Microsoft

Introduction to Business Continuity in Azure SQL

John Deardurff
Senior Cloud Solution Architect

John Deardurff

Senior Cloud Solution Architect (Global Delivery)
Worldwide Product Lead – SQL Server Workshops

Email: John.Deardurff@Microsoft.com
Website: www.SQLMCT.com
GitHub: github.com\SQLMCT
LinkedIn: /in/johndeardurff/

What does this session cover?

What is Business Continuity?

Service Tiers

High Availability

Zone Redundancy

Replicas and Failover Groups

Microsoft

What is Business Continuity in Azure SQL?

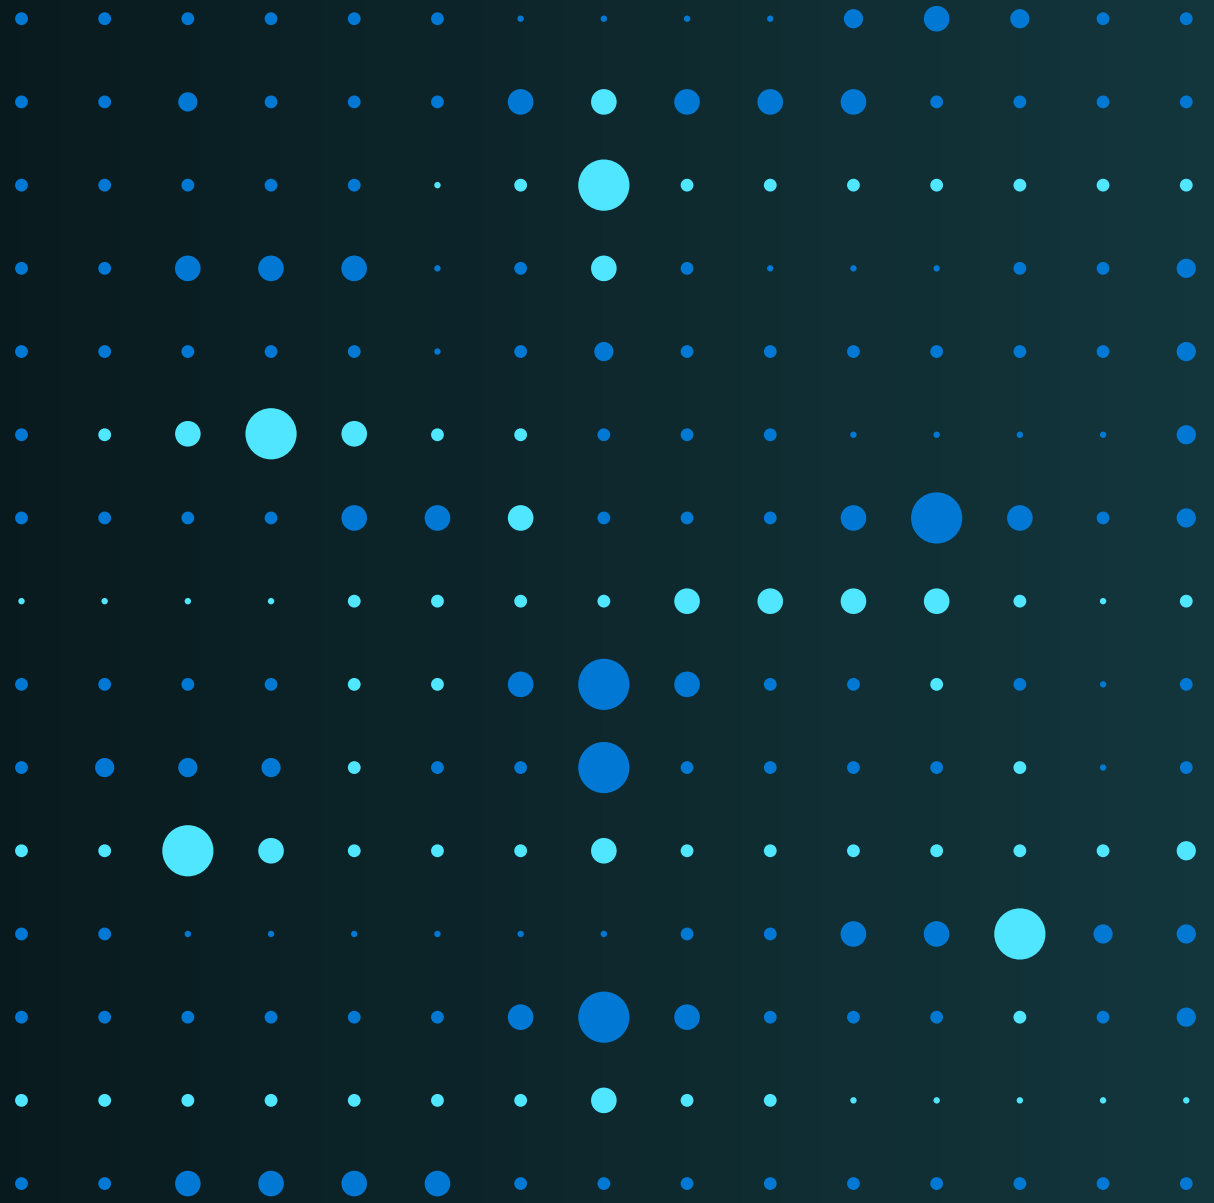

Business Continuity Problem

Mechanisms, Policies and Procedures that enable a business to continue operating in face of disruption.

Customer Error

Customer deleted their data by mistake.

Equipment Failures

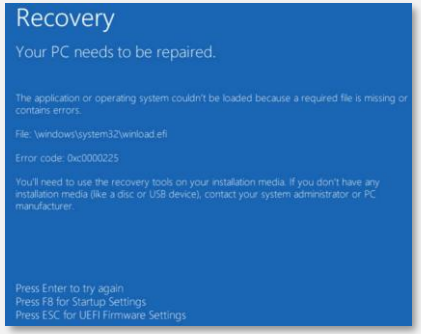

Hardware or other equipment failures

Natural Disasters

Flood, Hurricanes, Tornadoes, Earthquakes

Manmade Disasters

Wars or major ecological accidents

Availability vs High Availability vs Disaster Recovery

Availability

Every database comes with core **resiliency and availability built-in**, that protects it against software or hardware failures.

Industry leading financially backed SLA of 99.99% availability.

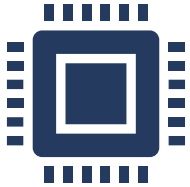

High Availability

Continuous availability of the database provided through **Zone Redundancy.**

Automatic in region recovery from zonal hardware and software failures that's transparent to applications.

Higher SLA of 99.995% availability.

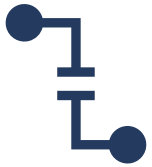

Disaster Recovery

Customer configured database replicas or failover groups

Ability to quickly recover the database from a catastrophic regional failure to provide business continuity.

Protect your business by maximizing application availability.

Microsoft

What are Service Tiers in Azure SQL?

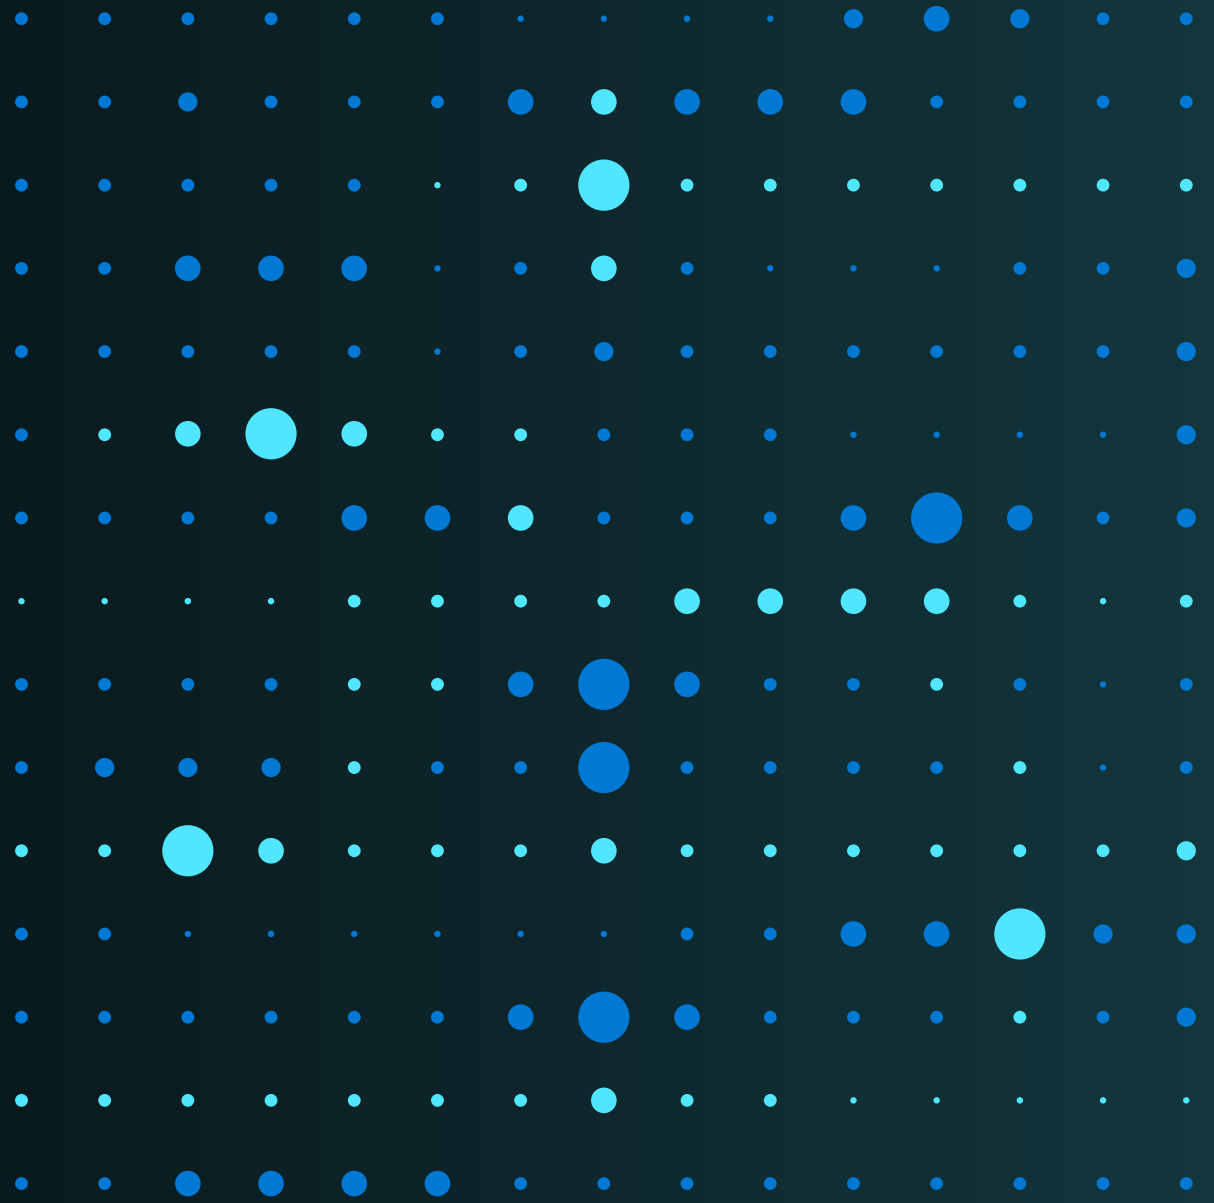

Overview of Service Tiers

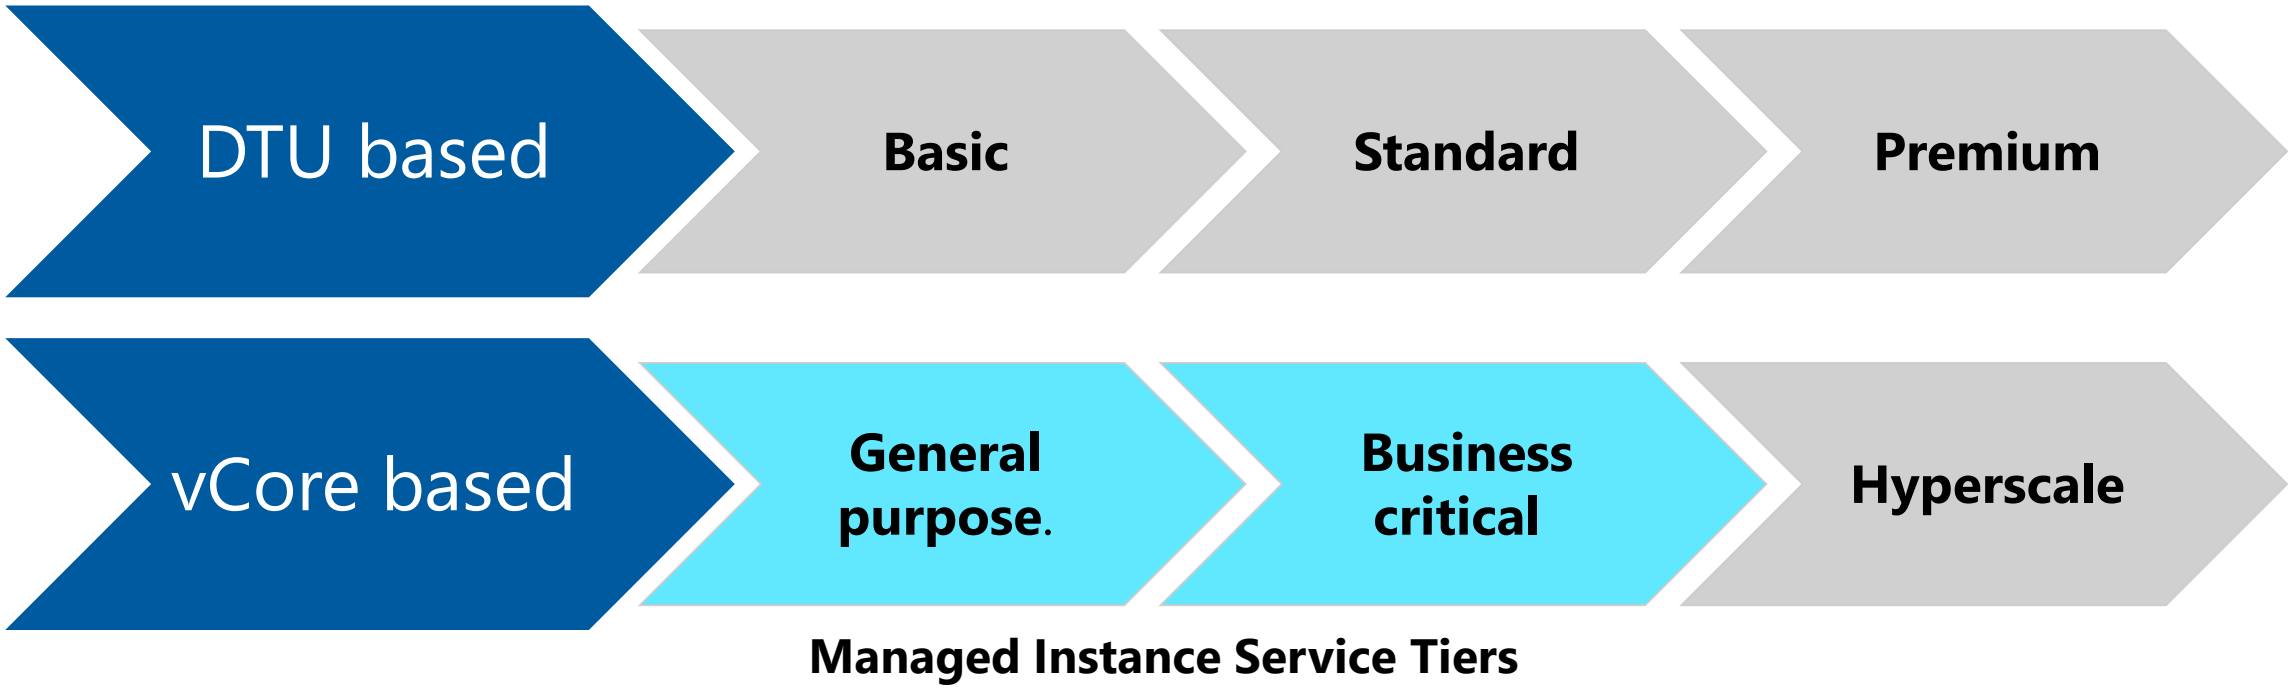

Changing Service Tiers

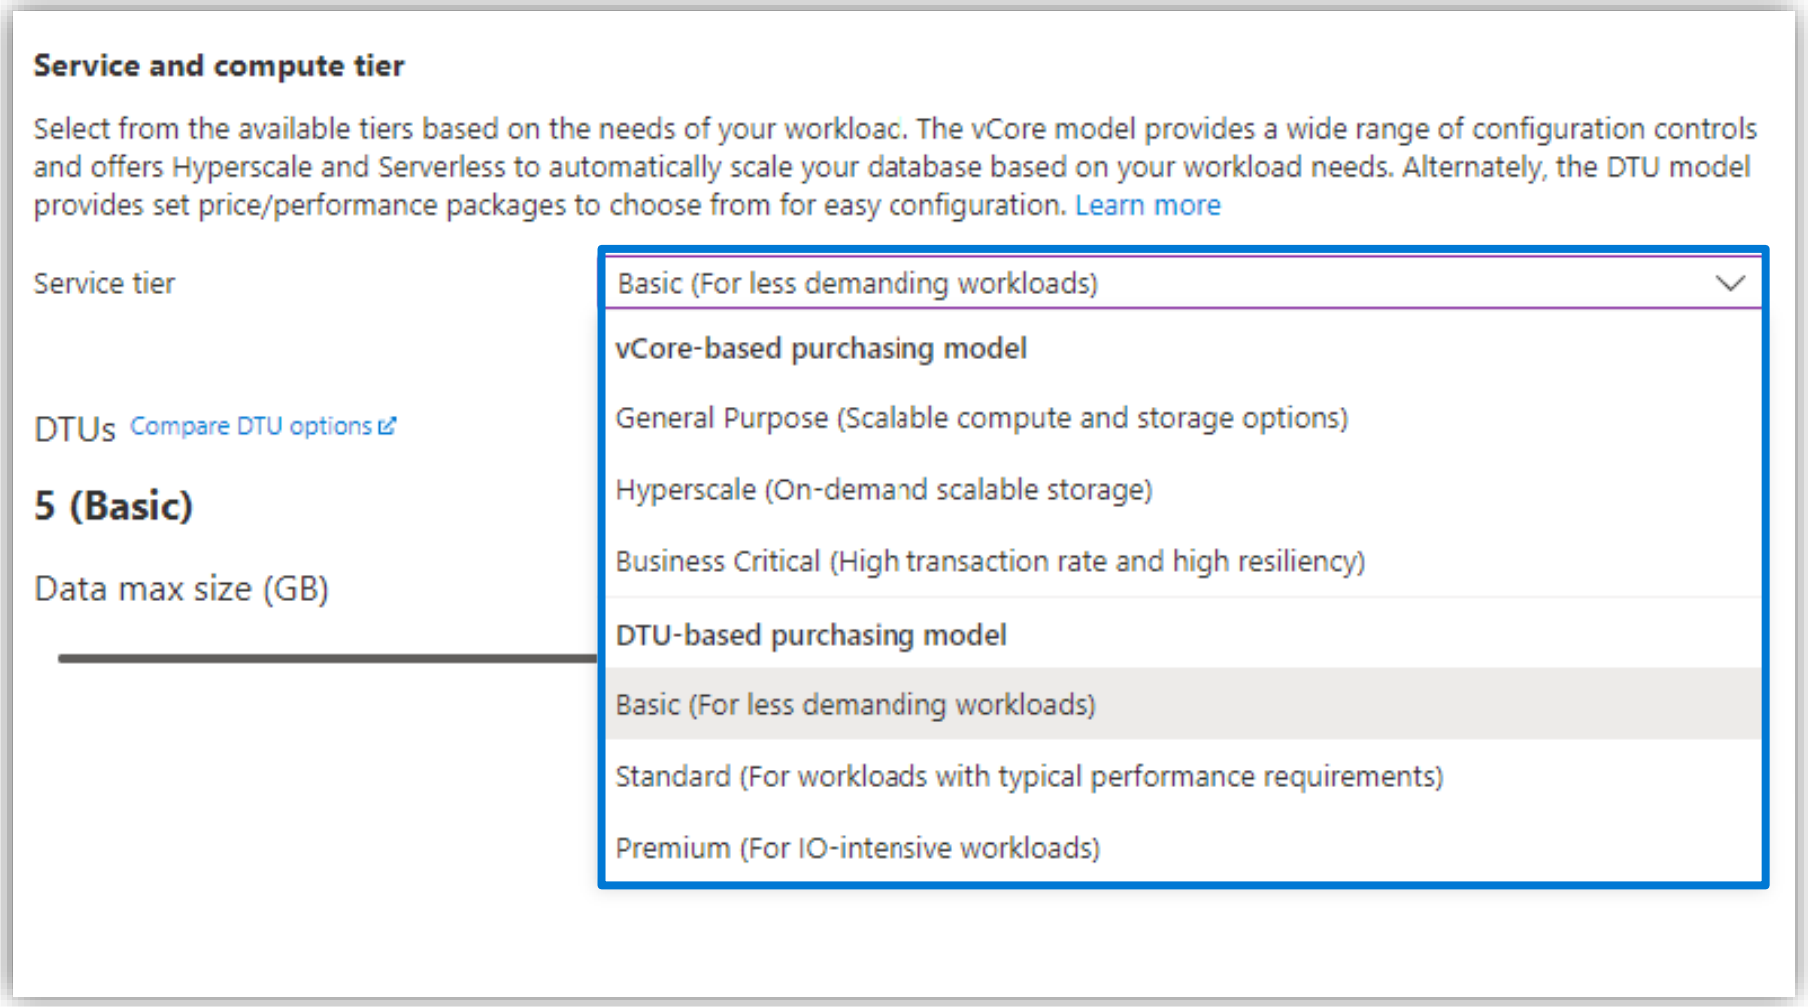

Microsoft

High Availability in Azure SQL

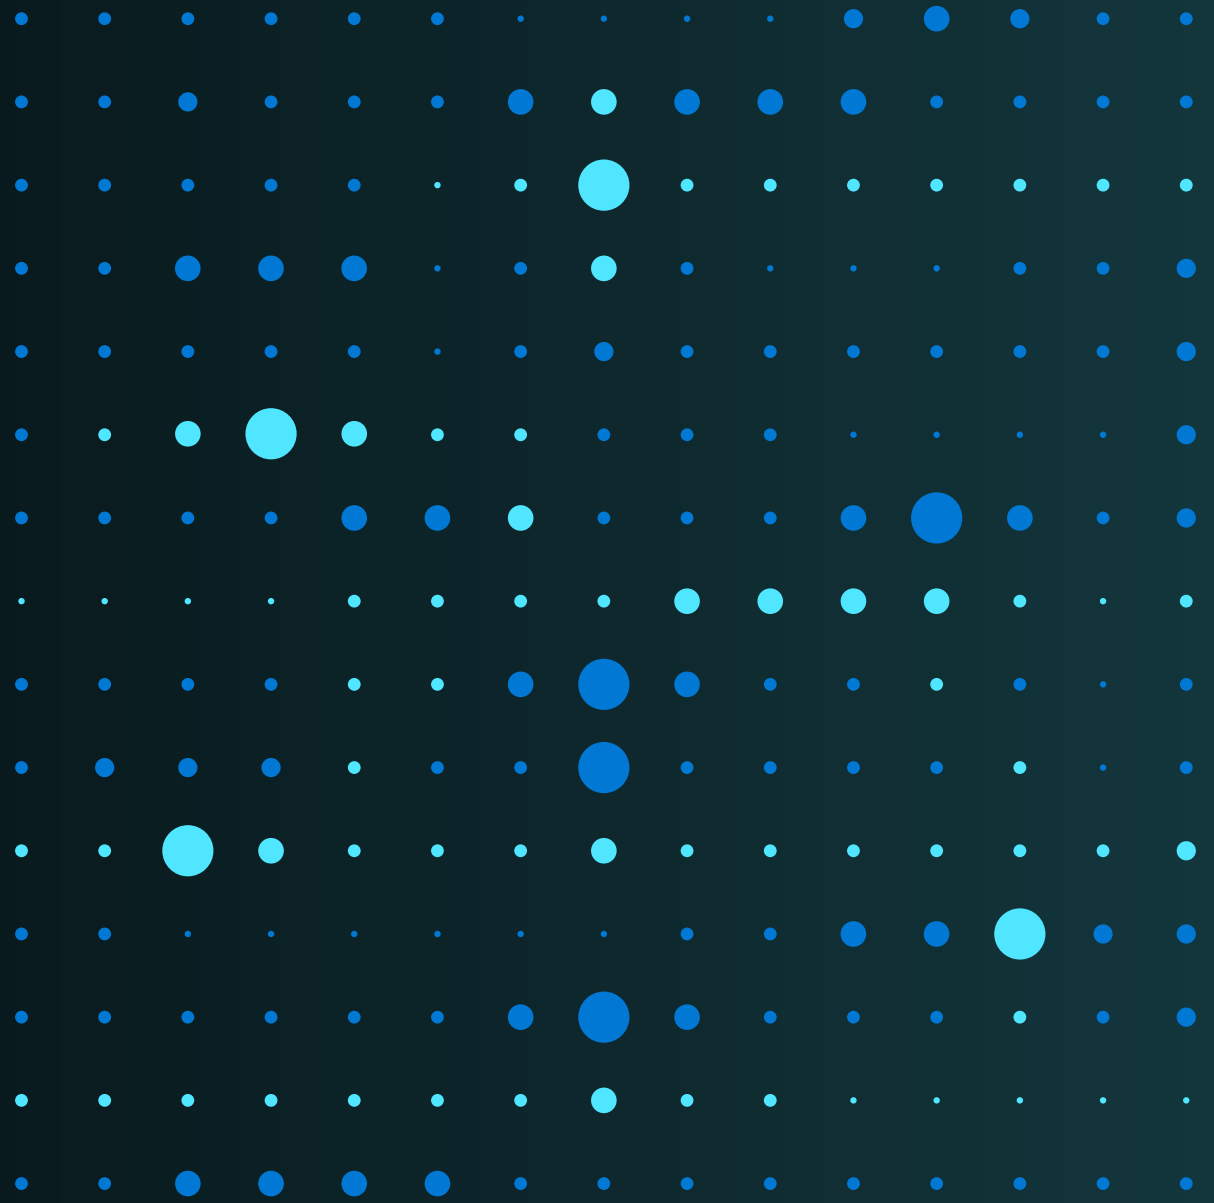

Availability Architecture: General Purpose vs Business Critical

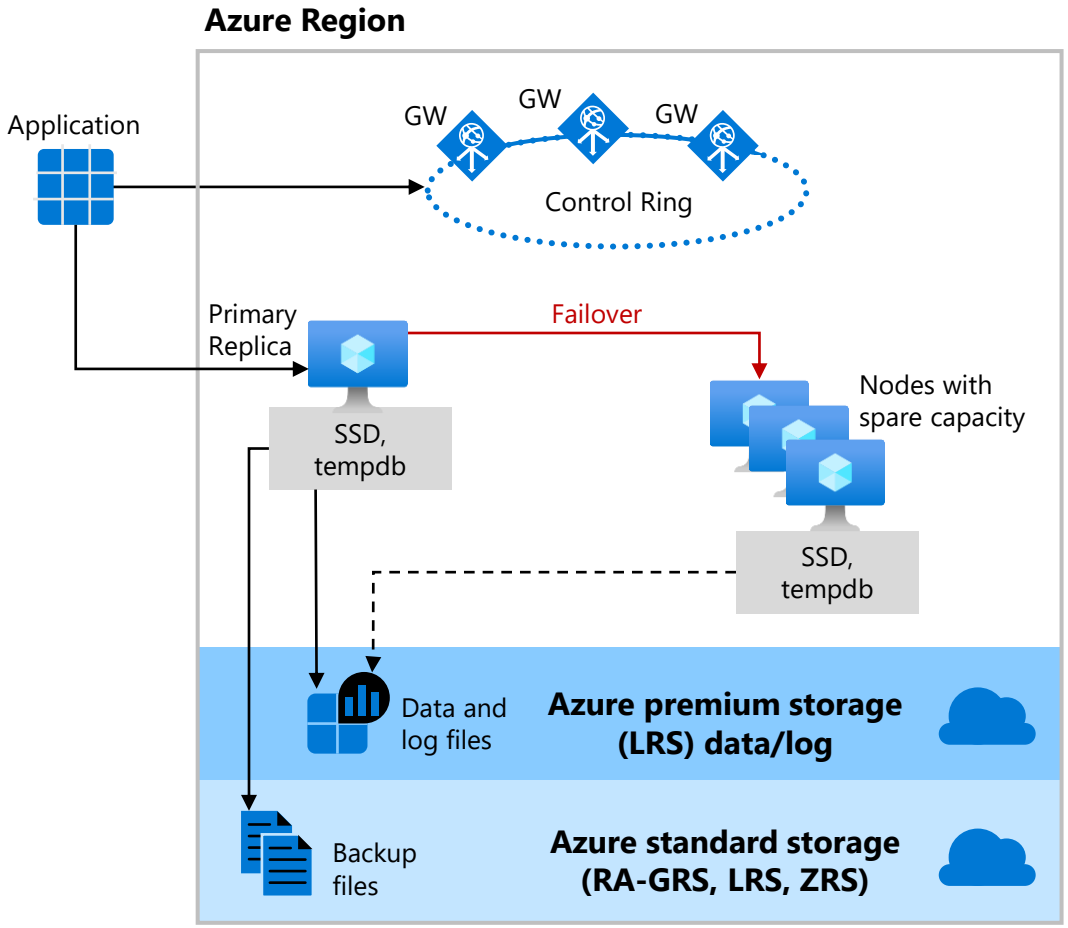

Basic (DTU), Standard (DTU), General Purpose (vCore)

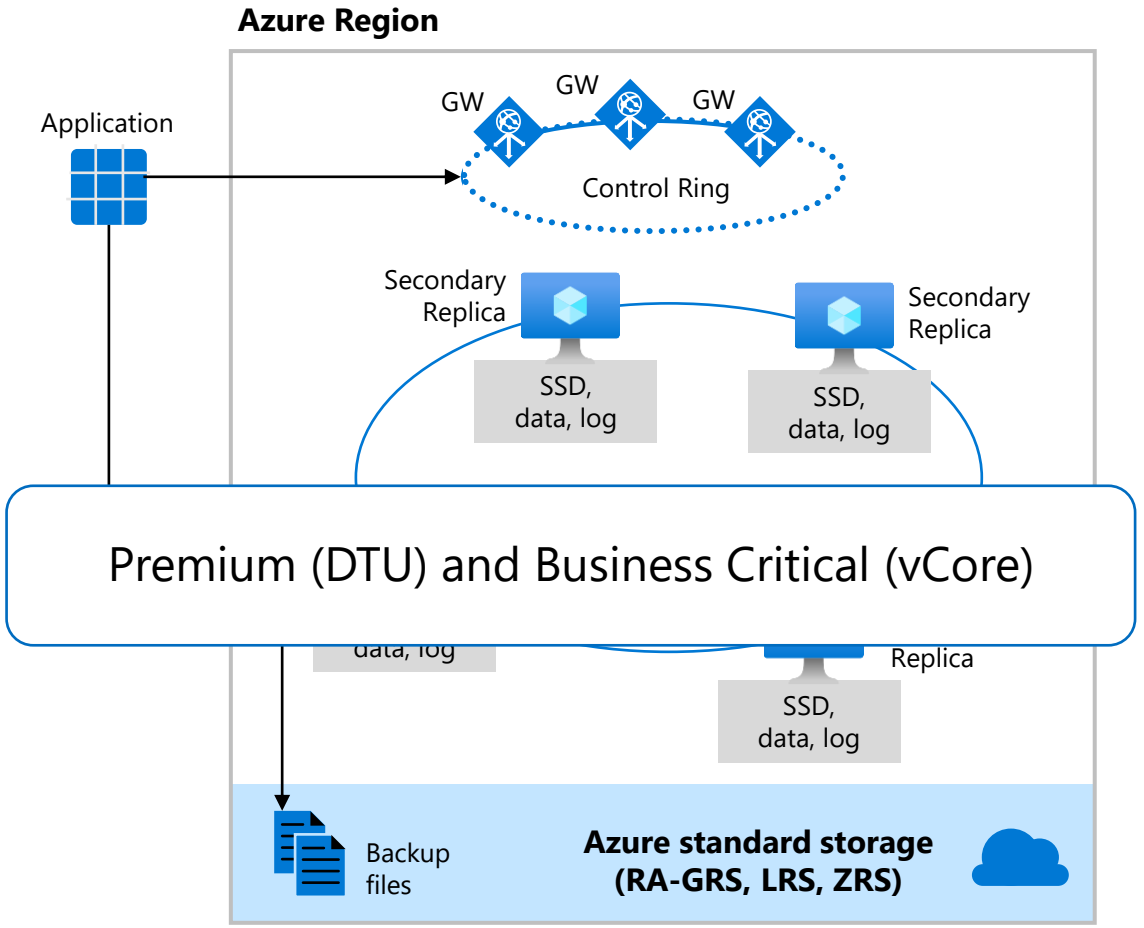

Business Critical (BC) service tier

High Availability - General Purpose

- Basic (DTU), Standard (DTU), General Purpose (vCore)
- Behaves like Failover Cluster Instance
- Remote storage provides data redundancy
- Backup files are in a different location with geo-redundancy
- Failover decisions based on SQL and Service Fabric
- Recovery time depends on spare capacity
- Connectivity redirection built-in

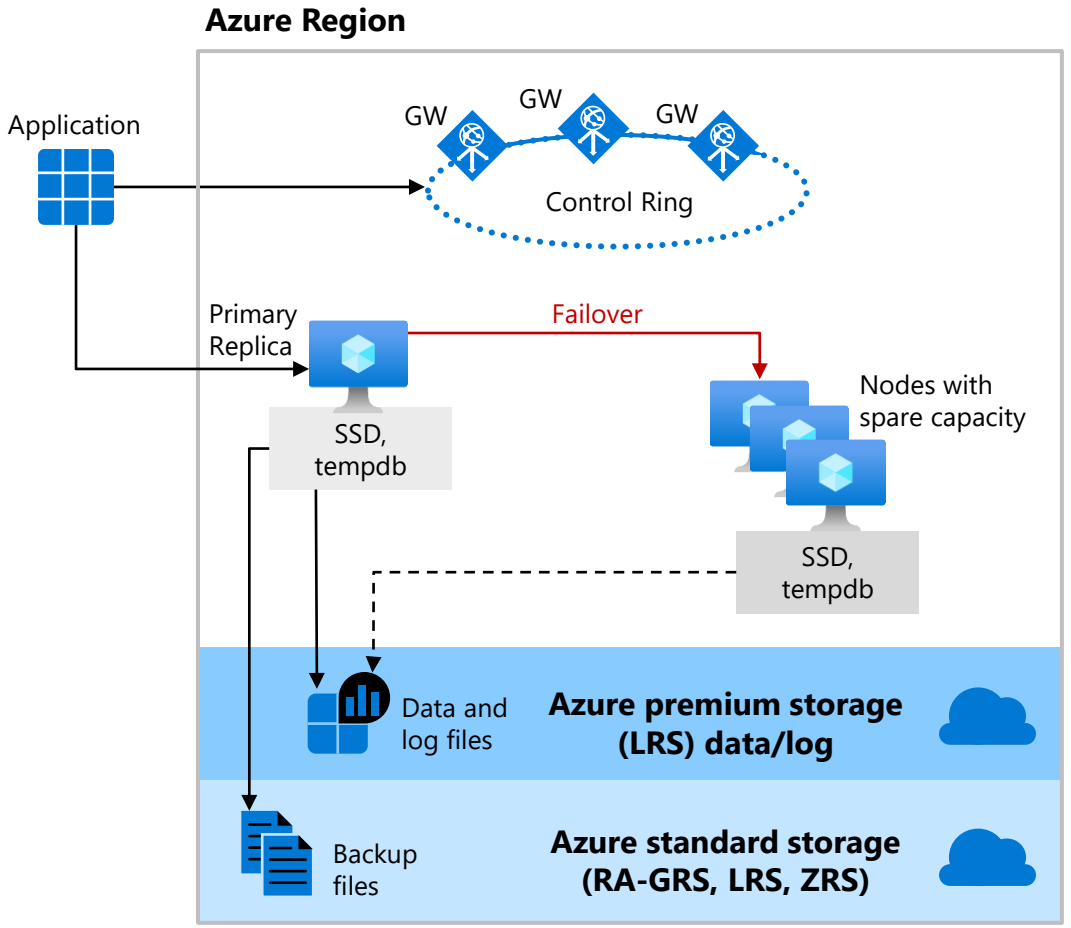

High Availability - Business Critical

Premium (DTU) and Business Critical (vCore)

High availability is implemented using a technology like SQL Server Always On Availability Groups.

High availability is achieved by replicating both compute and storage to additional nodes.

The cluster includes a single primary replica for read-write workloads, and up to three secondary replicas.

The failover is initiated by the Azure Service Fabric.

As an extra benefit, the premium availability model includes Read Scale-Out feature.

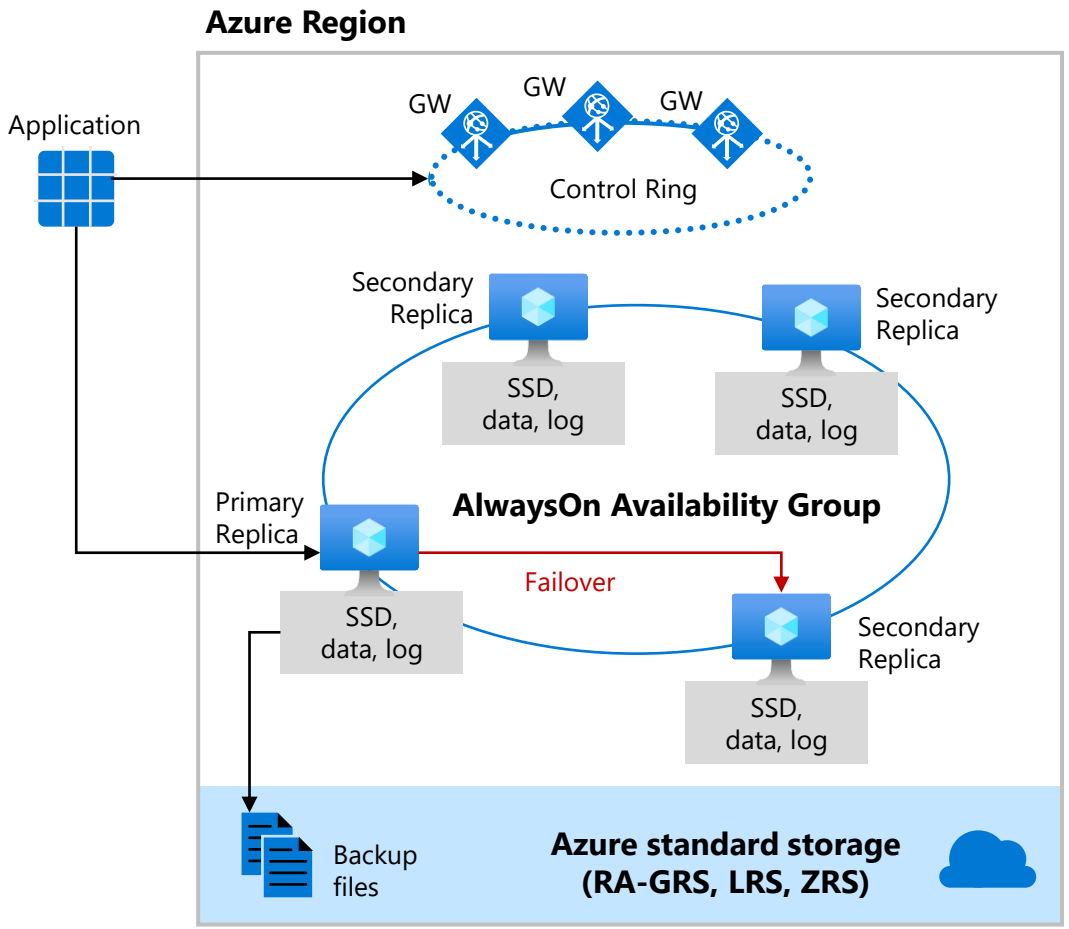

Connection Policy

Redirect (recommended): Clients establish connections directly to the node hosting the database, leading to reduced latency and improved throughput.

Proxy: In this mode, all connections are proxied via the Azure SQL Database gateways, leading to increased latency and reduced throughput.

Default: This is the connection policy in effect on all servers after creation unless you explicitly alter the connection policy to either Proxy or Redirect.

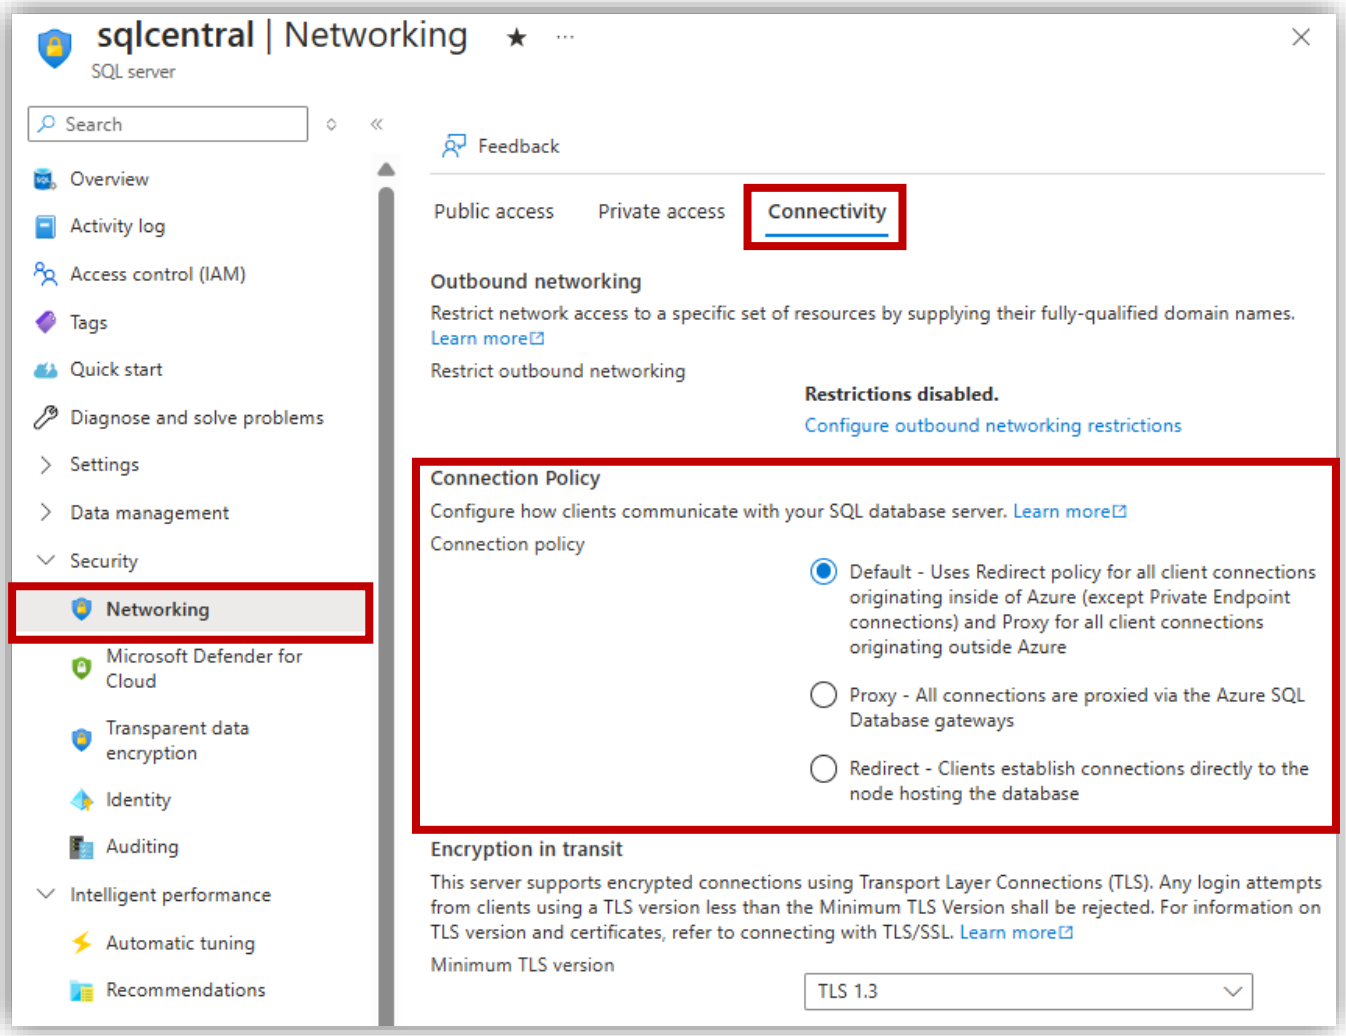

Storage redundancy

To enable high durability of backups several ways of replication are offered on instance creation.

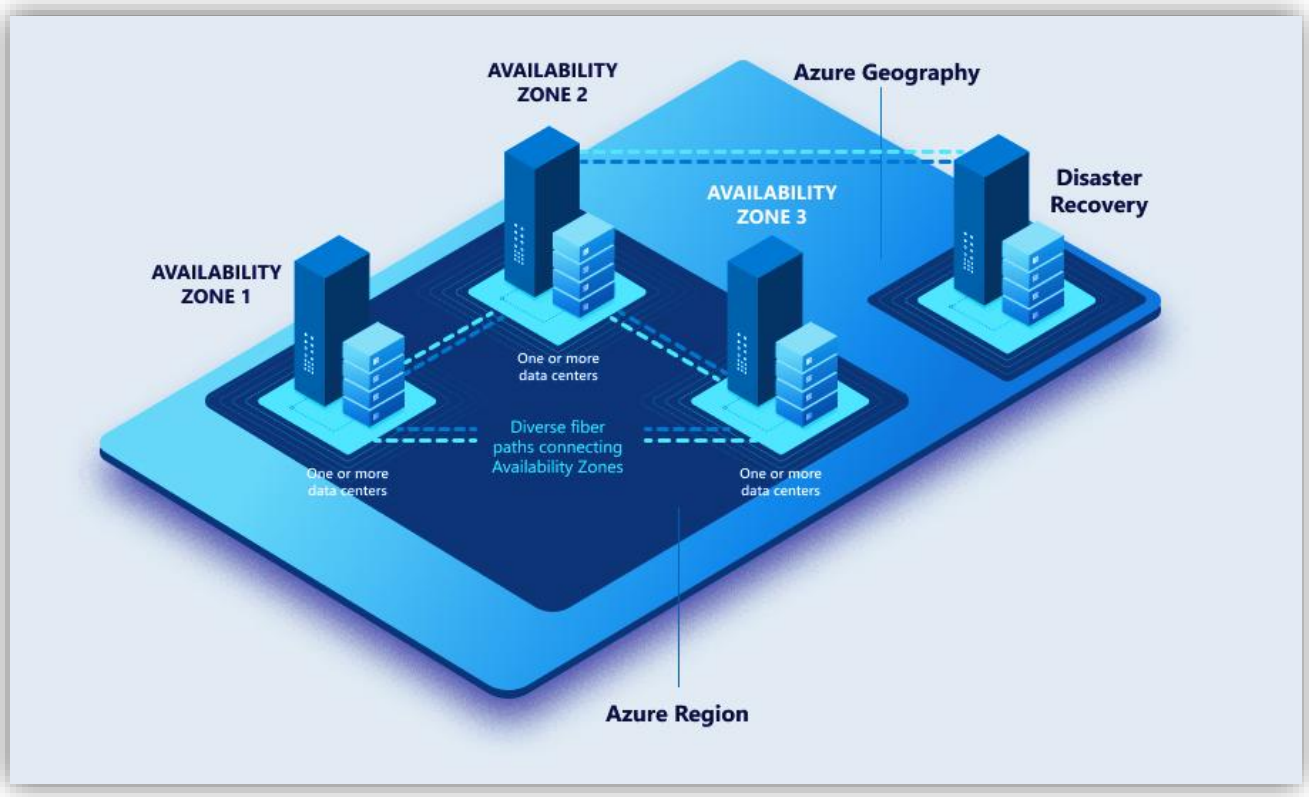

The backups can be all located within

1. LRS: The same building (Local)
2. ZRS: Same region, different buildings (Zone)
3. GRS: Across paired regions (Geo)

Read Scale-Out

Read Scale-Out redirects the read-only client connections to one of the ready-only replicas available instead of sharing the read-write replica.

Use **ApplicationIntent=ReadOnly**; to connect to the read-only replica.

Effectively isolate the read-only workload from the main read-write workload and doubles the compute capacity of the database or elastic pool at no additional charge.

This is ideal to scale-out for complex analytical workloads without affecting the primary OLTP workload.

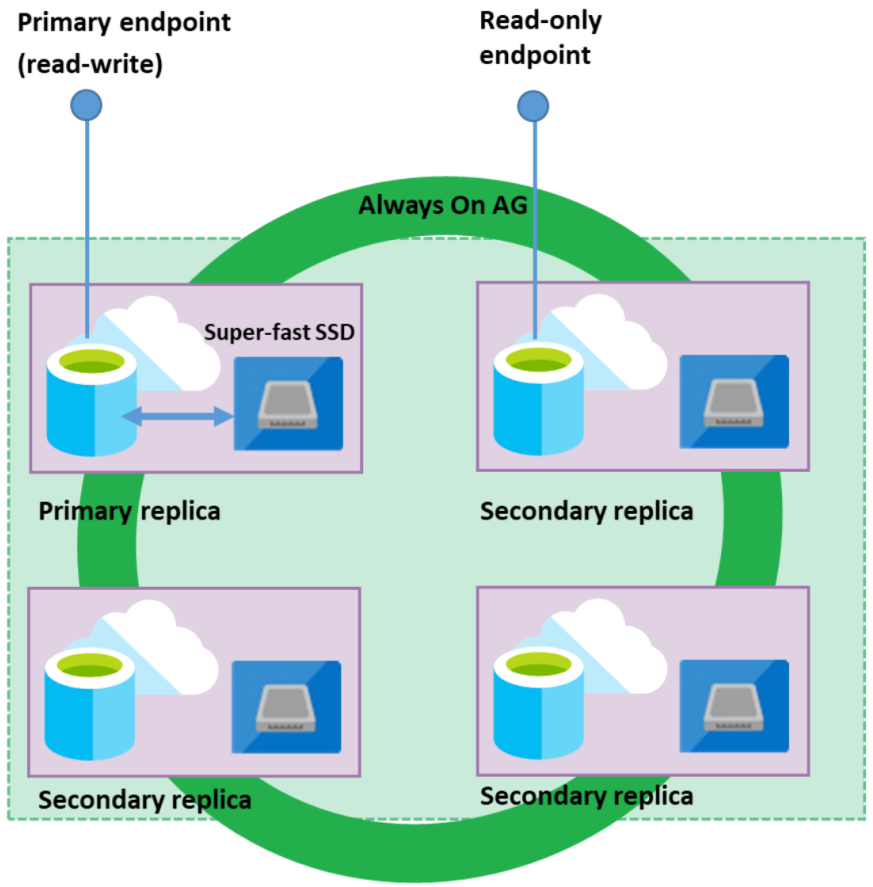

Microsoft

Zone Redundancy in Azure SQL?

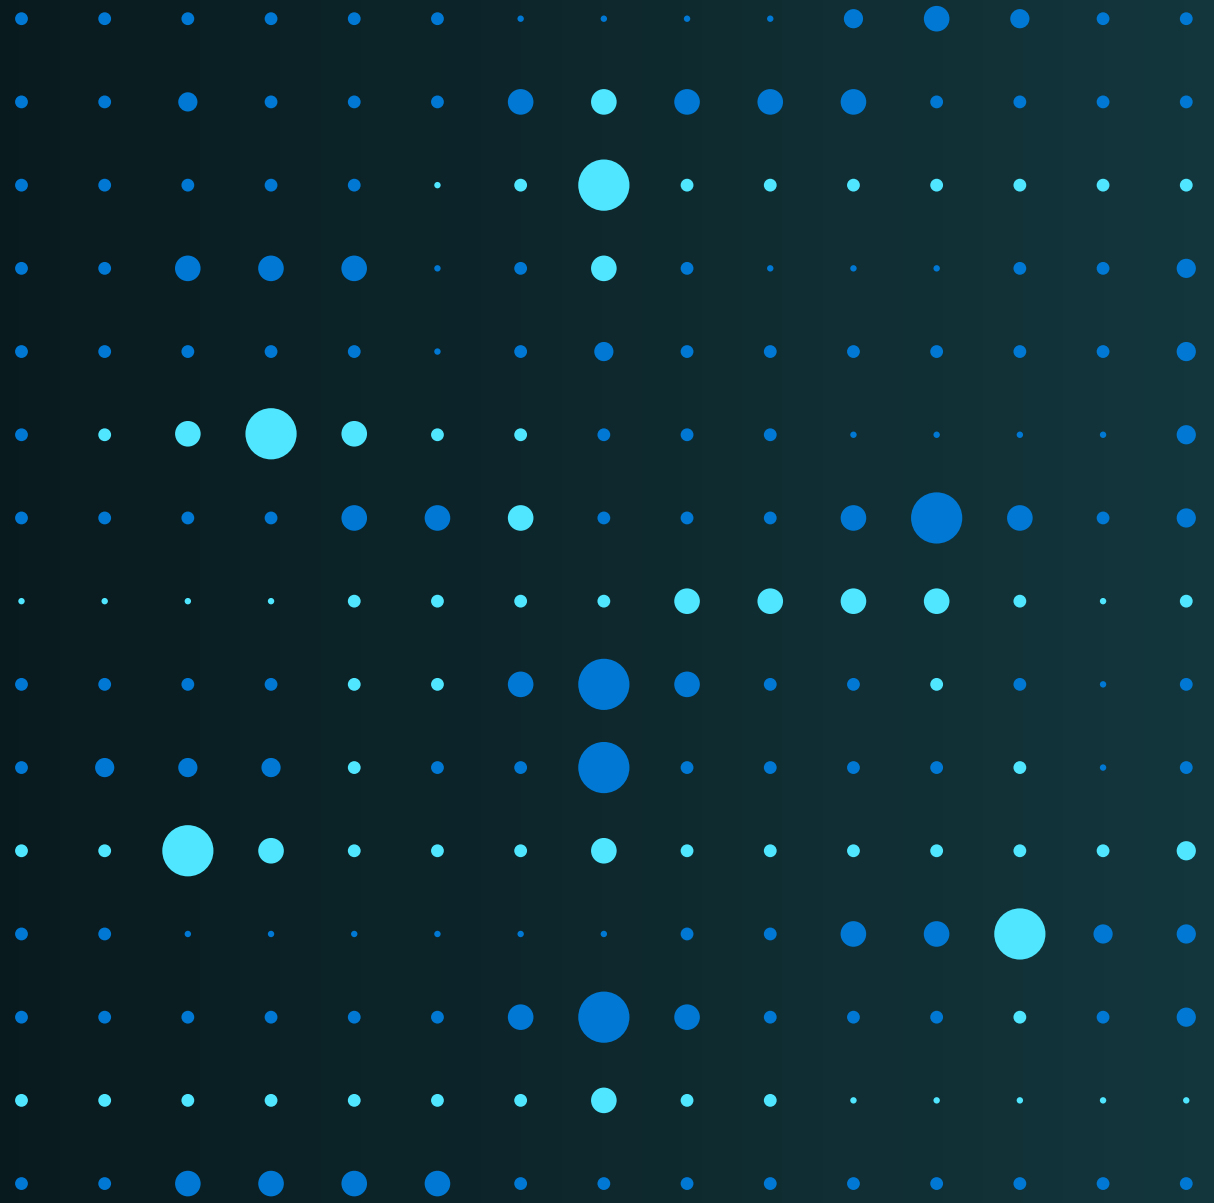

Zone redundant configuration – Premium and Business Critical

By default, the cluster of nodes for the premium availability model is created in the same datacenter.

SQL Database can place different replicas of the Business-Critical database to different availability zones in the same region.

The routing is controlled by Azure Traffic Manager (ATM).

The zone redundant databases have replicas in different datacenters with some distance between them, the increased network latency may impact the performance.

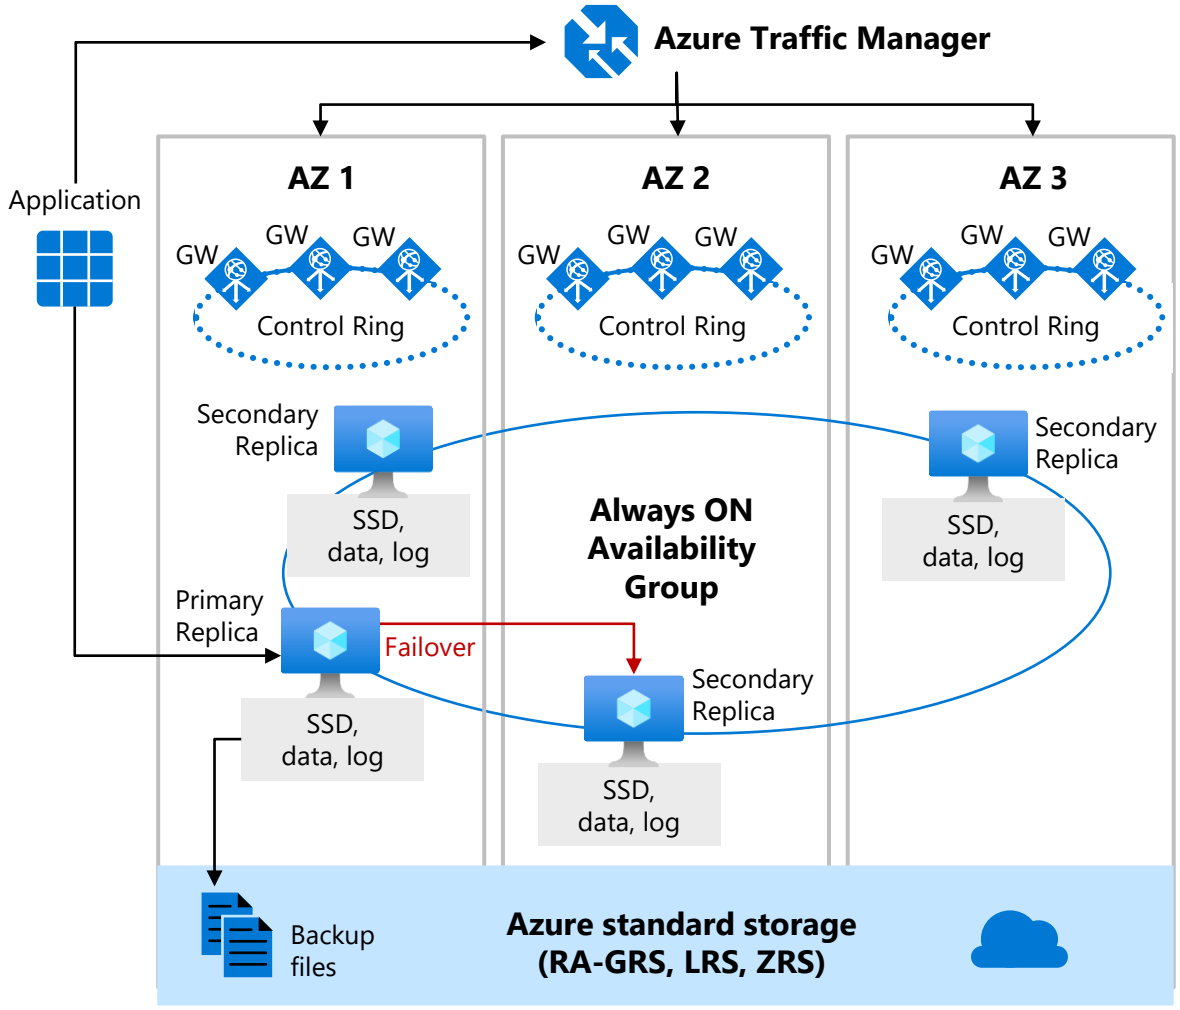

Zone redundant configuration – General Purpose

Zone redundant configurations for the General-Purpose service tier became Generally Available in December 2022. Only available in select regions.

A stateful data layer with the database files (.mdf/.ldf) that are stored in Zone Redundant Storage (ZRS) and are synchronously copied across three physically isolated Azure availability zones.

For zone-redundant serverless and provisioned General-Purpose databases, nodes with spare capacity are readily available in other Availability Zones for failover.

Zone redundancy for the General-Purpose service tier does come with an additional cost that includes charges for the extra resources required to maintain nodes with spare capcity across multiple Availability Zone

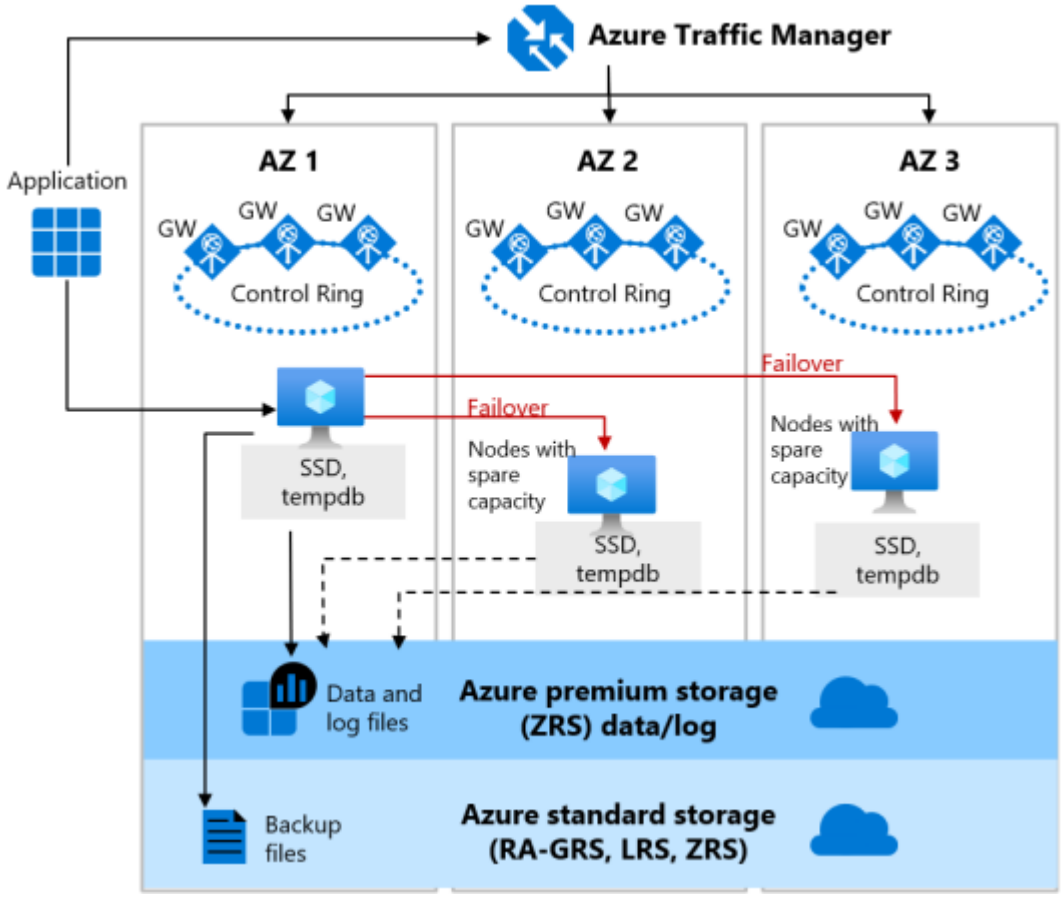

Demo Time

Microsoft

Disaster Recovery in Azure SQL?

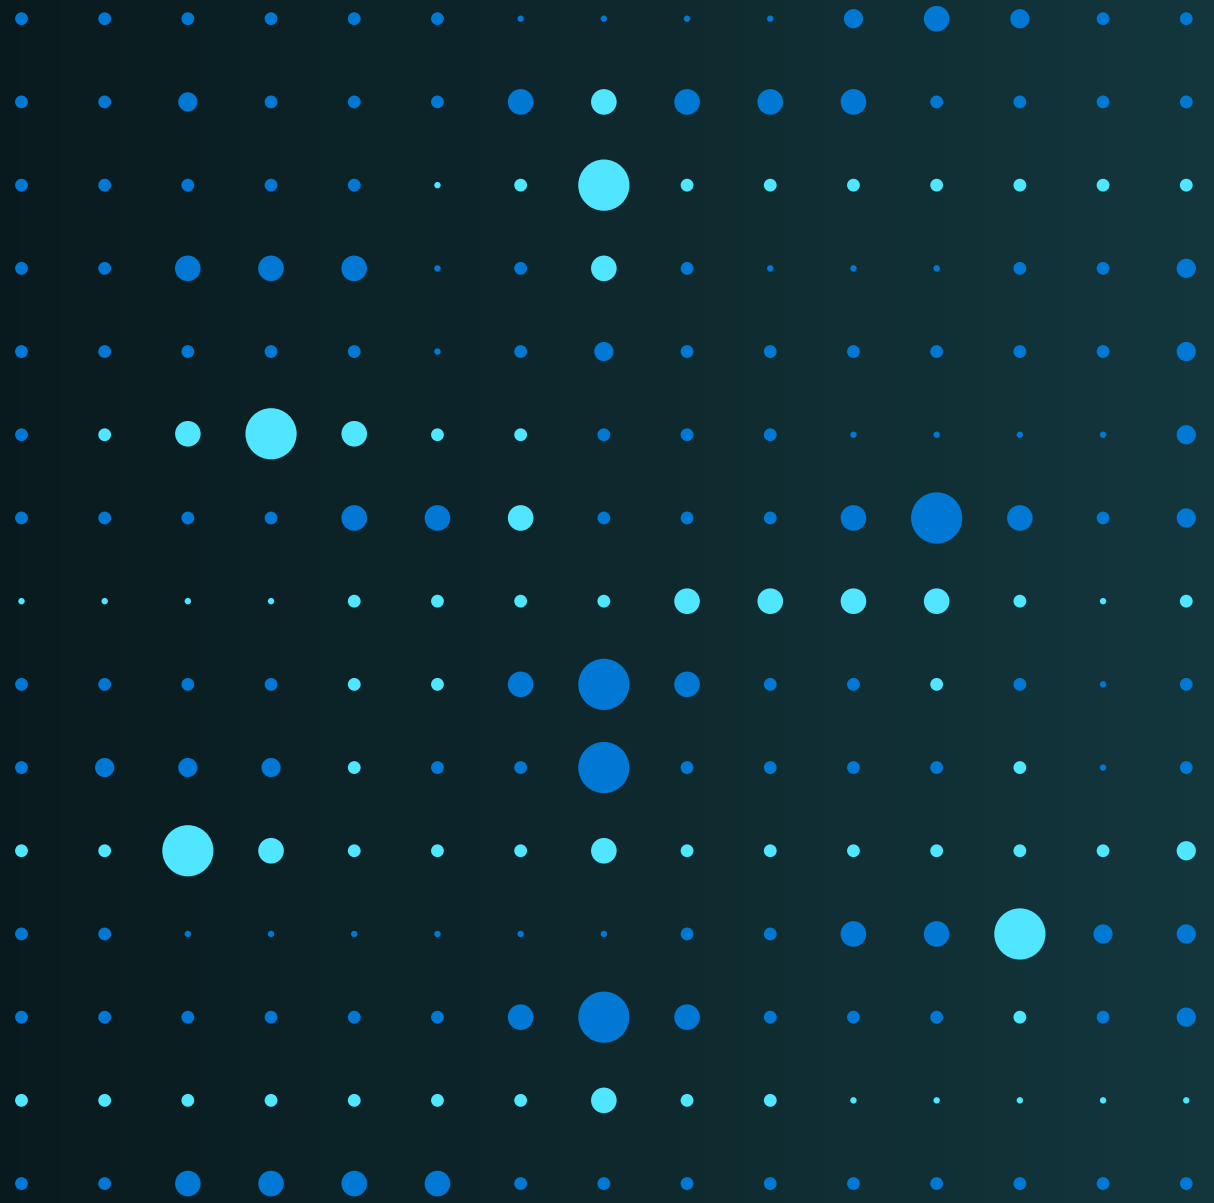

Active Geo-replicas

Service levels	Basic, standard, premium Self service
Readable secondaries	Up to 4
Regions available	Any Azure region
Replication	Automatic, asynchronous
Manageability tools	REST API, PowerShell, or Azure Portal
Recovery time objective (RTO)	<1 hour
Recovery point objective	<5 minutes
Failover	On demand

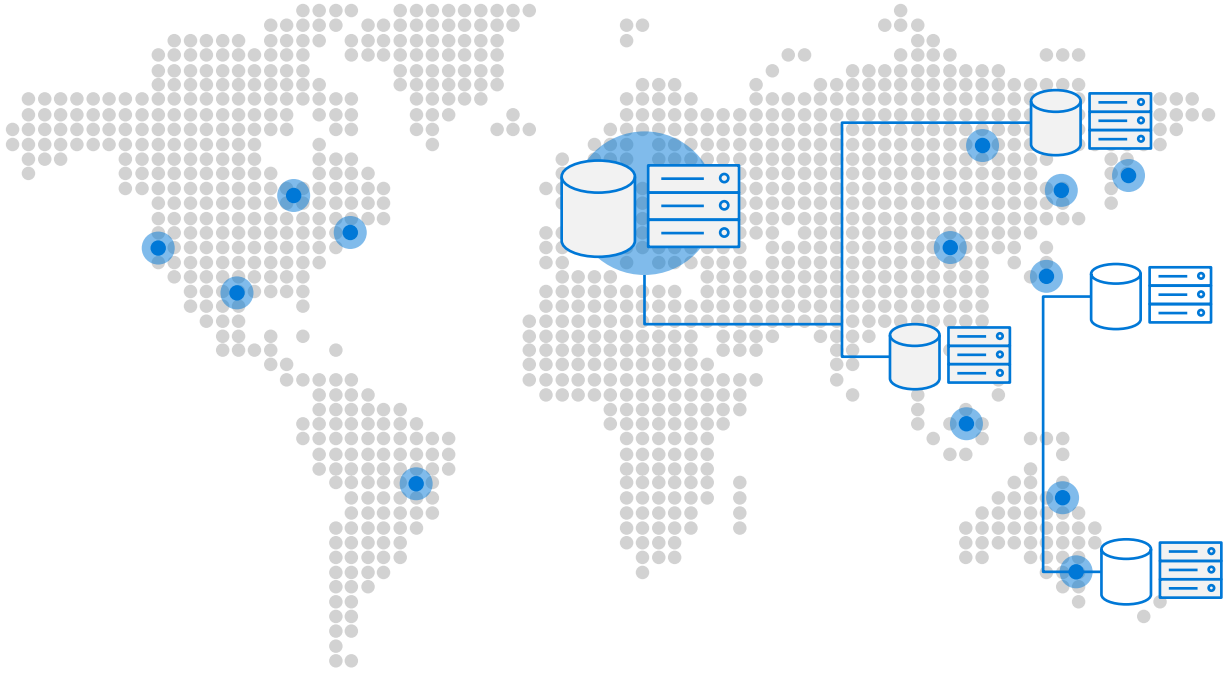

Up to 4 secondaries

Stand-by Replicas

Available for General Purpose or Business Critical service tiers.

A secondary database replica that is used *only* for disaster recovery. Cannot have any workloads running on it, or applications connecting to it.

Provides you with the number of vCores licensed to the primary database at no extra charge under the failover rights benefit.

Save on licensing costs up to 40%. You're still billed for the compute and storage that the secondary database uses.

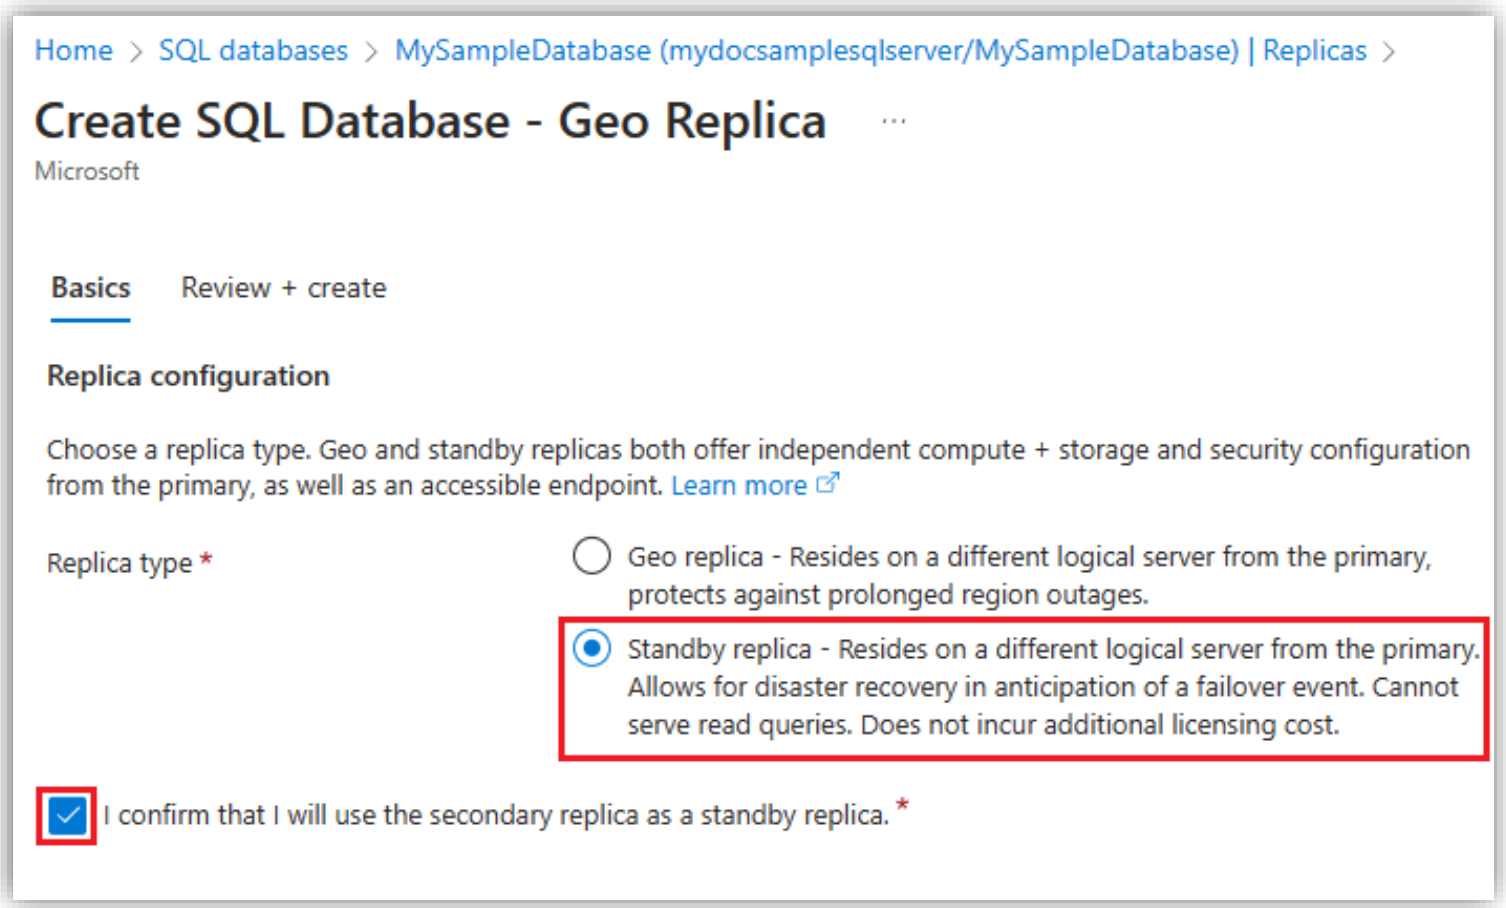

Failover groups extend geo-replication

- Enables geo-replication for a group of databases within a server.
- Automatically or manually failover a group of databases.
- Available for all service tiers.
- Configure the auto-failover policy that best meets your application needs.
- Usage of and listener end-points.
- DNS record is automatically updated.

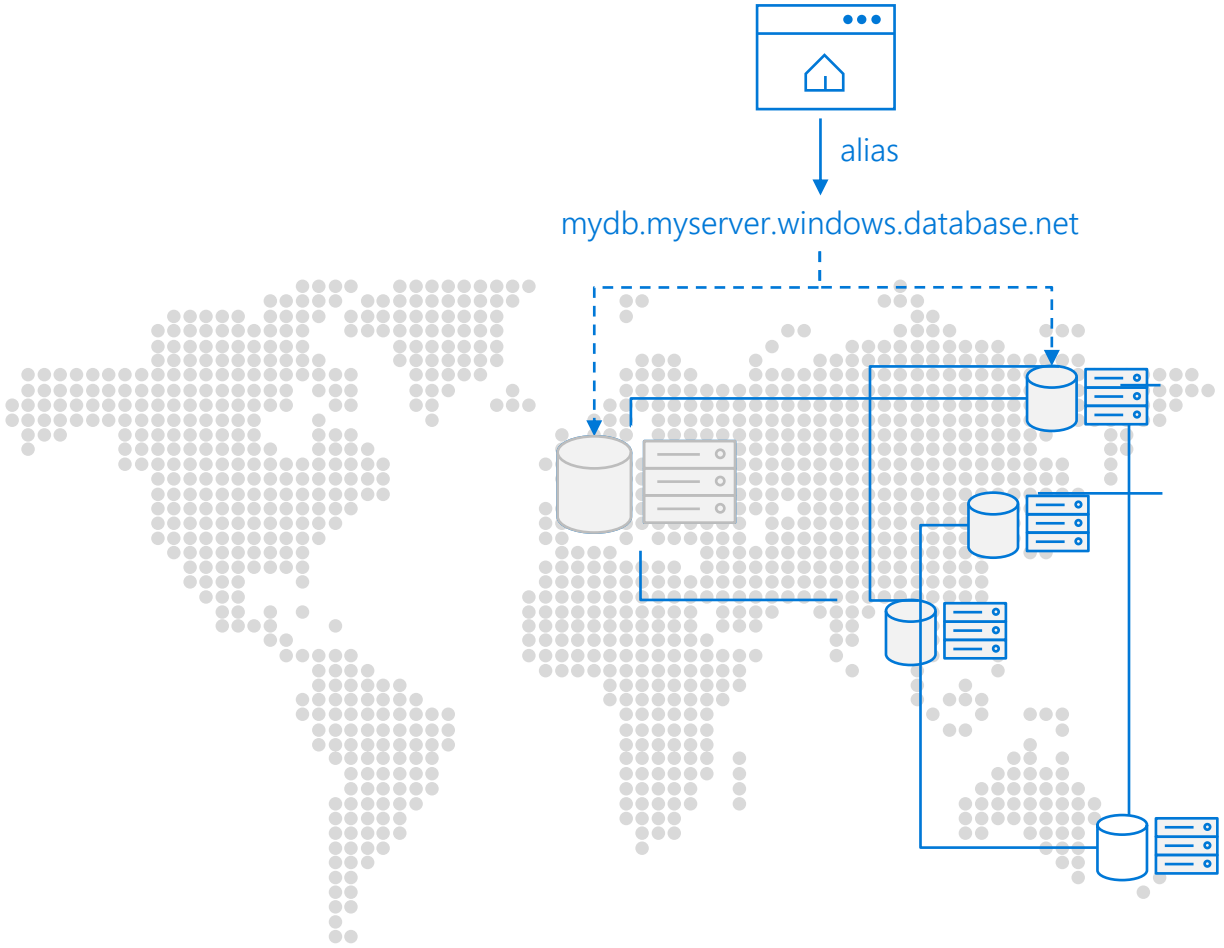

Active geo-replication vs auto-failover groups

	Geo-replication (Database)	Auto-failover groups (Server)
Automatic failover	No	Yes
Fail over multiple databases simultaneously	No	Yes
Update connection string after failover	Yes	No
Managed instance supported	No	Yes
Can be in same region as primary	Yes	No
Multiple replicas	Yes	No
Supports read-scale	Yes	Yes

Demo Time

Questions?

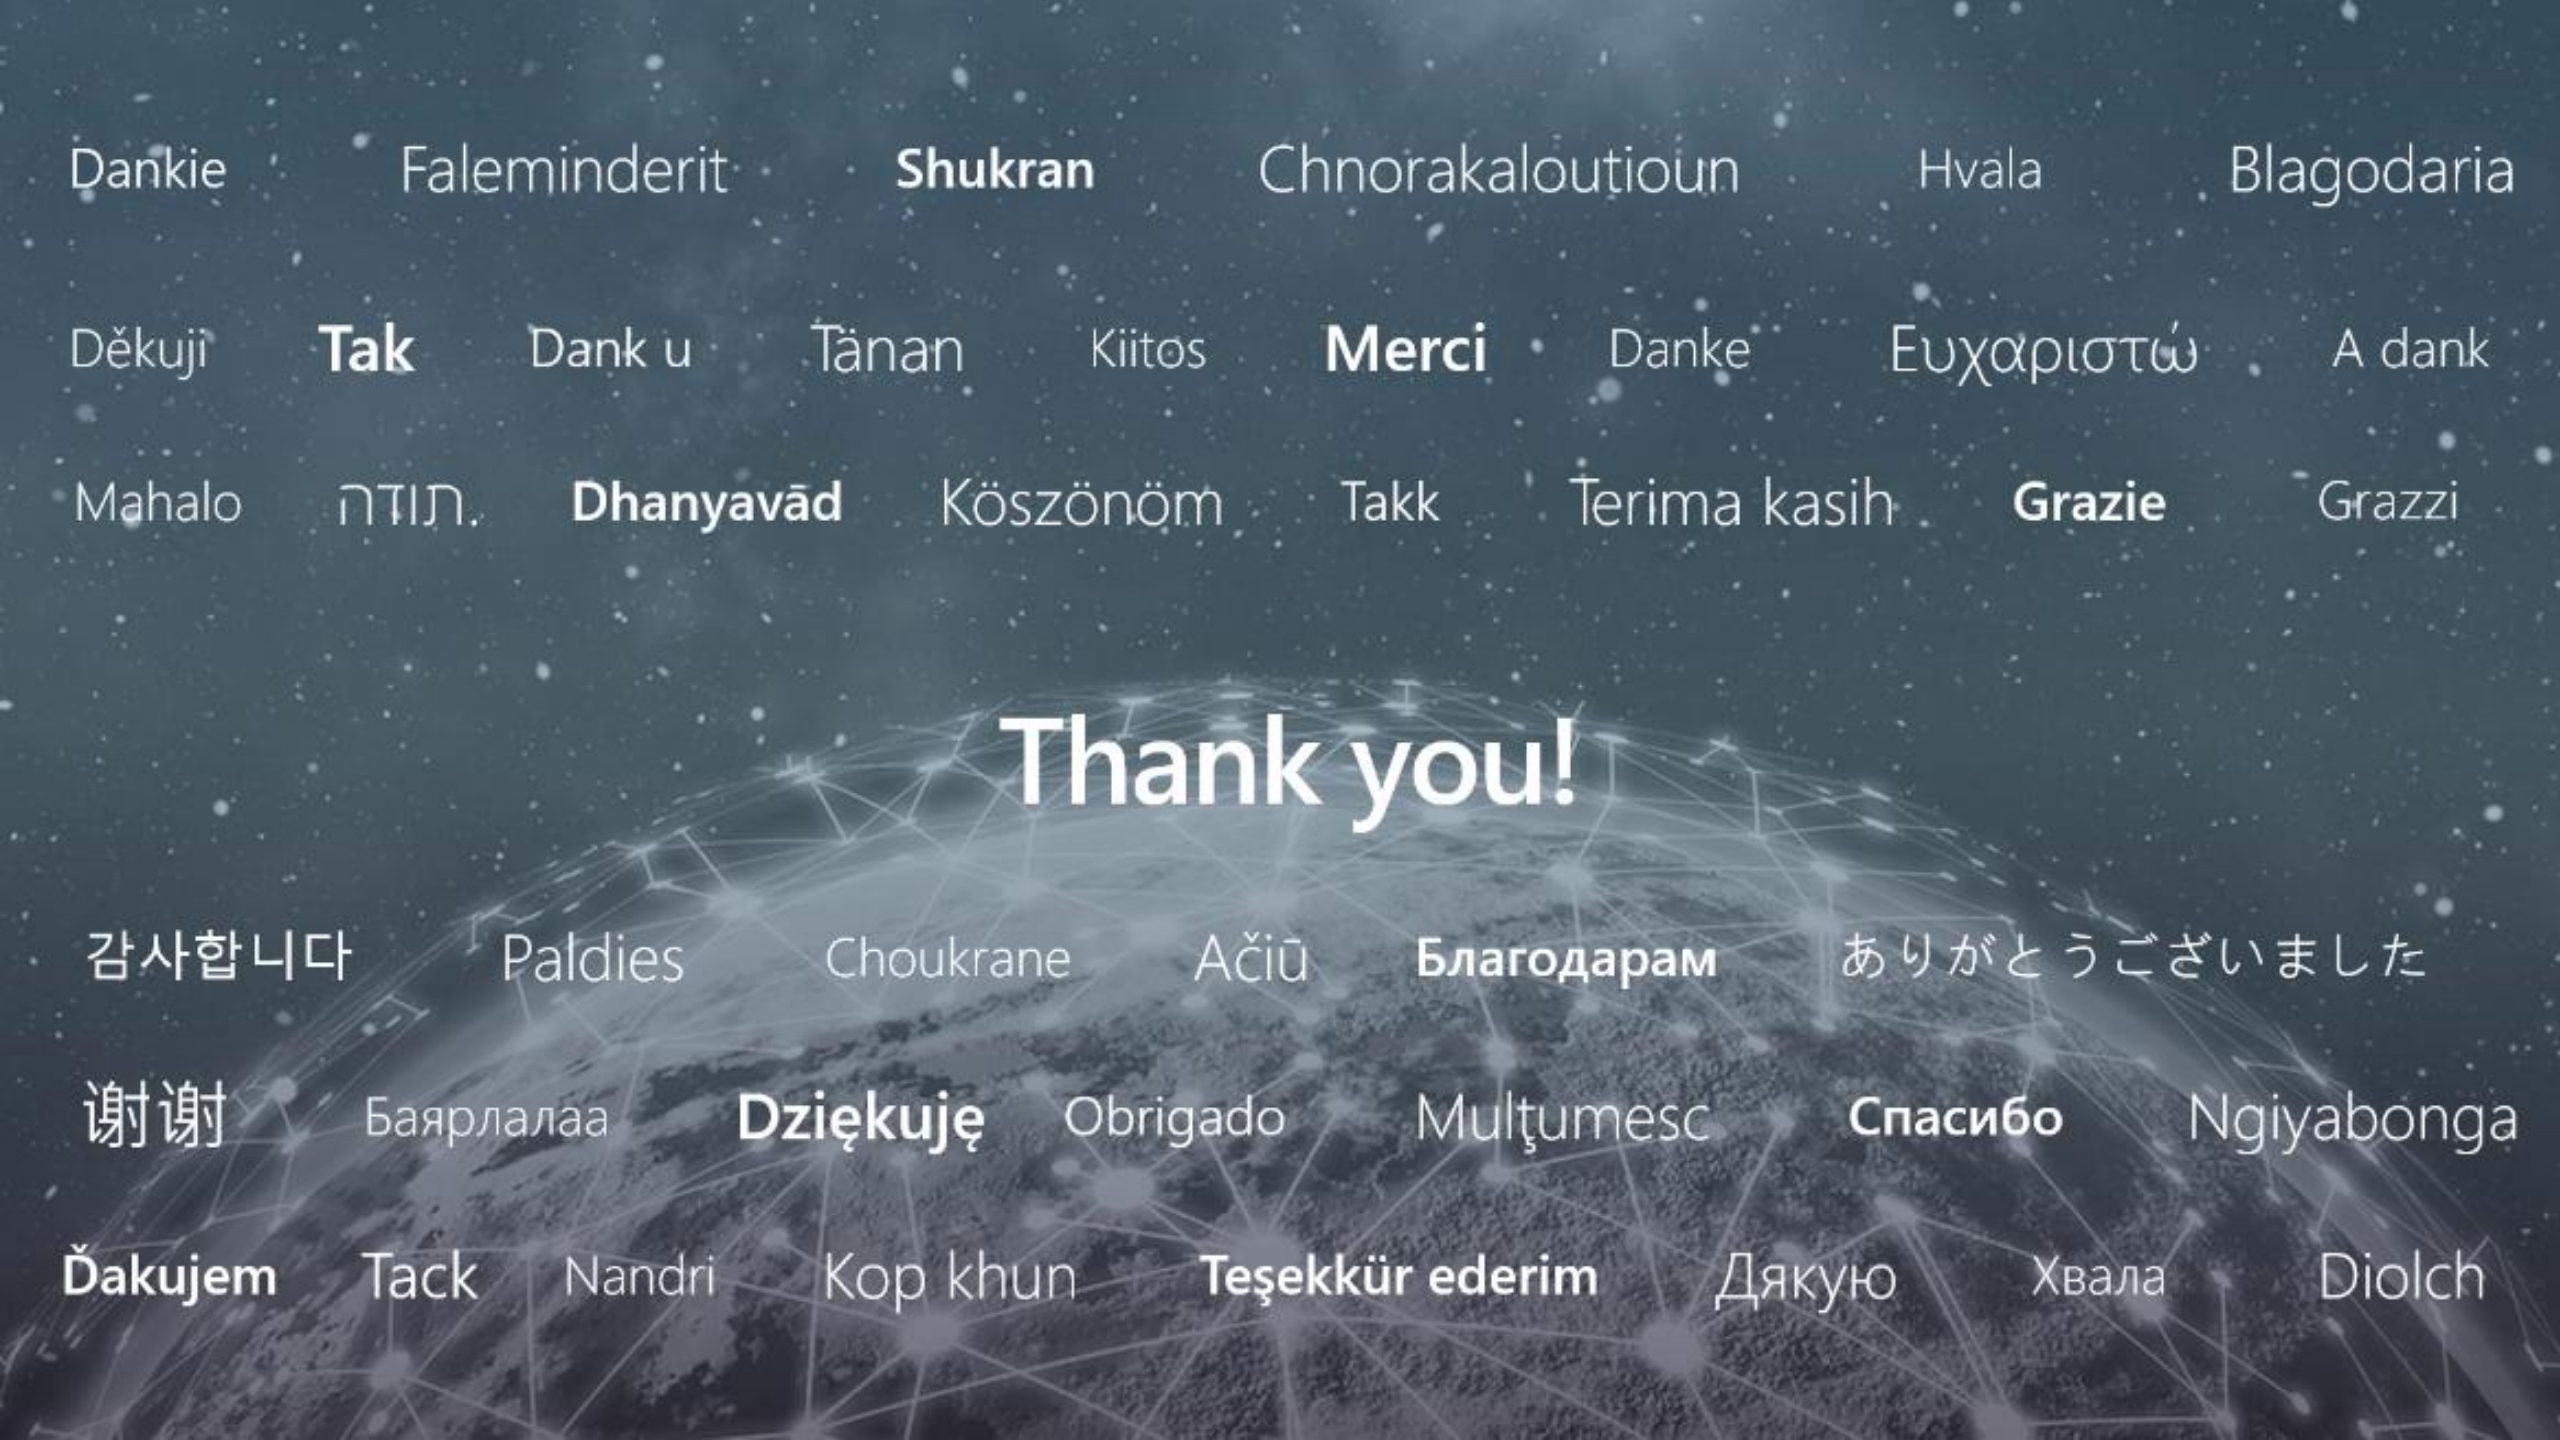
Dankie
Faleminderit
Shukran
Chnorakaloutioun
Hvala
Blagodaria
Děkuji
Tak
Dank u
Tänan
Kiitos
Merci
Danke
Ευχαριστώ
A dank
Mahalo
תודה.
Dhanyavād
Köszönöm
Takk
Terima kasih
Grazie
Grazzi
Thank you!
감사합니다
Paldies
Choukrane
Ačiū
Благодарам
ありがとうございました
谢谢
Баярлалаа
Dziękuję
Obrigado
Mulţumesc
Спасибо
Ngiyabonga
Ďakujem
Tack
Nandri
Kop khun
Teşekkür ederim
Дякую
Хвала
Diolch