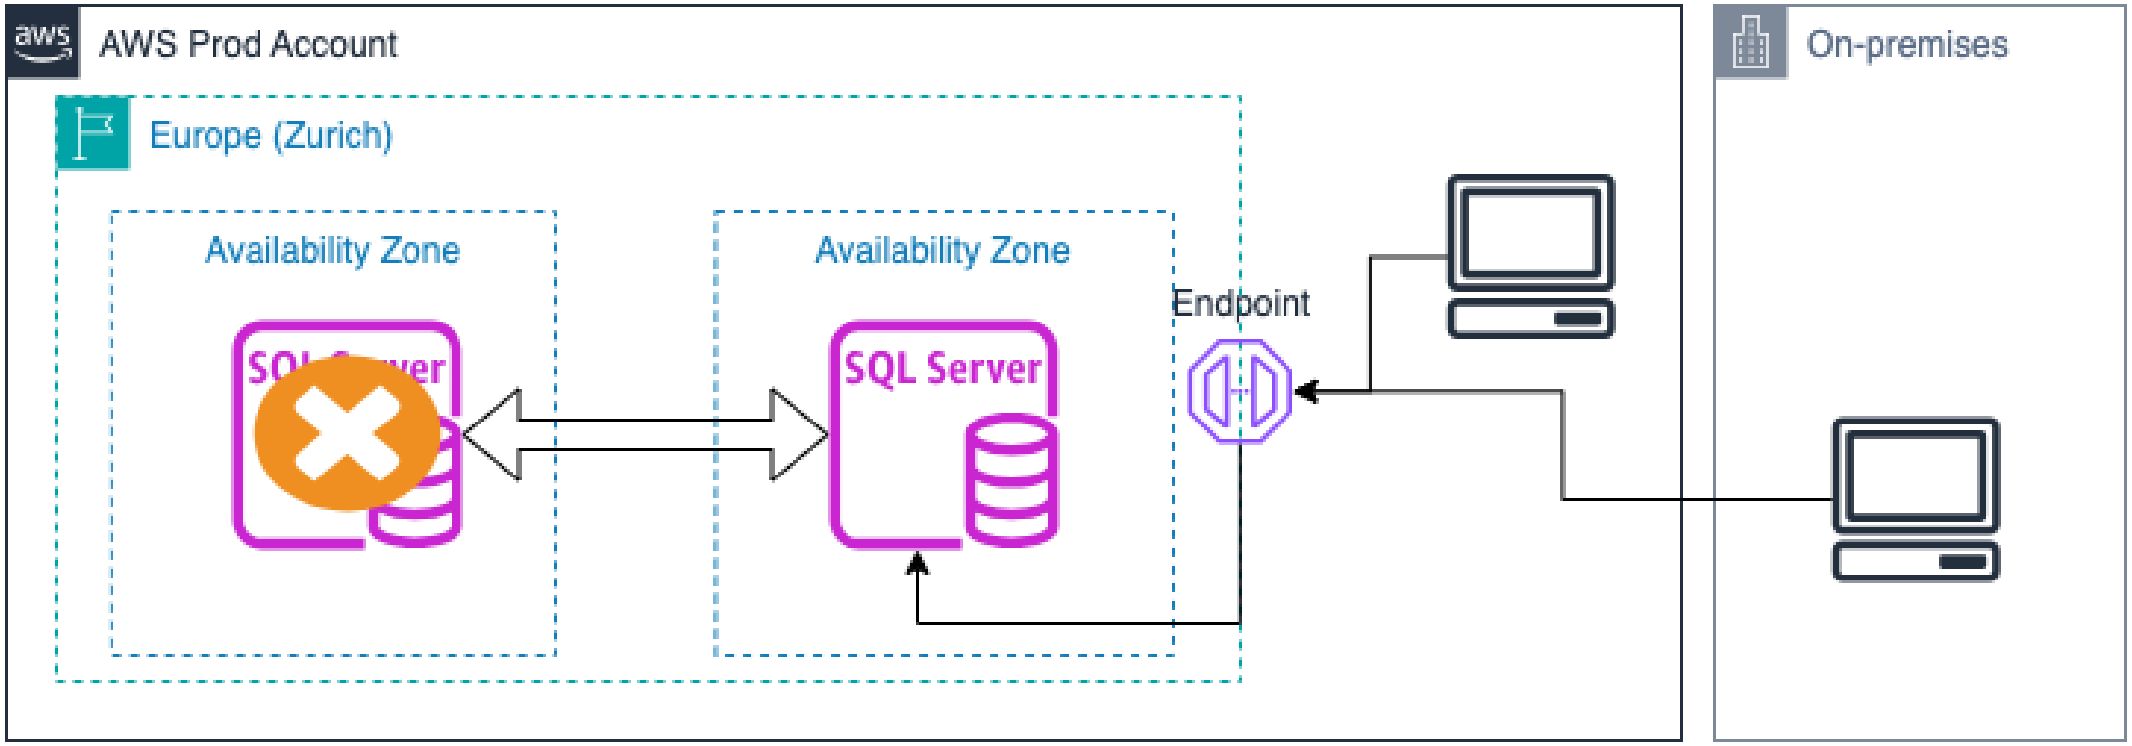

High Availability
Amazon FSx

Amazon FSx is available with Multi-AZ

> Data replicated automatically in 2 Availability zones
> Supports for SMB Continuously Available (CA) / SMB Transparent Failover file shares
> Data transfer between AZ for replication is free of charge

Build SQL Always On Failover Cluster Instance (FCI)

> Amazon FSx used twice
  > Multi-AZ file share for shared storage (Data)
  > Single-AZ file share for cluster witness (Quorum)
> Use the 2 same AZ for FSx than for EC2 to reduce latency and avoid cross AZ charges
> Use a 3rd AZ for file share witness to have quorum in case of AZ failure

> https://aws.amazon.com/blogs/storage/simplify-your-microsoft-sql-server-high-availability-deployments-using-amazon-fsx-for-windows-file-server/
> https://aws-ia.github.io/cfn-ps-microsoft-sql-fci-fsx/

High Availability
Amazon FSx

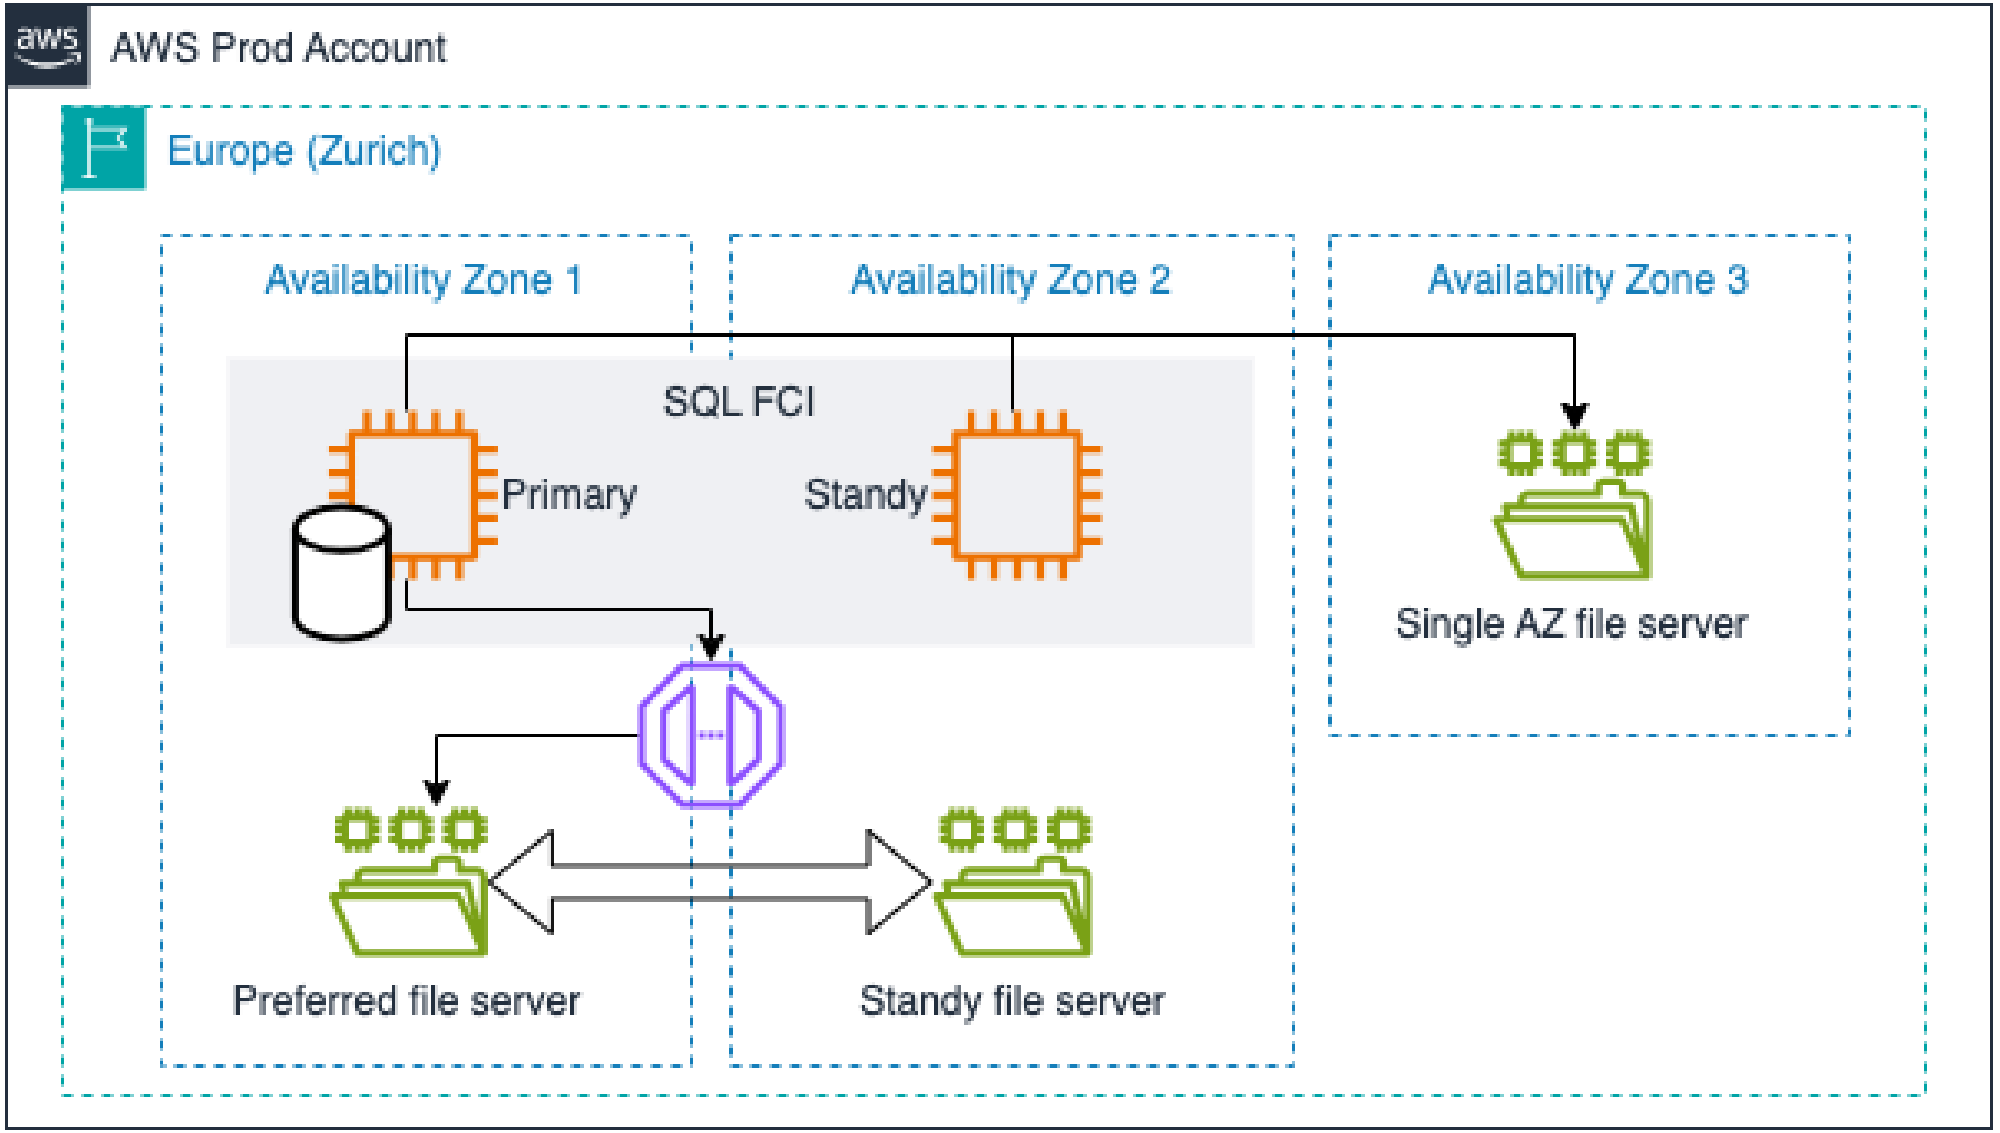

High Availability
Amazon FSx

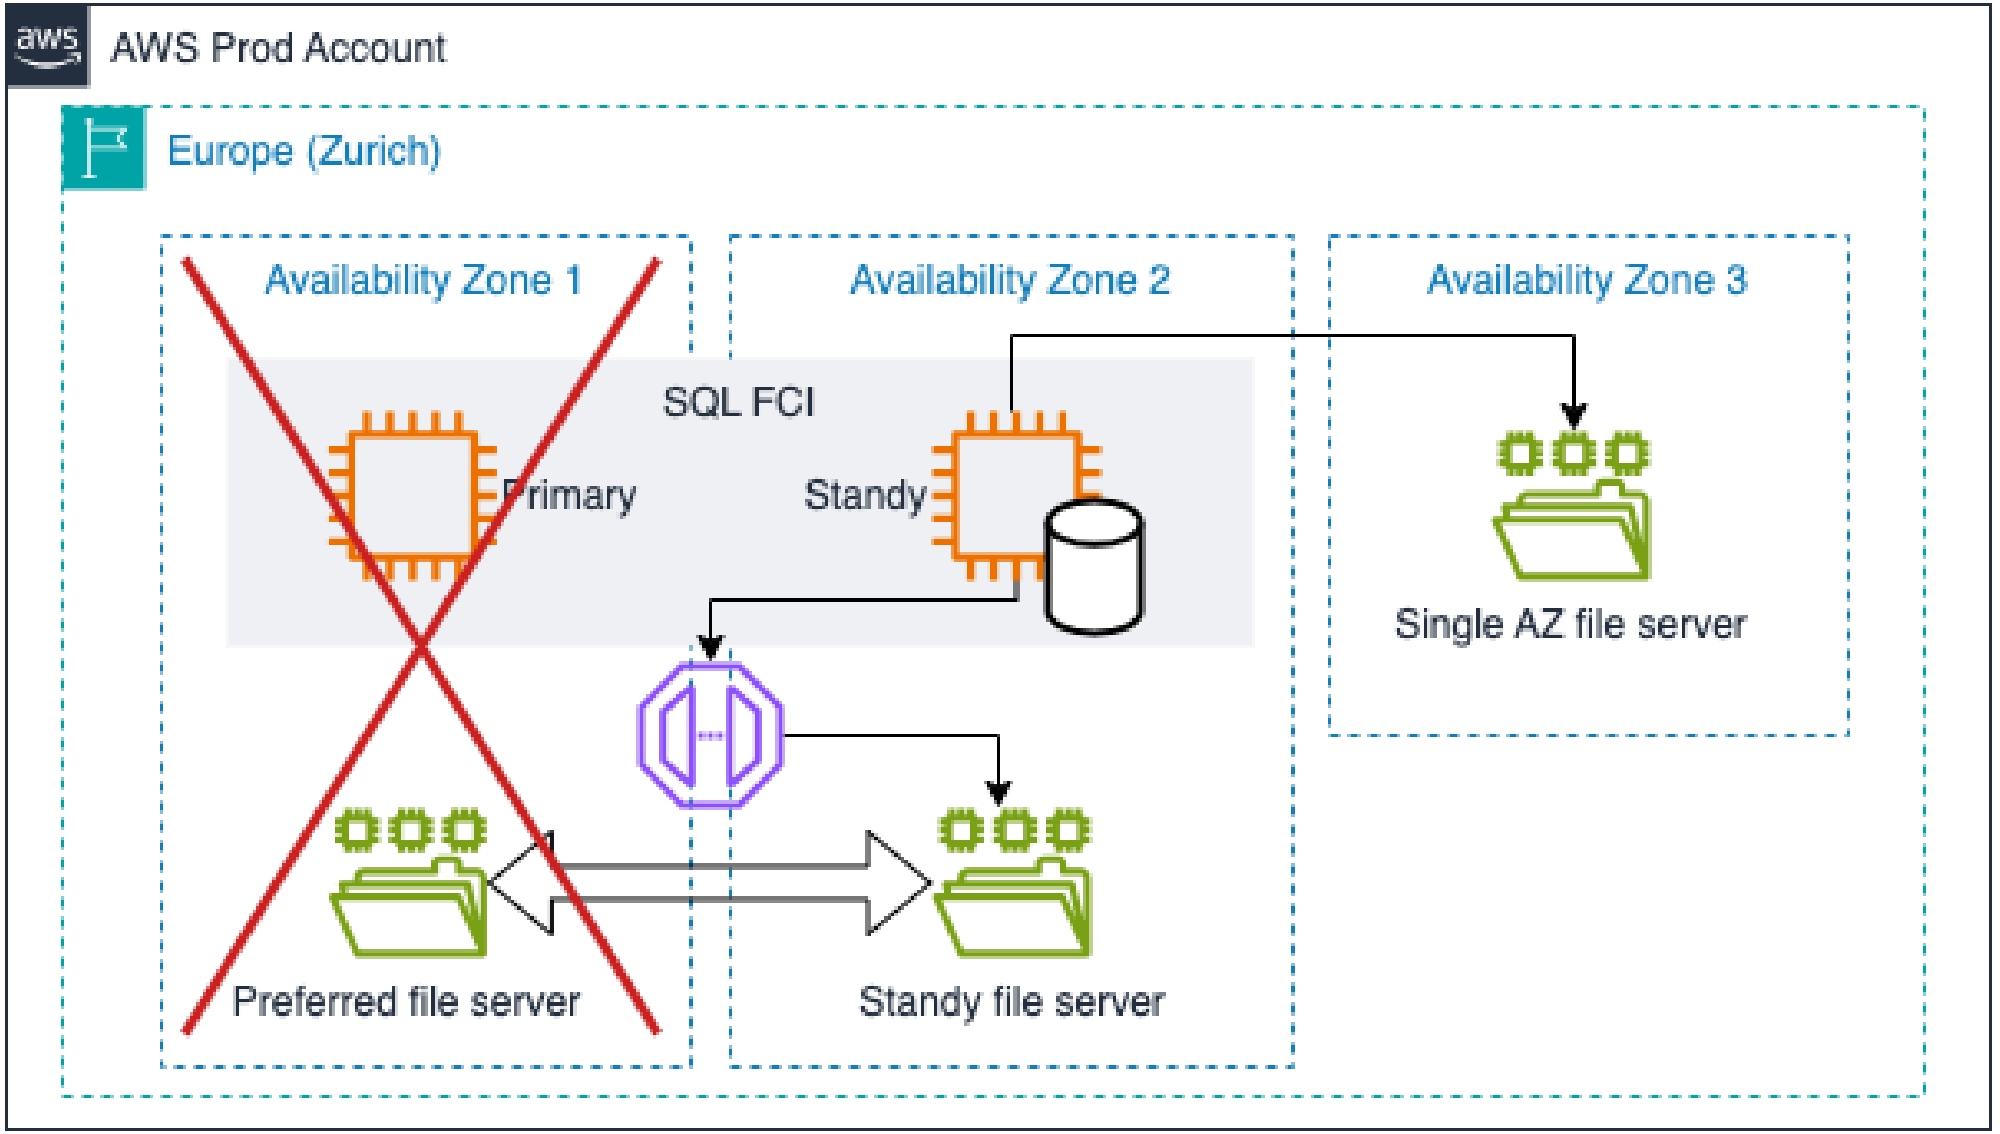

Any questions?

Please do ask!

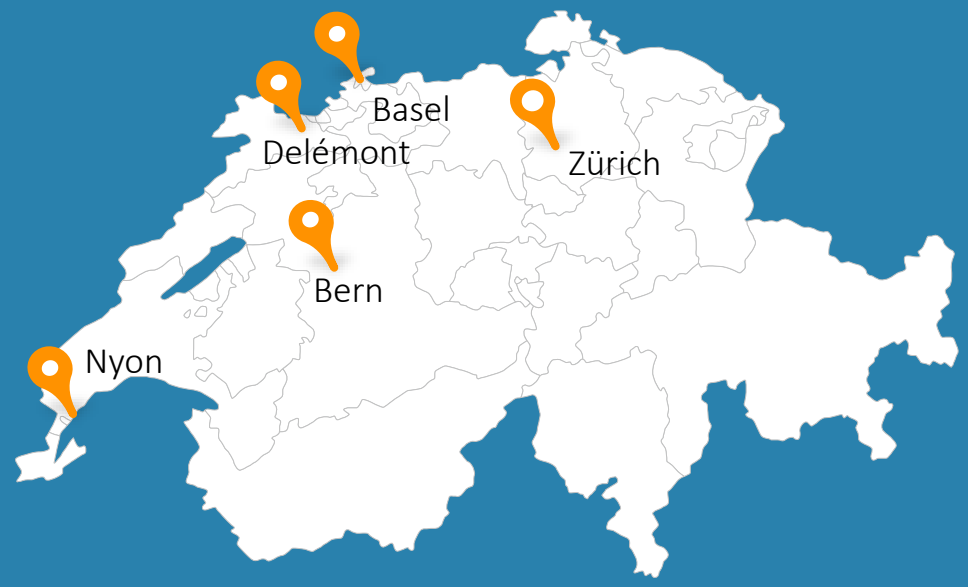

We would love to boost your IT-Infrastructure

How about you?