Azure SQL Database Resiliency and Availability

EXPLAINED

Rolf Tesmer

Senior Architect. Azure | Data | Analytics | AI

Linked In: https://www.linkedin.com/in/rolftesmer/

Blog: https://mrfoxsql.com.au/

Thursday 01 Dec 2022

Microsoft

My Top 5 Assumptions for Today

1. You have used the **Azure** cloud
2. You have **knowledge** of **SQL DB Services**
3. You have **awareness** of **core Azure Services**
4. You have used the **Azure Portal** and **SSMS**
5. You will **rate** my session **10 / 10**

(ie. stick around for the feedback survey at the end!)

Today's Basic Program

```
10 Azure Cloud Primer
20 Azure SQL Database Primer
30 Explore HA/DR Options
40 Activating HA/DR
50 Wrap Up & Summary
```

10. Azure Cloud Primer

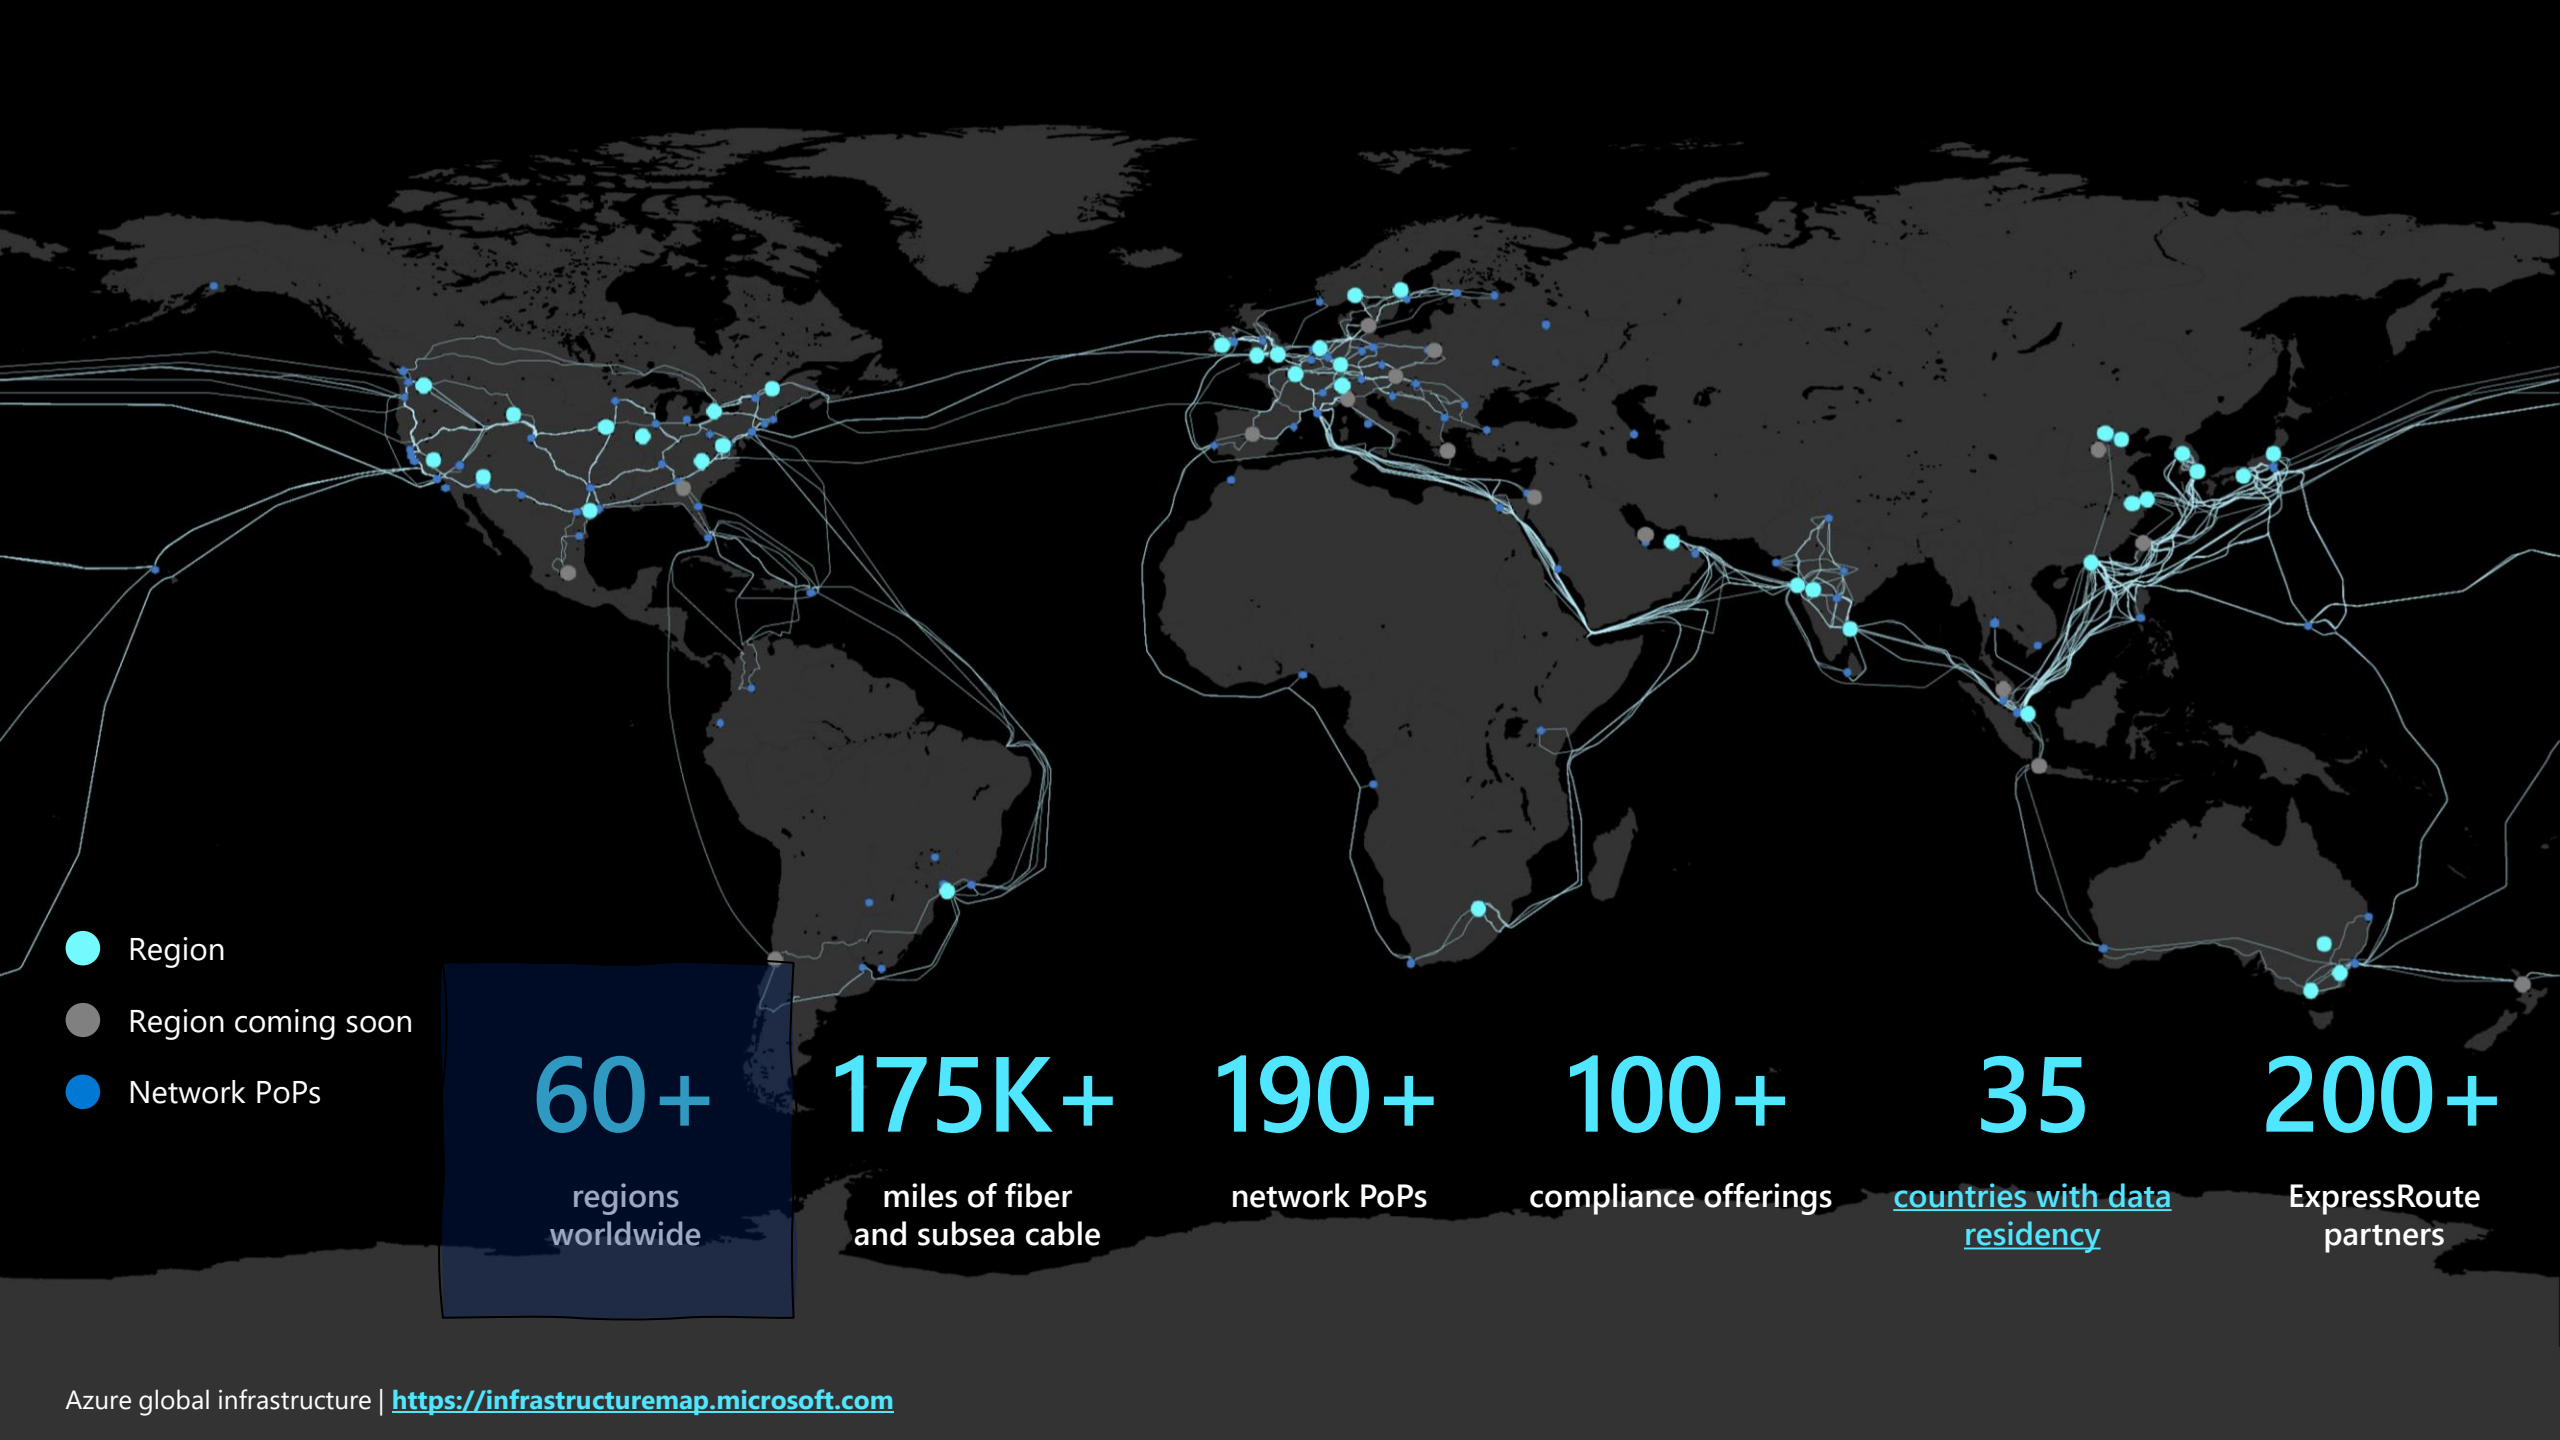

Azure global infrastructure | https://infrastructuremap.microsoft.com

Cloud Service Models

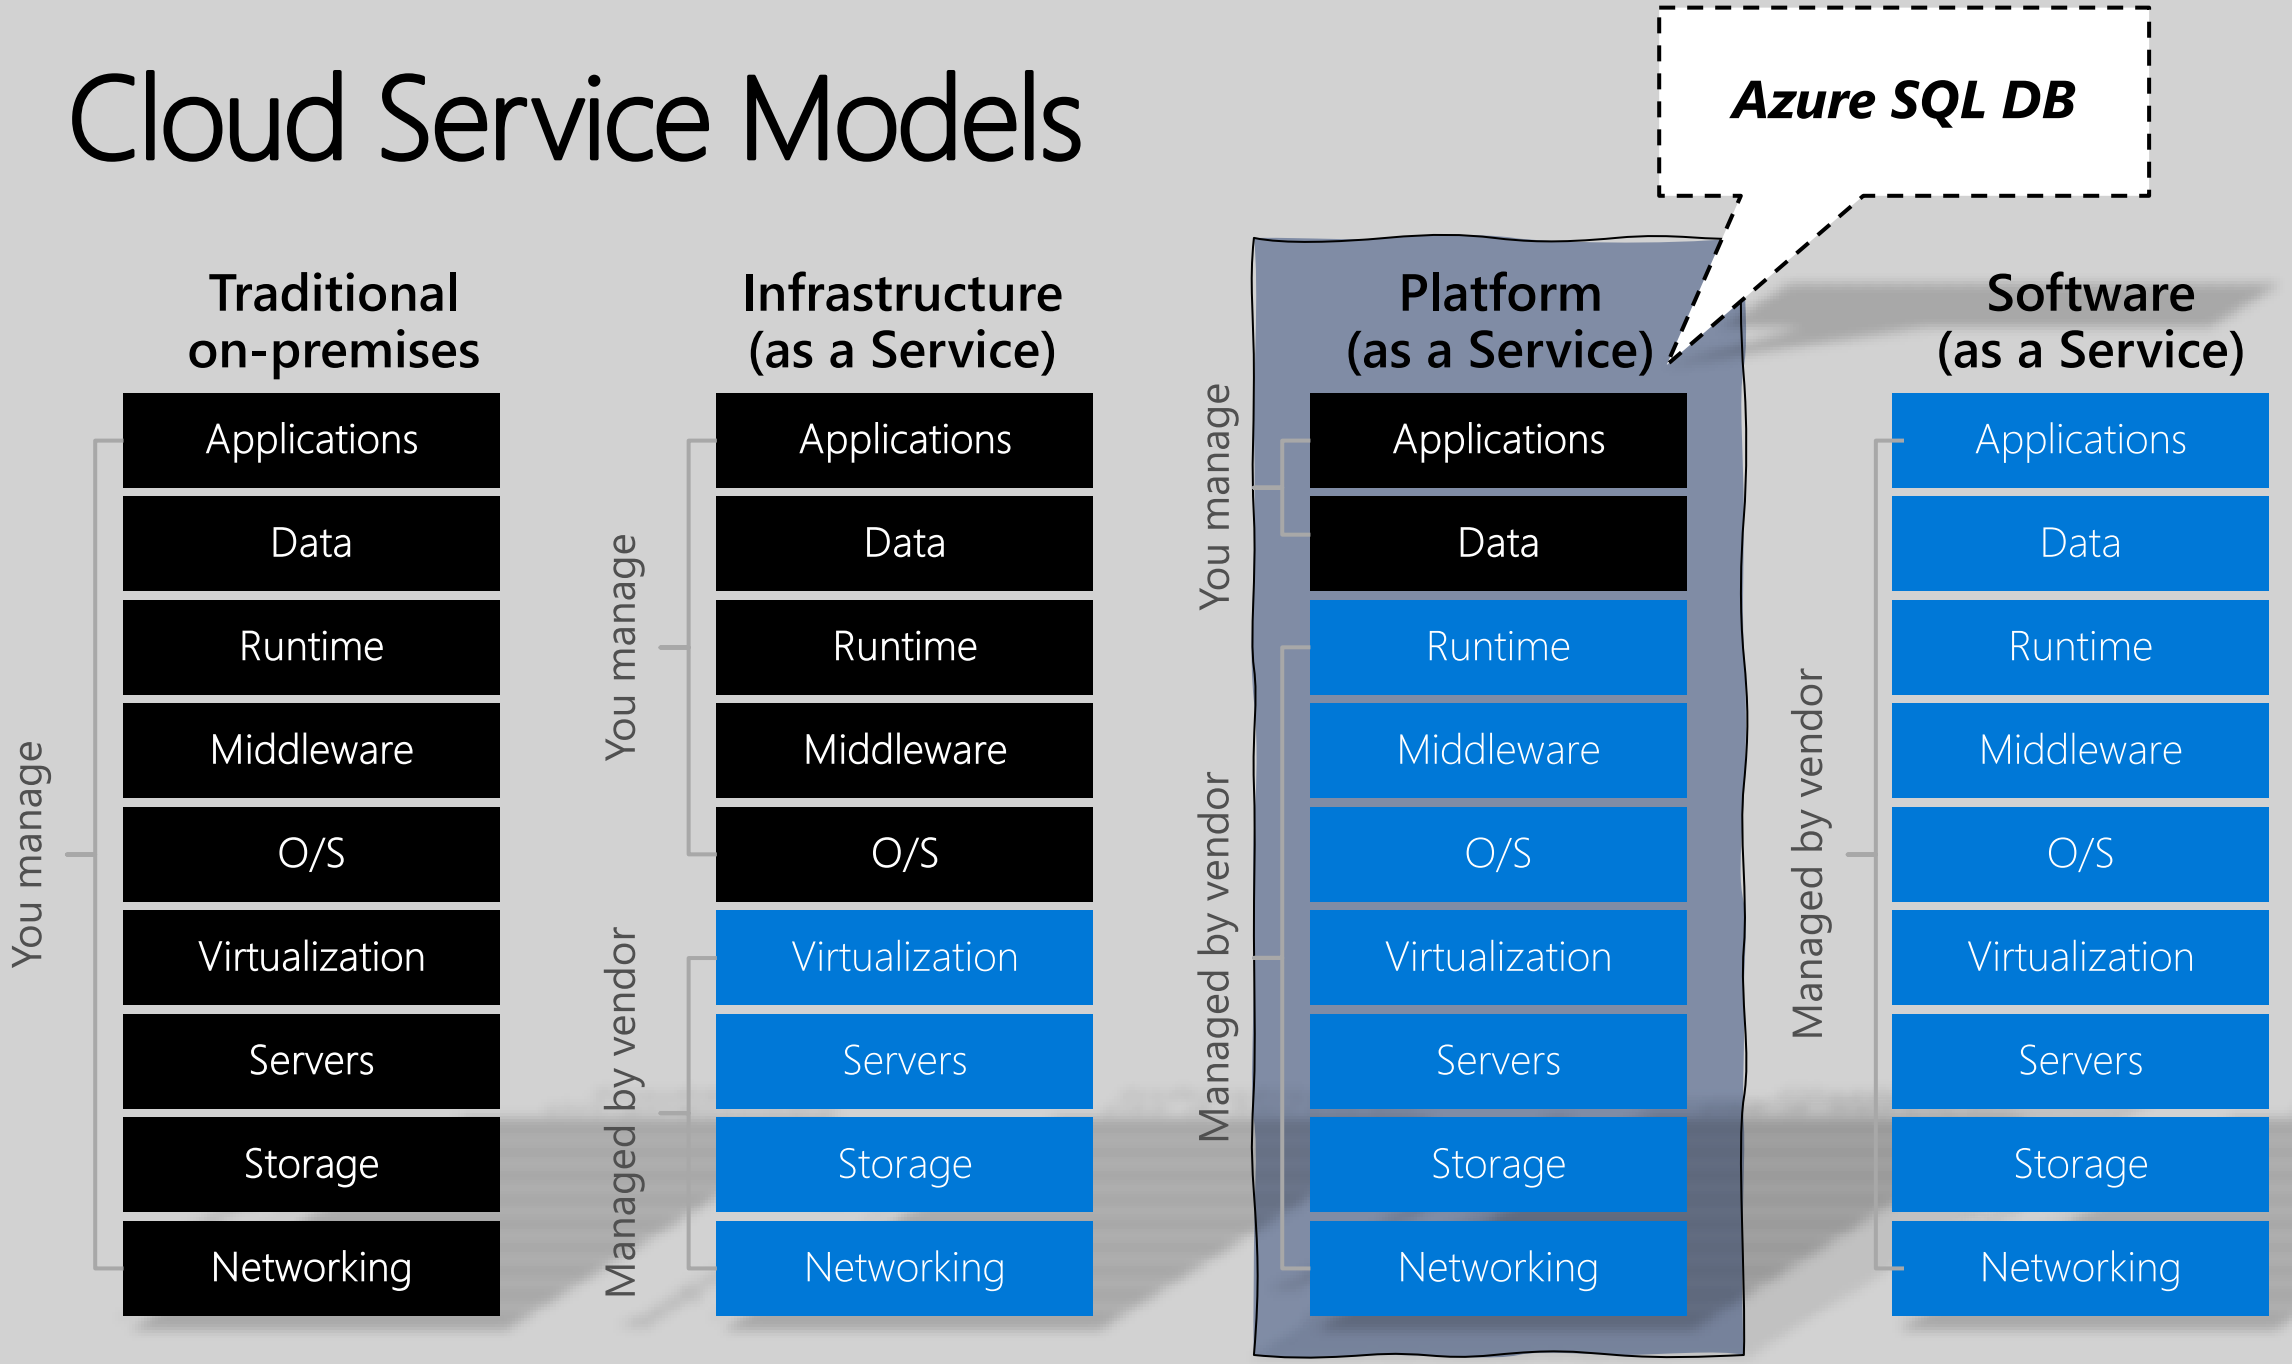

Solutions for SQL Server

Just talking about this stack

SQL Server on Azure VMs	Azure SQL Managed Instance	Azure SQL Database / Hyperscale / Serverless / Elastic Pools	Azure SQL Edge*
	Intelligent performance/security	Intelligent performance/security	Intelligent performance/security
Applications	Applications	Applications	Applications
Data	Data	Data	Data
Database	Database	Database	Database
SQL instance-level features	SQL instance-level features		
High Availability /DR/Backups	High Availability/ DR/Backups	High Availability/ DR/Backups	High Availability/ DR/Backups
Database provision/ Patch/Scaling	Database provision/ Patch/Scaling	Database provision/ Patch/Scaling	Database provision/ Patch/Scaling
Operating system	Operating system	Operating system	Operating system (container)
Virtualization	Virtualization	Virtualization	Container Platform
Hardware	Hardware	Hardware	Hardware & Operating System
Datacenter management	Datacenter management	Datacenter management	Device management (IoT Hub)

Managed by customer

Managed by Microsoft

Machine learning capability

*in connected scenario

20. Azure SQL DB Primer

Deployment Models: Single vs Elastic Pools

Azure SQL Database
Database scoped deployments

Single database

Simplest database scoped model provides resource guarantees.

Best for modern cloud born apps requiring the most management simplicity

Elastic pool

Shared resource model lowers TCO of multi-database environments

Best for multi-tenant SaaS apps requiring database isolation per tenant

Deployment Models: DTU vs vCore

DTU model

Simple, preconfigured

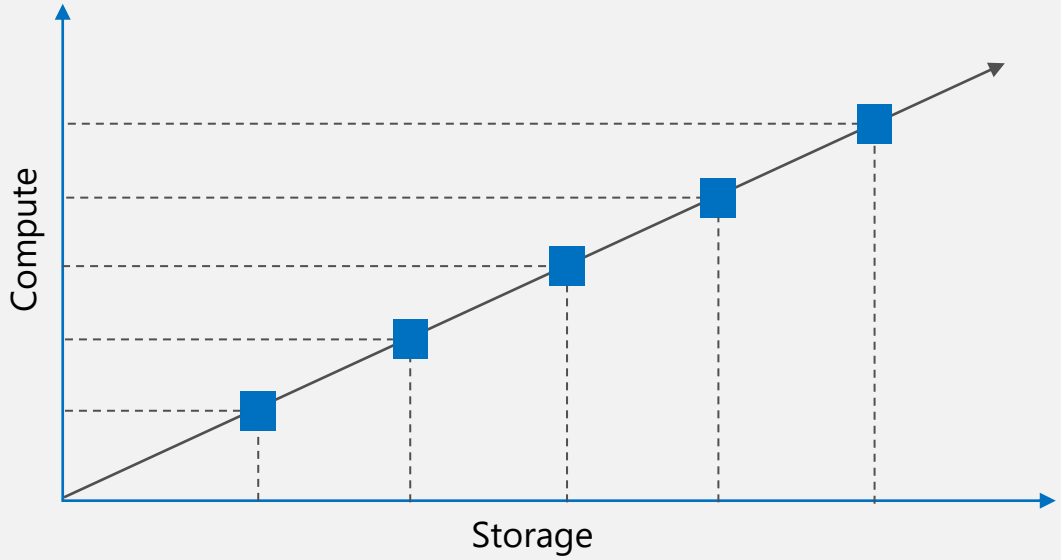

Pre-packaged, bundled unit that represents the allocated "database power" assigned to the resource

DTU sizing offers simplicity of choice – slider gives compute and storage allocation – but cannot align AHUB for SQL Server

Designed for predictable performance, but somewhat inflexible and limited in options

vCore model

Independent scalability

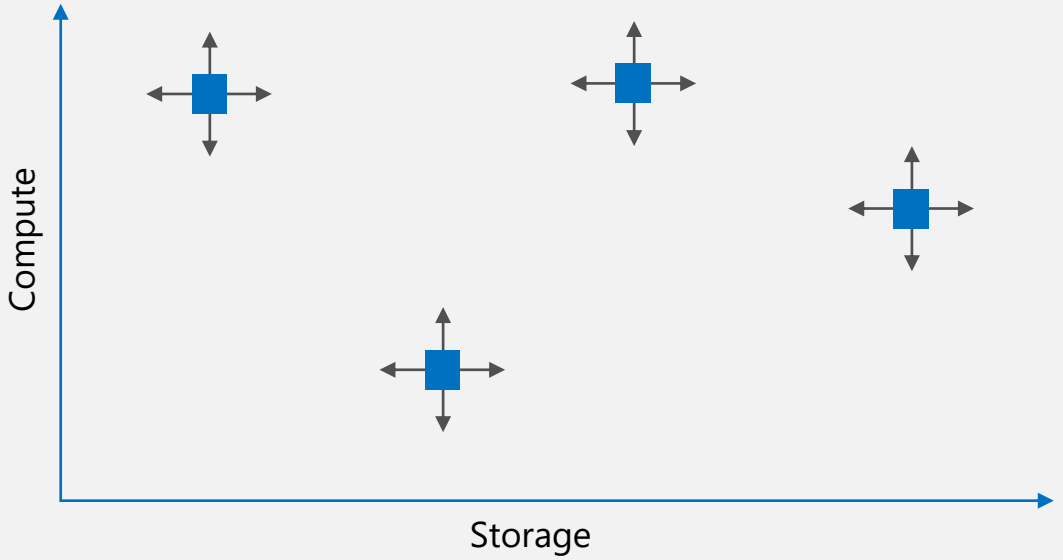

This model allows you to independently choose compute and storage resources.

It also allows you to use AHUB for SQL Server to gain cost savings (as the vCores are directly exposed)

Storage allocation covers backups and live database files

Customers who value flexibility; control and transparency

SQL Database SKU's and Deployment Types

[Purchasing models - Azure SQL Database | Microsoft Learn]

Serverless Compute Option *Serverless SKU currently only supported in General Purpose tier on Generation 5 hardware – and only in the vCore model.*

DTUs			vCores		
Basic	Standard	Premium	General Purpose	Business Critical	Hyperscale
Small databases particularly those in development phases	***General purpose databases with moderate performance requirements***	***Mission-critical databases with high performance and high-availability requirements***	***Data applications with basic IO and basic availability requirements***	***Business critical data applications with fast IO and high availability requirements***	***VLDB OLTP and HTAP workloads with highly scalable storage and read-scale requirements***

Elastic scale and performance: Three service tiers within DTU-based model, and two tiers within vCore-based model let you scale up and down based on throughput needs, and offer better resource isolation and an improved billing experience

Business continuity and data protection: A spectrum of business-continuity features across tiers lets you dial up control over data recovery and failover. Its important to note not all SKU's provide the same levels of availability and resiliency options

Familiar and fully-managed: Near-complete SQL Server compatibility and unprecedented efficiencies as your applications scale with a near-zero maintenance service and a variety of familiar management tools and programmatic APIs

Azure SQL Database Availability at 1000 Feet

What's in a ***0.005% ?***

Yearly SLA Difference...
99.99% = 52min 9.8s
99.995% = 26min 4.9s
https://uptime.is/

Reliability Built Into the Azure Cloud Fabric

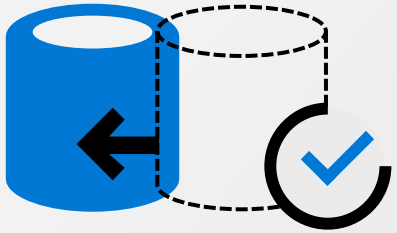

Built-in: service fabric availability

Optional: geo-replication and auto-failover groups powered by Always-On technology

Optional: Zone redundancy improves resilience against in region failures without app changes

Reliability Built Into the Azure Backup Processes

Built-in: Automated backups and point-in-time restore up to 35 days. Backups can be LRS, ZRS or GRS.

Built-in: Accelerated database recovery enables fast and consistent recovery regardless of number of active transactions

Optional: LTR Backups

All Backed by Industry-leading SLA

99.995%

Industry-leading availability SLA up to **99.995% with zone-redundant deployment**

Non-zone-redundant is 99.99%

Optional: With geo-replication deployment
@ RPO of 5 seconds - 100% deployed hours
@ RTO of 30 seconds - 100% deployed hours

Active geo-replication maintains a continuously synchronized secondary in the same or different region.

[Service Level Agreements Summary | Microsoft Azure]

30. Exploring HA & DR Options

What's a Typical Business Continuity Requirement?

Customer Objective:

We want to enable our business application to continuously operate during unplanned disruptive events - and any planned maintenance activities.

Make it so.

What typical scenarios *could* occur which make your life difficult?

1. Microsoft upgrade/maintain their Service Stack, or same per Customer on their App Stack
2. Local network, hardware or software failures within the service stack
3. Datacenter outage (ie power, network, etc) - possibly caused by a regional disaster.
4. Data corruption or deletion (ie caused by application bug, human error, hack, etc)

Cloud business continuity - database recovery - Azure SQL Database & SQL Managed Instance | Microsoft Learn

Azure SQL Database – HA/DR Options

1. In-Region Standard Availability
2. In-Region Zone Availability (AZ)
3. Backups, and Restore Options
4. Geo-replication – with secondary
5. Geo-replication – with failover groups

https://docs.microsoft.com/en-us/azure/sql-database/sql-database-business-continuity#sql-database-features-that-you-can-use-to-provide-business-continuity

Azure SQL Database - HA/DR Options

1. In-Region Standard Availability

https://docs.microsoft.com/en-us/azure/sql-database/sql-database-business-continuity#sql-database-features-that-you-can-use-to-provide-business-continuity

High Availability (Built-In Availability Models)

Standard Availability Model → DTU (Basic/Standard), vCore (General Purpose)

based on a separation of compute and storage

relies on high availability and reliability of the remote storage tier

targets budget-oriented business applications that can tolerate some performance degradation during maintenance activities.

Premium Availability Model → DTU (Premium), vCore (Business Critical)

based on a cluster of database engine processes

relies on the fact that there is always a quorum of available database engine nodes

targets mission-critical applications with high IO performance, high transaction rate and guarantees minimal performance impact to your workload during maintenance activities.

https://docs.microsoft.com/en-us/azure/sql-database/sql-database-high-availability

High Availability (the Standard Availability Model)

DTU (Basic, Standard), vCore (General Purpose)

Behaves like Failover Cluster Instance (FCI)

ie no ability to read from the replicas

Remote storage provides data redundancy within an Azure DC

Backups in different location with geo-redundancy (optional)

Client connectivity redirection built-in

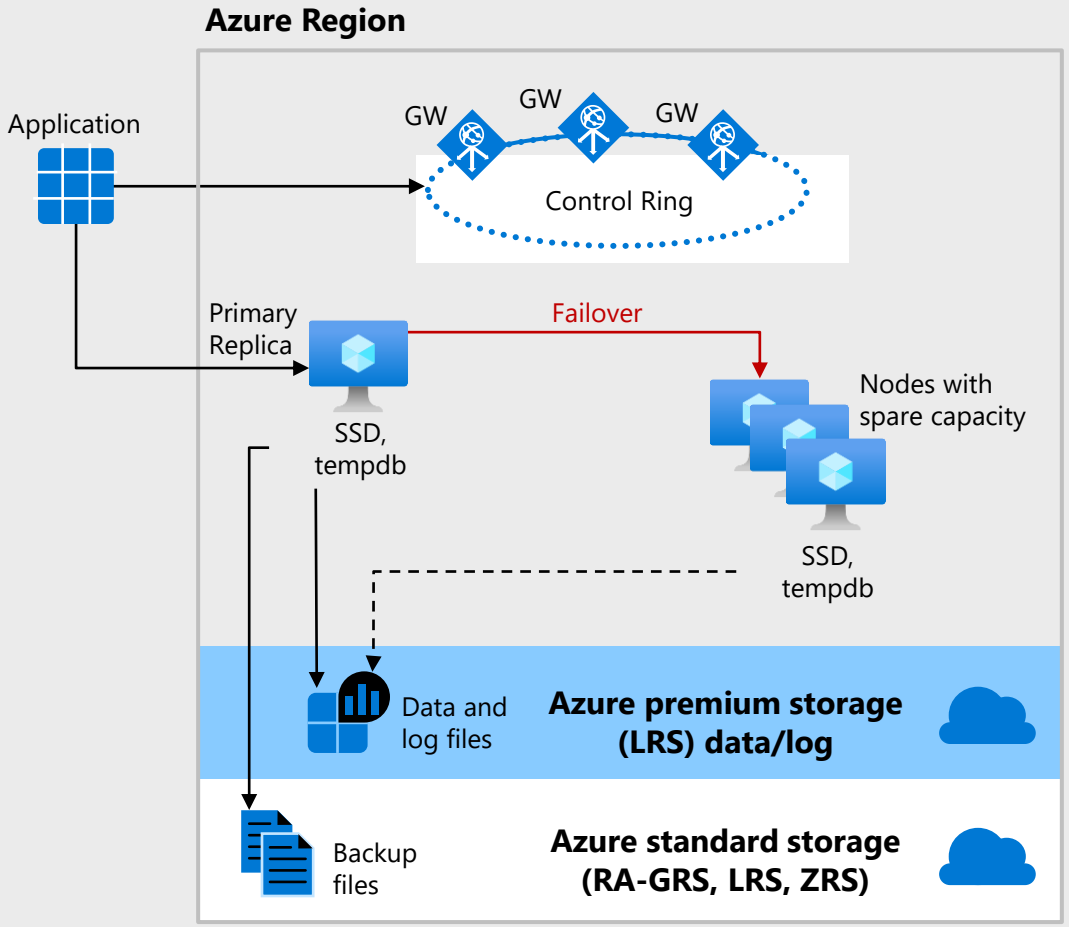

https://docs.microsoft.com/en-us/azure/sql-database/sql-database-high-availability

High Availability (the Premium Availability Model)

DTU (Premium), vCore (Business Critical)

High Availability achieved by replicating both compute and storage to additional nodes (ie no shared storage)

implemented using SQL Always-On Availability Groups

AAG cluster includes **single primary replica** (read-write workloads), and up to three **secondary replicas**.

By default includes “Read Scale-Out” feature.

ie provides ability to read from secondary replicas using “ApplicationIntent=ReadOnly”

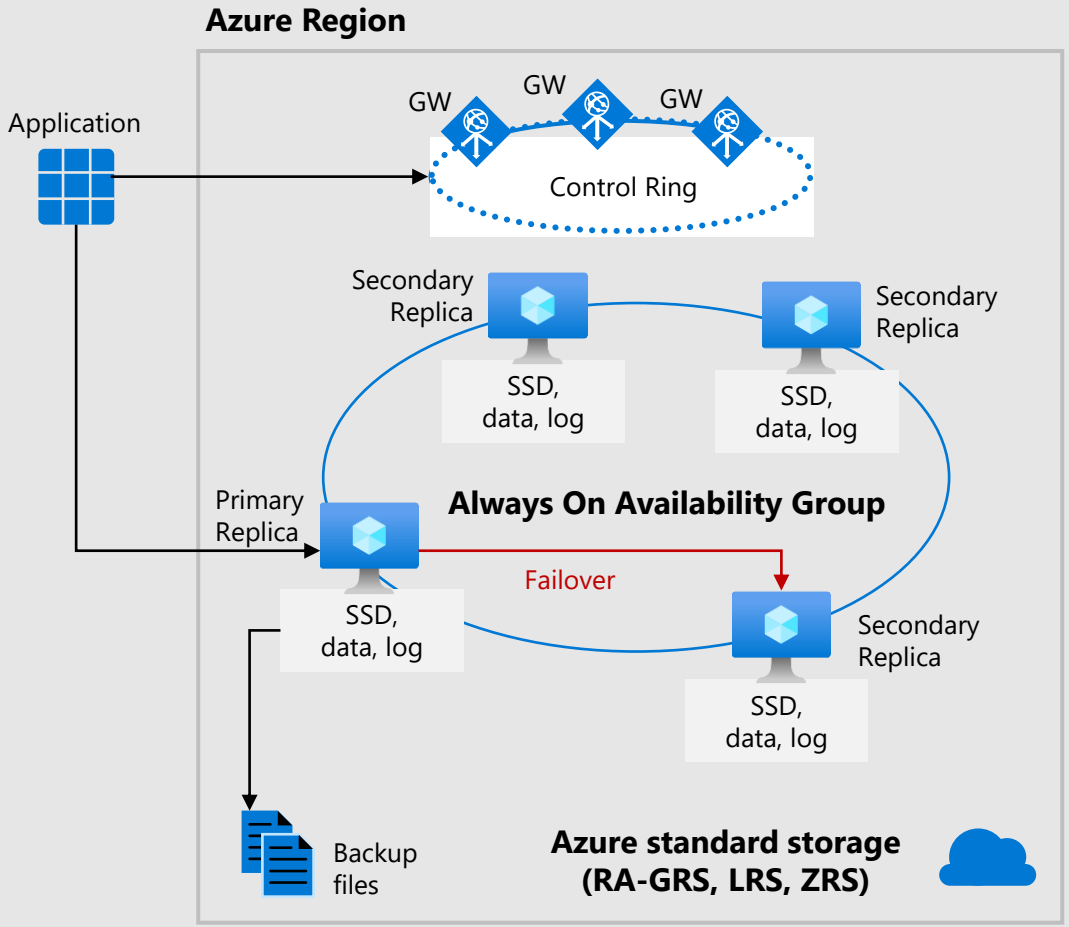

https://learn.microsoft.com/en-us/azure/azure-sql/database/read-scale-out?view=azuresql

https://docs.microsoft.com/en-us/azure/sql-database/sql-database-high-availability

Demo #1

- Reading from a secondary replica database using "*ApplicationIntent=ReadOnly*"
- [Read queries on replicas - Azure SQL Database & SQL Managed Instance | Microsoft Learn]

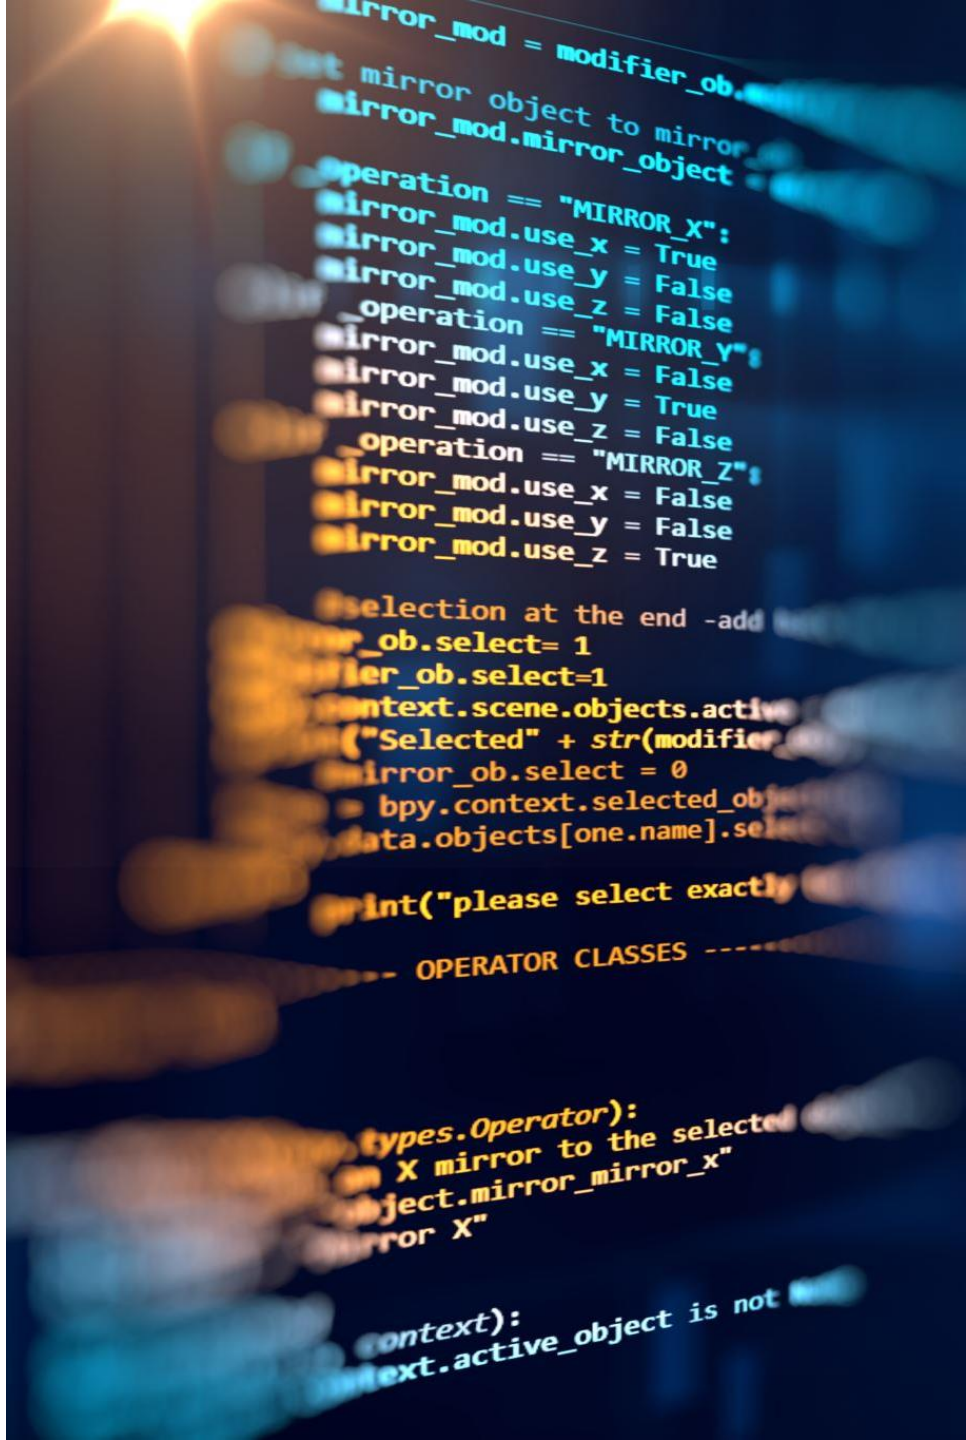

Azure SQL Database - HA/DR Options

2. In-Region Zone Availability (AZ)

https://docs.microsoft.com/en-us/azure/sql-database/sql-database-business-continuity#sql-database-features-that-you-can-use-to-provide-business-continuity

SQL DB in Availability Zone Redundant Configuration

Different database replicas can be placed in different availability zones in the same region.

Replication between zones is always Synchronous.

Zone redundancy configuration is free to implement, however may have performance overhead (ie distance between DC's)

Client routing controlled by Azure Traffic Manager

Supported in Premium (DTU), Business Critical (vCore), or Hyperscale (vCore)

NOTE - General Purpose (vCore) is currently "*Public Preview*"

Creating a Zone Redundant database on a Logical Server also adds the "master" database to the zone to protect server level resources

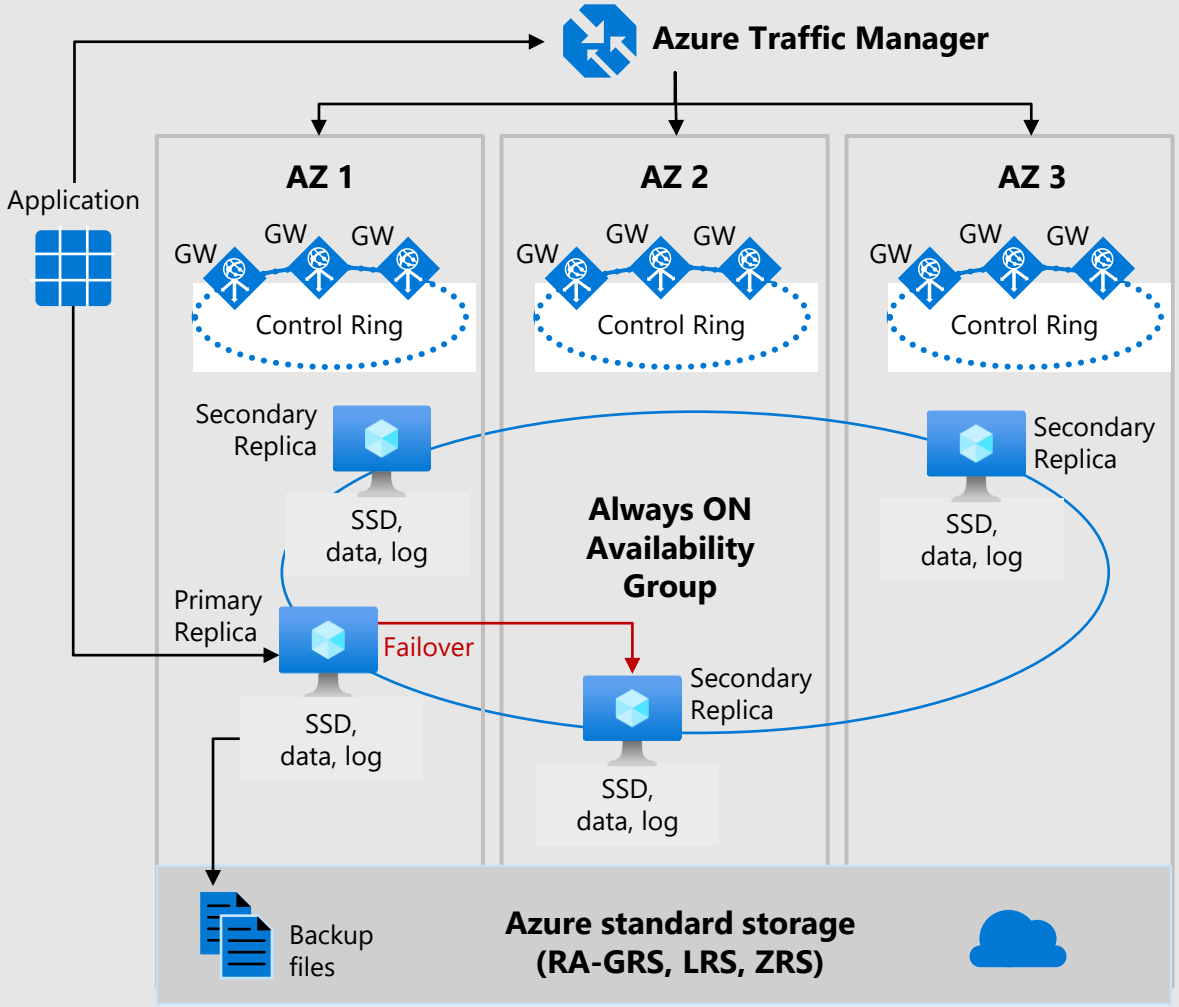

[High availability - Azure SQL Database and SQL Managed Instance | Microsoft Learn]

Demo #2

- Check the AZ status of databases
- [Get-AzSqlDatabase (Az.Sql) | Microsoft Learn]

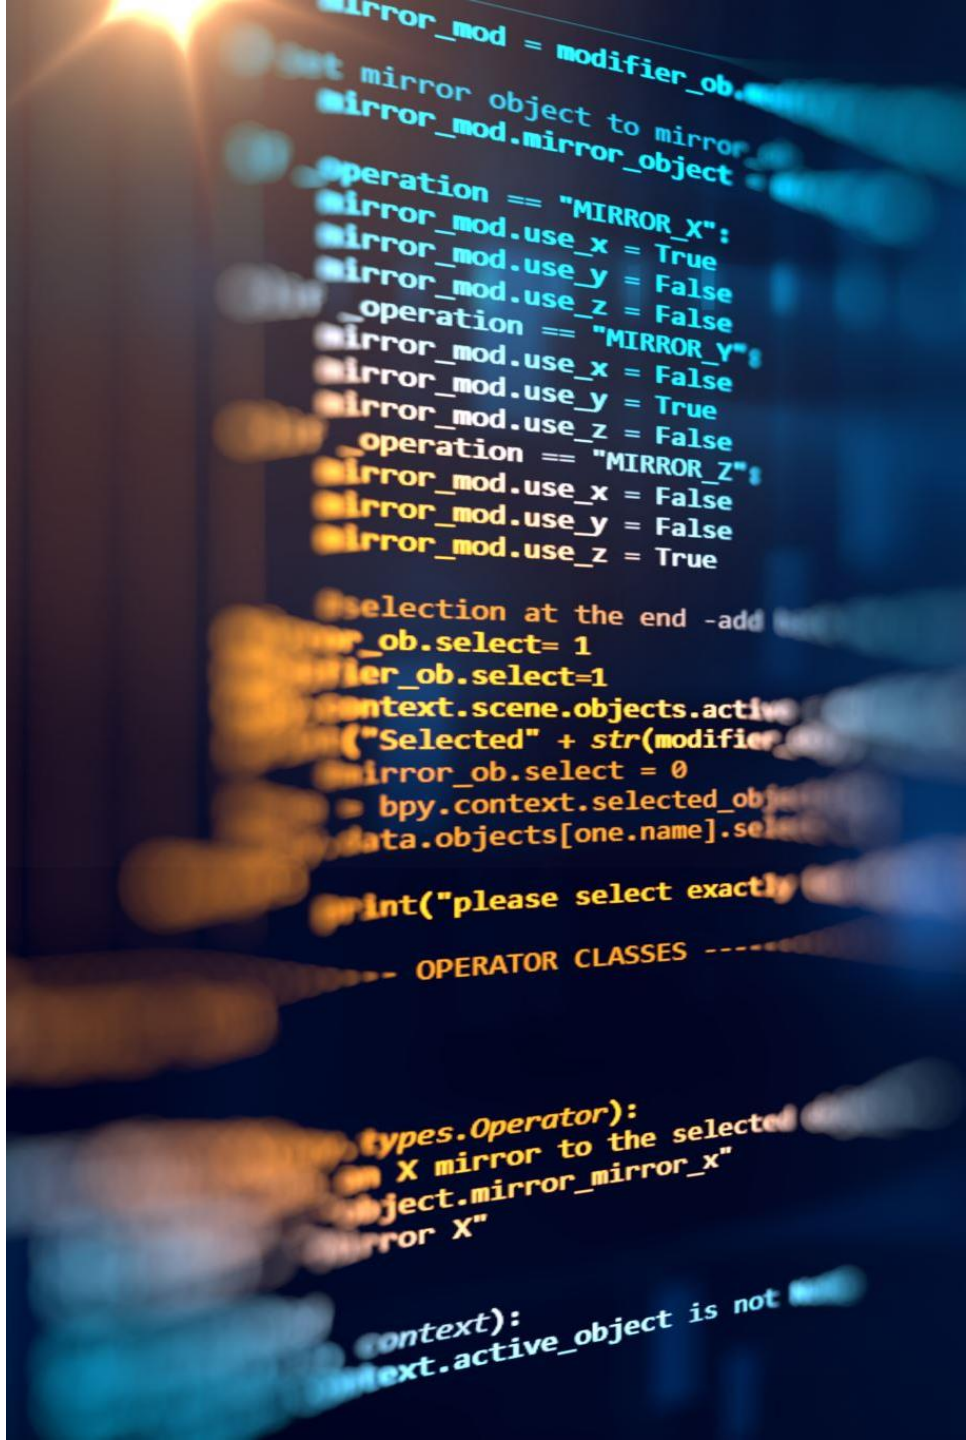

Azure SQL Database - HA/DR Options

3. Backups, and Restore Options

https://docs.microsoft.com/en-us/azure/sql-database/sql-database-business-continuity#sql-database-features-that-you-can-use-to-provide-business-continuity

Automated Backup and PITR Overview

Automatic backup

Full backups weekly, differential backup daily, log backups every 5 minutes

Daily and weekly backups copied to Azure Storage (Storage can be LRS, ZRS, or GRS by default)

Self-service database restore

Point-in-time up to a second granularity

REST API, PowerShell, or Portal

Creates a new database in the same logical server

Backup retention policy

Default = 7 days >> Maximum = 35 days
* DTU SKU's = No additional cost to retain backups
* vCore SKU's = Backup Costs Calculated Hourly

(optional) Long Term Retention (LTR) – 10 years max

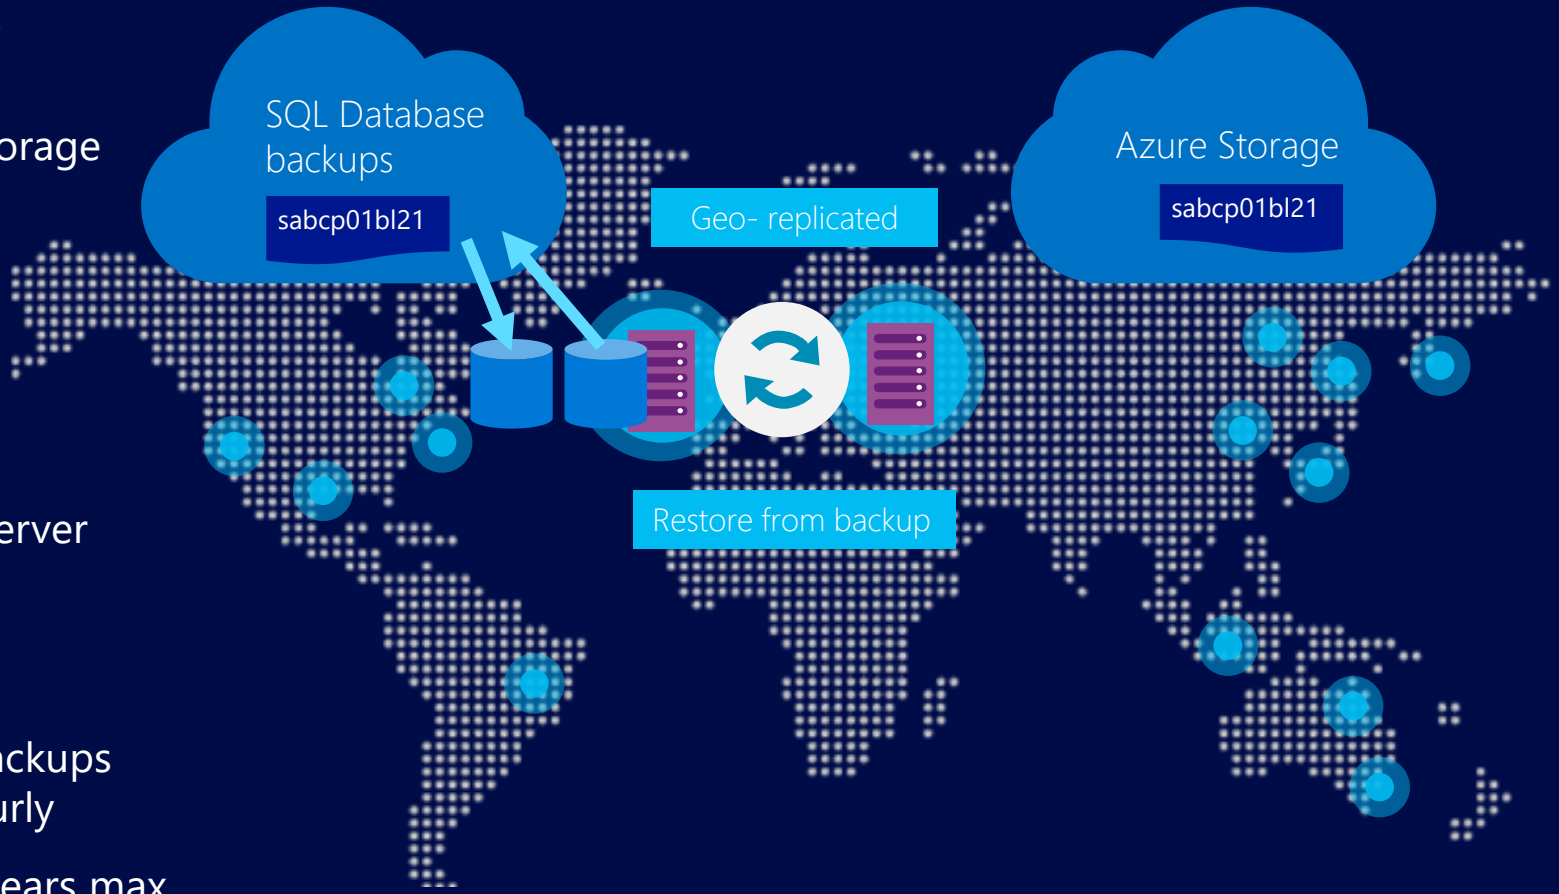

Automatic, geo-redundant backups - Azure SQL Database | Microsoft Learn

What SQL DB Recovery is Possible via Database Backups?

Options available for database recovery with database backups.

1. Restore to a **new database** on **same SQL Server** *or* **any SQL Server in <u>same region</u>**, recovered to specified point in time within the backup retention period.
2. Restore a **database** on **same SQL Server**, recovered to deletion time for a **deleted database**
3. Restore to a **new database** on **any SQL Server in <u>any region</u>**, recovered to point of most recent geo-replicated backups. *If geo-backup not enabled, this option is not available*
4. For Long-Term Retention (LTR) – can Restore to a **new database** from any LTR backup to any SQL Server in the same subscription (*even if the original server has been deleted*).

SQL DB Geo-Restore for Regional Disaster Protection

- Built on geo-redundant Azure Storage
- You MUST have selected GRS for your backups
- Restores last replicated backup to any Azure region as a new database
- No extra cost -- but also no capacity guarantee
- RTO≥12h, RPO=1h (up to… *worst case*)
- Self-service restore process

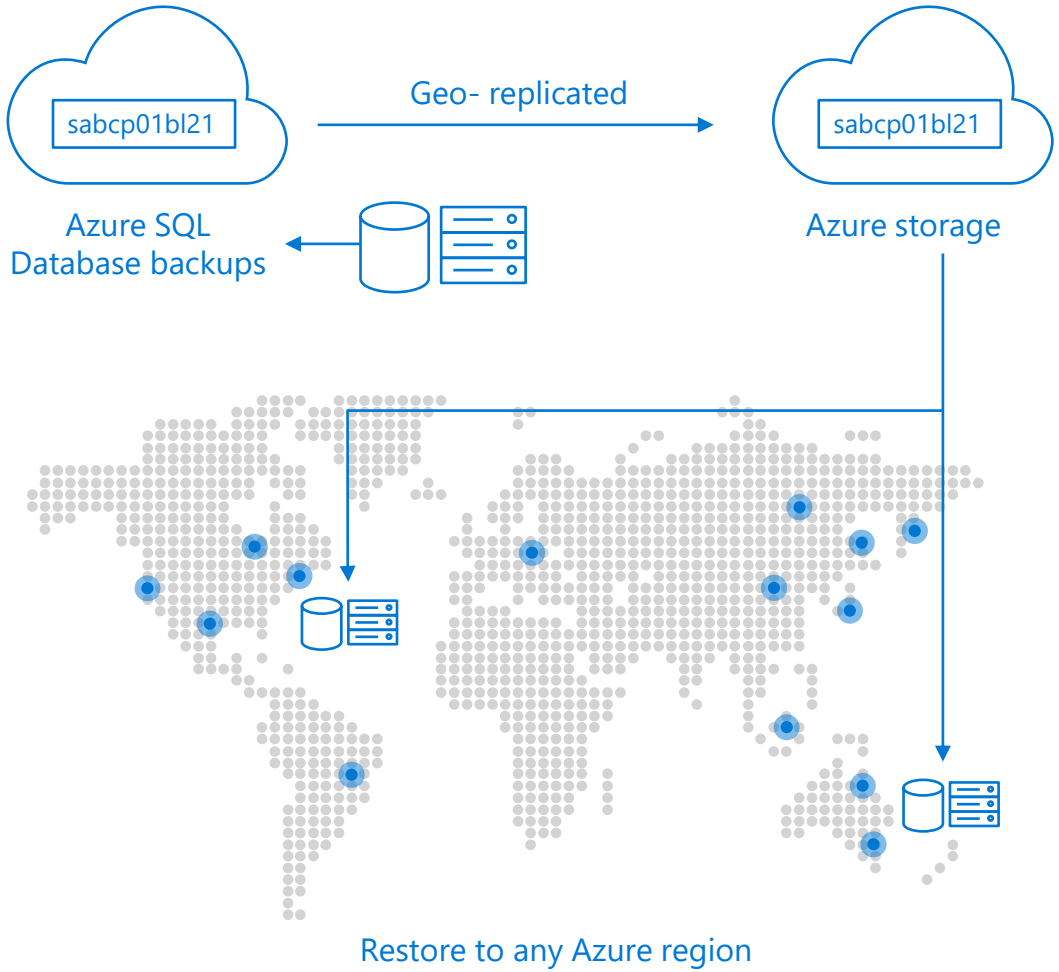

[Automatic, geo-redundant backups - Azure SQL Database | Microsoft Learn]

SQL DB Long Term Retention (LTR) - Extending Backups

Can configure single or pooled database with LTR Backup policy

Will automatically retain backups in separate Azure Blob Storage Containers

Retention up to max 10 years

Can recover any database using LTR backups via Azure portal or PowerShell

Can delete the original source server for the LTR backup, and the backup itself is not deleted (ie retention will continue to the LTR policy)

However - deleting an actual LTR backup itself is non-reversible.

Long-term backup retention - Azure SQL Database & Azure SQL Managed Instance | Microsoft Learn

SQL DB – How Does LTR Work?

Long-term backup retention leverages the automatic SQL Database backups created to enable point-time restore (PITR).

Specify for each SQL database how frequently you need to copy the backups to the long-term storage.

- Weekly backup retention (W)
- Monthly backup retention (M)
- Yearly backup retention (Y)
- Week of year (WeekOfYear)

W=0, M=0, Y=5, WeekOfYear=3

The 3rd full backup of each year will be kept for 5 years.

W=0, M=3, Y=0

The first full backup of each month will be kept for 3 months.

W=12, M=0, Y=0

Each weekly full backup will be kept for 12 weeks.

SQL DB - Restoring Deleted Databases

Restores the database to any point in time within the retention period

Creates a new database on the same server as was used by the original database

You can choose to failover to the restored database
or use scripts to recover data

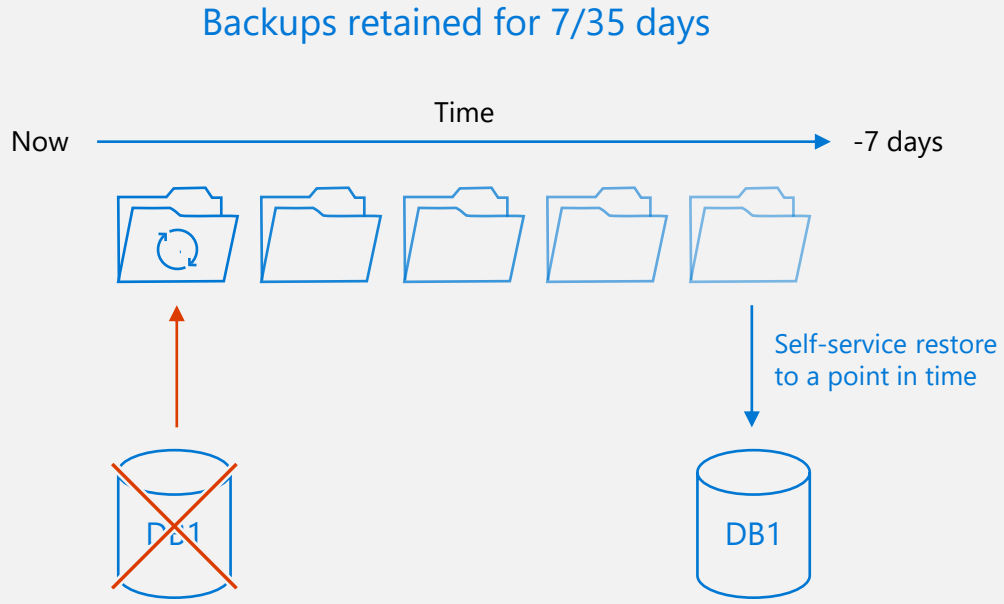

What Factors Affect SQL Database Recovery Time?

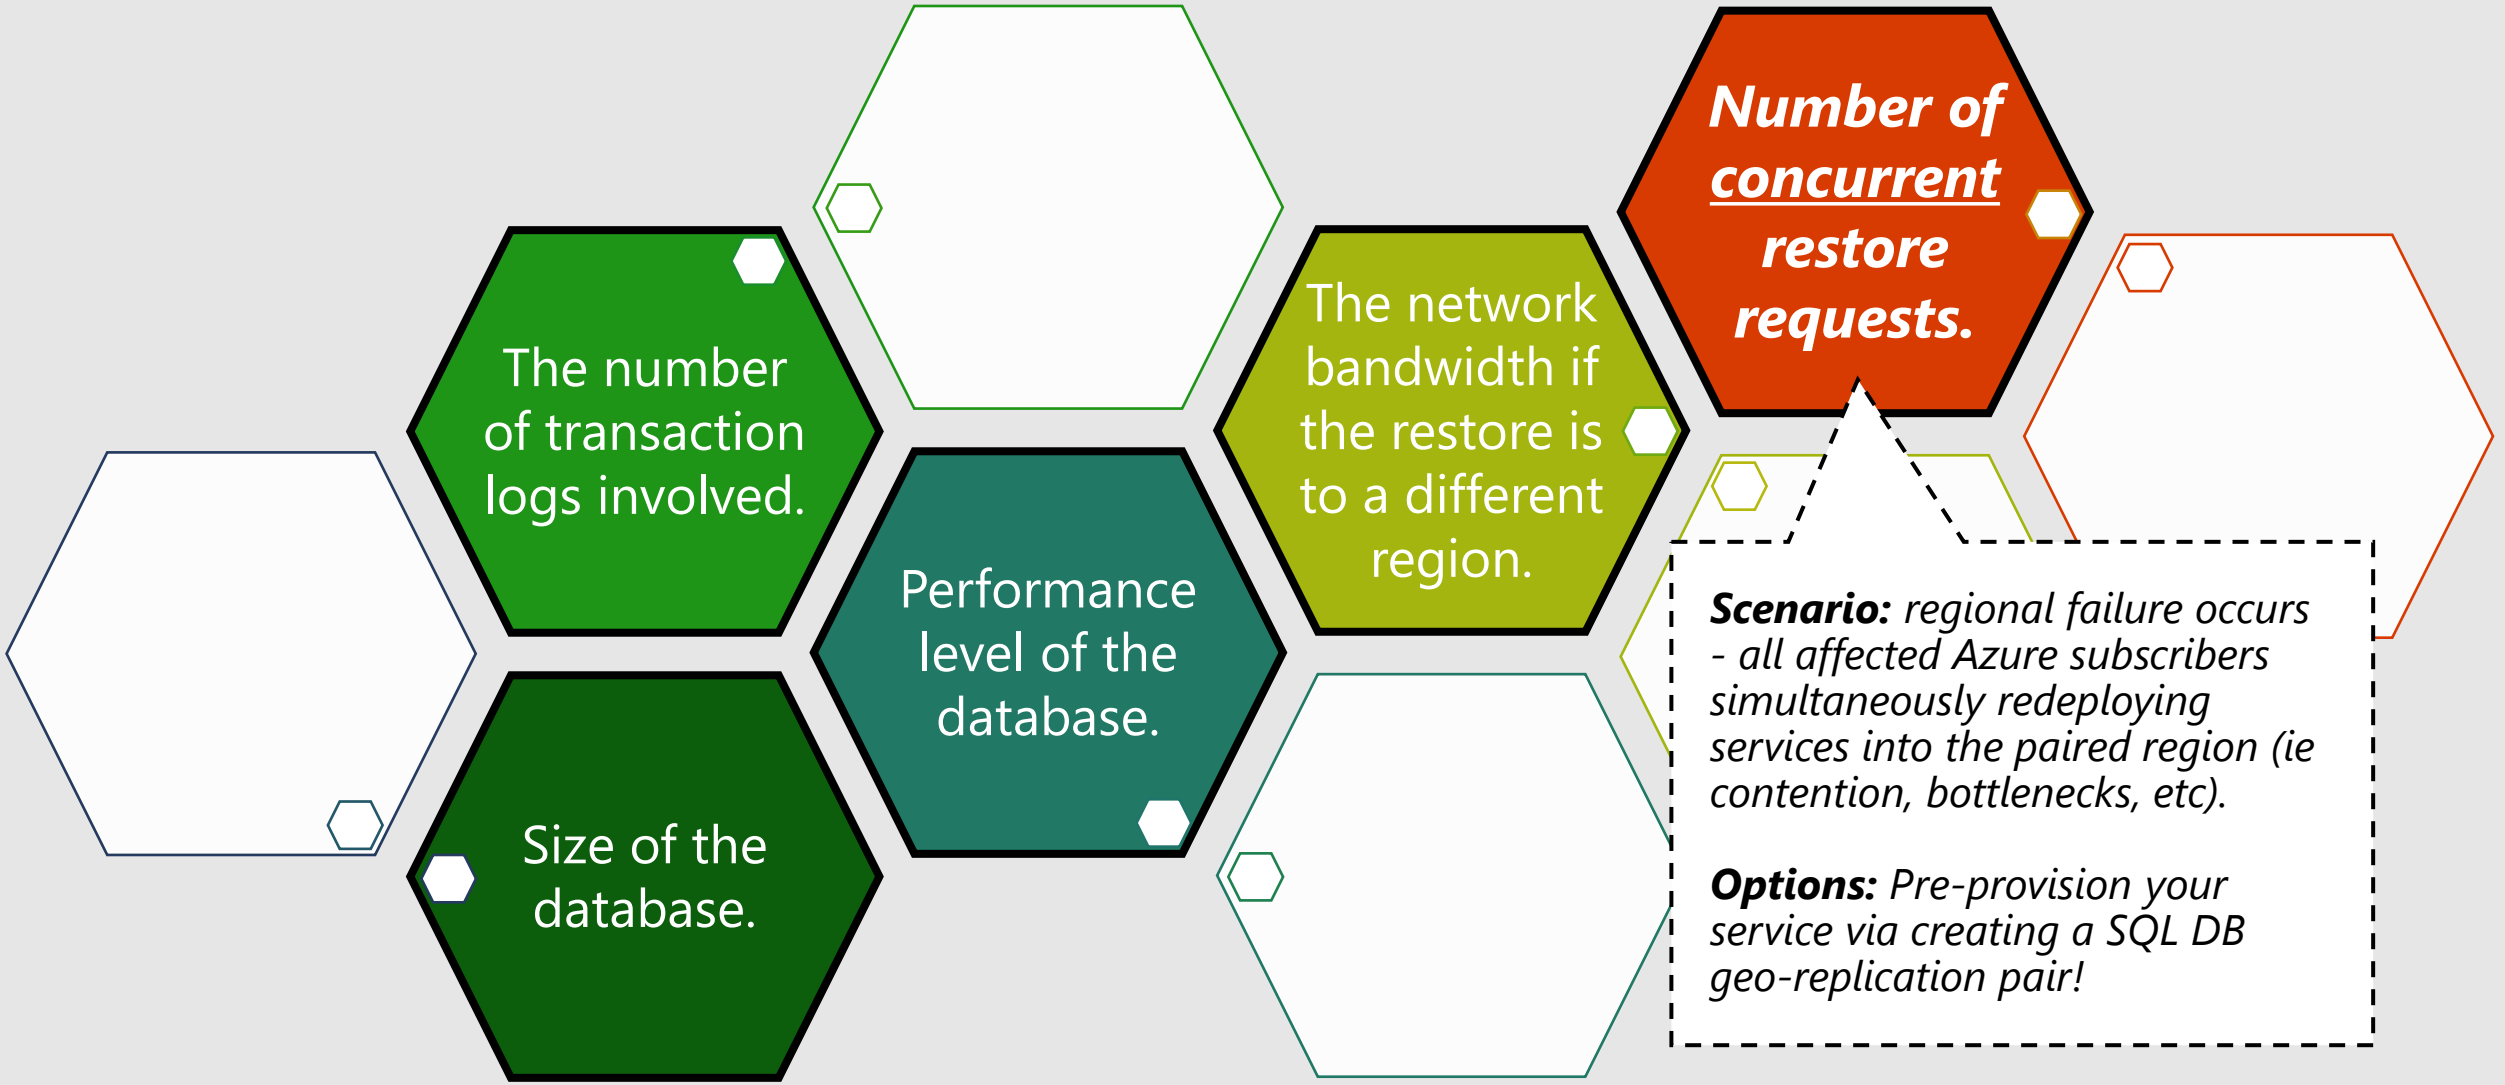

Azure SQL Database - Copy & Export Options

Copy

- Copies an Azure SQL Database to another DB on same or different Server.
- Essentially Creates a snapshot.

Export

- Creates a BACPAC file.
- Use Portal, SSMS, PowerShell or SQLPackage Utility.

NOTE: Neither of these options are designed for ongoing backup operations.

Azure SQL Database - HA/DR Options

4. Geo-replication – with secondary

https://docs.microsoft.com/en-us/azure/sql-database/sql-database-business-continuity#sql-database-features-that-you-can-use-to-provide-business-continuity

SQL DB Active Geo-Replication

Supported Service Levels	Available on All SQL DB SKU's
Readable Secondaries	Up to 4 replicas Can "*chain*" replicas (not via Portal)
Replication	**Asynchronous Only** Can "force" replication sync using *sp_wait_for_database_copy_sync*
Replication Lag	Query SQL DMV on Primary SQL DB *sys.dm_geo_replication_link_status*
Data Seeding on Create	Initial seeding @ ~500GB/Hour
Manageability tools	REST API, PowerShell, T-SQL, Portal, etc
Recovery Time Objective (RTO)	< 30 sec
Recovery Point Objective (RPO)	< 5 sec
Failover	Manual On Demand Only

Up to 4 secondaries

[Active geo-replication - Azure SQL Database | Microsoft Learn]

[Azure Region to Region Latency - Azure Speed Test]

Demo #3

- Check geo-replication lag
- [sys.dm_geo_replication_link_status - Azure SQL Database and Azure SQL Managed Instance | Microsoft Learn]

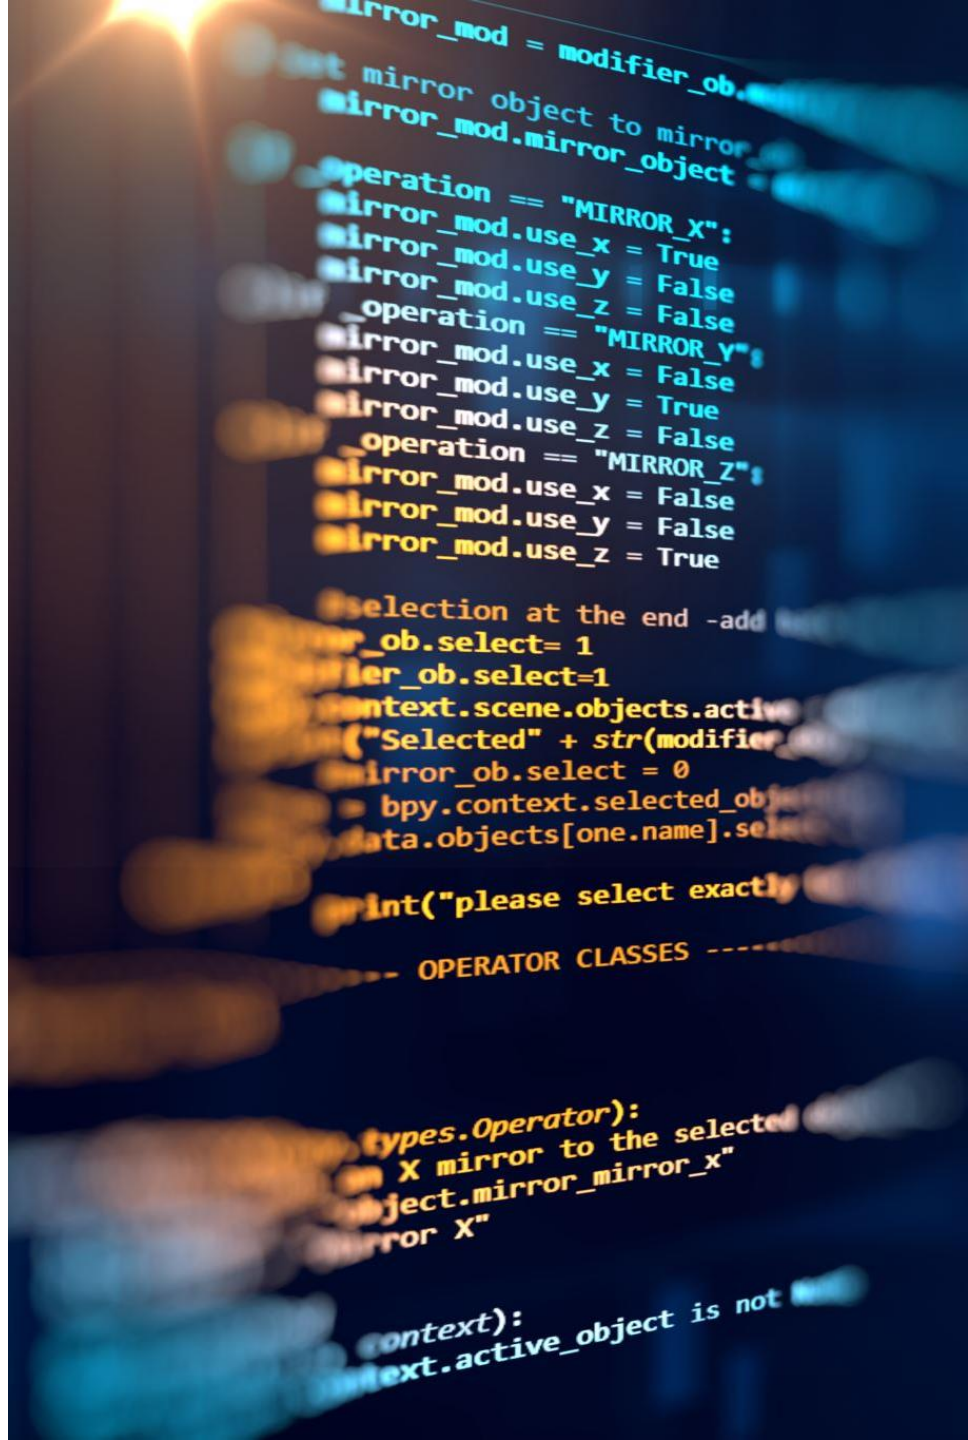

Azure SQL Database - HA/DR Options

5. Geo-replication – with failover groups

https://docs.microsoft.com/en-us/azure/sql-database/sql-database-business-continuity#sql-database-features-that-you-can-use-to-provide-business-continuity

SQL DB Failover Groups

"Failover Groups" Extends "Active Geo-Replication" Feature…

1. Enable geo-replication on **group of databases** (1+ DB's)
2. Provides a **single alias connection point** for RW + RO workloads
3. **Automatic DNS swap** during failover (ie *same connection strings*)
4. **Manual** or **Auto-Failover** options possible
5. **Manual-Failover** can be "*Forced*" (ie old primary offline, can result in data-loss) – or "*Friendly*" (ie old primary online, can get t-logs)
6. **Auto-Failover** has a "*grace period*" between 1-24 hours (*after which a "Forced" failover will occur, can result in data-loss*)
7. If **RO database offline**, all RO traffic can be directed to RW endpoint using "*AllowReadOnlyFailoverToPrimary*" property

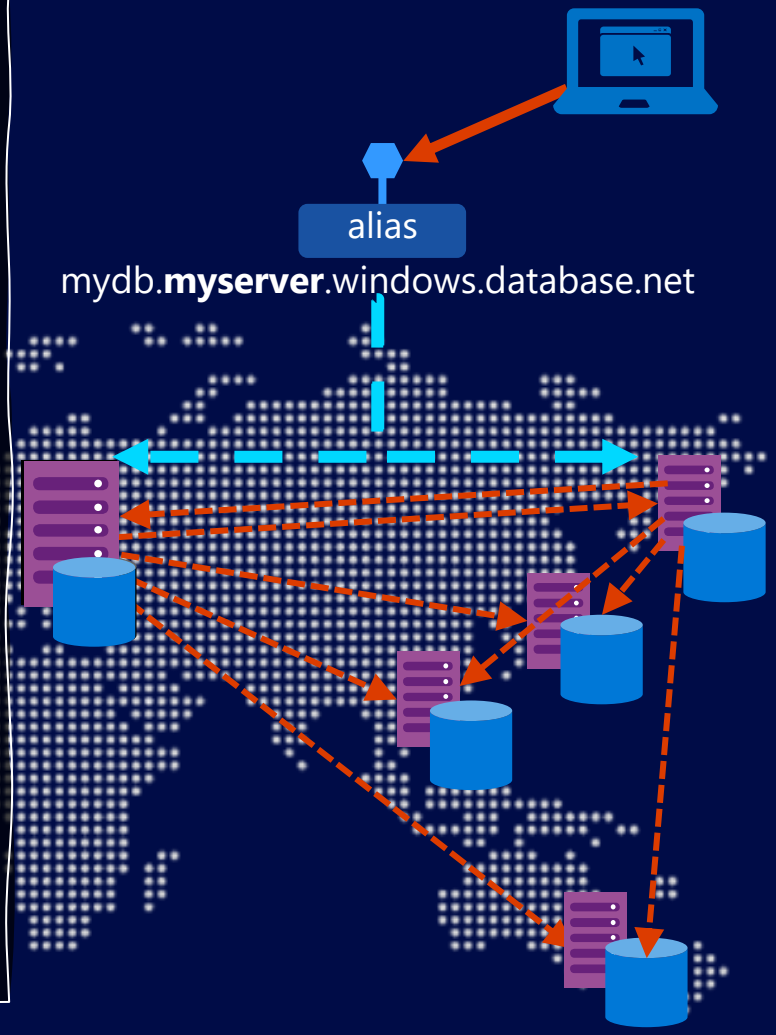

How to create and manage your cross subscription Auto-Failover groups (FoG)

Cloud business continuity - database recovery - Azure SQL Database & SQL Managed Instance | Microsoft Learn

Demo #4

- Access SQL DB Failover Groups with RW/RO endpoints
- [Auto-failover groups overview & best practices - Azure SQL Database | Microsoft Learn]

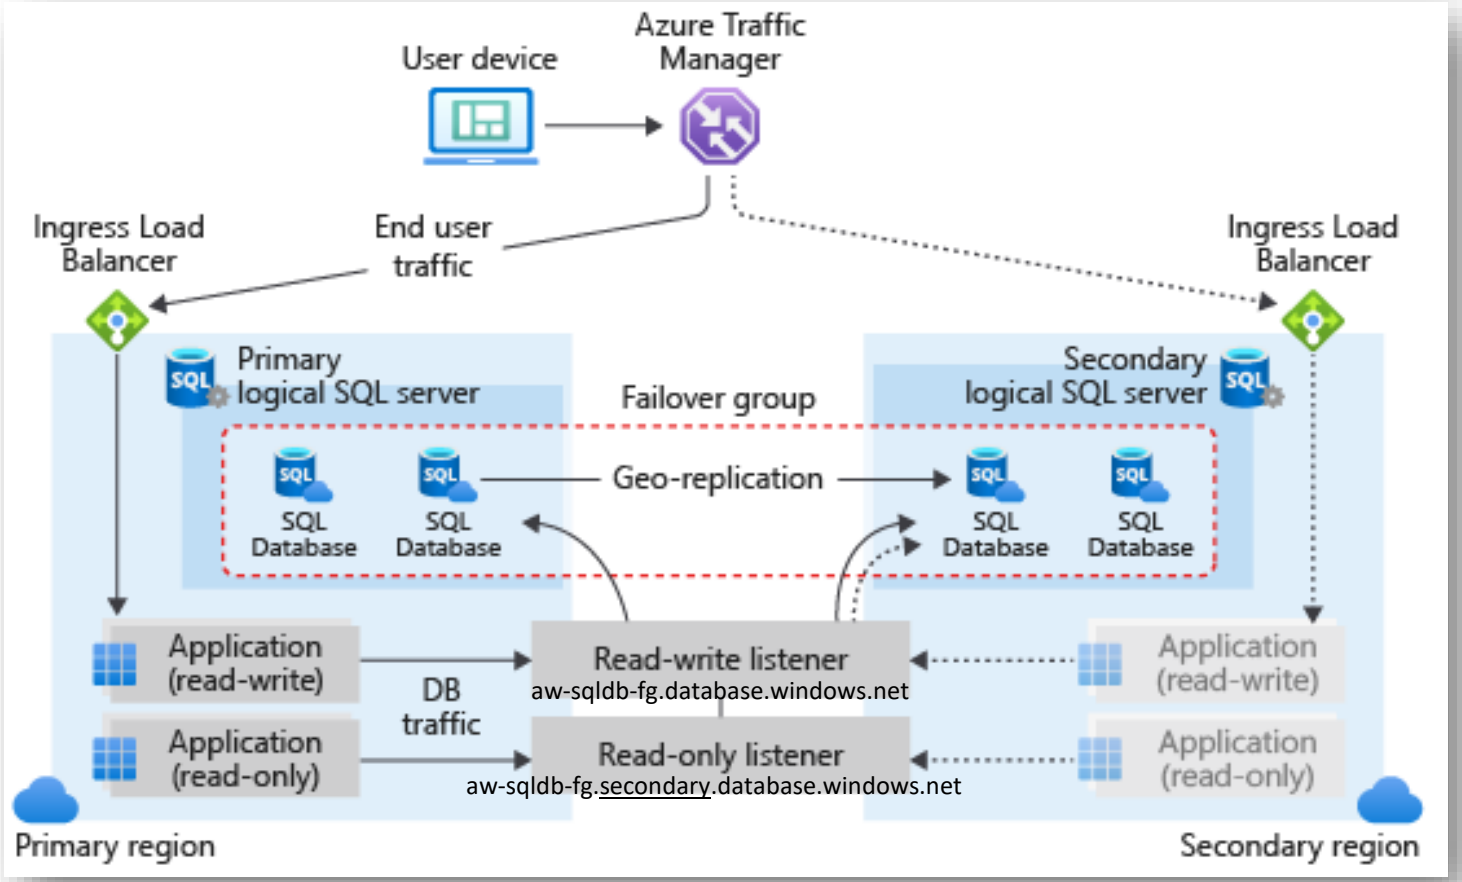

40. Activating HA & DR

DR Roles / Responsibilities

Microsoft

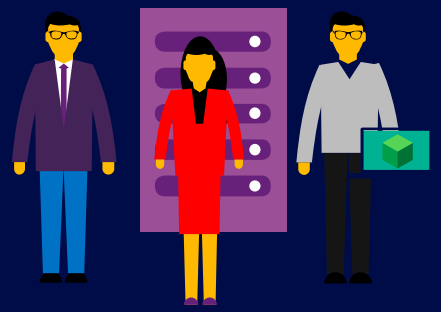

You

Azure – SQL Database Service

- Geo-distribution of SQL DB service
- Customer metadata protection and recovery
- Transparent high availability and protection from local failures
- Automatic geo-distributed of backups
- Automatic data synchronization of geo-replicated databases
- Platform compliance testing and certification
- Alert of impacted customers about their servers' degradation during regional failures

Customer

- Planning, design, regional selection for your service
- Detecting and alerting/surfacing application errors
- *Initiating* SQL DB point-in-time restore
- *Initiating* SQL DB geo-restore to selected region
- *Initiating* manual failover of geo-replicated SQL DB
- **Define your DR plans (Application & Database)**
- **Validate the plan via undertaking DR drills**
- **Note - Azure cannot simulate HA/DR failures**

Disaster recovery - Azure SQL Database | Microsoft Learn

DR for Azure Data Platform - Overview - Azure Architecture Center | Microsoft Learn

50. Wrap up + Summary

Summary: *The Key Takeaway Points*

1. SQL DB provides up to 99.995% (SLA)
2. SQL DB includes built-in availability
3. SQL DB offers several DR Options/Add-Ons
4. Some DR options allow use of RO replicas
5. Don’t forget you still need a DR plan

…and you need to test it!

Azure SQL Database Resiliency and Availability

EXPLAINED

[END OF SESSION]

Microsoft

Appendix and References

What Customers Say About SQL Versions

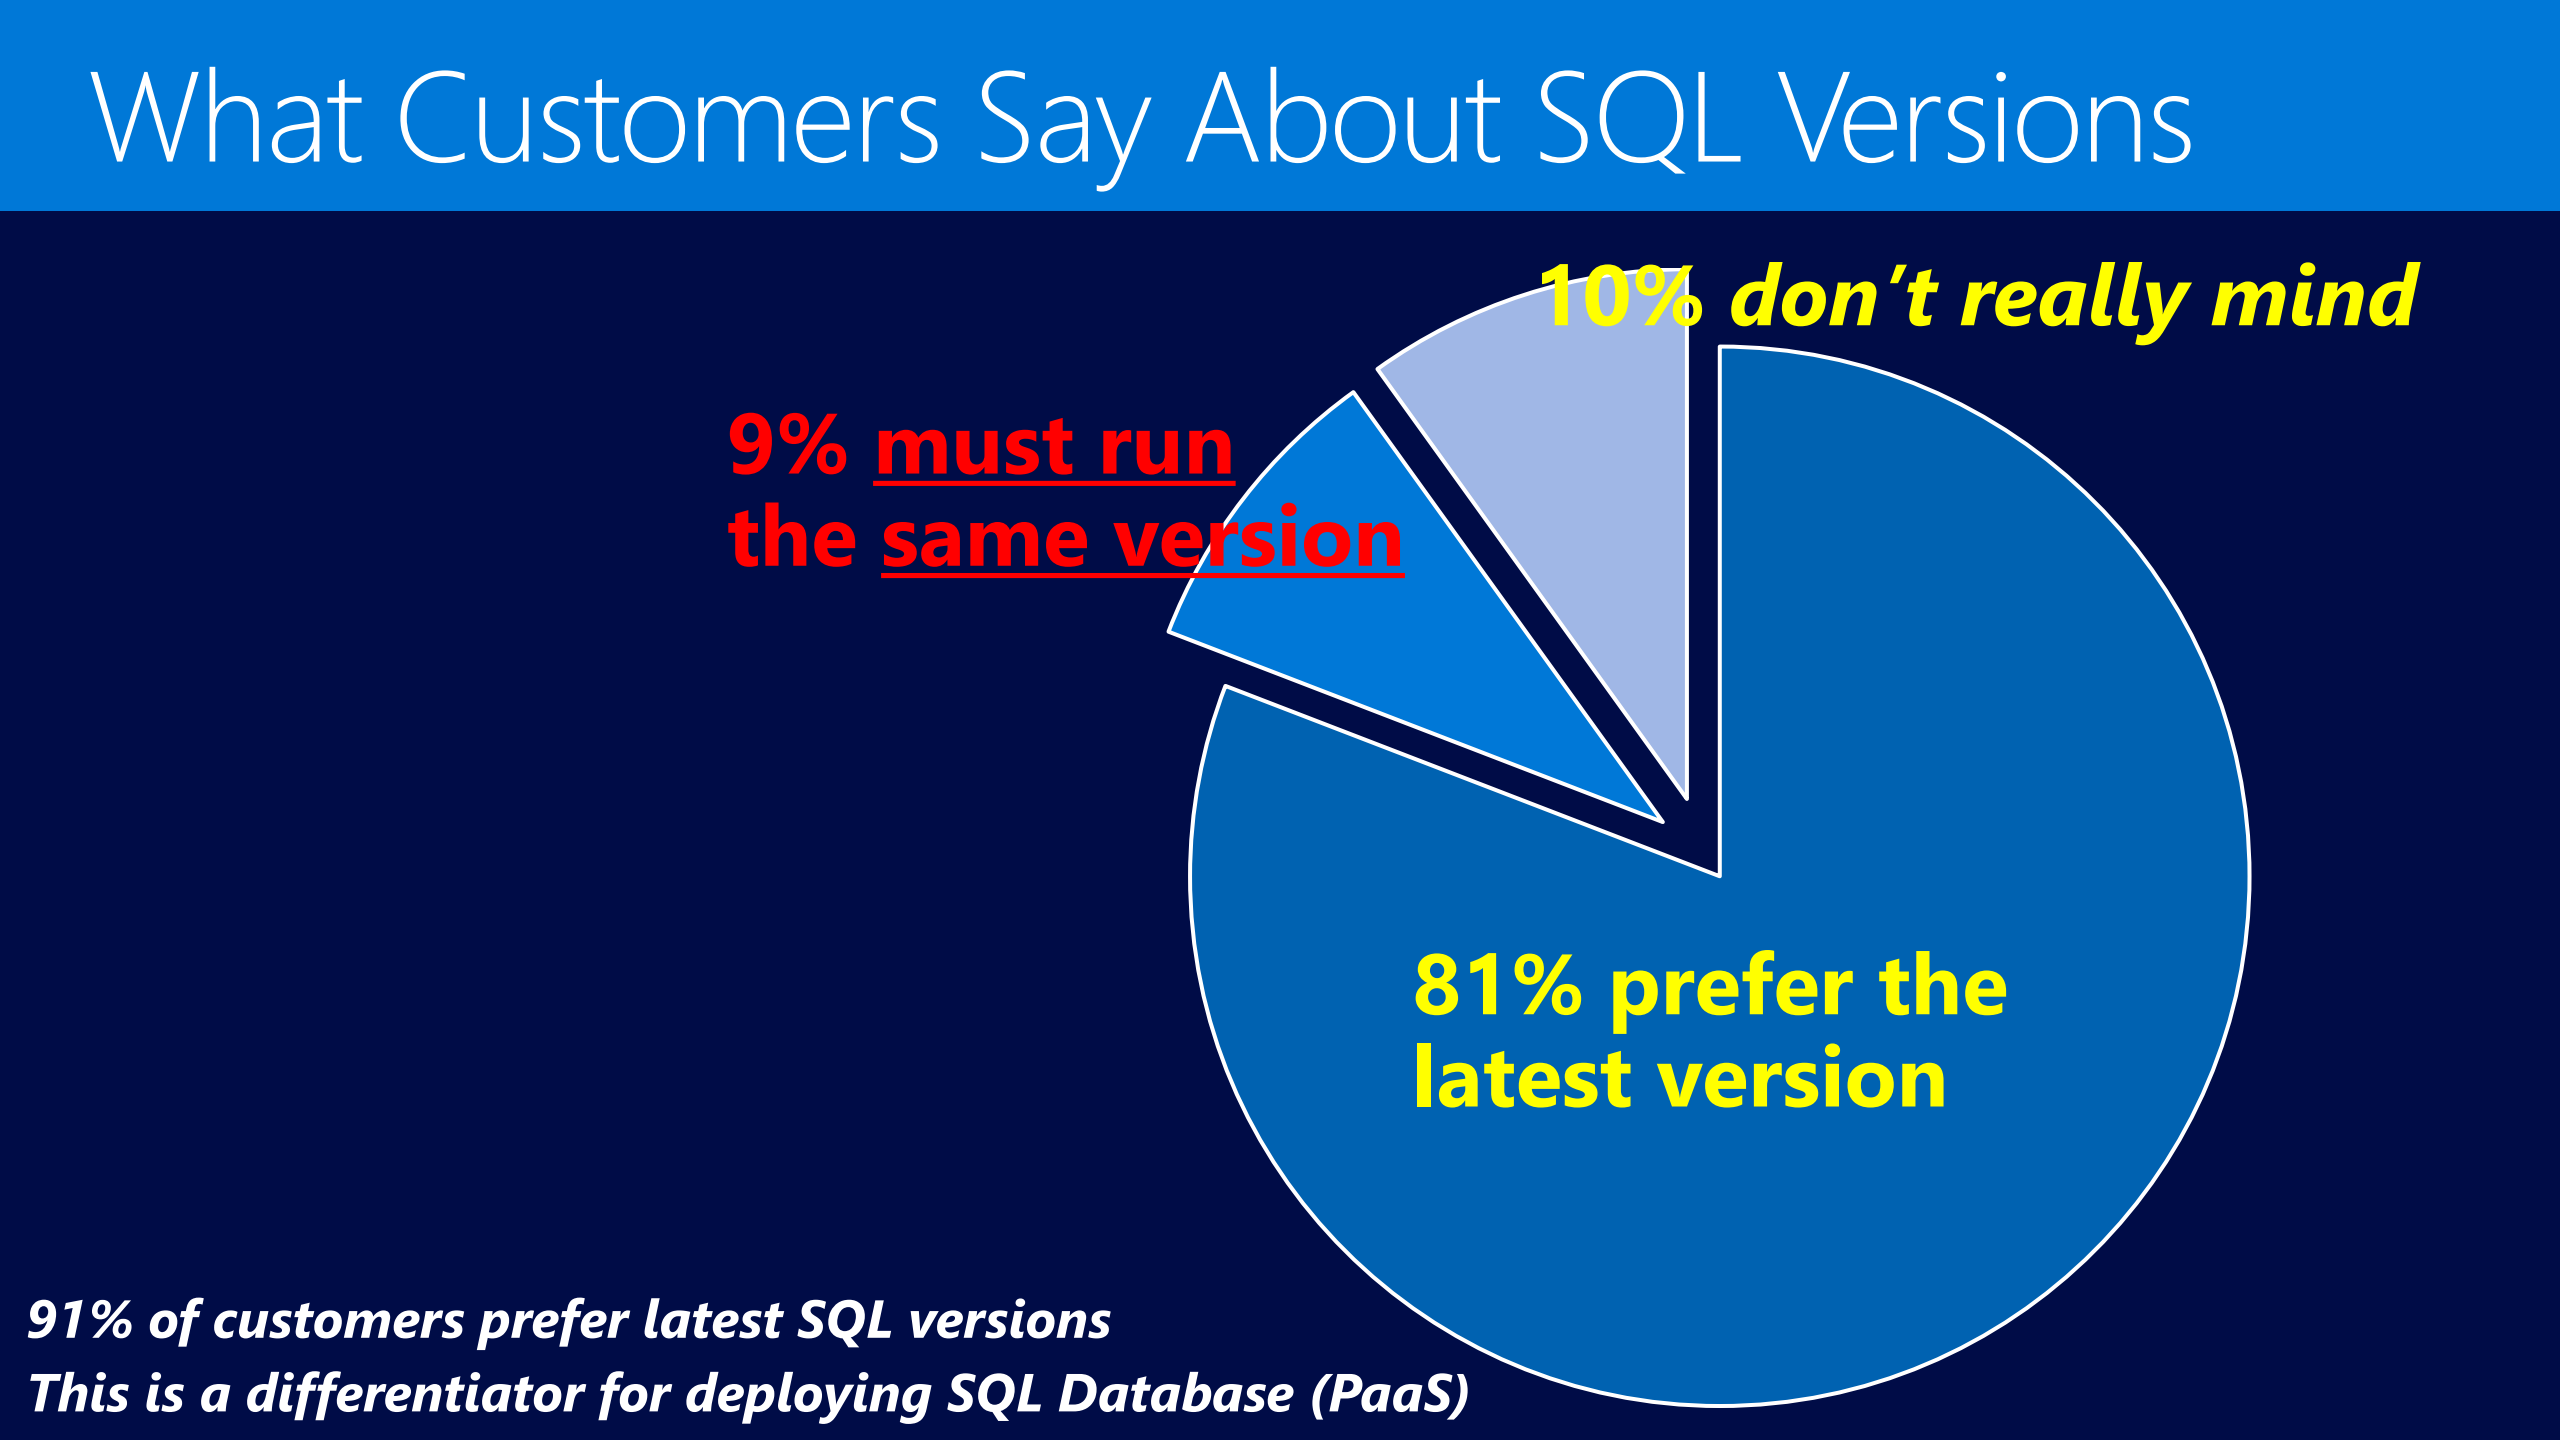

91% of customers prefer latest SQL versions
This is a differentiator for deploying SQL Database (PaaS)

Benefits of vCores vs DTU

- Choice between vCores and DTUs in Azure SQL Database as a unit of measure for available CPU
- Understand your compute requirements in the cloud vs. what you use on-premises today
- Easier to right-size the destination environment by removing the guesswork of DTUs
  - Gen 4 CPUs are based on Intel E5-2673 v3 (Haswell) 2.4 GHz processors. In Gen 4, 1 vCore = 1 physical CPU
  - 1-24 vCores
  - Gen 5 logical CPUs are based on Intel E5-2673 v4 (Broadwell)
    2.3 GHz processors. In Gen 5, 1 vCore = 1 hyper thread
  - 2-80 vCores

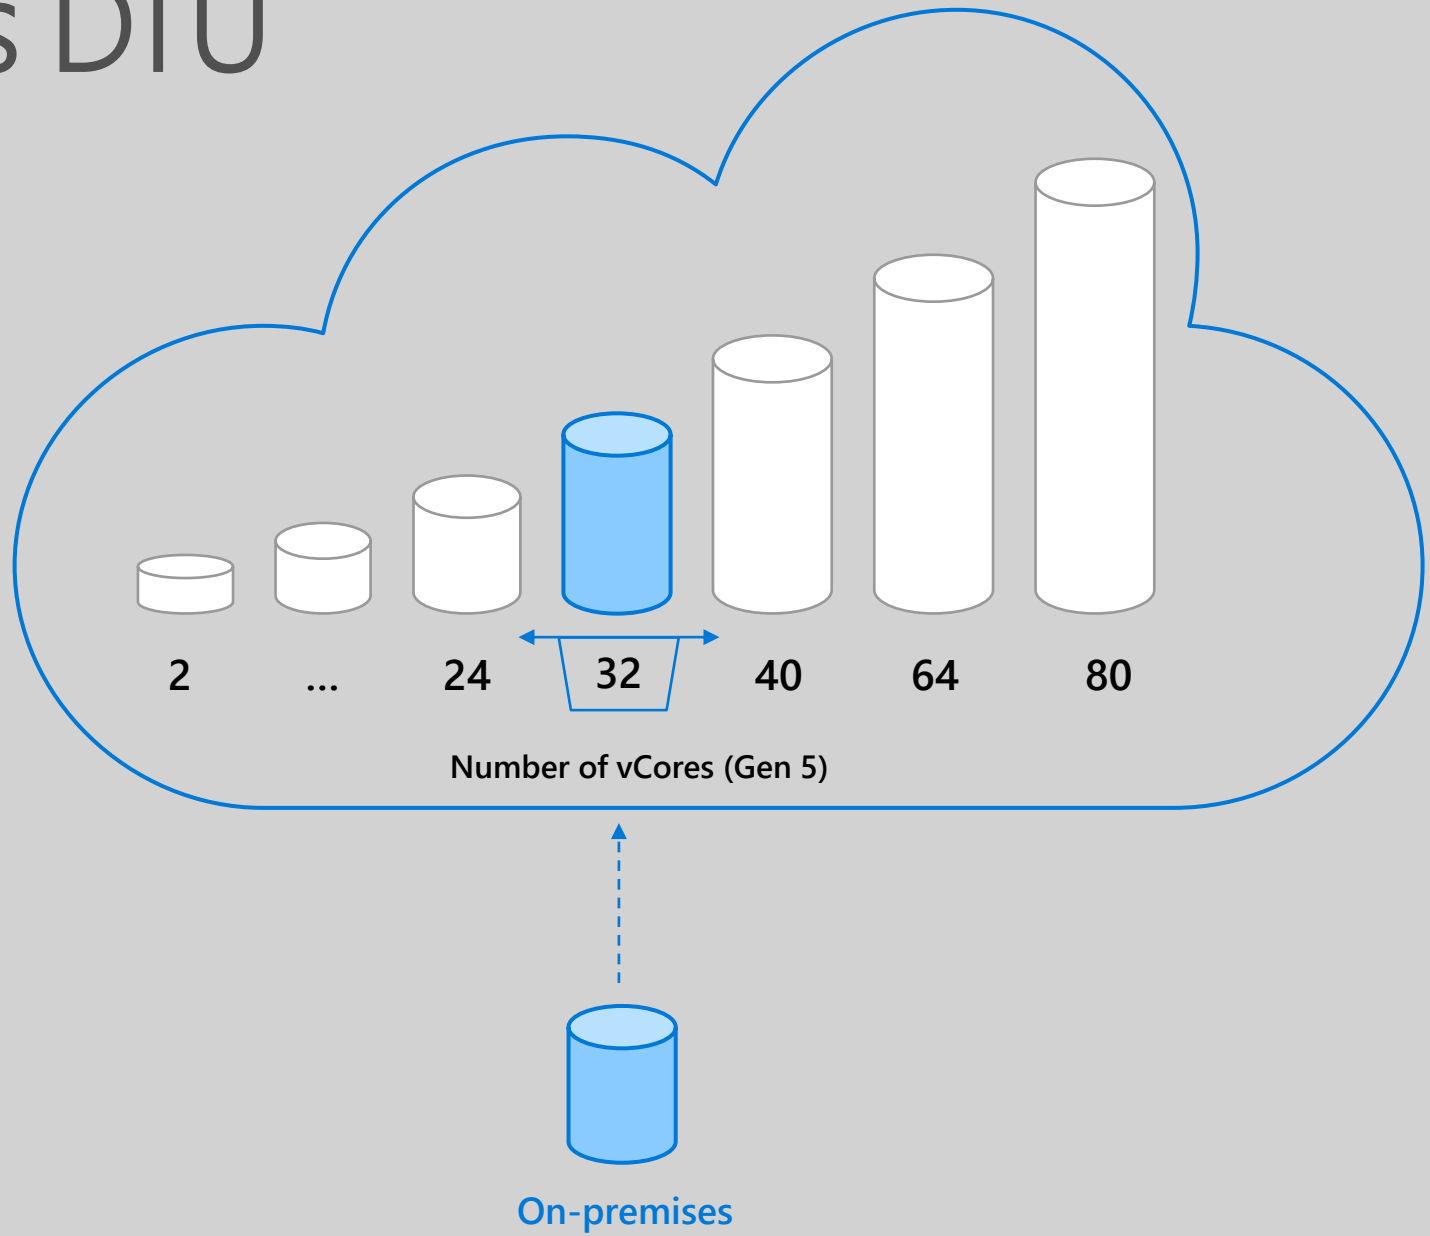

Azure covers 100 compliance offerings

Azure has the deepest and most comprehensive compliance coverage in the industry

Global

- CIS Benchmark
- CSA STAR attestation
- CSA STAR certification
- CSA STAR self-assessment
- ISO 20000-1
- ISO 22301
- ISO 27001
- ISO 27017
- ISO 27018
- ISO 27701
- ISO 9001
- SOC 1 Type 2
- SOC 2 Type 2
- SOC 3
- WCAG 2.0 (ISO 40500)

US Gov

- CJIS
- CNSSI 1253
- DFARS + CMMC
- DoD IL2
- DoD IL4
- DoD IL5
- DoD IL6
- DoE 10 CFR Part 810
- EAR
- FedRAMP
- FIPS 140-2
- IRS 1075
- ITAR
- NIST 800-171
- NIST 800-53
- NIST CSF
- Section 508 VPATs

Industry

Automotive
- TISAX (Germany)

Education
- FERPA (US)

Energy
- NERC (US)

Financial Services
- 23 NYCRR 500 (US)
- AFM + DNB (Netherlands)
- AMF + ACPR (France)
- APRA (Australia)
- CFTC 1.31 (US)
- EBA (EU)
- FCA + PRA (UK)
- FFIEC (US)
- FINMA (Switzerland)
- FINRA 4511 (US)
- FISC (Japan)
- FSA (Denmark)
- GLBA (US)
- KNF (Poland)
- MAS + ABS (Singapore)
- NBB + FSMA (Belgium)
- OSFI (Canada)

Financial Services (Cont.)
- OSPAR (Singapore)
- PCI DSS Level 1
- PCI 3DS
- RBI + IRDAI (India)
- SEC 17a-4 (US)
- SEC Regulation SCI (US)
- Shared assessments (US)
- SOX (US)
- TruSight

Healthcare and Life Sciences
- ASIP HDS (France)
- GxP (FDA 21 CFR Part 11)
- HIPAA (US)
- HITRUST
- MARS-E (US)
- NEN 7510 (Netherlands)

Media and Entertainment
- CDSA
- DPP (UK)
- FACT (UK)
- MPA

Telecommunications
- GSMA

Regional

Americas
- Argentina PDPA
- Canada Privacy Laws
- US CCPA

Asia Pacific
- Australia IRAP Unclassified
- China GB 18030:2005
- China DJCP (MLPS)
- China TRUCS/CCCPPF
- China TCS
- India MeitY
- Japan CS mark gold
- Japan My Number act
- Korea K-ISMS
- New Zealand ISPS
- Singapore MTCS Level 3

Europe and Middle East
- EU EN 301 549
- EU ENISA IAF
- EU GDPR
- EU model clauses
- Germany C5
- Germany IT—Grundschutz workbook
- Netherlands BIR 2012
- Russia personal data law
- Spain ENS High
- Spain LOPD
- UAE DESC
- UK cyber essentials plus
- UK G-Cloud
- UK PASF

https://aka.ms/AzureCompliance