



Cisco UCS C240 M6 and M7 Solution for Microsoft Azure Stack HCI

Published: March 2023



Document Version History

Date	Change
March 30, 2023	Original publication
November 10, 2023	Added C240 M7SN as Azure Stack HCI Validated node
November 10, 2023	Added steps for Azure Stack HCI deployment in non-converged network topology
March 14, 2024	Added storage switchless configuration/deployment, C220 M7 Validated Node, Nvidia ConnectX-6 Lx, and updated BIOS diagrams,

About the Cisco Validated Design Program

The Cisco Validated Design (CVD) program consists of systems and solutions designed, tested, and documented to facilitate faster, more reliable, and more predictable customer deployments. For more information, go to: <http://www.cisco.com/go/designzone>.

Executive Summary

Cisco Validated Designs (CVDs) include systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of customers. Cisco UCS Solution for Microsoft Azure HCI offers highly available and scalable software-defined hyperconverged solution that is enable by the purpose-built Azure Stack HCI 22H2 Operating System. The Azure Stack HCI 22H2 Operating System is an Azure hybrid cloud designed hyperconverged solution that is based on Microsoft Windows Server 2022 and includes Storage Spaces Direct, Windows Failover Clustering, and Hyper-V.

Azure Stack is a family of three solutions that include Azure Stack HCI, Azure Stack Hub, and Azure Stack Edge. Azure Stack HCI is focused on the following use cases:

- Datacenter consolidation
- Virtual desktop Infrastructure
- Business critical infrastructure
- Storage cost reduction
- High availability and disaster recovery
- Enterprise application virtualization
- Azure Kubernetes Services
- Remote branch office system
- Arc enabled services

This document describes the architecture, topology, and deployment of Azure Stack HCI on Cisco UCS C240 M6SN, C240 M7SN and C220 M7N with Cisco Nexus 9300 series switches. Following the deployment guidance as specified in this document will result in a solution that adheres to both Cisco and Microsoft best practices.

Solution Overview

This chapter contains the following:

- [Introduction](#)
- [Audience](#)
- [Purpose of this Document](#)

Introduction

Software defined data center solutions enable IT organizations to optimize resource efficiency and improve service delivery. It combines compute virtualization, software defined storage, and virtualized networking that meets or exceeds high availability, performance, and security requirements of the most demanding deployments. The solution uses a shared-nothing architecture and takes advantage of the compute, storage, and network resources that are available within individual server. The servers are connected with external switching fabric that provides reliable high throughput and low latency.

Audience

The audience of this document includes, but is not limited to, sales engineers, field consultants, professional services, IT managers, partner engineers, and customers who want to take advantage of an infrastructure that is built to deliver IT efficiency and enable IT innovation.

Purpose of this Document

This overview and step-by-step deployment document is intended to describe in detail the procedure used to deploy the Azure Stack HCI solution on a Cisco UCS C240 M6SN, C240 M7SN and C220 M7N Rack Server with the Mellanox ConnectX-6 Dx/Lx NIC and connected to Cisco Nexus 9300 series switches. The procedure in this document should be used for deploying and evaluating this solution in a lab environment prior to deploying the solution in production. The deployment details described in this document need to be implemented as described unless stated otherwise.

This document will be periodically updated with new content. The contents will include procedures for deploying additional capabilities as well as qualified Cisco UCS firmware and drivers that must be used for deploying this solution.

Technology Overview

This chapter contains the following:

- [Cisco UCS C240 M6 Rack Server](#)
- [Cisco UCS C240 M7 Rack Server](#)
- [Cisco UCS C220 M7 Rack Server](#)
- [NVIDIA/Mellanox ConnectX-6 Dx Ethernet SmartNIC](#)
- [NVIDIA/Mellanox ConnectX-6 Lx Ethernet SmartNIC](#)
- [Cisco Integrated Management Controller \(IMC\)](#)
- [Cisco Intersight](#)
- [AzureStack HCI](#)

Cisco UCS C240 M6 Rack Server



The Cisco UCS C240 M6 Rack Server is a 2-socket, 2-Rack-Unit (2RU) rack server offering industry-leading performance and expandability. It supports a wide range of storage and I/O-intensive infrastructure workloads, from big data and analytics to collaboration. Cisco UCS C-Series Rack Servers can be deployed as standalone servers or as part of a Cisco Unified Computing System (Cisco UCS) managed environment to take advantage of Cisco's standards-based unified computing innovations that help reduce customers' Total Cost of Ownership (TCO) and increase their business agility.

In response to ever-increasing computing and data-intensive real-time workloads, the enterprise-class Cisco UCS C240 M6 Rack Server extends the capabilities of the Cisco UCS portfolio in a 2RU form factor. It incorporates the 3rd Generation Intel Xeon Scalable processors, supporting up to 40 cores per socket and 33 percent more memory versus the previous generation.

The Cisco UCS C240 M6 Rack Server brings many new innovations to the Cisco UCS rack server portfolio. With the introduction of PCIe Gen 4.0 expansion slots for high-speed I/O, DDR4 memory bus, and expanded storage capabilities, the server delivers significant performance and efficiency gains that will improve your application performance. Its features including the following:

- Supports third-generation Intel Xeon Scalable CPUs, with up to 40 cores per socket
- Up to 32 DDR4 DIMMs for improved performance including higher density DDR4 DIMMs (16 DIMMs per socket)
- 16x DDR4 DIMMs + 16x Intel Optane persistent memory modules for up to 12 TB of memory

- Up to 8 PCIe Gen 4.0 expansion slots plus a modular LAN-on-motherboard (mLOM) slot
- Support for Cisco UCS VIC 1400 Series adapters as well as third-party options
- Up to 28 hot-swappable Small-Form-Factor (SFF) SAS/SATA/NVMe:
 - 28 SFF SAS/SATA (with up to 8x NVMe)
 - 26 NVMe in all NVMe SKU (SN)
 - 14 NVMe in all NVMe SKU (N)
 - 16 LFF drives with options 4 rear SAS/SATA/NVMe) disk drives, or 16 Large-Form-Factor (LFF) 3.5-inch drives plus 2 rear hot-swappable SFF drives
 - Support for a 12-Gbps SAS modular RAID controller in a dedicated slot, leaving the remaining PCIe Gen 4.0 expansion slots available for other expansion cards
 - Option for 26 NVMe drives at PCIe Gen4 x4 (2:1 oversubscribed)
- M.2 boot options:
 - Up to 960 GB with optional hardware RAID
- Up to five GPUs supported
- Modular LAN-on-motherboard (mLOM) slot that can be used to install a Cisco UCS Virtual Interface Card (VIC) without consuming a PCIe slot, supporting quad port 10/40 Gbps or dual port 40/100 Gbps network connectivity
- Dual embedded Intel x550 10GBASE-T LAN-on-motherboard (LOM) ports
- Modular M.2 SATA SSDs for boot

Table 1. Item and Specification Details

Item	Specifications
Form factor	2RU rack server
Processors	3 rd Generation Intel Xeon Scalable processors (1 or 2)
Memory	32 DDR4 DIMM slots: 16, 32, 64, 128 and 256 GB and up to 3200 MHz Support for the Intel Optane DC Persistent Memory (128G, 256G, 512G)
PCIe expansion	8 PCIe 4.0 slots plus 1 dedicated 12-Gbps RAID controller slot and 1 dedicated mLOM slot
Storage controller	Internal controllers: Cisco 12-Gbps Modular SAS Host Bus Adapter (HBA)
Internal storage	Backplane options: <ul style="list-style-type: none"> • Up to 28 x 2.5-inch SAS and SATA HDDs and SSDs (up to 4 NVMe PCIe drives) • Up to 26 x 2.5-inch NVMe PCIe SSDs (All direct attach Gen4 x4) • Up to 16 x 3.5-inch SAS and SATA HDDs and SSDs, and optional 2 rear 2.5-inch HDDs and SSDs (up to 4 NVMe PCIe drives)

Item	Specifications
Embedded Network Interface Cards (NICs)	Dual 10GBASE-T Intel x550 Ethernet ports
mLOM	Dedicated mLOM slot that can flexibly accommodate 1-, 10-, 25-, 40-, and 100-Gbps adapters
Power supplies	Hot-pluggable, redundant 1050W AC, 1050W DC, 1600W AC and 2300W AC
Other storage	Dedicated Baseboard Management Controller (BMC) FlexMMC for utilities (on board) Dual M.2 SATA SSDs with HW Raid support
Management	Cisco Intersight Cisco Integrated Management Controller (IMC) Cisco UCS Manager
Rack options	Cisco ball-bearing rail kit with optional reversible cable management farm
Hardware and software interoperability	See the Cisco Hardware and Software Interoperability List for a complete listing of supported operating systems and peripheral options.

Cisco UCS C240 M7 Rack Server



The Cisco UCS C240 M7 Rack Server is well-suited for a wide range of storage and I/O-intensive applications such as big data analytics, databases, collaboration, virtualization, consolidation, and high-performance computing in its two-socket, 2RU form factor.

The Cisco UCS C240 M7 Rack Server extends the capabilities of the Cisco UCS Rack Server portfolio. It incorporates the 4th Gen Intel Xeon Scalable Processors with 50 percent more cores per socket advanced features like Intel Advanced Matrix Extensions (AMX), Data Streaming Accelerator (DSA), In-Memory Analytics Accelerator (IAA), and QuickAssist Technology (QAT), many applications will see significant performance improvements.

You can deploy the Cisco UCS C-Series Rack Servers as standalone servers or as part of the Cisco Unified Computing System managed by Cisco Intersight or Cisco UCS Manager to take advantage of Cisco standards-based unified computing innovations that can help reduce your Total Cost of Ownership (TCO) and increase your business agility.

The Cisco UCS C240 M7 Rack Server brings many new innovations to the Cisco UCS rack server portfolio. With the introduction of PCIe Gen 5.0 expansion slots for high-speed I/O, a DDR5 memory bus, and expanded storage capabilities, the server delivers significant performance and efficiency gains that will improve your application performance. Its features including the following:

- Supports up to two 4th Gen Intel Xeon Scalable CPU, with up to 60 cores per socket
- Up to 32 DDR5 DIMMs for up to 8 TB of capacity using 256 GB DIMMs (16 DIMMs per socket)
- 4800 MT/s DDR5 memory plus other speeds depending on the CPU installed
- Up to 8 PCIe 4.0 slots or up to 4 PCIe 5.0 slots, plus a hybrid modular LAN on motherboard (mLOM) / OCP 3.0 slot
- Support for Cisco UCS VIC 15000 Series adapters as well as third-party options
- Up to 28 hot-swappable Small-Form-Factor (SFF) SAS/SATA or NVMe drives (with up to 8 direct-attach NVMe drives):
 - New tri-mode RAID controller supports SAS4 plus NVMe hardware RAID
 - Option for 28 NVMe drives at PCIe Gen4 x2 each
- M.2 boot options:
 - Up to two 960GB SATA M.2 drives with hardware RAID, or
 - Up to two 960GB NVMe M.2 drives with NVMe hardware RAID
- Up to five GPUs supported
- Modular LOM / OCP 3.0:
 - One dedicated PCIe Gen4x16 slot that can be used to add an mLOM or OCP 3.0 card for additional rear-panel connectivity
 - mLOM slot that can be used to install a Cisco UCS Virtual Interface Card (VIC) without consuming a PCIe slot, supporting quad port 10/25/50 Gbps or dual port 40/100/200 Gbps network connectivity
 - OCP 3.0 slot features full out-of-band management for select adapters

Table 2. Item and Specification Details

Item	Specifications
Form factor	2RU rack server
Processors	4th Generation Intel Xeon Scalable processors (1 or 2)
Memory	32 DDR5 DIMM slots: 16, 32, 64, 128 and 256 GB and up to 4800 MT/s
PCIe expansion	Up to 8 PCIe 4.0 slots or up to 4 PCIe 5.0 slots plus 1 dedicated 24-Gbps RAID controller slot and 1 dedicated mLOM slot
RAID controllers	Internal controllers: Cisco 12-Gbps Modular SAS Host Bus Adapter (HBA)
Internal storage	Backplane options:

Item	Specifications
	<ul style="list-style-type: none"> Up to 28 x 2.5-inch SAS and SATA HDDs/ SSDs and NVMe SSDs (up to 8 NVMe direct attach) Up to 28 x 2.5-inch NVMe PCIe SSDs (all direct-attach Gen4 x2)
mLOM/OCP 3.0	<p>One dedicated socket (x16 PCIe lane) that can be used to add an mLOM or OCP 3.0 card for additional rear-panel connectivity.</p> <p>mLOM slot can flexibly accommodate 10/25/50 and 40/100/200 100-Gbps Cisco VIC adapters.</p> <p>OCP 3.0 slot features full out-of-band manageability that supports Intel X710 OCP Dual 10GBase-T via mLOM interposer.</p>
Power supplies	<p>Hot-pluggable, redundant platinum and titanium options:</p> <ul style="list-style-type: none"> Platinum: 1050W DC, and 1600W AC Titanium: 1200W AC, and 2300W AC
Other storage	<p>Dedicated Baseboard Management Controller (BMC) FlexMMC for utilities (on board)</p> <p>Dual M.2 SATA/NVMe SSDs with HW Raid support</p>
Management	<p>Cisco Intersight</p> <p>Cisco Integrated Management Controller (IMC)</p> <p>Cisco UCS Manager</p>
Rack options	<p>Cisco ball-bearing rail kit with optional reversible cable management farm</p>
Hardware and software interoperability	<p>See the Cisco Hardware and Software Interoperability List for a complete listing of supported operating systems and peripheral options.</p>

Cisco UCS C220 M7 Rack Server



The Cisco UCS C220 M7 Rack Server is a versatile general-purpose infrastructure and application server. This high-density, 1RU, 2-socket rack server delivers industry-leading performance and efficiency for a wide range of workloads, including virtualization, collaboration, and bare-metal applications.

The Cisco UCS C220 M7 Rack Server extends the capabilities of the Cisco UCS rack server portfolio. It incorporates the 4th and 5th Gen Intel Xeon Scalable Processors with 50 percent more cores per socket, advanced features such as Intel Advanced Matrix Extensions (AMX), Data Streaming Accelerator (DSA), In-Memory Analytics

Accelerator (IAA), and QuickAssist Technology (QAT). Many applications will see significant performance improvements.

You can deploy the Cisco UCS C-Series rack servers as standalone servers or as part of the Cisco Unified Computing System managed by Cisco Intersight or Cisco UCS Manager to take advantage of Cisco standards-based unified computing innovations that can help reduce your Total Cost of Ownership (TCO) and increase your business agility.

The Cisco UCS C220 M7 rack server brings many new innovations to the Cisco UCS rack server portfolio. With the introduction of PCIe Gen 5.0 for high-speed I/O, a DDR5 memory bus, and expanded storage capabilities, the server delivers significant performance and efficiency gains that will improve your application performance.

- Supports up to two 4th Gen Intel Xeon Scalable CPUs, with up to 60 cores per socket.
- Up to 32 DDR5 DIMMs for up to 4 TB of capacity using 128 GB DIMMs (16 DIMMs per socket).
- Up to 5600 MT/s DDR5 memory plus other speeds depending on the CPU installed.
- Up to 3 PCIe 4.0 slots or up to 2 PCIe 5.0 slots, plus a modular LAN on motherboard (mLOM) slot.
- Support for Cisco UCS VIC 15000 Series adapters as well as third-party options.
- Up to 10 SAS/SATA or NVMe disk drives:
 - New tri-mode RAID controller supports SAS4 RAID or NVMe hardware RAID with optional up to four direct-attach NVMe drives.
 - Option for 10 direct-attach NVMe drives at PCIe Gen4x4 each.
- M.2 boot options:
 - Up to two 960GB SATA M.2 drives with hardware RAID.
 - Or
 - Up to two 960GB NVMe M.2 drives with NVMe hardware RAID.
- Up to three GPUs supported.
- Hybrid modular LOM/OCP 3.0:
 - One dedicated Gen 4.0 x16 slot that can be used to add an mLOM or OCP 3.0 card for additional rear-panel connectivity.
 - mLOM allows for Cisco UCS Virtual Interface Cards (VICs) without consuming a PCIe slot, supporting quad port 10/25/50 Gbps or dual port 40/100/200 Gbps network connectivity.
 - OCP 3.0 slot features full out-of-band management for select adapters.

Table 3. Item and Specification Details

Item	Specifications
Form factor	1RU rack server
Processors	Up to 2x 5th Gen Intel Xeon Scalable processors (1 or 2) or

Item	Specifications
	Up to 2x 4th Gen Intel Xeon Scalable processors (1 or 2)
Memory	32 DDR5-5600 DIMM slots (16 DIMMS per CPU): 16, 32, 48, 64, 96, 128GB at up to 5600 MT/s for up to 4TB of memory with 5th Gen Intel Xeon Scalable processors or 32 DDR5-4800 DIMM slots (16 DIMMS per CPU): 16, 32, 64, 128GB at up to 4800 MT/s for up to 4TB of memory with 4th Gen Intel Xeon Scalable processors
PCIe expansion	Up to 3 PCIe 4.0 slots or up to 2 PCIe 5.0 slots plus 1 dedicated 24-Gbps RAID controller slot and 1 dedicated mLOM/OCP 3.0 slot
RAID controllers	Internal controllers: <ul style="list-style-type: none"> • Cisco 24-Gbps modular tri-mode controller supports SAS 4 or NVMe hard-ware RAID • Cisco 12-Gbps modular RAID controller (PCIe 4.0) with 4-GB Flash-Backed Write Cache (FBWC) or <ul style="list-style-type: none"> • Cisco 12-Gbps modular SAS Host Bus Adapter (HBA) • External controller: Cisco 12-Gbps 9500-8e SAS HBA
Internal storage	Backplane options: <ul style="list-style-type: none"> • Up to 10 x 2.5-inch SAS and SATA HDDs, SSD, NVMe drives, with the option of up to 4 direct-attach NVMe drives • Up to 10 x 2.5-inch NVMe PCIe SSDs (all direct-attach PCIe Gen4x4)
mLOM/OCP 3.0	<ul style="list-style-type: none"> • One dedicated PCIe Gen4x16 slot that can be used to add an mLOM or OCP 3.0 card for additional rear-panel connectivity • mLOM slot can flexibly accommodate 10/25/50/100/25/40, and 40/100/200 100-Gbps Cisco VIC adapters • OCP 3.0 slot features full out-of-band manageability that supports Intel X710 OCP Dual 10GBase-T via mLOM interposer
Power supplies	Hot-pluggable, redundant platinum and titanium options: <ul style="list-style-type: none"> • Platinum: 770W AC, 1050W DC, and 1600W AC • Titanium: 1200W AC, and 2300W AC
Other storage	Dedicated Baseboard Management Controller (BMC) FlexMMC for utilities (on board) Dual M.2 SATA/NVMe SSDs with HW RAID support
Management	Cisco Intersight Cisco Integrated Management Controller (IMC) Cisco UCS Manager
Rack options	Cisco ball-bearing rail kit with optional reversible cable management arm
Hardware and software	See the Cisco Hardware and Software Interoperability List for a complete listing of supported

Item	Specifications
interoperability	operating systems and peripheral options.

NVIDIA/Mellanox ConnectX-6 Dx Ethernet SmartNIC


NVIDIA ConnectX-6 Dx is a highly secure and advanced smart network interface card (SmartNIC) that accelerates mission-critical cloud and data center applications, including security, virtualization, SDN/NFV, big data, machine learning, and storage. ConnectX-6 Dx provides up to two ports of 100Gb/s or a single port of 200Gb/s Ethernet connectivity and is powered by 50Gb/s (PAM4) or 25/10 Gb/s (NRZ) SerDes technology.

ConnectX-6 Dx features virtual switch (vSwitch) and virtual router (vRouter) hardware accelerations delivering orders-of-magnitude higher performance than software-based solutions. ConnectX-6 Dx supports a choice of single-root I/O virtualization (SR-IOV) and VirtIO in hardware, enabling customers to best address their application needs. By offloading cloud networking workloads, ConnectX-6 Dx frees up CPU cores for business applications while reducing total cost-of-ownership.

In an era where data privacy is key, ConnectX-6 Dx provides built-in inline encryption/decryption, stateful packet filtering, and other capabilities, bringing advanced security down to every node with unprecedented performance and scalability.

Built on the solid foundation of NVIDIA's ConnectX line of SmartNICs, ConnectX-6 Dx offers best-in-class RDMA over Converged Ethernet (RoCE) capabilities, enabling scalable, resilient, and easy-to-deploy RoCE solutions. For data storage, ConnectX-6 Dx optimizes a suite of storage accelerations, bringing NVMe-oF target and initiator offloads.

Figure 1. NVIDIA/Mellanox ConnectX-6 Dx



NVIDIA/Mellanox ConnectX-6 Lx Ethernet SmartNIC

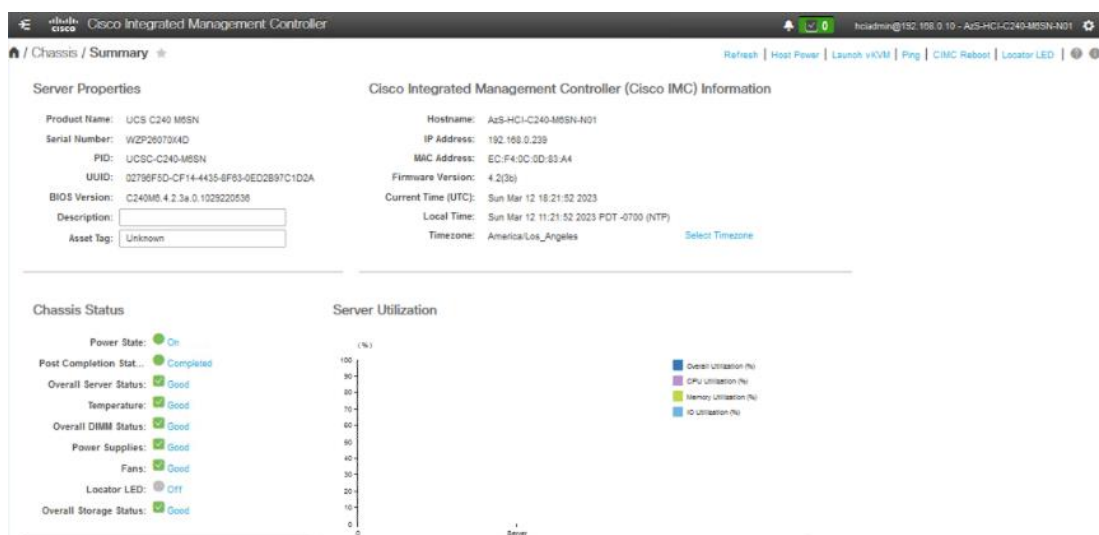
NVIDIA ConnectX-6 Lx Ethernet network interface cards (NIC) deliver high-performance network connectivity at 25GbE speeds coupled with advanced security and the best total cost of ownership for enterprise data centers. The NICs support up to two ports of 25GbE or a single port of 50GbE networking, along with PCI Express (PCIe) Gen3 and Gen4 x8 host connectivity to meet the demands of modern workloads in the cloud, in the data center, and at the edge.

Figure 2. NVIDIA/Mellanox ConnectX-6 Lx



Cisco Integrated Management Controller (IMC)

The Cisco Integrated Management Controller (IMC) is a baseboard management controller that provides embedded server management for Cisco UCS C-Series Rack Servers and Cisco UCS S-Series Storage Servers. The Cisco IMC enables system management in the data center and across distributed branch-office locations. It supports multiple management interfaces, including a Web User Interface (Web UI), a Command-Line Interface (CLI), and an XML API that is consistent with the one used by Cisco UCS Manager. IMC also supports industry-standard management protocols, including Redfish, Simple Network Management Protocol Version 3 (SNMPv3), and Intelligent Platform Management Interface Version 2.0 (IPMIv2.0). The figure below shows a sample Cisco IMC screen.



Cisco Intersight


Cisco Intersight Overview

Cisco Intersight is Cisco's systems management platform that delivers intuitive computing through cloud-powered intelligence. This platform offers a more intelligent level of management that enables IT organizations to analyze, simplify, and automate their environments in ways that were not possible with prior generations of tools. This capability empowers organizations to achieve significant savings in Total Cost of Ownership (TCO) and to deliver applications faster in support of new business initiatives. The advantages of the model-based management of the Cisco UCS platform plus Cisco Intersight are extended to Cisco UCS servers and Cisco HyperFlex, including Cisco HyperFlex Edge systems. Cisco HyperFlex Edge is optimized for remote sites, branch offices, and edge environments.

Endpoints supported by Cisco Intersight use model-based management to provision servers and associated storage and fabric automatically, regardless of form factor. Cisco Intersight works in conjunction with Cisco UCS Manager and the Cisco Integrated Management Controller (IMC). By simply associating a model-based configuration with a resource through server profiles, your IT staff can consistently align policy, server personality, and workloads. These policies can be created once and used by IT staff with minimal effort to deploy servers. The result is improved productivity and compliance and lower risk of failures due to inconsistent configuration.

Cisco Intersight will be integrated with data-center and hybrid-cloud platforms and services to securely deploy and manage infrastructure resources across data-center and edge environments. Additionally, Cisco provides integrations to third-party operations tools, starting with ServiceNow allowing you to use your existing solutions more effectively.

Cisco Intersight offers flexible deployment either as Software as a Service (SaaS) on Intersight.com or running on your premises with the Cisco Intersight virtual appliance. The virtual appliance provides users with the benefits of Cisco Intersight while allowing more flexibility for those with additional data locality and security requirements.



Cisco Intersight Features and Benefits

[Table 4](#) lists the main features and benefits of Cisco Intersight.

Table 4. Cisco Intersight Features and Benefits

Feature	Benefit
Unified management	<p>Simplify Cisco UCS, Cisco HyperFlex, Pure Storage, and Cisco Network Insights management from a single management platform.</p> <p>Increase scale across data centers and remote locations without additional complexity.</p> <p>Use a single dashboard to monitor Cisco UCS and Cisco HyperFlex systems.</p> <p>Cisco UCS Manager, Cisco IMC software, Cisco HyperFlex Connect, and Cisco UCS Director tunneling allow access to element managers that do not have local network access.</p>
Configuration, provisioning, and server profiles	<p>Treat Cisco UCS servers and storage as infrastructure resources that can be allocated and reallocated among application workloads for more dynamic and efficient use of server capacity.</p> <p>Create multiple server profiles with just a few clicks or through the available API, automating the provisioning process.</p> <p>Clone profiles to quickly provision Cisco UCS C-Series Rack Servers in standalone mode.</p> <p>Create, deploy, and manage your Cisco HyperFlex configurations.</p> <p>Help ensure consistency and eliminate configuration drift, maintaining standardization across many systems.</p>
Inventory information and status	<p>Display and report inventory information for Cisco UCS and Cisco HyperFlex systems.</p> <p>Use global search to rapidly identify systems based on names, identifiers, and other information.</p> <p>Use tagging to associate custom attributes with systems.</p> <p>Monitor Cisco UCS and Cisco HyperFlex server alerts and health status across data</p>


Feature	Benefit
	<p>centers and remote locations.</p> <p>View your Cisco HyperFlex configurations.</p> <p>Track and manage firmware versions across all connected Cisco UCS and Cisco HyperFlex systems.</p> <p>Track and manage software versions and automated patch updates for all claimed Cisco UCS Director software installations.</p>
Enhanced support experience	<p>Get centralized alerts about failure notifications.</p> <p>Automate the generation, forwarding, and analysis of technical support files to the Cisco Technical Assistance Center (TAC) to accelerate the troubleshooting process.</p>
Open API	<p>A RESTful API that supports the OpenAPI Specification (OAS) to provide full programmability and deep integrations systems.</p> <p>The Python and PowerShell SDKs will enable integrations with Ansible, Chef, Puppet, and other DevOps and IT Operations Management (ITOM) tools.</p> <p>ServiceNow integration to provide inventory and alerts to the IT Service Management platform.</p>
Seamless integration and upgrades	<p>Upgrades are available for Cisco UCS, Cisco HyperFlex systems, and Cisco UCS Director software running supported firmware and software versions.</p> <p>Upgrades to Cisco Intersight are delivered automatically without requiring the resources of traditional management tool upgrades and disruption to your operations.</p>

Azure Stack HCI

Azure Stack HCI 22H2 is a hyper-converged Windows Server 2022 cluster that uses validated hardware to run virtualized workloads on-premises. Optionally, you can also connect to Azure services for cloud-based backup, site-recovery, and more. Azure Stack HCI solutions use Microsoft-validated hardware to ensure optimal performance and reliability, and include support for technologies such as NVMe drives, persistent memory, and remote-direct memory access (RDMA) networking.

Azure Stack HCI is a solution that combines several products:

- Hardware from an OEM partner
- Azure Stack HCI OS 22H2
- Windows Admin Center
- Azure services (optional)



Azure Stack HCI is Microsoft’s hyperconverged solution available from a wide range of hardware partners. Consider the following scenarios for a hyperconverged solution to help you determine if Azure Stack HCI is the solution that best suits your needs:

- Refresh aging hardware. Replace older servers and storage infrastructure and run Windows and Linux virtual machines on-premises and at the edge with existing IT skills and tools.
- Consolidate virtualized workloads. Consolidate legacy apps on an efficient, hyperconverged infrastructure. Tap into the same types of cloud efficiencies used to run hyper-scale datacenters such as Microsoft Azure.
- Connect to Azure for hybrid cloud services. Streamline access to cloud management and security services in Azure, including offsite backup, site recovery, cloud-based monitoring, and more.

Hyperconverged Efficiencies

Azure Stack HCI solutions bring together highly virtualized compute, storage, and networking on industry-standard x86 servers and components. Combining resources in the same cluster makes it easier for you to deploy, manage, and scale. Manage with your choice of command-line automation or Windows Admin Center.

Achieve industry-leading virtual machine performance for your server applications with Hyper-V, the foundational hypervisor technology of the Microsoft cloud, and Storage Spaces Direct technology with built-in support for NVMe, persistent memory, and remote-direct memory access (RDMA) networking.

It helps keep apps and data secure with shielded virtual machines, network micro segmentation, and native encryption.

Hybrid Cloud Capabilities

You can take advantage of cloud and on-premises working together with a hyperconverged infrastructure platform in public cloud. Your team can start building cloud skills with built-in integration to Azure infrastructure management services:

- Azure Site Recovery for high availability and disaster recovery as a service (DRaaS).
- Azure Monitor, a centralized hub to track what's happening across your applications, network, and infrastructure - with advanced analytics powered by AI.
- Cloud Witness, to use Azure as the lightweight tie breaker for cluster quorum.
- Azure Backup for offsite data protection and to protect against ransomware.
- Azure Update Management for update assessment and update deployments for Windows VMs running in Azure and on-premises.
- Azure Network Adapter to connect resources on-premises with your VMs in Azure via a point-to-site VPN.
- Sync your file server with the cloud, using Azure File Sync.

Management Tools

Azure Stack HCI uses the same virtualization and software-defined storage and networking software as Azure Stack Hub. However, with Azure Stack HCI you have full admin rights on the cluster and can manage any of its technologies directly:

- [Hyper-V](#)
- [Storage Spaces Direct](#)
- [Failover Clustering](#)

To manage these technologies, you can use the following management tools:

- [PowerShell](#)
- [Windows Admin Center](#) (optional)
- [System Center](#) (Optional)
- Other management tools such as [Server Manager](#), and MMC snap-ins (Optional)
- Non-Microsoft tools such as 5Nine Manager (Optional)

If you choose to use System Center to deploy and manage your infrastructure, you'll use System Center Virtual Machine Management (VMM) and System Center Operations Manager. With VMM, you provision and manage the resources needed to create and deploy virtual machines and services to private clouds.

Hyper-V


Hyper-V is Microsoft's hardware virtualization product. It lets you create and run a software version of a computer, called a *virtual machine*. Each virtual machine acts like a complete computer, running an operating system and programs. When you need computing resources, virtual machines give you more flexibility, help save time and money, and are a more efficient way to use hardware than just running one operating system on physical hardware.

Hyper-V runs each virtual machine in its own isolated space, which means you can run more than one virtual machine on the same hardware at the same time. You might want to do this to avoid problems such as a crash affecting the other workloads, or to give different people, groups, or services access to different systems.

Storage Spaces Direct


Storage Spaces Direct uses industry-standard servers with local-attached drives to create highly available, highly scalable software-defined storage at a fraction of the cost of traditional SAN or NAS arrays. The hyper-converged architecture radically simplifies procurement and deployment, while features such as caching, storage tiers, and erasure coding, together with the latest hardware innovations such as RDMA networking and NVMe drives, deliver unrivaled efficiency and performance.

One cluster for compute and storage. The hyper-converged deployment option runs Hyper-V virtual machines directly on the servers providing the storage, storing their files on the local volumes. This eliminates the need to configure file server access and permissions and reduces hardware costs for small-to-medium business or remote office/branch office deployments.



Storage Spaces Direct is the evolution of Storage Spaces, first introduced in Windows Server 2012. It leverages many of the features you know today in Windows Server, such as Failover Clustering, the Cluster Shared Volume (CSV) file system, Server Message Block (SMB) 3, and of course Storage Spaces. It also introduces new technology, most notably the Software Storage Bus.

Figure 3. Overview of the Storage Spaces Direct Stack



Networking Hardware. Storage Spaces Direct uses SMB3, including SMB Direct and SMB Multichannel, over Ethernet to communicate between servers. Microsoft strongly recommends 10+ GbE with remote-direct memory access (RDMA).

Storage Hardware. From 1 to 16 servers with local-attached SATA, SAS, or NVMe drives. Each server must have at least 2 solid-state drives, and at least 4 additional drives. The SATA and SAS devices should be behind a host-bus adapter (HBA) and SAS expander. We strongly recommend the meticulously engineered and extensively validated platforms from our partners (coming soon).

Failover Clustering. The built-in clustering feature of Windows Server is used to connect the servers.

Software Storage Bus. The Software Storage Bus is new in Storage Spaces Direct. It spans the cluster and establishes a software-defined storage fabric whereby all the servers can see all of each other's local drives. You can think of it as replacing costly and restrictive Fibre Channel or Shared SAS cabling.

Storage Bus Layer Cache. The Software Storage Bus dynamically binds the fastest drives present (e.g. SSD) to slower drives (e.g. HDDs) to provide server-side read/write caching that accelerates IO and boosts throughput.

Storage Pool. The collection of drives that form the basis of Storage Spaces is called the storage pool. It is automatically created, and all eligible drives are automatically discovered and added to it.

Note: We strongly recommend you use one pool per cluster, with the default settings. Read our [Deep Dive into the Storage Pool](#) to learn more.

Storage Spaces. Storage Spaces provides fault tolerance to virtual "disks" using [mirroring, erasure coding, or both](#). You can think of it as distributed, software-defined RAID using the drives in the pool. In Storage Spaces Direct, these virtual disks typically have resiliency to two simultaneous drive or server failures (for example, 3-way mirroring, with each data copy in a different server) though chassis and rack fault tolerance is also available.

Resilient File System (ReFS). ReFS is the premier filesystem purpose-built for virtualization. It includes dramatic accelerations for .vhdx file operations such as creation, expansion, and checkpoint merging, and built-in checksums to detect and correct bit errors. It also introduces real-time tiers that rotate data between so-called "hot" and "cold" storage tiers in real-time based on usage.

Cluster Shared Volumes. The CSV file system unifies all the ReFS volumes into a single namespace accessible through any server, so that to each server, every volume looks and acts like it's mounted locally.

Failover Clustering

A failover cluster is a group of independent computers that work together to increase the availability and scalability of clustered roles (formerly called clustered applications and services). The clustered servers (called nodes) are connected by physical cables and by software. If one or more of the cluster nodes fail, other nodes begin to provide service (a process known as failover). In addition, the clustered roles are proactively monitored to verify that they are working properly. If they are not working, they are restarted or moved to another node.

Failover clusters also provide Cluster Shared Volume (CSV) functionality that provides a consistent, distributed namespace that clustered roles can use to access shared storage from all nodes. With the Failover Clustering feature, users experience a minimum of disruptions in service.

Failover Clustering has many practical applications, including:

- Highly available or continuously available file share storage for applications such as Microsoft SQL Server and Hyper-V virtual machines

-
- Highly available clustered roles that run on physical servers or on virtual machines that are installed on servers running Hyper-V

Solution Design

This chapter contains the following:

- [Architecture](#)
- [Physical Topology](#)
- [Azure Stack HCI Components](#)
- [Logical Topology](#)

Architecture

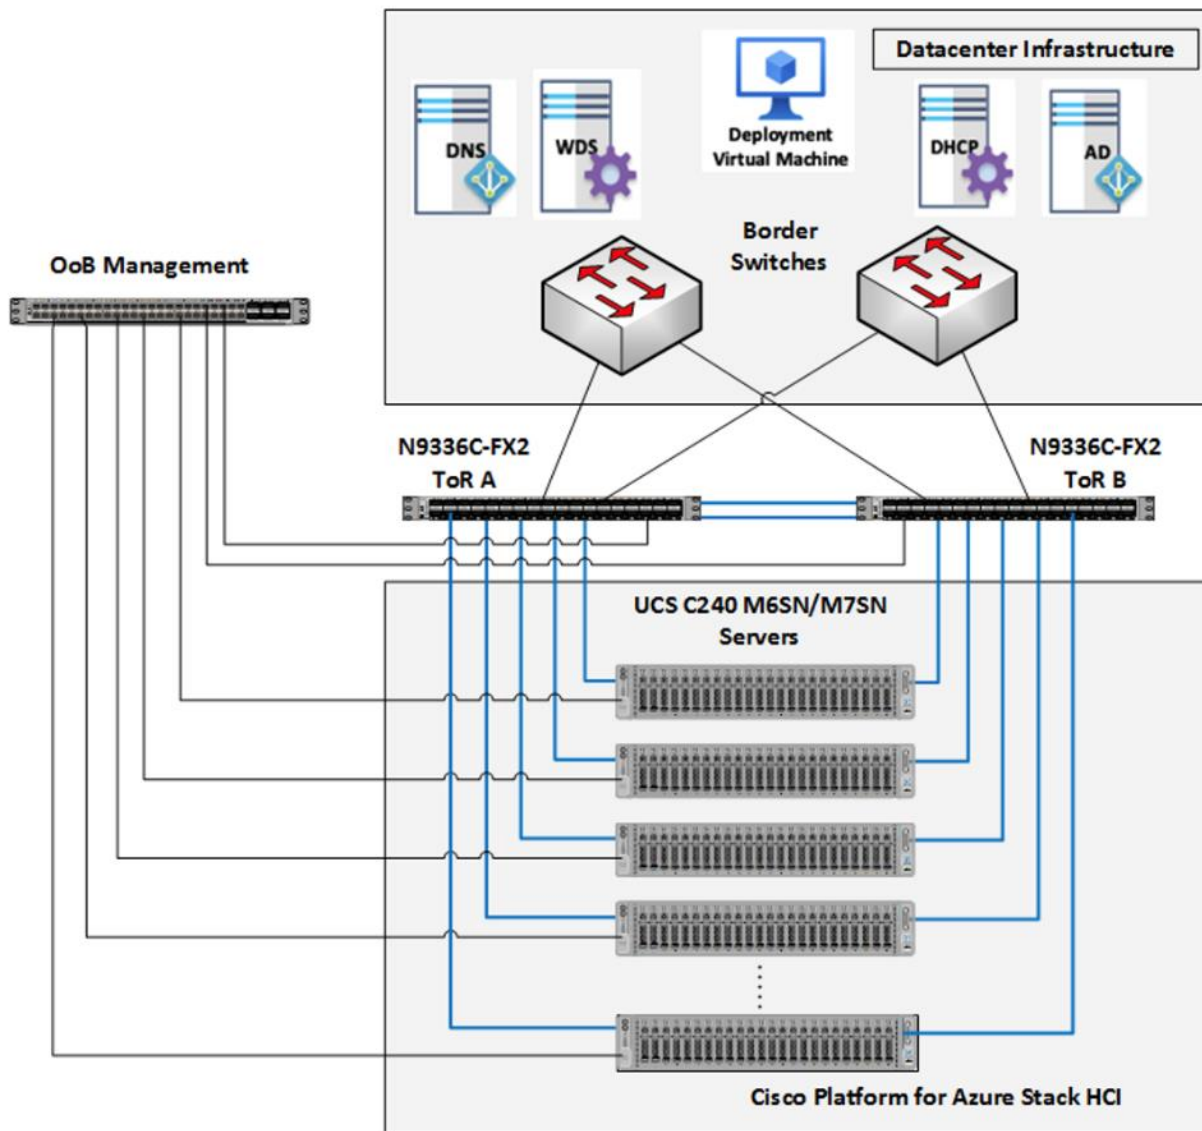
The Cisco solution for Azure Stack HCI architecture must be implemented as described in this document. Cisco provides a specific PID for ordering the configuration. The PID includes all of the required components that comprise the solution. The Azure Stack HCI cluster can be scaled from 1 to 16 servers. The architecture for the deployment of Azure Stack HCI solution consists of a storage switched configuration using two TOR switches with either fully converged or non-converged host network adapters.

The architecture has a data fabric and a management fabric. The servers connect to the data fabric using dual 100Gb connections. This data fabric is provided by the Cisco 9300 series switches which provide layer 2 connectivity and carries all the Azure Stack HCI network traffic (management, compute, and RDMA storage traffic). Server management is facilitated through an Out-of-band (OOB) management network that connects the server's dedicated management port to an OOB management switch with 1GbE links. The servers Azure Stack HCI OS 22H2 provides a rich set of software defined services that are core to this solution.

Physical Topology

The data center is expected to have infrastructure services such as DNS and Active Directory. WDS (Windows Deployment Service) and DHCP are also recommended to expedite deployments. These services must be accessible through the ToR (Top of Rack) or EoR (End of Row) network switches that connect the Cisco UCS C240 M6 and M7 Servers that are part of the Cisco solution for Azure Stack HCI to the datacenter infrastructure.

Figure 4. Physical Topology



Azure Stack HCI Components

Cisco UCS C240 M6SN, C240 M7SN, or C220 M7N Servers

The Cisco UCS C240 M6SN or C240 M7SN, or C220 M7N server configuration consists of a one or two dual-port 100GbE NVIDIA ConnectX-6 DX network interface cards, teamed with each port connecting to two different ToR switches and a single 1GbE dedicated management port which connects to an OOB management switch for communication with the Cisco Integrated Management Controller in each server.

The ToR switches, in this case Cisco Nexus 9300 Series Switches, carry both Azure Stack HCI cluster traffic and management network traffic to the Cisco UCS C240 M6SN or M7SN servers. The Azure Stack HCI cluster traffic flows through 100GbE links to the NVIDIA ConnectX-6 DX network interface card/s in each server. Out-of-band

management traffic is facilitated by a 1GbE connection to each of the Cisco UCS C240 M6SN, C240 M7SN, or C220 M7N servers.

ToR Switch

The ToR (Top of Rack) switches can be any Cisco Nexus switches that have confirmed support for the Azure Stack HCI requirements. The list of supported Cisco Nexus series switches and the NX-OS version can be viewed [here](#). The ToR switch provides layer 2 and layer 3 connectivity to the Azure Stack HCI cluster nodes. The ToR switches should include a security focused configuration that is standardized within the datacenter network. Two ToR switches in Virtual Port Channel (VPC) configuration provide high availability and redundancy for the network traffics.

The [Appendix](#) of this document has sample configurations that can be implemented in the ToR switch. These sample configurations include vPC, SVI, HSRP, and DHCP Relay.

Out-of-Band Management Switch

It is expected that the datacenter has a secure OoB (Out-of-Band) management network that is used to managed network devices in the datacenter. Cisco UCS C240 M6SN and M7SN servers and the ToR switches are directly connected to the out-of-band management switches and a disjoint layer-2 configuration is used to keep the management network path separate from the data network path. The OoB network needs to have internet access in order for Cisco Intersight to be able to access the UCS C240 M6/M7 servers.

Logical Topology

The logical topology is comprised of the following:

- Tenant/Compute Network

The Tenant network is a VLAN trunk that carries one or more VLANs that provide access to the tenant virtual machines. Those VLANs are configured on the ToR switch's port in trunk mode. To connect VMs to these VLANs, the corresponding VLAN tags are defined on the virtual network adapter. Each tenant VLAN is expected to have an IP subnet assigned to it.

- Management Network

The management network is a VLAN that carries network traffic to the parent partition. This network is used to access the host operating system. The connectivity to the management network is provided by the management (Mgmt) vNIC in the parent partition. Fault tolerance for the management vNIC is provided by the SET switch. A bandwidth limit can be assigned to the management, as necessary.

- Storage Network

The storage network carries RoCEv2 RDMA network traffic that is used for Storage Spaces Direct, storage replication, and Live Migration network traffic. This network is also used for cluster management communication. The storage network has a Storage A and Storage B segment, each with its own IP subnet. This design keeps the east-west RDMA isolated to the ToR switches and avoids the need for the upstream switches to be configured for supporting RoCEv2 traffic.

DCB (Data Center Bridging) is required for RoCE. If DCB is used, PFC and ETS configuration is implemented properly across every network port, including network switches. RoCE-based Azure Stack HCI imple-

mentations require the configuration of three PFC traffic classes, including the default traffic class, across the fabric and all hosts.

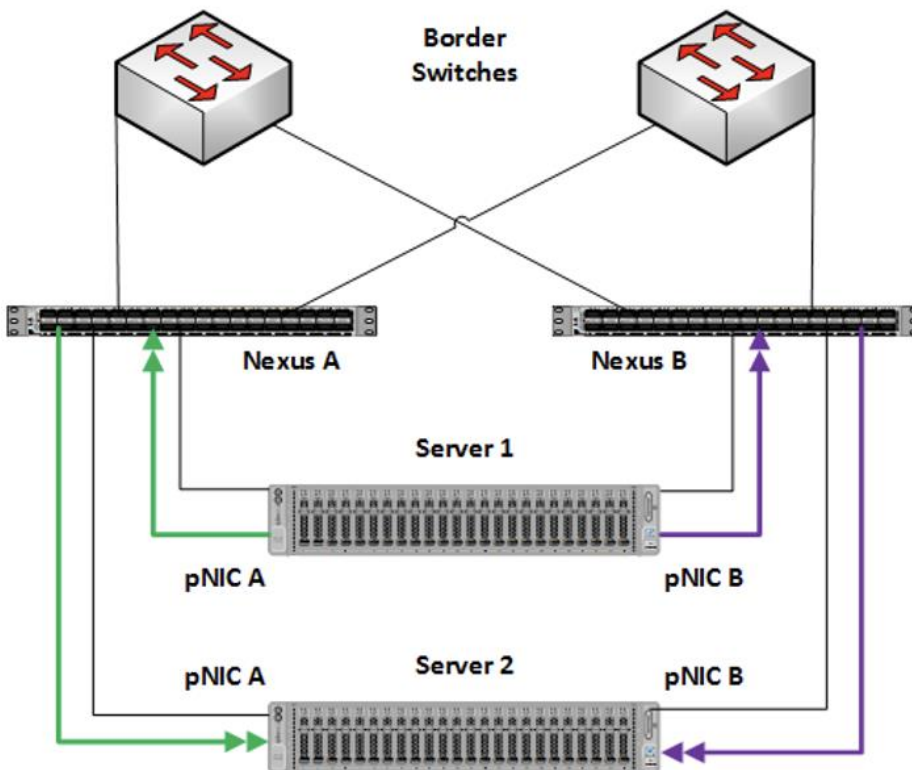
[Table 5](#) lists the QoS configuration used in this document.

Table 5. QoS Configuration

Purpose	Cluster Traffic	Storage (RDMA) traffic	Default (Tenant and Management Networks)
Flow Control (PFC enabled)	No	Yes	No
Traffic Class	5	4	0 (default)
Bandwidth reservation	1% for 25GbE or higher RDMA networks	50%	Default (no host configuration required)

[Figure 5](#) illustrates the east-west RDMA traffic isolation.

Figure 5. East-West RDMA Traffic Isolation




- SET Switch

This is a virtual switch with embedded teaming capabilities. The SET Switch provides teaming capabilities for network traffic that does not use SMB-Multichannel. SMB Direct (RDMA) traffic uses SMB-Multichannel for link aggregation and redundancy instead of the teaming feature in the SET switch.

MAC addresses for virtual NICs are randomly assigned to one on the physical NIC ports on the host. This MAC address assignment can be moved from one physical NIC to another at any time by the SET switch. This behavior provides load balancing and fault tolerance. A consequence of this behavior is that some of the east-west network traffic that is not storage SMB Direct (RDMA) traffic will transverse the upstream router/switch. An example of this is when virtual machine A with a virtual NIC MAC address assigned to physical NIC A communicates with virtual machine B that has virtual NIC MAC assigned to physical NIC B. [Figure 6](#) illustrates this behavior.

Figure 6. MAC Address Assignment



- Guest Partition

The tenant virtual machines run in the guest partition on the Hyper-V host. Each virtual machine runs in isolation from others and does not have direct access to physical hardware in the host. Network connectivity is provided to the tenant virtual machine by connecting synthetic NIC in the virtual machine to the SET switch on the host.

- Parent Partition

The parent partition is the host operating system that runs the virtualization management stack and has access to the physical server hardware. The parent partition has one management vNIC and two storage vNICs. An optional dedicated vNIC for backup operations can be added as needed.

Network ATC, a feature of Azure Stack HCI simplifies the deployment and network configuration management for Azure Stack HCI clusters. Network ATC provides an intent-based approach to host network deployment. By specifying one or more intents (management, compute, or storage) for a network adapter, you can automate the deployment of the intended configuration.

This section explains the logical topology for the below two network reference patterns:

- Fully converged network topology
- Non-converged network topology

Fully Converged Network Topology

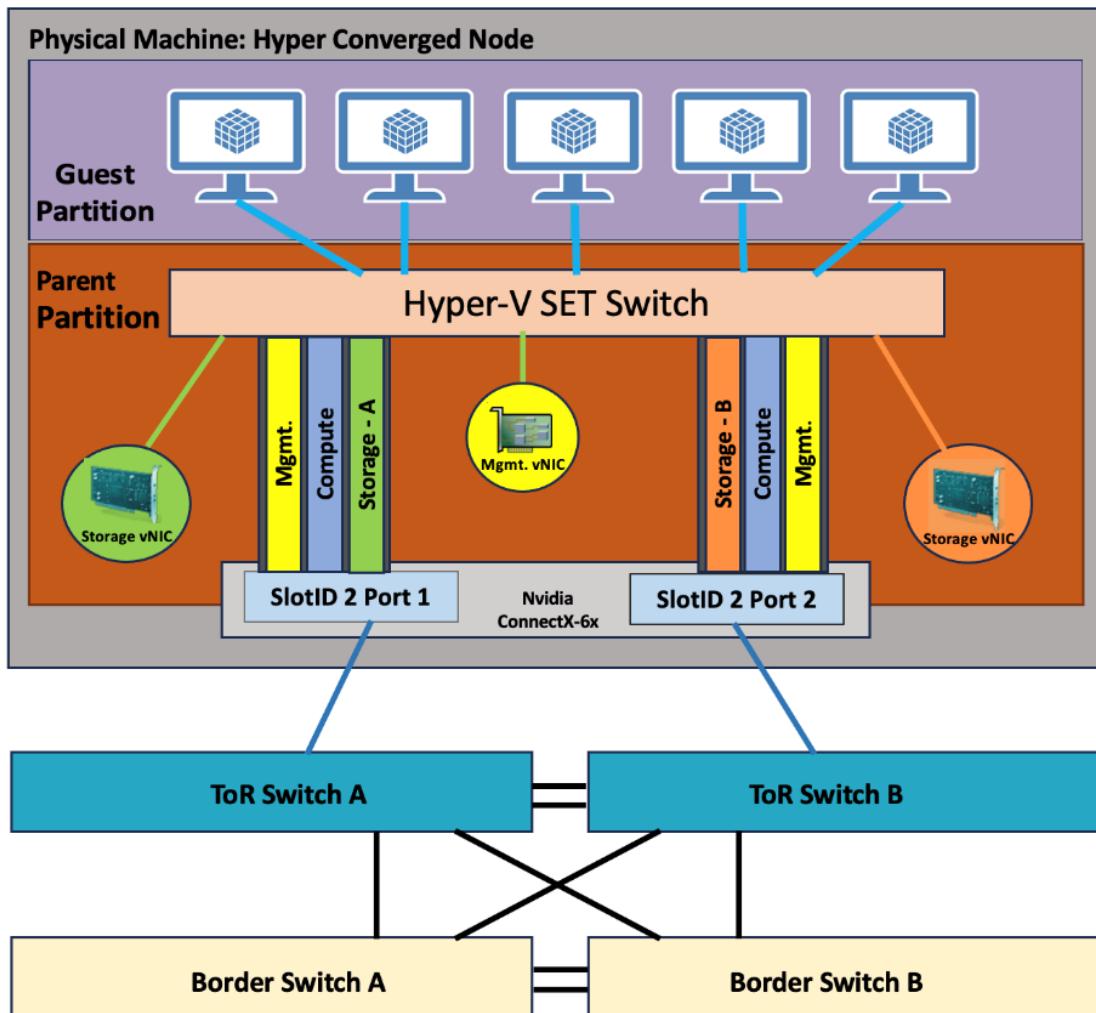
In this fully converged network topology, a single intent for compute, storage and management networks are deployed across all cluster nodes.

Consider this network topology if:

- Bandwidth requirements for north-south traffic don't require dedicated adapters.
- Physical switch ports are scarce and for cost reductions.

As illustrated in [Figure 7](#), this fully converged network pattern has the following logical network components.

Figure 7. Fully converged logical topology



Storage Network VLANs

The storage intent-based traffic in this pattern shares the physical network adapters with management and compute.

The storage network operates in different IP subnets. Each storage network uses the ATC predefined VLANs by default (711 and 712). However, these VLANs can be customized if necessary. If the default subnet defined by ATC is not usable, then manually assign all storage IP addresses in the cluster.

Management VLAN

All physical compute hosts require access to the management logical network. For IP address planning, each physical compute host must have at least one IP address assigned from the management logical network. A DHCP server can automatically assign IP addresses for the management network, or you can manually assign static IP addresses. When DHCP is the preferred IP assignment method, we recommend that you use DHCP reservations without expiration.

The management network supports the following VLAN configurations:

- Native VLAN – not required to supply VLAN IDs. This is required for solution-based installations.
- Tagged VLAN – supply VLAN IDs at the time of deployment.

The management network supports all traffic used for management of the cluster, including Remote Desktop, Windows Admin Center, and Active Directory.

Compute VLANs

Traditional VLANs are used to isolate your tenant workloads. Those VLANs are configured on the TOR switch's port in trunk mode. When connecting new VMs to these VLANs, the corresponding VLAN tag is defined on the virtual network adapter.

Non-Converged and Storage Switchless Network Topology

In this non-converged and storage switchless network topology, two intents – one intent for compute and management networks and a separate intent for storage network are deployed across all cluster nodes.

Consider this network topology:

- For enhanced network performance
- East-West storage traffic replication won't interfere or compete with north-south traffic dedicated for management and compute
- Storage switchless is ideal only for smaller deployments (2 or 3-Node cluster) such as edge, ROBO, etc.

As illustrated in [Figure 8](#) for non-converged and [Figure 9](#) for storage switchless network, the pattern has the following logical network components:

Figure 8. Non-converged logical topology

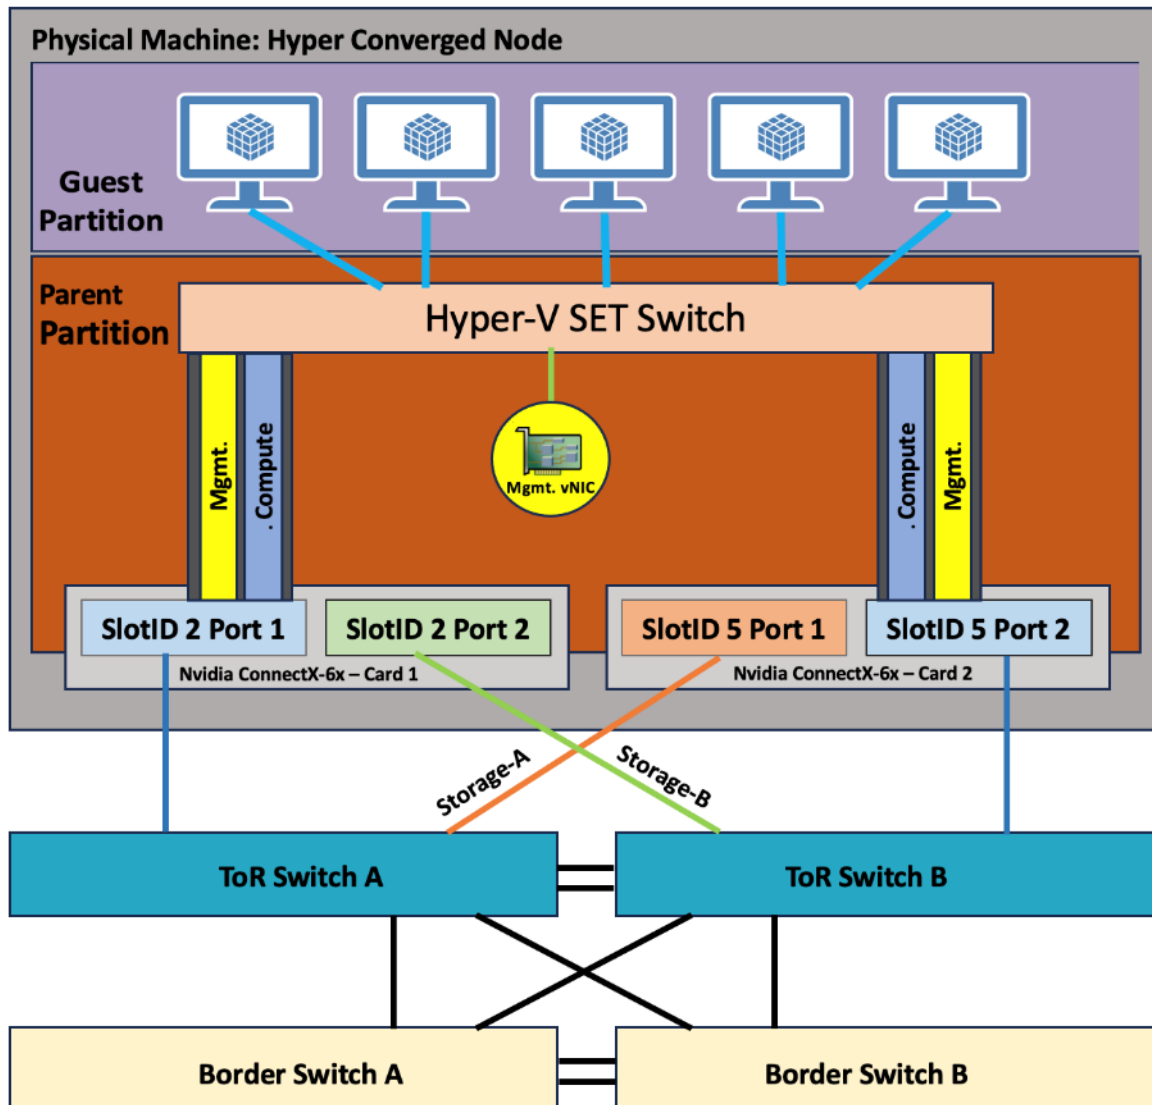
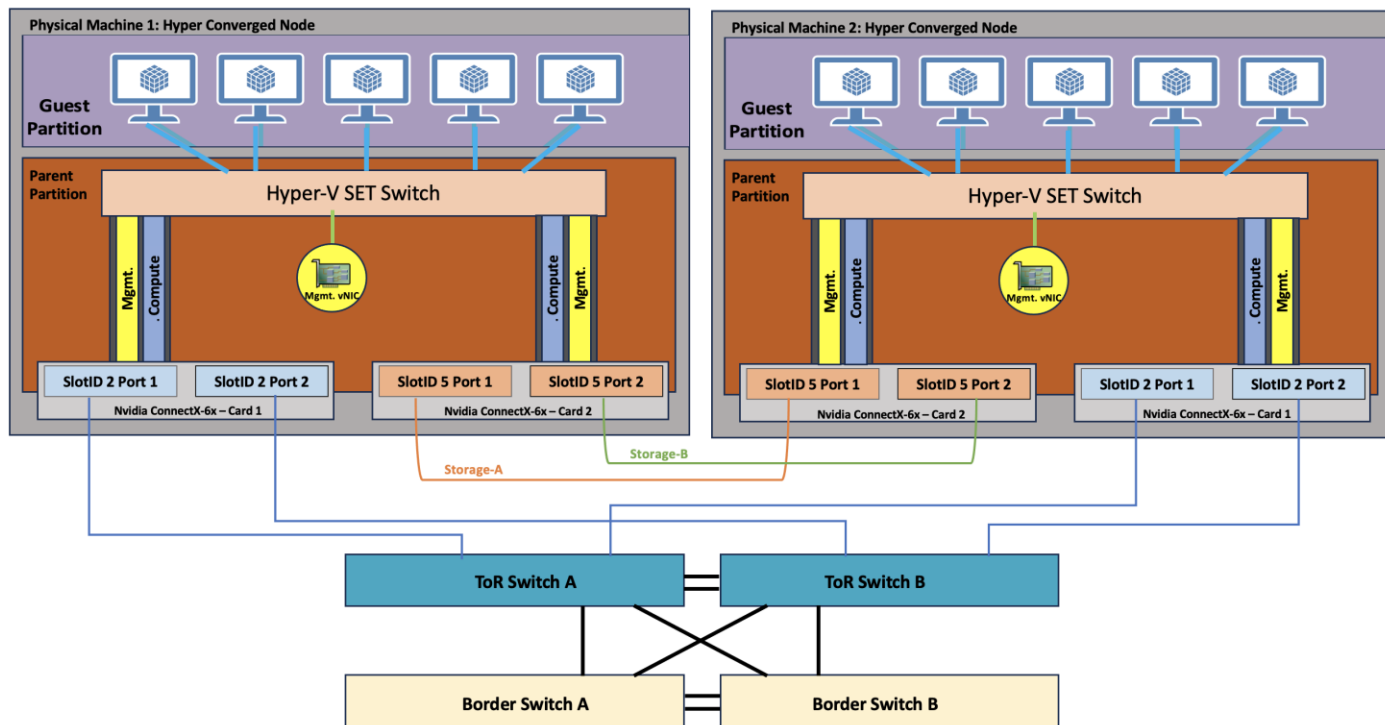


Figure 9. Storage Switchless converged logical topology



Storage Network VLANs

The storage intent-based traffic in this pattern consists of two individual networks supporting RDMA traffic. Each interface is dedicated to a separate storage network, and both can use the same VLAN tag.

The storage network operates in different IP subnets. Each storage network uses the ATC predefined VLANs by default (711 and 712). However, these VLANs can be customized if necessary. If the default subnet defined by ATC is not usable, then manually assign all storage IP addresses in the cluster.

Management VLAN

All physical compute hosts require access to the management logical network. For IP address planning, each physical compute host must have at least one IP address assigned from the management logical network. A DHCP server can automatically assign IP addresses for the management network, or you can manually assign static IP addresses. When DHCP is the preferred IP assignment method, we recommend that you use DHCP reservations without expiration.

The management network supports the following VLAN configurations:

- Native VLAN - not required to supply VLAN IDs. This is required for solution-based installations.
- Tagged VLAN - supply VLAN IDs at the time of deployment.

The management network supports all traffic used for management of the cluster, including Remote Desktop, Windows Admin Center, and Active Directory.

Compute VLANs

Traditional VLANs are used to isolate your tenant workloads. Those VLANs are configured on the TOR switch's port in trunk mode. When connecting new VMs to these VLANs, the corresponding VLAN tag is defined on the virtual network adapter.

[Table 6](#) lists the VLANs used in this document where the default storage VLANs (711 and 712) used by Network ATC are overridden.

Table 6. VLAN Names and IDs

VLAN Name	VLAN ID
Management	126
Tenant	100
Storage-A	107
Storage-B	207

Deployment Hardware and Software

This chapter contains the following:

- [Firmware and Drivers](#)
- [Deployment Checklist](#)
- [Bill of Materials](#)
- [Customer Support Requirements](#)

Firmware and Drivers

Firmware and drivers can be found on the Cisco download portal for Windows Server 2022 (Azure Stack HCI 22H2). These components will be periodically updated. Please sign up for notification at this download portal to receive notifications emails when updates are available.

The Cisco UCS C240 M6 or M7 standalone server platform for Microsoft Azure Stack HCI 22H2 firmware download portal can be accessed from the [Cisco UCS C-Series Rack-Mount Standalone Server Software Download](#) page. Also, it can be set up to notify you about the availability of the new firmware. Cisco highly recommends that you sign up for these notifications.

[Table 7](#) lists the software components hosted on the Cisco download portal and are required for the firmware upgrade procedure.

Table 7. Software Components

Component	Description
ucs-c240m6-huu-4.2.3e.iso	Cisco UCS C240 M6 Rack Server Software
ucs-cxxx-drivers-windows.4.2.3e.iso	Azure Stack HCI 22H2 (Win 2022) drivers for Cisco UCS C240 M6SN servers
ucs-c240m7-huu-4.3.2.240002.iso	Cisco UCS C240 M7 Rack Server Software
ucs-cxxx-drivers-windows.4.3.2f.iso	Azure Stack HCI 22H2 (Win 2022) drivers for Cisco UCS C240 M7SN servers
ucs-c220m7-huu-4.3.2.240002.iso	Cisco UCS C220 M7 Rack Server Software
ucs-cxxx-drivers-windows.4.3.2f.iso	

The following tables list the individual component version that are part of the respective firmware bundles and driver package:

Cisco UCS C-Series Rack-Mount Standalone Server	
Component	Firmware Version

Cisco UCS C-Series Rack-Mount Standalone Server

Cisco UCS C240 M6SN	4.2(3e)	
Cisco UCS C240 M7SN	4.3(2.240002)	
Cisco UCS C220 M7N	4.3(2.240002)	

Cisco UCS C240 M6SN Servers

Component	C-Series Rack-Mount	Firmware Version	Driver Version
CIMC (BMC)	4.2(3e)	4.2.3e	
BIOS	4.2(3e)	C240M6.4.2.3c.0.042023 0316	
Cisco-MLNX MCX623106AS-CDAT 2x100GbE QSFP56 PCIe	4.2(3e)	22.38.1900	3.10.51000
Cisco UCS-M2-HWRAID	4.2(3e)	2.3.17.1014	
Boot SSD (UCS-M2-960GB)	4.2(3e)	D0MH077	10.0.17763.887 (inbox)
U.2 Intel P5500 NVMe	4.2(3e)	2CV1C033	10.0.20348.1547 (inbox)

Cisco UCS C240 M7SN and C220 M7SN Servers

Component	C-Series Rack-Mount	Firmware Version	Driver Version
CIMC (BMC)	4.3(2.240002)	4.3(2.240002)	
BIOS	4.3(2.240002)	C240M7.4.3.2d.0.110123 2037	
Cisco-MLNX MCX623106AS-CDAT 2x100GbE QSFP56 PCIe	4.3(2.240002)	22.38.1900	3.10.51000
Cisco UCS-M2-HWRAID	4.3(2.240002)	2.3.17.1014	

Cisco UCS C240 M7SN and C220 M7SN Servers

Boot SSD (UCS-M2-960GB)	4.3(2.240002)	D3MC000	10.0.17763.887 (inbox)
U.2 Intel P5520 NVMe	4.3(2.240002)	9CV10200	10.0.20348.1 (inbox)

Host Operating System

Host OS Version	Azure Stack HCI OS 22H2 with current updates
-----------------	--

Physical Infrastructure

[Figure10](#) illustrates the physical topology of an Azure Stack HCI deployment on Cisco UCS C240 M6/M7 servers with Cisco Nexus 9300 series switches. The cabling information can be found in the [Appendix](#) of this document.

Figure 10. Physical Infrastructure

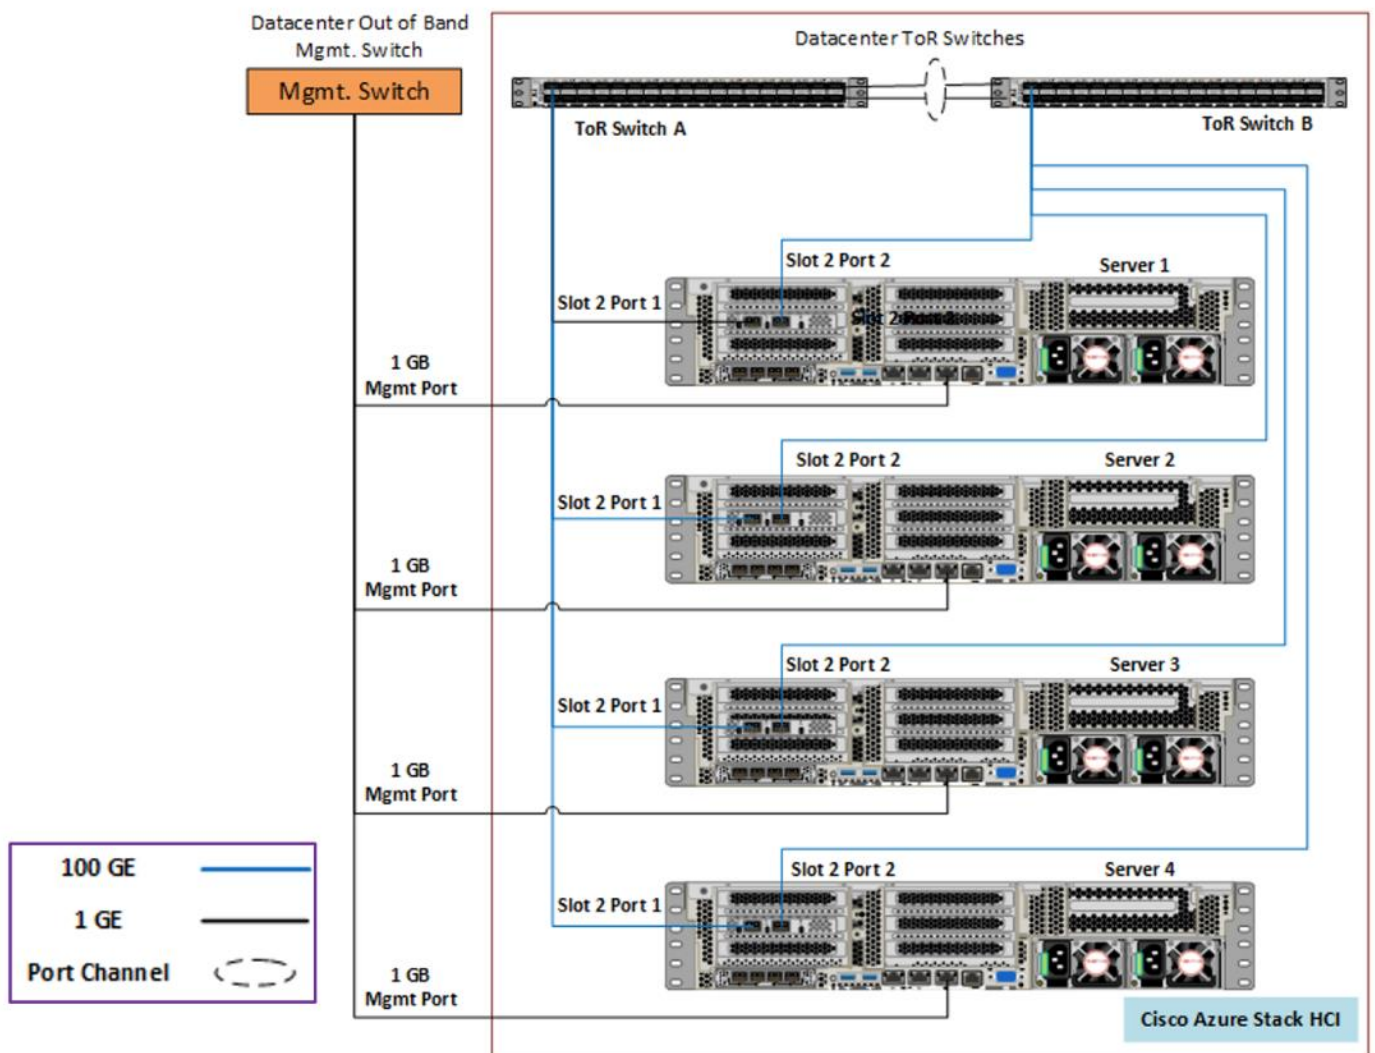


Figure 11 illustrates the data ports and management ports on the back of each server with one dual-port network adapter in a fully converged network topology.

- In this network reference pattern, the ToR switches in redundant configuration handle north-bound/southbound traffic.
- Two ports teamed on the host side handle the management, compute, and RDMA storage traffic connected to the ToR switches. Each port in the network adapter on each server is connected to a different ToR switch. SMB multichannel capability provides path aggregation and fault tolerance.

In this example, Server 1 has these two 100Gb data ports connected to ports eth1/1 port on ToR A and B switches. The single dedicated out-of-band management port is connected to an OOB management switch.

Figure 11. Fully converged network topology with one dual-port network adapter




Figure 12 illustrates the data ports and management ports on the back of each server with two dual-port network adapters in a fully converged network topology.

- In this network reference pattern, the ToR switches in redundant configuration handle north-bound/southbound traffic.
- Four ports teamed on the host side handle the management, compute, and RDMA storage traffic connected to the ToR switches. Each port on the two network adapters on each server is connected to a different ToR switch. SMB multichannel capability provides path aggregation and fault tolerance.

In this example, Server 1 has two 100Gb data ports from the first network adapter connected to ports eth1/1 port on ToR A and B switches and, the other two 100Gb data ports from the second network adapter connected to ports eth1/2 port on ToR A and B switches. The single dedicated out-of-band management port is connected to an OOB management switch.

Figure 12. Fully converged network topology with two dual-port network adapters




Figure 13 illustrates the data ports and management ports on the back of each server with two dual-port network adapters in a non-converged network topology.

- In this network reference pattern, the ToR switches in redundant configuration handle north-bound/southbound traffic.
- Two ports, one from each network adapter ('SlotID 2 Port 1' and 'SlotID 5 Port 2') teamed on the host side handle the management and compute traffic connected to the ToR switches. These two ports on each server are connected to a different ToR switch.
- Two ports, one from each network adapter ('SlotID 2 Port 2' and 'SlotID 5 Port 1') in standalone configuration are used for RDMA storage traffic. These two ports on each server are connected to a different ToR switch. SMB multichannel capability provides path aggregation and fault tolerance.

In this example, Server 1 has two 100Gb data ports from the first network adapter connected to ports eth1/1 port on ToR A and B switches and, the other two 100Gb data ports from the second network adapter connected to ports eth1/2 port on ToR A and B switches. The single dedicated out-of-band management port is connected to an OOB management switch.

Figure 13. Non-converged network topology with two dual-port network adapters


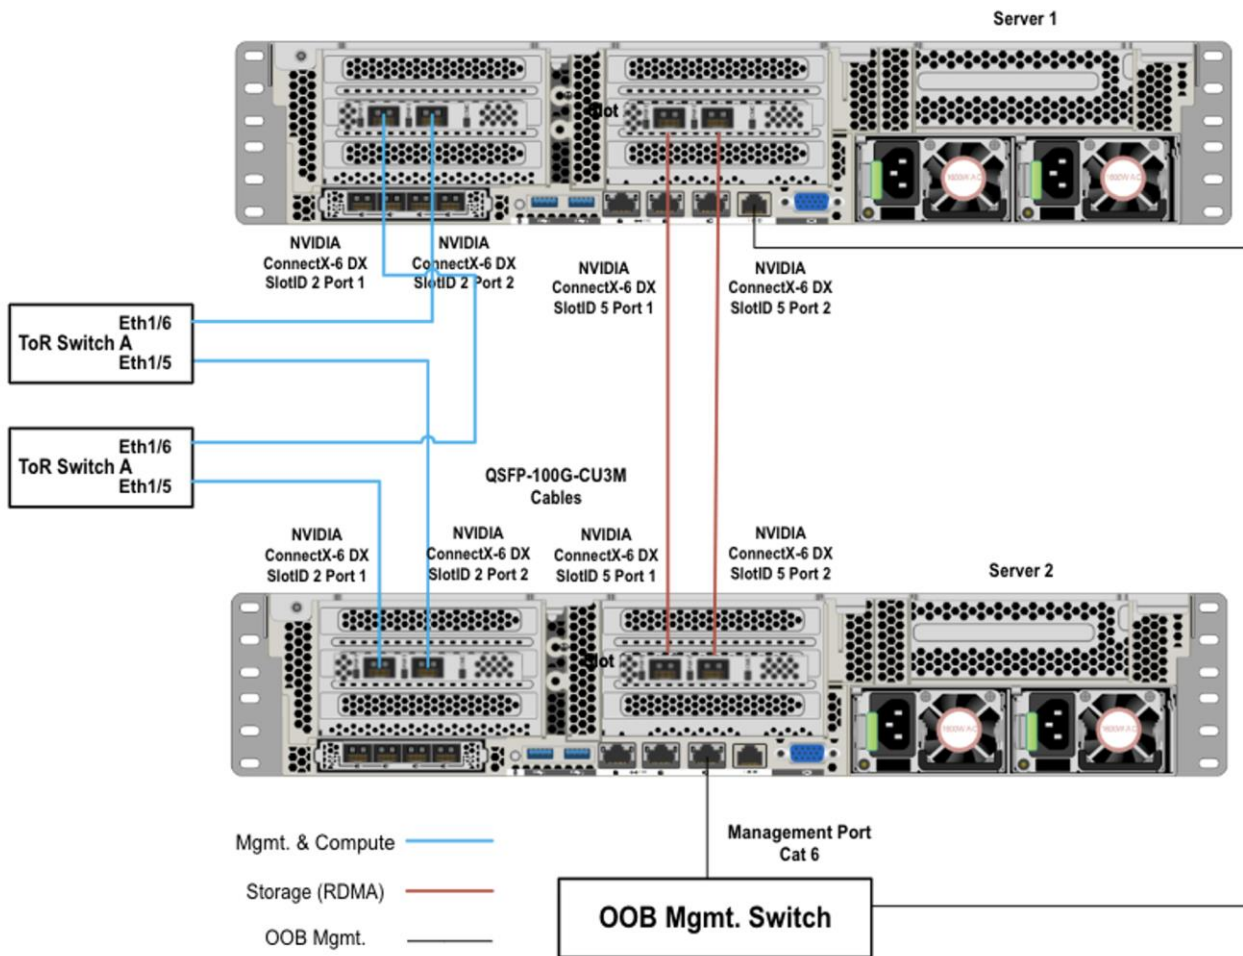


Figure 14 illustrates the data ports and management ports on the back of each server with two dual-port network adapters in a storage switchless network topology:

- In this network reference pattern, the ToR switches in redundant configuration handle north-bound/southbound traffic.
- Two ports ('SlotID 2 Port 1' and 'SlotID 2 Port 2') teamed on the host side handle the management and compute traffic connected to the ToR switches. These two ports on each server are connected to a different ToR switch.
- Two RDMA ports ('SlotID 5 Port 1' and 'SlotID 5 Port 2') in standalone configuration are used for East-West storage traffic. These two ports on each server are directly connected in a full-mesh configuration. SMB multichannel capability provides path aggregation and fault tolerance.

Figure 14. Storage Switchless network topology with two dual-port network adapters



Deployment Checklist

The following is the checklist for the deployment of a 4-node Azure Stack HCI cluster:

- ToR switch must support the [Azure Stack HCI requirements](#)
- ToR switch must implement L2 and L3 configuration for transporting northbound host and tenant traffic
- Out-of-Band management switch must be provided for connecting the ToR switches and Cisco UCS C240 M6 and M7 servers
- 3 IP addresses are required on the Out-of-Band Management Network for the ToR Cisco Nexus switches
- 1 IP address must be provided for each host (server) on the Out-of-Band Management Network
- VLANs
 - 1 Management
 - 2 Storage
 - 1 or more tenant

- IP subnets and addresses for all endpoints for the above VLANs
- Storage VLANs and Storage subnets need to be configured on the ToR switches
- Host operating system must have access to Azure
- Datacenter infrastructure that includes Active Directory Services, DNS, and NTP
- Cluster Quorum Witness
 - Can be Files Share or Cloud Witness
 - Required for Cluster with fewer than 5 cluster nodes
- Recommended for clusters with 3 or greater number of nodes
- Deployment host must be provided with access to the Out-of-Band Managed network and host management network
 - See the [Remote Management Host](#) configuration in the [Appendix](#)
- Deployment host must be running Windows Server 2019 or Windows Server 2022 and be domain joined to the same domain as the Azure Stack HCI hosts
- Account used to deploy Azure Stack HCI must have administrative rights on the Azure stack hosts and permissions to join the domain, add cluster securing principle to the domain, update the DNS A records for the computer joining the domain and Cluster Aware Updating services, and store Bitlocker keys in the domain
- Azure Account for registering Azure Stack HCI
- Download Azure Stack HCI OS 22H2 from Microsoft download site
- Download Cisco Drivers for Azure Stack HCI 22H2 deployment from Cisco download portal (link to be added)
- Download Cisco UCS Manager configuration script for Azure Stack HCI 22H2 deployments from Cisco download portal (link to be added)
- Recommended Items
 - Windows Deployment Service for PXE boot OS installation (Can be running on deployment host)
 - DHCP server with scope for management subnet to support PXE booting. Scope is temporary and only needed during PXE boot installation phase. (Can be running on deployment host)

Bill of Materials

This solution must be purchased using the Cisco UCS product ID **UCS-MAH-B00R00-M6**. This product ID includes all of the required hardware to build the solution as well as the Cisco Solution Support for this solution. A sample BoM is documented in the Cisco UCS for Microsoft Azure Stack HCI Datasheet at the following link: <https://www.cisco.com/c/en/us/solutions/data-center-virtualization/microsoft-applications-on-cisco-ucs/microsoft-azure-stack-hci.html>

Customer Support Requirements

The solution must adhere to Cisco Guidance for deploying Azure Stack HCI on Cisco UCS product ID **UCS-MAH-B00R00**.

Firmware and driver version must match the versions specified in this document. This document will be update periodically with more current firmware and driver versions. Customers are required to update their systems to the latest recommended firmware and driver version for this Azure Stack HCI solution.

Note: The current firmware and drivers can be downloaded from the Cisco download portal for Azure Stack HCI. The link to the download portal is in the [Firmware and Drivers](#) section.

Note: You must obtain an Azure Stack HCI support contract from Microsoft. The following is an example of this type of support contract:

- Unified Support for Enterprise
- Premier Support for Enterprise

For support option details, go to: [Get support for Azure Stack HCI - Azure Stack HCI | Microsoft Docs](#)

Solution Configuration

This chapter contains the following:

- [Configure Cisco Integrated Management Controller for Cisco UCS C240 Servers](#)
- [Claim Cisco UCS C240 Standalone Servers in Cisco Intersight](#)
- [Configure Cisco UCS C240 Standalone Servers using Cisco Intersight](#)
- [Configure Policies to Create Server Profile](#)
- [Configure Bitlocker for System Volume](#)
- [Configure Secured-Core on Hosts](#)
- [Fully Converged Network Topology Configuration](#)
- [Non-Converged Network Topology Configuration](#)
- [Storage Switchless Network Topology Configuration for 2 Nodes](#)
- [Configure QoS](#)
- [Prepare Server for Storage Spaces Direct](#)
- [Configure Windows Failover Cluster](#)
- [Configure Storage Spaces Direct](#)

Configure Cisco Integrated Management Controller for Cisco UCS C240 Servers

Procedure 1. Configure Cisco Integrated Management Controller (IMC)

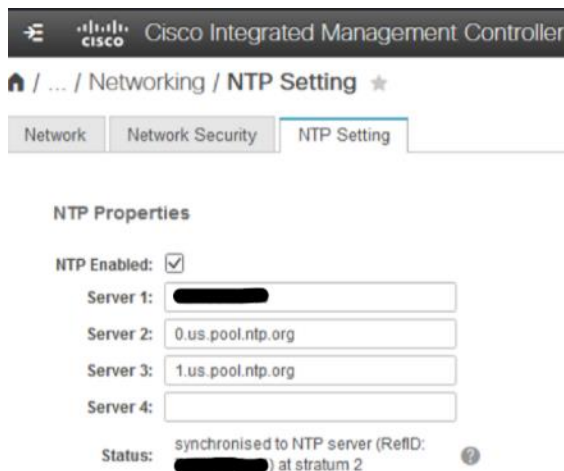
- Step 1. In the BIOS POST screen, press **F8** to display the CIMC configuration screen.
- Step 2. A prompt displays to enter the default password and provide the user password (only first time).
- Step 3. Select **Dedicated NIC** mode.
- Step 4. Select **Static** or **DHCP** assignment.
- Step 5. For Static mode, configure the IP address, Netmask and Gateway for the IPv4 setting of the CIMC.
- Step 6. Select **None** for NIC redundancy.
- Step 7. Press **F10** to save the configuration and exit the utility.
- Step 8. Open a web browser on a computer on the same network.
- Step 9. Enter the IMC IP address of the Cisco UCS C240 M6M7 Server: http://<var_cimc_ip_address>.
- Step 10. Enter the login credentials as updated in the IMC configuration.



Procedure 2. Synchronize Cisco UCS C240 Servers to NTP

Note: These steps provide the details for synchronizing the Cisco UCS environment to the NTP server.

- Step 1. Log back into Cisco IMC using a URL that starts with https://.
- Step 2. Select the Admin at the bottom of the left window and expand.
- Step 3. Select Networking > NTP Setting.
- Step 4. Select **NTP Enabled** check box to enable and enter the NTP server addresses.



Claim Cisco UCS C240 Standalone Servers in Cisco Intersight

Procedure 1. Cisco Intersight Device Claim – Register Cisco IMC to Cisco Intersight

Step 1. From the Cisco IMC, go to **Admin > Device connector**.

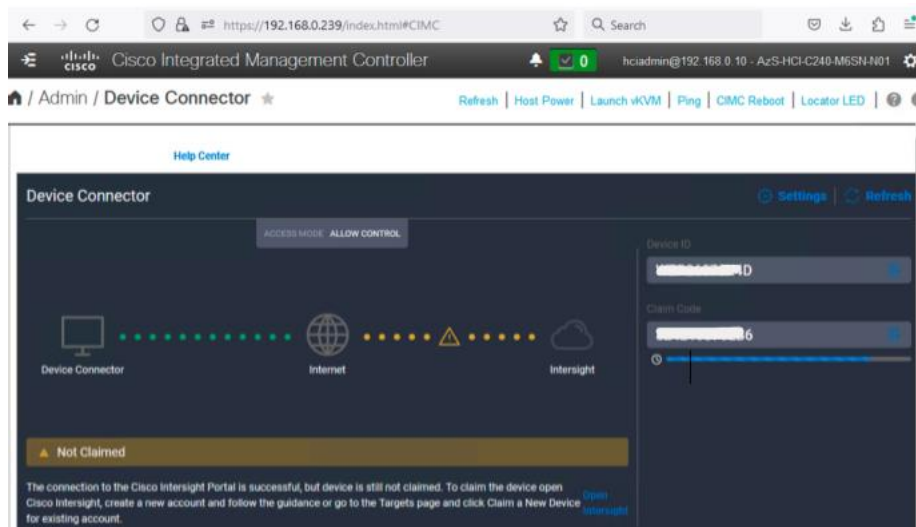
Step 2. On the right side of the screen, click **Settings**.

Step 3. From **Settings**, go to the **General** tab and enable the **Device connector**. For the Access Mode, select **Allow control** and enable **Tunneled vKVM**.

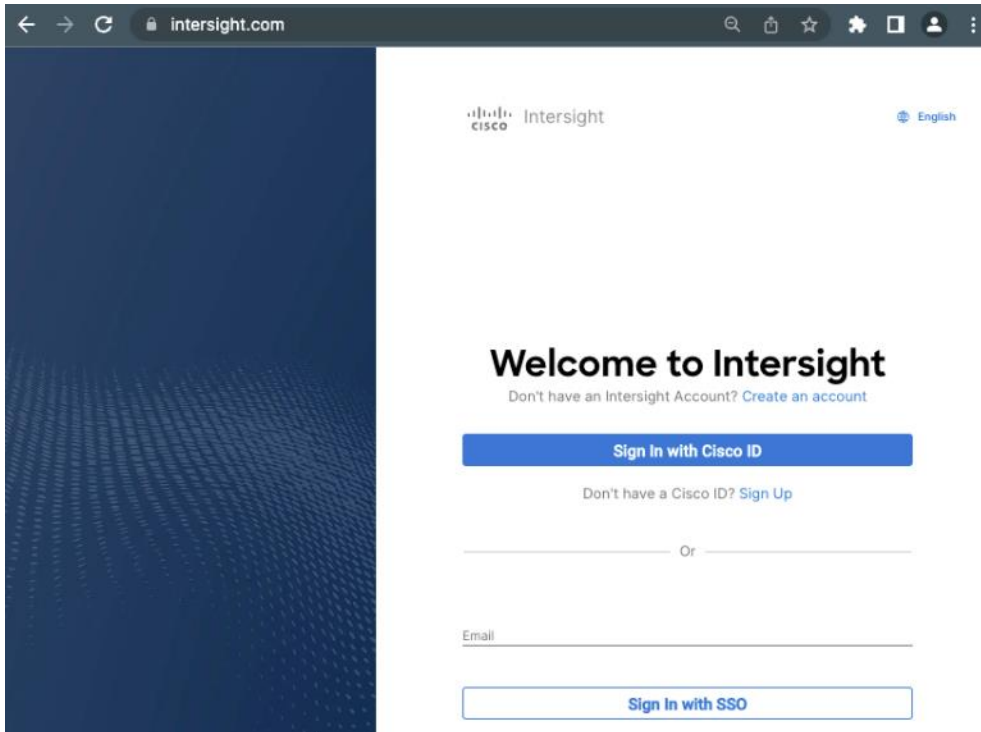
Note: Tunneled vKVM is supported only for Cisco UCS C-Series servers with an Advantage or Premier license. Tunneled vKVM can be launched to complete OS installation from Cisco Intersight.

Step 4. Verify reachability to Cisco Intersight is updated after configuring DNS, NTP and Proxy Settings.


Step 5. Copy the Device ID and Claim Code.




Step 6. Create a Cisco Intersight account—go to <https://intersight.com/> to create your Intersight account. You must have a valid Cisco ID to create a Cisco Intersight account. If you do not have a Cisco ID, create one by clicking **Sign Up**.



Step 7. After Logging in, from the Service Selector drop-down list, select **System** as shown below:



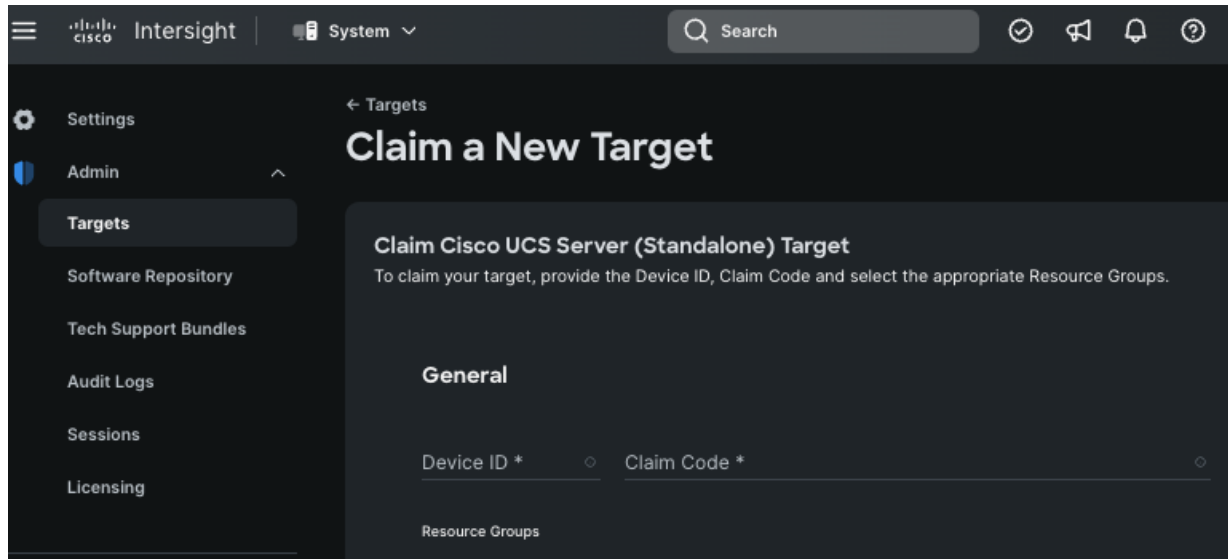
Step 8. Navigate to ADMIN > Targets and click Claim a New Target.



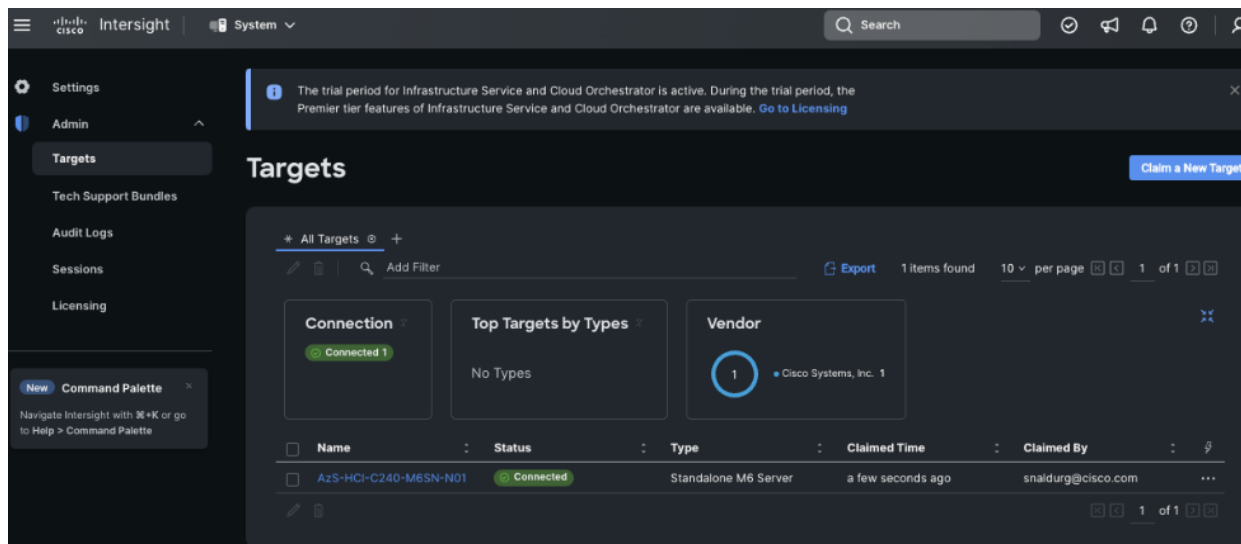
The Select Target Type window is displayed.

Step 9. In the filter column, select **Compute / Fabric** and select **Cisco UCS Server (Standalone)**, and then click **Start**.

Step 10. Enter the **Device ID** and **Claim Code** obtained from Cisco IMC.




Step 11. Click **Claim**.

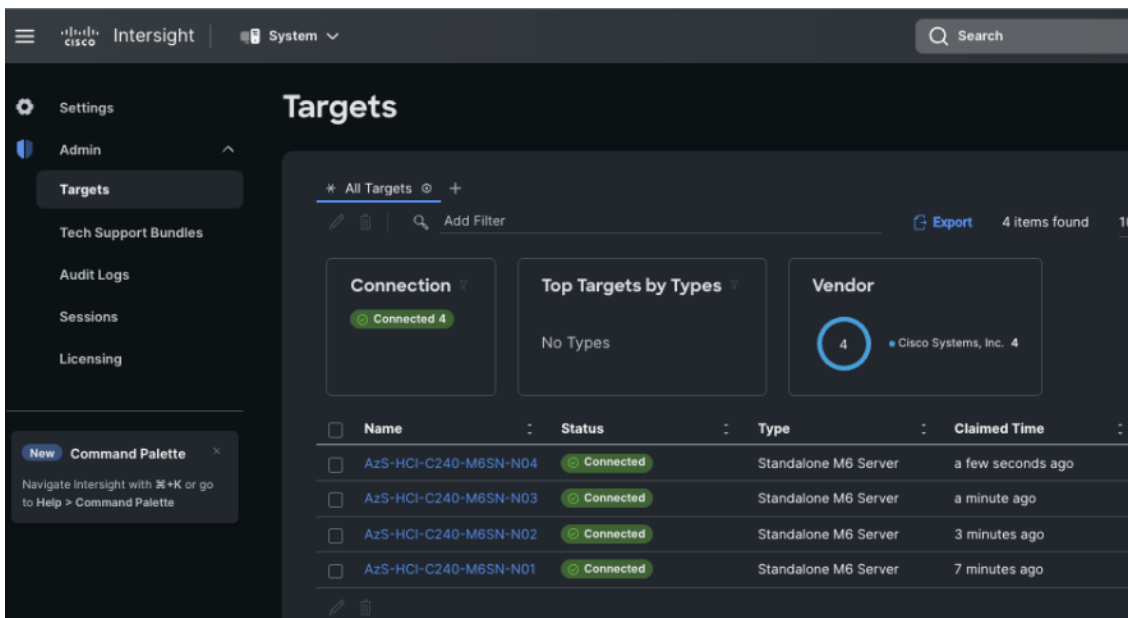


The Cisco UCS Server instance will be added to Intersight.

Step 12. Go back to **Cisco IMC** to confirm that the device is claimed. Click **Refresh** to update the status.



Step 13. Repeat steps 1 - 12 to claim other devices. After the targets are claimed, you can view the managed targets in the Targets table view.



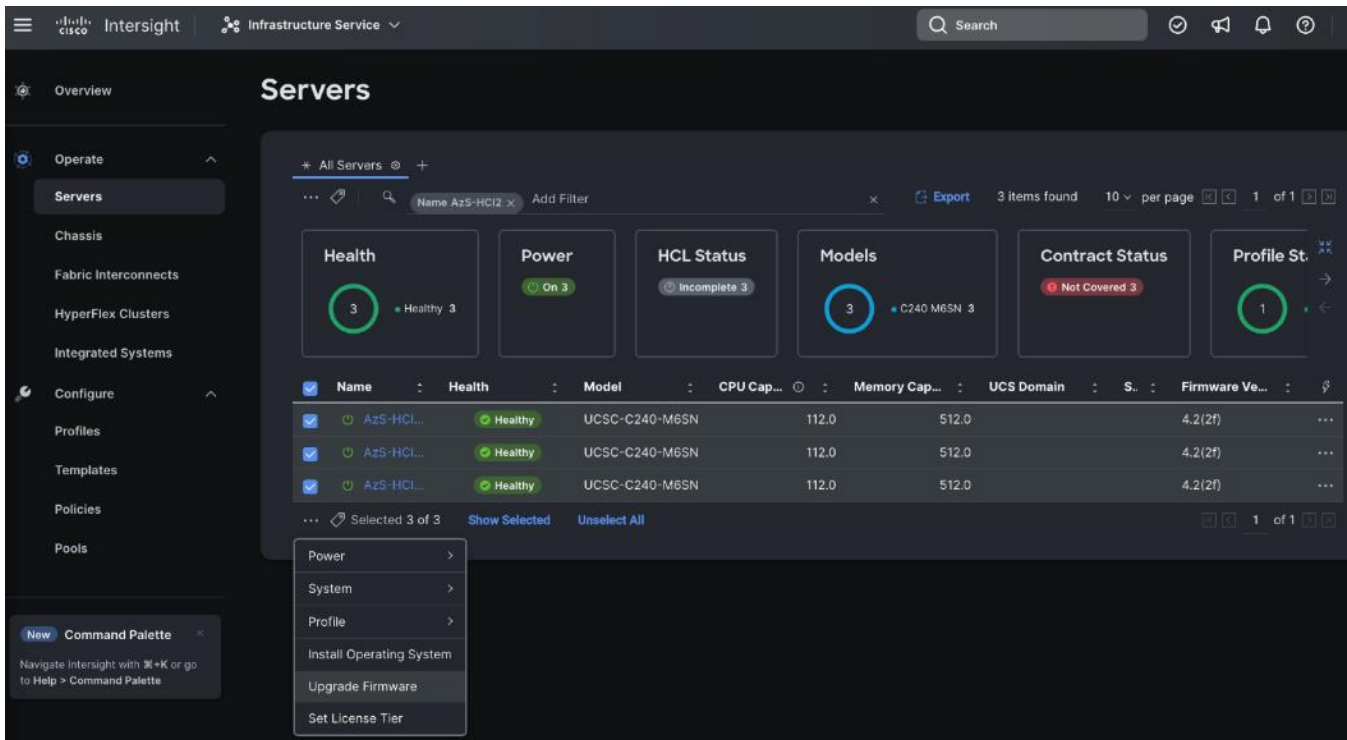
Step 14. Navigate to **Settings > Admin > Licensing** and register the license to assign Essential, Advanced, or Premier license for Cisco Intersight.

For more information about the different license tiers for Cisco Intersight, go to: [Cisco Intersight License Management](#).

Configure Cisco UCS C240 Standalone Servers using Cisco Intersight

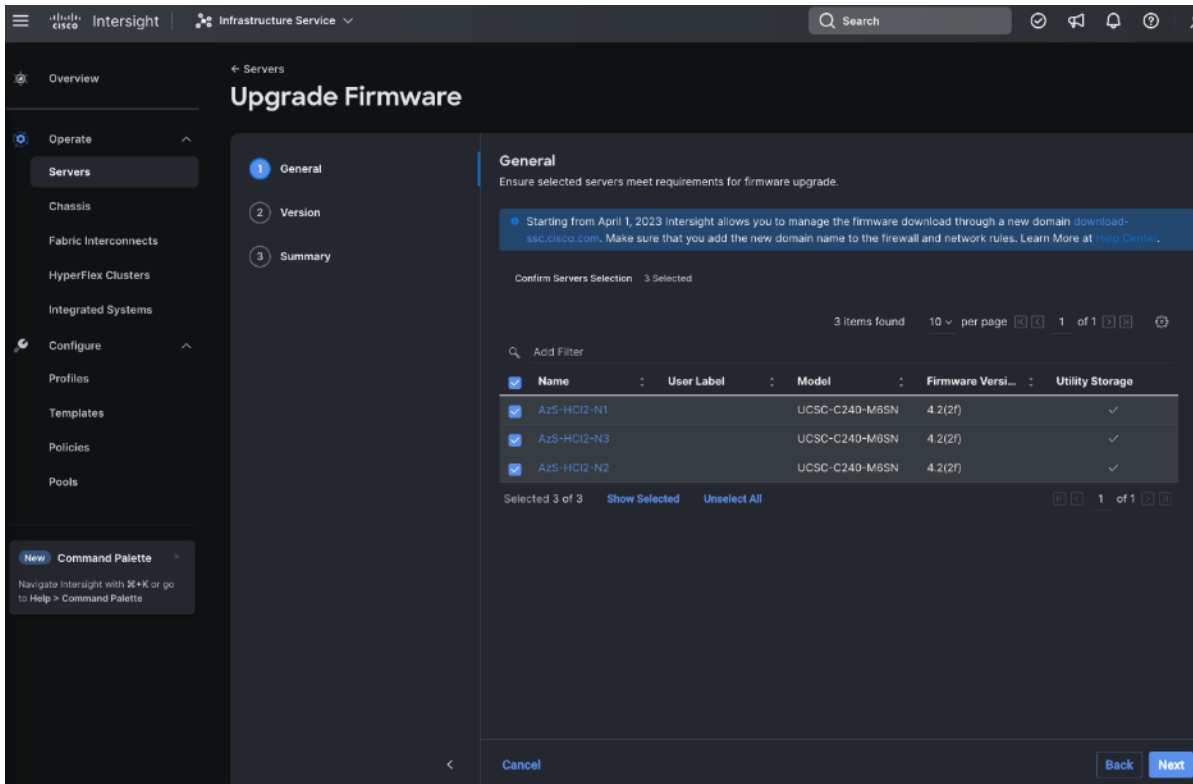
Procedure 1. Upgrade Cisco IMC firmware for Cisco UCS C240 from Cisco Intersight

- Step 1. From the Service Selector drop-down list, select **Infrastructure Service**.
- Step 2. Navigate to **Operate > Servers**, to launch the Servers Table view and select all the servers that require CIMC firmware upgrade.
- Step 3. Click the ellipses below the selected servers and click **Upgrade Firmware**.



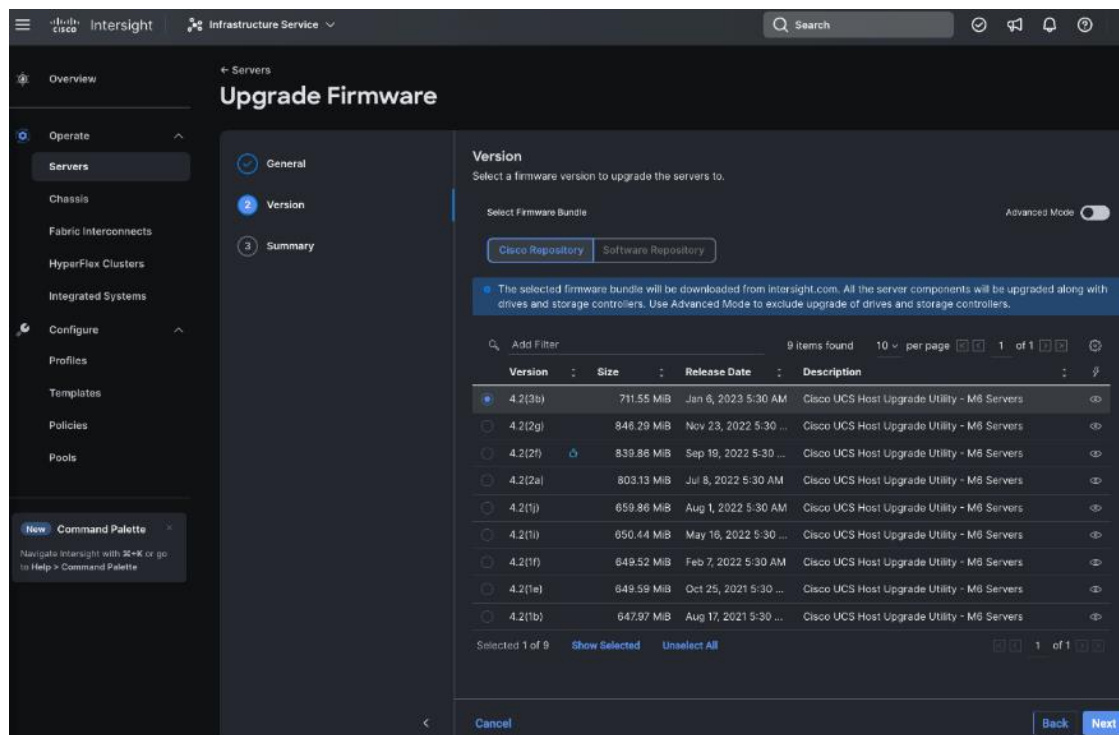
Step 4. On the Upgrade Firmware page, click **Start**.

Step 5. On the **General** page, select all the Servers and click **Next**.



Step 6. On the Version page, enable the **Advanced Mode** to exclude upgrade of drives and storage controllers:


- Exclude Drives—Check this box to exclude upgrade of drives.
- Exclude Storage Controllers—Check this box to exclude upgrade of storage controllers.



Note: To exclude storage controller, ensure that the firmware version of Cisco IMC and the target upgrade firmware version is 4.1(3a) or later release.

Step 7. On the Version page under Cisco Repository, select a firmware version bundle from the list below to upgrade the servers to and click **Next**.

Step 8. On the **Summary** page, confirm the configuration and click **Upgrade** to initiate the upgrade.



For more information on upgrading Cisco UCS C-Series Standalone Servers Firmware, go to: [Before you begin](#).

The upgrade workflow proceeds based on the selected reboot option.

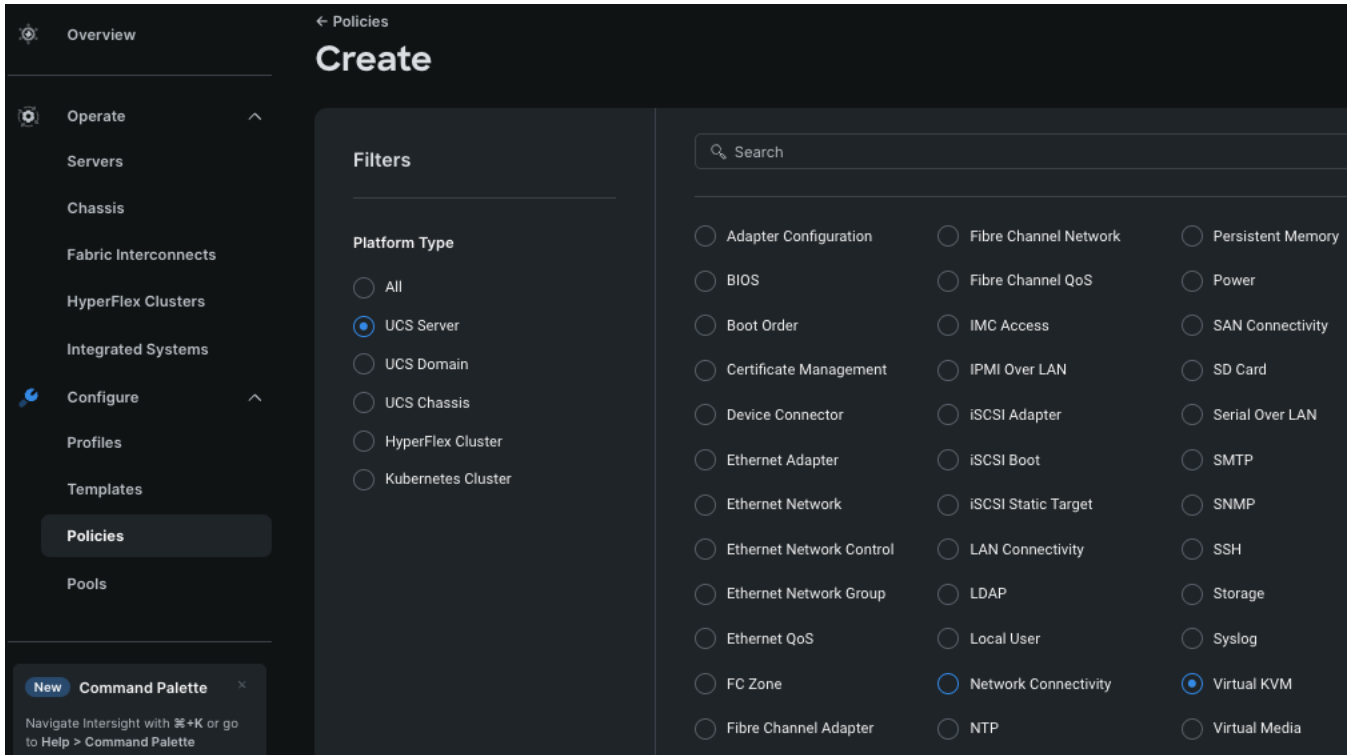
Configure Policies to Create Server Profile

Note: These steps can also be completed at the time of the Server Profile creation.

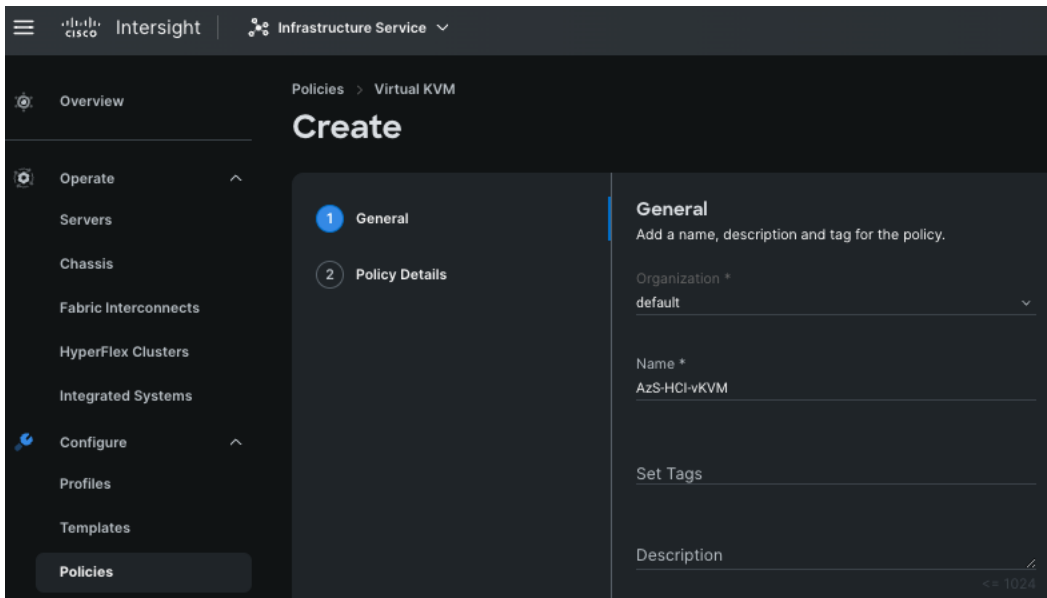
Procedure 1. Create Virtual KVM Policy

Step 1. From the Service Selector drop-down list, select **Infrastructure Services** and navigate to **Configure > Policies** and click **Create Policy**.

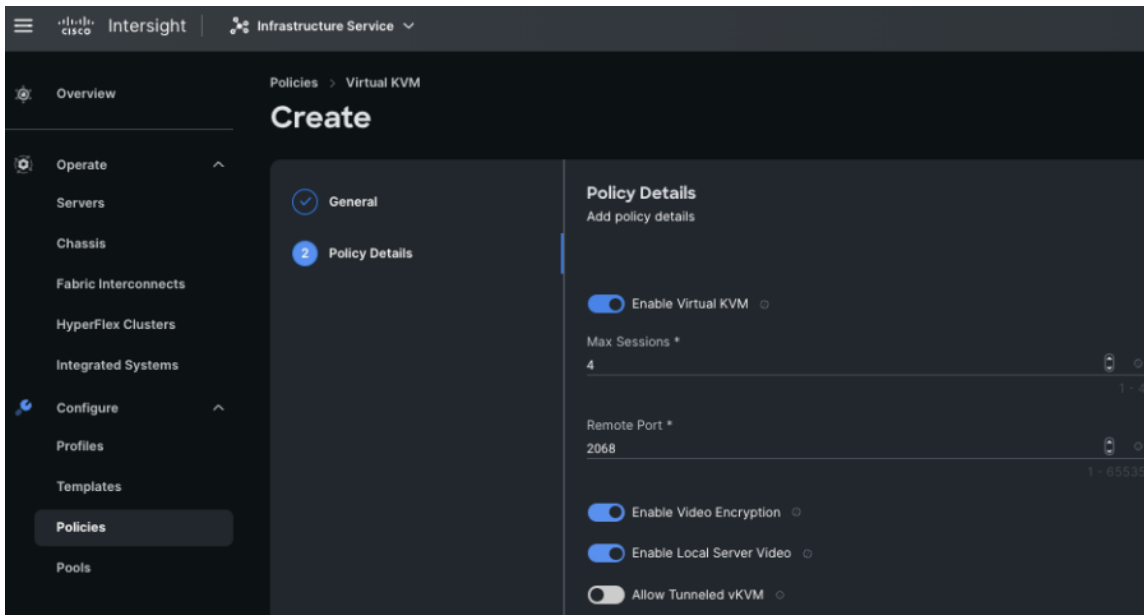
Step 2. On the Create page for Policies, select **UCS Server > Virtual KVM** and click **Start**.



Step 3. On the Virtual KVM Create page, enter the Organization, Name, Description and create a new tag or assign an existing tag and click **Next**.

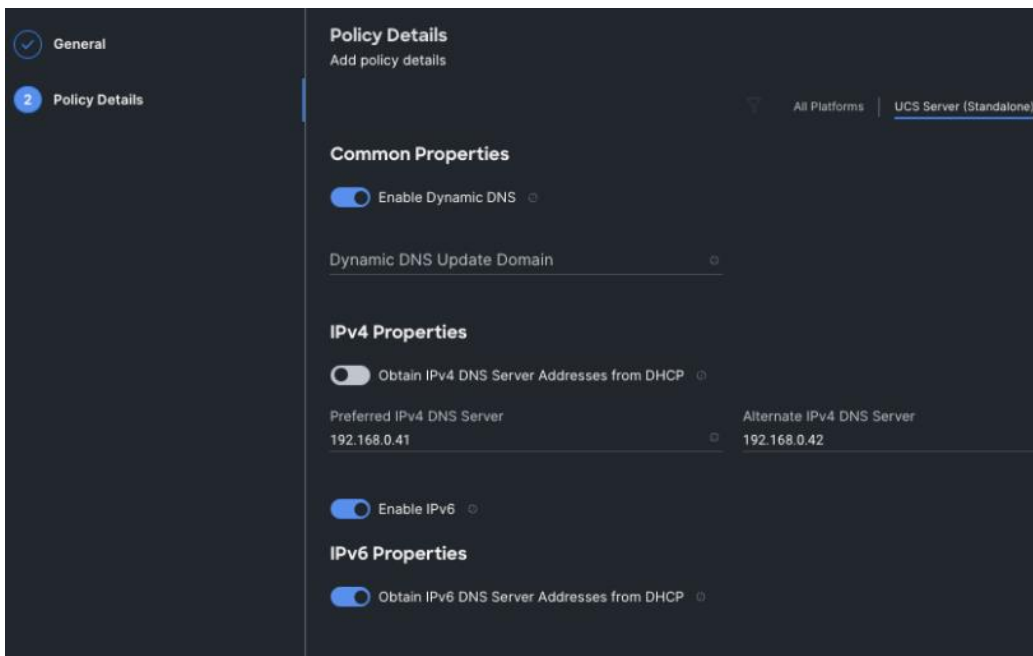


Step 4. On the Policy Details page, enable **Allow Tunneled vKVM**, and other options as shown below and click **Create**.



Procedure 2. Create Network Connectivity Policy

- Step 1. From the Service Selector drop-down list, select **Infrastructure Services** and navigate to **Configure > Policies** and click **Create Policy**.
- Step 2. On the Create page for Policies, select **UCS Server > Network Connectivity** and click **Start**.
- Step 3. On the Network Connectivity Create page, enter the Organization, Name, Description and create a new tag or assign an existing tag and click **Next**.
- Step 4. On the Policy Details page, enter the preferred IPv4 DNS server addresses and configure other options as shown below and click **Create**.



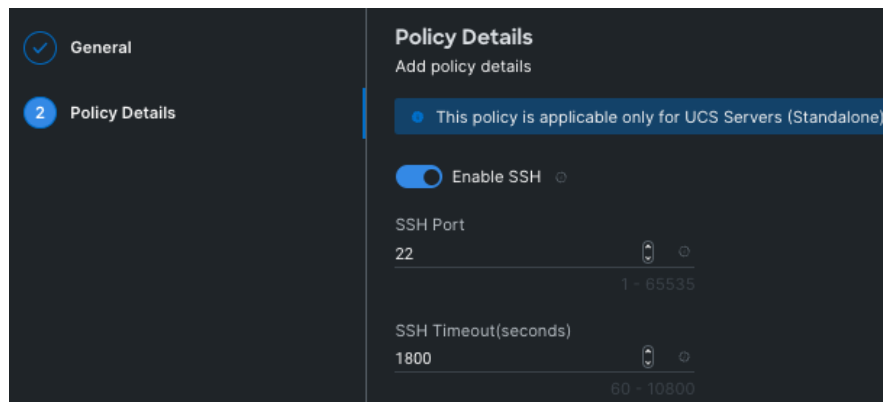
Procedure 3. Create SSH Policy

Step 1. From the Service Selector drop-down list, select **Infrastructure Services** and navigate to **Configure > Policies** and click **Create Policy**.

Step 2. On the Create page for Policies, select **UCS Server > SSH** and click **Start**.

Step 3. On the SSH Create page, enter the Organization, Name, Description and create a new tag or assign an existing tag and click **Next**.

Step 4. On the Policy Details page, **Enable SSH** and click **Create**.



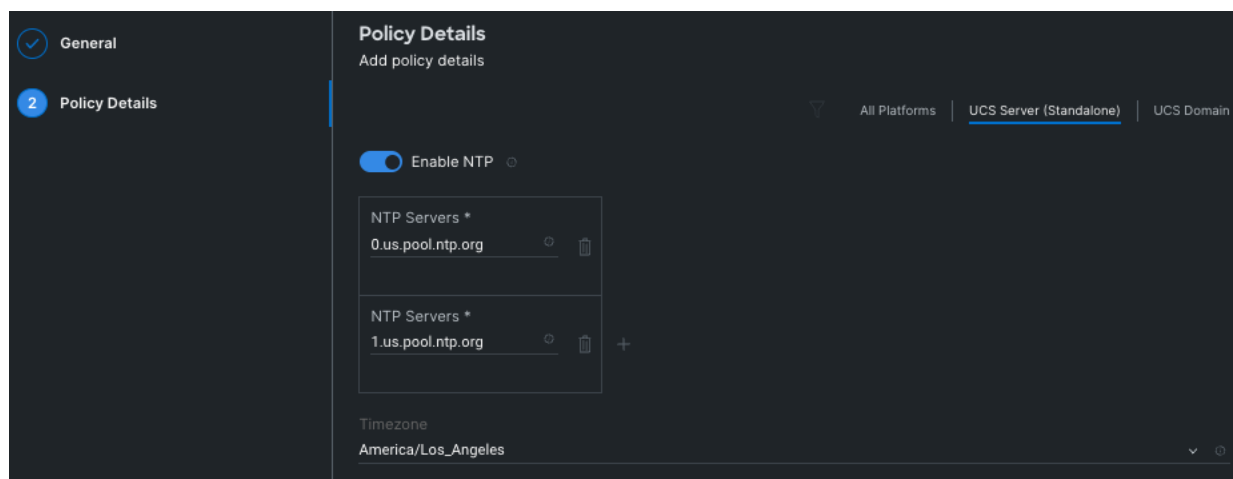
Procedure 4. Create NTP Policy

Step 1. From the Service Selector drop-down list, select **Infrastructure Services** and navigate to **Configure > Policies** and click **Create Policy**.

Step 2. On the Create page for Policies, select **UCS Server > NTP** and click **Start**.

Step 3. On the NTP Create page, enter the Organization, Name, Description and create a new tag or assign an existing tag and click **Next**.

Step 4. On the Policy Details page, Enable NTP, enter the NTP Server addresses and select a TimeZone. Click **Create**.



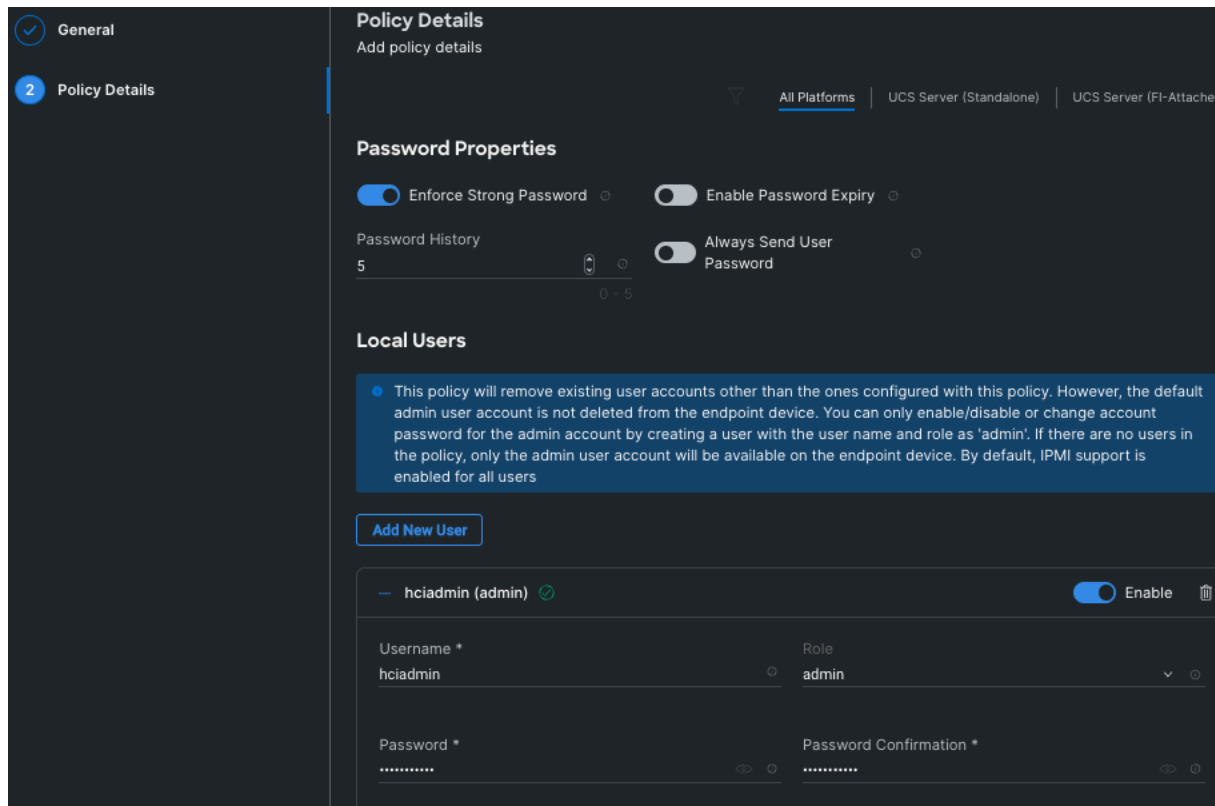
Procedure 5. Create Local User Policy

Step 1. From the Service Selector drop-down list, select **Infrastructure Services** and navigate to **Configure > Policies** and click **Create Policy**.

Step 2. On the Create page for Policies, select **UCS Server > Local User** and click **Start**.

Step 3. On the Local User Create page, enter the Organization, Name, Description and create a new tag or assign an existing tag and click **Next**.

Step 4. On the Policy Details page, Configure Password Properties and Add New User. Click **Create**.



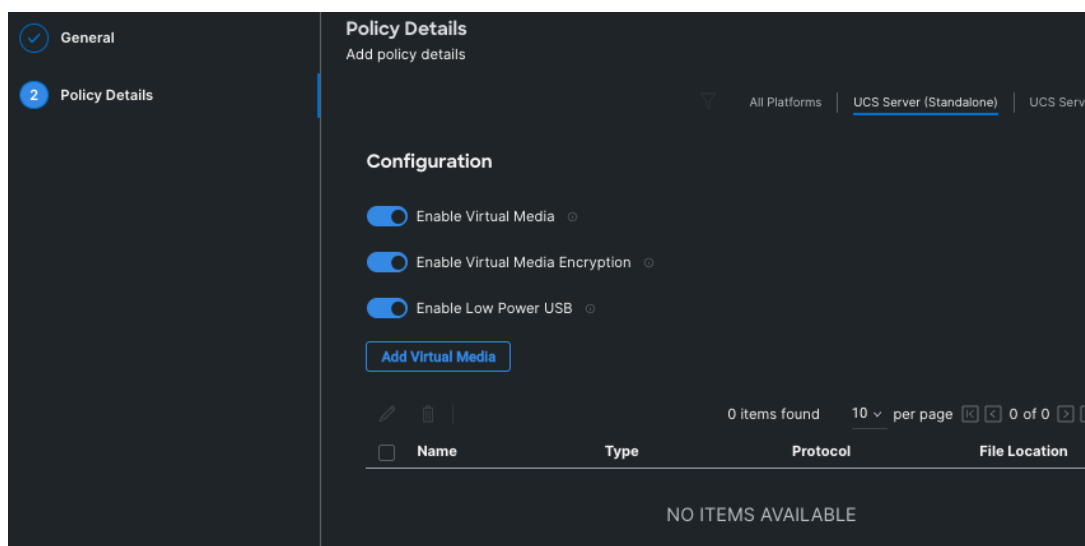
Procedure 6. Create Local User Policy

Step 1. From the Service Selector drop-down list, select **Infrastructure Services** and navigate to **Configure > Policies** and click **Create Policy**.

Step 2. On the Create page for Policies, select **UCS Server > Virtual Media** and click **Start**.

Step 3. On the Virtual Media Create page, enter the Organization, Name, Description and create a new tag or assign an existing tag and click **Next**.

Step 4. On the Policy Details page, Enable Virtual Media and other properties if required. Click **Create**.



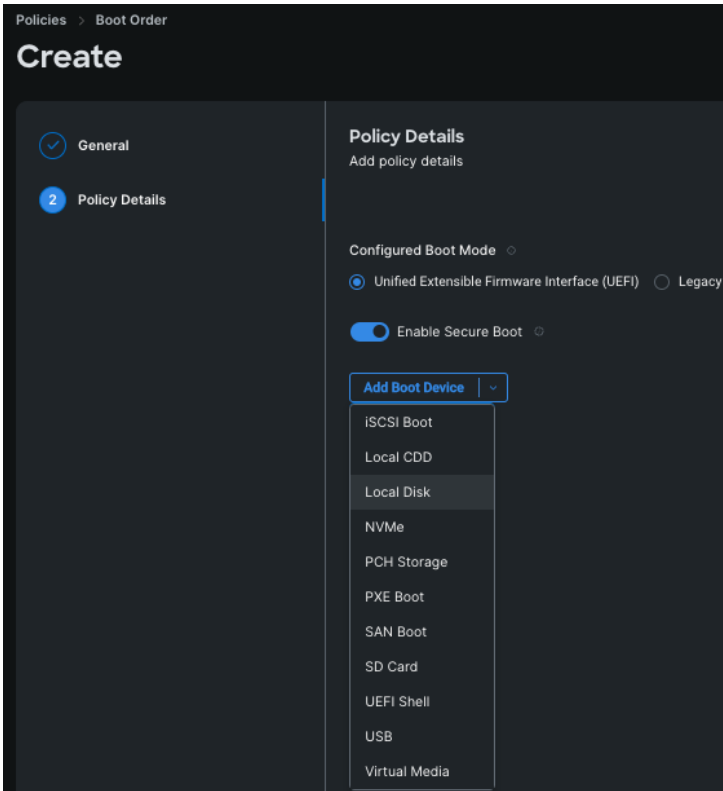
Procedure 7. Create Boot Order Policy

Step 1. From the Service Selector drop-down list, select **Infrastructure Services** and navigate to **Configure > Policies** and click **Create Policy**.

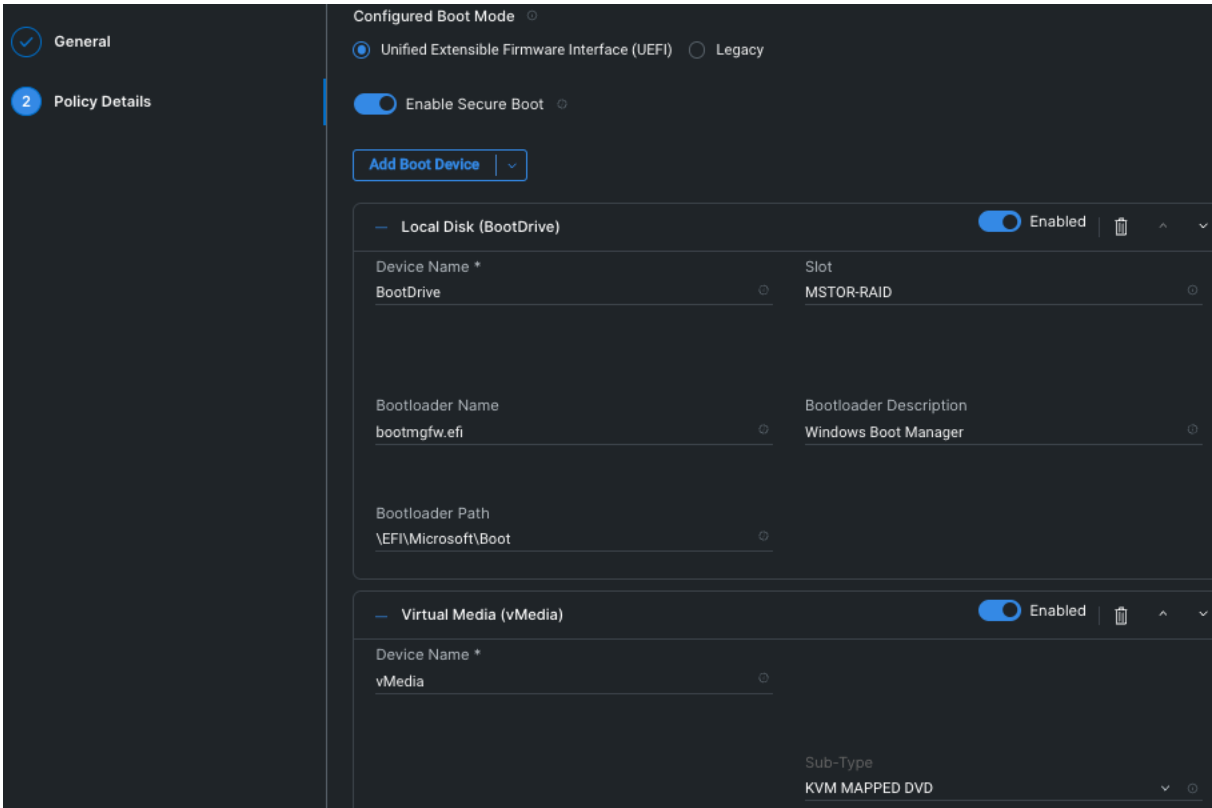
Step 2. On the Create page for Policies, select **UCS Server > Boot Order** and click **Start**.

Step 3. On the Boot Order Create page, enter the Organization, Name, Description and create a new tag or assign an existing tag and click **Next**.

Step 4. On the Policy Details page, **Enable Secure Boot** and from the Add Boot Device drop-down list, select the boot devices.



Step 5. Select Local Disk and Virtual Media and enter the details as shown in the below and click **Create**.



Procedure 8. Create BIOS Policy

Step 1. From the Service Selector drop-down list, select **Infrastructure Services** and navigate to **Configure > Policies** and click **Create Policy**.

Step 2. On the Create page for Policies, select **UCS Server > BIOS** and click **Start**.

Step 3. On the BIOS Create page, enter the Organization, Name, Description and create a new tag or assign an existing tag and click **Next**.

Step 4. On the Policy Details page, configure the tokens as shown in the following images, leaving the rest to defaults:

The screenshot shows the 'Policy Details' page for a BIOS policy. The left sidebar has 'General' selected and 'Policy Details' highlighted. The main content area is titled 'Boot Options' and contains the following configuration items:

Configuration Item	Value
Number of Retries	platform-default
Cool Down Time (sec)	platform-default
Boot Option Retry	platform-default
IPv4 HTTP Support	disabled
IPv4 PXE Support	enabled
IPv6 HTTP Support	disabled
IPv6 PXE Support	disabled
Network Stack	enabled
Onboard SCU Storage Support	platform-default
Onboard SCU Storage SW Stack	platform-default
Power ON Password	disabled
P-SATA Mode	platform-default
SATA Mode	platform-default
VMD Enablement	disabled

The screenshot shows the 'Policy Details' page for a BIOS policy, specifically the 'Intel Directed IO' section. The left sidebar has 'General' selected and 'Policy Details' highlighted. The main content area is titled 'Intel Directed IO' and contains the following configuration items:

Configuration Item	Value
Intel VT for Directed IO	enabled
Intel(R) VT-d Coherency Support	disabled
Intel(R) VT-d Interrupt Remapping	platform-default
Intel(R) VT-d PassThrough DMA Support	platform-default
Intel VTD ATS Support	enabled

General

Policy Details

Main

PCIe Slots CDN Control [ⓘ]
enabled

POST Error Pause [ⓘ]
platform-default

General

Policy Details

Power And Performance

C1 Auto Demotion [ⓘ]
enabled

C1 Auto UnDemotion [ⓘ]
enabled

Core Performance Boost [ⓘ]
platform-default

Global C State Control [ⓘ]
platform-default

L1 Stream HW Prefetcher [ⓘ]
platform-default

L2 Stream HW Prefetcher [ⓘ]
platform-default

Determinism Slider [ⓘ]
platform-default

Efficiency Mode Enable [ⓘ]
platform-default

CPPC [ⓘ]
platform-default

cTDP Control [ⓘ]
platform-default

Enhanced CPU Performance [ⓘ]
Disabled

LLC Dead Line [ⓘ]
enabled

Optimized Power Mode [ⓘ]
platform-default

UPI Link Enablement [ⓘ]
Auto

General

Policy Details

UPI Power Management [ⓘ]
disabled

Virtual NUMA [ⓘ]
disabled

XPT Remote Prefetch [ⓘ]
Auto

Processor

Adjacent Cache Line Prefetcher ⓘ enabled	Altitude ⓘ platform-default
Autonomous Core C State ⓘ disabled	CPU Autonomous C State ⓘ platform-default
Boot Performance Mode ⓘ Max Performance	APBDIS ⓘ platform-default
Downcore Control ⓘ platform-default	Streaming Stores Control ⓘ platform-default
Fixed SOC P-State ⓘ platform-default	DF C-States ⓘ platform-default
CCD Control ⓘ platform-default	CPU Downcore control ⓘ platform-default
CPU SMT Mode ⓘ platform-default	ACPI SRAT L3 Cache As NUMA Domain ⓘ platform-default

Channel Interleaving ⓘ platform-default	Cisco xGMI Max Speed ⓘ platform-default
Closed Loop Thermal Throttling ⓘ platform-default	Processor CMCI ⓘ enabled
Config TDP ⓘ platform-default	Configurable TDP Level ⓘ Normal
Core Multi Processing ⓘ all	Energy Performance ⓘ balanced-performance
Frequency Floor Override ⓘ platform-default	CPU Performance ⓘ custom
Power Technology ⓘ platform-default	Demand Scrub ⓘ platform-default
Direct Cache Access Support ⓘ platform-default	DRAM Clock Throttling ⓘ platform-default
Energy Efficient Turbo ⓘ disabled	Energy Performance Tuning ⓘ platform-default

General

2 Policy Details

Enhanced Intel Speedstep(R) Technology ⓘ	enabled	Processor EPP Enable ⓘ	platform-default
EPP Profile ⓘ	Balanced Performance	Execute Disable Bit ⓘ	platform-default
Local X2 Apic ⓘ	platform-default	Hardware Prefetcher ⓘ	enabled
CPU Hardware Power Management ⓘ	Disabled	IMC Interleaving ⓘ	platform-default
Intel Dynamic Speed Select ⓘ	disabled	Intel HyperThreading Tech ⓘ	enabled
Intel Speed Select ⓘ	Base	Intel Turbo Boost Tech ⓘ	enabled
Intel(R) VT ⓘ	enabled	IIO Error Enable ⓘ	platform-default
DCU IP Prefetcher ⓘ	enabled	KTI Prefetch ⓘ	enabled

General

2 Policy Details

LLC Prefetch ⓘ	enabled	Intel Memory Interleaving ⓘ	platform-default
Package C State Limit ⓘ	C0 C1 State	Patrol Scrub ⓘ	enabled
Patrol Scrub Interval * ⓘ	platform-default	PRMRR Size ⓘ	platform-default
Processor C1E ⓘ	disabled	Processor C3 Report ⓘ	platform-default
Processor C6 Report ⓘ	platform-default	CPU C State ⓘ	platform-default
P-STATE Coordination ⓘ	HW ALL	Power Performance Tuning ⓘ	os
UPI Link Frequency Select ⓘ	Auto	Rank Interleaving ⓘ	platform-default

General

2 Policy Details

Single PCTL ⓘ platform-default

SMT Mode ⓘ platform-default

Sub Numa Clustering ⓘ disabled

DCU Streamer Prefetch ⓘ enabled

SVM Mode ⓘ platform-default

Uncore Frequency Scaling ⓘ enabled

Workload Configuration ⓘ I/O Sensitive

X2APIC Opt-Out Flag ⓘ disabled

XPT Prefetch ⓘ Auto

General

2 Policy Details

Trusted Platform

Limit CPU PA to 46 Bits ⓘ enabled

DMA Control Opt-In Flag ⓘ enabled

Multikey Total Memory Encryption (MK-TME) ⓘ disabled

Software Guard Extensions (SGX) ⓘ disabled

Total Memory Encryption (TME) ⓘ enabled

Select Owner EPOCH Input Type ⓘ Manual User Defined Owner EPOCHs

SGX Auto MP Registration Agent ⓘ disabled

SGX Epoch 0* ⓘ 0

SGX Epoch 1* ⓘ 0

SGX Factory Reset ⓘ disabled

SGX PubKey Hash0* ⓘ 0

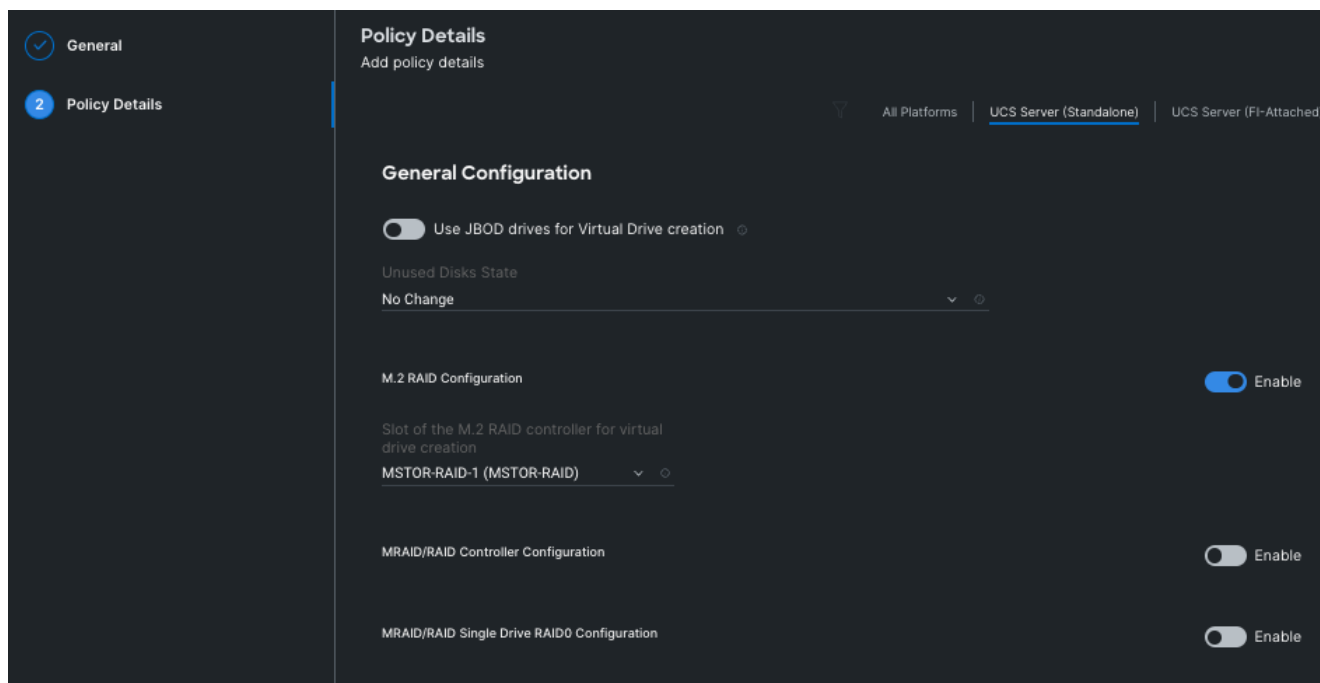
SGX PubKey Hash1* ⓘ 0

The screenshot displays the 'Policy Details' configuration page. On the left, a sidebar shows 'General' and 'Policy Details' (selected). The main area contains the following settings:

- SGX PubKey Hash2: 0
- SGX PubKey Hash3: 0
- SGX Write Enable: enabled
- SGX Package Information In-Band Access: disabled
- SGX QoS: enabled
- SHA-1 PCR Bank: enabled
- SHA256 PCR Bank: enabled
- SHA384 PCR Bank: platform-default
- Trusted Platform Module State: enabled
- TPM Pending Operation: None
- TPM Minimal Physical Presence: enabled
- Security Device Support: enabled
- Intel Trusted Execution Technology Support: enabled

Procedure 9. Create Storage Policy

- Step 1. From the Service Selector drop-down list, select **Infrastructure Services** and navigate to **Configure > Policies** and click **Create Policy**.
- Step 2. On the Create page for Policies, select **UCS Server > Storage** and click **Start**.
- Step 3. On the Storage Create page, enter the Organization, Name, Description and create a new tag or assign an existing tag and click **Next**.
- Step 4. On the Policy Details page, Enable the **M.2 RAID Configuration** and select the **MSTOR RAID-1 (MSTOR RAID)** from the drop-down list as shown in the following figure:



Procedure 10. Create Thermal Policy

Step 1. From the Service Selector drop-down list, select **Infrastructure Services** and navigate to **Configure > Policies** and click **Create Policy**.

Step 2. On the Create page for Policies, select **UCS Chassis > Thermal** and click **Start**.

Step 3. On the Thermal Create General page, enter the Organization, Name, Description and create a new tag or assign an existing tag and click **Next**.

Step 4. On the Thermal Create Policy Details page, click the drop-down list under **Fan Control Mode** and select **Balanced** or higher fan speed.

On the Policy Details page, Enable the **M.2 RAID Configuration** and select the **MSTOR RAID-1 (MSTOR RAID)** from the drop-down list.

Procedure 11. Create UCS Server Profile

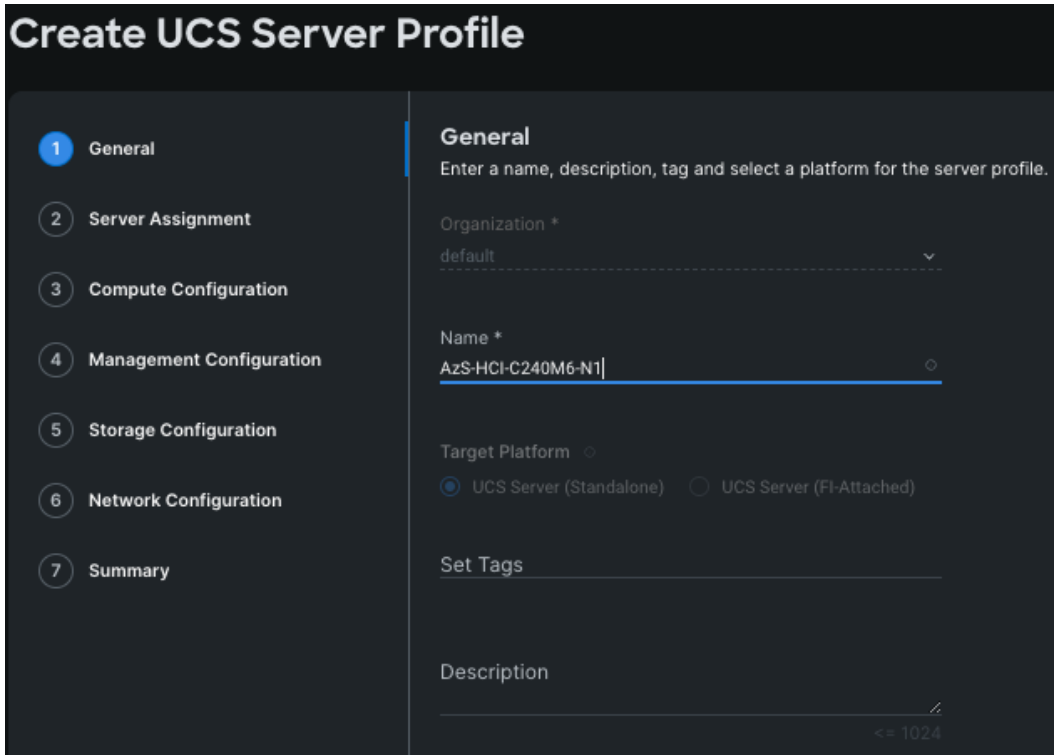
This procedure explains how to create a Cisco UCS server profile, clone it, and deploy servers.

Alternatively, you can create a **server profile template** from which multiple server profiles can be derived and deployed on servers. For more information on server profile templates, go to: https://intersight.com/help/saas/resources/cisco_intersight_managed_mode_configuration#server_profile_templates

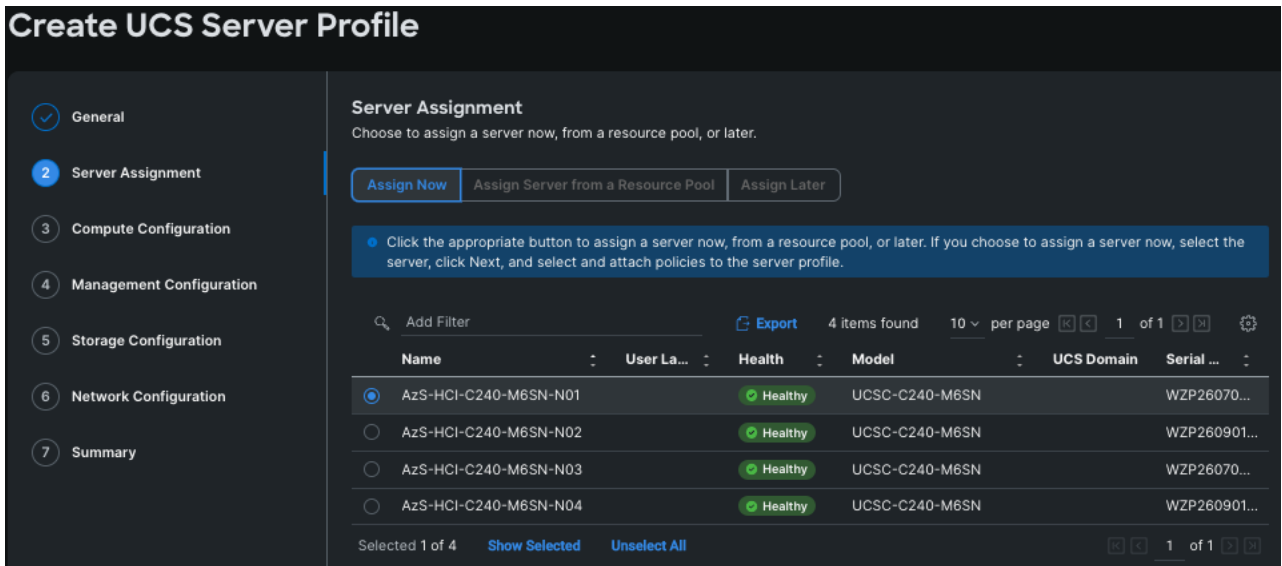
Step 1. From the Service Selector drop-down list, select **Infrastructure Services** and navigate to **Configure > Profiles** and click **Create UCS Server Profile**.

Step 2. On the Create UCS Server Profile page, click **Start**.


Step 3. On the General page, enter the Organization, Name, Description and create a new tag or assign an existing tag. For Target Platform, select **UCS Standalone** under and click **Next**.




Step 4. On the Server Assignment page, click **Assign Now** and select a server from the list shown below:



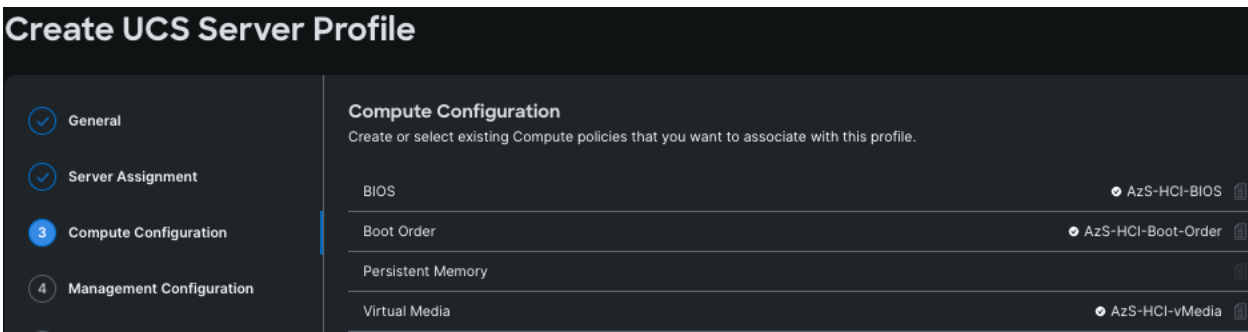
Step 5. On the Compute Configuration page, hover the mouse cursor over right-side of the row next to BIOS and click **Select Policy**.



Step 6. Select the policy created for BIOS in the previous section.



Step 7. Select the respective policies created in the previous sections for Boot Order and Virtual Media as shown below and click **Next**.



Step 8. Repeat steps 1 - 7 and complete the Management, Storage, and Network configuration and click **Next**.

Create UCS Server Profile

The screenshot shows the 'Management Configuration' step of the 'Create UCS Server Profile' wizard. The left sidebar contains a list of steps: General, Server Assignment, Compute Configuration, Management Configuration (highlighted with a blue circle and number 4), Storage Configuration, Network Configuration, and Summary. The main content area is titled 'Management Configuration' and includes the instruction: 'Create or select existing Management policies that you want to associate with this profile.' Below this, there is a list of configuration options, each with a radio button and a document icon:

- Device Connector
- IPMI Over LAN
- LDAP
- Local User (selected with radio button) - AzS-HCI-Local-User
- Network Connectivity (selected with radio button) - AzS-HCI-Network-Connectivity
- NTP (selected with radio button) - AzS-HCI-NTP
- Serial Over LAN
- SMTP
- SNMP
- SSH (selected with radio button) - AzS-HCI-SSH
- Syslog
- Virtual KVM (selected with radio button) - AzS-HCI-vKVM

Create UCS Server Profile

The screenshot shows the 'Storage Configuration' step of the 'Create UCS Server Profile' wizard. The left sidebar contains a list of steps: General, Server Assignment, Compute Configuration, Management Configuration, Storage Configuration (highlighted with a blue circle and number 5), Network Configuration, and Summary. The main content area is titled 'Storage Configuration' and includes the instruction: 'Create or select existing Storage policies that you want to associate with this profile.' Below this, there is a list of configuration options, each with a radio button and a document icon:

- SD Card
- Storage (selected with radio button) - AzS-HCI-Storage

Create UCS Server Profile

The screenshot shows the 'Network Configuration' step of the 'Create UCS Server Profile' wizard. The left sidebar contains a list of steps: General, Server Assignment, Compute Configuration, Management Configuration, Storage Configuration, Network Configuration (highlighted with a blue circle and number 6), and Summary. The main content area is titled 'Network Configuration' and includes the instruction: 'Create or select existing Network Configuration policies that you want to associate with this profile.' Below this, there is a list of configuration options, each with a radio button and a document icon:

- Adapter Configuration
- LAN Connectivity
- SAN Connectivity
- Auto Placement Configuration for vNICs & vHBAs (expanded section)

Under the 'Auto Placement Configuration for vNICs & vHBAs' section, there is a blue bar with the text: 'Graphical representation of vNICs & vHBAs placement is only applicable for Auto Configuration mode.'

Step 9. On the Summary page, verify the configuration and click **Deploy**.

Create UCS Server Profile

- General
- Server Assignment
- Compute Configuration
- Management Configuration
- Storage Configuration
- Network Configuration
- 7 Summary**

Summary

Verify details of the profile and the policies, resolve errors and deploy.

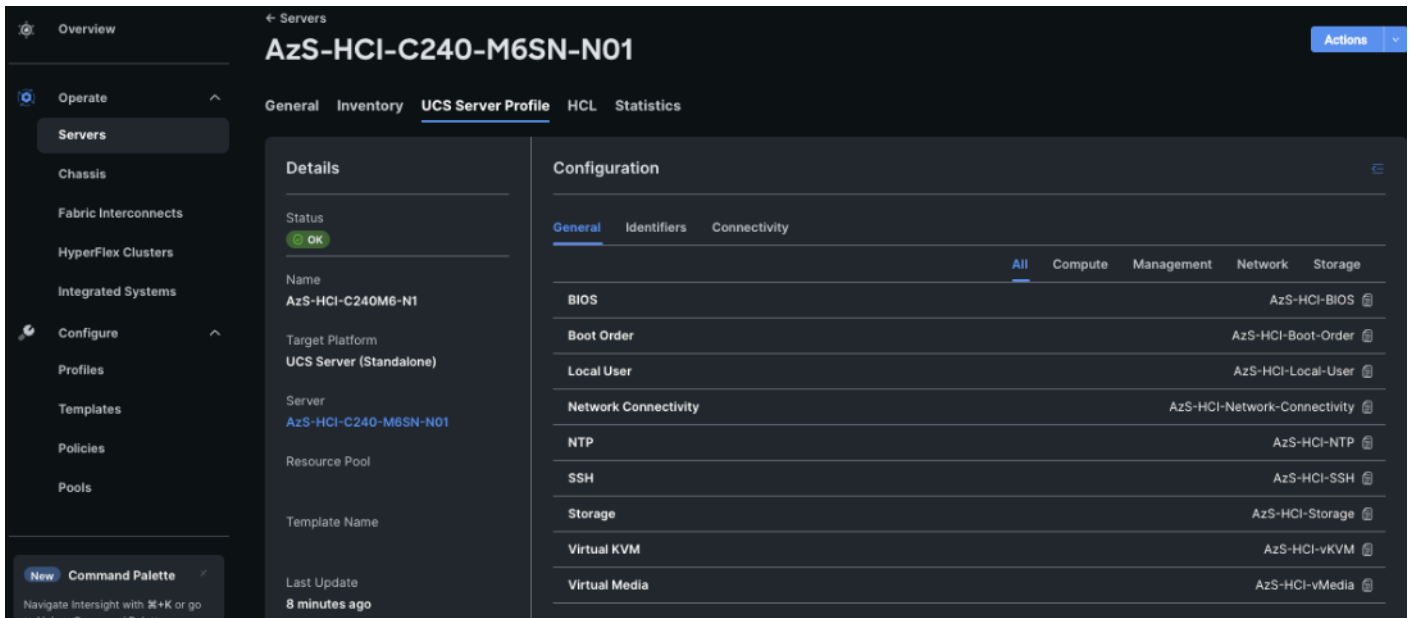

General

Organization	default	Status	Not Deployed
Name	AzS-HCI-C240M6-N1	Management IP	192.168.0.239
Assigned Server	AzS-HCI-C240-M6SN-N01		
Target Platform	UCS Server (Standalone)		


Compute Configuration	Management Configuration	Storage Configuration	Network Configuration	Errors/Warnings (0)
BIOS				AzS-HCI-BIOS
Boot Order				AzS-HCI-Boot-Order
Virtual Media				AzS-HCI-vMedia

[Close](#) [Back](#) [Deploy](#)

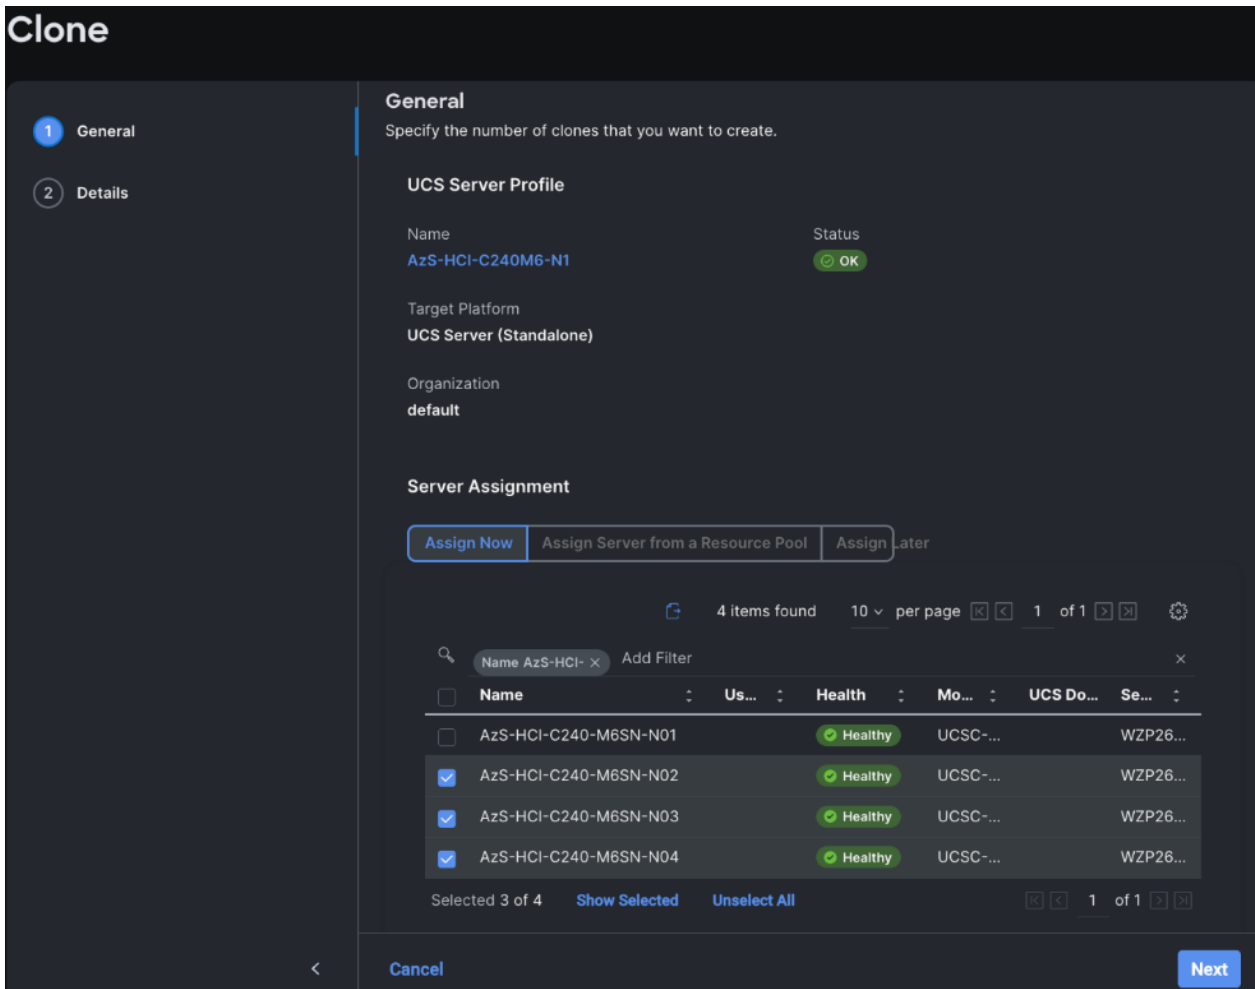
Deployment will take few minutes to complete, and the progress can be seen by clicking the Requests icon next to the Search field. The following figures show the status of the successfully deployed profile from Profile and Servers tab:



Step 10. Clone the Profile created in the previous steps, by clicking the ellipsis and selecting **Clone** as shown below:



Step 11. On the General page, click **Assign Now** and select the remaining unassigned servers and click **Next**.



Step 12. On the Details page, edit the name under **Clone Name Prefix** and the number under the **Start Index for Suffix** as shown below and click **Clone**.

Clone

Details
Edit the description, tags, and auto-generated names of the clones.

General

Organization *
default

Target Platform
UCS Server (Standalone)

Description
<= 1024

Set Tags

Clone Details

Clone Name Prefix	Digits Count	Start Index for Suffix
AzS-HCI-C240M6-N	1	2

1 Clone Name *
AzS-HCI-C240M6-N2

Assigned Server
AzS-HCI-C240-M6SN-N02

2 Clone Name *
AzS-HCI-C240M6-N3

Assigned Server
AzS-HCI-C240-M6SN-N03

3 Clone Name *
AzS-HCI-C240M6-N4

Assigned Server
AzS-HCI-C240-M6SN-N04

Close Back Clone

Step 13. On the Profiles page, select all the newly created profiles with Not Deployed status and click the ellipses. Click **Deploy**.

Profiles

HyperFlex Cluster Profiles UCS Chassis Profiles UCS Domain Profiles UCS Server Profiles

[Create UCS Server Profile](#)

* All UCS Server Prof... +

... Add Filter [Export](#) 4 items found 21 per page 1 of 1

<input type="checkbox"/>	Name	Status	Target Platform	UCS Server Template	Server	Last Update	
<input checked="" type="checkbox"/>	AzS-HCI-C240M6-N4	Not Deployed	UCS Server (Standal...		AzS-HCI-C240-M6S...	a few seconds ago	...
<input checked="" type="checkbox"/>	AzS-HCI-C240M6-N3	Not Deployed	UCS Server (Standal...		AzS-HCI-C240-M6S...	a few seconds ago	...
<input checked="" type="checkbox"/>	AzS-HCI-C240M6-N2	Not Deployed	UCS Server (Standal...		AzS-HCI-C240-M6S...	a few seconds ago	...
<input type="checkbox"/>	AzS-HCI-C240M6-N1	OK	UCS Server (Standal...		AzS-HCI-C240-M6S...	Mar 12, 2023 1:15 AM	...

... Selected 3 of 4 [Show Selected](#) [Unselect All](#) 1 of 1

Deploy

Unassign Server

Step 14. On the Deploy pop-up page, click **More Details** to confirm, and click **Deploy**.

Deploy (3 UCS Server Profiles) ✕

Selected UCS server profiles will be deployed to their assigned servers.

[^ More Details](#)

3 items found 21 per page 1 of 1


Deploy (3 UCS Server Profiles)

Server Profile Name	Server Name
AzS-HCI-C240M6-N4	AzS-HCI-C240-M6SN-N04
AzS-HCI-C240M6-N3	AzS-HCI-C240-M6SN-N03
AzS-HCI-C240M6-N2	AzS-HCI-C240-M6SN-N02

1 of 1

Cancel
Deploy

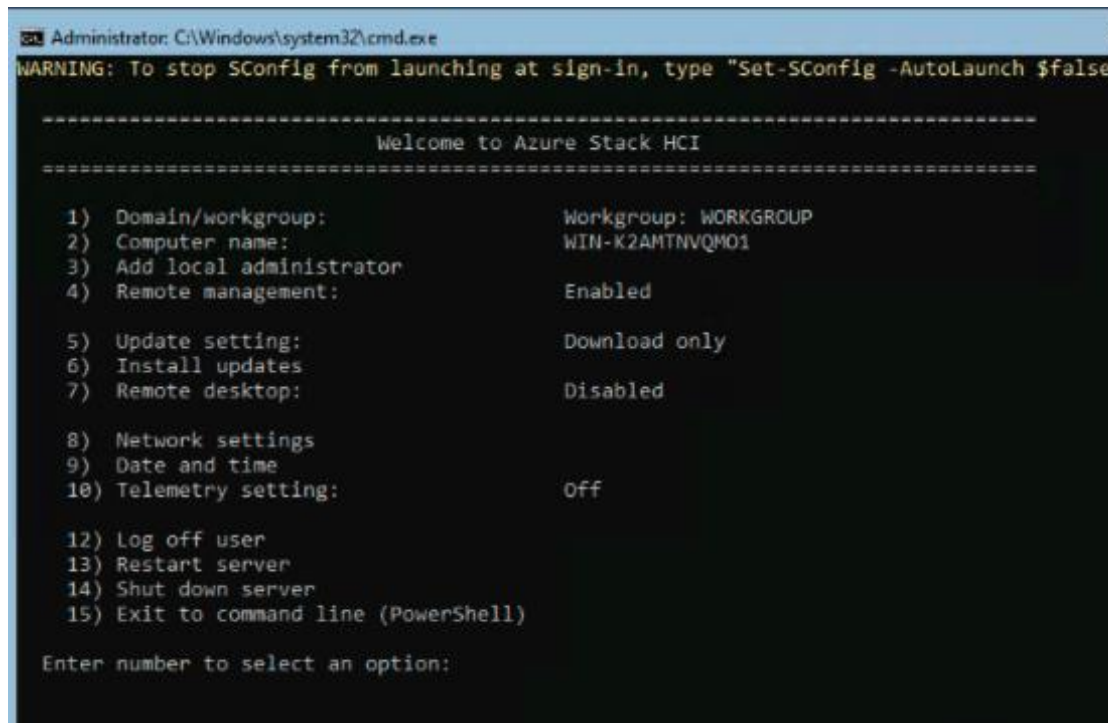
The following image shows the successfully deployed profiles on the assigned servers:



Procedure 12. Launch Server KVM Instance to Install the Operating System

Launch KVM to each server after the service profile association is complete. Install the Azur Stack HCI OS 22H2 using PXE boot or a vMedia mapped installation ISO. PXE boot for OS installation is a better choice because the installation process will run much faster where multiple servers can perform OS installation concurrently.

Step 1. Open a KVM session to each host and perform the following configuration to enable remote access to each host. After logging in, start PowerShell by selecting option 15 (Exit to command line (PowerShell)) in the SConfig screen.



Note: Each host must have a unique host name and IP address for your environment. The following is a table of host names and IP addresses used in this deployment.

Host Name	IP Address
AzS-HCI1-N1	192.168.126.21
AzS-HCI1-N2	192.168.126.22
AzS-HCI1-N3	192.168.126.23
AzS-HCI1-N4	192.168.126.24

Procedure 13. Verify the Operating System Version

Step 1. Run the command `Get-ComputerInfo | fl -Property OSDisplayVersion`:

```
WARNING: To launch Server Configuration tool again, run "SConfig"
PS C:\Users\Administrator> Get-ComputerInfo | fl -Property OSDisplayVersion

OSDisplayVersion : 22H2
```

Procedure 14. Verify Available NICs Seen by the Operating System

Step 1. Run the command `Get-NetAdapter -ifDesc "Mellanox*" | ft -AutoSize`:

```
PS C:\Users\hciadmin> Get-NetAdapter -ifDesc "Mellanox*" | ft -AutoSize

Name                InterfaceDescription          ifIndex Status MacAddress          LinkSpeed
----                -
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #2      8 Up      08-C0-EB-7E-D0-7C  100 Gbps
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter          4 Up      08-C0-EB-7E-D0-7D  100 Gbps
```

Procedure 15. Disable DHCP on Port 2 of the NIC and Verify the Setting

Step 1. Run the commands `Set-NetIPInterface -InterfaceAlias "SlotID 2 Port 2" -Dhcp Disabled` and `Get-NetIPInterface -InterfaceAlias "SlotID 2 Port 2" -Dhcp Disabled -AddressFamily IPv4 | ft -AutoSize`:

```
PS C:\> Get-NetIPInterface -InterfaceAlias "SlotID 2 Port 2" -Dhcp Disabled -AddressFamily IPv4 | ft -AutoSize

ifIndex InterfaceAlias AddressFamily NlMtu(Bytes) InterfaceMetric Dhcp ConnectionState PolicyStore
-----
4        SlotID 2 Port 2 IPv4          1500          10 Disabled Connected ActiveStore
```

Procedure 16. Configure Static NIC IP Address for Management NIC's

Note: Replace the IP address with the address specific to your environment.

Note: The VLAN for this subnet must be set to Native because VLAN tagging is not configured for this physical interface.

Step 1. Run the following command:

```
New-NetIPAddress -InterfaceAlias "SlotID 2 Port 1" -IPAddress 192.168.100.71 -PrefixLength 24  
-DefaultGateway 192.168.100.1
```

```
PS C:\> New-NetIPAddress -InterfaceAlias "SlotID 2 Port 1" -IPAddress 192.168.100.71 -PrefixLength 24 -DefaultGateway 192.168.100.1  
  
IPAddress      : 192.168.100.71  
InterfaceIndex : 5  
InterfaceAlias : SlotID 2 Port 1  
AddressFamily  : IPv4  
Type           : Unicast  
PrefixLength   : 24  
PrefixOrigin   : Manual  
SuffixOrigin   : Manual  
AddressState   : Tentative  
ValidLifetime  : Infinite ([TimeSpan]::MaxValue)  
PreferredLifetime : Infinite ([TimeSpan]::MaxValue)  
SkipAsSource   : False  
PolicyStore    : ActiveStore  
  
IPAddress      : 192.168.100.71  
InterfaceIndex : 5  
InterfaceAlias : SlotID 2 Port 1  
AddressFamily  : IPv4  
Type           : Unicast  
PrefixLength   : 24  
PrefixOrigin   : Manual  
SuffixOrigin   : Manual  
AddressState   : Invalid  
ValidLifetime  : Infinite ([TimeSpan]::MaxValue)  
PreferredLifetime : Infinite ([TimeSpan]::MaxValue)  
SkipAsSource   : False  
PolicyStore    : PersistentStore
```

Procedure 17. Configure DNS Client Server IP Address

Note: Replace the DNS Server IP address with the address specific to your environment.

Step 1. Run the following commands:

```
Set-DnsClientServerAddress -InterfaceAlias "SlotID 2 Port 1" -ServerAddresses  
192.168.0.41,192.168.0.42
```

```
Get-DnsClientServerAddress -InterfaceAlias "SlotID 2 Port 1"
```

```
PS C:\> Set-DnsClientServerAddress -InterfaceAlias "SlotID 2 Port 1" -ServerAddresses 192.168.0.41,192.168.0.42  
PS C:\> Get-DnsClientServerAddress -InterfaceAlias "SlotID 2 Port 1"
```

InterfaceAlias	Interface Index	Address Family	ServerAddresses
SlotID 2 Port 1	5	IPv4	{192.168.0.41, 192.168.0.42}
SlotID 2 Port 1	5	IPv6	{}

Procedure 18. Install Operating System Updates

Step 1. Select option 6 Install Updates from the SConfig Menu.

```
=====
                          Welcome to Azure Stack HCI
=====

1) Domain/workgroup:           Workgroup: WORKGROUP
2) Computer name:             WIN-KA1EKAC2L82
3) Add local administrator
4) Remote management:         Enabled
5) Update setting:            Download only
6) Install updates
7) Remote desktop:            Enabled (all clients)

8) Network settings
9) Date and time
10) Telemetry setting:         Off

12) Log off user
13) Restart server
14) Shut down server
15) Exit to command line (PowerShell)

Enter number to select an option: 6_
```

Step 2. Select option 2 All recommended quality updates only from the Install Updates menu.

Step 3. Select the option A to install all recommended quality updates.

```
=====
                          Install updates
=====

Search for:

1) All quality updates
2) Recommended quality updates only
3) Feature updates

Select an update category (Blank=Cancel): 2
Searching for recommended updates...

Available update(s):
1) 2022-12 Cumulative Update for .NET Framework 3.5 and 4.8 for Microsoft server operating system, version 22H2 for x64 (KB5021084)
2) Security Intelligence Update for Microsoft Defender Antivirus - KB2267602 (Version 1.381.3210.0)
3) 2023-01 Cumulative Update for Microsoft server operating system version 22H2 for x64-based Systems (KB5022291)

Install (A)ll updates, (N)o updates or select a (S)ingle update? (Blank=Cancel): A_
```

The updates will start downloading and installing.

Step 4. Select the option Y to reboot the server if a reboot is required after the update is installed.

```
-----
                        Install updates
-----

Search for:

 1) All quality updates
 2) Recommended quality updates only
 3) Feature updates

Select an update category (Blank=Cancel): 2
Searching for recommended updates...

Available update(s):
 1) 2022-12 Cumulative Update for .NET Framework 3.5 and 4.8 for Microsoft server operating system, version 22H2 for x64 (KB5021084)
 2) Security Intelligence Update for Microsoft Defender Antivirus - KB2267602 (Version 1.381.3210.0)
 3) 2023-01 Cumulative Update for Microsoft server operating system version 22H2 for x64-based Systems (KB5022291)

Install (A)ll updates, (N)o updates or select a (S)ingle update? (Blank=Cancel): A
Downloading update(s)...
Installing update(s)...

Installation results:
 1) Succeeded - 2022-12 Cumulative Update for .NET Framework 3.5 and 4.8 for Microsoft server operating system, version 22H2 for x64 (KB5021084)
 2) Succeeded - Security Intelligence Update for Microsoft Defender Antivirus - KB2267602 (Version 1.381.3210.0)
 3) Succeeded - 2023-01 Cumulative Update for Microsoft server operating system version 22H2 for x64-based Systems (KB5022291)

Summary:
  Installation: Succeeded
  Restart required: True

Restart now? (Y)es or (N)o: 
```

Step 5. After the server reboots, login again and select option 6 Install Updates again from the SConfig Menu.

Step 6. Select option 1 All quality updates from the Install Updates menu.

```
-----
                        Install updates
-----

Search for:

 1) All quality updates
 2) Recommended quality updates only
 3) Feature updates

Select an update category (Blank=Cancel): 1
```

Step 7. Select option A to install all updates if there are any available listed and reboot the server

Step 8. Repeat steps 1 - 7 after the server reboots to install any remaining updates

Note: The Cisco update installation may result in an error condition. This error can safely be ignored.

Step 9. After the server reboots, login again and select option 6 Install Updates again from the SConfig Menu.


```
WARNING: To stop SConfig from launching at sign-in, type "Set-SConfig -AutoLaunch $false"

-----
                        Welcome to Azure Stack HCI
-----

1) Domain/workgroup:                Domain: ucs-spaces.lab
2) Computer name:                   AZSHCI-C1-HOST4
3) Add local administrator
4) Remote management:               Enabled

5) Update setting:                  Manual
6) Install updates
7) Remote desktop:                  Enabled (more secure clients)

8) Network settings
9) Date and time
10) Telemetry setting:              Off

12) Log off user
13) Restart server
14) Shut down server
15) Exit to command line (PowerShell)

Enter number to select an option: 6_
```

Step 10. Select option 1 All quality updates form the Install Updates menu.

```
-----
                        Install updates
-----

Search for:

1) All quality updates
2) Recommended quality updates only
3) Feature updates

Select an update category (Blank=Cancel): 1_
```

Step 11. Verify that no other quality updates are available for installation. Install any remaining quality updates.

```
-----
                        Install updates
-----

Search for:

1) All quality updates
2) Recommended quality updates only
3) Feature updates

Select an update category (Blank=Cancel): 1
Searching for all applicable updates...

Available update(s):
1) Security Intelligence Update for Microsoft Defender Antivirus - KB2267602 (Version 1.359.1319.0)

Install update? (Y)es or (N)o: y_
```

Step 12. Return to the main SConfig menu.

Step 13. Select option 15 Exit to command line (PowerShell) in the SConfig screen.

```
=====
Welcome to Azure Stack HCI
=====
1) Domain/workgroup:           Workgroup: WORKGROUP
2) Computer name:             WIN-DHDTHBRP2BM
3) Add local administrator
4) Remote management:         Enabled
5) Update setting:            Download only
6) Install updates
7) Remote desktop:            Disabled
8) Network settings
9) Date and time
10) Telemetry setting:         Security
12) Log off user
13) Restart server
14) Shut down server
15) Exit to command line (PowerShell)

Enter number to select an option:
```

Procedure 19. Rename Computer

Step 1. Run the command:

```
Rename-Computer -NewName AzS-HCI1-N11 -Restart
```

```
PS C:\Users\hciadmin> Rename-Computer -NewName AzS-HCI1-N1 -Restart _
```

The server restarts after renaming the computer.

Procedure 20. Join the Windows Server to a Domain


Note: Replace the Active Directory Domain name with the domain name and account with domain admin privileges that is specific to your environment. Login with administrative privileges after the server reboot and enter option 15 to start a PowerShell session in the SConfig screen.

Note: The local computer time must be within 5 minutes of the domain controller time in order for the computer to join the active directory domain. The local computer date and time can be checked and adjusted using option 9 “Date and Time” in SConfig or by using the PowerShell Get-Date and Set-Date cmdlet.

Step 1. Run the following command to join the computer to the Active Directory domain:

```
Add-Computer -DomainName ucs-spaces.lab -Credential ucs-spaces.lab\HCIAdmin -Restart
```

```
PS C:\Users\Administrator> Add-Computer -DomainName ucs-spaces.lab -Credential ucs-spaces\hciadmin -Restart _
```



The server restarts after joining the domain.

Note: The following procedures are preformed from a domain joined remote management Host. See the Appendix for [Remote Management Host](#) configuration requirements.

Procedure 21. Configure Windows Memory Crashdump

Note: Hyper-V hosts allocate typically contain a considerable amount of physical memory, but the majority of the physical memory is allocated to virtual machines. For this reason, the parent partition of a Hyper-V host uses a relatively small amount of memory as compared to the total amount of memory installed in the system. The memory dump of the parent partition can provide vital debugging information in the rare case that an unexpected bugcheck (bluescreen) occurs on host.

The following setting enables the creation of a memory dump file and when a bugcheck occurs and use the Active Dump setting to optimize the amount of memory used when a memory dump is created:

```
$Creds = Get-Credential -Message "Enter Login Credentials" -User ucs-spaces\hciadmin
$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Cofiguring Memory Crashdump Registry settings " -ForegroundColor Yellow
Set-ItemProperty -Path HKLM:\System\CurrentControlSet\Control\CrashControl -Name CrashDumpEnabled -value 1
Set-ItemProperty -Path HKLM:\System\CurrentControlSet\Control\CrashControl -Name FilterPages -value 1
Get-ItemProperty -Path HKLM:\System\CurrentControlSet\Control\CrashControl -Name CrashDumpEnabled
Get-ItemProperty -Path HKLM:\System\CurrentControlSet\Control\CrashControl -Name FilterPages

}
}
```

```
}
```

```
Host Name: AZS-HCI1-N1
Configuring Memory Crashdump Registry settings

CrashDumpEnabled : 1
PSPath           : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\CrashControl
PSParentPath     : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control
PSChildName      : CrashControl
PSDrive          : HKLM
PSProvider       : Microsoft.PowerShell.Core\Registry
PSComputerName   : AZS-HCI1-N1
RunspaceId      : b85cbe6c-1a47-4d6a-a817-6613fadbc757

FilterPages      : 1
PSPath           : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\CrashControl
PSParentPath     : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control
PSChildName      : CrashControl
PSDrive          : HKLM
PSProvider       : Microsoft.PowerShell.Core\Registry
PSComputerName   : AZS-HCI1-N1
RunspaceId      : b85cbe6c-1a47-4d6a-a817-6613fadbc757

Host Name: AZS-HCI1-N2
Configuring Memory Crashdump Registry settings

CrashDumpEnabled : 1
PSPath           : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\CrashControl
PSParentPath     : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control
PSChildName      : CrashControl
PSDrive          : HKLM
PSProvider       : Microsoft.PowerShell.Core\Registry
PSComputerName   : AZS-HCI1-N2
RunspaceId      : fae15fb8-0210-4fea-b3c0-a4aaa9b06da5

FilterPages      : 1
PSPath           : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\CrashControl
PSParentPath     : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control
PSChildName      : CrashControl
PSDrive          : HKLM
PSProvider       : Microsoft.PowerShell.Core\Registry
PSComputerName   : AZS-HCI1-N2
RunspaceId      : fae15fb8-0210-4fea-b3c0-a4aaa9b06da5
```

```

Host Name: AzS-HCI1-N3
Configuring Memory Crashdump Registry settings
CrashDumpEnabled : 1
PSPath           : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\CrashControl
PSParentPath     : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control
PSChildName     : CrashControl
PSDrive         : HKLM
PSProvider      : Microsoft.PowerShell.Core\Registry
PSComputerName  : AzS-HCI1-N3
RunspaceId     : b0e513f9-a2a1-4576-8f54-a7d526ef385b

FilterPages     : 1
PSPath           : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\CrashControl
PSParentPath     : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control
PSChildName     : CrashControl
PSDrive         : HKLM
PSProvider      : Microsoft.PowerShell.Core\Registry
PSComputerName  : AzS-HCI1-N3
RunspaceId     : b0e513f9-a2a1-4576-8f54-a7d526ef385b

Host Name: AzS-HCI1-N4
Configuring Memory Crashdump Registry settings
CrashDumpEnabled : 1
PSPath           : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\CrashControl
PSParentPath     : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control
PSChildName     : CrashControl
PSDrive         : HKLM
PSProvider      : Microsoft.PowerShell.Core\Registry
PSComputerName  : AzS-HCI1-N4
RunspaceId     : 395ff0ce-e868-45d3-8abe-9651a71385cd

FilterPages     : 1
PSPath           : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\CrashControl
PSParentPath     : Microsoft.PowerShell.Core\Registry::HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control
PSChildName     : CrashControl
PSDrive         : HKLM
PSProvider      : Microsoft.PowerShell.Core\Registry
PSComputerName  : AzS-HCI1-N4
RunspaceId     : 395ff0ce-e868-45d3-8abe-9651a71385cd

```

Procedure 22. Configure Time Zone

Step 1. Time zone must have the same setting on all cluster nodes. The following script block configures the time zone:

```

$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -ScriptBlock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host "Configuring time zone..." -ForegroundColor Yellow
Set-Timezone -Name "Pacific Standard Time"

}
}

```

Note: The time zone is specific to the region. The following command lists available time zones.

```
Get-TimeZone -ListAvailable | ft StandardName, ID
```

Procedure 23. Enable Remote Desktop Access on the Host Servers

Step 1. Run the following:

```
$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
```

```
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -ScriptBlock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Enabling Remote Desktop access..." -ForegroundColor Yellow

Set-ItemProperty -Path "HKLM:\System\CurrentControlSet\Control\Terminal Server" -Name "fDenyTSConnections" -
Value 0

Enable-NetFirewallRule -DisplayGroup "Remote Desktop"

}

}
```

```
Host Name: AZS-HCI1-N1
Enabling Remote Desktop access...
Host Name: AZS-HCI1-N2
Enabling Remote Desktop access...
Host Name: AZS-HCI1-N3
Enabling Remote Desktop access...
Host Name: AZS-HCI1-N4
Enabling Remote Desktop access...
```

Procedure 24. Install Windows Features

The following Windows features are installed:

- Bitlocker
- Data Center Bridging
- Failover Clustering
- Hyper-V
- Hyper-V PowerShell
- Active Directory Remote Management PowerShell
- Cluster Management PowerShell
- File Server
- SMB Bandwidth Limit
- NetworkATC
- NetworkHUD
- FS-Data-Deduplication

```
$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {
```

```

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Enabling Required Windows Features and Restarting Host Server..." -ForegroundColor Yellow

Add-WindowsFeature -Name Hyper-V,Failover-Clustering,Data-Center-Bridging,Bitlocker,FS-FileServer, FS-SMBBW,
Hyper-V-PowerShell,RSAT-AD-Powershell,RSAT-Clustering-PowerShell,NetworkATC,NetworkHUD,FS-DATA-Deduplication
-IncludeAllSubFeature -IncludeManagementTools -Restart

}
}

```

```

Host Name: AZS-HCI1-N1
Enabling Required Windows Features and Restarting Host Server...

PSCoMpuTerName : AzS-HCI1-N1
RunspaceId     : dbae6316-0f35-4b72-8ec7-6e868097aaeb
Success        : True
RestartNeeded  : Yes
FeatureResult  : {BitLocker Drive Encryption, Data Center Bridging, Enhanced Storage, Failover Clustering...}
ExitCode       : SuccessRestartRequired

WARNING: You must restart this server to finish the installation process.
Host Name: AZS-HCI1-N2
Enabling Required Windows Features and Restarting Host Server...
PSCoMpuTerName : AzS-HCI1-N2
RunspaceId     : b4737863-7584-4f13-b225-1e8cc3f72119
Success        : True
RestartNeeded  : Yes
FeatureResult  : {BitLocker Drive Encryption, Data Center Bridging, Enhanced Storage, Failover Clustering...}
ExitCode       : SuccessRestartRequired

WARNING: You must restart this server to finish the installation process.
Host Name: AZS-HCI1-N3
Enabling Required Windows Features and Restarting Host Server...
PSCoMpuTerName : AzS-HCI1-N3
RunspaceId     : ea0abbbf-3704-4578-91fa-625f78eeb101
Success        : True
RestartNeeded  : Yes
FeatureResult  : {BitLocker Drive Encryption, Data Center Bridging, Enhanced Storage, Failover Clustering...}
ExitCode       : SuccessRestartRequired

WARNING: You must restart this server to finish the installation process.
Host Name: AZS-HCI1-N4
Enabling Required Windows Features and Restarting Host Server...
PSCoMpuTerName : AzS-HCI1-N4
RunspaceId     : 91a923f1-0ed8-4847-9ef5-0753b916983f
Success        : True
RestartNeeded  : Yes
FeatureResult  : {BitLocker Drive Encryption, Data Center Bridging, Enhanced Storage, Failover Clustering...}
ExitCode       : SuccessRestartRequired

WARNING: You must restart this server to finish the installation process.

```

Note: Each server node will reboot automatically to complete the feature installation process. Confirm that each server reboots successfully.

Procedure 25. Verify installed Windows Features

Step 1. Run the following:

```

$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

```

```

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verifying Required Windows Features..." -ForegroundColor Yellow

Get-WindowsFeature -Name Hyper-V,Failover-Clustering,Data-Center-Bridging,Bitlocker,FS-FileServer, FS-
SMBBW,Hyper-V-PowerShell,RSAT-AD-Powershell,RSAT-Clustering-PowerShell,NetworkATC,NetworkHUD,FS-DATA-
Deduplication | ft -AutoSize

}
}

```

```

Host Name: AZS-HCI1-N1
Verifying Required Windows Features...

Display Name                                     Name                                     Install State
-----
[X] File Server                                 FS-FileServer                           Installed
[X] Data Deduplication                         FS-Data-Deduplication                   Installed
[X] Hyper-V                                    Hyper-V                                  Installed
[X] BitLocker Drive Encryption                 BitLocker                                Installed
[X] Data Center Bridging                      Data-Center-Bridging                    Installed
[X] Failover Clustering                       Failover-Clustering                     Installed
[X] Network ATC                               NetworkATC                               Installed
[X] Network HUD                               NetworkHUD                               Installed
[X] Failover Cluster Module for Windows PowerShell RSAT-Clustering-PowerShell             Installed
[X] Active Directory module for Windows PowerShell RSAT-AD-PowerShell                     Installed
[X] Hyper-V Module for Windows PowerShell      Hyper-V-PowerShell                       Installed
[X] SMB Bandwidth Limit                       FS-SMBBW                                 Installed

Host Name: AZS-HCI1-N2
Verifying Required Windows Features...

Display Name                                     Name                                     Install State
-----
[X] File Server                                 FS-FileServer                           Installed
[X] Data Deduplication                         FS-Data-Deduplication                   Installed
[X] Hyper-V                                    Hyper-V                                  Installed
[X] BitLocker Drive Encryption                 BitLocker                                Installed
[X] Data Center Bridging                      Data-Center-Bridging                    Installed
[X] Failover Clustering                       Failover-Clustering                     Installed
[X] Network ATC                               NetworkATC                               Installed
[X] Network HUD                               NetworkHUD                               Installed
[X] Failover Cluster Module for Windows PowerShell RSAT-Clustering-PowerShell             Installed
[X] Active Directory module for Windows PowerShell RSAT-AD-PowerShell                     Installed
[X] Hyper-V Module for Windows PowerShell      Hyper-V-PowerShell                       Installed
[X] SMB Bandwidth Limit                       FS-SMBBW                                 Installed

```



```

Host Name: AZS-HCI1-N3
Verifying Required Windows Features...

Display Name                                     Name                                     Install State
-----
      [X] File Server                             FS-FileServer                           Installed
      [X] Data Deduplication                     FS-Data-Deduplication                   Installed
[X] Hyper-V                                     Hyper-V                                 Installed
[X] BitLocker Drive Encryption                 BitLocker                               Installed
[X] Data Center Bridging                       Data-Center-Bridging                    Installed
[X] Failover Clustering                       Failover-Clustering                     Installed
[X] Network ATC                               NetworkATC                              Installed
[X] Network HUD                               NetworkHUD                               Installed
      [X] Failover Cluster Module for Windows PowerShell RSAT-Clustering-PowerShell             Installed
      [X] Active Directory module for Windows PowerShell RSAT-AD-PowerShell                     Installed
      [X] Hyper-V Module for Windows PowerShell   Hyper-V-PowerShell                      Installed
[X] SMB Bandwidth Limit                       FS-SMBBW                                Installed

Host Name: AZS-HCI1-N4
Verifying Required Windows Features...

Display Name                                     Name                                     Install State
-----
      [X] File Server                             FS-FileServer                           Installed
      [X] Data Deduplication                     FS-Data-Deduplication                   Installed
[X] Hyper-V                                     Hyper-V                                 Installed
[X] BitLocker Drive Encryption                 BitLocker                               Installed
[X] Data Center Bridging                       Data-Center-Bridging                    Installed
[X] Failover Clustering                       Failover-Clustering                     Installed
[X] Network ATC                               NetworkATC                              Installed
[X] Network HUD                               NetworkHUD                               Installed
      [X] Failover Cluster Module for Windows PowerShell RSAT-Clustering-PowerShell             Installed
      [X] Active Directory module for Windows PowerShell RSAT-AD-PowerShell                     Installed
      [X] Hyper-V Module for Windows PowerShell   Hyper-V-PowerShell                      Installed
[X] SMB Bandwidth Limit                       FS-SMBBW                                Installed

```

Configure Bitlocker for System Volume

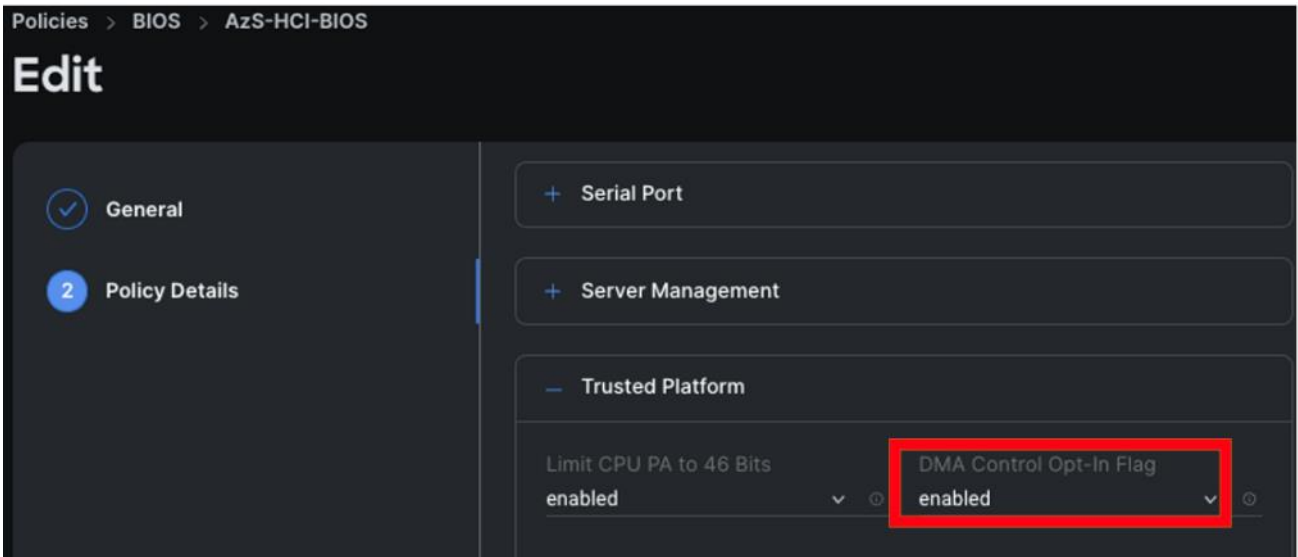
Using Bitlocker to encrypt system volume is an optional procedure in the deployment. TPM will be the primary key protector for the encrypted volume. The TPM will automatically decrypt the system volume at boot time. A recovery password will be an additional key protector in case the TPM fails. The recovery password will be backed up and stored in Active Directory Domain Service. Refer to the Appendix section to configure [Bitlocker](#).

Configure Secured-Core on Hosts

Cisco UCS C240 M6 and M7 validated node is designed for Secured-Core Server, which allows your Azure Stack HCI investment to run workloads on a highly secure infrastructure with unparalleled levels of host security enabled with TPM 2.0, Secure boot, virtualization based security (VBS), boot DMA guard, and DRTPM protection. This section explains the how-to build an infrastructure for the Secured-Core Server on Azure Stack HCI.

Procedure 1. Verify BIOS Setting

Step 1. Make sure the BIOS token **DMA Control Opt-In Flag** is enabled in the BIOS policy.



Note: The other required BIOS tokens for Secured-Core is enabled by default.

Procedure 2. Verify that Secure Boot is Enabled

Step 1. Run the following:

```
$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME
Write-Host "Checking Secure Boot Status " -ForegroundColor Yellow

Confirm-SecureBootUEFI

}
}
```

```
Host Name: AZS-HCI1-N1
Checking Secure Boot Status
True
Host Name: AZS-HCI1-N2
Checking Secure Boot Status
True
Host Name: AZS-HCI1-N3
Checking Secure Boot Status
True
Host Name: AZS-HCI1-N4
Checking Secure Boot Status
True
```

Procedure 3. Configure Secure Core

Step 1. Run the following to restart servers:

```
$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {
Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Enable Secure Core" $env:COMPUTERNAME -ForegroundColor Green
reg add "HKLM\SYSTEM\CurrentControlSet\Control\DeviceGuard\Scenarios\HypervisorEnforcedCodeIntegrity" /v
"Enabled" /t REG_DWORD /d 1 /f
reg add "HKLM\SYSTEM\CurrentControlSet\Control\DeviceGuard\Scenarios\HypervisorEnforcedCodeIntegrity" /v
"WasEnabledBy" /t REG_DWORD /d 0 /f
reg add "HKLM\SYSTEM\CurrentControlSet\Control\DeviceGuard\Scenarios\SystemGuard" /v "Enabled" /t REG_DWORD
/d 1 /f
}
}
```

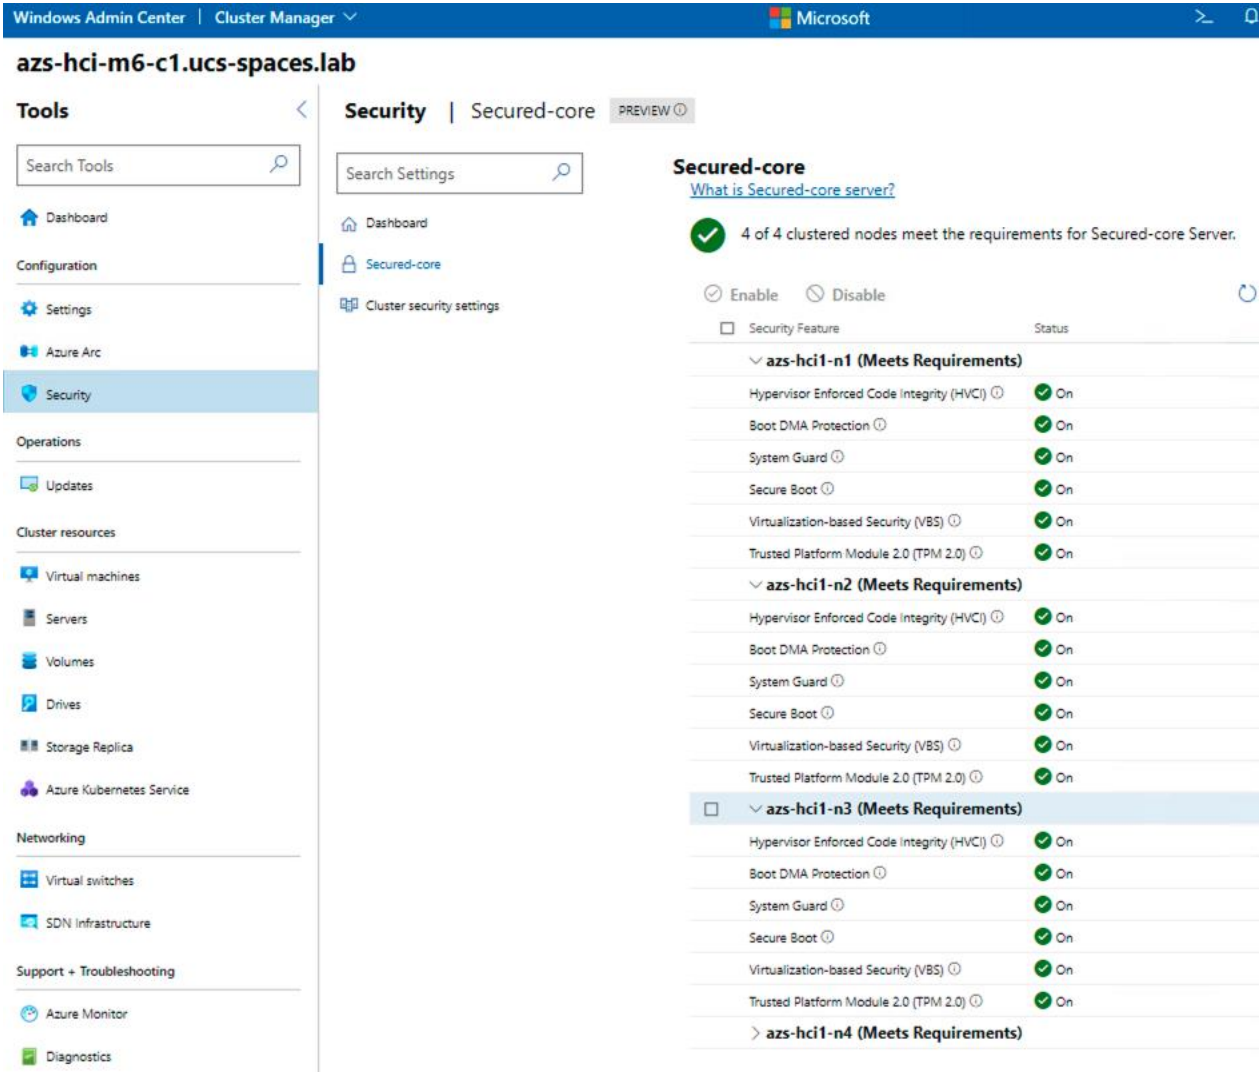
```
Enable Secure Core AZS-HCI1-N1
The operation completed successfully.
The operation completed successfully.
The operation completed successfully.
Enable Secure Core AZS-HCI1-N2
The operation completed successfully.
The operation completed successfully.
The operation completed successfully.
Enable Secure Core AZS-HCI1-N3
The operation completed successfully.
The operation completed successfully.
The operation completed successfully.
Enable Secure Core AZS-HCI1-N4
The operation completed successfully.
The operation completed successfully.
The operation completed successfully.
```

Step 2. Verify the Secured-core configuration:

- Launch **msinfo32** from the command prompt and confirm the following values:

Secure Boot State	On
Kernel DMA Protection	On
Virtualization-based security	Running
Virtualization-based security Required Security Properties	Base Virtualization Support
Virtualization-based security Available Security Properties	Base Virtualization Support, Secure Boot, DMA Protection,
Virtualization-based security Services Configured	Hypervisor enforced Code Integrity, Secure Launch
Virtualization-based security Services Running	Hypervisor enforced Code Integrity, Secure Launch

- Or download and install [Windows Admin Center](#) (WAC). Add target cluster (or server) for management in WAC and from the Server/Cluster Manager view, select Security (left-pane) under Configuration and click **Secured-core**:



Fully Converged Network Topology Configuration

Note: Skip the steps in this section if the deployment model chosen for your environment is a non-converged topology. Only one deployment model needs to be followed; either fully converged or non-converged or storage switchless topology.

This section contains the following procedures to create and deploy standalone Network ATC intent for a fully converged topology. Here you create a single intent for “Management, Compute and Storage” networks.

- [Identify Physical Network Card Port Names](#)
- [Create and Deploy Standalone Network ATC Intent](#)
- [Verify Network ATC Intent Status](#)

- [Verify Virtual Switch and Virtual NIC Creation in the Parent Partition](#)
- [Verify SET Switch Team Load Balancing Algorithm](#)
- [Configure Default Route Metric for Management NIC in Parent Partition](#)
- [Configure Static NIC IP Address for Storage NICs](#)
- [Verify NIC IP Address for Storage NICs](#)
- [Verify DNS Registration is Removed for Storage Interfaces](#)
- [Enable Preserving 802.1p Priority Marking to Pass Through the vSwitch](#)
- [Verify the Storage vNIC VLANs](#)
- [Verify Network Adapters](#)
- [Verify RDMA and RoCEv2 Protocol is Enabled on Physical NICs](#)
- [Verify that RDMA is Enabled on the Storage vNIC Adapters](#)
- [Verify the Mapping of each SMB-Direct NIC to the respective Fabric](#)
- [Verify RDMA Capabilities](#)

Procedure 1. Identify Physical Network Card Port Names

Step 1. Run the following:

```
$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Retrieving physical NIC port names " -ForegroundColor Yellow

Get-netadapter | ft Name, InterfaceDescription, Status, MacAddress, LinkSpeed

}

}
```

```

Host Name: AZS-HCI1-N1
Retrieving physical NIC port names

Name                InterfaceDescription                Status MacAddress                LinkSpeed
-----
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #2 Up      08-C0-EB-7E-D0-C4 100 Gbps
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter Up      08-C0-EB-7E-D0-C5 100 Gbps

Host Name: AZS-HCI1-N2
Retrieving physical NIC port names

Name                InterfaceDescription                Status MacAddress                LinkSpeed
-----
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #2 Up      08-C0-EB-7E-D2-ED 100 Gbps
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter Up      08-C0-EB-7E-D2-EC 100 Gbps

Host Name: AZS-HCI1-N3
Retrieving physical NIC port names

Name                InterfaceDescription                Status MacAddress                LinkSpeed
-----
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #2 Up      08-C0-EB-7E-D0-BC 100 Gbps
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter Up      08-C0-EB-7E-D0-BD 100 Gbps

Host Name: AZS-HCI1-N4
Retrieving physical NIC port names

Name                InterfaceDescription                Status MacAddress                LinkSpeed
-----
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter Up      08-C0-EB-7E-D3-2C 100 Gbps
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #2 Up      08-C0-EB-7E-D3-2D 100 Gbps

```

Note: If the NIC port names are “Ethernet” and “Ethernet 2”, CDN is not enabled. CDN (Consistent Device Naming) must be enabled for correct physical to virtual NIC mapping later in this guide.

Procedure 2. Create and Deploy Standalone Network ATC Intent

Step 1. Run the following script block to create a virtual switch with SET enabled and three virtual NICs:

```

$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Create and Deploy Standalone Network ATC Intent " -ForegroundColor Yellow

$QoSOverride = New-NetIntentQoSPropertyOverrides
$AdapterOverride = New-NetIntentAdapterPropertyOverrides
$StorageOverride = new-NetIntentStorageOverrides

$QoSOverride.PriorityValue8021Action_SMB = 4

```

```

$QoSOverride.PriorityValue8021Action_Cluster = 5
$AdapterOverride.NetworkDirectTechnology = 4
$storageOverride.EnableAutomaticIPGeneration = $false

$QoSOverride
$AdapterOverride
$storageOverride

Add-NetIntent -AdapterName "SlotID 2 Port 1", "SlotID 2 Port 2" -Management -Compute -Storage -StorageVlans
107, 207 -QoSPolicyOverrides $QoSOverride -AdapterPropertyOverrides $AdapterOverride -StorageOverrides
$storageoverride -Name Mgmt_Compute_Storage

}
}

```

Note: For proxy configuration, go to: <https://learn.microsoft.com/en-us/azure-stack/hci/manage/manage-network-atc?tabs=21H2>

Procedure 3. Verify Network ATC Intent Status

Step 1. Run the following:

```

$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Checking Network ATC Intent Status" -ForegroundColor Yellow

Get-netIntentStatus -ComputerName $node | ft
Host,IntentName,ConfigurationStatus,ProvisioningStatus,IsComputeIntentSet,IsManagementIntentSet,IsStorageInte
ntset,IsStretchIntentSet

}
}

```

```

Host Name: AZS-HCI1-N1
Checking Network ATC Intent Status
WARNING: If you are running Windows PowerShell remotely, note that some failover clustering cmdlets do not work remotely. When possible, run the cmdlet locally and specify a remote computer as the target. To run the cmdlet remotely, try using the Credential Security Service Provider (CredSSP). All additional errors or warnings from this cmdlet might be caused by running it remotely.

Host      IntentName      ConfigurationStatus ProvisioningStatus IsComputeIntentSet IsManagementIntentSet IsStorageIntentSet IsStretchIntentSet
-----
AZS-HCI1-N1 mgmt_compute_storage Success           Completed           True                True                True                False

Host Name: AZS-HCI1-N2
Checking Network ATC Intent Status
WARNING: If you are running Windows PowerShell remotely, note that some failover clustering cmdlets do not work remotely. When possible, run the cmdlet locally and specify a remote computer as the target. To run the cmdlet remotely, try using the Credential Security Service Provider (CredSSP). All additional errors or warnings from this cmdlet might be caused by running it remotely.

Host      IntentName      ConfigurationStatus ProvisioningStatus IsComputeIntentSet IsManagementIntentSet IsStorageIntentSet IsStretchIntentSet
-----
AZS-HCI1-N2 mgmt_compute_storage Success           Completed           True                True                True                False

Host Name: AZS-HCI1-N3
Checking Network ATC Intent Status
WARNING: If you are running Windows PowerShell remotely, note that some failover clustering cmdlets do not work remotely. When possible, run the cmdlet locally and specify a remote computer as the target. To run the cmdlet remotely, try using the Credential Security Service Provider (CredSSP). All additional errors or warnings from this cmdlet might be caused by running it remotely.

Host      IntentName      ConfigurationStatus ProvisioningStatus IsComputeIntentSet IsManagementIntentSet IsStorageIntentSet IsStretchIntentSet
-----
AZS-HCI1-N3 mgmt_compute_storage Success           Completed           True                True                True                False

Host Name: AZS-HCI1-N4
Checking Network ATC Intent Status
WARNING: If you are running Windows PowerShell remotely, note that some failover clustering cmdlets do not work remotely. When possible, run the cmdlet locally and specify a remote computer as the target. To run the cmdlet remotely, try using the Credential Security Service Provider (CredSSP). All additional errors or warnings from this cmdlet might be caused by running it remotely.

Host      IntentName      ConfigurationStatus ProvisioningStatus IsComputeIntentSet IsManagementIntentSet IsStorageIntentSet IsStretchIntentSet
-----
AZS-HCI1-N4 mgmt_compute_storage Success           Completed           True                True                True                False

```

Procedure 4. Verify Virtual Switch and Virtual NIC Creation in the Parent Partition

Step 1. Run the following:

```

$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Verifying Virtual Switch " -ForegroundColor Yellow
Get-VMSwitch | fl Name, SwitchType, NetAdapterInterfaceDescription, NetAdapterInterfaceDescriptions

Write-Host " Verifying Management vNIC in parent partition " -ForegroundColor Yellow
Get-netadapter | ft Name, InterfaceDescription, Status, MacAddress, LinkSpeed

}
}

```



```

Host Name: AZS-HCI1-N1
Verifying Virtual Switch

Name : ConvergedSwitch(mgmt_compute_storage)
SwitchType : External
NetAdapterInterfaceDescription : Teamed-Interface
NetAdapterInterfaceDescriptions : {Mellanox ConnectX-6 Dx Adapter, Mellanox ConnectX-6 Dx Adapter #2}

Verifying Management vNIC in parent partition

Name InterfaceDescription Status MacAddress LinkSpeed
----
vManagement(mgmt_compute_storage) Hyper-V Virtual Ethernet Adapter Up 08-C0-EB-7E-D0-C5 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 2) Hyper-V Virtual Ethernet Adapter #3 Up 00-15-5D-7E-15-01 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 1) Hyper-V Virtual Ethernet Adapter #2 Up 00-15-5D-7E-15-00 100 Gbps
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #2 Up 08-C0-EB-7E-D0-C4 100 Gbps
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter Up 08-C0-EB-7E-D0-C5 100 Gbps

```

```

Host Name: AZS-HCI1-N2
Verifying Virtual Switch

Name : ConvergedSwitch(mgmt_compute_storage)
SwitchType : External
NetAdapterInterfaceDescription : Teamed-Interface
NetAdapterInterfaceDescriptions : {Mellanox ConnectX-6 Dx Adapter #2, Mellanox ConnectX-6 Dx Adapter}

Verifying Management vNIC in parent partition

Name InterfaceDescription Status MacAddress LinkSpeed
----
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #2 Up 08-C0-EB-7E-D2-ED 100 Gbps
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter Up 08-C0-EB-7E-D2-EC 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 1) Hyper-V Virtual Ethernet Adapter #2 Up 00-15-5D-7E-16-00 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 2) Hyper-V Virtual Ethernet Adapter #3 Up 00-15-5D-7E-16-01 100 Gbps
vManagement(mgmt_compute_storage) Hyper-V Virtual Ethernet Adapter Up 08-C0-EB-7E-D2-ED 100 Gbps

```

```

Host Name: AZS-HCI1-N3
Verifying Virtual Switch

Name : ConvergedSwitch(mgmt_compute_storage)
SwitchType : External
NetAdapterInterfaceDescription : Teamed-Interface
NetAdapterInterfaceDescriptions : {Mellanox ConnectX-6 Dx Adapter #2, Mellanox ConnectX-6 Dx Adapter}

Verifying Management vNIC in parent partition

Name InterfaceDescription Status MacAddress LinkSpeed
----
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #2 Up 08-C0-EB-7E-D0-BC 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 1) Hyper-V Virtual Ethernet Adapter #2 Up 00-15-5D-7E-17-00 100 Gbps
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter Up 08-C0-EB-7E-D0-BD 100 Gbps
vManagement(mgmt_compute_storage) Hyper-V Virtual Ethernet Adapter Up 08-C0-EB-7E-D0-BC 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 2) Hyper-V Virtual Ethernet Adapter #3 Up 00-15-5D-7E-17-01 100 Gbps

```

```

Host Name: AZS-HCI1-N4
Verifying Virtual Switch

Name                : ConvergedSwitch(mgmt_compute_storage)
SwitchType          : External
NetAdapterInterfaceDescription : Teamed-Interface
NetAdapterInterfaceDescriptions : {Mellanox ConnectX-6 Dx Adapter, Mellanox ConnectX-6 Dx Adapter #2}

Verifying Management vNIC in parent partition

Name                InterfaceDescription      Status MacAddress      LinkSpeed
----                -
SlotID 2 Port 1    Mellanox ConnectX-6 Dx Adapter Up    08-C0-EB-7E-D3-2C 100 Gbps
SlotID 2 Port 2    Mellanox ConnectX-6 Dx Adapter #2 Up    08-C0-EB-7E-D3-2D 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 1) Hyper-V Virtual Ethernet Adapter #2 Up    00-15-5D-7E-18-00 100 Gbps
vManagement(mgmt_compute_storage) Hyper-V Virtual Ethernet Adapter Up    08-C0-EB-7E-D3-2C 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 2) Hyper-V Virtual Ethernet Adapter #3 Up    00-15-5D-7E-18-01 100 Gbps

```

Note: There will be a brief network disconnect on each server node when VM switch binds to the physical adapters.

Procedure 5. Verify SET Switch Team Load Balancing Algorithm

Note: The load balancing algorithm must be a Hyper-V Port. Each VM switch must be bound to both physical network adapters.

```

$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Verifying SET Switch Load Balancing Algorithm " -ForegroundColor Yellow
Get-VMSwitch | Get-VMSwitchTeam | fl

}

}

```

```

Host Name: AZS-HCI1-N1
Verifying SET Switch Load Balancing Algorithm

Name : ConvergedSwitch(mgmt_compute_storage)
Id : a35d7d75-0183-41b7-81b3-4537de25af4c
NetAdapterInterfaceDescription : {Mellanox ConnectX-6 Dx Adapter, Mellanox ConnectX-6 Dx Adapter #2}
NetAdapterInterfaceGuid : {2402f59e-a5a6-4e97-b45c-d5b1aaf9cb46, 3d44a866-1da9-4e3e-bccc-32972aabcb07}
TeamingMode : SwitchIndependent
LoadBalancingAlgorithm : HyperVPort

Host Name: AZS-HCI1-N2
Verifying SET Switch Load Balancing Algorithm

Name : ConvergedSwitch(mgmt_compute_storage)
Id : 08815ed7-a594-4524-9097-0124bfd94806
NetAdapterInterfaceDescription : {Mellanox ConnectX-6 Dx Adapter #2, Mellanox ConnectX-6 Dx Adapter}
NetAdapterInterfaceGuid : {d2a73545-37e4-4412-a3ef-80241e62ad1b, c26ea84a-da98-477e-8803-d2665edf81bc}
TeamingMode : SwitchIndependent
LoadBalancingAlgorithm : HyperVPort

Host Name: AZS-HCI1-N3
Verifying SET Switch Load Balancing Algorithm

Name : ConvergedSwitch(mgmt_compute_storage)
Id : 38fe406c-bde5-4a44-8258-39a3e0e21fbb
NetAdapterInterfaceDescription : {Mellanox ConnectX-6 Dx Adapter #2, Mellanox ConnectX-6 Dx Adapter}
NetAdapterInterfaceGuid : {c64238fb-edaa-4d2e-94fc-5208b2113ea4, 79f66362-567d-47e7-b3a6-ff1540422bfc}
TeamingMode : SwitchIndependent
LoadBalancingAlgorithm : HyperVPort

Host Name: AZS-HCI1-N4
Verifying SET Switch Load Balancing Algorithm

Name : ConvergedSwitch(mgmt_compute_storage)
Id : 732d6512-aa07-4771-a0c1-5e6354381953
NetAdapterInterfaceDescription : {Mellanox ConnectX-6 Dx Adapter, Mellanox ConnectX-6 Dx Adapter #2}
NetAdapterInterfaceGuid : {fa4cf808-4ba6-4f94-b866-ce5d03461d06, f5d0f7c4-c8af-48e1-b6c8-bf05fe53801a}
TeamingMode : SwitchIndependent
LoadBalancingAlgorithm : HyperVPort

```

Procedure 6. Configure Default Route Metric for Management NIC in Parent Partition

Step 1. Run the following:

```

$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Configuring default route metric for Management NIC " -ForegroundColor Yellow
netsh in ipv4 set ro 0.0.0.0/0 "vManagement(mgmt_compute_storage)" met=10
route print -4

}
}

```

```

Host Name: AZS-HCI1-N1
Configuring default route metric for Management NIC
Ok.

=====
Interface List
13...08 c0 eb 7e d0 c5 .....Hyper-V Virtual Ethernet Adapter
20...00 15 5d 7e 15 00 .....Hyper-V Virtual Ethernet Adapter #2
24...00 15 5d 7e 15 01 .....Hyper-V Virtual Ethernet Adapter #3
8...02 5c d5 62 7f bc .....Microsoft Failover Cluster Virtual Adapter
1.....Software Loopback Interface 1
=====

IPv4 Route Table
=====
Active Routes:
Network Destination        Netmask          Gateway          Interface        Metric
0.0.0.0                    0.0.0.0          192.168.126.1   192.168.126.21   15
127.0.0.0                  255.0.0.0        On-link         127.0.0.1        331
127.0.0.1                  255.255.255.255  On-link         127.0.0.1        331
127.255.255.255           255.255.255.255  On-link         127.0.0.1        331
169.254.0.0                255.255.255.0    On-link         169.254.209.93   261
169.254.0.0                255.255.255.0    On-link         169.254.53.86    261
169.254.53.86             255.255.255.255  On-link         169.254.53.86    261
169.254.209.93            255.255.255.255  On-link         169.254.209.93   261
169.254.255.255           255.255.255.255  On-link         169.254.209.93   261
169.254.255.255           255.255.255.255  On-link         169.254.53.86    261
192.168.126.0             255.255.255.192  On-link         192.168.126.21   261
192.168.126.21           255.255.255.255  On-link         192.168.126.21   261
192.168.126.63           255.255.255.255  On-link         192.168.126.21   261
224.0.0.0                  240.0.0.0        On-link         127.0.0.1        331
224.0.0.0                  240.0.0.0        On-link         192.168.126.21   261
224.0.0.0                  240.0.0.0        On-link         169.254.209.93   261
224.0.0.0                  240.0.0.0        On-link         169.254.53.86    261
255.255.255.255           255.255.255.255  On-link         127.0.0.1        331
255.255.255.255           255.255.255.255  On-link         192.168.126.21   261
255.255.255.255           255.255.255.255  On-link         169.254.209.93   261
255.255.255.255           255.255.255.255  On-link         169.254.53.86    261
=====
Persistent Routes:
Network Address            Netmask          Gateway Address  Metric
0.0.0.0                    0.0.0.0          192.168.126.1   Default
0.0.0.0                    0.0.0.0          192.168.126.1   256
0.0.0.0                    0.0.0.0          On-link          10
=====

```

Procedure 7. Configure Static NIC IP Address for Storage NICs

Note: Leave the gateway unconfigured for storage NICs.

Host	SMB NIC Name	SMB NIC IP Address
AzS-HCI1-N1	SMB-A	192.168.107.21
	SMB-B	192.168.207.21
AzS-HCI1-N2	SMB-A	192.168.107.22
	SMB-B	192.168.207.22

Host	SMB NIC Name	SMB NIC IP Address
AzS-HCI1-N3	SMB-A	192.168.107.23
	SMB-B	192.168.207.23
AzS-HCI1-N4	SMB-A	192.168.107.24
	SMB-B	192.168.207.24

```

$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")

$IPStorageNetA = "192.168.107." #vSMB(mgmt_compute_storage#SlotID 2 Port 1)networkaddress
$IPStorageNetB = "192.168.207." #vSMB(mgmt_compute_storage#SlotID 2 Port 2)networkaddress
$IPHostAddr = 21 #Starting host address

foreach ($node in $nodes) {
    $session = New-CimSession -ComputerName $node
    New-NetIPAddress -CimSession $session -InterfaceAlias "vSMB(mgmt_compute_storage#SlotID 2 Port 1)" -IPAddress
    ($IPStorageNetA+$IPHostAddr.ToString()) -PrefixLength 24

    New-NetIPAddress -CimSession $session -InterfaceAlias "vSMB(mgmt_compute_storage#SlotID 2 Port 2)" -IPAddress
    ($IPStorageNetB+$IPHostAddr.ToString()) -PrefixLength 24
    $IPHostAddr++
}

Get-CimSession | Remove-CimSession
Remove-Variable session

```

Note: Network connectivity may be temporarily disrupted during the following configuration operations, but connectivity will automatically recover.

Procedure 8. Verify NIC IP Address for Storage NICs

Step 1. Run the following:

```

$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

    Invoke-Command $node -Credential $Creds -scriptblock {
        write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

        Write-Host "Verifying Storage NIC IP Address " -ForegroundColor Yellow
        Get-NetIPConfiguration -InterfaceAlias vSMB* | fl InterfaceAlias, IPv4Address, IPv4DefaultGateway
    }
}

```

```
}
```

```
Host Name: AZS-HCI1-N1
Verifying Storage NIC IP Address

InterfaceAlias      : vsmb(mgmt_compute_storage#SlotID 2 Port 1)
IPv4Address         : {192.168.107.21}
IPv4DefaultGateway :

InterfaceAlias      : vsmb(mgmt_compute_storage#SlotID 2 Port 2)
IPv4Address         : {192.168.207.21}
IPv4DefaultGateway :

Host Name: AZS-HCI1-N2
Verifying Storage NIC IP Address

InterfaceAlias      : vsmb(mgmt_compute_storage#SlotID 2 Port 1)
IPv4Address         : {192.168.107.22}
IPv4DefaultGateway :

InterfaceAlias      : vsmb(mgmt_compute_storage#SlotID 2 Port 2)
IPv4Address         : {192.168.207.22}
IPv4DefaultGateway :

Host Name: AZS-HCI1-N3
Verifying Storage NIC IP Address

InterfaceAlias      : vsmb(mgmt_compute_storage#SlotID 2 Port 1)
IPv4Address         : {192.168.107.23}
IPv4DefaultGateway :

InterfaceAlias      : vsmb(mgmt_compute_storage#SlotID 2 Port 2)
IPv4Address         : {192.168.207.23}
IPv4DefaultGateway :

Host Name: AZS-HCI1-N4
Verifying Storage NIC IP Address

InterfaceAlias      : vsmb(mgmt_compute_storage#SlotID 2 Port 1)
IPv4Address         : {192.168.107.24}
IPv4DefaultGateway :

InterfaceAlias      : vsmb(mgmt_compute_storage#SlotID 2 Port 2)
IPv4Address         : {192.168.207.24}
IPv4DefaultGateway :
```

Procedure 9. Verify DNS Registration is Removed for Storage Interfaces

Step 1. Run the following:

```
$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Removing DNS Restistration from Storage NICs " -ForegroundColor Yellow
Set-DnsClient -InterfaceAlias "vSMB(mgmt_compute_storage#SlotID 2 Port 1)" -
RegisterThisConnectionsAddress:$false
Set-DnsClient -InterfaceAlias "vSMB(mgmt_compute_storage#SlotID 2 Port 2)" -
RegisterThisConnectionsAddress:$false
Get-DnsClient -InterfaceAlias "vSMB(mgmt_compute_storage#SlotID 2 Port 1)" | ft
InterfaceAlias,RegisterThisConnectionsAddress
Get-DnsClient -InterfaceAlias "vSMB(mgmt_compute_storage#SlotID 2 Port 2)" | ft
InterfaceAlias,RegisterThisConnectionsAddress

}
}
```

```

Host Name: AZS-HCI1-N1
Removing DNS Restistration from Storage NICs

InterfaceAlias                                RegisterThisConnectionsAddress
-----
vSMB(mgmt_compute_storage#SlotID 2 Port 1)    False

InterfaceAlias                                RegisterThisConnectionsAddress
-----
vSMB(mgmt_compute_storage#SlotID 2 Port 2)    False

Host Name: AZS-HCI1-N2
Removing DNS Restistration from Storage NICs

InterfaceAlias                                RegisterThisConnectionsAddress
-----
vSMB(mgmt_compute_storage#SlotID 2 Port 1)    False

InterfaceAlias                                RegisterThisConnectionsAddress
-----
vSMB(mgmt_compute_storage#SlotID 2 Port 2)    False

Host Name: AZS-HCI1-N3
Removing DNS Restistration from Storage NICs

InterfaceAlias                                RegisterThisConnectionsAddress
-----
vSMB(mgmt_compute_storage#SlotID 2 Port 1)    False

InterfaceAlias                                RegisterThisConnectionsAddress
-----
vSMB(mgmt_compute_storage#SlotID 2 Port 2)    False

Host Name: AZS-HCI1-N4
Removing DNS Restistration from Storage NICs

InterfaceAlias                                RegisterThisConnectionsAddress
-----
vSMB(mgmt_compute_storage#SlotID 2 Port 1)    False

InterfaceAlias                                RegisterThisConnectionsAddress
-----
vSMB(mgmt_compute_storage#SlotID 2 Port 2)    False

```

Procedure 10. Enable Preserving 802.1p Priority Marking to Pass Through the vSwitch

Note: The virtual switch zeros-out 802.1p priority marking in the packet header. This is the default behavior. Preserving the 802.1p priority marking in the packet header is required for classifying and prioritizing network traffic in the fabric and other northbound switches that have QoS policies configured. This setting affects prioritized network traffic traversing the virtual switch. This setting is required prioritizing Cluster Communication network traffic. RDMA traffic passing through RDMA enabled vNICs is not

affected by this setting because this traffic bypasses the virtual switch and goes directly to the physical NIC.

Step 1. Run the following:

```
$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Configure vSwitch to pass 802.1p priority marking " -ForegroundColor Yellow
Set-VMNetworkAdapter -Name "vManagement(mgmt_compute_storage)" -ManagementOS -IeeePriorityTag On
Get-VMNetworkAdapter -ManagementOS | ft Name,IeeePriorityTag

}
}
```

```
Host Name: AZS-HCI1-N1
Configure vSwitch to pass 802.1p priority marking

Name                                IeeePriorityTag
----                                -
vManagement(mgmt_compute_storage)    On
vSMB(mgmt_compute_storage#SlotID 2 Port 1)    On
vSMB(mgmt_compute_storage#SlotID 2 Port 2)    On

Host Name: AZS-HCI1-N2
Configure vSwitch to pass 802.1p priority marking

Name                                IeeePriorityTag
----                                -
vManagement(mgmt_compute_storage)    On
vSMB(mgmt_compute_storage#SlotID 2 Port 1)    On
vSMB(mgmt_compute_storage#SlotID 2 Port 2)    On

Host Name: AZS-HCI1-N3
Configure vSwitch to pass 802.1p priority marking

Name                                IeeePriorityTag
----                                -
vManagement(mgmt_compute_storage)    On
vSMB(mgmt_compute_storage#SlotID 2 Port 1)    On
vSMB(mgmt_compute_storage#SlotID 2 Port 2)    On

Host Name: AZS-HCI1-N4
Configure vSwitch to pass 802.1p priority marking

Name                                IeeePriorityTag
----                                -
vManagement(mgmt_compute_storage)    On
vSMB(mgmt_compute_storage#SlotID 2 Port 1)    On
vSMB(mgmt_compute_storage#SlotID 2 Port 2)    On
```

Procedure 11. Verify the Storage vNIC VLANs

Step 1. Run the following:

```
$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verify vNIC VLANs Configuration " -ForegroundColor Yellow

Get-VMNetworkAdapter -ManagementOS | Get-VMNetworkAdapterIsolation | FT IsolationMode, DefaultIsolationID,
ParentAdapter -AutoSize

}
}
```

```
Host Name: AZS-HCI1-N1
Verify vNIC VLANs Configuration

IsolationMode DefaultIsolationID ParentAdapter
-----
None          0 VMInternalNetworkAdapter, Name = 'vManagement(mgmt_compute_storage)'
Vlan         107 VMInternalNetworkAdapter, Name = 'vSMB(mgmt_compute_storage#SlotID 2 Port 1)'
Vlan         207 VMInternalNetworkAdapter, Name = 'vSMB(mgmt_compute_storage#SlotID 2 Port 2)'
```

```
Host Name: AZS-HCI1-N2
Verify vNIC VLANs Configuration

IsolationMode DefaultIsolationID ParentAdapter
-----
None          0 VMInternalNetworkAdapter, Name = 'vManagement(mgmt_compute_storage)'
Vlan         107 VMInternalNetworkAdapter, Name = 'vSMB(mgmt_compute_storage#SlotID 2 Port 1)'
Vlan         207 VMInternalNetworkAdapter, Name = 'vSMB(mgmt_compute_storage#SlotID 2 Port 2)'
```

```
Host Name: AZS-HCI1-N3
Verify vNIC VLANs Configuration

IsolationMode DefaultIsolationID ParentAdapter
-----
None          0 VMInternalNetworkAdapter, Name = 'vManagement(mgmt_compute_storage)'
Vlan         107 VMInternalNetworkAdapter, Name = 'vSMB(mgmt_compute_storage#SlotID 2 Port 1)'
Vlan         207 VMInternalNetworkAdapter, Name = 'vSMB(mgmt_compute_storage#SlotID 2 Port 2)'
```

```
Host Name: AZS-HCI1-N4
Verify vNIC VLANs Configuration

IsolationMode DefaultIsolationID ParentAdapter
-----
None          0 VMInternalNetworkAdapter, Name = 'vManagement(mgmt_compute_storage)'
Vlan         107 VMInternalNetworkAdapter, Name = 'vSMB(mgmt_compute_storage#SlotID 2 Port 1)'
Vlan         207 VMInternalNetworkAdapter, Name = 'vSMB(mgmt_compute_storage#SlotID 2 Port 2)'
```

Procedure 12. Verify Network Adapters

Step 1. Run the following:

```
$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {
```

```

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verifying NIC status " -ForegroundColor Yellow

Get-NetAdapter | sort Name | ft Name,InterfaceDescription,Status,MTUSize,LinkSpeed

}
}

```

```

Host Name: AZS-HCI1-N1
Verifying NIC status

Name                InterfaceDescription      Status MTUSize LinkSpeed
----                -
SlotID 2 Port 1     Mellanox ConnectX-6 Dx Adapter      Up    1660 100 Gbps
SlotID 2 Port 2     Mellanox ConnectX-6 Dx Adapter #2   Up    1660 100 Gbps
vManagement(mgmt_compute_storage)    Hyper-V Virtual Ethernet Adapter    Up    1500 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 1) Hyper-V Virtual Ethernet Adapter #2 Up    1500 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 2) Hyper-V Virtual Ethernet Adapter #3 Up    1500 100 Gbps

Host Name: AZS-HCI1-N2
Verifying NIC status

Name                InterfaceDescription      Status MTUSize LinkSpeed
----                -
SlotID 2 Port 1     Mellanox ConnectX-6 Dx Adapter #2   Up    1660 100 Gbps
SlotID 2 Port 2     Mellanox ConnectX-6 Dx Adapter      Up    1660 100 Gbps
vManagement(mgmt_compute_storage)    Hyper-V Virtual Ethernet Adapter    Up    1500 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 1) Hyper-V Virtual Ethernet Adapter #2 Up    1500 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 2) Hyper-V Virtual Ethernet Adapter #3 Up    1500 100 Gbps

Host Name: AZS-HCI1-N3
Verifying NIC status

Name                InterfaceDescription      Status MTUSize LinkSpeed
----                -
SlotID 2 Port 1     Mellanox ConnectX-6 Dx Adapter #2   Up    1660 100 Gbps
SlotID 2 Port 2     Mellanox ConnectX-6 Dx Adapter      Up    1660 100 Gbps
vManagement(mgmt_compute_storage)    Hyper-V Virtual Ethernet Adapter    Up    1500 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 1) Hyper-V Virtual Ethernet Adapter #2 Up    1500 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 2) Hyper-V Virtual Ethernet Adapter #3 Up    1500 100 Gbps

Host Name: AZS-HCI1-N4
Verifying NIC status

Name                InterfaceDescription      Status MTUSize LinkSpeed
----                -
SlotID 2 Port 1     Mellanox ConnectX-6 Dx Adapter      Up    1660 100 Gbps
SlotID 2 Port 2     Mellanox ConnectX-6 Dx Adapter #2   Up    1660 100 Gbps
vManagement(mgmt_compute_storage)    Hyper-V Virtual Ethernet Adapter    Up    1500 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 1) Hyper-V Virtual Ethernet Adapter #2 Up    1500 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 2) Hyper-V Virtual Ethernet Adapter #3 Up    1500 100 Gbps

```

Procedure 13. Verify RDMA and RoCEv2 Protocol is Enabled on Physical NICs

Step 1. Run the following:

```

$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

```

```

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verifying RDMA and RoCEv2 status on physical NICS " -ForegroundColor Yellow

Get-NetAdapterAdvancedProperty -InterfaceDescription "Mellanox ConnectX*" -DisplayName "NetworkDirect*" | ft
Name, InterfaceDescription, DisplayName, DisplayValue

}
}

```

```

Host Name: AZS-HCI1-N1
Verifying RDMA and RoCEv2 status on physical NICS

Name          InterfaceDescription          DisplayName          DisplayValue
-----
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Functionality Enabled
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Technology RoCEv2
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter NetworkDirect Functionality Enabled
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter NetworkDirect Technology RoCEv2

Host Name: AZS-HCI1-N2
Verifying RDMA and RoCEv2 status on physical NICS

Name          InterfaceDescription          DisplayName          DisplayValue
-----
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Functionality Enabled
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Technology RoCEv2
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter NetworkDirect Functionality Enabled
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter NetworkDirect Technology RoCEv2

Host Name: AZS-HCI1-N3
Verifying RDMA and RoCEv2 status on physical NICS

Name          InterfaceDescription          DisplayName          DisplayValue
-----
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Functionality Enabled
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Technology RoCEv2
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter NetworkDirect Functionality Enabled
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter NetworkDirect Technology RoCEv2

Host Name: AZS-HCI1-N4
Verifying RDMA and RoCEv2 status on physical NICS

Name          InterfaceDescription          DisplayName          DisplayValue
-----
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter NetworkDirect Functionality Enabled
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter NetworkDirect Technology RoCEv2
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Functionality Enabled
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Technology RoCEv2

```

Procedure 14. Verify that RDMA is Enabled on the Storage vNIC Adapters

Step 1. Run the following:

```

$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {

```

```
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
```

```
Write-Host "Verifying that RDMA is enabled on the Storage vNICs" -ForegroundColor Yellow
```

```
Get-NetAdapterRdma | ft
```

```
}  
}
```

```
Host Name: AZS-HCI1-N1  
Verifying that RDMA is enabled on the Storage vNICs  
  
Name                InterfaceDescription      Enabled  Operational  PFC      ETS  
----                -  
vManagement(mgmt_compu... Hyper-V Virtual Ethernet Adapter  False   False        NA       NA  
vSMB(mgmt_compute_stor... Hyper-V Virtual Ethernet Adapter #3  True    True         NA       NA  
vSMB(mgmt_compute_stor... Hyper-V Virtual Ethernet Adapter #2  True    True         NA       NA  
SlotID 2 Port 2      Mellanox ConnectX-6 Dx Adapter #2    True    True         True     True  
SlotID 2 Port 1      Mellanox ConnectX-6 Dx Adapter      True    True         True     True  
  
Host Name: AZS-HCI1-N2  
Verifying that RDMA is enabled on the Storage vNICs  
  
Name                InterfaceDescription      Enabled  Operational  PFC      ETS  
----                -  
SlotID 2 Port 1      Mellanox ConnectX-6 Dx Adapter #2    True    True         True     True  
SlotID 2 Port 2      Mellanox ConnectX-6 Dx Adapter      True    True         True     True  
vSMB(mgmt_compute_stor... Hyper-V Virtual Ethernet Adapter #2  True    True         NA       NA  
vSMB(mgmt_compute_stor... Hyper-V Virtual Ethernet Adapter #3  True    True         NA       NA  
vManagement(mgmt_compu... Hyper-V Virtual Ethernet Adapter  False   False        NA       NA  
  
Host Name: AZS-HCI1-N3  
Verifying that RDMA is enabled on the Storage vNICs  
  
Name                InterfaceDescription      Enabled  Operational  PFC      ETS  
----                -  
SlotID 2 Port 1      Mellanox ConnectX-6 Dx Adapter #2    True    True         True     True  
vSMB(mgmt_compute_stor... Hyper-V Virtual Ethernet Adapter #2  True    True         NA       NA  
SlotID 2 Port 2      Mellanox ConnectX-6 Dx Adapter      True    True         True     True  
vManagement(mgmt_compu... Hyper-V Virtual Ethernet Adapter  False   False        NA       NA  
vSMB(mgmt_compute_stor... Hyper-V Virtual Ethernet Adapter #3  True    True         NA       NA  
  
Host Name: AZS-HCI1-N4  
Verifying that RDMA is enabled on the Storage vNICs  
  
Name                InterfaceDescription      Enabled  Operational  PFC      ETS  
----                -  
SlotID 2 Port 1      Mellanox ConnectX-6 Dx Adapter      True    True         True     True  
SlotID 2 Port 2      Mellanox ConnectX-6 Dx Adapter #2    True    True         True     True  
vSMB(mgmt_compute_stor... Hyper-V Virtual Ethernet Adapter #2  True    True         NA       NA  
vManagement(mgmt_compu... Hyper-V Virtual Ethernet Adapter  False   False        NA       NA  
vSMB(mgmt_compute_stor... Hyper-V Virtual Ethernet Adapter #3  True    True         NA       NA
```

Procedure 15. Verify the Mapping of each SMB-Direct NIC to the respective Fabric

Step 1. Run the following:

```
$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")  
foreach ($node in $nodes) {  
  
Invoke-Command $node -Credential $Creds -scriptblock {  
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
```

```

Write-Host "Verify Mapping of each storage vNIC to the respective fabric " -ForegroundColor Yellow
Get-VMNetworkAdapterTeamMapping -ManagementOS | ft ComputerName,NetAdapterName,ParentAdapter

}

}

```

```

Host Name: AZS-HCI1-N1
Verify Mapping of each storage vNIC to the respective fabric

ComputerName NetAdapterName ParentAdapter
-----
AZS-HCI1-N1 SlotID 2 Port 1 VMInternalNetworkAdapter, Name = 'vSMB(mgmt_compute_storage#SlotID 2 Port 1)'
AZS-HCI1-N1 SlotID 2 Port 2 VMInternalNetworkAdapter, Name = 'vSMB(mgmt_compute_storage#SlotID 2 Port 2)'

Host Name: AZS-HCI1-N2
Verify Mapping of each storage vNIC to the respective fabric

ComputerName NetAdapterName ParentAdapter
-----
AZS-HCI1-N2 SlotID 2 Port 1 VMInternalNetworkAdapter, Name = 'vSMB(mgmt_compute_storage#SlotID 2 Port 1)'
AZS-HCI1-N2 SlotID 2 Port 2 VMInternalNetworkAdapter, Name = 'vSMB(mgmt_compute_storage#SlotID 2 Port 2)'

Host Name: AZS-HCI1-N3
Verify Mapping of each storage vNIC to the respective fabric

ComputerName NetAdapterName ParentAdapter
-----
AZS-HCI1-N3 SlotID 2 Port 1 VMInternalNetworkAdapter, Name = 'vSMB(mgmt_compute_storage#SlotID 2 Port 1)'
AZS-HCI1-N3 SlotID 2 Port 2 VMInternalNetworkAdapter, Name = 'vSMB(mgmt_compute_storage#SlotID 2 Port 2)'

Host Name: AZS-HCI1-N4
Verify Mapping of each storage vNIC to the respective fabric

ComputerName NetAdapterName ParentAdapter
-----
AZS-HCI1-N4 SlotID 2 Port 1 VMInternalNetworkAdapter, Name = 'vSMB(mgmt_compute_storage#SlotID 2 Port 1)'
AZS-HCI1-N4 SlotID 2 Port 2 VMInternalNetworkAdapter, Name = 'vSMB(mgmt_compute_storage#SlotID 2 Port 2)'

```

Procedure 16. Verify RDMA Capabilities

Step 1. Run the following:

```

$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verify Storage vNIC RDMA operational status " -ForegroundColor Yellow
Get-SmbClientNetworkInterface | ft FriendlyName, RDMAcapable

}

}

```

```

Host Name: AZS-HCI1-N1
Verify Storage vNIC RDMA operational status

FriendlyName                                RDMACapable
-----
vManagement(mgmt_compute_storage)          False
vSMB(mgmt_compute_storage#SlotID 2 Port 1)  True
vSMB(mgmt_compute_storage#SlotID 2 Port 2)  True
SlotID 2 Port 1                             False
SlotID 2 Port 2                             False
Local Area Connection* 1                    False

Host Name: AZS-HCI1-N2
Verify Storage vNIC RDMA operational status

FriendlyName                                RDMACapable
-----
SlotID 2 Port 1                             False
SlotID 2 Port 2                             False
vManagement(mgmt_compute_storage)          False
vSMB(mgmt_compute_storage#SlotID 2 Port 1)  True
vSMB(mgmt_compute_storage#SlotID 2 Port 2)  True
Local Area Connection* 1                    False

Host Name: AZS-HCI1-N3
Verify Storage vNIC RDMA operational status

FriendlyName                                RDMACapable
-----
SlotID 2 Port 2                             False
vSMB(mgmt_compute_storage#SlotID 2 Port 1)  True
SlotID 2 Port 1                             False
Local Area Connection* 1                    False
vManagement(mgmt_compute_storage)          False
vSMB(mgmt_compute_storage#SlotID 2 Port 2)  True

Host Name: AZS-HCI1-N4
Verify Storage vNIC RDMA operational status

FriendlyName                                RDMACapable
-----
SlotID 2 Port 1                             False
SlotID 2 Port 2                             False
vManagement(mgmt_compute_storage)          False
vSMB(mgmt_compute_storage#SlotID 2 Port 1)  True
vSMB(mgmt_compute_storage#SlotID 2 Port 2)  True
Local Area Connection* 1                    False

```

Non-Converged Network Topology Configuration

Note: Skip the steps in this section if the deployment model chosen for your environment, is a fully converged topology. Only one deployment model needs to be followed, i.e., either fully converged or non-converged or storage switchless topology.

This subject contains the following procedures to create and deploy standalone Network ATC intent for a non-converged topology. Here you create two intents – one intent for “Management and Compute” networks, and a separate intent for “Storage” Network:

- [Identify Physical Network Card Port Names](#)

- [Create and Deploy Standalone Network ATC Intent - Management and Compute Only](#)
- [Verify Network ATC Intent Status](#)
- [Create and Deploy Standalone Network ATC Intent - Storage Only](#)
- [Verify Network ATC Intent Status](#)
- [Verify Virtual Switch and Virtual NIC Creation in the Parent Partition](#)
- [Verify SET Switch Team Load Balancing Algorithm](#)
- [Configure Default Route Metric for Management NIC in Parent Partition](#)
- [Configure Static NIC IP Address for Storage NICs](#)
- [Verify NIC IP Address for Storage NICs](#)
- [Verify DNS Registration is Removed for Storage Interfaces](#)
- [Enable Preserving 802.1p Priority Marking to Pass Through the vSwitch](#)
- [Verify the Storage vNIC VLANs](#)
- [Verify Network Adapters](#)
- [Verify RDMA and RoCEv2 Protocol is Enabled on Physical NICs](#)
- [Verify RDMA is Enabled on the Storage NIC Adapters](#)
- [Verify RDMA Capabilities](#)

Procedure 1. Identify Physical Network Card Port Names

Step 1. Run the following:

```
$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")
foreach ($node in $nodes) {
Invoke-Command $node -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Retrieving physical NIC port names " -ForegroundColor Yellow

Get-netadapter | ft Name, InterfaceDescription, Status, MacAddress, LinkSpeed

}
}
```



```

Host Name: AZSHCI1-M6-N1
Retrieving physical NIC port names

Name          InterfaceDescription      Status MacAddress      LinkSpeed
-----
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #2 Up    08-C0-EB-7E-D0-C4 100 Gbps
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #1 Up    08-C0-EB-7E-D0-C5 100 Gbps
SlotID 5 Port 1 Mellanox ConnectX-6 Dx Adapter #4 Up    94-6D-AE-75-F1-23 100 Gbps
SlotID 5 Port 2 Mellanox ConnectX-6 Dx Adapter #3 Up    94-6D-AE-75-F1-22 100 Gbps

Host Name: AZSHCI1-M6-N2
Retrieving physical NIC port names

Name          InterfaceDescription      Status MacAddress      LinkSpeed
-----
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #2 Up    08-C0-EB-7E-D2-EC 100 Gbps
SlotID 5 Port 2 Mellanox ConnectX-6 Dx Adapter #3 Up    94-6D-AE-75-E6-1A 100 Gbps
SlotID 5 Port 1 Mellanox ConnectX-6 Dx Adapter #4 Up    94-6D-AE-75-E6-18 100 Gbps
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #1 Up    08-C0-EB-7E-D2-ED 100 Gbps

Host Name: AZSHCI1-M6-N3
Retrieving physical NIC port names

Name          InterfaceDescription      Status MacAddress      LinkSpeed
-----
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #1 Up    08-C0-EB-7E-D0-BD 100 Gbps
SlotID 5 Port 1 Mellanox ConnectX-6 Dx Adapter #4 Up    94-6D-AE-75-E8-FB 100 Gbps
SlotID 5 Port 2 Mellanox ConnectX-6 Dx Adapter #3 Up    94-6D-AE-75-E8-FA 100 Gbps
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #2 Up    08-C0-EB-7E-D0-BC 100 Gbps

```

Note: If the NIC port names are “Ethernet” and “Ethernet 2”, CDN is not enabled. CDN (Consistent Device Naming) must be enabled for correct physical to virtual NIC mapping later in this guide.

Procedure 2. Create and Deploy Standalone Network ATC Intent – Management and Compute only

Step 1. Run the following script block to create a virtual switch with SET enabled and three virtual NICs:

```

$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Create and Deploy Standalone NetworkATC Intent - Mgmt and Compute only " -ForegroundColor Yellow

$Mgmt_Compute_NICs = @("SlotID 2 Port 1", "SlotID 5 Port 2")

Add-NetIntent -Name Mgmt_Compute -AdapterName $Mgmt_Compute_NICs -Management -Compute -verbose

}
}

```

Note: There will be a brief network disconnect on each server node when VM switch binds to the physical adapters.

Procedure 3. Verify Network ATC Intent Status

Step 1. Run the following:

```

$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")
foreach ($node in $nodes) {

```

```

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Checking Network ATC Intent Status" -ForegroundColor Yellow

Get-netIntentStatus -ComputerName $node | ft
Host,IntentName,ConfigurationStatus,ProvisioningStatus,IsComputeIntentSet,IsManagementIntentSet,IsStorageInte
ntset,IsStretchIntentSet

}
}

```

```

Host Name: AZSHCI1-M6-N1
Checking Network ATC Intent Status
WARNING: If you are running Windows PowerShell remotely, note that some failover clustering cmdlets do not work remotely. When possible,
run the cmdlet locally and specify a remote computer as the target. To run the cmdlet remotely, try using the Credential Security Service
Provider (CredSSP). All additional errors or warnings from this cmdlet might be caused by running it remotely.

Host          IntentName  ConfigurationStatus  ProvisioningStatus  IsComputeIntentSet  IsManagementIntentSet  IsStorageIntentSet  IsStretchI
ntentSet
-----
AZSHCI1-M6-N1 mgmt_compute Success          Completed           True                True                 False              False

Host Name: AZSHCI1-M6-N2
Checking Network ATC Intent Status
WARNING: If you are running Windows PowerShell remotely, note that some failover clustering cmdlets do not work remotely. When possible,
run the cmdlet locally and specify a remote computer as the target. To run the cmdlet remotely, try using the Credential Security Service
Provider (CredSSP). All additional errors or warnings from this cmdlet might be caused by running it remotely.

Host          IntentName  ConfigurationStatus  ProvisioningStatus  IsComputeIntentSet  IsManagementIntentSet  IsStorageIntentSet  IsStretchI
ntentSet
-----
AZSHCI1-M6-N2 mgmt_compute Success          Completed           True                True                 False              False

Host Name: AZSHCI1-M6-N3
Checking Network ATC Intent Status
WARNING: If you are running Windows PowerShell remotely, note that some failover clustering cmdlets do not work remotely. When possible,
run the cmdlet locally and specify a remote computer as the target. To run the cmdlet remotely, try using the Credential Security Service
Provider (CredSSP). All additional errors or warnings from this cmdlet might be caused by running it remotely.

Host          IntentName  ConfigurationStatus  ProvisioningStatus  IsComputeIntentSet  IsManagementIntentSet  IsStorageIntentSet  IsStretchI
ntentSet
-----
AZSHCI1-M6-N3 mgmt_compute Success          Completed           True                True                 False              False

```

Procedure 4. Create and Deploy Standalone Network ATC Intent – Storage only

Step 1. Run the following script block to create a virtual switch with SET enabled and three virtual NICs:

```

$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Create and Deploy Standalone Network ATC Intent - Storage Only " -ForegroundColor Yellow

$Storage_NICs = @("SlotID 2 Port 2", "SlotID 5 Port 1")
$Storage_VLANS = @(207,107)

$QoSOverride = New-NetIntentQoSPolicyOverRides

```

```

$QoSOverride.PriorityValue8021Action_SMB = 4
$QoSOverride.PriorityValue8021Action_Cluster = 5

$AdapterOverride = New-NetIntentAdapterPropertyOverrides
$AdapterOverride.NetworkDirectTechnology = 4
$AdapterOverride.JumboPacket = 9014

$storageOverride = new-NetIntentStorageOverrides
$storageOverride.EnableAutomaticIPGeneration = $false

Add-NetIntent -Name Storage -AdapterName $Storage_NICs -Storage -StorageVlans $Storage_VLANS -
QoSPolicyOverrides $QoSOverride -AdapterPropertyOverrides $AdapterOverride -StorageOverrides $storageoverride
}
}

```

Procedure 5. Verify Network ATC Intent Status

Step 1. Run the following:

```

$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Checking Network ATC Intent Status" -ForegroundColor Yellow

Get-netIntentStatus -ComputerName $node | ft
Host,IntentName,ConfigurationStatus,ProvisioningStatus,IsComputeIntentSet,IsManagementIntentSet,IsStorageInte
ntset,IsStretchIntentSet

}
}

```

```

Host Name: AZSHCI1-M6-N1
Checking Network ATC Intent Status
WARNING: If you are running Windows PowerShell remotely, note that some failover clustering cmdlets do not work remotely. When possible, run the cmdlet locally and specify a remote computer as the target. To run the cmdlet remotely, try using the Credential Security Service Provider (CredSSP). All additional errors or warnings from this cmdlet might be caused by running it remotely.

Completed

Host      IntentName      ConfigurationStatus      ProvisioningStatus      IsComputeIntentSet      IsManagementIntentSet      IsStorageIntentSet      IsStretchIntentSet
-----
AZSHCI1-M6-N1 mgmt_compute      Success                  Completed                True                     True                      False                   False
AZSHCI1-M6-N1 storage      Success                  Completed                False                    False                     True                    False

Host Name: AZSHCI1-M6-N2
Checking Network ATC Intent Status
WARNING: If you are running Windows PowerShell remotely, note that some failover clustering cmdlets do not work remotely. When possible, run the cmdlet locally and specify a remote computer as the target. To run the cmdlet remotely, try using the Credential Security Service Provider (CredSSP). All additional errors or warnings from this cmdlet might be caused by running it remotely.

Completed

Host      IntentName      ConfigurationStatus      ProvisioningStatus      IsComputeIntentSet      IsManagementIntentSet      IsStorageIntentSet      IsStretchIntentSet
-----
AZSHCI1-M6-N2 mgmt_compute      Success                  Completed                True                     True                      False                   False
AZSHCI1-M6-N2 storage      Success                  Completed                False                    False                     True                    False

Host Name: AZSHCI1-M6-N3
Checking Network ATC Intent Status
WARNING: If you are running Windows PowerShell remotely, note that some failover clustering cmdlets do not work remotely. When possible, run the cmdlet locally and specify a remote computer as the target. To run the cmdlet remotely, try using the Credential Security Service Provider (CredSSP). All additional errors or warnings from this cmdlet might be caused by running it remotely.

Completed

Host      IntentName      ConfigurationStatus      ProvisioningStatus      IsComputeIntentSet      IsManagementIntentSet      IsStorageIntentSet      IsStretchIntentSet
-----
AZSHCI1-M6-N3 mgmt_compute      Success                  Completed                True                     True                      False                   False
AZSHCI1-M6-N3 storage      Success                  Completed                False                    False                     True                    False

```

Procedure 6. Verify Virtual Switch and Virtual NIC Creation in the Parent Partition

Step 1. Run the following:

```

$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Verifying Virtual Switch " -ForegroundColor Yellow
Get-VMSwitch | fl Name, SwitchType, NetAdapterInterfaceDescription, NetAdapterInterfaceDescriptions

Write-Host " Verifying Management vNIC in parent partition " -ForegroundColor Yellow
Get-netadapter | ft Name, InterfaceDescription, Status, MacAddress, LinkSpeed

}
}

```

```

Host Name: AZSHCI1-M6-N1
Verifying Virtual Switch

Name                : ConvergedSwitch(mgmt_compute)
SwitchType          : External
NetAdapterInterfaceDescription : Teamed-Interface
NetAdapterInterfaceDescriptions : {Mellanox ConnectX-6 Dx Adapter, Mellanox ConnectX-6 Dx Adapter #3}

Verifying Management vNIC in parent partition

Name                InterfaceDescription      Status MacAddress      LinkSpeed
-----
SlotID 2 Port 2    Mellanox ConnectX-6 Dx Adapter #2 Up    08-C0-EB-7E-D0-C4 100 Gbps
SlotID 2 Port 1    Mellanox ConnectX-6 Dx Adapter Up    08-C0-EB-7E-D0-C5 100 Gbps
vManagement(mgmt_compute) Hyper-V Virtual Ethernet Adapter Up    08-C0-EB-7E-D0-C5 100 Gbps
SlotID 5 Port 1    Mellanox ConnectX-6 Dx Adapter #4 Up    94-6D-AE-75-F1-23 100 Gbps
SlotID 5 Port 2    Mellanox ConnectX-6 Dx Adapter #3 Up    94-6D-AE-75-F1-22 100 Gbps

Host Name: AZSHCI1-M6-N2
Verifying Virtual Switch

Name                : ConvergedSwitch(mgmt_compute)
SwitchType          : External
NetAdapterInterfaceDescription : Teamed-Interface
NetAdapterInterfaceDescriptions : {Mellanox ConnectX-6 Dx Adapter, Mellanox ConnectX-6 Dx Adapter #3}

Verifying Management vNIC in parent partition

Name                InterfaceDescription      Status MacAddress      LinkSpeed
-----
SlotID 2 Port 2    Mellanox ConnectX-6 Dx Adapter #2 Up    08-C0-EB-7E-D2-EC 100 Gbps
SlotID 5 Port 2    Mellanox ConnectX-6 Dx Adapter #3 Up    94-6D-AE-75-E6-1A 100 Gbps
SlotID 5 Port 1    Mellanox ConnectX-6 Dx Adapter #4 Up    94-6D-AE-75-E6-1B 100 Gbps
vManagement(mgmt_compute) Hyper-V Virtual Ethernet Adapter Up    08-C0-EB-7E-D2-ED 100 Gbps
SlotID 2 Port 1    Mellanox ConnectX-6 Dx Adapter Up    08-C0-EB-7E-D2-ED 100 Gbps

Host Name: AZSHCI1-M6-N3
Verifying Virtual Switch

Name                : ConvergedSwitch(mgmt_compute)
SwitchType          : External
NetAdapterInterfaceDescription : Teamed-Interface
NetAdapterInterfaceDescriptions : {Mellanox ConnectX-6 Dx Adapter #2, Mellanox ConnectX-6 Dx Adapter #3}

Verifying Management vNIC in parent partition

Name                InterfaceDescription      Status MacAddress      LinkSpeed
-----
SlotID 2 Port 2    Mellanox ConnectX-6 Dx Adapter Up    08-C0-EB-7E-D0-BD 100 Gbps
SlotID 5 Port 1    Mellanox ConnectX-6 Dx Adapter #4 Up    94-6D-AE-75-E8-FB 100 Gbps
SlotID 5 Port 2    Mellanox ConnectX-6 Dx Adapter #3 Up    94-6D-AE-75-E8-FA 100 Gbps
vManagement(mgmt_compute) Hyper-V Virtual Ethernet Adapter Up    08-C0-EB-7E-D0-BC 100 Gbps
SlotID 2 Port 1    Mellanox ConnectX-6 Dx Adapter #2 Up    08-C0-EB-7E-D0-BC 100 Gbps

```

Procedure 7. Verify SET Switch Team Load Balancing Algorithm

Note: The load balancing algorithm must be a Hyper-V Port. Each VM switch must be bound to both physical network adapters.

```

$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Verifying SET Switch Load Balancing Algorithm " -ForegroundColor Yellow
Get-VMSwitch | Get-VMSwitchTeam | fl

}
}

```

```

Host Name: AZSHCI1-M6-N1
Verifying SET Switch Load Balancing Algorithm

Name           : ConvergedSwitch(mgmt_compute)
Id             : 77c58161-731b-4a30-985d-a313a9bb7849
NetAdapterInterfaceDescription : {Mellanox ConnectX-6 Dx Adapter, Mellanox ConnectX-6 Dx Adapter #3}
NetAdapterInterfaceGuid : {6f4f6750-163a-49b9-b2ce-7420a3a7b018, 34acdad1-b66a-42ae-bcdb-24ef0590c665}
TeamingMode    : SwitchIndependent
LoadBalancingAlgorithm : HyperVPort

Host Name: AZSHCI1-M6-N2
Verifying SET Switch Load Balancing Algorithm

Name           : ConvergedSwitch(mgmt_compute)
Id             : cff44939-4dcd-45a3-98ff-6b4e04b60ac6
NetAdapterInterfaceDescription : {Mellanox ConnectX-6 Dx Adapter, Mellanox ConnectX-6 Dx Adapter #3}
NetAdapterInterfaceGuid : {09299664-3225-41bd-8325-8716f66415a1, a2040103-eea8-4efd-aea4-d0778bf85879}
TeamingMode    : SwitchIndependent
LoadBalancingAlgorithm : HyperVPort

Host Name: AZSHCI1-M6-N3
Verifying SET Switch Load Balancing Algorithm

Name           : ConvergedSwitch(mgmt_compute)
Id             : 1d9c623b-dc47-4435-83ef-f51165bfb558
NetAdapterInterfaceDescription : {Mellanox ConnectX-6 Dx Adapter #2, Mellanox ConnectX-6 Dx Adapter #3}
NetAdapterInterfaceGuid : {344fb850-563a-4efd-9164-45fda9f70ed0, 58f2d14d-8147-415b-bd10-3641118fd2fc}
TeamingMode    : SwitchIndependent
LoadBalancingAlgorithm : HyperVPort

```

Procedure 8. Configure Default Route Metric for Management NIC in Parent Partition

Step 1. Run the following:

```

$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Configuring default route metric for Management NIC " -ForegroundColor Yellow
netsh in ipv4 set ro 0.0.0.0/0 "vManagement(mgmt_compute_storage)" met=10
route print -4

}
}

```

```

Host Name: AZS-HCI1-N1
Configuring default route metric for Management NIC
Ok.

=====
Interface List
13...08 c0 eb 7e d0 c5 .....Hyper-V Virtual Ethernet Adapter
20...00 15 5d 7e 15 00 .....Hyper-V Virtual Ethernet Adapter #2
24...00 15 5d 7e 15 01 .....Hyper-V Virtual Ethernet Adapter #3
8...02 5c d5 62 7f bc .....Microsoft Failover Cluster Virtual Adapter
1.....Software Loopback Interface 1
=====

IPv4 Route Table
=====
Active Routes:
Network Destination        Netmask          Gateway          Interface        Metric
0.0.0.0                    0.0.0.0         192.168.126.1   192.168.126.21   15
127.0.0.0                  255.0.0.0       On-link         127.0.0.1        331
127.0.0.1                  255.255.255.255 On-link         127.0.0.1        331
127.255.255.255           255.255.255.255 On-link         127.0.0.1        331
169.254.0.0                255.255.255.0   On-link         169.254.209.93   261
169.254.0.0                255.255.255.0   On-link         169.254.53.86    261
169.254.53.86             255.255.255.255 On-link         169.254.53.86    261
169.254.209.93           255.255.255.255 On-link         169.254.209.93   261
169.254.255.255          255.255.255.255 On-link         169.254.209.93   261
169.254.255.255          255.255.255.255 On-link         169.254.53.86    261
192.168.126.0             255.255.255.192 On-link         192.168.126.21   261
192.168.126.21           255.255.255.255 On-link         192.168.126.21   261
192.168.126.63           255.255.255.255 On-link         192.168.126.21   261
224.0.0.0                 240.0.0.0       On-link         127.0.0.1        331
224.0.0.0                 240.0.0.0       On-link         192.168.126.21   261
224.0.0.0                 240.0.0.0       On-link         169.254.209.93   261
224.0.0.0                 240.0.0.0       On-link         169.254.53.86    261
255.255.255.255           255.255.255.255 On-link         127.0.0.1        331
255.255.255.255           255.255.255.255 On-link         192.168.126.21   261
255.255.255.255           255.255.255.255 On-link         169.254.209.93   261
255.255.255.255           255.255.255.255 On-link         169.254.53.86    261
=====
Persistent Routes:
Network Address            Netmask          Gateway Address  Metric
0.0.0.0                    0.0.0.0         192.168.126.1   Default
0.0.0.0                    0.0.0.0         192.168.126.1   256
0.0.0.0                    0.0.0.0         On-link         10
=====

```

Procedure 9. Configure Static NIC IP Address for Storage NICs

Note: Leave the gateway unconfigured for storage NICs.

Host	SMB NIC Name	SMB NIC IP Address
AZSHCI1-M6-N1	SlotID 5 Port 1 (Storage-A)	192.168.107.21
	SlotID 2 Port 2 (Storage-B)	192.168.207.21
AZSHCI1-M6-N2	SlotID 5 Port 1 (Storage-A)	192.168.107.22
	SlotID 2 Port 2 (Storage-B)	192.168.207.22

Host	SMB NIC Name	SMB NIC IP Address
AZSHCI1-M6-N3	SlotID 5 Port 1 (Storage-A)	192.168.107.23
	SlotID 2 Port 2 (Storage-B)	192.168.207.23

```

$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")

$IPStorageNetA = "192.168.107." #Storage-A(SlotID 5 Port 1)networkaddress
$IPStorageNetB = "192.168.207." #Storage-B(SlotID 2 Port 2)networkaddress
$IPIHostAddr = 21 #Starting host address

foreach ($node in $nodes) {
    $session = New-CimSession -ComputerName $node
    New-NetIPAddress -CimSession $session -InterfaceAlias "SlotID 5 Port 1" -IPAddress
    ($IPStorageNetA+$IPIHostAddr.ToString()) -PrefixLength 24

    New-NetIPAddress -CimSession $session -InterfaceAlias "SlotID 2 Port 2" -IPAddress
    ($IPStorageNetB+$IPIHostAddr.ToString()) -PrefixLength 24
    $IPIHostAddr++
}

Get-CimSession | Remove-CimSession
Remove-Variable session

```

Note: Network connectivity may be temporarily disrupted during the following configuration operations, but connectivity will automatically recover.

Procedure 10. Verify NIC IP Address for Storage NICs

Step 1. Run the following:

```

$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")
foreach ($node in $nodes) {

    Invoke-Command $node -Credential $Creds -scriptblock {
        write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

        Write-Host "Verifying Storage NIC IP Address " -ForegroundColor Yellow
        Get-NetIPConfiguration -InterfaceAlias vSMB* | fl InterfaceAlias, IPv4Address, IPv4DefaultGateway
    }
}

```



```

Host Name: AZSHCI1-M6-N1
Verifying Storage NIC IP Address

InterfaceAlias      : SlotID 5 Port 1
IPv4Address         : {192.168.107.21}
IPv4DefaultGateway :

InterfaceAlias      : SlotID 2 Port 2
IPv4Address         : {192.168.207.21}
IPv4DefaultGateway :

Host Name: AZSHCI1-M6-N2
Verifying Storage NIC IP Address

InterfaceAlias      : SlotID 5 Port 1
IPv4Address         : {192.168.107.22}
IPv4DefaultGateway :

InterfaceAlias      : SlotID 2 Port 2
IPv4Address         : {192.168.207.22}
IPv4DefaultGateway :

Host Name: AZSHCI1-M6-N3
Verifying Storage NIC IP Address

InterfaceAlias      : SlotID 5 Port 1
IPv4Address         : {192.168.107.23}
IPv4DefaultGateway :

InterfaceAlias      : SlotID 2 Port 2
IPv4Address         : {192.168.207.23}
IPv4DefaultGateway :

```

Procedure 11. Verify DNS Registration is Removed for Storage Interfaces

Step 1. Run the following:

```

$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Removing DNS Registration from Storage NICs " -ForegroundColor Yellow
Set-DnsClient -InterfaceAlias "SlotID 2 Port 2" -RegisterThisConnectionsAddress:$false
Set-DnsClient -InterfaceAlias "SlotID 5 Port 1" -RegisterThisConnectionsAddress:$false
Get-DnsClient -InterfaceAlias "SlotID 2 Port 2"| ft InterfaceAlias,RegisterThisConnectionsAddress
Get-DnsClient -InterfaceAlias "SlotID 5 Port 1"| ft InterfaceAlias,RegisterThisConnectionsAddress

}
}

```

```

Host Name: AZSHCI1-M6-N1
Removing DNS Registration from Storage NICs

InterfaceAlias RegisterThisConnectionsAddress
-----
SlotID 2 Port 2 False

InterfaceAlias RegisterThisConnectionsAddress
-----
SlotID 5 Port 1 False

Host Name: AZSHCI1-M6-N2
Removing DNS Registration from Storage NICs

InterfaceAlias RegisterThisConnectionsAddress
-----
SlotID 2 Port 2 False

InterfaceAlias RegisterThisConnectionsAddress
-----
SlotID 5 Port 1 False

Host Name: AZSHCI1-M6-N3
Removing DNS Registration from Storage NICs

InterfaceAlias RegisterThisConnectionsAddress
-----
SlotID 2 Port 2 False

InterfaceAlias RegisterThisConnectionsAddress
-----
SlotID 5 Port 1 False

```

Procedure 12. Enable Preserving 802.1p Priority Marking to Pass Through the vSwitch

Note: The virtual switch zeros-out 802.1p priority marking in the packet header. This is the default behavior. Preserving the 802.1p priority marking in the packet header is required for classifying and prioritizing network traffic in the fabric and other northbound switches that have QoS policies configured. This setting affects prioritized network traffic traversing the virtual switch. This setting is required for prioritizing Cluster Communication network traffic. RDMA traffic passing through RDMA enabled vNICs is not affected by this setting because this traffic bypasses the virtual switch and goes directly to the physical NIC.

Step 1. Run the following:

```

$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Configure vSwitch to pass 802.1p priority marking " -ForegroundColor Yellow
Set-VMNetworkAdapter -Name "vManagement(mgmt_compute_storage)" -ManagementOS -IeeePriorityTag On
Get-VMNetworkAdapter -ManagementOS | ft Name,IeeePriorityTag

}
}

```

```

Host Name: AZSHCI1-M6-N1
Configure vSwitch to pass 802.1p priority marking

Name                IeeePriorityTag
----                -
vManagement(mgmt_compute)  On

Host Name: AZSHCI1-M6-N2
Configure vSwitch to pass 802.1p priority marking

Name                IeeePriorityTag
----                -
vManagement(mgmt_compute)  On

Host Name: AZSHCI1-M6-N3
Configure vSwitch to pass 802.1p priority marking

Name                IeeePriorityTag
----                -
vManagement(mgmt_compute)  On

```

Procedure 13. Verify the Storage NIC VLANs

Step 1. Run the following:

```

$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verify NIC VLANs Configuration " -ForegroundColor Yellow
Get-VMNetworkAdapter -ManagementOS | Get-VMNetworkAdapterIsolation | FT IsolationMode, DefaultIsolationID,
ParentAdapter -AutoSize
Get-NetAdapter | select Name,InterfaceDescription,Status,MacAddress,VlanID | ft -AutoSize

}
}

```

```

Host Name: AZSHCI1-M6-N1
Verify vNIC VLANs Configuration

IsolationMode DefaultIsolationID ParentAdapter
-----
None 0 VMInternalNetworkAdapter, Name = 'vManagement(mgmt_compute)'
```

Name	InterfaceDescription	Status	MacAddress	VlanID
SlotID 2 Port 2	Mellanox ConnectX-6 Dx Adapter #2	Up	08-C0-EB-7E-D0-C4	207
SlotID 2 Port 1	Mellanox ConnectX-6 Dx Adapter	Up	08-C0-EB-7E-D0-C5	0
vManagement(mgmt_compute)	Hyper-V Virtual Ethernet Adapter	Up	08-C0-EB-7E-D0-C5	0
SlotID 5 Port 1	Mellanox ConnectX-6 Dx Adapter #4	Up	94-6D-AE-75-F1-23	107
SlotID 5 Port 2	Mellanox ConnectX-6 Dx Adapter #3	Up	94-6D-AE-75-F1-22	0

```

Host Name: AZSHCI1-M6-N2
Verify vNIC VLANs Configuration

IsolationMode DefaultIsolationID ParentAdapter
-----
None 0 VMInternalNetworkAdapter, Name = 'vManagement(mgmt_compute)'
```

Name	InterfaceDescription	Status	MacAddress	VlanID
SlotID 2 Port 2	Mellanox ConnectX-6 Dx Adapter #2	Up	08-C0-EB-7E-D2-EC	207
SlotID 5 Port 2	Mellanox ConnectX-6 Dx Adapter #3	Up	94-6D-AE-75-E6-1A	0
SlotID 5 Port 1	Mellanox ConnectX-6 Dx Adapter #4	Up	94-6D-AE-75-E6-1B	107
vManagement(mgmt_compute)	Hyper-V Virtual Ethernet Adapter	Up	08-C0-EB-7E-D2-ED	0
SlotID 2 Port 1	Mellanox ConnectX-6 Dx Adapter	Up	08-C0-EB-7E-D2-ED	0

```

Host Name: AZSHCI1-M6-N3
Verify vNIC VLANs Configuration

IsolationMode DefaultIsolationID ParentAdapter
-----
None 0 VMInternalNetworkAdapter, Name = 'vManagement(mgmt_compute)'
```

Name	InterfaceDescription	Status	MacAddress	VlanID
SlotID 2 Port 2	Mellanox ConnectX-6 Dx Adapter	Up	08-C0-EB-7E-D0-BD	207
SlotID 5 Port 1	Mellanox ConnectX-6 Dx Adapter #4	Up	94-6D-AE-75-E8-FB	107
SlotID 5 Port 2	Mellanox ConnectX-6 Dx Adapter #3	Up	94-6D-AE-75-E8-FA	0
vManagement(mgmt_compute)	Hyper-V Virtual Ethernet Adapter	Up	08-C0-EB-7E-D0-BC	0
SlotID 2 Port 1	Mellanox ConnectX-6 Dx Adapter #2	Up	08-C0-EB-7E-D0-BC	0

Procedure 14. Verify Network Adapters

Step 1. Run the following:

```

$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verifying NIC status " -ForegroundColor Yellow
Get-NetAdapter | sort Name | ft Name,InterfaceDescription,Status,MTUSize,LinkSpeed

}
}

```

```

Host Name: AZSHCI1-M6-N1
Verifying Virtual Switch

Name                : ConvergedSwitch(mgmt_compute)
SwitchType          : External
NetAdapterInterfaceDescription : Teamed-Interface
NetAdapterInterfaceDescriptions : {Mellanox ConnectX-6 Dx Adapter, Mellanox ConnectX-6 Dx Adapter #3}

Verifying Management vNIC in parent partition

Name                InterfaceDescription      Status MacAddress      LinkSpeed
-----
SlotID 2 Port 2    Mellanox ConnectX-6 Dx Adapter #2 Up    08-C0-EB-7E-D0-C4 100 Gbps
SlotID 2 Port 1    Mellanox ConnectX-6 Dx Adapter Up    08-C0-EB-7E-D0-C5 100 Gbps
vManagement(mgmt_compute) Hyper-V Virtual Ethernet Adapter Up    08-C0-EB-7E-D0-C5 100 Gbps
SlotID 5 Port 1    Mellanox ConnectX-6 Dx Adapter #4 Up    94-6D-AE-75-F1-23 100 Gbps
SlotID 5 Port 2    Mellanox ConnectX-6 Dx Adapter #3 Up    94-6D-AE-75-F1-22 100 Gbps

Host Name: AZSHCI1-M6-N2
Verifying Virtual Switch

Name                : ConvergedSwitch(mgmt_compute)
SwitchType          : External
NetAdapterInterfaceDescription : Teamed-Interface
NetAdapterInterfaceDescriptions : {Mellanox ConnectX-6 Dx Adapter, Mellanox ConnectX-6 Dx Adapter #3}

Verifying Management vNIC in parent partition

Name                InterfaceDescription      Status MacAddress      LinkSpeed
-----
SlotID 2 Port 2    Mellanox ConnectX-6 Dx Adapter #2 Up    08-C0-EB-7E-D2-EC 100 Gbps
SlotID 5 Port 2    Mellanox ConnectX-6 Dx Adapter #3 Up    94-6D-AE-75-E6-1A 100 Gbps
SlotID 5 Port 1    Mellanox ConnectX-6 Dx Adapter #4 Up    94-6D-AE-75-E6-1B 100 Gbps
vManagement(mgmt_compute) Hyper-V Virtual Ethernet Adapter Up    08-C0-EB-7E-D2-ED 100 Gbps
SlotID 2 Port 1    Mellanox ConnectX-6 Dx Adapter Up    08-C0-EB-7E-D2-ED 100 Gbps

Host Name: AZSHCI1-M6-N3
Verifying Virtual Switch

Name                : ConvergedSwitch(mgmt_compute)
SwitchType          : External
NetAdapterInterfaceDescription : Teamed-Interface
NetAdapterInterfaceDescriptions : {Mellanox ConnectX-6 Dx Adapter #2, Mellanox ConnectX-6 Dx Adapter #3}

Verifying Management vNIC in parent partition

Name                InterfaceDescription      Status MacAddress      LinkSpeed
-----
SlotID 2 Port 2    Mellanox ConnectX-6 Dx Adapter Up    08-C0-EB-7E-D0-BD 100 Gbps
SlotID 5 Port 1    Mellanox ConnectX-6 Dx Adapter #4 Up    94-6D-AE-75-E8-FB 100 Gbps
SlotID 5 Port 2    Mellanox ConnectX-6 Dx Adapter #3 Up    94-6D-AE-75-E8-FA 100 Gbps
vManagement(mgmt_compute) Hyper-V Virtual Ethernet Adapter Up    08-C0-EB-7E-D0-BC 100 Gbps
SlotID 2 Port 1    Mellanox ConnectX-6 Dx Adapter #2 Up    08-C0-EB-7E-D0-BC 100 Gbps

```

Procedure 15. Verify RDMA and RoCEv2 Protocol is Enabled on Physical NICs

Step 1. Run the following:

```

$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verifying RDMA and RoCEv2 status on physical NICS " -ForegroundColor Yellow

Get-NetAdapterAdvancedProperty -InterfaceDescription "Mellanox ConnectX*" -DisplayName "NetworkDirect*" | ft
Name, InterfaceDescription, DisplayName, DisplayValue

}
}

```

```

Host Name: AZSHCI1-M6-N1
Verifying RDMA and RoCEv2 status on physical NICs
-----
Name                InterfaceDescription                DisplayName                DisplayValue
-----
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Functionality Enabled
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Technology RoCEv2
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter NetworkDirect Functionality Enabled
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter NetworkDirect Technology RoCEv2
SlotID 5 Port 1 Mellanox ConnectX-6 Dx Adapter #4 NetworkDirect Functionality Enabled
SlotID 5 Port 1 Mellanox ConnectX-6 Dx Adapter #4 NetworkDirect Technology RoCEv2
SlotID 5 Port 2 Mellanox ConnectX-6 Dx Adapter #3 NetworkDirect Functionality Enabled
SlotID 5 Port 2 Mellanox ConnectX-6 Dx Adapter #3 NetworkDirect Technology RoCEv2

Host Name: AZSHCI1-M6-N2
Verifying RDMA and RoCEv2 status on physical NICs
-----
Name                InterfaceDescription                DisplayName                DisplayValue
-----
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Functionality Enabled
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Technology RoCEv2
SlotID 5 Port 2 Mellanox ConnectX-6 Dx Adapter #3 NetworkDirect Functionality Enabled
SlotID 5 Port 2 Mellanox ConnectX-6 Dx Adapter #3 NetworkDirect Technology RoCEv2
SlotID 5 Port 1 Mellanox ConnectX-6 Dx Adapter #4 NetworkDirect Functionality Enabled
SlotID 5 Port 1 Mellanox ConnectX-6 Dx Adapter #4 NetworkDirect Technology RoCEv2
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter NetworkDirect Functionality Enabled
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter NetworkDirect Technology RoCEv2

Host Name: AZSHCI1-M6-N3
Verifying RDMA and RoCEv2 status on physical NICs
-----
Name                InterfaceDescription                DisplayName                DisplayValue
-----
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter NetworkDirect Functionality Enabled
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter NetworkDirect Technology RoCEv2
SlotID 5 Port 1 Mellanox ConnectX-6 Dx Adapter #4 NetworkDirect Functionality Enabled
SlotID 5 Port 1 Mellanox ConnectX-6 Dx Adapter #4 NetworkDirect Technology RoCEv2
SlotID 5 Port 2 Mellanox ConnectX-6 Dx Adapter #3 NetworkDirect Functionality Enabled
SlotID 5 Port 2 Mellanox ConnectX-6 Dx Adapter #3 NetworkDirect Technology RoCEv2
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Functionality Enabled
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Technology RoCEv2

```

Procedure 16. Verify RDMA is Enabled on the Storage vNIC Adapters

Step 1. Run the following:

```

$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verifying that RDMA is enabled on the Storage vNICs" -ForegroundColor Yellow
Get-NetAdapterRdma | ft

}

}

```

```

Host Name: AZSHCI1-M6-N1
Verifying that RDMA is enabled on the Storage vNICs
Name                InterfaceDescription      Enabled  Operational  PFC     ETS
-----
SlotID 2 Port 2    Mellanox ConnectX-6 Dx Adapter #2  True    True         True    True
SlotID 2 Port 1    Mellanox ConnectX-6 Dx Adapter      True    True         True    True
VManagement(mgmt_compute) Hyper-V Virtual Ethernet Adapter    False   False        NA     NA
SlotID 5 Port 1    Mellanox ConnectX-6 Dx Adapter #4   True    True         True    True
SlotID 5 Port 2    Mellanox ConnectX-6 Dx Adapter #3   True    True         True    True

Host Name: AZSHCI1-M6-N2
Verifying that RDMA is enabled on the Storage vNICs
Name                InterfaceDescription      Enabled  Operational  PFC     ETS
-----
SlotID 2 Port 2    Mellanox ConnectX-6 Dx Adapter #2  True    True         True    True
SlotID 5 Port 2    Mellanox ConnectX-6 Dx Adapter #3   True    True         True    True
SlotID 5 Port 1    Mellanox ConnectX-6 Dx Adapter #4   True    True         True    True
VManagement(mgmt_compute) Hyper-V Virtual Ethernet Adapter    False   False        NA     NA
SlotID 2 Port 1    Mellanox ConnectX-6 Dx Adapter      True    True         True    True

Host Name: AZSHCI1-M6-N3
Verifying that RDMA is enabled on the Storage vNICs
Name                InterfaceDescription      Enabled  Operational  PFC     ETS
-----
SlotID 2 Port 2    Mellanox ConnectX-6 Dx Adapter      True    True         True    True
SlotID 5 Port 1    Mellanox ConnectX-6 Dx Adapter #4   True    True         True    True
SlotID 5 Port 2    Mellanox ConnectX-6 Dx Adapter #3   True    True         True    True
VManagement(mgmt_compute) Hyper-V Virtual Ethernet Adapter    False   False        NA     NA
SlotID 2 Port 1    Mellanox ConnectX-6 Dx Adapter #2   True    True         True    True

```

Procedure 17. Verify RDMA Capabilities

Step 1. Run the following:

```

$nodes = ("AZSHCI1-M6-N1", "AZSHCI1-M6-N2", "AZSHCI1-M6-N3")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verify Storage vNIC RDMA operational status " -ForegroundColor Yellow
Get-SmbClientNetworkInterface | ft FriendlyName, RDMACapable

}
}

```

```
Host Name: AZSHCI1-M6-N1
Verify Storage vNIC RDMA operational status

FriendlyName          RDMACapable
-----
vManagement(mgmt_compute)  False
SlotID 5 Port 1          True
SlotID 2 Port 2          True
Local Area Connection* 1  False

Host Name: AZSHCI1-M6-N2
Verify Storage vNIC RDMA operational status

FriendlyName          RDMACapable
-----
vManagement(mgmt_compute)  False
SlotID 2 Port 2          True
SlotID 5 Port 1          True
Local Area Connection* 1  False

Host Name: AZSHCI1-M6-N3
Verify Storage vNIC RDMA operational status

FriendlyName          RDMACapable
-----
SlotID 5 Port 1          True
vManagement(mgmt_compute)  False
SlotID 2 Port 2          True
Local Area Connection* 1  False
```

Storage Switchless Network Topology Configuration for 2 Nodes

Note: Skip the steps in this section if the deployment model chosen for your environment is a fully converged or non-converged topology. Only one deployment model needs to be followed, i.e., either fully converged or non-converged or storage switchless topology.

This subject contains the following procedures to create and deploy standalone Network ATC intent for a non-converged topology. Here you create two intents - one intent for “Management and Compute” networks, and a separate intent for “Storage” Network:

- [Identify Physical Network Card Port Names](#)
- [Create and Deploy Standalone Network ATC Intent - Management and Compute Only](#)
- [Verify Network ATC Intent Status](#)
- [Create and Deploy Standalone Network ATC Intent - Storage Only](#)
- [Verify Network ATC Intent Status](#)
- [Verify Virtual Switch and Virtual NIC Creation in the Parent Partition](#)
- [Verify SET Switch Team Load Balancing Algorithm](#)
- [Configure Default Route Metric for Management NIC in Parent Partition](#)
- [Configure Static NIC IP Address for Storage NICs](#)
- [Verify NIC IP Address for Storage NICs](#)
- [Verify DNS Registration is Removed for Storage Interfaces](#)
- [Enable Preserving 802.1p Priority Marking to Pass Through the vSwitch](#)
- [Verify the Storage vNIC VLANs](#)

- [Verify Network Adapters](#)
- [Verify RDMA and RoCEv2 Protocol is Enabled on Physical NICs](#)
- [Verify RDMA is Enabled on the Storage NIC Adapters](#)
- [Verify RDMA Capabilities](#)

Procedure 1. Identify Physical Network Card Port Names

Step 1. Run the following:

```
$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")
foreach ($node in $nodes) {
Invoke-Command $node -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Retrieving physical NIC port names " -ForegroundColor Yellow

Get-netadapter | ft Name, InterfaceDescription, Status, MacAddress, LinkSpeed

}
}
```

```
Host Name: AZSHCI-SL-N3
Retrieving physical NIC port names

Name                InterfaceDescription      Status MacAddress      LinkSpeed
-----
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter Up      08-C0-EB-7E-D1-38 100 Gbps
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #3 Up      08-C0-EB-7E-D1-39 100 Gbps
SlotID 5 Port 1 Mellanox ConnectX-6 Dx Adapter #4 Up      94-6D-AE-75-E6-0E 100 Gbps
SlotID 5 Port 2 Mellanox ConnectX-6 Dx Adapter #2 Up      94-6D-AE-75-E6-0F 100 Gbps

Host Name: AZSHCI-SL-N4
Retrieving physical NIC port names

Name                InterfaceDescription      Status MacAddress      LinkSpeed
-----
SlotID 5 Port 2 Mellanox ConnectX-6 Dx Adapter #3 Up      94-6D-AE-75-E5-F3 100 Gbps
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter Up      08-C0-EB-7E-D1-24 100 Gbps
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #2 Up      08-C0-EB-7E-D1-25 100 Gbps
SlotID 5 Port 1 Mellanox ConnectX-6 Dx Adapter #4 Up      94-6D-AE-75-E5-F2 100 Gbps
```

Note: If the NIC port names are “Ethernet” and “Ethernet 2”, CDN is not enabled. CDN (Consistent Device Naming) must be enabled for correct physical to virtual NIC mapping later in this guide.

Procedure 2. Create and Deploy Standalone Network ATC Intent – Management and Compute only

Step 1. Run the following script block to create a virtual switch with SET enabled and three virtual NICs:

```
$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Create and Deploy Standalone NetworkATC Intent - Mgmt and Compute only " -ForegroundColor Yellow
```

```

$Mgmt_Compute_NICs = @("SlotID 2 Port 1", "SlotID 2 Port 2")

Add-NetIntent -Name Mgmt_Compute -AdapterName $Mgmt_Compute_NICs -Management -Compute -verbose

}

}

```

Note: There will be a brief network disconnect on each server node when VM switch binds to the physical adapters.

Procedure 3. Verify Network ATC Intent Status

Step 1. Run the following:

```

$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Checking Network ATC Intent Status" -ForegroundColor Yellow

Get-netIntentStatus -ComputerName $node | ft
Host,IntentName,ConfigurationStatus,ProvisioningStatus,IsComputeIntentSet,IsManagementIntentSet,IsStorageInte
ntset,IsStretchIntentSet

}

}

```

```

Host Name: AZSHCI-SL-N3
Checking Network ATC Intent Status
WARNING: If you are running Windows PowerShell remotely, note that some failover clustering cmdlets do not work remotely. When possible, run the cmdlet locally and specify a remote computer as the target. To run the cmdlet remotely, try using the Credential Security Service Provider (CredSSP). All additional errors or warnings from this cmdlet might be caused by running it remotely.

Host      IntentName  ConfigurationStatus  ProvisioningStatus  IsComputeIntentSet  IsManagementIntentSet  IsStorageIntentSet  IsStretchIntentSet
-----
AZSHCI-SL-N3  mgmt_compute  Success              Completed           True                 True                    False                False

Host Name: AZSHCI-SL-N4
Checking Network ATC Intent Status
WARNING: If you are running Windows PowerShell remotely, note that some failover clustering cmdlets do not work remotely. When possible, run the cmdlet locally and specify a remote computer as the target. To run the cmdlet remotely, try using the Credential Security Service Provider (CredSSP). All additional errors or warnings from this cmdlet might be caused by running it remotely.

Host      IntentName  ConfigurationStatus  ProvisioningStatus  IsComputeIntentSet  IsManagementIntentSet  IsStorageIntentSet  IsStretchIntentSet
-----
AZSHCI-SL-N4  mgmt_compute  Success              Completed           True                 True                    False                False

```

Procedure 4. Create and Deploy Standalone Network ATC Intent - Storage only

Step 1. Run the following script block to create a virtual switch with SET enabled and three virtual NICs:

```

$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Create and Deploy Standalone Network ATC Intent - Storage Only " -ForegroundColor Yellow

$Storage_NICs = @("SlotID 5 Port 1", "SlotID 5 Port 2")
$Storage_VLANS = @(107,207)

$QoSOverride = New-NetIntentQoSPolicyOverRides
$QoSOverride.PriorityValue8021Action_SMB = 4
$QoSOverride.PriorityValue8021Action_Cluster = 5

$AdapterOverride = New-NetIntentAdapterPropertyOverrides
$AdapterOverride.NetworkDirectTechnology = 4
$AdapterOverride.JumboPacket = 9014

$storageOverride = new-NetIntentStorageOverrides
$storageOverride.EnableAutomaticIPGeneration = $false

Add-NetIntent -Name Storage -AdapterName $Storage_NICs -Storage -StorageVlans $Storage_VLANS -
QoSPolicyOverrides $QoSOverride -AdapterPropertyOverrides $AdapterOverride -StorageOverrides $storageoverride

}
}

```

Procedure 5. Verify Network ATC Intent Status

Step 1. Run the following:

```

$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Checking Network ATC Intent Status" -ForegroundColor Yellow

Get-netIntentStatus -ComputerName $node | ft
Host,IntentName,ConfigurationStatus,ProvisioningStatus,IsComputeIntentSet,IsManagementIntentSet,IsStorageInte
ntset,IsStretchIntentSet

```

```
}  
}
```

```
Host Name: AZSHCI-SL-N3  
Checking Network ATC Intent Status  
WARNING: If you are running Windows PowerShell remotely, note that some failover clustering cmdlets do not work remotely. When possible, run the cmdlet locally and specify a remote computer as the target. To run the cmdlet remotely, try using the Credential Security Service Provider (CredSSP). All additional errors or warnings from this cmdlet might be caused by running it remotely.  
  
Completed  


| Host         | IntentName   | ConfigurationStatus | ProvisioningStatus | IsComputeIntentSet | IsManagementIntentSet | IsStorageIntentSet | IsStretchIntentSet |
|--------------|--------------|---------------------|--------------------|--------------------|-----------------------|--------------------|--------------------|
| AZSHCI-SL-N3 | mgmt_compute | Success             | Completed          | True               | True                  | False              | False              |
| AZSHCI-SL-N3 | storage      | Success             | Completed          | False              | False                 | True               | False              |

  
Host Name: AZSHCI-SL-N4  
Checking Network ATC Intent Status  
WARNING: If you are running Windows PowerShell remotely, note that some failover clustering cmdlets do not work remotely. When possible, run the cmdlet locally and specify a remote computer as the target. To run the cmdlet remotely, try using the Credential Security Service Provider (CredSSP). All additional errors or warnings from this cmdlet might be caused by running it remotely.  
  
Completed  


| Host         | IntentName   | ConfigurationStatus | ProvisioningStatus | IsComputeIntentSet | IsManagementIntentSet | IsStorageIntentSet | IsStretchIntentSet |
|--------------|--------------|---------------------|--------------------|--------------------|-----------------------|--------------------|--------------------|
| AZSHCI-SL-N4 | mgmt_compute | Success             | Completed          | True               | True                  | False              | False              |
| AZSHCI-SL-N4 | storage      | Success             | Completed          | False              | False                 | True               | False              |


```

Procedure 6. Verify Virtual Switch and Virtual NIC Creation in the Parent Partition

Step 1. Run the following:

```
$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")  
foreach ($node in $nodes) {  
  
Invoke-Command $node -Credential $Creds -scriptblock {  
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green  
  
Write-Host " Verifying Virtual Switch " -ForegroundColor Yellow  
Get-VMSwitch | fl Name, SwitchType, NetAdapterInterfaceDescription, NetAdapterInterfaceDescriptions  
  
Write-Host " Verifying Management vNIC in parent partition " -ForegroundColor Yellow  
Get-netadapter | ft Name, InterfaceDescription, Status, MacAddress, LinkSpeed  
  
}  
}
```

```

Host Name: AZSHCI-SL-N3
Verifying Virtual Switch

Name                : ConvergedSwitch(mgmt_compute)
SwitchType          : External
NetAdapterInterfaceDescription : Teamed-Interface
NetAdapterInterfaceDescriptions : {Mellanox ConnectX-6 Dx Adapter #3, Mellanox ConnectX-6 Dx Adapter}

Verifying Management vNIC in parent partition

Name                InterfaceDescription      Status MacAddress      LinkSpeed
----                -
SlotID 2 Port 2    Mellanox ConnectX-6 Dx Adapter #3 Up    08-C0-EB-7E-D1-38 100 Gbps
SlotID 2 Port 1    Mellanox ConnectX-6 Dx Adapter #3 Up    08-C0-EB-7E-D1-39 100 Gbps
SlotID 5 Port 1    Mellanox ConnectX-6 Dx Adapter #4 Up    94-6D-AE-75-E6-0E 100 Gbps
SlotID 5 Port 2    Mellanox ConnectX-6 Dx Adapter #2 Up    94-6D-AE-75-E6-0F 100 Gbps
vManagement(mgmt_compute) Hyper-V Virtual Ethernet Adapter Up    08-C0-EB-7E-D1-39 100 Gbps

Host Name: AZSHCI-SL-N4
Verifying Virtual Switch

Name                : ConvergedSwitch(mgmt_compute)
SwitchType          : External
NetAdapterInterfaceDescription : Teamed-Interface
NetAdapterInterfaceDescriptions : {Mellanox ConnectX-6 Dx Adapter #2, Mellanox ConnectX-6 Dx Adapter}

Verifying Management vNIC in parent partition

Name                InterfaceDescription      Status MacAddress      LinkSpeed
----                -
SlotID 5 Port 2    Mellanox ConnectX-6 Dx Adapter #3 Up    94-6D-AE-75-E5-F3 100 Gbps
SlotID 2 Port 2    Mellanox ConnectX-6 Dx Adapter #3 Up    08-C0-EB-7E-D1-24 100 Gbps
SlotID 2 Port 1    Mellanox ConnectX-6 Dx Adapter #2 Up    08-C0-EB-7E-D1-25 100 Gbps
vManagement(mgmt_compute) Hyper-V Virtual Ethernet Adapter Up    08-C0-EB-7E-D1-25 100 Gbps
SlotID 5 Port 1    Mellanox ConnectX-6 Dx Adapter #4 Up    94-6D-AE-75-E5-F2 100 Gbps

```

Procedure 7. Verify SET Switch Team Load Balancing Algorithm

Note: The load balancing algorithm must be a Hyper-V Port. Each VM switch must be bound to both physical network adapters.

Step 1. Run the following:

```

$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Verifying SET Switch Load Balancing Algorithm " -ForegroundColor Yellow
Get-VMSwitch | Get-VMSwitchTeam | fl

}
}

```

```

Host Name: AZSHCI-SL-N3
Verifying SET Switch Load Balancing Algorithm

Name           : ConvergedSwitch(mgmt_compute)
Id             : a3965fdc-e0d1-4316-8fe0-26ffb466cda8
NetAdapterInterfaceDescription : {Mellanox ConnectX-6 Dx Adapter #3, Mellanox ConnectX-6 Dx Adapter}
NetAdapterInterfaceGuid : {9e3557ff-a22f-4a63-b62d-a6fd2d151878, a28b2744-bdba-4596-88e7-b63a4dad69f3}
TeamingMode    : SwitchIndependent
LoadBalancingAlgorithm : HyperVPort

Host Name: AZSHCI-SL-N4
Verifying SET Switch Load Balancing Algorithm

Name           : ConvergedSwitch(mgmt_compute)
Id             : d69ed292-40ba-4851-816a-19922016130d
NetAdapterInterfaceDescription : {Mellanox ConnectX-6 Dx Adapter #2, Mellanox ConnectX-6 Dx Adapter}
NetAdapterInterfaceGuid : {455aedfd-f6f5-4cfe-a265-1b3ce0f5b215, 5e6dd61a-baff-480c-8eae-b8d10db9dd2f}
TeamingMode    : SwitchIndependent
LoadBalancingAlgorithm : HyperVPort

```

Procedure 8. Configure Default Route Metric for Management NIC in Parent Partition

Step 1. Run the following:

```

$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Configuring default route metric for Management NIC " -ForegroundColor Yellow
netsh in ipv4 set ro 0.0.0.0/0 "vManagement(mgmt_compute_storage)" met=10
route print -4

}

}

```

```

Host Name: AZSHCI-SL-N3
Configuring default route metric for Management NIC
Ok.

=====
Interface List
5...94 6d ae 75 e6 0f .....Mellanox ConnectX-6 Dx Adapter #2
7...94 6d ae 75 e6 0e .....Mellanox ConnectX-6 Dx Adapter #4
15...08 c0 eb 7e d1 39 .....Hyper-V Virtual Ethernet Adapter
4...02 ee 5b 3d 52 ce .....Microsoft Failover Cluster Virtual Adapter
1.....Software Loopback Interface 1
=====

IPv4 Route Table
=====
Active Routes:
Network Destination        Netmask          Gateway          Interface        Metric
0.0.0.0                    0.0.0.0          192.168.126.1   192.168.126.13   15
127.0.0.0                  255.0.0.0        On-link         127.0.0.1        331
127.0.0.1                  255.255.255.255 On-link         127.0.0.1        331
127.255.255.255           255.255.255.255 On-link         127.0.0.1        331
169.254.0.0                255.255.0.0      On-link         169.254.25.1     261
169.254.25.1              255.255.255.255 On-link         169.254.25.1     261
169.254.255.255           255.255.255.255 On-link         169.254.25.1     261
192.168.107.0              255.255.255.0    On-link         192.168.107.13   261
192.168.107.13            255.255.255.255 On-link         192.168.107.13   261
192.168.107.255           255.255.255.255 On-link         192.168.107.13   261
192.168.126.0              255.255.255.192 On-link         192.168.126.13   261
192.168.126.13            255.255.255.255 On-link         192.168.126.13   261
192.168.126.63            255.255.255.255 On-link         192.168.126.13   261
224.0.0.0                  240.0.0.0        On-link         127.0.0.1        331
224.0.0.0                  240.0.0.0        On-link         192.168.126.13   261
224.0.0.0                  240.0.0.0        On-link         192.168.107.13   261
224.0.0.0                  240.0.0.0        On-link         169.254.25.1     261
255.255.255.255           255.255.255.255 On-link         127.0.0.1        331
255.255.255.255           255.255.255.255 On-link         192.168.126.13   261
255.255.255.255           255.255.255.255 On-link         192.168.107.13   261
255.255.255.255           255.255.255.255 On-link         169.254.25.1     261
=====

Persistent Routes:
Network Address          Netmask  Gateway Address  Metric
0.0.0.0                  0.0.0.0  192.168.126.1   Default
0.0.0.0                  0.0.0.0  192.168.126.1   256
0.0.0.0                  0.0.0.0  On-link         10
=====

```

Procedure 9. Configure Static NIC IP Address for Storage NICs

Note: Leave the gateway unconfigured for storage NICs.

Host	SMB NIC Name	SMB NIC IP Address
AZSHCI-SL-N3	SlotID 5 Port 1 (Storage-A)	192.168.107.13
	SlotID 5 Port 2 (Storage-B)	192.168.207.13
AZSHCI1-M6-N2	SlotID 5 Port 1 (Storage-A)	192.168.107.14
	SlotID 5 Port 2 (Storage-B)	192.168.207.14

```

$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")

$IPStorageNetA = "192.168.107." #Storage-A(SlotID 5 Port 1)networkaddress
$IPStorageNetB = "192.168.207." #Storage-B(SlotID 5 Port 2)networkaddress
$IPHostAddr = 13 #Starting host address

foreach ($node in $nodes) {
    $session = New-CimSession -ComputerName $node
    New-NetIPAddress -CimSession $session -InterfaceAlias "SlotID 5 Port 1" -IPAddress
    ($IPStorageNetA+$IPHostAddr.ToString()) -PrefixLength 24
}

```

```

New-NetIPAddress -CimSession $session -InterfaceAlias "SlotID 5 Port 2" -IPAddress
($IPStorageNetB+$IPHostAddr.ToString()) -PrefixLength 24

$IPHostAddr++

}

Get-CimSession | Remove-CimSession

Remove-Variable session

```

Note: Network connectivity may be temporarily disrupted during the following configuration operations, but connectivity will automatically recover.

Procedure 10. Verify NIC IP Address for Storage NICs

Step 1. Run the following:

```

$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verifying Storage NIC IP Address " -ForegroundColor Yellow
Get-NetIPConfiguration -InterfaceAlias vSMB* | fl InterfaceAlias, IPv4Address, IPv4DefaultGateway

}

}

```

```

Host Name: AZSHCI-SL-N3
Verifying Storage NIC IP Address

InterfaceAlias      : SlotID 5 Port 1
IPv4Address         : {192.168.107.13}
IPv4DefaultGateway :

InterfaceAlias      : SlotID 5 Port 2
IPv4Address         : {192.168.207.13}
IPv4DefaultGateway :

Host Name: AZSHCI-SL-N4
Verifying Storage NIC IP Address

InterfaceAlias      : SlotID 5 Port 1
IPv4Address         : {192.168.107.14}
IPv4DefaultGateway :

InterfaceAlias      : SlotID 5 Port 2
IPv4Address         : {192.168.207.14}
IPv4DefaultGateway :

```

Procedure 11. Verify DNS Registration is Removed for Storage Interfaces

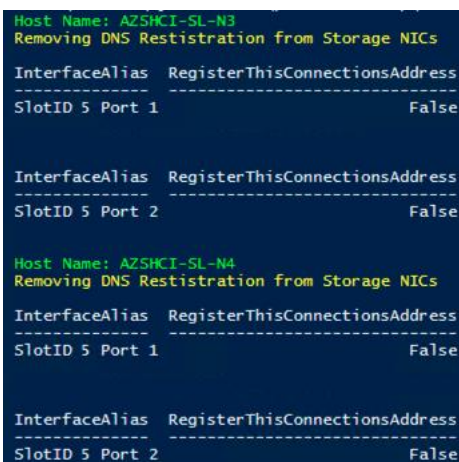
Step 1. Run the following:

```
$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Removing DNS Restistration from Storage NICs " -ForegroundColor Yellow
Set-DnsClient -InterfaceAlias "SlotID 5 Port 1" -RegisterThisConnectionsAddress:$false
Set-DnsClient -InterfaceAlias "SlotID 5 Port 2" -RegisterThisConnectionsAddress:$false
Get-DnsClient -InterfaceAlias "SlotID 5 Port 1"| ft InterfaceAlias,RegisterThisConnectionsAddress
Get-DnsClient -InterfaceAlias "SlotID 5 Port 2"| ft InterfaceAlias,RegisterThisConnectionsAddress

}
}
```



```
Host Name: AZSHCI-SL-N3
Removing DNS Restistration from Storage NICs

InterfaceAlias  RegisterThisConnectionsAddress
-----
SlotID 5 Port 1                False

InterfaceAlias  RegisterThisConnectionsAddress
-----
SlotID 5 Port 2                False

Host Name: AZSHCI-SL-N4
Removing DNS Restistration from Storage NICs

InterfaceAlias  RegisterThisConnectionsAddress
-----
SlotID 5 Port 1                False

InterfaceAlias  RegisterThisConnectionsAddress
-----
SlotID 5 Port 2                False
```

Procedure 12. Enable Preserving 802.1p Priority Marking to Pass Through the vSwitch

Note: The virtual switch zeros-out 802.1p priority marking in the packet header. This is the default behavior. Preserving the 802.1p priority marking in the packet header is required for classifying and prioritizing network traffic in the fabric and other northbound switches that have QoS policies configured. This setting affects prioritized network traffic traversing the virtual switch. This setting is required prioritizing Cluster Communication network traffic. RDMA traffic passing through RDMA enabled vNICs is not affected by this setting because this traffic bypasses the virtual switch and goes directly to the physical NIC.

Step 1. Run the following:

```
$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
```

```

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Configure vSwitch to pass 802.1p priority marking " -ForegroundColor Yellow
Set-VMNetworkAdapter -Name "vManagement(mgmt_compute)" -ManagementOS -IeeePriorityTag On
Get-VMNetworkAdapter -ManagementOS | ft Name,IeeePriorityTag

}

}

```

```

Host Name: AZSHCI-SL-N3
Configure vSwitch to pass 802.1p priority marking

Name                IeeePriorityTag
----                -
vManagement(mgmt_compute)  On

Host Name: AZSHCI-SL-N4
Configure vSwitch to pass 802.1p priority marking

Name                IeeePriorityTag
----                -
vManagement(mgmt_compute)  On

```

Procedure 13. Verify the Storage NIC VLANs

Step 1. Run the following:

```

$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verify NIC VLANs Configuration " -ForegroundColor Yellow
Get-VMNetworkAdapter -ManagementOS | Get-VMNetworkAdapterIsolation | FT IsolationMode, DefaultIsolationID,
ParentAdapter -AutoSize
Get-NetAdapter | select Name,InterfaceDescription,Status,MacAddress,VlanID | ft -AutoSize

}

}

```

```

Host Name: AZSHCI-SL-N3
Verify vNIC VLANs Configuration

IsolationMode DefaultIsolationID ParentAdapter
-----
None 0 VMInternalNetworkAdapter, Name = 'vManagement(mgmt_compute)'
```

Name	InterfaceDescription	Status	MacAddress	VlanID
SlotID 2 Port 2	Mellanox ConnectX-6 Dx Adapter	Up	08-C0-EB-7E-D1-38	0
SlotID 2 Port 1	Mellanox ConnectX-6 Dx Adapter #3	Up	08-C0-EB-7E-D1-39	0
SlotID 5 Port 1	Mellanox ConnectX-6 Dx Adapter #4	Up	94-6D-AE-75-E6-0E	107
SlotID 5 Port 2	Mellanox ConnectX-6 Dx Adapter #2	Up	94-6D-AE-75-E6-0F	207
vManagement(mgmt_compute)	Hyper-V Virtual Ethernet Adapter	Up	08-C0-EB-7E-D1-39	

```

Host Name: AZSHCI-SL-N4
Verify vNIC VLANs Configuration

IsolationMode DefaultIsolationID ParentAdapter
-----
None 0 VMInternalNetworkAdapter, Name = 'vManagement(mgmt_compute)'
```

Name	InterfaceDescription	Status	MacAddress	VlanID
SlotID 5 Port 2	Mellanox ConnectX-6 Dx Adapter #3	Up	94-6D-AE-75-E5-F3	207
SlotID 2 Port 2	Mellanox ConnectX-6 Dx Adapter	Up	08-C0-EB-7E-D1-24	0
SlotID 2 Port 1	Mellanox ConnectX-6 Dx Adapter #2	Up	08-C0-EB-7E-D1-25	0
vManagement(mgmt_compute)	Hyper-V Virtual Ethernet Adapter	Up	08-C0-EB-7E-D1-25	
SlotID 5 Port 1	Mellanox ConnectX-6 Dx Adapter #4	Up	94-6D-AE-75-E5-F2	107

Procedure 14. Verify Network Adapters

Step 1. Run the following:

```

$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verifying NIC status " -ForegroundColor Yellow

Get-NetAdapter | sort Name | ft Name,InterfaceDescription,Status,MTUSize,LinkSpeed

}

}

```

```

Host Name: AZSHCI-SL-N3
Verifying NIC status
```

Name	InterfaceDescription	Status	MTUSize	LinkSpeed
SlotID 2 Port 1	Mellanox ConnectX-6 Dx Adapter #3	Up	1500	100 Gbps
SlotID 2 Port 2	Mellanox ConnectX-6 Dx Adapter	Up	1500	100 Gbps
SlotID 5 Port 1	Mellanox ConnectX-6 Dx Adapter #4	Up	9000	100 Gbps
SlotID 5 Port 2	Mellanox ConnectX-6 Dx Adapter #2	Up	9000	100 Gbps
vManagement(mgmt_compute)	Hyper-V Virtual Ethernet Adapter	Up	1500	100 Gbps

```

Host Name: AZSHCI-SL-N4
Verifying NIC status
```

Name	InterfaceDescription	Status	MTUSize	LinkSpeed
SlotID 2 Port 1	Mellanox ConnectX-6 Dx Adapter #2	Up	1500	100 Gbps
SlotID 2 Port 2	Mellanox ConnectX-6 Dx Adapter	Up	1500	100 Gbps
SlotID 5 Port 1	Mellanox ConnectX-6 Dx Adapter #4	Up	9000	100 Gbps
SlotID 5 Port 2	Mellanox ConnectX-6 Dx Adapter #3	Up	9000	100 Gbps
vManagement(mgmt_compute)	Hyper-V Virtual Ethernet Adapter	Up	1500	100 Gbps

Procedure 15. Verify RDMA and RoCEv2 Protocol is Enabled on Physical NICs

Step 1. Run the following:

```
$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verifying RDMA and RoCEv2 status on physical NICS " -ForegroundColor Yellow

Get-NetAdapterAdvancedProperty -InterfaceDescription "Mellanox ConnectX*" -DisplayName "NetworkDirect*" | ft
Name, InterfaceDescription, DisplayName, DisplayValue

}
}
```

```
Host Name: AZSHCI-SL-N3
Verifying RDMA and RoCEv2 status on physical NICS

Name                InterfaceDescription          DisplayName          DisplayValue
-----
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #3 NetworkDirect Functionality Enabled
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #3 NetworkDirect Technology RoCEv2
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #3 NetworkDirect Functionality Enabled
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #3 NetworkDirect Technology RoCEv2
SlotID 5 Port 1 Mellanox ConnectX-6 Dx Adapter #4 NetworkDirect Functionality Enabled
SlotID 5 Port 1 Mellanox ConnectX-6 Dx Adapter #4 NetworkDirect Technology RoCEv2
SlotID 5 Port 2 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Functionality Enabled
SlotID 5 Port 2 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Technology RoCEv2

Host Name: AZSHCI-SL-N4
Verifying RDMA and RoCEv2 status on physical NICS

Name                InterfaceDescription          DisplayName          DisplayValue
-----
SlotID 5 Port 2 Mellanox ConnectX-6 Dx Adapter #3 NetworkDirect Functionality Enabled
SlotID 5 Port 2 Mellanox ConnectX-6 Dx Adapter #3 NetworkDirect Technology RoCEv2
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #3 NetworkDirect Functionality Enabled
SlotID 2 Port 2 Mellanox ConnectX-6 Dx Adapter #3 NetworkDirect Technology RoCEv2
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Functionality Enabled
SlotID 2 Port 1 Mellanox ConnectX-6 Dx Adapter #2 NetworkDirect Technology RoCEv2
SlotID 5 Port 1 Mellanox ConnectX-6 Dx Adapter #4 NetworkDirect Functionality Enabled
SlotID 5 Port 1 Mellanox ConnectX-6 Dx Adapter #4 NetworkDirect Technology RoCEv2
```

Procedure 16. Verify RDMA is Enabled on the Storage vNIC Adapters

Step 1. Run the following:

```
$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verifying that RDMA is enabled on the Storage vNICs" -ForegroundColor Yellow

Get-NetAdapterRdma | ft

}
}
```

```

Host Name: AZSHCI-SL-N3
Verifying that RDMA is enabled on the Storage vNICs

Name                InterfaceDescription      Enabled  Operational  PFC     ETS
-----            -
SlotID 2 Port 2     Mellanox ConnectX-6 Dx Adapter  True    True         True   True
SlotID 2 Port 1     Mellanox ConnectX-6 Dx Adapter #3  True    True         True   True
SlotID 5 Port 1     Mellanox ConnectX-6 Dx Adapter #4  True    True         True   True
SlotID 5 Port 2     Mellanox ConnectX-6 Dx Adapter #2  True    True         True   True
vManagement(mgmt_compute) Hyper-V Virtual Ethernet Adapter  False   False        NA     NA

Host Name: AZSHCI-SL-N4
Verifying that RDMA is enabled on the Storage vNICs

Name                InterfaceDescription      Enabled  Operational  PFC     ETS
-----            -
SlotID 5 Port 2     Mellanox ConnectX-6 Dx Adapter #3  True    True         True   True
SlotID 2 Port 2     Mellanox ConnectX-6 Dx Adapter      True    True         True   True
SlotID 2 Port 1     Mellanox ConnectX-6 Dx Adapter #2  True    True         True   True
vManagement(mgmt_compute) Hyper-V Virtual Ethernet Adapter  False   False        NA     NA
SlotID 5 Port 1     Mellanox ConnectX-6 Dx Adapter #4  True    True         True   True

```

Procedure 17. Verify RDMA Capabilities

Step 1. Run the following:

```

$nodes = ("AZSHCI-SL-N3", "AZSHCI-SL-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verify Storage vNIC RDMA operational status " -ForegroundColor Yellow

Get-SmbClientNetworkInterface | ft FriendlyName, RDMACapable

}
}

```

```

Host Name: AZSHCI-SL-N3
Verify Storage vNIC RDMA operational status

FriendlyName          RDMACapable
-----
SlotID 2 Port 1       False
SlotID 2 Port 2       False
SlotID 5 Port 1       True
SlotID 5 Port 2       True
Local Area Connection* 1 False
vManagement(mgmt_compute) False

Host Name: AZSHCI-SL-N4
Verify Storage vNIC RDMA operational status

FriendlyName          RDMACapable
-----
SlotID 2 Port 1       False
SlotID 2 Port 2       False
vManagement(mgmt_compute) False
SlotID 5 Port 1       True
SlotID 5 Port 2       True
Local Area Connection* 1 False

```

Configure QoS

This section contains the following procedures and applies to both fully converged and non-converged topologies:

- [Verify Traffic Class Configuration on all Nodes](#)

- [Set DCBX Not Willing Mode on all Nodes](#)

Procedure 1. Verify Traffic Class Configuration on all Nodes

Step 1. Run the following:

```
$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Verifying Traffic Class Configuration " -ForegroundColor Yellow
Get-NetQoSTrafficClass | ft -AutoSize

}
}
```

```
Host Name: AZS-HCI1-N1
Verifying Traffic Class Configuration

Name      Algorithm Bandwidth(%) Priority PolicySet IfIndex IfAlias
-----
[Default] ETS       49         0-3,6-7 Global
SMB_Direct ETS       50         4       Global
Cluster   ETS       1          5       Global

Host Name: AZS-HCI1-N2
Verifying Traffic Class Configuration

Name      Algorithm Bandwidth(%) Priority PolicySet IfIndex IfAlias
-----
[Default] ETS       49         0-3,6-7 Global
SMB_Direct ETS       50         4       Global
Cluster   ETS       1          5       Global

Host Name: AZS-HCI1-N3
Verifying Traffic Class Configuration

Name      Algorithm Bandwidth(%) Priority PolicySet IfIndex IfAlias
-----
[Default] ETS       49         0-3,6-7 Global
SMB_Direct ETS       50         4       Global
Cluster   ETS       1          5       Global

Host Name: AZS-HCI1-N4
Verifying Traffic Class Configuration

Name      Algorithm Bandwidth(%) Priority PolicySet IfIndex IfAlias
-----
[Default] ETS       49         0-3,6-7 Global
SMB_Direct ETS       50         4       Global
Cluster   ETS       1          5       Global
```

Procedure 2. Set DCBX Willing Mode on all Nodes

Step 1. Run the following:

Note: Server nodes need to be in Willing mode in order for DCBX auto negotiation to take place and Priority Flow Control to be enabled on the ToR Switch ports where Nodes are connected.

```

$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verifying that DCBX is set to Not Willing mode" -ForegroundColor Yellow
Get-netadapter | Get-NetQosDcbxSetting | ft InterfaceAlias, PolicySet, Willing

}
}

```

```

Host Name: AZS-HCI1-N1
Verifying that DCBX is set to Not Willing mode

InterfaceAlias      PolicySet Willing
-----
SlotID 2 Port 1 AdapterSpecific  False
SlotID 2 Port 2 AdapterSpecific  False

Host Name: AZS-HCI1-N2
Verifying that DCBX is set to Not Willing mode

InterfaceAlias      PolicySet Willing
-----
SlotID 2 Port 1 AdapterSpecific  False
SlotID 2 Port 2 AdapterSpecific  False

Host Name: AZS-HCI1-N3
Verifying that DCBX is set to Not Willing mode

InterfaceAlias      PolicySet Willing
-----
SlotID 2 Port 2 AdapterSpecific  False
SlotID 2 Port 1 AdapterSpecific  False

Host Name: AZS-HCI1-N4
Verifying that DCBX is set to Not Willing mode

InterfaceAlias      PolicySet Willing
-----
SlotID 2 Port 1 AdapterSpecific  False
SlotID 2 Port 2 AdapterSpecific  False

```

Prepare Server for Storage Spaces Direct

This section contains the following procedures and applies to both fully converged and non-converged topologies.

- [Run Windows Updates](#)
- [Clean Inventory Storage Drives that will be used by Storage Spaces Direct](#)

- [Verify the Servers are ready for Storage Spaces Direct](#)

Procedure 1. Run Windows Updates

IMPORTANT! It is extremely important to install the latest updated for Failover Cluster, Scale-Out Files Server, and Storage Spaces. Run Windows Update to install the latest updates after installing the Windows Features.

Note: The Cluster-Aware Updating role will be installed after the cluster is created. The cluster-aware updating is a feature that automates downloading and installing Windows Server updates on all cluster nodes.

Procedure 2. Clean Inventory Storage Drives that will be used by Storage Spaces Direct

Step 1. Run the following:

```
$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Preparing disk for Storage Spaces Direct" -ForegroundColor Yellow
Write-Host Cleaning Storage Drives....
#Remove Existing virtual disks and storage pools
Update-StorageProviderCache
    Get-StoragePool | ? IsPrimordial -eq $false | Set-StoragePool -IsReadOnly:$false -ErrorAction SilentlyContinue
    Get-StoragePool | ? IsPrimordial -eq $false | Get-VirtualDisk | Remove-VirtualDisk -Confirm:$false -ErrorAction SilentlyContinue
    Get-StoragePool | ? IsPrimordial -eq $false | Remove-StoragePool -Confirm:$false -ErrorAction SilentlyContinue
    Get-PhysicalDisk | Reset-PhysicalDisk -ErrorAction SilentlyContinue
    Get-Disk | ? Number -ne $null | ? IsBoot -ne $true | ? IsSystem -ne $true | ? PartitionStyle -ne RAW | %
    {
        $_ | Set-Disk -isoffline:$false
        $_ | Set-Disk -isreadonly:$false
        $_ | Clear-Disk -RemoveData -RemoveOEM -Confirm:$false
        $_ | Set-Disk -isreadonly:$true
        $_ | Set-Disk -isoffline:$true
    }

#Inventory Storage Disks
Get-Disk | Where Number -Ne $Null | Where IsBoot -Ne $True | Where IsSystem -Ne $True | Where PartitionStyle -Eq RAW | Group -NoElement -Property FriendlyName | ft
}
```



```
}
```

```
Host Name: AZS-HCI1-N1  
Preparing disk for Storage Spaces Direct  
Cleaning Storage Drives....
```

```
Count Name  
-----  
12 INTEL SSDPF2KX038T9K
```

```
Host Name: AZS-HCI1-N2  
Preparing disk for Storage Spaces Direct  
Cleaning Storage Drives....
```

```
Count Name  
-----  
12 INTEL SSDPF2KX038T9K
```

```
Host Name: AZS-HCI1-N3  
Preparing disk for Storage Spaces Direct  
Cleaning Storage Drives....
```

```
Count Name  
-----  
12 INTEL SSDPF2KX038T9K
```

```
Host Name: AZS-HCI1-N4  
Preparing disk for Storage Spaces Direct  
Cleaning Storage Drives....
```

```
Count Name  
-----  
12 INTEL SSDPF2KX038T9K
```

Procedure 3. Verify the Servers are ready for Storage Spaces Direct

Step 1. Run the following:

```
$CandidateClusterNode = "AzS-HCI1-N1"  
Invoke-Command $CandidateClusterNode -Credential $Creds -scriptblock {  
  
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green  
Write-Host " Enabling CredSSP" -ForegroundColor Yellow  
$Void = Enable-WSManCredSSP -Role Server -Force  
}  
  
Invoke-Command $CandidateClusterNode -Credential $Creds -authentication Credssp -scriptblock {  
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green  
  
$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")
```

```

Write-Host " Validating Cluster Nodes..." -ForegroundColor Yellow
Test-Cluster -Node $nodes -Include "System Configuration",Networking,Inventory, "Storage Spaces Direct"

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WSManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
Get-WSManCredSSP
}

```

```

Host Name: AZS-HCI-HOST01
Enabling CredSSP
Host Name: AZS-HCI-HOST01
Validating Cluster Nodes...

Mode                LastWriteTime         Length Name                                           PSComputerName
----                -
-a-----          3/7/2022  11:54 AM         1691940 Validation Report 2022.03.07 At 11.49.27.htm      AzS-HCI-Host01
Disabling CredSSP
Verifying that CredSSP are disabled on target server...
The machine is configured to allow delegating fresh credentials to the following target(s): wsman/*
This computer is not configured to receive credentials from a remote client computer.

```

```

Host Name: AZS-HCI1-N1
Enabling CredSSP
Host Name: AZS-HCI1-N1
Validating Cluster Nodes...
WARNING: System Configuration - Validate Software Update Levels: The test reported some warnings..
WARNING:
Test Result:
HadUnselectedTests, ClusterConditionallyApproved
Testing has completed for the tests you selected. You should review the warnings in the Report. A cluster solution is supported by Microsoft only if you run all cluster
validation tests, and all tests succeed (with or without warnings).
Test report file path: C:\Users\hciadmin\AppData\Local\Temp\Validation Report 2023.03.08 At 10.41.40.htm

Mode                LastWriteTime         Length Name                                           PSComputerName
----                -
-a-----          3/8/2023  10:45 AM         1611336 Validation Report 2023.03.08 At 10.41.40.htm      AzS-HCI1-N1
Disabling CredSSP
Verifying that CredSSP are disabled on target server...
The machine is not configured to allow delegating fresh credentials.
This computer is not configured to receive credentials from a remote client computer.

```

Step 2. Review the validation report and resolve all errors and warning before proceeding to create the Storage Spaces Direct Cluster:






Failover Cluster Validation Report

Node:	Azs-HCI1-N1.ucs-spaces.lab	Validated
Node:	AzS-HCI1-N2.ucs-spaces.lab	Validated
Node:	AzS-HCI1-N3.ucs-spaces.lab	Validated
Node:	AzS-HCI1-N4.ucs-spaces.lab	Validated
Started	3/8/2023 10:41:40 AM	
Completed	3/8/2023 10:45:22 AM	

The Validate a Configuration Wizard must be run after any change is made to the configuration of the cluster or hardware. For more information, see <https://go.microsoft.com/fwlink/p/?LinkId=280145>.

Results by Category

Name	Result Summary	Description
Inventory		Success
Network		Success
Storage Spaces Direct		Success
System Configuration		Success

Configure Windows Failover Cluster

Note: In this section, choose the appropriate deployment model (fully converged or non-converged network topology configuration) under the “Create and Deploy Clustered Network ATC Intent” section and skip the others as per your requirement.

This section contains the following procedures:

- [Create the Cluster](#)
- [Verify Status for Cluster Nodes after creating the Cluster](#)
- [Remove Standalone Network ATC Intent](#)
- [Create and Deploy Clustered Network ATC Intent- Fully Converged Network Topology](#)
- [Create and Deploy Clustered Network ATC Intent - Non-Converged Network Topology](#)
- [Verify Cluster Network Interfaces](#)
- [Verify Live Migration Exclusion List](#)
- [Configure Live Migration to use SMB Protocol](#)
- [Configure Live Migration Bandwidth Limit](#)
- [Create Maximum Bandwidth Limit for Management vNIC](#)
- [Create the File Share for the Cluster Witness](#)
- [Configure File Share Witness](#)
- [Additional Cluster Quorum Witness Options](#)
- [Configure Cluster-Aware Updating](#)
- [Configure Kernel Soft Reboot for Cluster Aware Updating](#)

Procedure 1. Create the Cluster

Step 1. Create the cluster with a static IP Address:

```
$CandidateClusterNode = "AzS-HCI1-N1"
Invoke-Command $CandidateClusterNode -Credential $Creds -scriptblock {
```

```

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WSManCredSSP -Role Server -Force
}

Invoke-Command $CandidateClusterNode -Credential $Creds -authentication Credssp -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

$nodes = ("AzS-HCI1-N1", "AzS-HCI1-N2", "AzS-HCI1-N3", "AzS-HCI1-N4")

Write-Host " Creating the cluster..." -ForegroundColor Yellow
$Cluster = "AzS-HCI-M6-C1"
New-Cluster -Name $Cluster -Node $nodes -StaticAddress 192.168.126.25 -NoStorage
Get-Cluster | fl Name, SharedVolumesRoot

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WSManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
Get-WSManCredSSP
}

```

```

Host Name: AZS-HCI1-N1
Enabling CredSSP
Host Name: AZS-HCI1-N1
Creating the cluster...

Name                PSComputerName
----                -
AzS-HCI-M6-C1      AzS-HCI1-N1

Name                : AzS-HCI-M6-C1
SharedVolumesRoot  : C:\ClusterStorage

Disabling CredSSP
Verifying that CredSSP are disabled on target server...
The machine is not configured to allow delegating fresh credentials.
This computer is not configured to receive credentials from a remote client computer.

```

Procedure 2. Verify Status for Cluster Nodes after creating the Cluster

Step 1. Run the following:

```

$Cluster = "AzS-HCI-M6-C1"

Invoke-Command $Cluster -Credential $Creds -scriptblock {

```

```

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WSManCredSSP -Role Server -Force
}

Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {

$Cluster = "AzS-HCI-M6-C1"

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Checking cluster nodes..." -ForegroundColor Yellow
Get-ClusterNode -Cluster $Cluster | ft Name, State, Type

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WSManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
Get-WSManCredSSP
}

```

```

Host Name: AZS-HCI1-N4
Enabling CredSSP
Host Name: AZS-HCI1-N4
Checking cluster nodes...

Name          State Type
----          -
Azs-HCI1-N1   Up   Node
Azs-HCI1-N2   Up   Node
Azs-HCI1-N3   Up   Node
Azs-HCI1-N4   Up   Node

Disabling CredSSP
Verifying that CredSSP are disabled on target server...
The machine is not configured to allow delegating fresh credentials.
This computer is not configured to receive credentials from a remote client computer.

```

Procedure 3. Remove Standalone Network ATC Intent

Step 1. Run the following:

```

$Cluster = "AzS-HCI-M6-C1"
$nodes = (Get-ClusterNode -Cluster $Cluster).Name
foreach ($node in $nodes) {

Invoke-Command $node -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

```

```

Write-Host " Identifying and Removing Standalone Network ATC Intent " -ForegroundColor Yellow
$intent = Get-NetIntent | Where-Object {$_.Scope -Like 'Host' -and $_.IntentName -EQ 'mgmt_compute_storage'}

Write-Host "Removing Standalone Network ATC Intent $intent" -ForegroundColor Yellow
Remove-NetIntent -Name $intent.IntentName

}
}

```

Procedure 4. Create and Deploy Clustered Network ATC Intent – Fully Converged Network Topology

Note: Ignore the steps in this section if deploying non-converged network topology configuration in your environment.

Step 1. Run the following to create and deploy clustered Network ATC intent for fully converged topology configuration:

```

$Cluster = "AzS-HCI-M6-C1"

Invoke-Command $Cluster -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WSManCredSSP -Role Server -Force
}

Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Create and Deploy Clustered Network ATC Intent " -ForegroundColor Yellow

$ClusterName = Get-cluster
$QoSOverride = New-NetIntentQoSPolicyOverRides
$AdapterOverride = New-NetIntentAdapterPropertyOverrides
$storageOverride = new-NetIntentStorageOverrides

$QoSOverride.PriorityValue8021Action_SMB = 4
$QoSOverride.PriorityValue8021Action_Cluster = 5
$AdapterOverride.NetworkDirectTechnology = 4
$storageOverride.EnableAutomaticIPGeneration = $false

$QoSOverride
$AdapterOverride
$storageOverride

```

```
Add-NetIntent -AdapterName "SlotID 2 Port 1", "SlotID 2 Port 2" -Management -Compute -Storage -StorageVlans 107, 207 -QoSPolicyOverrides $QoSOverride -AdapterPropertyOverrides $AdapterOverride -StorageOverrides $StorageOverride -Name Mgmt_Compute_Storage
```

```
Write-Host " Disabling CredSSP" -ForegroundColor Yellow
```

```
Disable-WSManCredSSP -Role Server
```

```
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
```

```
Get-WSManCredSSP
```

```
}
```

```
-- Creating a new intent with name Mgmt_Compute_Storage
-- Compute intent was submitted
-- Management intent was submitted
-- Storage intent was submitted
-- Override found for Adapter Properties
-- Override found for QoS Policy
-- Override found for Storage Parameters
-- The specified Storage Vlan for SlotID 2 Port 1 was: 107
-- The specified Storage Vlan for SlotID 2 Port 2 was: 207
-- Checking if exact intent request 'mgmt_compute_storage' already exists
-- Checking if specified physical adapters conflict with an existing intent
-- Validating if physical NICs with the name exist on the remote server(s) and are status 'Up'
-- Validating network adapters and virtual switch on all the following nodes
Azs-HCI1-N1 Azs-HCI1-N2 Azs-HCI1-N3 Azs-HCI1-N4
-- Found SlotID 2 Port 1 on Azs-HCI1-N1
-- Found SlotID 2 Port 2 on Azs-HCI1-N1
-- Validating physical NICs on Azs-HCI1-N1 are symmetric
-- Found SlotID 2 Port 1 on Azs-HCI1-N2
-- Found SlotID 2 Port 2 on Azs-HCI1-N2
-- Validating physical NICs on Azs-HCI1-N2 are symmetric
-- Found SlotID 2 Port 1 on Azs-HCI1-N3
-- Found SlotID 2 Port 2 on Azs-HCI1-N3
-- Validating physical NICs on Azs-HCI1-N3 are symmetric
-- Found SlotID 2 Port 1 on Azs-HCI1-N4
-- Found SlotID 2 Port 2 on Azs-HCI1-N4
-- Validating physical NICs on Azs-HCI1-N4 are symmetric
-- Submitting Intent request for mgmt_compute_storage
-- SUCCESS: Intent request for mgmt_compute_storage submitted
-- Checking for existing global intent
-- No existing global intent. Putting global intent with default parameters
```

Please check Get-NetIntentStatus to see provisioning status. Deployment can take several minutes to complete.

Disabling CredSSP

Verifying that CredSSP are disabled on target server...

The machine is not configured to allow delegating fresh credentials.

This computer is not configured to receive credentials from a remote client computer.

Step 2. Run the following to verify clustered Network ATC deployment and status:

```
$Cluster = "Azs-HCI-M6-C1"
```

```
Invoke-Command $Cluster -Credential $Creds -scriptblock {
```

```
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
```

```
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
```

```
$Void = Enable-WSManCredSSP -Role Server -Force
```

```
}
```

```
Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verify Clustered Network ATC Intent Status" -ForegroundColor Yellow

$ClusterName = (Get-cluster).Name

Get-NetIntent -ClusterName $ClusterName | Select IntentName,scope
Get-NetIntentStatus -ClusterName $ClusterName | Select Host, IntentName, ConfigurationStatus,
ProvisioningStatus

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WSManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
Get-WSManCredSSP

}
```



```

Host Name: AZS-HCI1-N4
Enabling CredSSP
Host Name: AZS-HCI1-N4
Verify Clustered Network ATC Intent Status

IntentName      : mgmt_compute_storage
Scope           : Cluster
PSComputerName  : AzS-HCI-M6-C1
RunspaceId      : aaad49c7-9cf0-4f38-ac37-68f142a68420

Host            : azs-hci1-n1
IntentName      : mgmt_compute_storage
ConfigurationStatus : Success
ProvisioningStatus : Completed
PSComputerName  : AzS-HCI-M6-C1
RunspaceId      : aaad49c7-9cf0-4f38-ac37-68f142a68420

Host            : azs-hci1-n2
IntentName      : mgmt_compute_storage
ConfigurationStatus : Success
ProvisioningStatus : Completed
PSComputerName  : AzS-HCI-M6-C1
RunspaceId      : aaad49c7-9cf0-4f38-ac37-68f142a68420

Host            : azs-hci1-n3
IntentName      : mgmt_compute_storage
ConfigurationStatus : Success
ProvisioningStatus : Completed
PSComputerName  : AzS-HCI-M6-C1
RunspaceId      : aaad49c7-9cf0-4f38-ac37-68f142a68420

Host            : azs-hci1-n4
IntentName      : mgmt_compute_storage
ConfigurationStatus : Success
ProvisioningStatus : Completed
PSComputerName  : AzS-HCI-M6-C1
RunspaceId      : aaad49c7-9cf0-4f38-ac37-68f142a68420

Disabling CredSSP
Verifying that CredSSP are disabled on target server...
The machine is not configured to allow delegating fresh credentials.
This computer is not configured to receive credentials from a remote client computer.

```

Note: It may take a few minutes for the network intent application to complete.

Step 3. Run the following verify network adapter status after network intent has been applied:

```

$nodes = (Get-ClusterNode -Cluster $Cluster).Name
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WSManCredSSP -Role Server -Force

```

```
Write-Host "Verifying NIC Port Status " -ForegroundColor Yellow
```

```
Get-netadapter | ft Name, InterfaceDescription, Status, MTUSize, MacAddress, LinkSpeed
```

```
Write-Host " Disabling CredSSP" -ForegroundColor Yellow
```

```
Disable-WSManCredSSP -Role Server
```

```
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
```

```
Get-WSManCredSSP
```

```
}
```

```
}
```

```
Host Name: AZS-HCI1-N1
Enabling CredSSP
Verifying NIC Port Status

Name                               InterfaceDescription              Status MTUSize MacAddress      LinkSpeed
----                               -
vManagement(mgmt_compute_storage)  Hyper-V Virtual Ethernet Adapter Up     1500 08-C0-EB-7E-D0-C5 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 2) Hyper-V Virtual Ethernet Adapter #3 Up     1500 00-15-5D-7E-15-01 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 1) Hyper-V Virtual Ethernet Adapter #2 Up     1500 00-15-5D-7E-15-00 100 Gbps
SlotID 2 Port 2                    Mellanox ConnectX-6 Dx Adapter #2 Up     1660 08-C0-EB-7E-D0-C4 100 Gbps
SlotID 2 Port 1                    Mellanox ConnectX-6 Dx Adapter    Up     1660 08-C0-EB-7E-D0-C5 100 Gbps

Disabling CredSSP
Verifying that CredSSP are disabled on target server...
The machine is not configured to allow delegating fresh credentials.
This computer is not configured to receive credentials from a remote client computer.
Host Name: AZS-HCI1-N2
Enabling CredSSP
Verifying NIC Port Status

Name                               InterfaceDescription              Status MTUSize MacAddress      LinkSpeed
----                               -
SlotID 2 Port 1                    Mellanox ConnectX-6 Dx Adapter #2 Up     1660 08-C0-EB-7E-D2-ED 100 Gbps
SlotID 2 Port 2                    Mellanox ConnectX-6 Dx Adapter    Up     1660 08-C0-EB-7E-D2-EC 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 1) Hyper-V Virtual Ethernet Adapter #2 Up     1500 00-15-5D-7E-16-00 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 2) Hyper-V Virtual Ethernet Adapter #3 Up     1500 00-15-5D-7E-16-01 100 Gbps
vManagement(mgmt_compute_storage)  Hyper-V Virtual Ethernet Adapter Up     1500 08-C0-EB-7E-D2-ED 100 Gbps

Disabling CredSSP
Verifying that CredSSP are disabled on target server...
The machine is not configured to allow delegating fresh credentials.
This computer is not configured to receive credentials from a remote client computer.
Host Name: AZS-HCI1-N3
Enabling CredSSP
Verifying NIC Port Status

Name                               InterfaceDescription              Status MTUSize MacAddress      LinkSpeed
----                               -
SlotID 2 Port 1                    Mellanox ConnectX-6 Dx Adapter #2 Up     1660 08-C0-EB-7E-D0-BC 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 1) Hyper-V Virtual Ethernet Adapter #2 Up     1500 00-15-5D-7E-17-00 100 Gbps
SlotID 2 Port 2                    Mellanox ConnectX-6 Dx Adapter    Up     1660 08-C0-EB-7E-D0-BD 100 Gbps
vManagement(mgmt_compute_storage)  Hyper-V Virtual Ethernet Adapter Up     1500 08-C0-EB-7E-D0-BC 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 2) Hyper-V Virtual Ethernet Adapter #3 Up     1500 00-15-5D-7E-17-01 100 Gbps

Disabling CredSSP
Verifying that CredSSP are disabled on target server...
The machine is not configured to allow delegating fresh credentials.
This computer is not configured to receive credentials from a remote client computer.
Host Name: AZS-HCI1-N4
Enabling CredSSP
Verifying NIC Port Status

Name                               InterfaceDescription              Status MTUSize MacAddress      LinkSpeed
----                               -
SlotID 2 Port 1                    Mellanox ConnectX-6 Dx Adapter    Up     1660 08-C0-EB-7E-D3-2C 100 Gbps
SlotID 2 Port 2                    Mellanox ConnectX-6 Dx Adapter #2 Up     1660 08-C0-EB-7E-D3-2D 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 1) Hyper-V Virtual Ethernet Adapter #2 Up     1500 00-15-5D-7E-18-00 100 Gbps
vManagement(mgmt_compute_storage)  Hyper-V Virtual Ethernet Adapter Up     1500 08-C0-EB-7E-D3-2C 100 Gbps
vSMB(mgmt_compute_storage#SlotID 2 Port 2) Hyper-V Virtual Ethernet Adapter #3 Up     1500 00-15-5D-7E-18-01 100 Gbps
```

Step 4. Check the cluster networks:

```
$Cluster = "AzS-HCI-M6-C1"
```

```

Invoke-Command $Cluster -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WSManCredSSP -Role Server -Force
}

Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Checking cluster networks " -ForegroundColor Yellow

$ClusterName = (Get-cluster).Name

Get-ClusterNetwork -Cluster $ClusterName | ft name,address,state,role -autosize

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WSManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
Get-WSManCredSSP
}

```

```

Host Name: AZS-HCI1-N4
Enabling CredSSP
Host Name: AZS-HCI1-N4
Checking cluster networks

Name                                Address                State  Role
----                                -
mgmt_compute_storage(Management)    192.168.126.0          Up     ClusterAndClient
mgmt_compute_storage(Storage_VLAN107) 192.168.107.0         Up     Cluster
mgmt_compute_storage(Storage_VLAN207) 192.168.207.0         Up     Cluster

Disabling CredSSP
Verifying that CredSSP are disabled on target server...
The machine is not configured to allow delegating fresh credentials.
This computer is not configured to receive credentials from a remote client computer.

```

Step 5. Run the following to verify cluster network interfaces:

```

$Cluster = "AzS-HCI-M6-C1"

Invoke-Command $Cluster -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WSManCredSSP -Role Server -Force
}

```

```

Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Verifying cluster network interfaces " -ForegroundColor Yellow

$ClusterName = (Get-cluster).Name

Get-ClusterNetworkInterface -Cluster $ClusterName | sort Name | ft Network, Name

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WSManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
Get-WSManCredSSP
}

```

```

Host Name: AZS-HCI1-N4
Enabling CredSSP
Host Name: AZS-HCI1-N4
Verifying cluster network interfaces

Network                                     Name
-----
mgmt_compute_storage(Management)          AzS-HCI1-N1 - vManagement(mgmt_compute_storage)
mgmt_compute_storage(Storage_VLAN107)     AzS-HCI1-N1 - vSMB(mgmt_compute_storage#SlotID 2 Port 1)
mgmt_compute_storage(Storage_VLAN207)     AzS-HCI1-N1 - vSMB(mgmt_compute_storage#SlotID 2 Port 2)
mgmt_compute_storage(Management)          AzS-HCI1-N2 - vManagement(mgmt_compute_storage)
mgmt_compute_storage(Storage_VLAN107)     AzS-HCI1-N2 - vSMB(mgmt_compute_storage#SlotID 2 Port 1)
mgmt_compute_storage(Storage_VLAN207)     AzS-HCI1-N2 - vSMB(mgmt_compute_storage#SlotID 2 Port 2)
mgmt_compute_storage(Management)          AzS-HCI1-N3 - vManagement(mgmt_compute_storage)
mgmt_compute_storage(Storage_VLAN107)     AzS-HCI1-N3 - vSMB(mgmt_compute_storage#SlotID 2 Port 1)
mgmt_compute_storage(Storage_VLAN207)     AzS-HCI1-N3 - vSMB(mgmt_compute_storage#SlotID 2 Port 2)
mgmt_compute_storage(Management)          AzS-HCI1-N4 - vManagement(mgmt_compute_storage)
mgmt_compute_storage(Storage_VLAN107)     AzS-HCI1-N4 - vSMB(mgmt_compute_storage#SlotID 2 Port 1)
mgmt_compute_storage(Storage_VLAN207)     AzS-HCI1-N4 - vSMB(mgmt_compute_storage#SlotID 2 Port 2)

Disabling CredSSP
Verifying that CredSSP are disabled on target server...
The machine is not configured to allow delegating fresh credentials.
This computer is not configured to receive credentials from a remote client computer.

```

Procedure 5. Create and Deploy Clustered Network ATC Intent - Non-Converged Network Topology

Note: Ignore the steps in this section if deploying fully converged network topology configuration in your environment.

Step 1. Run the following to create and deploy clustered Network ATC intent for non-converged network topology configuration:

```

$Cluster = "AZSHCI1-M6-C"

Invoke-Command $Cluster -Credential $Creds -scriptblock {

```

```

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WSManCredSSP -Role Server -Force
}

Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Create and Deploy Clustered Network ATC Intent for non-Converged " -ForegroundColor Yellow

$ClusterName = Get-cluster
$Mgmt_Compute_NICs = @("SlotID 2 Port 1", "SlotID 5 Port 2")

$Storage_NICs = @("SlotID 2 Port 2", "SlotID 5 Port 1")
$Storage_VLANS = @(207,107)

$QoSOverride = New-NetIntentQoSPolicyOverRides
$QoSOverride.PriorityValue8021Action_SMB = 4
$QoSOverride.PriorityValue8021Action_Cluster = 5

$AdapterOverride = New-NetIntentAdapterPropertyOverrides
$AdapterOverride.NetworkDirectTechnology = 4
$AdapterOverride.JumboPacket = 9014

$StorageOverride = new-NetIntentStorageOverrides
$StorageOverride.EnableAutomaticIPGeneration = $false

Write-Host " Creating Mgmt and Compute intent with "SlotID 2 Port 1" & "SlotID 5 Port 2" " -ForegroundColor
Yellow
Add-NetIntent -Name Mgmt_Compute -AdapterName $Mgmt_Compute_NICs -Management -Compute -verbose

Write-Host " Creating Storage intent with "SlotID 2 Port 2" & "SlotID 5 Port 1" " -ForegroundColor Yellow
Add-NetIntent -Name Storage -AdapterName $Storage_NICs -Storage -StorageVlans $Storage_VLANS -
QoSPolicyOverrides $QoSOverride -AdapterPropertyOverrides $AdapterOverride -StorageOverrides $StorageOverride

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WSManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
Get-WSManCredSSP
}

```

```

Enabling CredSSP
Host Name: AZSHCI1-M6-N3
Create and Deploy Clustered Network ATC Intent for non-Converged
Creating Mgmt and Compute intent with SlotID 2 Port 1 & SlotID 5 Port 2
VERBOSE: Connecting to cluster on local computer AZSHCI1-M6-N3.
VERBOSE: Connecting to cluster on local computer AZSHCI1-M6-N3.
-- Creating a new intent with name Mgmt_Compute
-- Compute intent was submitted
-- Management intent was submitted
-- The specified Storage Vlan for SlotID 2 Port 1 was: 711
-- The specified Storage Vlan for SlotID 5 Port 2 was: 712
VERBOSE: Connecting to cluster on local computer AZSHCI1-M6-N3.
VERBOSE: Connecting to cluster on local computer AZSHCI1-M6-N3.
-- Checking if exact intent request 'mgmt_compute' already exists
-- Checking if specified physical adapters conflict with an existing intent
-- Validating if physical NICs with the name exist on the remote server(s) and are status 'Up'
VERBOSE: Connecting to cluster on local computer AZSHCI1-M6-N3.
-- Validating network adapters and virtual switch on all the following nodes
AZSHCI1-M6-N1 AZSHCI1-M6-N2 AZSHCI1-M6-N3
-- Found SlotID 2 Port 1 on AZSHCI1-M6-N1
-- Found SlotID 5 Port 2 on AZSHCI1-M6-N1
-- Validating physical NICs on AZSHCI1-M6-N1 are symmetric
-- Found SlotID 2 Port 1 on AZSHCI1-M6-N2
-- Found SlotID 5 Port 2 on AZSHCI1-M6-N2
-- Validating physical NICs on AZSHCI1-M6-N2 are symmetric
-- Found SlotID 2 Port 1 on AZSHCI1-M6-N3
-- Found SlotID 5 Port 2 on AZSHCI1-M6-N3
-- Validating physical NICs on AZSHCI1-M6-N3 are symmetric
-- Submitting Intent request for mgmt_compute
-- SUCCESS: Intent request for mgmt_compute submitted
-- Checking for existing global intent
VERBOSE: Connecting to cluster on local computer AZSHCI1-M6-N3.
VERBOSE: Connecting to cluster on local computer AZSHCI1-M6-N3.
-- No existing global intent. Putting global intent with default parameters

Please check Get-NetIntentStatus to see provisioning status. Deployment can take several minutes to complete.
Creating Storage intent with SlotID 2 Port 2 & SlotID 5 Port 1
-- Creating a new intent with name Storage
-- Storage intent was submitted
-- Override found for Adapter Properties
-- Override found for QoS Policy
-- Override found for Storage Parameters
-- The specified Storage Vlan for SlotID 2 Port 2 was: 207
-- The specified Storage Vlan for SlotID 5 Port 1 was: 107
-- Checking if exact intent request 'storage' already exists
-- Checking if specified physical adapters conflict with an existing intent
-- Validating if physical NICs with the name exist on the remote server(s) and are status 'Up'
-- Validating network adapters and virtual switch on all the following nodes
AZSHCI1-M6-N1 AZSHCI1-M6-N2 AZSHCI1-M6-N3
-- Found SlotID 2 Port 2 on AZSHCI1-M6-N1
-- Found SlotID 5 Port 1 on AZSHCI1-M6-N1
-- Validating physical NICs on AZSHCI1-M6-N1 are symmetric
-- Found SlotID 2 Port 2 on AZSHCI1-M6-N2
-- Found SlotID 5 Port 1 on AZSHCI1-M6-N2
-- Validating physical NICs on AZSHCI1-M6-N2 are symmetric
-- Found SlotID 2 Port 2 on AZSHCI1-M6-N3
-- Found SlotID 5 Port 1 on AZSHCI1-M6-N3
-- Validating physical NICs on AZSHCI1-M6-N3 are symmetric
-- Submitting Intent request for storage
-- SUCCESS: Intent request for storage submitted
-- Checking for existing global intent

```

Step 2. Run the following to verify clustered Network ATC deployment and status:

```

$Cluster = "AZSHCI1-M6-C"

Invoke-Command $Cluster -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WSManCredSSP -Role Server -Force
}

Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host "Verify Clustered Network ATC Intent Status for non-converged" -ForegroundColor Yellow

$ClusterName = (Get-cluster).Name

```

```

Get-NetIntent -ClusterName $ClusterName | Select IntentName,scope
Get-NetIntentStatus -ClusterName $ClusterName | Select Host, IntentName, ConfigurationStatus,
ProvisioningStatus,IsComputeIntentSet,IsManagementIntentSet,IsStorageIntentset | ft -AutoSize

Write-Host " Disabling CredSSP" -ForegroundColor Yellow

Disable-WManCredSSP -Role Server

Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow

Get-WManCredSSP

}

```

```

Host Name: AZSHCI1-M6-N3
Enabling CredSSP
Host Name: AZSHCI1-M6-N3
Verify Clustered Network ATC Intent Status

IntentName : mgmt_compute
Scope      : Cluster
PSCoMputerName : AZSHCI1-M6-C
RunspaceId : 95e22f3e-7d71-49d2-ae9f-0955ad9af119

IntentName : storage
Scope      : Cluster
PSCoMputerName : AZSHCI1-M6-C
RunspaceId : 95e22f3e-7d71-49d2-ae9f-0955ad9af119

Host      IntentName ConfigurationStatus ProvisioningStatus IsComputeIntentSet IsManagementIntentSet IsStorageIntentSet
-----
azshci1-m6-n1 mgmt_compute Success Completed True True False
azshci1-m6-n2 mgmt_compute Success Completed True True False
azshci1-m6-n3 mgmt_compute Success Completed True True False
azshci1-m6-n1 storage Success Completed False False True
azshci1-m6-n2 storage Success Completed False False True
azshci1-m6-n3 storage Success Completed False False True

```

Note: It may take a few minutes for the network intent application to complete.

Step 3. Run the following verify network adapter status after network intent has been applied:

```

$nodes = (Get-ClusterNode -Cluster $Cluster).Name
foreach ($node in $nodes) {

Invoke-Command $node -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WManCredSSP -Role Server -Force

Write-Host "Verifying NIC Port Status " -ForegroundColor Yellow

Get-netadapter | ft Name, InterfaceDescription, Status, MTUSize, MacAddress, LinkSpeed

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
Get-WManCredSSP

}
}

```

```
}
```

```
Host Name: AZSHCI1-M6-N1
Enabling CredSSP
Verifying NIC Port Status

Name                InterfaceDescription      Status MTUSize MacAddress      LinkSpeed
-----
SlotID 2 Port 2     Mellanox ConnectX-6 Dx Adapter #2 Up      9000 08-C0-EB-7E-D0-C4 100 Gbps
SlotID 2 Port 1     Mellanox ConnectX-6 Dx Adapter Up       1500 08-C0-EB-7E-D0-C5 100 Gbps
vManagement(mgmt_compute) Hyper-V Virtual Ethernet Adapter Up      1500 08-C0-EB-7E-D0-C5 100 Gbps
SlotID 5 Port 1     Mellanox ConnectX-6 Dx Adapter #4 Up      9000 94-6D-AE-75-F1-23 100 Gbps
SlotID 5 Port 2     Mellanox ConnectX-6 Dx Adapter #3 Up      1500 94-6D-AE-75-F1-22 100 Gbps

Disabling CredSSP
Verifying that CredSSP are disabled on target server...
The machine is not configured to allow delegating fresh credentials.
This computer is not configured to receive credentials from a remote client computer.
Host Name: AZSHCI1-M6-N2
Enabling CredSSP
Verifying NIC Port Status

Name                InterfaceDescription      Status MTUSize MacAddress      LinkSpeed
-----
SlotID 2 Port 2     Mellanox ConnectX-6 Dx Adapter #2 Up      9000 08-C0-EB-7E-D2-EC 100 Gbps
SlotID 5 Port 2     Mellanox ConnectX-6 Dx Adapter #3 Up      1500 94-6D-AE-75-E6-1A 100 Gbps
SlotID 5 Port 1     Mellanox ConnectX-6 Dx Adapter #4 Up      9000 94-6D-AE-75-E6-1B 100 Gbps
vManagement(mgmt_compute) Hyper-V Virtual Ethernet Adapter Up      1500 08-C0-EB-7E-D2-ED 100 Gbps
SlotID 2 Port 1     Mellanox ConnectX-6 Dx Adapter Up       1500 08-C0-EB-7E-D2-ED 100 Gbps

Disabling CredSSP
Verifying that CredSSP are disabled on target server...
The machine is not configured to allow delegating fresh credentials.
This computer is not configured to receive credentials from a remote client computer.
Host Name: AZSHCI1-M6-N3
Enabling CredSSP
Verifying NIC Port Status

Name                InterfaceDescription      Status MTUSize MacAddress      LinkSpeed
-----
SlotID 2 Port 2     Mellanox ConnectX-6 Dx Adapter Up      9000 08-C0-EB-7E-D0-BD 100 Gbps
SlotID 5 Port 1     Mellanox ConnectX-6 Dx Adapter #4 Up      9000 94-6D-AE-75-E8-FB 100 Gbps
SlotID 5 Port 2     Mellanox ConnectX-6 Dx Adapter #3 Up      1500 94-6D-AE-75-E8-FA 100 Gbps
vManagement(mgmt_compute) Hyper-V Virtual Ethernet Adapter Up      1500 08-C0-EB-7E-D0-BC 100 Gbps
SlotID 2 Port 1     Mellanox ConnectX-6 Dx Adapter #2 Up      1500 08-C0-EB-7E-D0-BC 100 Gbps
```

Step 4. Check cluster networks:

```
$Cluster = "AZSHCI1-M6-C"

Invoke-Command $Cluster -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WSManCredSSP -Role Server -Force
}

Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Checking cluster networks " -ForegroundColor Yellow

$ClusterName = (Get-cluster).Name

Get-ClusterNetwork -Cluster $ClusterName | ft name,address,state,role -autosize

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WSManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
```



```
Get-WSManCredSSP
}
```

```
Host Name: AZSHCI1-M6-N3
Enabling CredSSP
Host Name: AZSHCI1-M6-N3
Checking cluster networks
```

Name	Address	State	Role
mgmt_compute(Management)	192.168.126.0	Up	ClusterAndClient
storage(Storage_VLAN107)	192.168.107.0	Up	Cluster
storage(Storage_VLAN207)	192.168.207.0	Up	Cluster

Step 5. Run the following to verify cluster network interfaces:

```
$Cluster = "AZSHCI1-M6-C"

Invoke-Command $Cluster -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WSManCredSSP -Role Server -Force
}

Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Verifying cluster network interfaces " -ForegroundColor Yellow

$ClusterName = (Get-cluster).Name

Get-ClusterNetworkInterface -Cluster $ClusterName | sort Name | ft Network, Name

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WSManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
Get-WSManCredSSP
}
```

```

Host Name: AZSHCI1-M6-N3
Enabling CredSSP
Host Name: AZSHCI1-M6-N3
Verifying cluster network interfaces

Network          Name
-----
storage(Storage_VLAN207) AZSHCI1-M6-N1 - SlotID 2 Port 2
storage(Storage_VLAN107) AZSHCI1-M6-N1 - SlotID 5 Port 1
mgmt_compute(Management) AZSHCI1-M6-N1 - vManagement(mgmt_compute)
storage(Storage_VLAN207) AZSHCI1-M6-N2 - SlotID 2 Port 2
storage(Storage_VLAN107) AZSHCI1-M6-N2 - SlotID 5 Port 1
mgmt_compute(Management) AZSHCI1-M6-N2 - vManagement(mgmt_compute)
storage(Storage_VLAN207) AZSHCI1-M6-N3 - SlotID 2 Port 2
storage(Storage_VLAN107) AZSHCI1-M6-N3 - SlotID 5 Port 1
mgmt_compute(Management) AZSHCI1-M6-N3 - vManagement(mgmt_compute)

```

Procedure 6. Verify Live Migration Exclusion List

Step 1. Run the following to get Management Cluster Network ID:

```

$Cluster = "AzS-HCI-M6-C1"

Invoke-Command $Cluster -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WSManCredSSP -Role Server -Force
}

Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Checking Management cluster network settings " -ForegroundColor Yellow

$ClusterName = (Get-cluster).Name

Get-ClusterNetwork -Cluster $ClusterName -Name "mgmt_compute_storage(Management)" | fl *

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WSManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
Get-WSManCredSSP
}

```

```

Address      : 192.168.126.0
AddressMask  : 255.255.255.192
AutoMetric   : True
Cluster      : AzS-HCI-M6-C1
Description  :
Id           : 88c30671-5333-494e-8536-fda117b88a13
Ipv4Addresses : {192.168.126.0}
Ipv4PrefixLengths : {26}
Ipv6Addresses : {}
Ipv6PrefixLengths : {}
Metric       : 68800
Name         : Management
Role         : ClusterAndClient
State        : Up

```

Step 2. Run the following to Verify Live Migration Exclusion List:

```

$Cluster = "AzS-HCI-M6-C1"

Invoke-Command $Cluster -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WSManCredSSP -Role Server -Force
}

Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Verifying Management network exclusion from Live Migration Network list " -ForegroundColor
Yellow

$ClusterName = (Get-cluster).Name

Get-ClusterResourceType -Cluster $ClusterName -Name "Virtual Machine" | Get-ClusterParameter -Name
MigrationExcludeNetworks | fl *

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WSManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
Get-WSManCredSSP
}

```

```

ClusterObject : Virtual Machine
Name           : MigrationExcludeNetworks
IsReadOnly     : False
ParameterType  : String
Value         : 88c30671-5333-494e-8536-fda117b88a13

```

For more information, go to: [https://technet.microsoft.com/en-us/library/dn550728\(v=ws.11\).aspx](https://technet.microsoft.com/en-us/library/dn550728(v=ws.11).aspx)

Procedure 7. Configure Live Migration to use SMB Protocol

Note: SMB protocol provides the best throughput for Live Migration. The default setting is Compression which is best for constrained networks.

Step 1. Run the following:

```
$Cluster = "AzS-HCI-M6-C1"
$nodes = (Get-ClusterNode -Cluster $Cluster).Name
foreach ($node in $nodes) {

Invoke-Command $node -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Configuring Live Migration to use SMB protocol" -ForegroundColor Yellow
Set-VMHost -VirtualMachineMigrationPerformanceOption SMB
Get-VMHost | fl VirtualMachineMigrationPerformanceOption

}

}
```

```
Host Name: AZS-HCI1-N1
Configuring Live Migration to use SMB protocol

VirtualMachineMigrationPerformanceOption : SMB

Host Name: AZS-HCI1-N2
Configuring Live Migration to use SMB protocol

VirtualMachineMigrationPerformanceOption : SMB

Host Name: AZS-HCI1-N3
Configuring Live Migration to use SMB protocol

VirtualMachineMigrationPerformanceOption : SMB

Host Name: AZS-HCI1-N4
Configuring Live Migration to use SMB protocol

VirtualMachineMigrationPerformanceOption : SMB
```

Procedure 8. Configure Live Migration Bandwidth Limit

Note: SMB Direct is allocated 50% of the link speed bandwidth. The following configuration parameter limits SMB Direct bandwidth allowed for Live Migration to 29%. The remaining SMB Direct bandwidth is allocated to Storage Bus Layer and Cluster Shared Volume network traffic.

Step 1. Run the following:

```
$nodes = (Get-ClusterNode -Cluster $Cluster).Name

foreach ($node in $nodes) {

Invoke-Command $node -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host "Configuring Live Migration Bandwidth Limit: 3625MB" -ForegroundColor Yellow

Set-SMBBandwidthLimit -Category LiveMigration -BytesPerSecond 3625MB

Get-SMBBandwidthLimit -Category LiveMigration

}
```

```
}
```

```
Host Name: AZS-HCI1-N1
Configuring Live Migration Bandwidth Limit: 3625MB

Host Name: AZS-HCI1-N2
Configuring Live Migration Bandwidth Limit: 3625MB
Category      Bytes Per Second PSComputerName
-----
LiveMigration 3801088000      AZS-HCI1-N1
LiveMigration 3801088000      AZS-HCI1-N2
Host Name: AZS-HCI1-N3
Configuring Live Migration Bandwidth Limit: 3625MB
LiveMigration 3801088000      AZS-HCI1-N3
Host Name: AZS-HCI1-N4
Configuring Live Migration Bandwidth Limit: 3625MB
LiveMigration 3801088000      AZS-HCI1-N4
```

Procedure 9. Create Maximum Bandwidth Limit for Management vNIC

Note: This is an optional configuration item that limits the network bandwidth to the management vNIC. The management vNIC shares total bandwidth with the bandwidth allocated to tenant network traffic. The allocated tenant network traffic bandwidth is 50 percent of the total bandwidth in a fully converged network configuration. The following configuration example sets the maximum bandwidth limit 4Gb/s (10 percent of the tenant network traffic bandwidth) for the management vNIC. This value can be adjusted as needed.

```
$nodes = (Get-ClusterNode -Cluster $Cluster).Name

foreach ($node in $nodes) {
Invoke-Command $node -scriptblock {

$MgmtBandwidthLimit = "40000000"

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host "Configuring management vNIC maximum bandwidth Limit: $MgmtBandwidthLimit" -ForegroundColor Yellow
Set-VMNetworkAdapter -ManagementOS -Name "vManagement(mgmt_compute_storage)" -MaximumBandwidth
$MgmtBandwidthLimit

Write-Host "Verifying management vNIC maximum bandwidth Limit" -ForegroundColor Yellow
(Get-VMNetworkAdapter -ManagementOS -Name "vManagement(mgmt_compute_storage)").BandwidthSetting | ft
ParentAdapter, MaximumBandwidth

}
}
```

```

Host Name: AZS-HCI2-N1
Configuring management vNIC maximum bandwidth Limit: 40000000
Verifying management vNIC maximum bandwidth Limit

ParentAdapter                                     MaximumBandwidth
-----
VMInternalNetworkAdapter, Name = 'vManagement(mgmt_compute_storage)' 40000000

Host Name: AZS-HCI2-N2
Configuring management vNIC maximum bandwidth Limit: 40000000
Verifying management vNIC maximum bandwidth Limit

ParentAdapter                                     MaximumBandwidth
-----
VMInternalNetworkAdapter, Name = 'vManagement(mgmt_compute_storage)' 40000000

Host Name: AZS-HCI2-N3
Configuring management vNIC maximum bandwidth Limit: 40000000
Verifying management vNIC maximum bandwidth Limit

ParentAdapter                                     MaximumBandwidth
-----
VMInternalNetworkAdapter, Name = 'vManagement(mgmt_compute_storage)' 40000000

```

Procedure 10. Create the File Share for the Cluster Witness

Step 1. Run the following commands:

Note: These commands require the files share witness server to be a domain member. It's recommended that the witness share is placed on a highly available scale out file server. The “-ContinuouslyAvailable” command option should be used when creating a share on a highly available scale out file server.

```

$FSW = "fsw01.ucs-spaces.lab"

$FSWDomain = "ucs-spaces.lab"

$ShareName = "FSW-AzS-HCI-M6-C1"

$SharePath = "C:\FileShareWitness\FSW-AzS-HCI-M6-C1"

Invoke-Command -ComputerName $FSW -ScriptBlock {

#Create Directory on File Share Witness
Write-Host "Creating directory on files share witness"
mkdir $Using:SharePath

#Create file share on the file share witness
Write-Host "Creating file share on file share witness"

```

```

new-smbshare -Name $Using:ShareName -Path $Using:SharePath -FullAccess "ucs-spaces.lab\Domain Admins", "ucs-
spaces.lab\AzS-HCI-M6-C1$", "ucs-spaces.lab\AzS-HCI1-N1$", "ucs-spaces.lab\AzS-HCI1-N2$", "ucs-
spaces.lab\AzS-HCI1-N3$", "ucs-spaces.lab\AzS-HCI1-N4$"

#Verify file share on file share witness
Write-Host "Verifying file share on file share witness"
Get-SmbShare -Name $Using:ShareName | ft name,path -AutoSize

#Verify file share permissions on the file share witness
Write-Host "Verifying file share permissions on the file share witness"
Get-SmbShareAccess -Name $Using:ShareName | ft -AutoSize

#Set file level permissions on the file share directory that match the file share permissions
Write-Host "Setting file level permissions on the file share directory that match the file share permissions"
Set-SmbPathAcl -ShareName $Using:ShareName

#Verify file level permissions on the file share
Write-Host "Verifying file level permissions on the file share"
Get-Acl -Path $Using:SharePath | fl
}

```

```

name          path
----          -
FSW-AzS-HCI-M6-C1 C:\FileShareWitness\FSW-AzS-HCI-M6-C1

Verifying file share permissions on the file share witness

Name          ScopeName AccountName          AccessControlType AccessRight
----          -
FSW-AzS-HCI-M6-C1 * UCS-SPACES\Domain Admins Allow Full
FSW-AzS-HCI-M6-C1 * UCS-SPACES\AzS-HCI-M6-C1$ Allow Full
FSW-AzS-HCI-M6-C1 * UCS-SPACES\AzS-HCI1-N1$ Allow Full
FSW-AzS-HCI-M6-C1 * UCS-SPACES\AzS-HCI1-N2$ Allow Full
FSW-AzS-HCI-M6-C1 * UCS-SPACES\AzS-HCI1-N3$ Allow Full
FSW-AzS-HCI-M6-C1 * UCS-SPACES\AzS-HCI1-N4$ Allow Full

Setting file level permissions on the file share directory that match the file share permissions
Verifying file level permissions on the file share

Path : Microsoft.PowerShell.Core\FileSystem::C:\FileShareWitness\FSW-AzS-HCI-M6-C1
Owner : BUILTIN\Administrators
Group : UCS-SPACES\Domain Users
Access : UCS-SPACES\Domain Admins Allow FullControl
        UCS-SPACES\AzS-HCI1-N1$ Allow FullControl
        UCS-SPACES\AzS-HCI1-N4$ Allow FullControl
        UCS-SPACES\AzS-HCI1-N2$ Allow FullControl
        UCS-SPACES\AzS-HCI1-N3$ Allow FullControl
        UCS-SPACES\AzS-HCI-M6-C1$ Allow FullControl
        NT AUTHORITY\SYSTEM Allow FullControl
        BUILTIN\Administrators Allow FullControl
        BUILTIN\Users Allow ReadAndExecute, Synchronize
        BUILTIN\Users Allow AppendData
        BUILTIN\Users Allow CreateFiles
        CREATOR OWNER Allow 268435456

```

Procedure 11. Configure File Share Witness

Step 1. Run the following:

```
$Cluster = "AzS-HCI-M6-C1"
$FSW = "fsw01.ucs-spaces.lab"
$ShareName = "FSW-AzS-HCI-M6-C1"

Set-ClusterQuorum -Cluster $Cluster -FileShareWitness \\$FSW\$ShareName
```

```
Cluster           QuorumResource
-----
AzS-HCI-M6-C1     File Share Witness
```

Procedure 12. Verify File Share Witness Path

Step 1. Run the following:

```
$Cluster = "AzS-HCI-M6-C1"
Get-ClusterResource -Cluster $Cluster -Name "File Share Witness" | Get-ClusterParameter -Name SharePath
```

```
Object           Name      Value                                     Type
-----
File Share Witness SharePath \\fsw01.ucs-spaces.lab\FSW-AzS-HCI-M6-C1 String
```

Procedure 13. Additional Cluster Quorum Witness Options

Note: Cloud Witness and none domain join files share witness can be implemented as alternate cluster witness options. Implementation details for these options can be found that the following links:

<https://docs.microsoft.com/en-us/windows-server/failover-clustering/deploy-cloud-witness>

<https://techcommunity.microsoft.com/t5/Failover-Clustering/New-File-Share-Witness-Feature-in-Windows-Server-2019/ba-p/372149>

Procedure 14. Configure Cluster-Aware Updating

Note: The Cluster-Aware Updating role will be installed after the cluster is created. The cluster-aware updating is a feature that automates downloading and installing Windows Server updates on all cluster nodes.

For more information about Cluster-Aware Updating, see: <https://docs.microsoft.com/en-us/windows-server/failover-clustering/cluster-aware-updating>

Step 1. Run the following commands to configure Cluster-Aware Updating:

```
$Cluster = "AzS-HCI-M6-C1"

Invoke-Command $Cluster -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
```

```

$Void = Enable-WSManCredSSP -Role Server -Force
}

Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host " Configuring Cluster-Aware Updating ... " -ForegroundColor Yellow
$ClusterName = (Get-cluster).Name
Add-CauClusterRole -ClusterName $ClusterName -DaysOfWeek Tuesday,Saturday -IntervalWeeks 3 -MaxFailedNodes 1
-MaxRetriesPerNode 2 -EnableFirewallRules -Force

Write-Host " Verifying Cluster-Aware Updating configuraiton " -ForegroundColor Yellow
Get-CauClusterRole -ClusterName $ClusterName | ft

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WSManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
Get-WSManCredSSP

}

```

Note: This process might take several minutes.

Adding CAU clustered role on cluster "AzS-HCI-M6-C1".
Creating the clustered role and computer account (also known as the virtual computer object or VCO)...

Selecting CAU clustered role name.
Checking if name "CAUAzS-H9e3" is in use...

Name	Value
ResourceGroupName	CAUAzS-H9e34su8 (AzS-HCI1-N4)
Status	Online
StartDate	3/10/2023 3:00:00 AM
MaxFailedNodes	1
MaxRetriesPerNode	2
EnableFirewallRules	True
FailbackMode	Immediate
DaysOfWeek	Tuesday, Saturday
IntervalWeeks	3

Procedure 15. Configure Kernel Soft Reboot for Cluster Aware Updating

Note: Kernel Soft Reboot reduces the time required to reboot a server by bypassing BIOS and firmware initiation. Kernel Soft Reboot works with Cluster Aware Updating for applying software updates. Kernel Soft Reboot cannot be used to the server when BIOS and firmware updates need to be applied.

Step 1. Run the following:

```
$Cluster = "AzS-HCI-M6-C1"

Invoke-Command $Cluster -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WSManCredSSP -Role Server -Force

}

Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
$ClusterName = (Get-cluster).Name

Write-Host " Configuring Kernel Soft Reboot  for Cluster Aware Updating ..." -ForegroundColor Yellow
Get-Cluster -Name $ClusterName | Set-ClusterParameter -Name CauEnableSoftReboot -Value 1 -Create

Write-Host " Verifying Kernel Soft Reboot configuraiton " -ForegroundColor Yellow
Get-Cluster -Name $ClusterName | Get-ClusterParameter -Name CauEnableSoftReboot | ft Name, Value

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WSManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
Get-WSManCredSSP

}

}
```

```
Host Name: AZS-HCI1-N4
Enabling CredSSP
Host Name: AZS-HCI1-N4
Configuring Kernel Soft Reboot  for Cluster Aware Updating ...
Verifying Kernel Soft Reboot configuraiton

Name                Value
----                -
CauEnableSoftReboot 1

Disabling CredSSP
Verifying that CredSSP are disabled on target server...
The machine is not configured to allow delegating fresh credentials.
This computer is not configured to receive credentials from a remote client computer.
```

Configure Storage Spaces Direct

This section contains the following procedures:

- [Enable Storage Spaces Direct](#)
- [Verify the newly created Storage Pool, NVMe SSD Cache, and Storage Tiers](#)
- [Create a Virtual Disk with Mirror Resiliency by using the Performance Tier template](#)
- [Create Storage QoS Policy](#)
- [Register the Azure Stack HCI Cluster with Azure](#)
- [Create a Virtual Machine with Failover Capability](#)

Procedure 1. Enable Storage Spaces Direct

The following command automatically enables Storage Spaces Direct and configures the following:

- **Create a pool:** Creates a single large pool that has a name like “S2D on Cluster1”.
- **Configures the Storage Spaces Direct caches:** If there is more than one media (drive) type available for Storage Spaces Direct use, it enables the fastest as cache devices (read and write in most cases).
- **Tiers:** Creates 2 tiers as default tiers. One is called “Capacity” and the other called “Performance.” The cmdlet analyzes the devices and configures each tier with the mix of device types and resiliency.

```
$Cluster = "AzS-HCI-M6-C1"

Invoke-Command $Cluster -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WSManCredSSP -Role Server -Force

}

Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
$ClusterName = (Get-cluster).Name

Write-Host " Enabling Storage Spaces Direct " -ForegroundColor Yellow
Enable-ClusterStorageSpacesDirect -Confirm:$false

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WSManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
Get-WSManCredSSP

}
```

Enable-ClusterStorageSpacesDirect -Confirm:\$false.
0/1 completed.

Enabling cluster Storage Spaces Direct.
Starting health providers, 59% Complete.

```
CacheMetadataReserveBytes : 34359738368
CacheModeHDD              : ReadWrite
CacheModeSSD              : WriteOnly
CachePageSizeKBytes       : 16
CacheState                 : Enabled
State                     : Enabled
PSComputerName            : AzS-HCI-M6-C1
```

Procedure 2. Verify the newly created Storage Pool, NVMe SSD Cache, and Storage Tiers

Step 1. Run the following:

```
$Cluster = "AzS-HCI-M6-C1"

Invoke-Command $Cluster -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WManCredSSP -Role Server -Force
}

Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
$ClusterName = (Get-cluster).Name

Write-Host " Verifying Storage Pools " -ForegroundColor Yellow
Get-StoragePool | ft friendlyname, OperationalStatus, HealthStatus, IsPrimordial, IsReadOnly

Write-Host " Verifying NVMe SSD Cache Tier " -ForegroundColor Yellow
Get-PhysicalDisk | ? Usage -eq "Journal" | ft FriendlyName, CanPool, HealthStatus, Usage, Size

Write-Host " Verifying Storage Tier configuration " -ForegroundColor Yellow
Get-storagetier | ft FriendlyName, ResiliencySettingName, MediaType, NumberOfDataCopies,
PhysicalDiskRedundancy

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
```

```
Get-WSManCredSSP
}
```

```
Host Name: AZS-HCI1-N4
Enabling CredSSP
Host Name: AZS-HCI1-N4
Verifying Storage Pools

friendlyname      OperationalStatus HealthStatus IsPrimordial IsReadOnly
-----
Primordial        OK                Healthy      True         False
Primordial        OK                Healthy      True         False
S2D on AzS-HCI-M6-C1 OK                Healthy      False        False

Verifying NVMe SSD Cache Tier
Verifying Storage Tier configuration

FriendlyName ResiliencySettingName MediaType NumberOfDataCopies PhysicalDiskRedundancy
-----
ParityOnSSD Parity                SSD          1                2
Performance Mirror                SSD          3                2
MirrorOnSSD Mirror                SSD          3                2
Capacity Parity                SSD          1                2
```

Procedure 3. Create a Virtual Disk with Mirror Resiliency by using the Performance Tier template

It is optimal to create a virtual disk in multiples that match the number of cluster nodes that will run virtual machines. For example, the number of virtual disks for cluster with 4 nodes should be 4, 8, 12, and so on.

Note: The following link contains Microsoft recommendations for volume capacity planning:
<https://docs.microsoft.com/en-us/azure-stack/hci/concepts/plan-volumes>

The **New-Volume** cmdlet simplifies deployments as it ties together a long list of operations that would otherwise have to be done in individual commands such as creating the virtual disk, partitioning, and formatting the virtual disk, adding the virtual disk to the cluster, and converting it into CSVFS.

Step 1. Run the following command to create the multiple times. Update the Virtual Disk friendly name and size as required:

```
$Cluster = "AzS-HCI-M6-C1"
Invoke-Command $Cluster -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WSManCredSSP -Role Server -Force
}

Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
$ClusterName = (Get-cluster).Name

Write-Host " Creating Virtual Disk " -ForegroundColor Yellow
```

```

New-Volume -StoragePoolFriendlyName "S2D*" -FriendlyName VDisk01 -FileSystem CSVFS_ReFS -
ResiliencySettingName Mirror -Size 4TB

Write-Host " Disabling CredSSP" -ForegroundColor Yellow

Disable-WManCredSSP -Role Server

Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow

Get-WManCredSSP

}

```

```

Host Name: AZS-HCI1-N4
Enabling CredSSP
Host Name: AZS-HCI1-N4
Creating Virtual Disk

PSComputerName      : AzS-HCI-M6-C1
RunspaceID          : 33819760-37b1-42c1-a5f3-0a89e0dc722a
ObjectID            : {1}\AzS-HCI-M6-C1\root/Microsoft/Windows/Storage/Providers_v2\WSP_Volume.ObjectID
                   : abae-0e7eb0b18d74\

PassThroughClass    :
PassThroughIds      :
PassThroughNamespace :
PassThroughServer    :
UniqueID            : \\?\Volume{2f430749-3354-4c99-abae-0e7eb0b18d74}\
AllocationUnitSize  : 4096
DedupMode           : 0
DriveLetter         :
DriveType           : 3
FileSystem           : CSVFS
FileSystemLabel      : VDisk01
FileSystemType       : 32769
HealthStatus        : 0
OperationalStatus   : {2}
Path                : \\?\Volume{2f430749-3354-4c99-abae-0e7eb0b18d74}\
Size                : 4397979402240
SizeRemaining       : 4371716718592

Disabling CredSSP
Verifying that CredSSP are disabled on target server...
The machine is not configured to allow delegating fresh credentials.
This computer is not configured to receive credentials from a remote client computer.

```

Step 2. The virtual disk status can be viewed by running the following command:

```

$cluster = "AzS-HCI-M6-C1"
Invoke-Command $cluster -scriptblock {Get-VirtualDisk}

```

Step 3. Run the following command to view the path of the new virtual disk:

```

$cluster = "AzS-HCI-M6-C1"
Invoke-Command $cluster -scriptblock {Get-ClusterSharedVolume | fl Name,SharedVolumeInfo,OwnerNode}

```

```

Name           : Cluster Virtual Disk (VDisk01)
SharedVolumeInfo : {C:\ClusterStorage\VDisk01}
OwnerNode      : AZS-HCI1-N3

```

Note: The Cluster Shared Volume ownership can be realigned with the cluster nodes if desired. It is optimal when cluster virtual disk ownership is evenly distributed across the cluster nodes.

Procedure 4. Create Storage QoS Policy

Note: Storage QoS Policies limit the maximum IOPS that can be consumed by a virtual disk. These policies can prevent a “noisy neighbor” scenario where an individual virtual machine consumes an undesirable amount of storage IOPS and bandwidth, thus starving the available IOPS and bandwidth for other tenant virtual machines. The storage QoS policy is first created, and the policy ID is applied to a virtual disk (VHDX).

Step 1. The minimum and maximum IOPS values can be adjusted to as needed for the specific environment, by running the following command:

```
$Cluster = "AzS-HCI-M6-C1"
Invoke-Command $Cluster -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Creating and verifying Storage Polices" -ForegroundColor Yellow

New-StorageQoSPolicy -Name Copper -MinimumIops 50 -MaximumIops 100 -PolicyType Dedicated
New-StorageQoSPolicy -Name Bronze -MinimumIops 100 -MaximumIops 250 -PolicyType Dedicated
New-StorageQoSPolicy -Name Silver -MinimumIops 200 -MaximumIops 500 -PolicyType Dedicated
New-StorageQoSPolicy -Name Gold -MinimumIops 500 -MaximumIops 5000 -PolicyType Dedicated
New-StorageQoSPolicy -Name Platinum -MinimumIops 1000 -MaximumIops 10000 -PolicyType Dedicated

Get-StorageQoSPolicy | ft Name,Status, MinimumIops,MaximumIops,MaximumIOBandwidth,PolicyID

}
```

Name	Status	MinimumIops	MaximumIops	MaximumIOBandwidth	PolicyId
Default	Ok	0	0	0	00000000-0000-0000-0000-000000000000
Platinum	Ok	1000	10000	0	5e383722-6649-4a29-998b-a034cd2d6c1e
Silver	Ok	200	500	0	4c8fe979-b978-445d-a504-c703a5d038a1
Bronze	Ok	100	250	0	10f98289-6578-4465-89d5-f17130b72db8
Gold	Ok	500	5000	0	b311a2b0-c5e8-4092-a8c5-0c5686614496
Copper	Ok	50	100	0	bdf977ca-ffbe-4621-807f-54fe206d3da5

Note: The Maximum IOPS value is in units of 8KB-normalized. IO larger than 8KB is treated as multiple normalized IOPS. For example, 64KB IO is treated as 8 normalized IOPS.

Procedure 5. Register the Azure Stack HCI Cluster with Azure

Note: Follow the documentation at the following link to register the cluster with the Azure subscription. The registration must be completed successfully in order to create virtual machines in the Azure Stack HCI cluster: <https://docs.microsoft.com/en-us/azure-stack/hci/deploy/register-with-azure>

Procedure 6. Create a Virtual Machine with Failover Capability

Note: The following script is an example of creating a virtual machine with failover capability. This example includes creation of a VHDX file for the virtual machine with an attached storage QoS policy:


```

$Cluster = "AzS-HCI-M6-C1"
Invoke-Command $Cluster -Credential $Creds -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
Write-Host " Enabling CredSSP" -ForegroundColor Yellow
$Void = Enable-WSManCredSSP -Role Server -Force
}

Invoke-Command $Cluster -Credential $Creds -authentication Credssp -scriptblock {

write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
$CSVPath = ((Get-ClusterSharedVolume).SharedVolumeInfo).FriendlyVolumeName
$VHDPATH = "$CSVPath\VM01-Disk01.vhdx"
$VMSwitch = (Get-VMSwitch).Name
$VMName = "VM01"
$VMPath = "$CSVPath\VirtualMachines"
$VMMemoryCapacity = 8GB

Write-Host "Creating VHDX $VHDPATH ..." -ForegroundColor Yellow
New-VHD -Path $CSVPath\VM01-Disk01.vhdx -Fixed -SizeBytes 100GB

Write-Host "Creating virtual machine $VMName with memory capacity $VMMemoryCapacity ..." -ForegroundColor Yellow
New-VM -Name $VMName -Path $VMPath -MemoryStartupBytes $VMMemoryCapacity -VHDPATH $VHDPATH -Generation 2 -
SwitchName $VMSwitch

$BronzeStorageQoSPolicyID = (Get-StorageQoSPolicy -Name Silver).PolicyID

Write-Host "Setting QoS Plicy for virtual machine $VMName ..." -ForegroundColor Yellow
Get-VM -VMName $VMName | Get-VMHardDiskDrive | Set-VMHardDiskDrive -QoSPolicyID $BronzeStorageQoSPolicyID

Write-Host "Clustering the virtual machine $VMName ..." -ForegroundColor Yellow
Get-VM -Name $VMName | Add-ClusterVirtualMachineRole -Name $VMName

Write-Host " Disabling CredSSP" -ForegroundColor Yellow
Disable-WSManCredSSP -Role Server
Write-Host " Verifying that CredSSP are disabled on target server..." -ForegroundColor Yellow
Get-WSManCredSSP
}

```

```
Host Name: AZS-HCI1-N4
Creating VHDX C:\ClusterStorage\VDisk01\VM01-Disk01.vhdx ...

Number          :
PSComputerName  : AzS-HCI-M6-C1
RunspaceId     : 8fe207d4-4631-4f27-ad9c-c31275c03fb2
ComputerName   : AZS-HCI1-N4
Path           : C:\ClusterStorage\VDisk01\VM01-Disk01.vhdx
VhdFormat      : VHDX
VhdType        : Fixed
FileSize       : 107378376704
Size           : 107374182400
MinimumSize    :
LogicalSectorSize : 512
PhysicalSectorSize : 4096
BlockSize      : 0
ParentPath     :
DiskIdentifier  : 96875311-7641-4FEB-A76D-192EB7FCA536
FragmentationPercentage : 0
Alignment      : 1
Attached       : False
DiskNumber     :
IsPMEMCompatible : False
AddressAbstractionType : None

Creating virtual machine VM01 with memory capacity 8589934592 ...
Name           : VM01
State          : Off
CpuUsage       : 0
MemoryAssigned : 0
MemoryDemand   : 0
MemoryStatus   :
Uptime         : 00:00:00
Status         : Operating normally
ReplicationState : Disabled
Generation     : 2
PSComputerName : AzS-HCI-M6-C1

Setting QoS Policy for virtual machine VM01 ...
Clustering the virtual machine VM01 ...
Name           : VM01
OwnerNode      : AZS-HCI1-N4
State          : Offline
PSComputerName : AzS-HCI-M6-C1
```

Appendix

This appendix contains the following:

- [Reference Links](#)
- [Cabling Information](#)
- [Remote Management Host](#)
- [Launch Server KVM Instance to Install the Operating System](#)
- [Locate Windows Driver Required for Cisco UCS C240 M6 or M7 Server](#)
- [Add Drivers and Windows Updates to a Windows Installation Image](#)
- [Install and Configure DHCP Server Feature](#)
- [Configure Bitlocker for System Volume](#)
- [ToR Switch Configuration](#)

Reference Links

Cluster-Aware Updating: <https://docs.microsoft.com/en-us/windows-server/failover-clustering/cluster-aware-updating>

Active Memory Dump: <https://techcommunity.microsoft.com/t5/failover-clustering/windows-server-2016-failover-cluster-troubleshooting/ba-p/372008>

Cabling Information


In this section, detailed cabling connectivity information is described for all topologies. For specific details on supported cable information, refer to the following links:

<https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/nvidia-mellanox-connectx-6-ethernet-smartnic-data-sheet.pdf>

<https://tmgmatrix.cisco.com/?si=Mellanox>

Table 8. Cabling Map for 4-node (with 2 NICs) fully converged Azure Stack HCI cluster


Nexus 9336C-FX2 – ToR-A					Nexus 9336C-FX2 – ToR-B				
From		To		Connection Type	From		To		Connection Type
S-Device	Port	D-Device	NVIDIA ConnectX-6x Port #		S-Device	Port	D-Device	NVIDIA ConnectX-6x Port #	
ToR-A	eth1/1	Node 1	SlotID 2 Port 1	QSFP-100G-CU3M	ToR-B	eth1/1	Node 1	SlotID 2 Port 2	QSFP-100G-CU3M
ToR-A	eth1/2	Node 2	SlotID 2 Port 1	QSFP-100G-CU3M	ToR-B	eth1/2	Node 2	SlotID 2 Port 2	QSFP-100G-CU3M
ToR-A	eth1/3	Node 3	SlotID 2 Port 1	QSFP-100G-CU3M	ToR-B	eth1/3	Node 3	SlotID 2 Port 2	QSFP-100G-CU3M
ToR-A	eth1/4	Node 4	SlotID 2 Port 1	QSFP-100G-CU3M	ToR-B	eth1/4	Node 4	SlotID 2 Port 2	QSFP-100G-CU3M
ToR-A	eth1/31	ToR-B	eth1/31	QSFP-100G-CU3M	ToR-B	eth1/31	ToR-A	eth1/31	QSFP-100G-CU3M
ToR-A	eth1/32	ToR-B	eth1/32	QSFP-100G-CU3M	ToR-B	eth1/32	ToR-A	eth1/32	QSFP-100G-CU3M
ToR-A	MGMT	Cust. OOBM	NA	Cat6	ToR-B	MGMT	Cust. OOBM	NA	Cat6



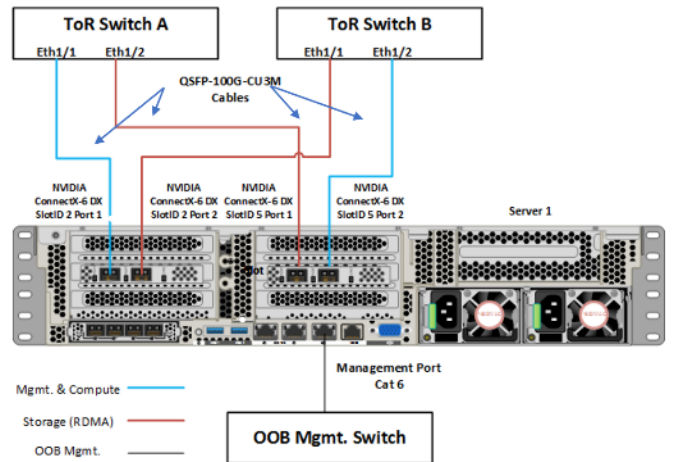
Fully Converged Network Topology (with 2 NICs)

Table 9. Cabling Map for 4-node (with 4 NICs) fully converged and non-converged Azure Stack HCI cluster

Nexus 9336C-FX2 – ToR-A					Nexus 9336C-FX2 – ToR-B				
From		To			From		To		
S-Device	Port	D-Device	NVIDIA ConnectX-6x Port #	Connection Type	S-Device	Port	D-Device	NVIDIA ConnectX-6x Port #	Connection Type
ToR-A	eth1/1	Node 1	SlotID 2 Port 1	QSFP-100G-CU3M	ToR-B	eth1/1	Node 1	SlotID 2 Port 2	QSFP-100G-CU3M
ToR-A	eth1/2	Node 1	SlotID 5 Port 1	QSFP-100G-CU3M	ToR-B	eth1/2	Node 1	SlotID 5 Port 2	QSFP-100G-CU3M
ToR-A	eth1/3	Node 2	SlotID 2 Port 1	QSFP-100G-CU3M	ToR-B	eth1/3	Node 2	SlotID 2 Port 2	QSFP-100G-CU3M
ToR-A	eth1/4	Node 2	SlotID 5 Port 1	QSFP-100G-CU3M	ToR-B	eth1/4	Node 2	SlotID 5 Port 2	QSFP-100G-CU3M
ToR-A	eth1/5	Node 3	SlotID 2 Port 1	QSFP-100G-CU3M	ToR-B	eth1/5	Node 3	SlotID 2 Port 2	QSFP-100G-CU3M
ToR-A	eth1/6	Node 3	SlotID 5 Port 1	QSFP-100G-CU3M	ToR-B	eth1/6	Node 3	SlotID 5 Port 2	QSFP-100G-CU3M
ToR-A	eth1/7	Node 4	SlotID 2 Port 1	QSFP-100G-CU3M	ToR-B	eth1/7	Node 4	SlotID 2 Port 2	QSFP-100G-CU3M
ToR-A	eth1/8	Node 4	SlotID 5 Port 1	QSFP-100G-CU3M	ToR-B	eth1/8	Node 4	SlotID 5 Port 2	QSFP-100G-CU3M
ToR-A	eth1/31	ToR-B	eth1/31	QSFP-100G-CU3M	ToR-B	eth1/31	ToR-A	eth1/31	QSFP-100G-CU3M
ToR-A	eth1/32	ToR-B	eth1/32	QSFP-100G-CU3M	ToR-B	eth1/32	ToR-A	eth1/32	QSFP-100G-CU3M
ToR-A	MGMT	Cust. OOBM	NA	Cat6	ToR-B	MGMT	Cust. OOBM	NA	Cat6



Fully Converged Network Topology with 4 NICs

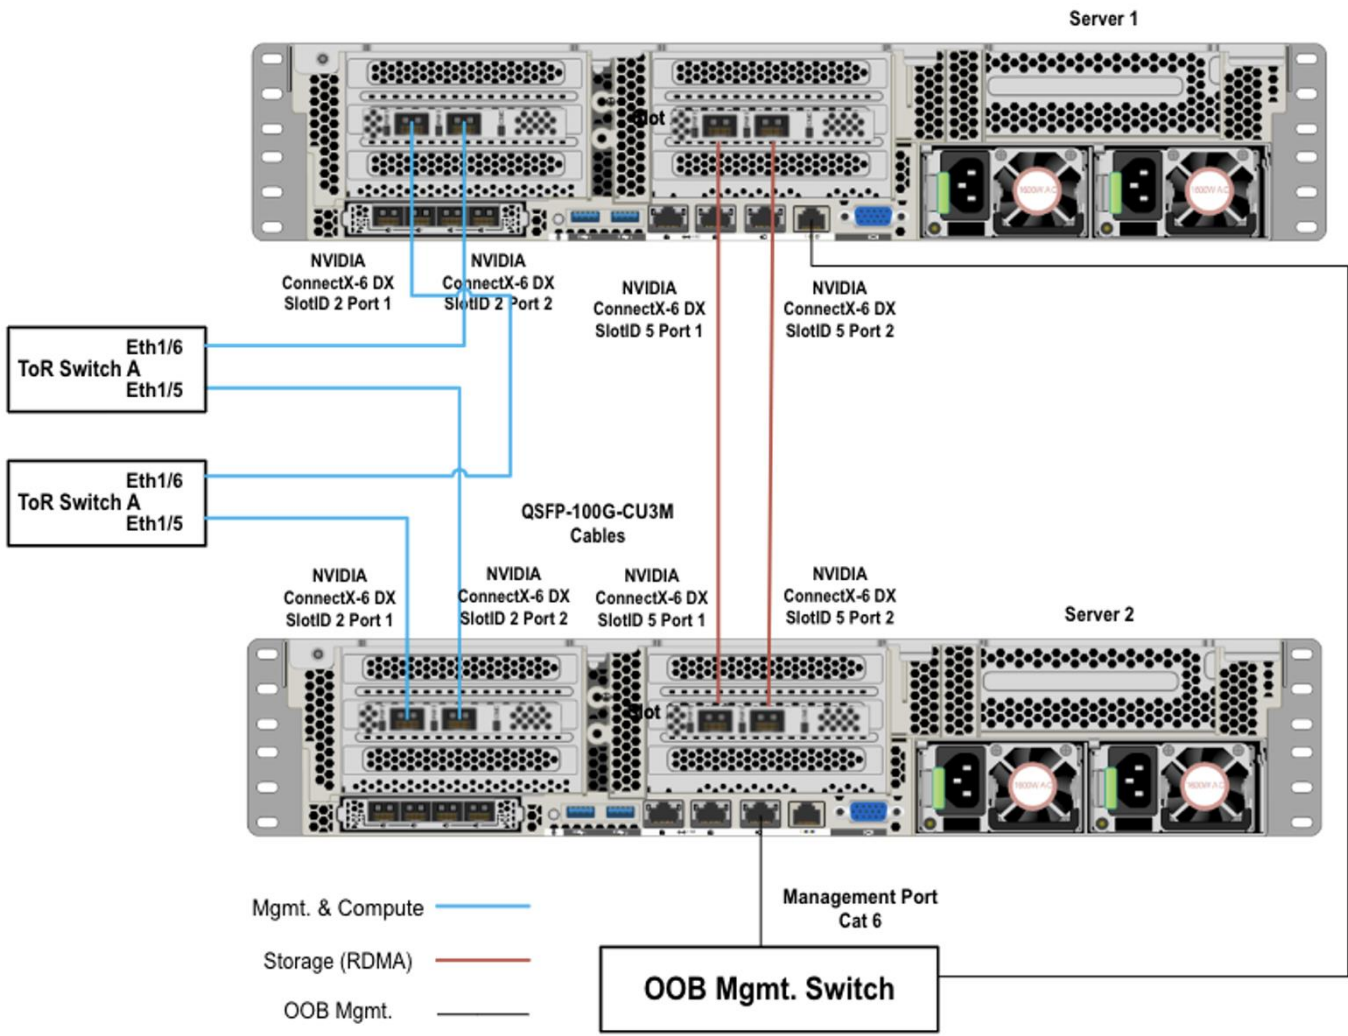


Non-Converged Network Topology with 4 NICs

Table 10. Cabling Map for 2-node (with 4 NICs) storage switchless topology for Azure Stack HCI cluster

Nexus 9336C-FX2 – ToR-A					Nexus 9336C-FX2 – ToR-B				
From		To			From		To		
S-Device	Port	D-Device	NVIDIA ConnectX-6x Port #	Connection Type	S-Device	Port	D-Device	NVIDIA ConnectX-6x Port #	Connection Type
ToR-A	eth1/1	Node 1	SlotID 2 Port 1	QSFP-100G-CU3M	ToR-B	eth1/1	Node 1	SlotID 2 Port 2	QSFP-100G-CU3M
ToR-A	eth1/2	Node 2	SlotID 2 Port 1	QSFP-100G-CU3M	ToR-B	eth1/2	Node 2	SlotID 2 Port 2	QSFP-100G-CU3M
ToR-A	eth1/31	ToR-B	eth1/31	QSFP-100G-CU3M	ToR-B	eth1/31	ToR-A	eth1/31	QSFP-100G-CU3M
ToR-A	eth1/32	ToR-B	eth1/32	QSFP-100G-CU3M	ToR-B	eth1/32	ToR-A	eth1/32	QSFP-100G-CU3M
ToR-A	MGMT	Cust. OOBM	NA	Cat6	ToR-B	MGMT	Cust. OOBM	NA	Cat6

Server-to-Server Connections				
From		To		
S-Device	NVIDIA ConnectX-6x Port #	D-Device	NVIDIA ConnectX-6x Port #	Connection Type
Node 1	SlotID 2 Port 1	Node 2	SlotID 2 Port 1	QSFP-100G-CU3M
Node 1	SlotID 2 Port 2	Node 2	SlotID 2 Port 2	QSFP-100G-CU3M



2-Node Storage Switchless Network Topology

Remote Management Host

The required Windows features are as follows:

- Clustering
- Hyper-V Management
- Group Policy Management
- Bitlocker Recovery Password Viewer
- Active Directory Management Tools

```
#Install required management modules
```

```
Add-WindowsFeature -Name RSAT-Hyper-V-Tools,RSAT-ADDS-Tools, RSAT-Clustering, RSAT-Clustering-MgmtRSAT-Clustering-PowerShell, RSAT-Feature-Tools-BitLocker-BdeAducExt,GPMC -IncludeManagementTools
```

```

Install-Module AZ.ConnectedMachine -force

#Update download provider modules for downloading modules from PSGallery
Set-PSRepository -Name "PSGallery" -InstallationPolicy Trusted
Install-PackageProvider -Name NuGet -Force
Install-Module -Name PowershellGet -Force -Confirm:$false
#Close and restart the PowerShell Windows before proceeding

#Configure WinRM for remote management of nodes
winrm quickconfig

#Enable sending remote management commands to the cluster nodes
$nodes = ("AzS-HCI1-N1", " AzS-HCI1-N2", " AzS-HCI1-N3", " AzS-HCI1-N4")
Enable-WSManCredSSP -Role "Client" -DelegateComputer $nodes

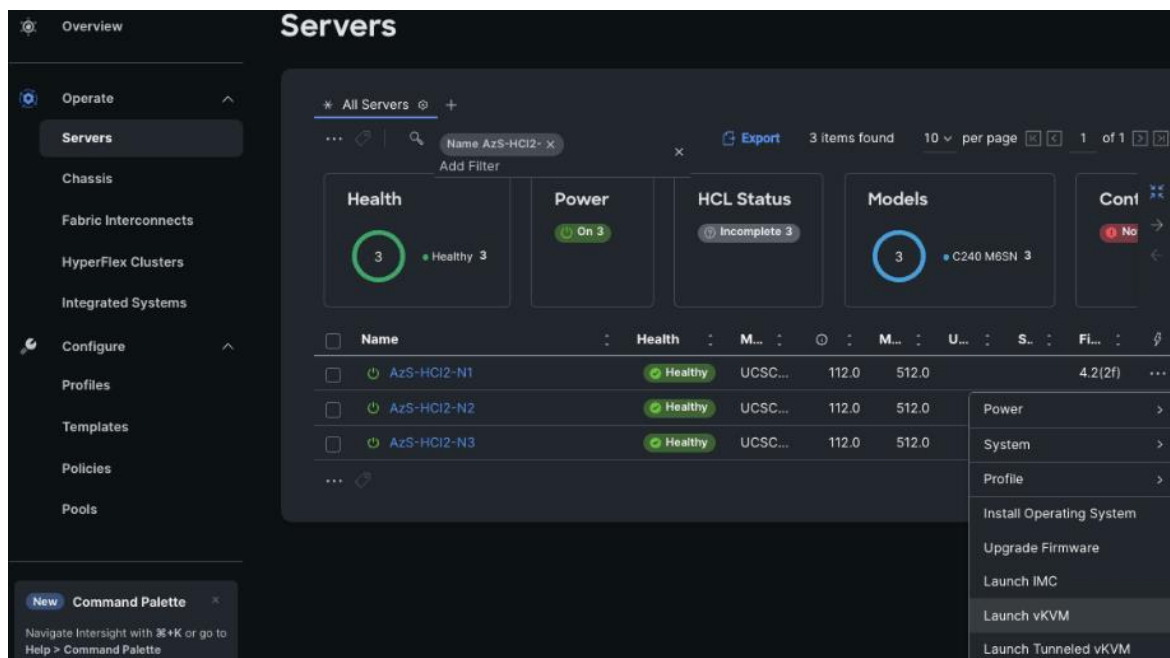
```

Launch Server KVM Instance to Install the Operating System

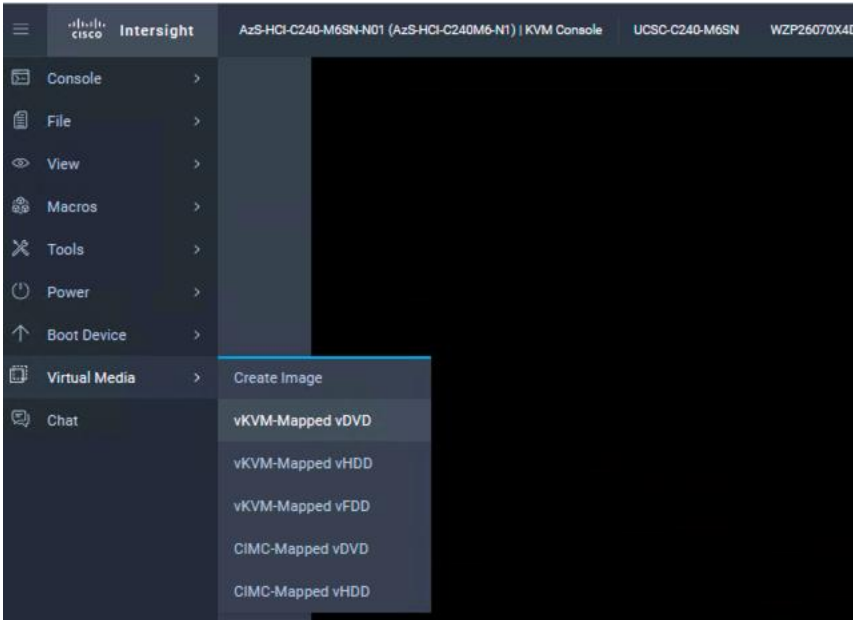
Procedure 1. Install the OS by launching the Server KVM Instance

Install the Azur Stack HCI OS 22H2 using PXE boot or a vMedia mapped installation ISO. This section explains the steps to install OS using vMedia method. Installation of OS using PXE boot is out of the scope of this document.

Step 1. From the Server tab in Cisco Intersight, select **Servers**. From the list of options select **Launch vKVM**.




Step 2. From the left pane of KVM page, navigate down to **Virtual Media** and select **vKVM-Mapped DVD**.

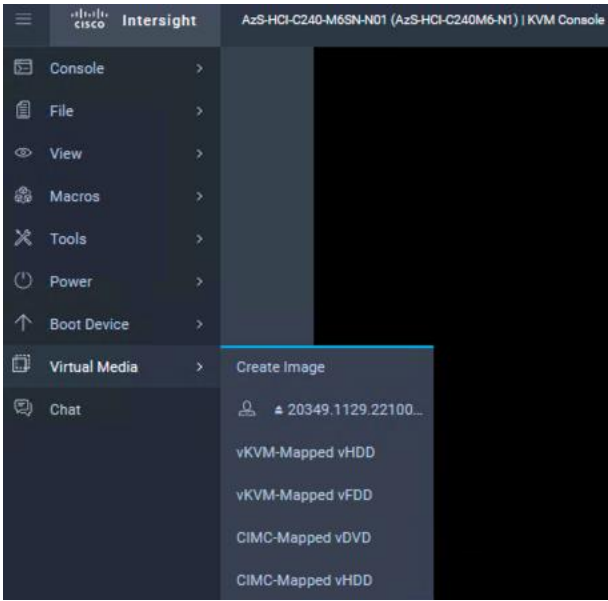


Step 3. A Map Virtual Media - CD/DVD window displays, click **Browse**.

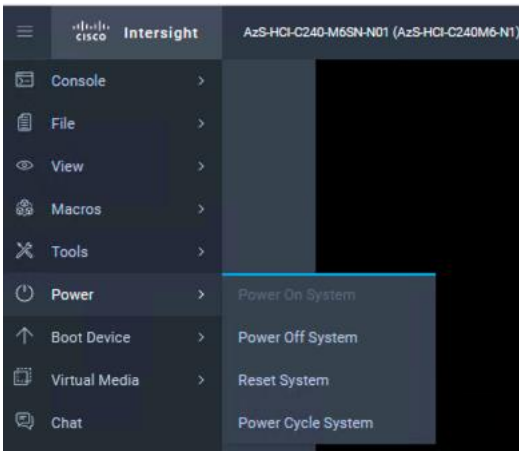
Step 4. Select the downloaded **Azure Stack HCI OS 22H2** and click **Map Drive**.



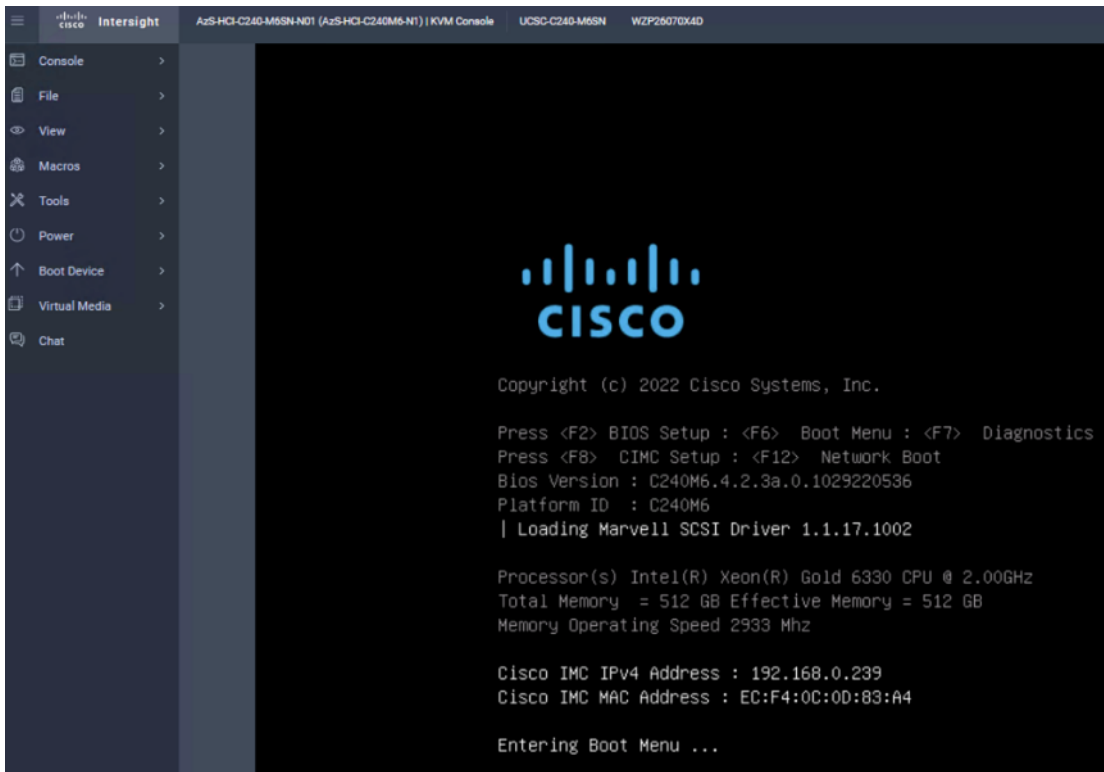
Step 5. Verify the file is selected by clicking **Virtual Media**.



Step 6. Navigate to **Power** and click **Power Cycle System** to restart the server.



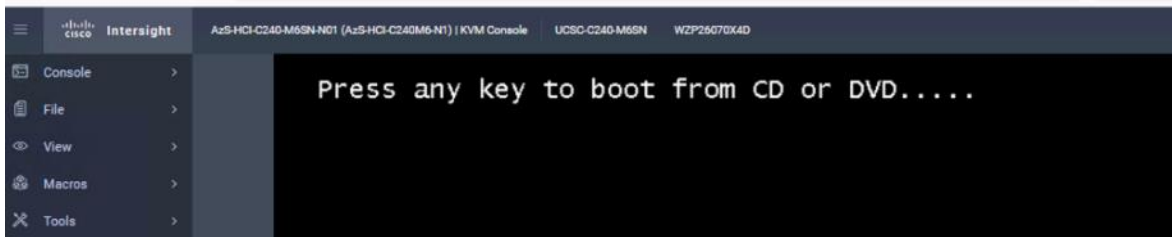
Step 7. During the POST, press **F6** to launch Boot Menu.




Step 8. In the Select boot device, select **Cisco vKVM-Mapped vDVD** and press **Enter**.



Step 9. Wait for Press any key to boot from CD or DVD on the screen and press any key to launch the OS installation.



Step 10. The installation begins. Next, follow the interactive steps and complete the installation.

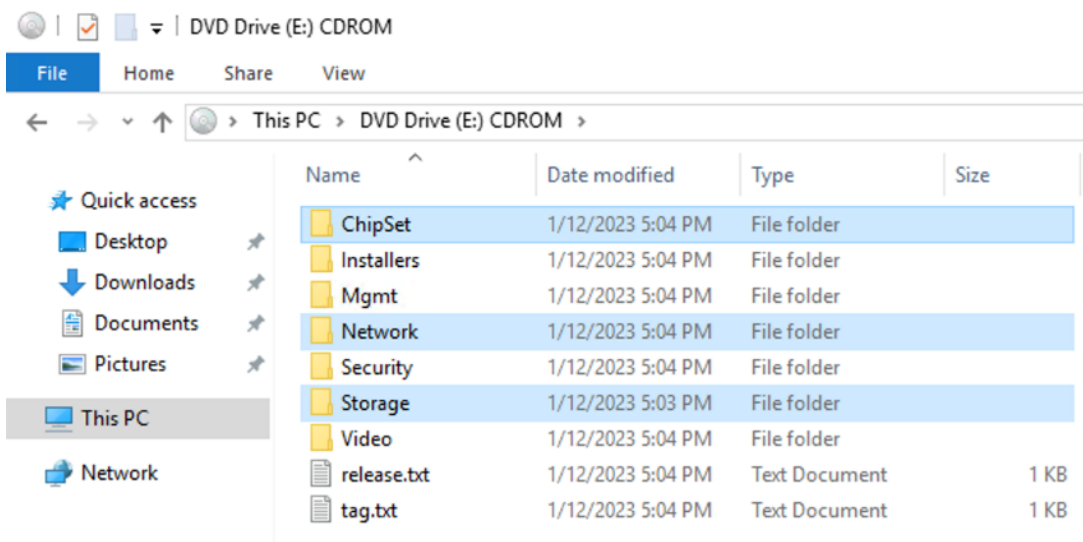


Locate Windows Driver Required for Cisco UCS C240 M6 or M7 Server

Procedure 1. Install Windows Drivers

Post OS installation, download the relevant Windows driver image for the Cisco UCS Standalone Server software (4.2.3x for M6 servers or 4.3.x for M7 servers) from the download portal [Software Download - Cisco Systems](#) and install the drivers for Chipset, Storage, and Network.

Step 1. Mount the downloaded iso image for Windows drivers only.



Step 2. Copy the following files from the mounted drive to a separate folder. Copy this folder with drivers to all the Cisco UCS C240 M6/M7 servers.

```
.\ChipSet\Intel\ChipsetSoftware\10.1.18807.8279\SetupChipset.exe
```

```
.\Network\Mellanox\ConnectX4-5-6\W2K22\MLNX_WinOF2-3_0_50000_All_x64.exe
```

```
.\Storage\Intel\C600\W2K22\*.*
```

Procedure 2. Intel Chipset Installation

Step 1. Run the following command on all the nodes to install the chipset drivers. The system will restart automatically in couple of minutes after the chipset installation in unattended silent mode. Monitor and wait for system to restart.

```
SetupChipset.exe -silent
```

```
PS C:\Users\Administrator> C:\Deploy\C240M6-4.2.2d-Drivers\Intel\SetupChipset.exe -silent
PS C:\Users\Administrator> _
```

Procedure 3. NVIDIA/Mellanox ConnectX-6 DX Driver Installation

Step 1. Run the following command on all the nodes to install the drivers for NVIDIA (Mellanox) in unattended mode.

Unattended install

```
MLNX_WinOF2-[Driver/Version]_<revision_version>_All_-Arch.exe /S /v/qn
```

Unattended install with Logs

```
MLNX_WinOF2-[Driver/Version]_<revision_version>_All_-Arch.exe /S /v/qn /v"/l*vx [LogFile]"
```

```
PS C:\Users\Administrator> C:\Deploy\C240M6-4.2.2d-Drivers\MLNX\2.80\MLNX_WinOF2-2_80_50000_All_x64.exe /S /v/q /v
"/l*vx c:\mlnx-log-2.80"
PS C:\Users\Administrator> dir c:\

Directory: C:\

Mode                LastWriteTime         Length Name
----                -
d-----          2/3/2023   4:25 AM          Deploy
d-----          2/3/2023   4:54 AM          msinfo32-before-any-driver-install
d-----          5/8/2021   1:15 AM          PerfLogs
d-r-----        2/3/2023   6:09 AM          Program Files
d-----          5/8/2021   2:34 AM          Program Files (x86)
d-r-----        2/2/2023   9:16 AM          Users
d-----          2/3/2023   4:52 AM          Windows
-a-----          2/3/2023   6:09 AM      875542 mlnx-log-2.80
```

Step 2. Alternatively, run the following command to extract only the driver files and use pnputil command to install the drivers:

```
MLNX_WinOF2-2_0_<revision_version>_All_x64.exe /a /vMT_DRIVERS_ONLY=1
```

Procedure 4. Storage Drivers Install

Step 1. Run the following command to install the Intel storage (MegaSR) drivers:

```
pnputil.exe /add-driver C:\Deploy\C240M6-4.2.3b-Drivers\Storage\Intel\C600\W2K22\*.inf /install
```

```
PS C:\Users\Administrator> pnputil.exe /add-driver C:\Deploy\C240M6-4.2.3b-Drivers\Storage\Intel\C600\W2K22\*.inf /install
Microsoft PnP Utility
Adding driver package: MegaSR1.inf
Driver package added successfully.
Published Name:          oem11.inf
Adding driver package: nodev.inf
Driver package added successfully.
Published Name:          oem12.inf
Total driver packages: 2
Added driver packages: 2
PS C:\Users\Administrator> _
```

Note: All drivers can be installed using PNPUtil.exe.

The following PNPUtile.exe example can be used to install drivers:

```
pnputil /add-driver C:\temp\drivers \*.inf
```

PNPUtil.exe documentation can be found here: <https://docs.microsoft.com/en-us/windows-hardware/drivers/devtest/pnputil>

Add Drivers and Windows Updates to a Windows Installation Image

A Windows ISO image includes boot.wim and install.wim files that are used for installation. The following are the PowerShell cmdlets to inject drivers into these .wim files.

- Get-WindowsImage: <https://docs.microsoft.com/en-us/powershell/module/dism/get-windowsimage?view=win10-ps>
- Mount-WindowsImage: <https://docs.microsoft.com/en-us/powershell/module/dism/mount-windowsimage?view=win10-ps>
- Add-WindowsDriver: <https://docs.microsoft.com/en-us/powershell/module/dism/add-windowsdriver?view=win10-ps>
- Dismount-WindowsImage: <https://docs.microsoft.com/en-us/powershell/module/dism/dismount-windowsimage?view=win10-ps>

Procedure 1. Prepare Driver Injection Computer

Step 1. Copy contents of Windows Server 2019 ISO distribution ISO, including boot.wim and install.wim, to a computer disk that will be used to inject the drivers.

Example:

Destination path = C:\temp\Source-ISO

Step 2. Copy required drivers into a subdirectory on the server. Each driver should have its own subdirectory. Each driver should include a .sys, .inf, and a .cat file at minimum. Drivers cannot be in a zip file or exe file. Chip-set drivers need to be extracted prior to injection.

Example:

Destination path: C:\temp\drivers

Step 3. Create a subdirectory for mounting the target image.

Example:

```
md C:\temp\offline
```

Procedure 2. Inject Drivers into boot.wim Images

Step 1. Identify available images in the boot file (there should be two).

Example:

```
Get-WindowsImage -ImagePath C:\temp\Source-ISO \boot.wim
```

Step 2. Identify the index for the index number of the image that needs drivers.

Step 3. Mount the target image.

Example:

```
Mount-WindowsImage -ImagePath C:\temp\Source-ISO \boot.wim -Index 2 -Path C:\temp\offline
```

Step 4. Add drivers to the mounted image. You only need to add the drivers for devices that need to be accessed during the preinstallation phase and are not in the Windows distribution. This may be the boot device drivers and network drivers.

Example:

```
Add-WindowsDriver -Path .\offline -Driver C:\temp\drivers\[NetworkDriver]
Add-WindowsDriver -Path .\offline -Driver C:\temp\drivers\[BootDeviceDriver]
```

Step 5. Save and dismount the image.

Example:

```
Dismount-WindowsImage -Path c:\temp\offline -save
```

Step 6. Repeat steps 1 - 5 for the other images in the boot.wim file if necessary.

Procedure 3. Inject Drivers into install.wim images

Step 1. Identify available images in the boot file (there should be two).

Example:

```
Get-WindowsImage -ImagePath C:\temp\Source-ISO\install.wim
```

Step 2. Identify the index for the index number of the image that needs drivers.

Step 3. Mount the target image.

Example:

```
Mount-WindowsImage -ImagePath C:\temp\Source-ISO\install.wim -Index 4 -Path C:\temp\offline
```

Step 4. Add drivers to the mounted image. You only need to add all required drivers.

Example:

```
Add-WindowsDriver -Path C:\temp \offline -Driver C:\temp \drivers -Recurse
```

Step 5. Save and dismount the image.

Example:

```
Dismount-WindowsImage -Path c:\temp\offline -save
```

Step 6. Repeat steps 1 - 5 for the other images in the install.wim file if necessary.

The updated install.wim and boot.wim can be copied to and PXE server that is used for deployment. WDS (Windows Deployment Service) is an example of a PXE server that can be used to deploy the Windows operating system.

Create an ISO image with Update .WIM Files

If a PXE server is unavailable for executing deployments, the operating system can be installed using and Windows installation ISO image. A new ISO image must be created with the updated .WIM installation files.

OSCDIMG.exe is a command line tool that can be used to create a new ISO installation image using the updated files. This tool is part of if the Automation Deployment Kit (ADK).

<https://docs.microsoft.com/en-us/windows-hardware/get-started/adk-install>

<https://docs.microsoft.com/en-us/windows-hardware/manufacture/desktop/oscdimg-command-line-options>

Example:

```
Oscdimg.exe -bC:\temp\Source-ISO\efi\microsoft\bootEfiSys.bin -pEF -u1 -udfver102 C:\temp\Source-ISO  
C:\temp\Updated-Server2019.iso
```

Install and Configure DHCP Server Feature

Procedure 1. Install and Configure the DHCP Server feature

Step 1. Run the following:

```
Install-WindowsFeature -Name DHCP -IncludeManagementTools  
netsh dhcp add securitygroups  
Restart-Service dhcpserver  
  
Add-DhcpServerv4Scope -name "HCI-Lab-P09-100.101.124.0" -StartRange 100.101.124.221 -EndRange 100.101.124.249  
-SubnetMask 255.255.255.0 -State Active  
  
Set-DhcpServerv4OptionValue -OptionID 3 -Value 100.101.124.1 -ScopeID 100.101.124.0  
Set-DhcpServerv4OptionValue -OptionID 4 -Value 10.10.240.20 -ScopeID 100.101.124.0  
Set-DhcpServerv4OptionValue -OptionID 42 -Value 10.10.240.20 -ScopeID 100.101.124.0  
Set-DhcpServerv4OptionValue -OptionID 6 -Value 110.10.240.23 -ScopeID 100.101.124.0  
  
Get-DhcpServerv4Scope -ScopeId 100.101.124.0  
Get-DhcpServerv4OptionValue -ScopeId 100.101.124.0  
  
#ScopeID 60 is required by WDS when DHCP is also running on the same server. ScopeID 60 is added as a DHCP a  
scope option when WDS is configured.  
  
#OptionId 3 (Router)  
#OptionId 4 (Time Server)
```

```
#OptionId 42 (NTP Server)
#OptionId 6 (DNS Server)

#Verify DHCP Scope
Get-DhcpServerv4Scope -ScopeId 100.101.124.0

#Verify DHCP Scope Option
Get-DhcpServerv4OptionValue -ScopeId 100.101.124.0
```

Configure Bitlocker for System Volume

Using Bitlocker to encrypt system volume is an optional procedure in the deployment. TPM will be the primary key protector for the encrypted volume. The TPM will automatically decrypt the system volume at boot time. A recovery password will be an additional key protector in case the TPM fails. The recovery password will be backed up and stored in Active Directory Domain Service.

Procedure 1. Verify that Secure Boot is Enabled

Step 1. Run the following:

```
$nodes = ("AzS-HCI-Host01", "AzS-HCI-Host02", "AzS-HCI-Host03", "AzS-HCI-Host04")
foreach ($node in $nodes) {
Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME
Write-Host "Checking Secure Boot Status " -ForegroundColor Yellow
Confirm-SecureBootUEFI
}
}
```

```
ASHC-HOST01
True
ASHC-HOST02
True
ASHC-HOST03
True
ASHC-HOST04
True
```

Procedure 2. Enable Bitlocker Group Policy

Note: A local group policy needs to be enabled. This local group policy allows the Recovery Password to be backed up to Active Directory Domain Service.

Step 1. Run the following:

```
$nodes = ("AzS-HCI-Host01", "AzS-HCI-Host02", "AzS-HCI-Host03", "AzS-HCI-Host04")
foreach ($node in $nodes) {
Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green
}
```



```

Write-Host "Configuring BitLocker Registry settings to allow recovery password backup to AD..." -
ForegroundColor Yellow

New-Item -Path HKLM:\SOFTWARE\Policies\Microsoft -Name FVE
New-ItemProperty -Path HKLM:\SOFTWARE\Policies\Microsoft\FVE -Name "OSRecovery" -Value "1" -PropertyType
DWORD
New-ItemProperty -Path HKLM:\SOFTWARE\Policies\Microsoft\FVE -Name "OSManageDRA" -Value "1" -PropertyType
DWORD
New-ItemProperty -Path HKLM:\SOFTWARE\Policies\Microsoft\FVE -Name "OSRecoveryPassword" -Value "2" -
PropertyType DWORD
New-ItemProperty -Path HKLM:\SOFTWARE\Policies\Microsoft\FVE -Name "OSRecoveryKey" -Value "2" -PropertyType
DWORD
New-ItemProperty -Path HKLM:\SOFTWARE\Policies\Microsoft\FVE -Name "OSHideRecoveryPage" -Value "0" -
PropertyType DWORD
New-ItemProperty -Path HKLM:\SOFTWARE\Policies\Microsoft\FVE -Name "OSActiveDirectoryBackup" -Value "1" -
PropertyType DWORD
New-ItemProperty -Path HKLM:\SOFTWARE\Policies\Microsoft\FVE -Name "OSActiveDirectoryInfoToStore" -Value "1" -
PropertyType DWORD
New-ItemProperty -Path HKLM:\SOFTWARE\Policies\Microsoft\FVE -Name "OSRequireActiveDirectoryBackup" -
Value "0" -PropertyType DWORD

Get-ItemProperty -path HKLM:\SOFTWARE\Policies\Microsoft\FVE -Name "OSRecovery"
Get-ItemProperty -path HKLM:\SOFTWARE\Policies\Microsoft\FVE -Name "OSManageDRA"
Get-ItemProperty -path HKLM:\SOFTWARE\Policies\Microsoft\FVE -Name "OSRecoveryPassword"
Get-ItemProperty -path HKLM:\SOFTWARE\Policies\Microsoft\FVE -Name "OSHideRecoveryPage"
Get-ItemProperty -path HKLM:\SOFTWARE\Policies\Microsoft\FVE -Name "OSActiveDirectoryBackup"
Get-ItemProperty -path HKLM:\SOFTWARE\Policies\Microsoft\FVE -Name "OSActiveDirectoryInfoToStore"
Get-ItemProperty -path HKLM:\SOFTWARE\Policies\Microsoft\FVE -Name "OSRequireActiveDirectoryBackup"

}
}

```

Procedure 3. Create and Backup Recovery Password

Note: Create the recover password key protector and back it up to Active Directory Domain Service.

Step 1. Run the following:

```

$nodes = ("AzS-HCI-Host01", "AzS-HCI-Host02", "AzS-HCI-Host03", "AzS-HCI-Host04")
foreach ($node in $nodes) {
Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Creating and Backing Up Recovery Key " -ForegroundColor Yellow
Add-BitLockerKeyProtector -MountPoint "C:" -RecoveryPasswordProtector
$KPID = ((Get-BitLockerVolume -MountPoint "C:").KeyProtector | ? KeyProtectorType -EQ
"RecoveryPassword").KeyProtectorId
Backup-BitLockerKeyProtector -MountPoint "C:" -KeyProtectorId $KPID
}
}

```

```
}  
}
```

ASHC-HOST01

WARNING: ACTIONS REQUIRED:

1. Save this numerical recovery password in a secure location away from your computer:

386936-198451-447381-250371-077506-360635-041558-111023

To prevent data loss, save this password immediately. This password helps ensure that you can unlock the encrypted volume.

```
PSComputerName      : ASHC-Host01  
RunspaceId         : 20454839-0ab4-4ba5-bb9d-6f9e85816510  
ComputerName       : ASHC-HOST01  
MountPoint         : C:  
EncryptionMethod   : None  
AutoUnlockEnabled  :  
AutoUnlockKeyStored : False  
MetadataVersion    : 2  
VolumeStatus       : FullyDecrypted  
ProtectionStatus   : Off  
LockStatus         : Unlocked  
EncryptionPercentage : 0  
WipePercentage     : 0  
VolumeType         : OperatingSystem  
CapacityGB         : 892.5361  
KeyProtector       : {RecoveryPassword}
```

```
PSComputerName      : ASHC-Host01  
RunspaceId         : 14ff4b7e-2985-4f5a-bedd-3376ed3e2df3  
ComputerName       : ASHC-HOST01  
MountPoint         : C:  
EncryptionMethod   : None  
AutoUnlockEnabled  :  
AutoUnlockKeyStored : False  
MetadataVersion    : 2  
VolumeStatus       : FullyDecrypted  
ProtectionStatus   : Off  
LockStatus         : Unlocked  
EncryptionPercentage : 0  
WipePercentage     : 0  
VolumeType         : OperatingSystem  
CapacityGB         : 892.5361  
KeyProtector       : {RecoveryPassword}
```

```

ASHC-HOST02
WARNING: ACTIONS REQUIRED:

1. Save this numerical recovery password in a secure location away from your computer:

608102-507760-408606-144562-351076-363583-605825-128865

To prevent data loss, save this password immediately. This password helps ensure that you can unlock the encrypted volume.
PSCoMputerName      : ASHC-Host02
RunspaceId          : da9b0452-4ae8-4de4-966f-449f0204f627
ComputerName        : ASHC-HOST02
MountPoint           : C:
EncryptionMethod     : None
AutoUnlockEnabled   :
AutoUnlockKeyStored : False
MetadataVersion     : 2
VolumeStatus        : FullyDecrypted
ProtectionStatus     : Off
LockStatus           : Unlocked
EncryptionPercentage : 0
WipePercentage       : 0
VolumeType           : OperatingSystem
CapacityGB           : 892.5361
KeyProtector         : {RecoveryPassword}

PSCoMputerName      : ASHC-Host02
RunspaceId          : 14012078-b552-4ddd-96be-670d6134c74a
ComputerName        : ASHC-HOST02
MountPoint           : C:
EncryptionMethod     : None
AutoUnlockEnabled   :
AutoUnlockKeyStored : False
MetadataVersion     : 2
VolumeStatus        : FullyDecrypted
ProtectionStatus     : Off
LockStatus           : Unlocked
EncryptionPercentage : 0
WipePercentage       : 0
VolumeType           : OperatingSystem
CapacityGB           : 892.5361
KeyProtector         : {RecoveryPassword}

```

```

ASHC-HOST03
WARNING: ACTIONS REQUIRED:

1. Save this numerical recovery password in a secure location away from your computer:

080333-463199-580701-488554-263890-068981-212509-627242

To prevent data loss, save this password immediately. This password helps ensure that you can unlock the encrypted volume.
PSCoMputerName      : ASHC-Host03
RunspaceId          : 9b0947ad-0b5a-4618-a85a-751910dba6a8
ComputerName        : ASHC-HOST03
MountPoint           : C:
EncryptionMethod     : None
AutoUnlockEnabled   :
AutoUnlockKeyStored : False
MetadataVersion     : 2
VolumeStatus        : FullyDecrypted
ProtectionStatus     : Off
LockStatus           : Unlocked
EncryptionPercentage : 0
WipePercentage       : 0
VolumeType           : OperatingSystem
CapacityGB           : 892.5361
KeyProtector         : {RecoveryPassword}

PSCoMputerName      : ASHC-Host03
RunspaceId          : 310315a2-2c06-4e2d-8a3c-750592be10df
ComputerName        : ASHC-HOST03
MountPoint           : C:
EncryptionMethod     : None
AutoUnlockEnabled   :
AutoUnlockKeyStored : False
MetadataVersion     : 2
VolumeStatus        : FullyDecrypted
ProtectionStatus     : Off
LockStatus           : Unlocked
EncryptionPercentage : 0
WipePercentage       : 0
VolumeType           : OperatingSystem
CapacityGB           : 892.5361
KeyProtector         : {RecoveryPassword}

```

```
ASHC-HOST04
WARNING: ACTIONS REQUIRED:

1. Save this numerical recovery password in a secure location away from your computer:

174284-621027-461373-145277-225137-356125-272382-289047

To prevent data loss, save this password immediately. This password helps ensure that you can unlock the encrypted volume.
PSComputerName      : ASHC-Host04
RunspaceId          : d0819314-ef69-41f7-bb20-c0dae456d799
ComputerName        : ASHC-HOST04
MountPoint           : C:
EncryptionMethod     : None
AutoUnlockEnabled   :
AutoUnlockKeyStored : False
MetadataVersion     : 2
VolumeStatus        : FullyDecrypted
ProtectionStatus     : Off
LockStatus           : Unlocked
EncryptionPercentage : 0
WipePercentage       : 0
VolumeType           : OperatingSystem
CapacityGB           : 892.5361
KeyProtector         : {RecoveryPassword}

PSComputerName      : ASHC-Host04
RunspaceId          : 209a2743-3877-423a-b431-137c6639bd12
ComputerName        : ASHC-HOST04
MountPoint           : C:
EncryptionMethod     : None
AutoUnlockEnabled   :
AutoUnlockKeyStored : False
MetadataVersion     : 2
VolumeStatus        : FullyDecrypted
ProtectionStatus     : Off
LockStatus           : Unlocked
EncryptionPercentage : 0
WipePercentage       : 0
VolumeType           : OperatingSystem
CapacityGB           : 892.5361
KeyProtector         : {RecoveryPassword}
```

Procedure 4. Enable Bitlocker

Note: Enable Bitlocker for the system volume and add the TPM protector. Encryption of the system volume will not start until the server is rebooted.

Step 1. Run the following:

```
$nodes = ("AzS-HCI-Host01", "AzS-HCI-Host02", "AzS-HCI-Host03", "AzS-HCI-Host04")
foreach ($node in $nodes) {
Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Enabling Bitlocker Protection " -ForegroundColor Yellow
Enable-BitLocker -MountPoint "C:" -EncryptionMethod XtsAes256 -UsedSpaceOnly -TpmProtector

}
}
```

```
ASHC-HOST01
WARNING: ACTIONS REQUIRED:

1. Save this numerical recovery password in a secure location away from your computer:
386936-198451-447381-250371-077506-360635-041558-111023

To prevent data loss, save this password immediately. This password helps ensure that you can unlock the encrypted volume.
2. Restart the computer to run a hardware test.
(Type: get-help Restart-Computer for command line instructions.)

PSComputerName      : ASHC-Host01
RunspaceId          : 7b263445-531b-4461-a1f2-ab75dacd240a
ComputerName        : ASHC-HOST01
MountPoint          : C:
EncryptionMethod    : XtsAes256
AutoUnlockEnabled   : 
AutoUnlockKeyStored : False
MetadataVersion     : 2
VolumeStatus        : FullyDecrypted
ProtectionStatus    : Off
LockStatus          : Unlocked
EncryptionPercentage : 0
WipePercentage      : 0
VolumeType          : OperatingSystem
CapacityGB          : 892.5361
KeyProtector        : {RecoveryPassword, Tpm}
```

```
ASHC-HOST02
WARNING: ACTIONS REQUIRED:

1. Save this numerical recovery password in a secure location away from your computer:
608102-507760-408606-144562-351076-363583-605825-128865

To prevent data loss, save this password immediately. This password helps ensure that you can unlock the encrypted volume.
2. Restart the computer to run a hardware test.
(Type: get-help Restart-Computer for command line instructions.)

PSComputerName      : ASHC-Host02
RunspaceId          : 3329db3b-1cd2-4dbc-ba5a-275114b31f11
ComputerName        : ASHC-HOST02
MountPoint          : C:
EncryptionMethod    : XtsAes256
AutoUnlockEnabled   : 
AutoUnlockKeyStored : False
MetadataVersion     : 2
VolumeStatus        : FullyDecrypted
ProtectionStatus    : Off
LockStatus          : Unlocked
EncryptionPercentage : 0
WipePercentage      : 0
VolumeType          : OperatingSystem
CapacityGB          : 892.5361
KeyProtector        : {RecoveryPassword, Tpm}
```

```
ASHC-HOST03
WARNING: ACTIONS REQUIRED:

1. Save this numerical recovery password in a secure location away from your computer:
080333-463199-580701-488554-263890-068981-212509-627242

To prevent data loss, save this password immediately. This password helps ensure that you can unlock the encrypted volume.
2. Restart the computer to run a hardware test.
(Type: get-help Restart-Computer for command line instructions.)

PSComputerName      : ASHC-Host03
RunspaceId          : c8779363-e539-408f-a186-a634332d0bb6
ComputerName        : ASHC-HOST03
MountPoint          : C:
EncryptionMethod    : XtsAes256
AutoUnlockEnabled   : 
AutoUnlockKeyStored : False
MetadataVersion     : 2
VolumeStatus        : FullyDecrypted
ProtectionStatus    : Off
LockStatus          : Unlocked
EncryptionPercentage : 0
WipePercentage      : 0
VolumeType          : OperatingSystem
CapacityGB          : 892.5361
KeyProtector        : {RecoveryPassword, Tpm}
```

```

ASHC-HOST04
WARNING: ACTIONS REQUIRED:

1. Save this numerical recovery password in a secure location away from your computer:

174284-621027-461373-145277-225137-356125-272382-289047

To prevent data loss, save this password immediately. This password helps ensure that you can unlock the encrypted volume.
2. Restart the computer to run a hardware test.
   (Type: get-help Restart-Computer for command line instructions.)
PSCoComputerName      : ASHC-Host04
RunspaceId           : ea816db5-8f29-4121-88bc-22baa59f00e3
ComputerName         : ASHC-HOST04
MountPoint           : C:
EncryptionMethod     : XtsAes256
AutoUnlockEnabled    :
AutoUnlockKeyStored  : False
MetadataVersion      : 2
VolumeStatus         : FullyDecrypted
ProtectionStatus     : Off
LockStatus           : Unlocked
EncryptionPercentage : 0
WipePercentage       : 0
VolumeType           : OperatingSystem
CapacityGB           : 892.5361
KeyProtector         : {RecoveryPassword, Tpm}

```

Procedure 5. Reboot Server to Enable Bitlocker for the System Volume

Note: Bitlocker will enable the when the server reboots. Bitlocker verifies that the key protectors are correctly configure. Volume encryption will take a few minutes to complete after the server reboots.

Step 1. Run the following:

```

$nodes = ("AzS-HCI-Host01", "AzS-HCI-Host02", "AzS-HCI-Host03", "AzS-HCI-Host04")
foreach ($node in $nodes) {
Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Restarting Host After Enabling Bitlocker Protection " -ForegroundColor Yellow

Restart-Computer -Force

}
}

```

Procedure 6. Verify Bitlocker Status

Note: Verify the Bitlocker Protection Status and Encryption Percentage.

Step 1. Run the following:

```

$nodes = ("AzS-HCI-Host01", "AzS-HCI-Host02", "AzS-HCI-Host03", "AzS-HCI-Host04")
foreach ($node in $nodes) {
Invoke-Command $node -Credential $Creds -scriptblock {
write-host "Host Name:" $env:COMPUTERNAME -ForegroundColor Green

Write-Host "Verifying Bitlocker Protection Status " -ForegroundColor Yellow

Get-BitlockerVolume -MountPoint "c:" | FT
}
}

```

```
}  
}
```

Note: Bitlocker protection status will indicate “Off” until encryption reaches 100%.

Procedure 7. Verity Bitlocker Recovery Password Backup

Note: Bitlocker Recovery Password View provides the ability to read the backup of the recovery password that that is backed up to Active Directory Domain Services. This is an optional Windows feature that can be installed by running the following PowerShell command on a system that will be sued to read the password from Active Directly Domain Services.

Step 1. Add-WindowsFeature -Name RSAT-Feature-Tools-BitLocker-BdeAduExt

The following PowerShell scriptblock retrieves the Bitlocker password that is backed up to Active Directory:

```
$nodes = ("AzS-HCI-Host01", "AzS-HCI-Host02", "AzS-HCI-Host03", "AzS-HCI-Host04")  
foreach ($node in $nodes) {  
  
$objComputer = Get-ADComputer $node  
write-host "Host Name:" $node -ForegroundColor Green  
  
$Bitlocker_Objects = Get-ADObject -Filter {objectclass -eq 'msFVE-RecoveryInformation'} -SearchBase  
$objComputer.DistinguishedName -Properties 'msFVE-RecoveryPassword'  
  
foreach ($Bitlocker_Object in $Bitlocker_Objects) {  
Write-Host "Date, Time, and Password ID:" ($Bitlocker_Objects).Name  
Write-Host "Recovery Password:" ($Bitlocker_Objects).'msFVE-RecoveryPassword' -ForegroundColor Cyan -  
Separator " " "  
Write-Host ""  
  
}  
}
```

Details on accessing the recovery password backup can be found here: <https://learn.microsoft.com/en-us/windows/security/operating-system-security/data-protection/bitlocker/recovery-process>. Recovery passwords backup should be verified as part of every deployment:

Organizations using Bitlocker should be familiar with Bitlocker recovery procedures in case recovering access to a Bitlocker protected volume is required. The Microsoft guide to Bitlocker recovery can be found here: <https://docs.microsoft.com/en-us/windows/security/information-protection/bitlocker/bitlocker-recovery-guide-plan>

ToR Switch Configuration Example

This section describes the ToR (Cisco Nexus 9336C-FX2) switches example configuration used for the deployment of this Azure Stack HCI solution. The Cisco Nexus switch configuration explains the basic L2 and L3 functionality and QoS configuration for the Azure Stack HCI solution environment used in the validation environment. The gateways required for this solution are hosted by the pair of Cisco Nexus switches, but the primary routing is passed onto an existing router that is upstream of the converged infrastructure. This upstream router will need

to be aware of any networks created on the Cisco Nexus switches, but configuration of an upstream router is beyond the scope of this deployment guide.

Check NXOS Version

For Azure Stack HCI solution, the supported NXOS version is 10.3(2)F or later and the supported Cisco Nexus switches are listed here: <https://learn.microsoft.com/en-us/azure-stack/hci/concepts/physical-network-requirements?tabs=Cisco%2C22H2reqs>

ToR-A	ToR-B
<pre>show version include 'NXOS Chassis' NXOS: version 10.3(2) [Feature Release] NXOS image file is: bootflash:///nxos64- cs.10.3.2.F.bin cisco Nexus9000 C9336C-FX2 Chassis</pre>	<pre>show version include 'NXOS Chassis' NXOS: version 10.3(2) [Feature Release] NXOS image file is: bootflash:///nxos64- cs.10.3.2.F.bin cisco Nexus9000 C9336C-FX2 Chassis</pre>

Enable Features

Some of the key NX-OS features implemented for this solution are:

- Feature interface-vlan – Allows for VLAN IP interfaces to be configured within the switch as gateways.
- Feature HSRP – Allows for Hot Standby Routing Protocol configuration for high availability.
- Feature LACP – Allows for the utilization of Link Aggregation Control Protocol (802.3ad) by the port channels configured on the switch.
- Feature vPC – Virtual Port-Channel (vPC) presents the two Nexus switches as a single “logical” port channel to the connecting upstream or downstream device.
- Feature LLDP – Link Layer Discovery Protocol (LLDP), a vendor-neutral device discovery protocol, allows the discovery of both Cisco devices and devices from other sources.
- Feature NX-API – NX-API improves the accessibility of CLI by making it available outside of the switch by using HTTP/HTTPS. This feature helps with configuring the Cisco Nexus switch remotely using the automation framework.
- Feature UDLD – Enables unidirectional link detection for various interfaces.
- Feature DHCP – Allows for the configuration of DHCP relay agent, DHCP snooping, or any of the features that depend on DHCP
- Feature scp-server – Enables the SCP (Secure Copy) server on the Cisco NX-OS device in order to copy files and from a remote device.
- Feature bgp (optional) – enables bgp to be used between ToR and aggregation router

ToR-A	ToR-B
<pre>feature nxapi feature scp-server</pre>	<pre>feature nxapi feature scp-server</pre>

ToR-A	ToR-B
<code>cfs eth distribute</code>	<code>cfs eth distribute</code>
<code>feature bgp</code>	<code>feature bgp</code>
<code>feature udld</code>	<code>feature udld</code>
<code>feature interface-vlan</code>	<code>feature interface-vlan</code>
<code>feature hsrp</code>	<code>feature hsrp</code>
<code>feature lacp</code>	<code>feature lacp</code>
<code>feature dhcp</code>	<code>feature dhcp</code>
<code>feature vpc</code>	<code>feature vpc</code>
<code>feature lldp</code>	<code>feature lldp</code>

Configure VLANs

The table below provides the VLANs created for different traffics used in this solution.

ToR-A	ToR-B
<code>vlan 2</code> <code> name Reserved_Port_Ethernet</code>	<code>vlan 2</code> <code> name Reserved_Port_Ethernet</code>
<code>vlan 101</code> <code> name Tenant</code>	<code>vlan 101</code> <code> name Tenant</code>
<code>vlan 107</code> <code> name StorageA</code>	<code>vlan 207</code> <code> name StorageB</code>
<code>vlan 126</code> <code> name Management</code>	<code>vlan 126</code> <code> name Management</code>

Create Hot Standby Router Protocol (HSRP) Switched Virtual Interfaces (SVI)

These interfaces can be considered optional if the subnets of the VLANs used within the environment are managed entirely by an upstream switch, but if that is the case, all managed VLANs will need to be carried up through the vPC to the Upstream switches.

Routing between the SVIs is directly connected between them as they reside in the same Virtual Routing and Forwarding in-instance (VRF), and traffic set to enter and exit the VRF will traverse the default gateway set for the switches.

ToR-A	ToR-B
<code>interface Vlan2</code> <code> description Unused_Ports</code> <code> mtu 9216</code>	<code>interface Vlan2</code> <code> description Unused_Ports</code> <code> mtu 9216</code>
<code>interface Vlan101</code> <code> description Tenant network</code>	<code>interface Vlan101</code> <code> description Tenant network</code>

ToR-A	ToR-B
<pre> no shutdown mtu 9216 no ip redirects ip address 192.168.101.2/24 ip directed-broadcast no ipv6 redirects hsrp version 2 hsrp 101 priority 150 forwarding-threshold lower 1 upper 150 ip 192.168.101.1 ip dhcp relay address 192.168.0.10 </pre>	<pre> no shutdown mtu 9216 no ip redirects ip address 192.168.101.3/24 ip directed-broadcast no ipv6 redirects hsrp version 2 hsrp 101 priority 140 forwarding-threshold lower 1 upper 140 ip 192.168.101.1 ip dhcp relay address 192.168.0.10 </pre>
<pre> interface Vlan126 description Management Network no shutdown mtu 9216 no ip redirects ip address 192.168.126.2/26 ip directed-broadcast no ipv6 redirects hsrp version 2 hsrp 126 priority 150 forwarding-threshold lower 1 upper 150 ip 192.168.126.1 ip dhcp relay address 192.168.0.10 </pre>	<pre> interface Vlan126 description Management Network no shutdown mtu 9216 no ip redirects ip address 192.168.126.3/26 ip directed-broadcast no ipv6 redirects hsrp version 2 hsrp 126 priority 140 forwarding-threshold lower 1 upper 140 ip 192.168.126.1 ip dhcp relay address 192.168.0.10 </pre>
<pre> interface Vlan200 description P2P_TOR1-ibgp-1_To_TOR2-ibgp- 1 no shutdown mtu 9216 no ip redirects ip address 192.168.200.45/30 no ipv6 redirects </pre>	<pre> interface Vlan200 description P2P_TOR1-ibgp-1_To_TOR2-ibgp- 1 no shutdown mtu 9216 no ip redirects ip address 192.168.200.46/30 no ipv6 redirects </pre>

Create the vPC Domain

Create a vPC domain ID with a unique number (from 1 -1000) and configure the role priority and peer-keepalive on both the switches. The vPC domain ID is same on both switches and this will handle the vPC settings specified within the switches. Note that the commands for role priority value and peer-keepalive varies slightly on both switches.

ToR-A	ToR-B
<pre>interface mgmt0 vrf member management ip address 192.168.0.24/24 vpc domain 120 peer-switch role priority 10 peer-keepalive destination 192.168.0.25 source 192.168.0.24 delay restore 150 peer-gateway layer3 peer-router auto-recovery</pre>	<pre>interface mgmt0 vrf member management ip address 192.168.0.25/24 vpc domain 120 peer-switch role priority 20 peer-keepalive destination 192.168.0.24 source 192.168.0.25 delay restore 150 peer-gateway layer3 peer-router auto-recovery</pre>

On each switch, configure the Port Channel member interfaces that will be part of the vPC Peer Link and configure the vPC Peer Link:

ToR-A	ToR-B
<pre>interface port-channel10 description vPC Peer-Link • switchport mode trunk switchport trunk allowed vlan 101,107,126,200,207 spanning-tree port type network service-policy type qos input AzS_HCI_QoS vpc peer-link</pre>	<pre>interface port-channel10 description vPC Peer-Link switchport mode trunk switchport trunk allowed vlan 101,107,126,200,207 spanning-tree port type network service-policy type qos input AzS_HCI_QoS vpc peer-link</pre>

QoS Configuration on ToR Switches (Cisco Nexus 9300 series switches)

This procedure explains the QoS configuration example for supporting RoCE (RDMA over Converged Ethernet) traffic on the ToR switches.

Using Cisco Modular Quality of Service Command Line Interface (MQC), you can define and configure QoS policies by following these steps:

1. Define a particular class of traffic.
2. After creating class-map, we put them in to a policy-map, where we mark (using bandwidth, policing, shaping, and so on) the traffic.
3. Use a service-policy command to apply that p-map to an interface in inbound or outbound direction.

Note: The QoS configuration in the host OS should match the QoS configuration performed in the Network switch (ToR) configuration

Create class-map type QoS and match based on CoS Value

In the following example, RDMA (for storage traffic) and CLUSTER-COMM (for cluster heartbeat traffic) traffic classes are defined and matched with layer 2 CoS 4 and CoS 5 respectively for classification.

ToR-A	ToR-B
<pre>class-map type qos match-all RDMA match cos 4 class-map type qos match-all CLUSTER-COMM match cos 5</pre>	<pre>class-map type qos match-all RDMA match cos 4 class-map type qos match-all CLUSTER-COMM match cos 5</pre>

Create policy-map type QoS and Set qos-group and add/or Policing Rule

A policy-map named AzS_HCI_QoS is created and referenced to RDMA and CLUSTER-COMM class-maps and set the qos-group accordingly as shown in the following example.

ToR-A	ToR-B
<pre>policy-map type qos AzS_HCI_QoS class RDMA set qos-group 4 class CLUSTER-COMM set qos-group 5</pre>	<pre>policy-map type qos AzS_HCI_QoS class RDMA set qos-group 4 class CLUSTER-COMM set qos-group 5</pre>

Attach policy-map type QoS as input to an Interface

The policy-map created in the previous step is now applied to interfaces port-channel 10 and interfaces ethernet 1/1-4, where all Azure Stack HCI cluster nodes are connected.

ToR-A	ToR-B
<pre>interface port-channel10 service-policy type qos input AzS_HCI_QoS interface Ethernet1/1 service-policy type qos input AzS_HCI_QoS</pre>	<pre>interface port-channel10 service-policy type qos input AzS_HCI_QoS interface Ethernet1/1 service-policy type qos input AzS_HCI_QoS</pre>

ToR-A	ToR-B
<pre>interface Ethernet1/2 service-policy type qos input AzS_HCI_QoS</pre>	<pre>interface Ethernet1/2 service-policy type qos input AzS_HCI_QoS</pre>
<pre>interface Ethernet1/3 service-policy type qos input AzS_HCI_QoS</pre>	<pre>interface Ethernet1/3 service-policy type qos input AzS_HCI_QoS</pre>
<pre>interface Ethernet1/4 service-policy type qos input AzS_HCI_QoS</pre>	<pre>interface Ethernet1/4 service-policy type qos input AzS_HCI_QoS</pre>

Create class-map type network-qos and match based on qos-group Value

The network QoS policy defines the characteristics of QoS properties network wide.

Two class-map type network-qos named RDMA_CL_Map_NetQoS and Cluster-Comm_CL_Map_NetQoS are created and matched with qos-group 4 and qos-group 5, respectively.

ToR-A	ToR-B
<pre>class-map type network-qos RDMA_CL_Map_NetQos match qos-group 4</pre>	<pre>class-map type network-qos RDMA_CL_Map_NetQos match qos-group 4</pre>
<pre>class-map type network-qos Cluster- Comm_CL_Map_NetQos match qos-group 5</pre>	<pre>class-map type network-qos Cluster- Comm_CL_Map_NetQos match qos-group 5</pre>

Create policy-map type network-qos and Define Actions

In this example, the QoS network policy created to set Jumbo MTU for both traffic classes and no-drop (pause) to only RoCE traffic. During congestion, PFC sends a pause frame that indicates which CoS values needs to be paused. This network-qos policy is then applied to the system.

ToR-A	ToR-B
<pre>policy-map type network-qos QOS_NETWORK class type network-qos RDMA_CL_Map_NetQos pause pfc-cos 4 mtu 9216 class type network-qos Cluster- Comm_CL_Map_NetQos mtu 9216 class type network-qos class-default mtu 9216</pre>	<pre>policy-map type network-qos QOS_NETWORK class type network-qos RDMA_CL_Map_NetQos pause pfc-cos 4 mtu 9216 class type network-qos Cluster- Comm_CL_Map_NetQos mtu 9216 class type network-qos class-default mtu 9216</pre>

ToR-A	ToR-B
<pre>system qos service-policy type network-qos QOS_NETWORK</pre>	<pre>system qos service-policy type network-qos QOS_NETWORK</pre>

Note: For the drop and no drop configuration, you also need to enable PFC per port.

Create policy-map type queuing referencing with the system-defined class-map type queuing and create actions

A policy map with minimum bandwidth percentage guarantee is specified to traffic class in periods of congestion - 50% is allocated to RDMA (storage) traffic, 49% is allocated to management and compute traffic and 1% is allocated to cluster heartbeat traffic.

Weighted random early detection (WRED) with minimum and maximum thresholds is also set to drop packets when the configured thresholds are exceeded. WRED configured with ECN (explicit congestion notification) marks packets instead of dropping them when the average queue length exceeds a specific threshold value. With WRED ECN feature, end hosts use this marking as signal that the network is congested to slow down sending packets.

ToR-A	ToR-B
<pre>policy-map type queuing QOS_EGRESS_PORT class type queuing c-out-8q-q-default bandwidth remaining percent 49 class type queuing c-out-8q-q1 bandwidth remaining percent 0 class type queuing c-out-8q-q2 bandwidth remaining percent 0 class type queuing c-out-8q-q3 bandwidth remaining percent 0 class type queuing c-out-8q-q4 bandwidth remaining percent 50 random-detect minimum-threshold 300 kbytes maximum-threshold 300 kbytes drop-probability 100 weight 0 ecn class type queuing c-out-8q-q5 bandwidth percent 1 class type queuing c-out-8q-q6 bandwidth remaining percent 0 class type queuing c-out-8q-q7 bandwidth remaining percent 0 system qos</pre>	<pre>policy-map type queuing QOS_EGRESS_PORT class type queuing c-out-8q-q-default bandwidth remaining percent 49 class type queuing c-out-8q-q1 bandwidth remaining percent 0 class type queuing c-out-8q-q2 bandwidth remaining percent 0 class type queuing c-out-8q-q3 bandwidth remaining percent 0 class type queuing c-out-8q-q4 bandwidth remaining percent 50 random-detect minimum-threshold 300 kbytes maximum-threshold 300 kbytes drop-probability 100 weight 0 ecn class type queuing c-out-8q-q5 bandwidth percent 1 class type queuing c-out-8q-q6 bandwidth remaining percent 0 class type queuing c-out-8q-q7 bandwidth remaining percent 0 system qos</pre>

ToR-A	ToR-B
<pre>service-policy type queuing output QOS_EGRESS_PORT</pre>	<pre>service-policy type queuing output QOS_EGRESS_PORT</pre>

Attach policy-map queuing to interfaces

The example below shows policy-map queuing and priority-flow-control on are applied to ethernet 1/1-4 interfaces. Azure Stack HCI cluster nodes are connected to these interfaces.

ToR-A	ToR-B
<pre>interface Ethernet1/1 priority-flow-control mode on service-policy type queuing output QOS_EGRESS_PORT</pre>	<pre>interface Ethernet1/1 priority-flow-control mode on service-policy type queuing output QOS_EGRESS_PORT</pre>
<pre>interface Ethernet1/2 priority-flow-control mode on service-policy type queuing output QOS_EGRESS_PORT</pre>	<pre>interface Ethernet1/2 priority-flow-control mode on service-policy type queuing output QOS_EGRESS_PORT</pre>
<pre>interface Ethernet1/3 priority-flow-control mode on service-policy type queuing output QOS_EGRESS_PORT</pre>	<pre>interface Ethernet1/3 priority-flow-control mode on service-policy type queuing output QOS_EGRESS_PORT</pre>
<pre>interface Ethernet1/4 priority-flow-control mode on service-policy type queuing output QOS_EGRESS_PORT</pre>	<pre>interface Ethernet1/4 priority-flow-control mode on service-policy type queuing output QOS_EGRESS_PORT</pre>

The example below shows the full running-configuration of both ToR switches used for this deployment in a fully converged network topology.

ToR-A	ToR-B
<pre>switchname AzS-HCI-ToR1 class-map type network-qos RDMA_CL_Map_NetQos match qos-group 4 class-map type network-qos Cluster- Comm_CL_Map_NetQos match qos-group 5 policy-map type network-qos QOS_NETWORK</pre>	<pre>switchname AzS-HCI-ToR2 class-map type network-qos RDMA_CL_Map_NetQos match qos-group 4 class-map type network-qos Cluster- Comm_CL_Map_NetQos match qos-group 5 policy-map type network-qos QOS_NETWORK</pre>

ToR-A	ToR-B
<pre> class type network-qos RDMA_CL_Map_NetQos pause pfc-cos 4 mtu 9216 class type network-qos Cluster- Comm_CL_Map_NetQos mtu 9216 class type network-qos class-default mtu 9216 vdc AzS-HCI-ToR1 id 1 limit-resource vlan minimum 16 maximum 4094 limit-resource vrf minimum 2 maximum 4097 limit-resource port-channel minimum 0 maximum 511 limit-resource m4route-mem minimum 58 maximum 58 limit-resource m6route-mem minimum 8 maximum 8 feature nxapi feature scp-server cfs eth distribute feature bgp feature udld feature interface-vlan feature hsrp feature lacp feature dhcp feature vpc feature lldp mac address-table aging-time 1510 ip domain-lookup spanning-tree mode mst class-map type qos match-all RDMA match cos 4 class-map type qos match-all CLUSTER-COMM match cos 5 </pre>	<pre> class type network-qos RDMA_CL_Map_NetQos pause pfc-cos 4 mtu 9216 class type network-qos Cluster- Comm_CL_Map_NetQos mtu 9216 class type network-qos class-default mtu 9216 vdc AzS-HCI-ToR2 id 1 limit-resource vlan minimum 16 maximum 4094 limit-resource vrf minimum 2 maximum 4097 limit-resource port-channel minimum 0 maximum 511 limit-resource m4route-mem minimum 58 maximum 58 limit-resource m6route-mem minimum 8 maximum 8 feature nxapi feature scp-server cfs eth distribute feature bgp feature udld feature interface-vlan feature hsrp feature lacp feature dhcp feature vpc feature lldp mac address-table aging-time 1510 ip domain-lookup spanning-tree mode mst class-map type qos match-all RDMA match cos 4 class-map type qos match-all CLUSTER-COMM match cos 5 </pre>

ToR-A	ToR-B
<pre> policy-map type qos AzS_HCI_QoS class RDMA set qos-group 4 class CLUSTER-COMM set qos-group 5 policy-map type queuing QOS_EGRESS_PORT class type queuing c-out-8q-q-default bandwidth remaining percent 49 class type queuing c-out-8q-q1 bandwidth remaining percent 0 class type queuing c-out-8q-q2 bandwidth remaining percent 0 class type queuing c-out-8q-q3 bandwidth remaining percent 0 class type queuing c-out-8q-q4 bandwidth remaining percent 50 random-detect minimum-threshold 300 kbytes maximum-threshold 300 kbytes drop-probability 100 weight 0 ecn class type queuing c-out-8q-q5 bandwidth percent 1 class type queuing c-out-8q-q6 bandwidth remaining percent 0 class type queuing c-out-8q-q7 bandwidth remaining percent 0 system qos service-policy type queuing output QOS_EGRESS_PORT service-policy type network-qos QOS_NETWORK copp profile strict snmp-server user admin network-admin auth md5 0x743ead09954eb506ae83f49f45f2db95 priv des 0x743ead09954eb 506ae83f49f45f2db95 localizedkey rmon event 1 description FATAL(1) owner PMON@FATAL rmon event 2 description CRITICAL(2) owner PMON@CRITICAL </pre>	<pre> policy-map type qos AzS_HCI_QoS class RDMA set qos-group 4 class CLUSTER-COMM set qos-group 5 policy-map type queuing QOS_EGRESS_PORT class type queuing c-out-8q-q-default bandwidth remaining percent 49 class type queuing c-out-8q-q1 bandwidth remaining percent 0 class type queuing c-out-8q-q2 bandwidth remaining percent 0 class type queuing c-out-8q-q3 bandwidth remaining percent 0 class type queuing c-out-8q-q4 bandwidth remaining percent 50 random-detect minimum-threshold 300 kbytes maximum-threshold 300 kbytes drop-probability 100 weight 0 ecn class type queuing c-out-8q-q5 bandwidth percent 1 class type queuing c-out-8q-q6 bandwidth remaining percent 0 class type queuing c-out-8q-q7 bandwidth remaining percent 0 system qos service-policy type queuing output QOS_EGRESS_PORT service-policy type network-qos QOS_NETWORK copp profile strict snmp-server user admin network-admin auth md5 0x4f03854fbf75be4bec6b38ed1223a54d priv des 0x4f03854fbf75b e4bec6b38ed1223a54d localizedkey rmon event 1 description FATAL(1) owner PMON@FATAL rmon event 2 description CRITICAL(2) owner PMON@CRITICAL </pre>

ToR-A	ToR-B
rmon event 3 description ERROR(3) owner PMON@ERROR	rmon event 3 description ERROR(3) owner PMON@ERROR
rmon event 4 description WARNING(4) owner PMON@WARNING	rmon event 4 description WARNING(4) owner PMON@WARNING
rmon event 5 description INFORMATION(5) owner PMON@INFO	rmon event 5 description INFORMATION(5) owner PMON@INFO
ntp server 72.163.32.44 use-vrf management system default switchport	ntp server 72.163.32.44 use-vrf management system default switchport
vlan 1-2,101,107,126,200,207	vlan 1-2,101,107,126,200,207
vlan 2	vlan 2
name Reserved_Port_Ethernet	name Reserved_Port_Ethernet
vlan 101	vlan 101
name Tenant	name Tenant
vlan 107	vlan 107
name StorageA	name StorageA
vlan 126	vlan 126
name Management	name Management
vlan 200	vlan 200
name iBGP-Link	name iBGP-Link
vlan 207	vlan 207
name StorageB	name StorageB
spanning-tree port type edge bpduguard default	spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default	spanning-tree port type edge bpdufilter default
spanning-tree port type network default	spanning-tree port type network default
service dhcp	service dhcp
ip dhcp relay	ip dhcp relay
ipv6 dhcp relay	ipv6 dhcp relay
vrf context management	vrf context management
ip route 0.0.0.0/0 192.168.0.1	ip route 0.0.0.0/0 192.168.0.1
congestion-control random-detect forward-nonecn	congestion-control random-detect forward-nonecn
vpc domain 120	vpc domain 120
peer-switch	peer-switch
role priority 10	role priority 20
peer-keepalive destination 192.168.0.25	peer-keepalive destination 192.168.0.24

ToR-A	ToR-B
<pre> source 192.168.0.24 delay restore 150 peer-gateway layer3 peer-router auto-recovery </pre>	<pre> source 192.168.0.25 delay restore 150 peer-gateway layer3 peer-router auto-recovery </pre>
<pre> interface Vlan1 no ip redirects no ipv6 redirects </pre>	<pre> interface Vlan1 no ip redirects no ipv6 redirects </pre>
<pre> interface Vlan2 description Unused_Ports mtu 9216 </pre>	<pre> interface Vlan2 description Unused_Ports mtu 9216 </pre>
<pre> interface Vlan101 description Tenant network no shutdown mtu 9216 no ip redirects ip address 192.168.101.2/24 ip directed-broadcast no ipv6 redirects hsrp version 2 hsrp 101 priority 150 forwarding-threshold lower 1 upper 150 ip 192.168.101.1 ip dhcp relay address 192.168.0.10 </pre>	<pre> interface Vlan101 description Tenant network no shutdown mtu 9216 no ip redirects ip address 192.168.101.3/24 ip directed-broadcast no ipv6 redirects hsrp version 2 hsrp 101 priority 140 forwarding-threshold lower 1 upper 140 ip 192.168.101.1 ip dhcp relay address 192.168.0.10 </pre>
<pre> interface Vlan126 description Management Network no shutdown mtu 9216 no ip redirects ip address 192.168.126.2/26 ip directed-broadcast no ipv6 redirects </pre>	<pre> interface Vlan126 description Management Network no shutdown mtu 9216 no ip redirects ip address 192.168.126.3/26 ip directed-broadcast no ipv6 redirects </pre>

ToR-A	ToR-B
<pre> hsrp version 2 hsrp 126 priority 150 forwarding-threshold lower 1 upper 150 ip 192.168.126.1 ip dhcp relay address 192.168.0.10 interface Vlan200 description P2P_TOR1-ibgp-1_To_TOR2-ibgp- 1 no shutdown mtu 9216 no ip redirects ip address 192.168.200.45/30 no ipv6 redirects interface port-channel10 description vPC Peer-Link switchport mode trunk switchport trunk allowed vlan 101,107,126,200,207 spanning-tree port type network service-policy type qos input AzS_HCI_QoS vpc peer-link interface Ethernet1/1 description AzS-HCI Fabric-A NIC Port switchport mode trunk switchport trunk native vlan 126 switchport trunk allowed vlan 101,107,126 priority-flow-control mode on spanning-tree port type edge trunk mtu 9216 service-policy type qos input AzS_HCI_QoS service-policy type queuing output QOS_EGRESS_PORT no shutdown interface Ethernet1/2 </pre>	<pre> hsrp version 2 hsrp 126 priority 140 forwarding-threshold lower 1 upper 140 ip 192.168.126.1 ip dhcp relay address 192.168.0.10 interface Vlan200 description P2P_TOR1-ibgp-1_To_TOR2-ibgp- 1 no shutdown mtu 9216 no ip redirects ip address 192.168.200.46/30 no ipv6 redirects interface port-channel10 description vPC Peer-Link switchport mode trunk switchport trunk allowed vlan 101,107,126,200,207 spanning-tree port type network service-policy type qos input AzS_HCI_QoS vpc peer-link interface Ethernet1/1 description AzS-HCI Fabric-B NIC Port switchport mode trunk switchport trunk native vlan 126 switchport trunk allowed vlan 101,126,207 priority-flow-control mode on spanning-tree port type edge trunk mtu 9216 service-policy type qos input AzS_HCI_QoS service-policy type queuing output QOS_EGRESS_PORT no shutdown interface Ethernet1/2 </pre>

ToR-A	ToR-B
<pre> description AzS-HCI Fabric-A NIC Port switchport mode trunk switchport trunk native vlan 126 switchport trunk allowed vlan 101,107,126 priority-flow-control mode on spanning-tree port type edge trunk mtu 9216 service-policy type qos input AzS_HCI_QoS service-policy type queuing output QOS_EGRESS_PORT no shutdown </pre>	<pre> description AzS-HCI Fabric-B NIC Port switchport mode trunk switchport trunk native vlan 126 switchport trunk allowed vlan 101,126,207 priority-flow-control mode on spanning-tree port type edge trunk mtu 9216 service-policy type qos input AzS_HCI_QoS service-policy type queuing output QOS_EGRESS_PORT no shutdown </pre>
<pre> interface Ethernet1/3 description AzS-HCI Fabric-A NIC Port switchport mode trunk switchport trunk native vlan 126 switchport trunk allowed vlan 101,107,126 priority-flow-control mode on spanning-tree port type edge trunk mtu 9216 service-policy type qos input AzS_HCI_QoS service-policy type queuing output QOS_EGRESS_PORT no shutdown </pre>	<pre> interface Ethernet1/3 description AzS-HCI Fabric-B NIC Port switchport mode trunk switchport trunk native vlan 126 switchport trunk allowed vlan 101,126,207 priority-flow-control mode on spanning-tree port type edge trunk mtu 9216 service-policy type qos input AzS_HCI_QoS service-policy type queuing output QOS_EGRESS_PORT no shutdown </pre>
<pre> interface Ethernet1/4 description AzS-HCI Fabric-A NIC Port switchport mode trunk switchport trunk native vlan 126 switchport trunk allowed vlan 101,107,126 priority-flow-control mode on spanning-tree port type edge trunk mtu 9216 service-policy type qos input AzS_HCI_QoS service-policy type queuing output QOS_EGRESS_PORT no shutdown </pre>	<pre> interface Ethernet1/4 description AzS-HCI Fabric-B NIC Port switchport mode trunk switchport trunk native vlan 126 switchport trunk allowed vlan 101,126,207 priority-flow-control mode on spanning-tree port type edge trunk mtu 9216 service-policy type qos input AzS_HCI_QoS service-policy type queuing output QOS_EGRESS_PORT no shutdown </pre>

ToR-A	ToR-B
interface Ethernet1/5	interface Ethernet1/5
interface Ethernet1/6	interface Ethernet1/6
interface Ethernet1/7	interface Ethernet1/7
interface Ethernet1/8	interface Ethernet1/8
interface Ethernet1/9	interface Ethernet1/9
interface Ethernet1/10	interface Ethernet1/10
interface Ethernet1/11	interface Ethernet1/11
interface Ethernet1/12	interface Ethernet1/12
interface Ethernet1/13	interface Ethernet1/13
interface Ethernet1/14	interface Ethernet1/14
interface Ethernet1/15	interface Ethernet1/15
interface Ethernet1/16	interface Ethernet1/16
interface Ethernet1/17	interface Ethernet1/17
interface Ethernet1/18	interface Ethernet1/18
interface Ethernet1/19	interface Ethernet1/19
interface Ethernet1/20	interface Ethernet1/20
interface Ethernet1/21	interface Ethernet1/21
interface Ethernet1/22	interface Ethernet1/22
interface Ethernet1/23	interface Ethernet1/23

ToR-A	ToR-B
interface Ethernet1/24	interface Ethernet1/24
interface Ethernet1/25	interface Ethernet1/25
interface Ethernet1/26	interface Ethernet1/26
interface Ethernet1/27	interface Ethernet1/27
interface Ethernet1/28	interface Ethernet1/28
interface Ethernet1/29	interface Ethernet1/29
interface Ethernet1/30	interface Ethernet1/30
interface Ethernet1/31 description VPC Peer to ToR-B:/1/31 switchport mode trunk switchport trunk allowed vlan 101,107,126,200,207 channel-group 10 mode active no shutdown	interface Ethernet1/31 description VPC Peer to ToR-A:/1/31 switchport mode trunk switchport trunk allowed vlan 101,107,126,200,207 channel-group 10 mode active no shutdown
interface Ethernet1/32 description VPC Peer to ToR-B:/1/32 switchport mode trunk switchport trunk allowed vlan 101,107,126,200,207 channel-group 10 mode active no shutdown	interface Ethernet1/32 description VPC Peer to ToR-A:/1/32 switchport mode trunk switchport trunk allowed vlan 101,107,126,200,207 channel-group 10 mode active no shutdown
interface Ethernet1/33 description P2P_Boarder1_To_ToR1 no switchport mtu 9216 ip address 192.168.200.49/30 no shutdown	interface Ethernet1/33 description P2P_Boarder1_To_ToR2 no switchport mtu 9216 ip address 192.168.200.57/30 no shutdown
interface Ethernet1/34	interface Ethernet1/34

ToR-A	ToR-B
<pre> description P2P_Boarder2_To_ToR1 no switchport mtu 9216 ip address 192.168.200.53/30 no shutdown interface Ethernet1/35 interface Ethernet1/36 interface mgmt0 vrf member management ip address 192.168.0.24/24 interface loopback0 description INFRA:Loopback_/ToR1:AzS-HCI- TOR-1:192.168.200.41 ip address 192.168.200.41/32 line console line vty boot nxos bootflash:/nxos64-cs.10.3.2.F.bin router bgp 64911 router-id 192.168.200.41 bestpath as-path multipath-relax log-neighbor-changes address-family ipv4 unicast network 192.168.101.0/24 network 192.168.126.0/26 network 192.168.200.41/32 network 192.168.200.44/30 network 192.168.200.56/30 network 192.168.200.60/30 maximum-paths 8 maximum-paths ibgp 8 template peer Border1-64821 remote-as 64821 address-family ipv4 unicast maximum-prefix 12000 warning-only </pre>	<pre> description P2P_Boarder2_To_ToR2 no switchport mtu 9216 ip address 192.168.200.61/30 no shutdown interface Ethernet1/35 interface Ethernet1/36 interface mgmt0 vrf member management ip address 192.168.0.25/24 interface loopback0 description INFRA:Loopback_/ToR1:AzS-HCI- TOR-1:192.168.200.42 ip address 192.168.200.42/32 line console line vty boot nxos bootflash:/nxos64-cs.10.3.2.F.bin router bgp 64911 router-id 192.168.200.42 bestpath as-path multipath-relax log-neighbor-changes address-family ipv4 unicast network 192.168.101.0/24 network 192.168.126.0/26 network 192.168.200.42/32 network 192.168.200.44/30 network 192.168.200.48/30 network 192.168.200.52/30 maximum-paths 8 maximum-paths ibgp 8 template peer Border1-64821 remote-as 64821 address-family ipv4 unicast maximum-prefix 12000 warning-only </pre>

ToR-A	ToR-B
<pre> template peer Border2-64821 remote-as 64821 address-family ipv4 unicast maximum-prefix 12000 warning-only template peer iBGPPeer-64911 remote-as 64911 address-family ipv4 unicast maximum-prefix 12000 warning-only neighbor 192.168.200.46 inherit peer iBGPPeer-64911 description 64811:P2P_TOR1-ibgp- 1_To_TOR2-ibgp-1:192.168.200.46 neighbor 192.168.200.50 inherit peer Border1-64821 description 64821:P2P_Boarder1_To_ToR1:192.168.200.5 0 address-family ipv4 unicast prefix-list ExternalPrefix in prefix-list ExternalPrefix out neighbor 192.168.200.54 inherit peer Border2-64821 description 64821:P2P_Boarder2_To_ToR1:192.168.200.5 4 address-family ipv4 unicast prefix-list ExternalPrefix in prefix-list ExternalPrefix out neighbor 192.168.101.0/24 inherit peer iBGPPeer-64911 description iBGPPeer-64911- Tenant:192.168.101.0 neighbor 192.168.126.0/26 inherit peer iBGPPeer-64911 description iBGPPeer-64911- Management:192.168.126.0 </pre>	<pre> template peer Border2-64821 remote-as 64821 address-family ipv4 unicast maximum-prefix 12000 warning-only template peer iBGPPeer-64911 remote-as 64911 address-family ipv4 unicast maximum-prefix 12000 warning-only neighbor 192.168.200.45 inherit peer iBGPPeer-64911 description 64811:P2P_TOR1-ibgp- 1_To_TOR2-ibgp-1:192.168.200.45 neighbor 192.168.200.58 inherit peer Border1-64821 description 64821:P2P_Boarder1_To_ToR1:192.168.200.5 8 address-family ipv4 unicast prefix-list ExternalPrefix in prefix-list ExternalPrefix out neighbor 192.168.200.62 inherit peer Border2-64821 description 64821:P2P_Boarder2_To_ToR1:192.168.200.6 2 address-family ipv4 unicast prefix-list ExternalPrefix in prefix-list ExternalPrefix out neighbor 192.168.101.0/24 inherit peer iBGPPeer-64911 description iBGPPeer-64911- Tenant:192.168.101.0 neighbor 192.168.126.0/26 inherit peer iBGPPeer-64911 description iBGPPeer-64911- Management:192.168.126.0 </pre>

About the Author

Sanjeev Naldurgkar, Technical Leader, Cisco Systems, Inc.

Sanjeev Naldurgkar is a technical leader on the Cisco UCS Solutions Engineering and Technical Marketing team, focusing on Microsoft solutions that include Azure Stack Hub, Azure Stack HCI, and Azure. His two decades of IT experience span multiple companies including Microsoft. Sanjeev holds a bachelor's degree in Electronics and Communications Engineering, along with leading industry certifications from Microsoft and VMware.

Acknowledgements

For their support and contribution to the design, validation, and creation of this Cisco Validated Design, we would like to thank:

- Bogna Trimouillat, Cisco Systems, Inc.
- Kirk Davidson, Cisco Systems, Inc.
- Babu Mahadevan, Cisco Systems, Inc.

Feedback

For comments and suggestions about this guide and related guides, join the discussion on [Cisco Community](https://cs.co/en-cvds) at <https://cs.co/en-cvds>.

CVD Program

ALL DESIGNS, SPECIFICATIONS, STATEMENTS, INFORMATION, AND RECOMMENDATIONS (COLLECTIVELY, "DESIGNS") IN THIS MANUAL ARE PRESENTED "AS IS," WITH ALL FAULTS. CISCO AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE. IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THE DESIGNS, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

THE DESIGNS ARE SUBJECT TO CHANGE WITHOUT NOTICE. USERS ARE SOLELY RESPONSIBLE FOR THEIR APPLICATION OF THE DESIGNS. THE DESIGNS DO NOT CONSTITUTE THE TECHNICAL OR OTHER PROFESSIONAL ADVICE OF CISCO, ITS SUPPLIERS OR PARTNERS. USERS SHOULD CONSULT THEIR OWN TECHNICAL ADVISORS BEFORE IMPLEMENTING THE DESIGNS. RESULTS MAY VARY DEPENDING ON FACTORS NOT TESTED BY CISCO.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unified Computing System (Cisco UCS), Cisco UCS B-Series Blade Servers, Cisco UCS C-Series Rack Servers, Cisco UCS S-Series Storage Servers, Cisco UCS X-Series, Cisco UCS Manager, Cisco UCS Management Software, Cisco Unified Fabric, Cisco Application Centric Infrastructure, Cisco Nexus 9000 Series, Cisco Nexus 7000 Series, Cisco Prime Data Center Network Manager, Cisco NX-OS Software, Cisco MDS Series, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. (LDW_UP5)

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

Americas Headquarters

Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters

Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters

Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at <https://www.cisco.com/go/offices>.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)