

AltosCloud HCI Solutions for Microsoft Azure Stack HCI

Hyper-converged Infrastructure (HCI) is the ongoing evolution of data centers architectures. The HCI solution provides pre-integrated system with virtualization software to deliver compute, storage, networking and management. Microsoft Azure Stack HCI on Altos BrainSphere™ server is a world-class, integrated HCI stack built on proven technologies that have already been deployed at scale, including Hyper-V, Storage Spaces Direct, and Azure-inspired software-defined networking (SDN). There are many reasons customers choose Azure Stack HCI, including:

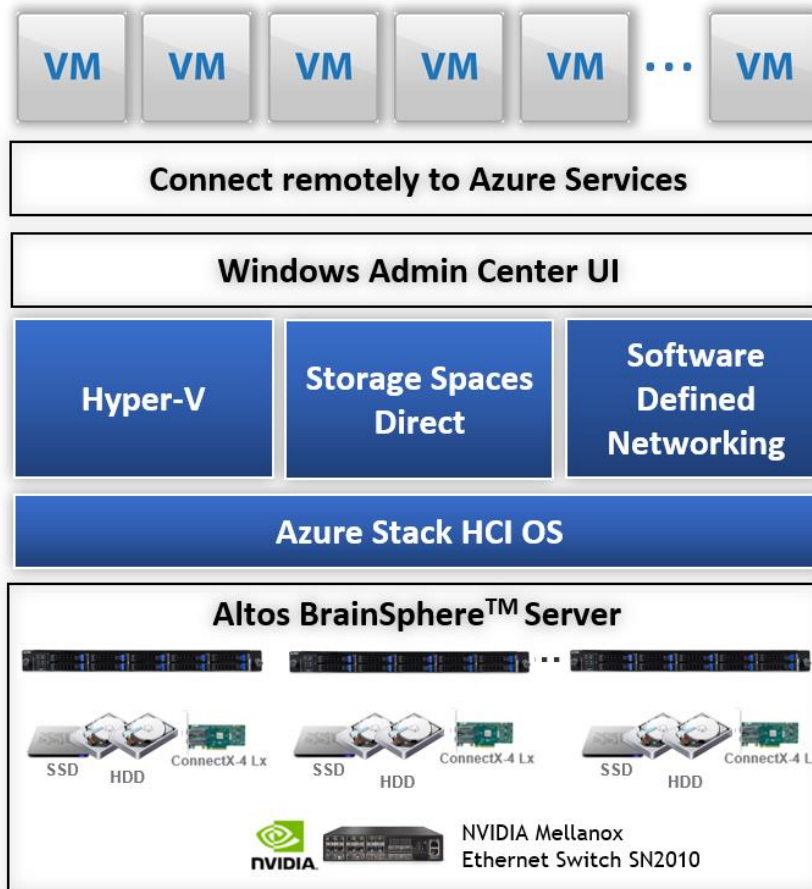
- Efficient, optimized infrastructure for on-premises IT
- Simplified operations by leveraging existing virtualization and storage concepts and skills.
- Easy management with existing data center processes and tools such as Microsoft Windows Admin Center, System Center, Active Directory, Group Policy, and PowerShell scripting.
- Joint support between Microsoft and Altos Computing Inc. improves the customer experience.

Microsoft Azure Stack HCI

Azure Stack HCI is a hyper-converged cluster solution that runs virtualized Windows and Linux workloads in a hybrid on-premises environment. Azure hybrid services enhance the cluster with enterprise-class capabilities such as cloud-based monitoring, site recovery, and VM backups, as well as a central view of all of your Azure Stack HCI deployments in the Azure portal. You can manage the cluster with your existing tools including Windows Admin Center, System Center, and PowerShell.

At its core, Azure Stack HCI on Altos BrainSphere™ server is a solution that combines the following:

- Microsoft certified hardware from Altos Computing Inc.
- Microsoft Azure Stack HCI OS
- Windows Admin Center
- Azure services
- End-to-end Ethernet networking by NVIDIA®



Hardware Overview

This section briefly describes the hardware used to validate the deployment examples in this document. A complete listing of hardware validated can be found at the end of this document.

AltosCloud HCI for Microsoft Azure Stack HCI relies on a dedicated Ethernet Storage Fabric (ESF) to communicate between nodes. The storage traffic communicates between nodes is through 25GbE network with RoCE (RDMA over Converged Ethernet) capable support. The front-end network is for OS management and users access.

NVIDIA Mellanox SN2010 Ethernet Switch

The migration to software-defined storage (SDS) and HCI is moving into the mainstream. Using the right networking to interconnect SDS and HCI clusters is critical to reap the promised benefits of fast storage and SDS/HCI in the most efficient way possible.

The NVIDIA Mellanox Spectrum® half-width, 1-Rack Unit (RU), SN2010 Top-of-Rack (TOR) switch is the best switch for SDS and HCI networks. It features 18-port 10/25GbE connectivity for storage/hyper-converged servers with 4-port 100GbE uplink connectivity to higher-speed networks.


Before NVIDIA introduced its unique half-width TOR switches, customers had to install a 48+4-port switch for TOR, with more than half of the switch ports unutilized.

To make the situation worse, as TOR switch redundancy is most enterprise deployments, the under-utilization is then exacerbated, not to mention the waste in rack space, power consumption, and cooling.

With the SN2010 switches, the user can put two units side-by-side in a 1RU rack space, achieving needed network connectivity and high availability (HA).

NVIDIA Ethernet switches are known for their superior performance for data flows in various applications, and fair bandwidth allocation. As both SDS and HCI run I/O intensive applications, high-performance TOR switches become more important. This guide uses SN2010s as TOR switches for in-band traffic.

NVIDIA SN2010 also simplifies network configuration for bringing up a HCI cluster, and makes networking management, monitoring and troubleshooting much more efficiency with single pane-of-class management and network telemetry software.



The NVIDIA Mellanox ConnectX®-4 Lx 10/25GbE network adapter cards are also used in Altos HCI servers for the best end-to-end connectivity and RoCE performance.

Azure Stack HCI Validated Node

Altos BrainSphere™ Server

AltosCloud HCI solutions for Microsoft Azure Stack HCI are pre-configured with certified components that simplify ordering and reduce deployment risks. Storage Spaces Direct (S2D) Ready Nodes are built on Altos BrainSphere™ servers and are available in Hybrid (mix of HDD and SSD) storage configurations.

Four Altos BrainSphere™ servers with Hybrid storage configurations are used in the examples in this document.




Best Practices

This section provides some best practices for configuring Azure Stack HCI. The procedures outlined in the examples are based on a 4-node Altos BrainSphere™ R369 F4 configuration. However, the same procedures, with some minor adjustments, can be applied to other configurations.


SSD/HDD Configuration

For each Altos BrainSphere™ R369 F4 server, install the SSD/HDD drives as shown in the following figure.



Storage and Network Adapters Configuration


For each Altos BrainSphere™ R369 F4 server, install the network adapters, and storage controller as shown in the following figure.



Network Design with Single Network Switch

Connect each server node to each switch as shown in the following figure.


A complete listing of required network switch and cables can be found at the end of this document.



Network Design with Dual Network Switches in High Availability

Connect each server node to each switch as shown in the following figure.

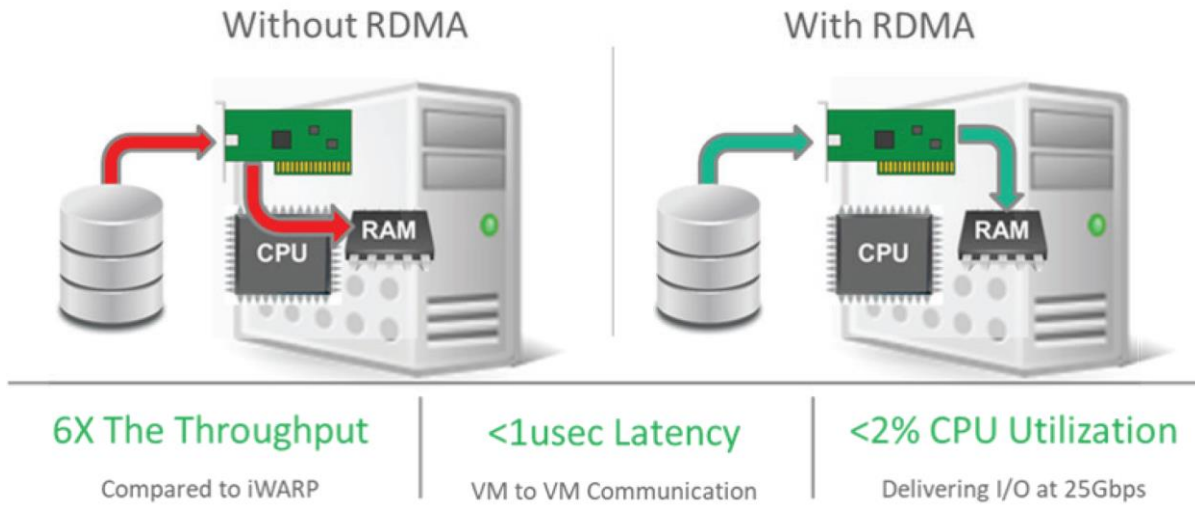
A complete listing of required network switch and cables can be found at the end of this document.



RoCE Increases Server Efficiency

AltosCloud HCI solutions use ConnectX-4 Lx 10/25GbE networks adapters that are RoCE capable. RoCE significantly increases throughput and lowers latency by performing direct memory transfers between servers.

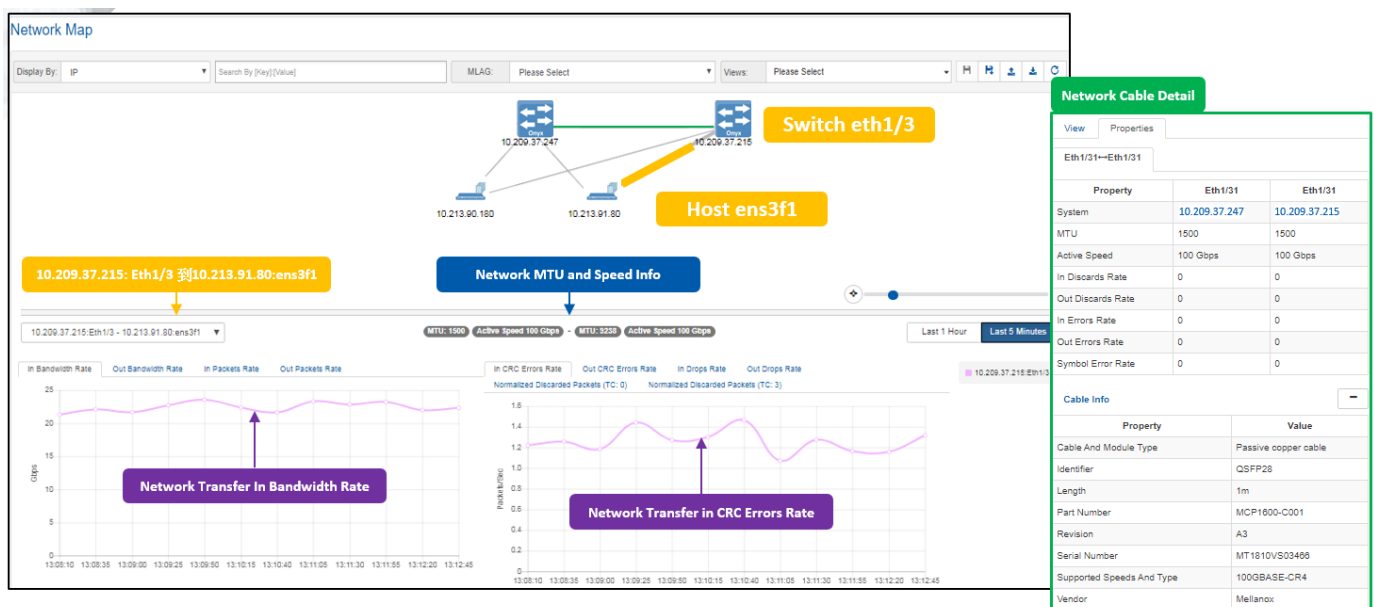
RoCE lowers latency and reduces CPU consumption in Windows S2D deployments. This is realized because RoCE uses hardware offloads on the adapter to bypass the CPU, eliminating its involvement in the data transfer process. The result is lower latencies and an improvement on CPU efficiencies. RoCE provides the absolutely best performance by offloading the CPU and allowing the maximum CPU power for application processes.



NVIDIA Mellanox NEO

NVIDIA Mellanox NEO® is a powerful network management platform for data-center network orchestration, designed to simplify network provisioning, monitoring and operations of the modern data center. NEO offers robust automation capabilities that extend existing tools, from network staging and bring-up, to day-to-day operations. NEO serves as a network API for NVIDIA Ethernet solutions.

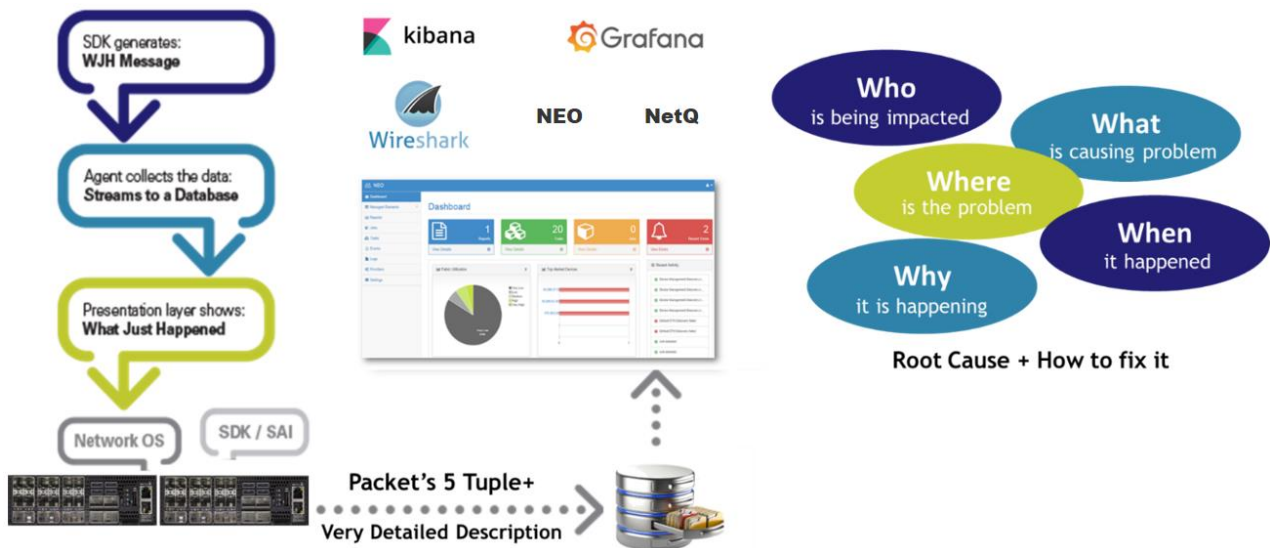
NEO simplifies fabric management, automates configuration of devices, provides deep visibility into traffic and network health, and enables early detection and auto-recovery of errors and failures. NEO leverages REST APIs to simplify integration with third-party orchestration and management platforms, and allows access to fabric-related data and provisioning operations.



NVIDIA Mellanox What Just Happened (WJH)

NVIDIA Mellanox What Just Happened™ (WJH) is an advanced streaming telemetry technology. With WJH, traffic inspection, filtering, and issue identification are done by the switch platform where the network context is readily available. As a result, only issue-relevant data is streamed out. When packets are dropped, for example, WJH provides packet headers along with a detailed reason as to why the packet was dropped, thus enabling more thorough analysis.

WJH also helps network operators by dramatically reducing mean time to innocence (MTTI) or issue resolution. Additionally, WJH provides insights to help improve resource utilization and capacity planning.



Conclusion

The integrated AltosCloud HCI and NVIDIA Ethernet networking solution provides a hyper-converged infrastructure for the modern datacenter that allows organizations to build the following:

- Efficient and fault-tolerant platform for databases and data analysis,
- Simplified operations with a cloud-like hybrid datacenter experience
- Reduced cost of IT through leveraging existing Microsoft licensing and Windows admin expertise.

AltosCloud HCI for Azure Stack HCI Configuration Options

Configuration	2 to 16 Nodes		
Product Name	BrainSphere™ R369 F4 HCI Hybrid		
OS	Microsoft Azure Stack HCI OS		
Solution Type	Hyper-converged Infrastructure, Hybrid		
Form Factor	Rack Mount Chassis, 1U 2.5" 10 Bays		
Trusted Platform Module (TPM)	2.0		
CPU	Dual Intel® Xeon® Scalable Processor Family		
Memory	256GB (up to 1.5TB per node)		
HBA	RS3UC080J		
NIC	Qty.	Model	RDMA type
	1	ConnectX®-4 Lx EN network interface card 25GbE dual-port SFP28 (MCX4121A)	RoCE
Drive	Type	Qty.	Form Factor
Cache	SSD	2	2.5"
Capacity	SAS	4 to 6	2.5"
OS	SSD	2 (RAID 1 protected)	2.5"

NVIDIA Mellanox Switches, Cables, Kit Options

Items	Vendor Part Number	QTY	Note
Switch Data Network	SN2010-CB2F	1	Ⓐ
Switch Management Network	AS4610	1	Ⓑ
NIC	MCX4121A-ACAT	4	Ⓒ
Cable (MLAG HA, 100G)	MCP1600-C002E30N	2	Ⓓ
Cable (Uplink, 10G)	MFM1T02A-SR	4	Ⓔ
Cable (Uplink, 25G)	MCP2M00-A002E30N	8	Ⓕ
Cable (Mgmt, 1G)	MC3208411-T	4	
Rack Kit	MTEF-KIT-D	1	Ⓖ

About Altos

Altos Computing Inc. (abbr. Altos) is established in 2017 and it is a subsidiary of Acer Inc. The business model of Altos is to provide the best streamlined and cost-effective integrated solutions thru in-house R&D working with ODM/IHV/ISV on servers, workstations, thin client, network and storage. In the era of demanding speed, Altos provides leading solution included, but not limited to High Performance Computing, Virtual Desktop Infrastructure, Cloud Infrastructure and Software Defined Storage, etc.

Altos provides solutions and services to government, academia, cloud service providers, datacenter operators and enterprises.

To find out more, visit our website: www.altoscomputing.com

In a continuing effort to improve the quality of our products, information in this brochure is subject to change without notice. Images appearing are only representations of some of the configurations available for this model. Availability may vary depending on region. Altos disclaims any liability for errors and omissions in product descriptions.