

Technical Use Cases for Azure Stack HCI



Branch office and edge




Virtual desktop infrastructure



High-performance SQL Server



Trusted enterprise virtualization



Scale-out storage

Microsoft Azure Stack HCI solutions provide industry-leading storage performance using leading-edge technologies in QCT servers. Technologies like NVMe, Intel Optane DC Memory, and RDMA networking combine to give you blazing fast performance that can scale from just 2 nodes all the way up to 16. Features like mirror-accelerated parity and data duplication ensure you get the most from your hardware investment. Connect it all to Azure File Sync, and you have a virtually bottomless storage solution. Below, you will find a how-to guide for Azure Stack HCI Scale-out storage configurations that includes:

- Plan and deploy QCT hardware and OS/tools to support Azure Stack HCI Scale-out storage
- Step by step documentation to enable Scale-out storage, including extending to the cloud with Azure File Sync

How to deploy Scale-out storage on Azure Stack HCI

1. Hardware and OS configuration for Scale-out storage configurations



QCT QuantaPlex T21P-4U

Scale:

- 2 to 16 nodes

Single Node Data:

- CPU: 12-36 cores (Intel)
- RAM: 256GB to 512GB
- Raw storage: 80TB to 640TB
- Storage type: SSD + HDD
- Network speed: Up to 25Gb

QCT QuantaGrid D52BQ-2U Balanced

Scale:

- 2 to 4 nodes

Single Node Data:

- CPU: 20-56 cores (Intel)
- RAM: 256GB to 768GB
- Raw storage: 32TB to 80TB
- Storage type: SSD + HDD
- Network speed: Up to 25Gb

[Learn more >](#)

QCT QuantaGrid D52T-1ULH

Scale:

- 2 to 4 nodes

Single Node Data:

- CPU: 20-56 cores (Intel)
- RAM: 256GB to 384GB
- Raw storage: 32TB to 120TB
- Storage type: SSD+HDD
- Network speed: Up to 25GB

[Learn more >](#)

❖ Server: [QuantaPlex T21P-4U](#): 2-4 nodes

- CPU: Intel® Xeon® Scalable processors with Silver, Gold and Platinum options
- RAM: 256 GB~512GB
- HBA: Avago SAS3 3008 Fury
- NIC: 2x Mellanox ConnectX-4, Dual Port 25G
- Cache: 8~16x SATA SSD 480GB/960GB/1.92TB
- Capacity: 8~64x SATA HDD 8TB/10TB

❖ 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
 - Cache: 2~4x SATA SSD 480GB/960GB/1.92TB
 - Capacity: 4~8x SATA HDD 8TB/10TB
- ✧ Server: [QuantaGrid D52T-1ULH](#): 2-4 nodes
- CPU: Intel® Xeon® Scalable processors with Silver, Gold and Platinum options
 - RAM: 256 GB~512GB
 - HBA: Avago SAS3 3008 Fury
 - NIC: 2x Mellanox ConnectX-4, Dual Port 25G
 - Cache: 2~4x SATA SSD 480GB/960GB/1.92TB
 - Capacity: 4~12x SATA HDD 8TB/10TB
- ✧ Switch: 2x TOR [QuantaMesh T4048-IX8D](#) and 1x BMC [QuantaMesh T1048-LY4R](#)
- ✧ OS: Windows Server 2019 Datacenter Edition

2. Plan Hardware Deployment

✧ Hardware:

The two to four nodes of servers were interconnected using a Mellanox based 25GbE Ethernet RDMA cards and support DCB/PFC/ETS Ethernet switch.

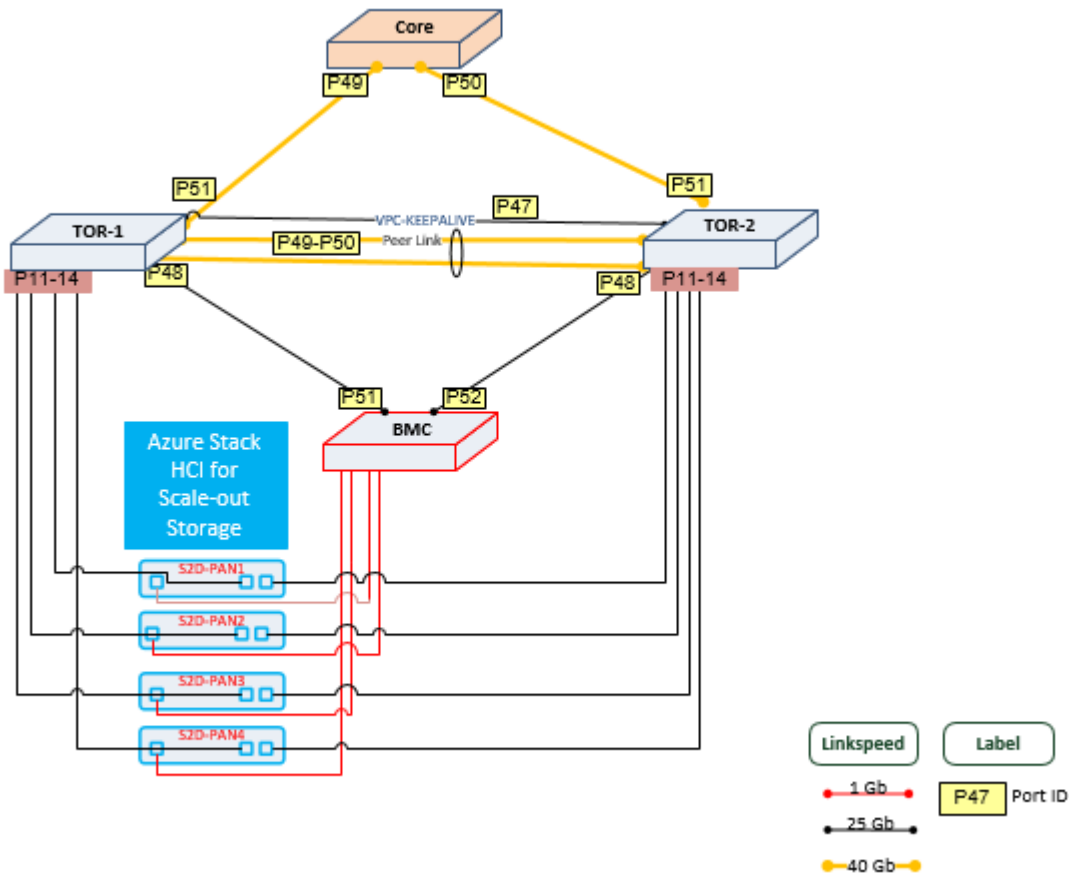
The SSD(cache tier) + HDD(capacity tier) drives were added to a single Storage Spaces Direct pool with multiple volumes based on the number of QCT S2D server nodes.

✧ 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.

✧ 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
5. [Install Windows Admin Center \(WAC\)](#)

From Windows Admin Center (WAC), set up Azure File Sync.

Additionally, you can Set Up additional Azure hybrid services such as Backup, File Sync, Site Recovery, Point-to-Site VPN, Update Management, and Security Center in WAC.



2. Enable general purpose file services on your Azure Stack HCI cluster

Once your Azure Stack HCI cluster setup has been deployed, follow the [steps in this article](#) to configure it for general purpose file services.

Windows Failover Clustering supports different kinds of file service roles. The primary focus of this document is the general purpose file server role, appropriate for serving files to information workers and other traditional file services roles. It enables key capabilities like quotas, data deduplication, accessible via multiple protocols like SMB and NFS.

The other is called the "Scale-Out File Server" role, which is designed for hosting virtual machines, high-performance SQL server databases, and other kinds of specialized, high-performance scenarios. The Scale-Out File Server role is not designed for the needs of typical user files, and is outside the scope of this document.

=====

3. Enable Azure File Sync to unify your on-premises file storage

Use Azure File Sync to centralize your organization's file shares in Azure Files, while keeping the flexibility, performance, and compatibility of an on-premises file server. You can use any protocol that's available on Windows Server to access your data locally, including SMB, NFS, and FTPS. You can have as many caches as you need across the world.

How do I deploy Azure File Sync on Azure Stack HCI?

Note: Updated guidance on how to install and configure Azure File Sync will [always be located here](#).

- a. If you do not already have an Azure account, get your [free account here](#)
- b. [Evaluate and prepare](#) your environment for Azure File Sync
- c. [Create an Azure File Share](#) in the Azure Portal
- d. [Deploy the Azure File Sync](#) locally on your Azure Stack HCI cluster

Summary

With completion of the Azure Stack HCI Scale-out storage deployment, you now have a scalable, flexible platform for all of your file serving needs. The addition of Azure File Sync will allow you to centralize and distribute your data quickly and easily among your various sites and users.