IDERA® | Cloud Whitepaper

AZURE SQL HYPERSCALE HIGH-PERFOMANCE DBAAS

BY RICHARD GILES

INTRODUCTION

Azure SQL Hyperscale is a new database as a service (DBaaS) offering from Microsoft. Microsoft built Azure SQL Hyperscale on the experiences learned by Microsoft Azure SQL Database. Azure SQL Hyperscale strives for improvements in security and availability plus lower and more flexible cost, and higher performance.

DESIGN GOALS

The Azure SQL Hyperscale database has the following design principles:

- It supports a database size of up to 100 TB
- It provides nearly instantaneous database backups (based on file snapshots stored in Azure Blob storage) regardless of size with no input/output impact on Compute
- It provides fast database restores (based on file snapshots) in minutes rather than hours or days
- It provides higher overall performance due to a higher log throughput and faster transaction commit times regardless of data volumes
- It provides rapid scale-out by provisioning one or more read-only nodes for offloading the read workload and for use as hot-standbys
- It provides rapid scale-up in constant time of the compute resources to accommodate heavy workloads when needed, and then scale the compute resources back down when not required

PERFORMANCE COMPARISON

The following chart shows a comparison of the Azure SQL Database service versus the Hyperscale service tier:

	Azure SQL Database	Azure SQL Hyperscale
Maximum Database Size (TB)	4	100
Availability (%)	99.99	99.999
Upsize/Downsize	O (data)	O (1)
Storage Impact	4x copies (+backup)	2x copies (+backup)
CPU Impact	4x single images	25% reduction
Recovery	O (1)	O (1)
Commit Latency (ms)	3	<0.5
Log Throughput (MB/s)	50	>100

For upsize/downsize, the table expresses the values in Big-O notation. O(data) or O(n) would be time dependent on the size of the data with larger values taking longer to complete ('linear time'). O(1) corresponds to a time step of 1 to complete regardless of size ('constant time').

SQL SERVER FEATURES USED

The design of Azure SQL Hyperscale made use of the following existing SQL Server features:

- Page Version Store maintains versions of database records
- Accelerated Database Recovery (ADR) coupled with Page Version Store helps to eliminate the undo phase in many cases providing the ‘constant time’ bounded by the checkpointing interval
- Resilient Buffer Pool Extension (BPE) was made resilient (recoverable after failure) for Azure SQL Hyperscale
- Remote Block I/O (RBIO) Protocol is a strongly typed stateless protocol that distributes components of the database engine across multiple tiers
- Snapshot Backup/Restore was extended to allow backup/restore to work on Azure Storage (XStore) and thus enables multi-terabyte databases to work in ‘constant time’ with impacting CPU or input/output cost in the Compute tier
- I/O Stack Virtualization was exploited to help hide the storage hierarchy from the compute process

ARCHITECTURAL OVERVIEW

The following diagram visualizes the architecture of Azure SQL Hyperscale.

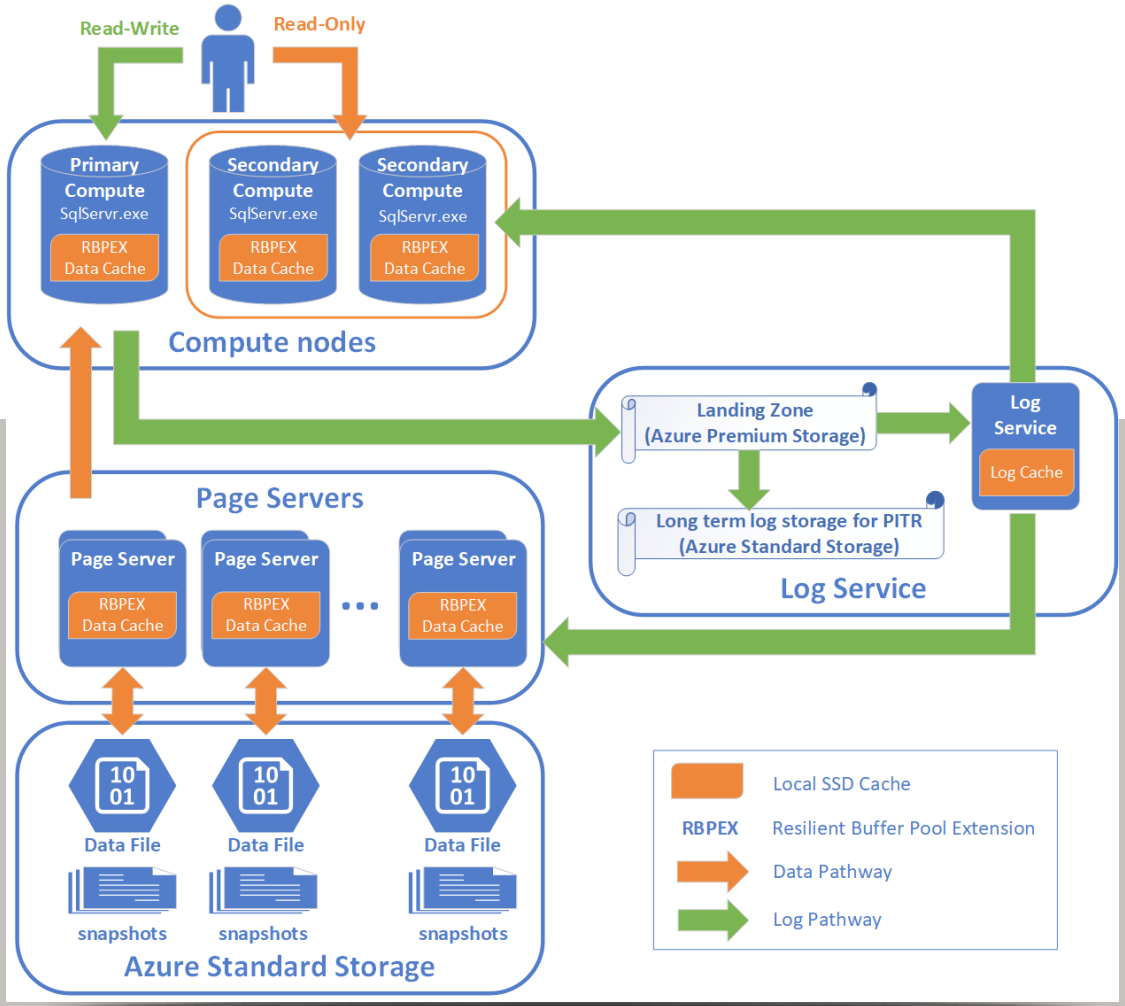

The core Service tiers are as follows:

- Compute nodes is where the relational engine lives
- Page Servers are systems representing a scaled-out storage engine
- Log Service node accepts log records from the primary compute node, persists them in a durable cache, and forwards the log records to the rest of the Compute nodes
- Azure Standard Storage node is the destination for data from Page Servers and is used for backup purposes as well as for replication

CONFIGURATION

Configure a database service tier for Azure SQL Hyperscale by selecting an Azure SQL Database. Then configure the database using the vCore-based purchasing model. The following screen capture image of the Azure Dashboard displays the option for Hyperscale on-demand scalable storage with data up to 200,000 input/output operations per second with a latency of 1 to 2 milliseconds and log up to 7,000 input/output operations per second with a latency of 5 to 10 milliseconds. However, the capability to change from Hyperscale to another service tier is not supported.

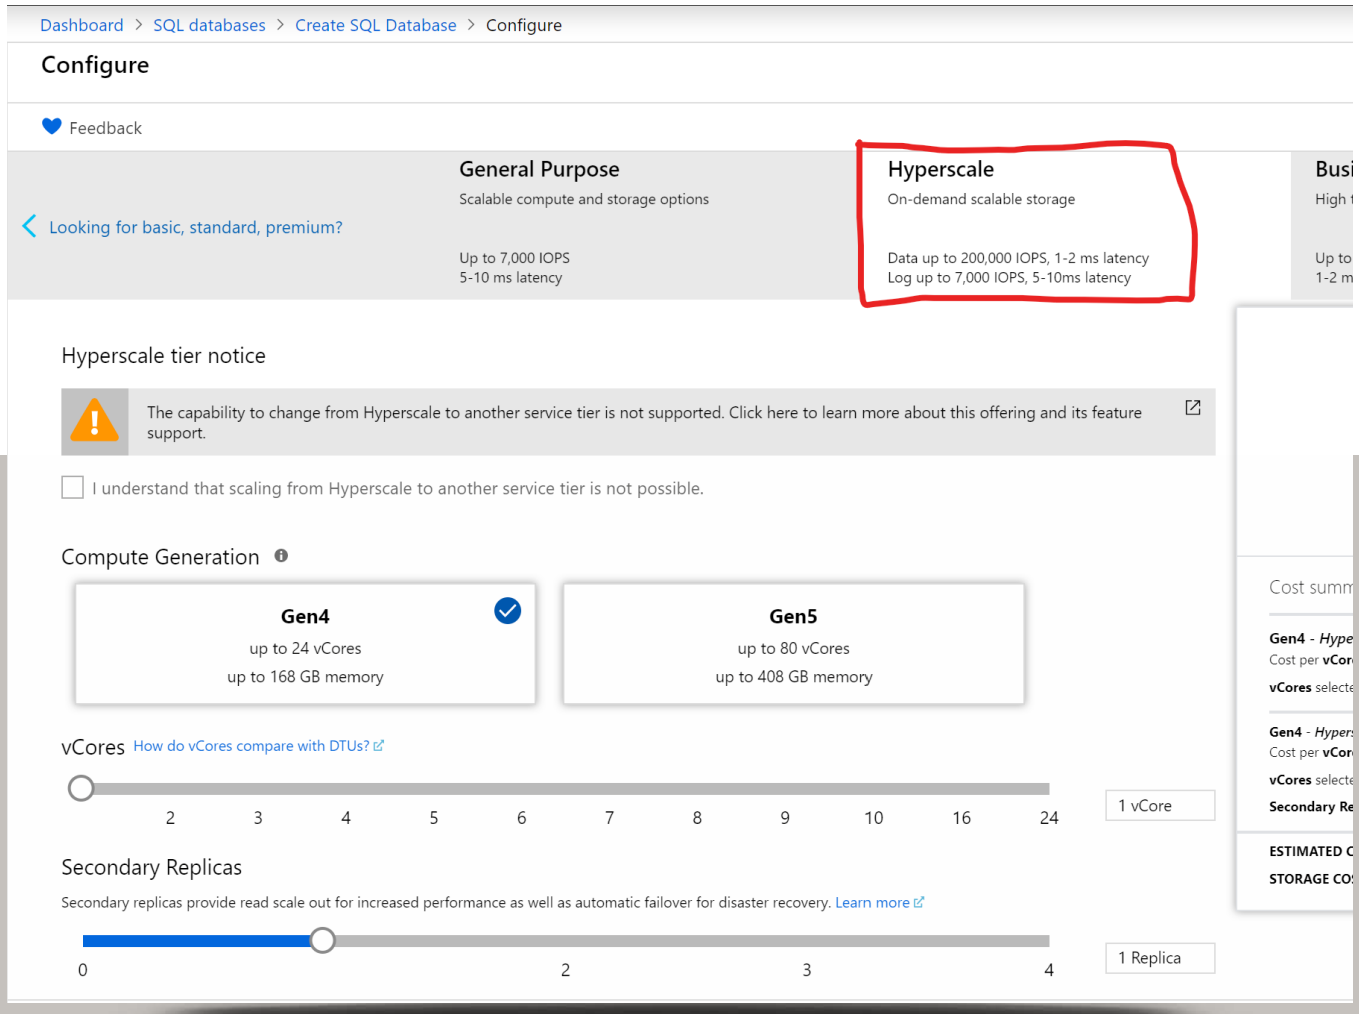

CONCLUSION

FAzure SQL Hyperscale is based on the idea of separating compute and storage that other vendors such as Amazon and Oracle have modeled products on.. However, Azure SQL Hyperscale architecture also separates durability and availability. That separation helps meet a customer need for the balancing of cost, performance, and availability.

Microsoft has released a single database version of Azure SQL Hyperscale. However, Microsoft is working on future improvements for bulk operations, multiple-master variants, better support for the hybrid transaction/analytical processing (HTAP) as can be found in MariaDB, and making use of the Log Service for other services such as audit and security.

Azure SQL Hyperscale is designed to be able to migrate from Azure SQL Databases. And Azure SQL Hyperscale is intended to have no impact on current applications and tools designed for Azure SQL Databases. The lack of adverse effects is good news for third-party vendors, including IDERA.

REFERENCES

‘Socrates: The New SQL Server in the Cloud’; P. Antonopoulos et al.; Microsoft Azure & Microsoft Research; 14 pages; May 2019

‘FAQ about Azure SQL Hyperscale databases’; Microsoft documentation; Azure > SQL Database

‘Hyperscale service tier for up to 100 TB’; Microsoft documentation; Azure > SQL Database

‘Azure SQL Database Hyperscale support for single databases is now available’; Microsoft updates; Azure > SQL Database

Note The Microsoft Research paper refers to Azure SQL Hyperscale by the code name ‘Socrates’. Microsoft does not reference this code name any longer and has replaced it with ‘Azure SQL Hyperscale’ service tier.

ABOUT THE AUTHOR

Richard Giles is an accomplished Software and Database Architect. He has been developing commercial software for over 30 years and has played key roles in building successful products including Microsoft Operations Manager and Idera SQL Server product line. Richard is currently providing Technical Product Consulting services as a Consultant. He has been involved with Microsoft Partner programs for over 20 years and has actively tested Microsoft products as part of the Insider program. Richard has a strong analytical ability, is thorough when researching emerging technology and takes great pride in producing innovative solutions. Always the consummate learner, he is been actively learning about new technologies for most of his career and has focused on Database Platforms, Data Science, Analytics and Cloud Computing.