

22nd June 2021

SQL Azure Hyperscale. O como Microsoft rompe los límites de SQL en la nube.





www.dotnet2021.com

ORGANIZATION

IN COOPERATION WITH

"Educando con una sonrisa."

Verónica Bas Cordobés

Cloud Solution Architect

ONLINE TECH CONFERENCE

#DotNet2021

Juan Carlos Marínez Riesgo

Customer Engineer

Azure SQL Database — Everything built-in

The intelligent relational cloud database service

Intelligent performance

Scales on the fly

Realize automatic performance improvements from continuous assessment and innovation

Change service tiers, performance levels, and storage dynamically without downtime

Business continuity

Easily manage and monitor business critical functions for reliable operations

Works in your environment

Develop your app and connect to SQL Database with the tools and platforms you prefer

Advanced threat protection

Build security-enhanced apps with built-in protection and industryleading compliance

ONLINE TEC

DotNet2021

Challenges with managing Very Large Databases (VLDBs)

Size of data

Operations take a LONG time (days in some cases) Ongoing operations degrade database performance Higher probability of failure for longer operations Provisioning more storage to expand the database can be painful

Logistics of moving to larger box Economics of sizing for max peaks

Hyperscale is the foundation for massive app growth

Hyperscale is a new, highly scalable service tier that adapts on-demand to your workload's needs, auto-scaling up to 100TB per database.

- Storage dynamically adapts to your workloads' needs, auto-scaling up to 100TB.
- Provision one or more additional compute nodes that can serve your read-only workload and use them as a hot-standby, in case of failover.
- Perform operations in constant time, regardless of the size of the data operation.
- Compute and storage resources scale rapidly and independently without sacrificing performance.

ONLINE TECH CONFERENCE

Azure SQL Database Hyperscale

Cloud native

Architected for cloud

Storage

Scalable new tiered storage architecture

Large database

Support for 100TB+ Infinite Log

Read scale with replicas

Supports 4 Read replicas

#DotNet2021

Seamless compatibility

Fully compatible with Azure SQL Database

Scale

Scale compute and storage separately Scale up in constant time.

Low latency, high throughput for large databases Snapshot-based backups – no impact on query performance Rapid database restore

Performance

DotNet2021 **Common Use Cases**

Cloud-native solution for VLDB management challenges Attractive platform for *databases of any size*, not just VLDBs

Real-time analytics/HTAP

OLTP for databases of any size, at better price/performance ratio than Premium/Business Critical Operational data stores, data marts, SMP data warehouses

#DotNet2021

Data ingestion at scale, IOT, logs

If you are currently running interactive analytics queries using SQL Server as a data warehouse, Hyperscale is a great option because you can host small and mid-size data warehouses (such as a few TB up to 100 TB) at a lower cost, and you can migrate your SQL Server data warehouse workloads to Hyperscale with minimal T-SQL code changes.

ONLINE TECH CONFERENCE

#DotNet2021

If you are running data analytics on a large scale with complex queries and sustained ingestion rates higher than 100 MB/s, or using Parallel Data Warehouse (PDW), Teradata, or other Massively Parallel Processing (MPP) data warehouses, Azure Synapse Analytics may be the best choice.

Choose from 3 service tiers

	General purpose	Hyperscale	Business cri
Best for	Most budget- oriented workloads	VLDB OLTP and HTAP workloads with highly scalable storage and read- scale requirements	Critical business applications wit high IO requirements .
Compute tiers		Gen4: 1 to 24 vCore Gen5: 2 to 80 vCore	•
Storage	Premium remote 32GB – 8TB per instance	Local SSD cache backed by Azure Storage Auto-scale up to 100TB of storage	Local SSD 32GB – 4TB per instance
HA	2 read replicas	Primary read/write replica + up to 4 read replicas	3 replicas, 1 rea scale replica, zo redundant HA
In-Memory	Not supported	Not supported	Supported
Read-write IO	Low	Medium	High

DotNet2021 ONLINE TECH CONFERENCE Hyperscale components

#DotNet2021

Log Service Local SSD Cache

DotNet2021

Hyperscale architecture

ONLINE TECH CONFERENCE

DotNet2021 Write IO **Secondary** Primary Tx commit Compute Compute <2.5ms Log landing zone Long term log storage

Log Service

Memory

cache

Primary compute write log to the log landing zone - <2.5ms latency

Commit transactions after hardening the log

Async log apply to the secondary compute

Async log apply to page servers

Page server writes data to data files at checkpoints

Read

request

#DotNet2021

Pre-build RBPEX when page server instance started

Two page-server replicas guarantee IO latency

RBPEX on compute nodes is proportional to # of vCores

Hit local RBPEX - < 0.5ms Hit page server RBPEX - < 2ms

Optimized for OLTP workload – operating on hot data

Use Column Store index to optimize HTAP/OLAP workload

AUTOMATIC Compute Node Scale-UP

- Create an alert in Azure Monitor CPU > 70%
- Create an action associated to the alert to execute a runbook.
- Execute a prepared workload to lead CPU to achieve the threshold.
- Keep the workload running at least 1 minute. 4.
- Show how SLO is changed from 2 Vcores to 8 Vcores in a very short time. 5.
- Show how CPU pressure is alleviated. 6.

ONLINE TECH CONFERENCE

DotNet2021 **Backup & Restore**

Data Migration Options

ONLINE TECH CONFERENCE

- Azure Data Migration Service
 - Bulk load/bulk copy
 - Using Azure Data Factory
 - Using Spark connector for SQL
 - Using SqlBulkCopy (sample: smartbulkcopy)
 - Transactional/snapshot replication
 - Migration from another service tier is just a service lacksquareobjective change
 - Migration is done online ۲
 - Avoid intense write workload on the source during migration lacksquare

- Currently cannot go back to prior service tier •
- Can restore to a point in time before migration, within retention period

Customer stories

Customer: ClearSale

Industry: Partner Professional Services

Size: 1,000–9,999 employees

Country: Brazil

Products and Services: Microsoft Azure Azure Event Hubs **Azure Functions** Azure SQL Database hyperscale service tier

Read full story here

on us when it counts."

—Jelther Goncalves, Data Engineer, ClearSale

Situation:

ClearSale helps major brands identify ClearSale needed to migrate its fraudulent transactions on their platform, applications, and databases to the cloud and e-commerce sites. But its onpremises platform had become chose the Microsoft Azure SQL increasingly difficult to scale for Database hyperscale service tier, Azure Event Hubs, and Azure customer growth or peak shopping demand. Functions to achieve greater scalability and flexibility.

"Working in the Azure SQL Database hyperscale tier helps us streamline upgrades to new and existing applications, and with instant, unlimited scaling, customers can rely

Solution:

Impact:

Since beginning its migration and building a new cloud-based fraud prevention platform, ClearSale can now process virtually an infinite number of transactions a minute with automatic scaling to accommodate organic growth and spikes in transaction rates or volumes.

Customer: Clearent

Industry: Banking and Capital Markets

Size: 50–999 employees

Country: United States

Products and services: Microsoft Azure Microsoft Azure SQL Database Microsoft Power BI

Read full story here

-Zach Fransen, Director of Data Management, Clearent

Situation:

Clearent processes around 500 million payment transactions a year—a number that's steadily growing. With data coming from dozens of microservices and customers wanting faster insights, the company needed to replace its outdated infrastructure.

"We have an enormous amount of data that needs to be available to a large number of employees and customers, and the hyperscale service tier is perfect for hitting that sweet spot of concurrency and scale."

Solution:

Clearent chose the Microsoft Azure SQL Database hyperscale service tier as its new data management solution, providing seamless transactional scalability for its modern microservices architecture and rapid insights with concurrent access for tens of thousands of users.

Impact:

Clearent uses the hyperscale tier's automatic scaling to handle any amount of data with confidence and, most importantly, to quickly serve that data back to its customers. The result? An improved developer experience, simpler IT management, and better business insights for everyone.

Customer: Big Red Cloud

Industry: Partner Professional Services

Size: 1 – 49 employees

Country: Ireland

Products and services: Microsoft Azure Azure SQL Database

Read full story here

glad we did."

—Marc O'Dwyer, Chief Executive Officer, Big Red Cloud

Situation:

A leading accounting software company in After deploying the Microsoft Ireland, Big Red Cloud lightens the Azure SQL Database Hyperscale bookkeeping load for small to mediumservice tier to elasticize its sized businesses (SMBs). With onexisting databases, Big Red premises or traditional cloud Cloud could focus on delivering infrastructure, its growth might have been timely, innovative features to too much of a good thing. It needed a help its customers stay cloud database solution that could scale compliant—and competitive. with its business.

"In its infancy, people talked about the 'bet' of the cloud. But we said, 'We're going to take that bet.' We decided to align ourselves with one company—Microsoft. We're

Solution:

Impact:

Big Red Cloud now has just under 25% of the Irish market share in accounting software for SMBs. It scaled from 5,000 to 6,500 customers in just two years. And with its SQL Database Hyperscale tier in place, the company knows it can continue to grow rapidly without negative impacts to its customers.

Thanks and ... See you soon!

Thanks also to the sponsors. Without whom this would not have been posible.

plain concepts

www.dotnet2021.com

#DotNet2021

My Public[®] Inbox

Devs DNA _m

