QCT AZURE STACK HCI:
HIGH-PERFORMANCE MICROSOFT SQL SERVER

Technical Use Cases for Azure Stack HCI

- Branch office and edge
- Virtual desktop infrastructure
- **High-performance Microsoft SQL Server**
- Trusted enterprise virtualization
- Scale-out storage

This solution leverages your Azure Stack HCI investment to run Microsoft SQL Server for highly available and highly performant enterprise database applications. It also provides customers easy backup to Azure with hybrid connectivity built in. Below, you will find a how-to guide for deploying Microsoft SQL Server on Azure Stack HCI that includes:

- Solution Overview
- Step by step documentation to deploy Microsoft SQL Server on Azure Stack HCI

Solution Overview

Azure Stack HCI provides enterprise customers a highly available, cost efficient, flexible platform to run a high-performance Microsoft SQL Server leveraging the power of state-of-the-art hardware and Storage Spaces Direct. Azure Stack HCI presents a highly competitive solution for delivering exceptionally performant Microsoft SQL Server. Whether running Online Transaction Processing (OLTP) workloads, or Data Warehouse and BI, to AI and advance analytics over Big Data, you will benefit from the resiliency that Azure Stack HCI offers. This is especially important for mission critical databases. Leveraging the flexibility to run SQL Server in VMs (Windows Server or Linux), it allows you to consolidate multiple database workloads and easily scale out by adding additional VMs to the Azure Stack HCI environment as needed.

Additionally, Azure Stack HCI enables you to integrate Microsoft SQL Server with Azure Backup service and Azure Blob Storage service to provide cloud-based backup solutions that are reliable and secure.

How to deploy Microsoft SQL Server on Azure Stack HCI

1. Hardware and OS configuration for Azure Stack HCI


QCT QuantaGrid D52BQ-2U Performance

Scale:
- 2 to 4 nodes

Single Node Data:
- CPU: 20-56 cores (Intel)
- RAM: 256GB to 768GB
- Raw storage: 2TB to 48TB
- Storage type: SSD
- Network speed: Up to 25Gb

Learn more >

QCT QuantaGrid D52B-1U

Scale:
- 2 to 4 nodes

Single Node Data:
- CPU: 20-56 cores (Intel)
- RAM: 256GB to 768GB
- Raw storage: 12.8TB to 38.4TB
- Storage type: NVMe
- Network speed: Up to 25GB

Learn more >

✧ Server: QuantaGrid D52BQ-2U; 2-4 nodes
- CPU: Intel® Xeon® Scalable processors with Silver, Gold and Platinum options
- RAM: 256 GB~768GB
- HBA: SAS 9305-16i
- NIC: 1x Quanta OCP Mezz CX4, Dual Port 25G or 1x Q.logic 41212, Dual Port 25G
- Capacity: 4~24x SATA SSD 480GB/960GB/1.92TB

✧

✧ Server: QuantaGrid D52B-2U; 2-4 nodes
- CPU: Intel® Xeon® Scalable processors with Silver, Gold and Platinum options
- RAM: 256 GB~768GB
- NIC: 1x Quanta OCP Mezz CX4, Dual Port 25G or 1x Q.logic 41212, Dual Port 25G
- Capacity: 4~12x NVMe SSD 3.2TB

✧ Switch: 2x TOR QuantaMesh T4048-IX8D and 1x BMC QuantaMesh T1048-LY4R
✧ OS: Windows Server 2019 Datacenter Edition

2. Plan Hardware Deployment
✧ Hardware:

The four servers are interconnected using a Mellanox based 25GbE Ethernet RDMA cards and support DCB/PFC/ETS Ethernet switch.
The SSD drives were added to a single Storage Spaces Direct pool with multiple volumes based on the number of QCT S2D server nodes.

✧ Software:
Each server ran Windows Server 2019 Datacenter Edition and participated in a Windows Failover Cluster (required for Storage Spaces Direct).

Volumes were configured for operating system (OS), data and log drives as 2-way mirrored volumes, resulting in one local copy of data, and one remote copy on other nodes.

SQL Server workloads ran in Hyper-V virtual machines, with Windows Server 2019 as the guest OS, running SQL Server 2019 Standard Edition. Each VM was configured with 4 virtual cores (mapped to 2 physical cores) and 10GB of RAM.

The disk included as show:

Drive	Size(GB)	Purpose	Note
C:	100	Windows OS	SQL VM installed OS with Sysprep
K:	100	Data	SQL data file with format for 64K
L:	10	Log	SQL log file with format for 8K

✧ Network and Switch Connectivity:

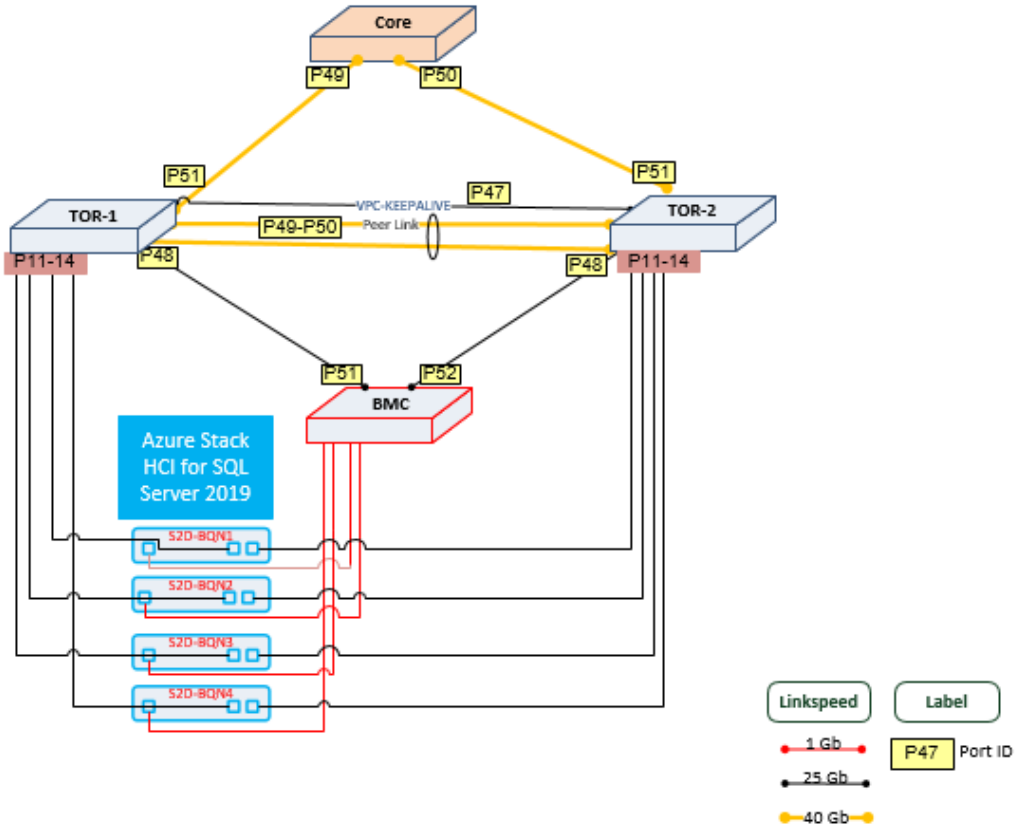

[Step by Step guide to deploy Azure Stack HCI]

1. Install Windows Server 2019 Datacenter (follow guidance above for network connectivity for Clustering)
2. Add Roles and Features
3. Setup Failover Clustering and enable a Cluster Witness
4. Setup Storage Spaces Direct

Install Windows Admin Center (WAC) to manage Windows Server and Windows Server VMs.

3. Set up Microsoft SQL Server on Azure Stack HCI
   Set up Windows Server VM, then
   install SQL Server on Windows

4. Monitoring and performance tuning
   To insure performance and health of your Microsoft SQL Server instances on Azure Stack HCI, it is important that appropriate monitoring and tuning is put in place. Additional SQL Server database engine tutorials are included here. For tuning SQL Server 2016/2017 for high performance, the following recommended practices are provided.

5. High Availability (HA)
   Azure Stack HCI leverages Windows Server Failover Clustering (WSFC) and can be utilized to support Microsoft SQL Server running in VMs (designed to help with hardware failure). Microsoft SQL Server also offers Always On availability groups

(AG) which provides database-level high availability and is designed to help with application and software faults. In addition to WSFC and AG, Azure Stack HCI can also leverage Always On Failover Cluster Instance (FCI) based on using Storage Spaces Direct technology for shared storage. All of these options can leverage the Microsoft Azure Cloud witness for quorum control. It is recommended that cluster AntiAffinity rules in WSFC be leveraged for the VMs to be placed on different physical nodes in order to maintain uptime for SQL Server in the event of host failures when you configure Always On availability groups.

6. Set up Azure hybrid scenarios
   Azure Backup supports backing up and restoring Microsoft SQL Server with application consistency. Install Azure Backup Server to start backup of your on-prem SQL data.

   Alternatively, you can also leverage Azure Blob Storage service for SQL Server to backup and restore to Azure Blob Storage service. This is suitable for off-site archiving. To manage the Azure Blob Storage backups, you can leverage the Managed SQL Backup feature included in Microsoft SQL Server.

   In addition to the backup scenario, you can set up other database services that Microsoft SQL Server (Microsoft SQL Server 2016/2017/2019) offers, connecting to Azure services such as (but not limited to) Azure Replica, Stretch Database, Azure Data Factory.

Summary

With completion of Microsoft SQL Server deployment using Azure Stack HCI, you now have a platform capable of running complex, highly available database workloads in VMs.