How to Deploy Azure Stack HCI clusters via Microsoft Windows Admin Center

Last Update: November 2021

Provides instructions to deploy ThinkAgile MX solutions via WAC wizard

Applies to Azure Stack HCI and Windows Server operating systems

Includes post-deployment configuration guidance

Includes comprehensive list of Lenovo and Microsoft reference links

Dave Feisthammel
Hussein Jammal
Mike Miller
David Ye

LENOVO PRESS

Table of Contents

1 Abstract

This document provides step-by-step instructions for deploying an Azure Stack HCI cluster on Lenovo ThinkAgile MX Certified Nodes and Appliances using the Microsoft Windows Admin Center (WAC) deployment wizard. It guides the reader through the wizard, explaining each step and providing real world examples from our labs. After completing the wizard, additional configuration steps are presented to ensure readiness of this HCI solution for production use.

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges. See our publications at http://lenovopress.com.

Do you have the latest version? We update our papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

2 Introduction

Deploying hyperconverged infrastructure has become the de-facto standard for organizations looking to modernize their aging infrastructure. Large storage deployments are increasingly being replaced by HCI-based solutions for most general-purpose workloads. HCI has proven to deliver better efficiency and price performance in the datacenter. Additionally, customers have been choosing a hybrid approach, migrating certain workloads to the cloud, while keeping other workloads on-premises.

Azure Stack HCI operating system is Microsoft's HCI solution for customers who wish to run workloads on-premises and extend easily to Microsoft Azure for hybrid capabilities such as back-up, site recovery, storage, cloud-based monitoring and more. Whether you prefer to deploy the Azure Stack HCI operating system or Windows Server, Lenovo® ThinkAgile™ MX solutions provide hardware that is certified for use in both scenarios.

The benefits of Lenovo HCI solutions include:

- Highly available and scale-on-demand compute/storage integrated solutions
- Easy to provision new IT services and reduce deployment time
- Better performance and lower Total Cost of Ownership (TCO)
- Flexible infrastructure and data centers

Lenovo has worked closely with Microsoft for many years to ensure our products perform smoothly and reliably with Microsoft operating systems and software. We have created Lenovo ThinkAgile MX Certified Node and Appliance solutions that contain only servers and server components that have been certified under the Microsoft Azure Stack HCI Program to run Microsoft Storage Spaces Direct (S2D) properly. These solutions provide a solid foundation for customers looking to consolidate both storage and compute capabilities on a single hardware platform. They provide outstanding performance, high availability protection and effortless scale-out growth potential to accommodate evolving business needs. ThinkAgile MX Series platforms offer the choice of Azure Stack HCI Appliances (called "Integrated Systems" by Microsoft) or Azure Stack HCI Certified Nodes (called "Validated Nodes" by Microsoft). These validated platforms help modernize on-premises infrastructure with pre-tested, pre-configured, and easy-to-order configurations, with seamless Azure integration.

This document provides step-by-step instructions for deploying an Azure Stack HCI cluster on Lenovo ThinkAgile MX Certified Nodes and Appliances using the WAC deployment wizard. It guides the reader through the wizard, explaining each step and providing real world examples. After completing the deployment wizard, additional configuration steps are presented to ensure readiness of this HCI solution for production use.

Other Azure Stack HCI deployment methods are available. Our Lenovo Press companion document, *Microsoft Storage Spaces Direct (S2D) Deployment Guide*, covers several deployment scenarios, as well as background discussions of the Microsoft Azure Stack HCI Certification Program. This document focuses on using PowerShell commands and scripting for deployment, rather than the WAC deployment wizard. You can find this document at the following URL:

https://lenovopress.com/lp0064

2.1 Architectural overview

Azure Stack HCI provides a hyperconverged infrastructure in which there is no separation between the resource pools for compute and storage. Instead, each server node provides hardware resources to support the running of VMs under Hyper-V, as well as the allocation of its internal storage to contribute to the cluster storage pool.

Figure demonstrates this all-in-one configuration for a four-node hyperconverged solution. When it comes to growth, each additional node added to the environment will mean both compute and storage resources are increased together. Perhaps workload metrics dictate that a specific resource increase is sufficient to cure a bottleneck (e.g., CPU resources). Nevertheless, any scaling will mean the addition of both compute and storage resources.

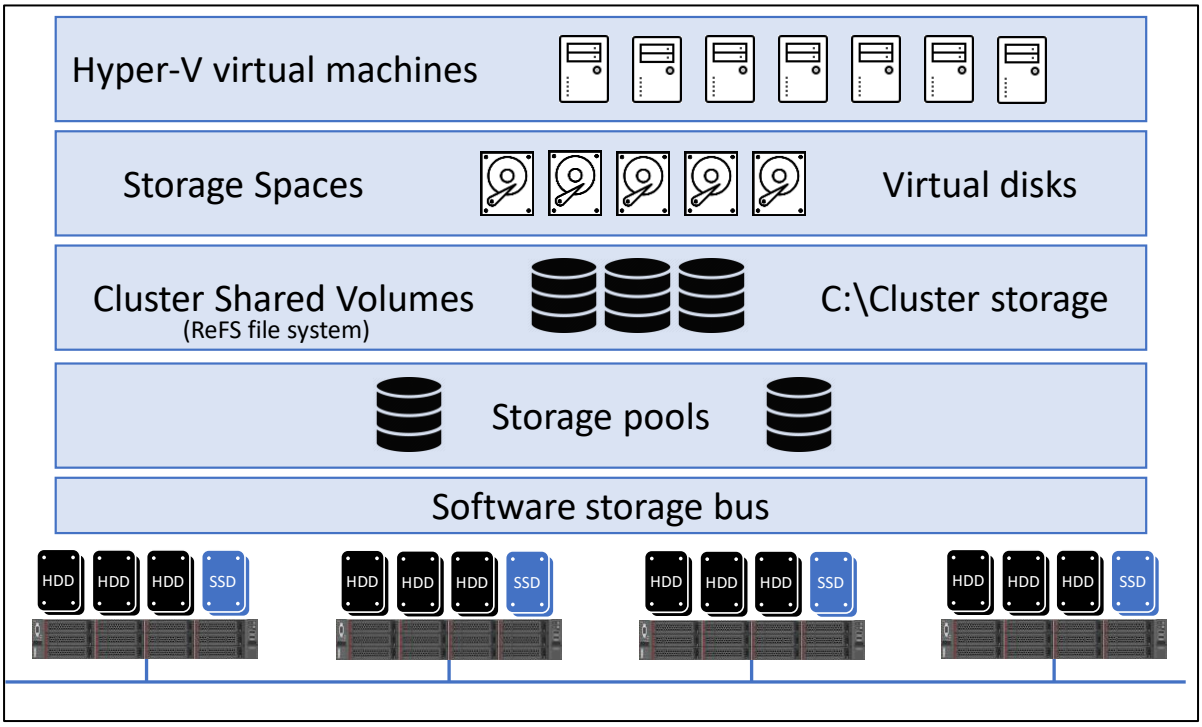

Figure 1. Azure Stack HCI block diagram

Azure Stack HCI functionality relies on Remote Direct Memory Access (RDMA) networking for storage (East-West) traffic inside the cluster. The two main implementations of RDMA that can be used for Azure Stack HCI are RDMA over Converged Ethernet version 2 (RoCEv2) and iWARP. Which implementation is chosen is primarily a personal preference. The key difference, in terms of the Azure Stack HCI deployment process, is that a RoCE implementation requires configuration of the network switches (if used) to enable Data Center Bridging (DCB), while iWARP does not require any special network switch configuration.

In Lenovo ThinkAgile MX solutions, Mellanox network adapters have been certified to carry RoCE storage traffic, while Marvell (Cavium/QLogic) network adapters have been certified to carry iWARP storage traffic. Although other network adapters can be used for North-South traffic into and out of the cluster (Microsoft refers to this as “management” traffic), only Mellanox or Marvell network adapters should be used for East-West storage traffic.

3 Solution deployment

Multiple methods can be used to deploy an Azure Stack HCI failover cluster, including GUI-based approaches, PowerShell methods, and via the WAC deployment wizard. If running the Azure Stack HCI operating system (or the Core version of Windows Server), GUI-based deployment is not possible since there is no GUI with which to work.

Although using PowerShell to deploy an HCI cluster is the most flexible method, Microsoft has built a cluster deployment wizard into WAC that simplifies the entire process. This document focuses on using the WAC deployment wizard to build and configure an HCI cluster. It also addresses several post-deployment steps that might be required for particular environments or to adhere to general best practices.

3.1 Deployment considerations

Before deployment of an Azure Stack HCI failover cluster can be initiated, the servers that will become cluster nodes must be prepared. In addition, if RoCEv2 is preferred for storage traffic, any network switches that will be used to carry this traffic must also be prepared. Switch configuration is outside the scope of this document. If you will be using existing Lenovo RackSwitch network switches to carry storage traffic, please refer to our companion document, *Microsoft Storage Spaces Direct (S2D) Deployment Guide* for details related to configuring Lenovo switches to carry storage traffic via RoCEv2.

There are several deployment scenarios available which are based on the implementation of RDMA being used (previously discussed), the total number of nodes to be added to the HCI cluster, and whether network switches will be used to carry East-West storage traffic. HCI cluster deployment via the WAC deployment wizard is very similar for all scenarios after initial system configuration is complete. The examples shown in this document are based on a 4-node HCI cluster deployment that uses a dual-port 25GbE Mellanox network adapter for RoCEv2 East-West (storage) traffic and two 10GbE LAN On Motherboard (LOM) ports configured in a team to carry North-South traffic into and out of the cluster (referred to as “management” traffic by WAC).

Although the WAC deployment wizard covers multiple HCI deployment scenarios, it does not handle certain situations, such as using one dual-port network adapter on each node to carry all traffic (that is, both East-West traffic and North-South traffic). If the WAC deployment wizard does not seem to support the desired deployment scenario, refer to our companion document *Microsoft Storage Spaces Direct (S2D) Deployment Guide*, at the following URL:

https://lenovopress.com/lp0064

The examples found in this document are based on creating a 4-node Azure Stack HCI cluster using Lenovo ThinkAgile MX3520-H Appliance nodes that are configured with a dual-port 25GbE Mellanox ConnectX-4 network adapter and two 10GbE LOM ports. Use of these four network interfaces present in our lab systems is shown in Figure 2 below.

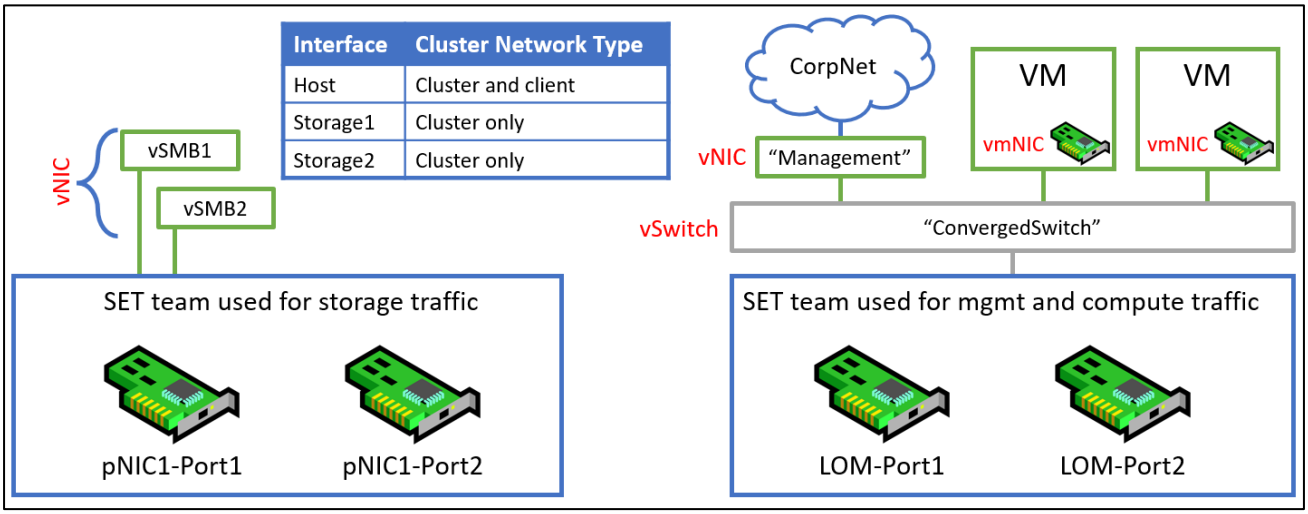

Figure 2. Network interfaces used in examples

As you can see from the above diagram, two SET teams will be created during HCI cluster deployment. The team on the left is built from two RDMA-capable network interfaces (using a dual-port 25GbE Mellanox NIC in our case) and is used for East-West storage traffic. The team on the right is built from two LOM ports and is used for North-South management and compute traffic.

3.2 General hardware preparation

Lenovo ThinkAgile MX solutions are built on top of multiple ThinkSystem rack servers, including the SR650, SR630 V2, and SR650 V2 rack servers. In addition, our ThinkAgile MX1020 Appliances and MX1021 Certified Nodes are based on the ThinkSystem SE350 edge server. For SE350-based solutions, we do not recommend using the WAC deployment wizard to build an HCI cluster due to the unusual network adapters (for example, the Wireless Network Module) that might not be recognized properly by the wizard. For detailed instructions on deploying HCI clusters based on the SE350 edge server, refer to our *ThinkAgile MX1021 on SE350 Azure Stack HCI (S2D) Deployment Guide* which can be found at the following URL:

https://lenovopress.com/lp1298

Before launching the WAC deployment wizard, all systems must be configured identically, including system settings, storage devices, and physical network adapters. We need to change a couple of system settings to optimize system performance and also to ensure that unneeded network interfaces do not cause any issues with cluster validation and creation later.

3.2.1 Operating Mode

The system Operating Mode should be changed to Maximum Performance to optimize system performance for its intended role as an Azure Stack HCI cluster node. To modify this system setting, follow these steps:

1. Reboot the server if necessary and enter the UEFI menu screen by pressing the F1 key when prompted at the bottom of the screen.
    a. If using the graphical system setup, navigate to UEFI Setup > System Settings and then select Operating Modes. Ensure that Choose Operating Mode is set to Maximum

Performance. Once this setting change has been made, click the Save icon on the right, and then click Back to return to the System Settings screen. Proceed with Step 2 below.

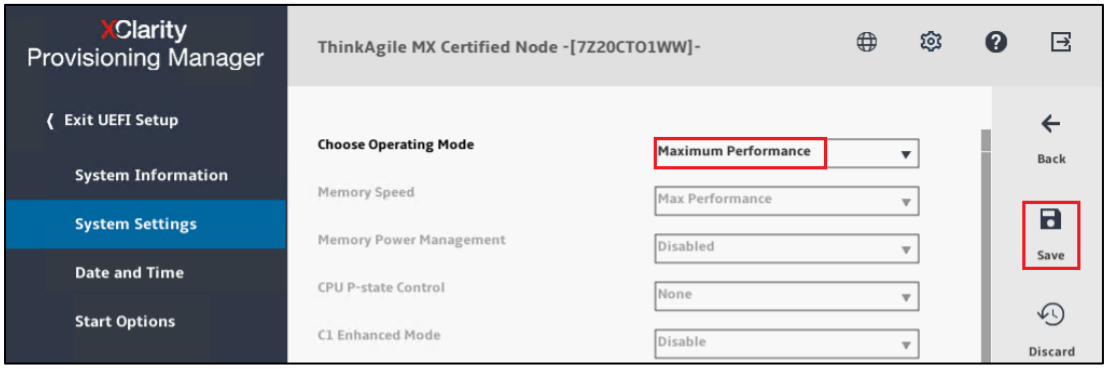

b. If using the text-based system setup, navigate to System Settings > Operating Modes. Ensure that Choose Operating Mode is set to Maximum Performance. If it is not, press Enter and use the arrow keys to select “Maximum Performance” before pressing Enter again. Once the setting change has been made, press the Esc key to return to the System Settings screen, and then proceed with Step 2 below.

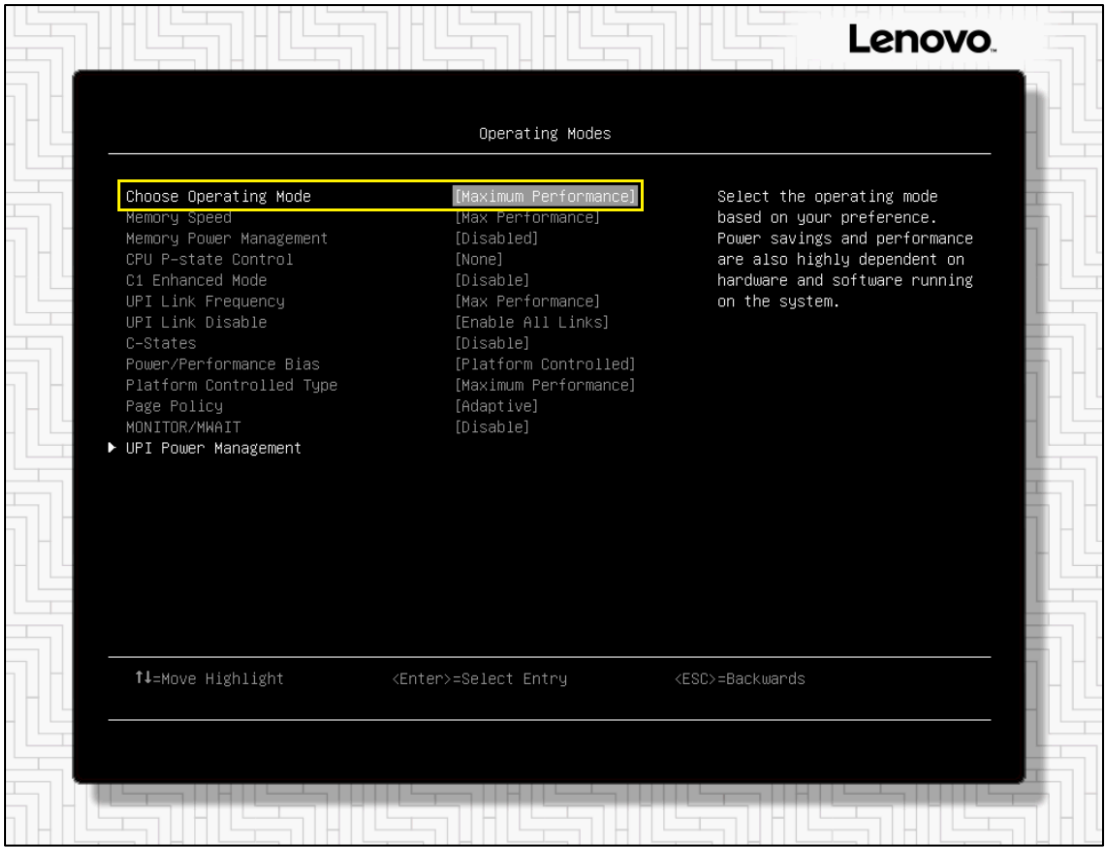

2. Once the Operating Mode has been set to Maximum Performance, continue with the next steps to disable unneeded network adapter ports in UEFI if necessary. If all installed network adapter ports will be used by Azure Stack HCI, exit System Setup, saving any changes made and reboot the system. Then proceed with the *Network preparation* section.

3.2.2 Disable unneeded network ports in UEFI

Lenovo ThinkSystem rack servers can be configured with network ports that are not associated with a conventional PCIe network adapter. For ThinkSystem SR630 and SR650 rack servers, these network ports are referred to as LAN on Motherboard (LOM) ports, which are backed by various Intel network adapters. None of these network adapters have been certified to carry storage (East-West) traffic, but all of them can be used for management (North-South) traffic.

However, ThinkSystem SR630 V2 and SR650 V2 rack servers can be configured with OCP network adapters, which are physically located in the same place as previous LOM ports. OCP network adapters for these servers include the Intel network adapters mentioned above, as well as OCP versions of both Mellanox and Marvell (Cavium/QLogic) network adapters.

The point to be made here is that any network ports that will not be used in the HCI cluster should not be visible to the operating system (or WAC) in order to avoid misconfiguration of the network adapters. For example, if a 4-port LOM/OCP network adapter is present in the system, but only 2 ports will be used by the HCI cluster, the other 2 ports should be disabled in system UEFI so they are not visible to WAC.

To disable unneeded LOM/OCP ports in system UEFI, follow these steps:

1. From the System Setup screen, follow the instructions below based on whether you are using the graphical or text-based system setup.
    a. If using the graphical system setup, in the main System Settings pane, navigate to Devices and I/O Ports > Enable/Disable Onboard Device(s) and scroll to the bottom of the device list.
    b. Disable either Onboard LAN to disable all ports, or disable each unneeded port individually, as necessary. The example below shows the configuration settings for using two ports from a 4-port LOM adapter. Once this setting change has been made, click the Save icon on the right, and then click Exit UEFI Setup to reboot the system.

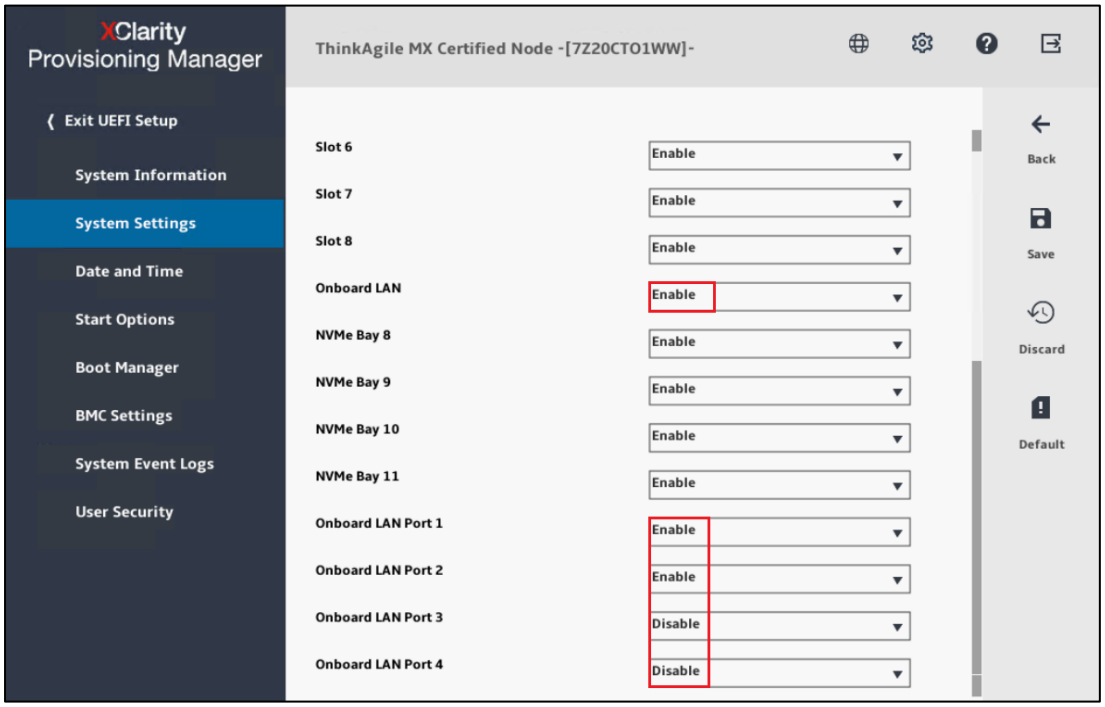

c. If using the text-based system setup, from the System Settings page, navigate to Devices and I/O Ports > Enable/Disable Onboard Device(s). Ensure that unneeded LOM/OCP ports are disabled. The example below shows the configuration settings for using two ports from a 4-port LOM adapter. Once this setting change has been made, press the Esc key repeatedly until prompted to save the new settings. Press the 'Y' key to save the settings and reboot the system.

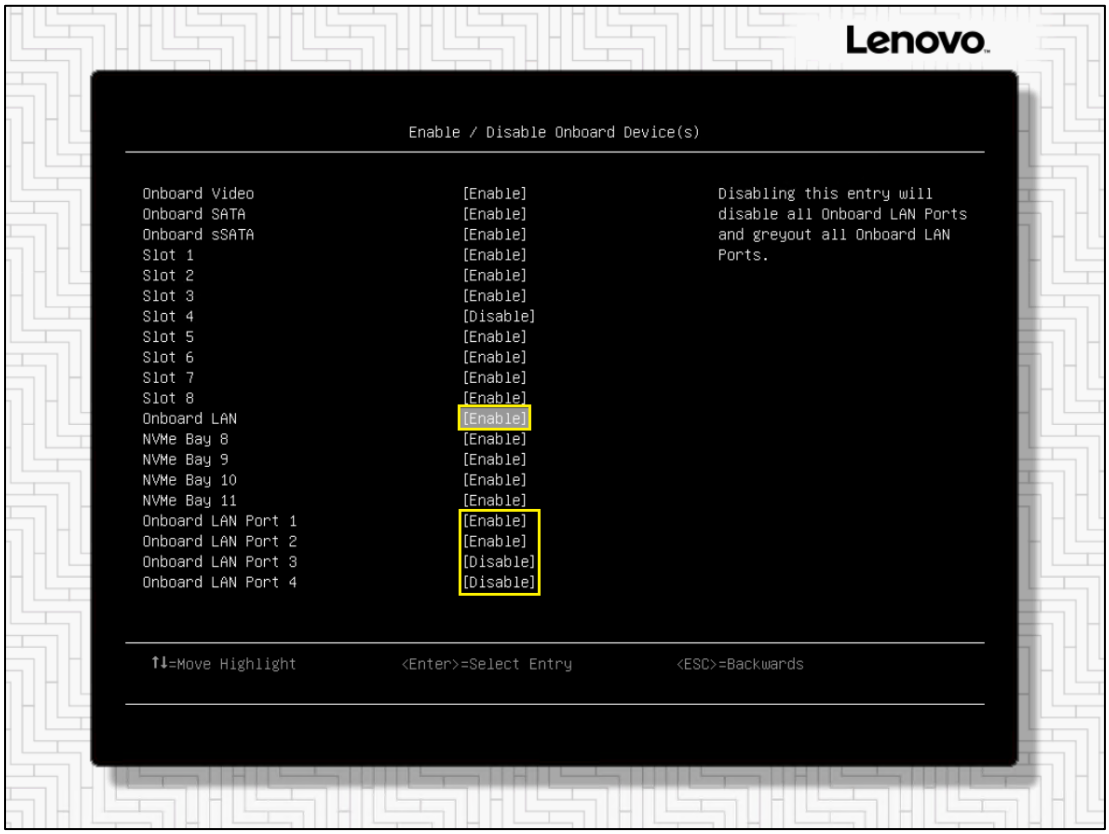

2. Once all system setting changes have been made, exit System Setup, saving changes when prompted. The system will reboot.

Note: Our instructions and example screen captures above are based on SR650 LOM ports, but disabling OCP ports in the SR650 V2 server is similar. For a V2 server, go to System Settings > Network > *First Port of OCP Adapter* > Port Enablement. If you do not see the Port Enablement option, you likely are not looking at the *first* port. Once in the Port Enablement option, press Enter and find an appropriate option, such as “Disable ports 3 & 4”. The example screenshot below shows how to disable Ports 3 and 4 in a 4-port OCP network adapter.

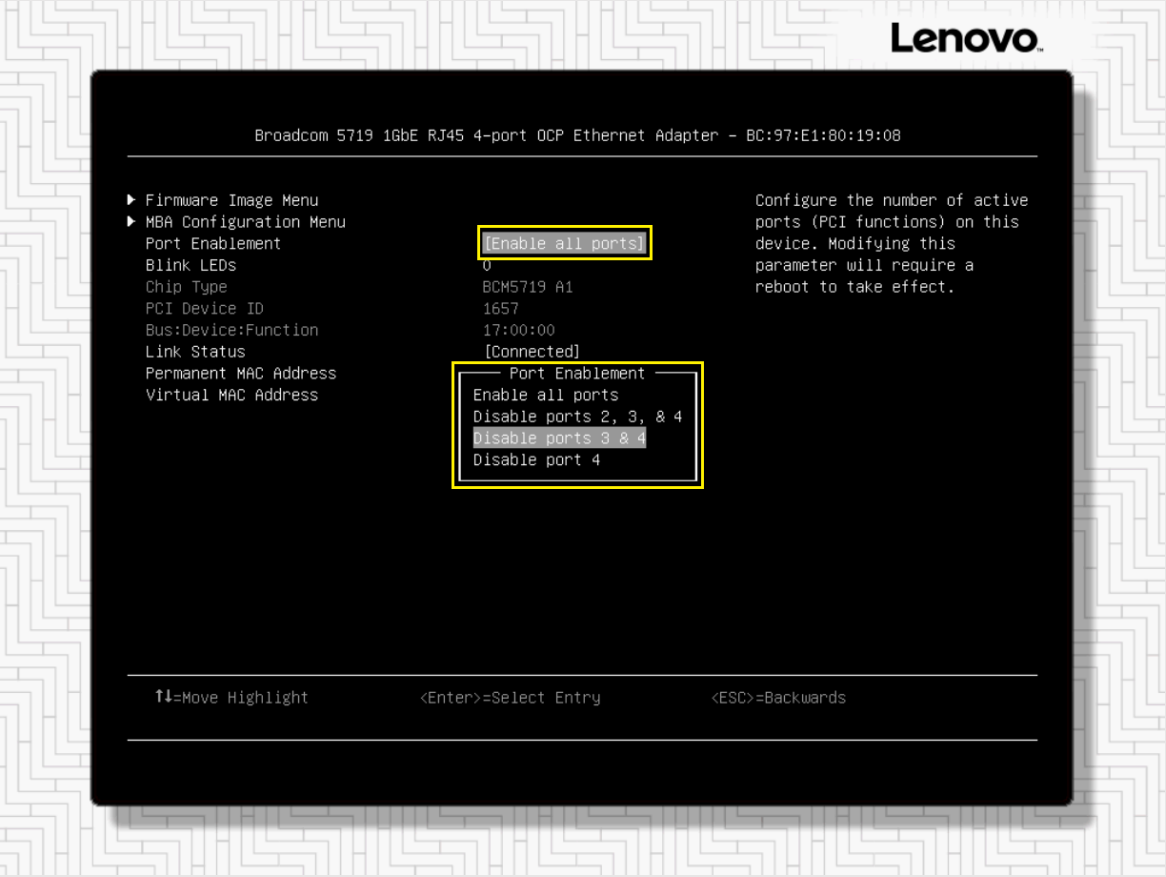

Once unneeded LOM/OCP ports have been disabled, proceed to the next section.

3.2.3 Firmware and driver updates

All ThinkAgile MX solutions make use of “Best Recipes” that define a set of firmware and device driver versions that have been tested as a unit and validated for use by Azure Stack HCI. For the latest set of ThinkAgile MX Best Recipes, refer to the following URL:

https://support.lenovo.com/us/en/solutions/HT507406

Firmware and device drivers must be identical between HCI cluster nodes. To make the process of updating firmware and drivers easy, the Lenovo XClarity Integrator (LXCI) for WAC can be used by the WAC deployment wizard to make appropriate updates during the deployment process. LXCI for WAC is a plug-in that integrates functions for managing, monitoring, and updating Lenovo ThinkAgile MX servers and their

components directly from the WAC user interface. For more information regarding LXCI for WAC, refer to the following URL:

https://support.lenovo.com/us/en/solutions/ht507549

For more information on installing LXCI for WAC, see section “Lenovo XClarity Integrator (LXCI) for Windows Admin Center.” If you do not install LXCI for WAC prior to HCI cluster deployment, you will need to ensure that all firmware and device drivers on all nodes are identical before launching the WAC deployment process.

3.3 Network preparation

Network switch configuration is outside the scope of this document. If you will be using existing Lenovo RackSwitch network switches to carry storage traffic, please refer to our companion document, *Microsoft Storage Spaces Direct (S2D) Deployment Guide* for details related to configuring Lenovo switches to carry storage traffic via RoCEv2.

Note: Although Lenovo network switches are no longer being offered, our companion deployment guide includes Lenovo switch configuration details to aid customers who already own these network switches.

It might be desirable to create a small HCI cluster (2-3 nodes) without using high-speed network switches to carry storage traffic. That is, each node is directly connected to every other node in the cluster to carry storage traffic. In this deployment scenario, network switches are needed only to carry North-South (management) traffic into and out of the cluster. Since network switches are not touched by WAC, its deployment wizard supports both direct-connected clusters and clusters that use network switches to carry storage traffic.

In all deployment scenarios, network adapters need to be at least partially configured before WAC is able to connect to the systems that will become nodes in the Azure Stack HCI cluster. This applies to the network adapter port(s) that will be used for management (North-South) traffic. Ports that will be used exclusively to carry storage (East-West) traffic do not need to be configured in advance.

3.3.1 Marvell network port configuration

Lenovo has certified QLogic/Cavium/Marvell QL41262 and QL41232 network adapters for use in ThinkAgile MX solutions to carry iWARP RDMA storage traffic. Although these network adapters support both iWARP and RoCEv2 implementations of RDMA, they are shipped with the RoCEv2 mode active by default. Since these network adapters have NOT been certified for RoCEv2 use, they must be configured to enable iWARP RDMA mode.

Making this change is a two-part process. First, a UEFI setting must be changed in System Setup and then PowerShell is leveraged to make the required advanced network adapter property change. These changes must be made on all systems that will become HCI cluster nodes. To configure the NIC for iWARP in server UEFI, follow these steps:

1. Reboot the server if necessary and enter the UEFI menu screen by pressing the F1 key when prompted at the bottom of the screen.
    a. If using the graphical system setup, navigate to UEFI Setup > System Settings > Network . Click on the item in the Network Device List that represents Port 1 on the QLogic QL41262

25Gb Ethernet adapter, and then click on Port Level Configuration. Change RDMA Operational Mode at the bottom of the center pane to iWARP.

b. Click Back twice to return to the Global Network Settings page and then click on the item in the Network Device List that represents Port 2 on the same network adapter. Repeat the setting change for this port. Repeat this step for both ports on the second QLogic QL41262 25Gb Ethernet adapter if it exists.

c. Once all ports on all Cavium/QLogic network adapters are configured for iWARP, click the Save icon on the right side of the page, and then click Exit UEFI Setup. Click to confirm and reboot the system. Proceed to Step 1 in the second part of this process below.

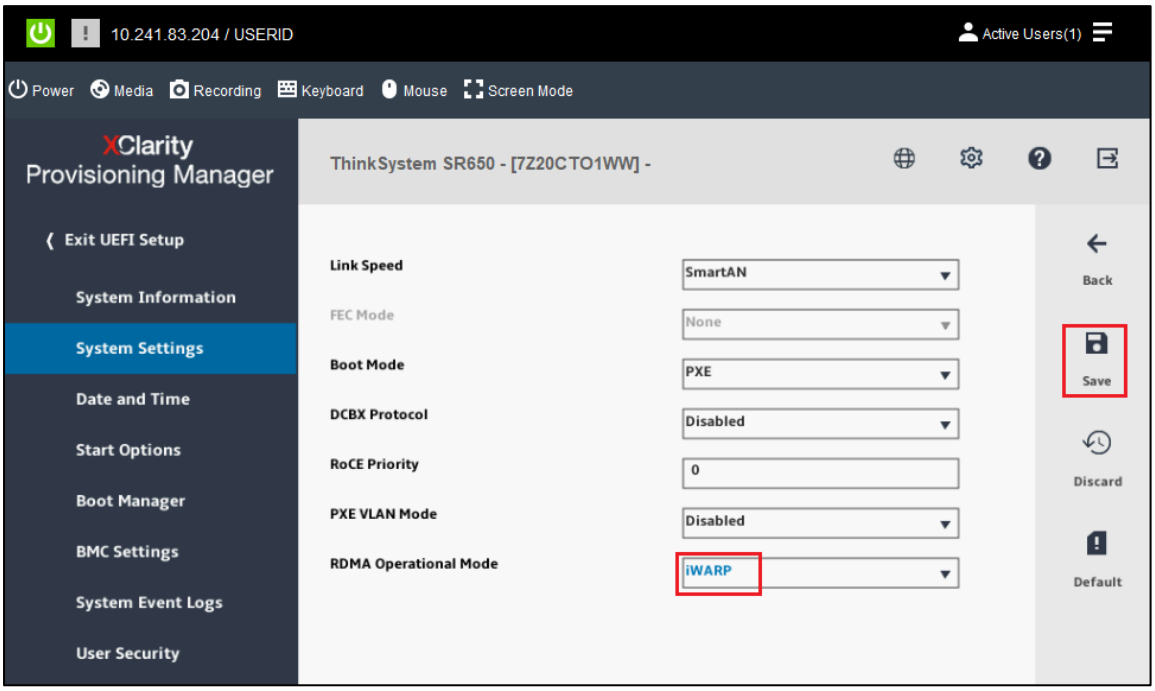

a. If using the text-based system setup, navigate to System Settings > Network and then select the item from the Network Device List that represents Port 1 on the QLogic QL41262 25Gb Ethernet adapter. Select Port Level Configuration and change RDMA Operational Mode to iWARP.

b. Press the ESC key twice to return to the Global Network Settings page and then select the item from the Network Device List that represents Port 2 on the same network adapter. Repeat the setting change for this port. Repeat this step for both ports on the second QLogic QL41262 25Gb Ethernet adapter if it exists.

c. Once all ports on all Cavium/QLogic network adapters are configured for iWARP, press the Esc key repeatedly to exit system settings, saving the changes that were made. The system will reboot. Proceed to Step 1 in the second part of this process below.

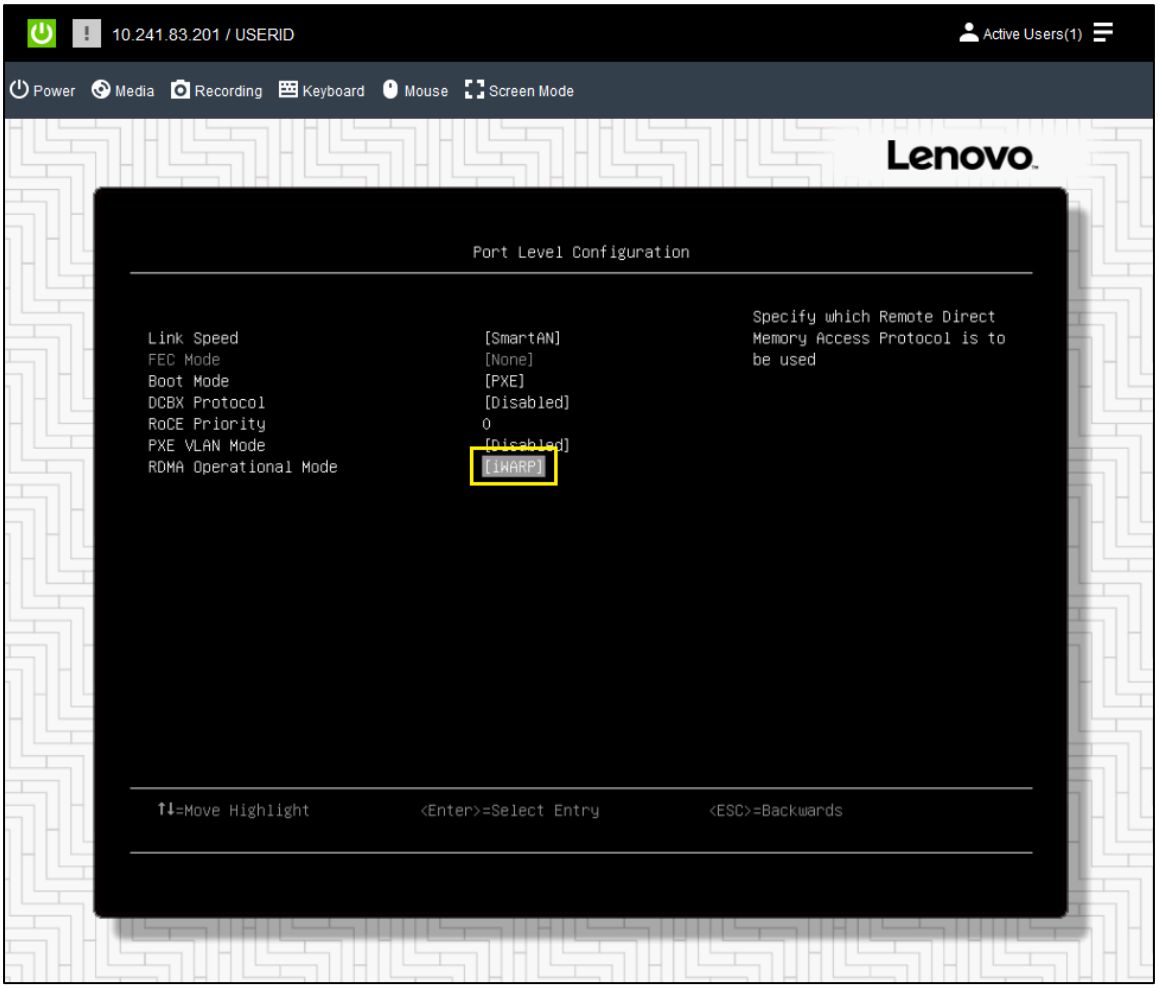

Once the RDMA Operational Mode has been set to iWARP in system UEFI and the system has rebooted, continue with the second part of this process to finish configuring the Cavium/QLogic NIC ports for iWARP in the operating system using PowerShell.

1. Log in to the Azure Stack HCI operating system on the server that will become the first HCI cluster node. This will launch SConfig, which is the main user interface for the OS.
2. Type 15 and press Enter to exit SConfig and open a PowerShell prompt.
3. Run the PowerShell commands shown in the example below to enable iWARP RDMA mode for each port on each network adapter. Make sure to change the network interface name shown in each command to the correct value for your system.

```
Set-NetAdapterAdvancedProperty -Name "Slot1 Port1" -DisplayName "NetworkDirect Technology"
-DisplayValue "iWarp"
Set-NetAdapterAdvancedProperty -Name "Slot1 Port2" -DisplayName "NetworkDirect Technology"
-DisplayValue "iWarp"
```

Once complete, repeat the steps in this entire section (both parts) for each of the systems that will become HCI cluster nodes.

3.3.2 Management network port configuration

After disabling any LOM/OCP ports in UEFI that will not be used (see *Disable unneeded network ports in UEFI* above for more information), the remaining enabled ports need to have appropriate IP configurations applied. In our example deployment, we use two LOM ports in an SR650 server to carry management traffic.

To configure these network interfaces, follow these steps:

1. Log in to the Azure Stack HCI operating system on the server that will become the first HCI cluster node. This will launch SConfig, which is the main user interface for the OS.
2. Type 8 and press Enter to see the Network settings screen.
3. Find the first network interface that needs to be configured in the Network settings screen, type the Index # for this network adapter and then press Enter.

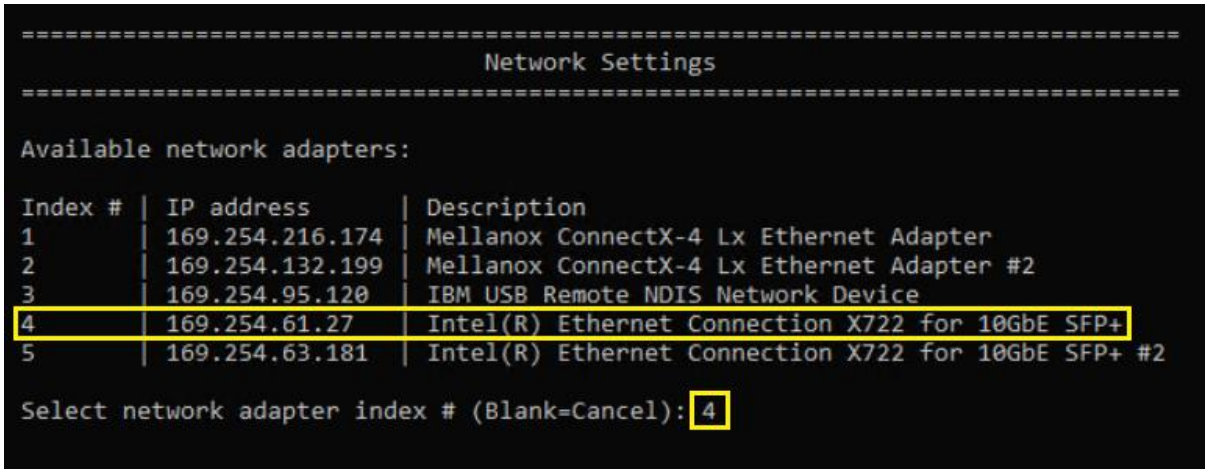

4. Type 1 and press Enter to change the IP address for the network interface.

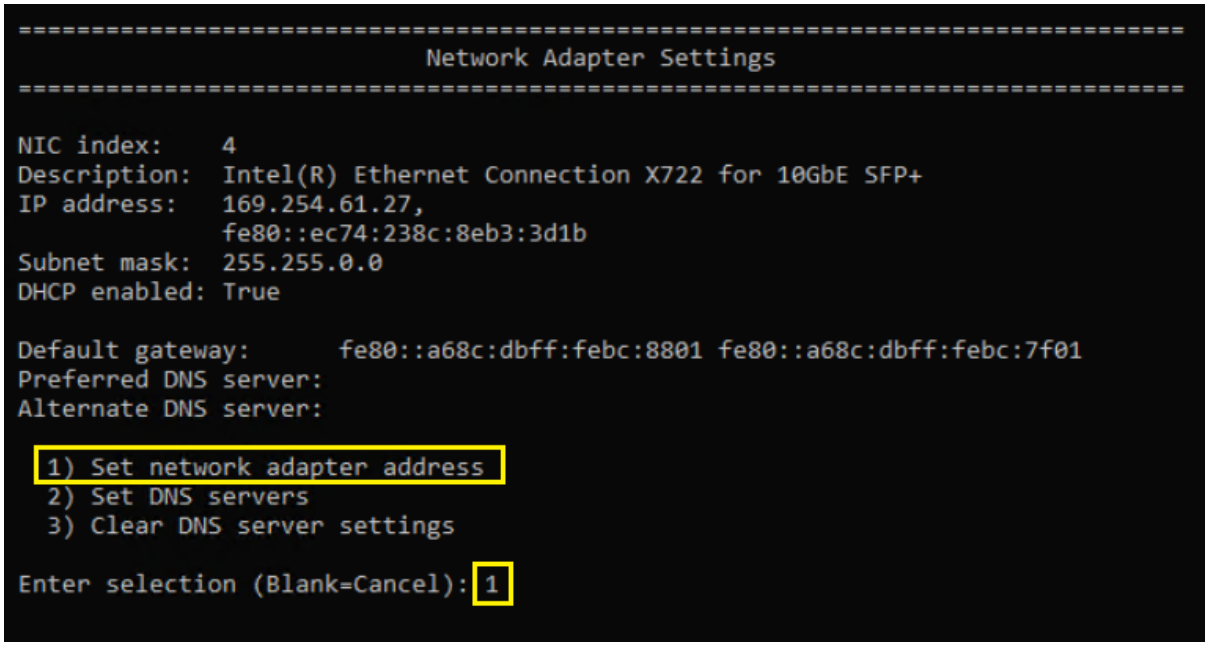

5. Select the appropriate entries and type the desired IP settings, including IP address, subnet mask, and default gateway (for static IP configuration). After entering the default gateway, press Enter again to begin IP configuration. Status is displayed as the configuration is applied. Once complete, press Enter to continue.

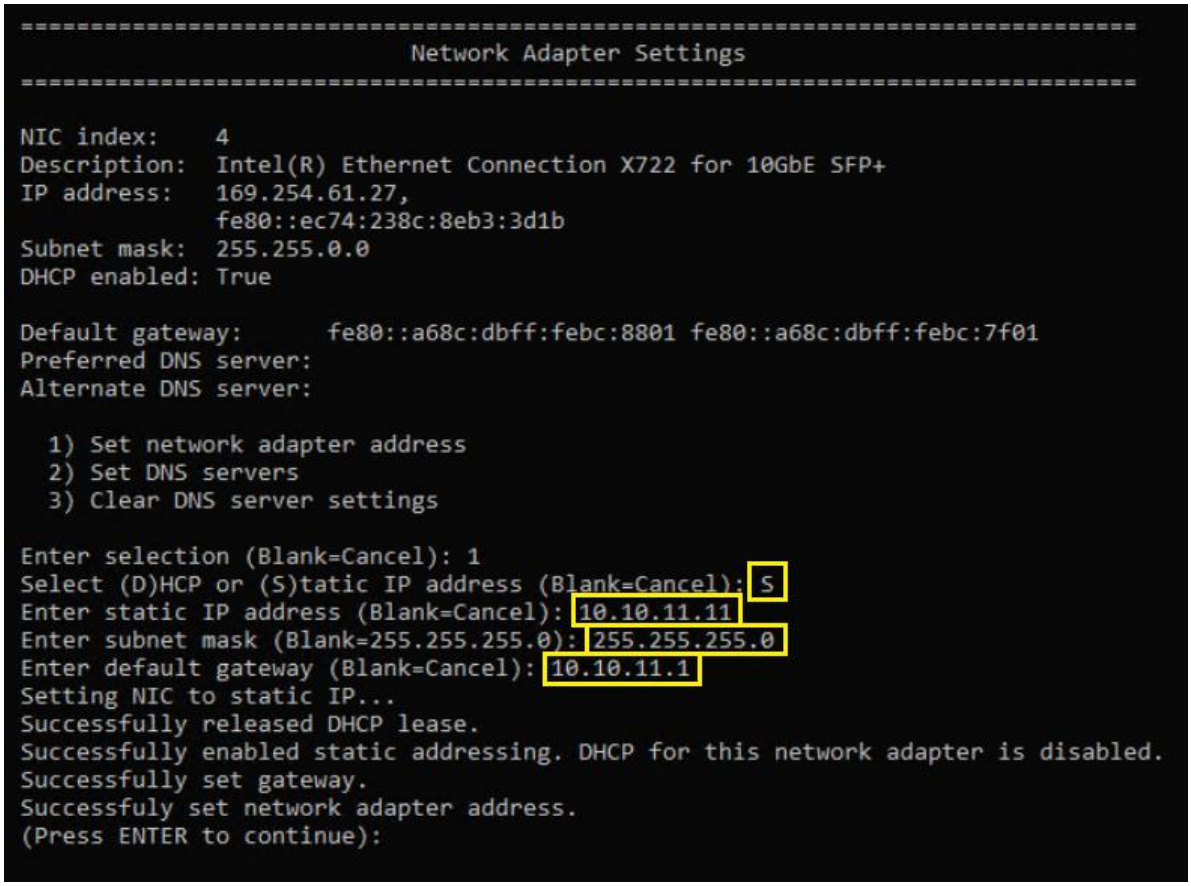

6. Back in SConfig, use a similar sequence to set DNS server(s) for this network interface.

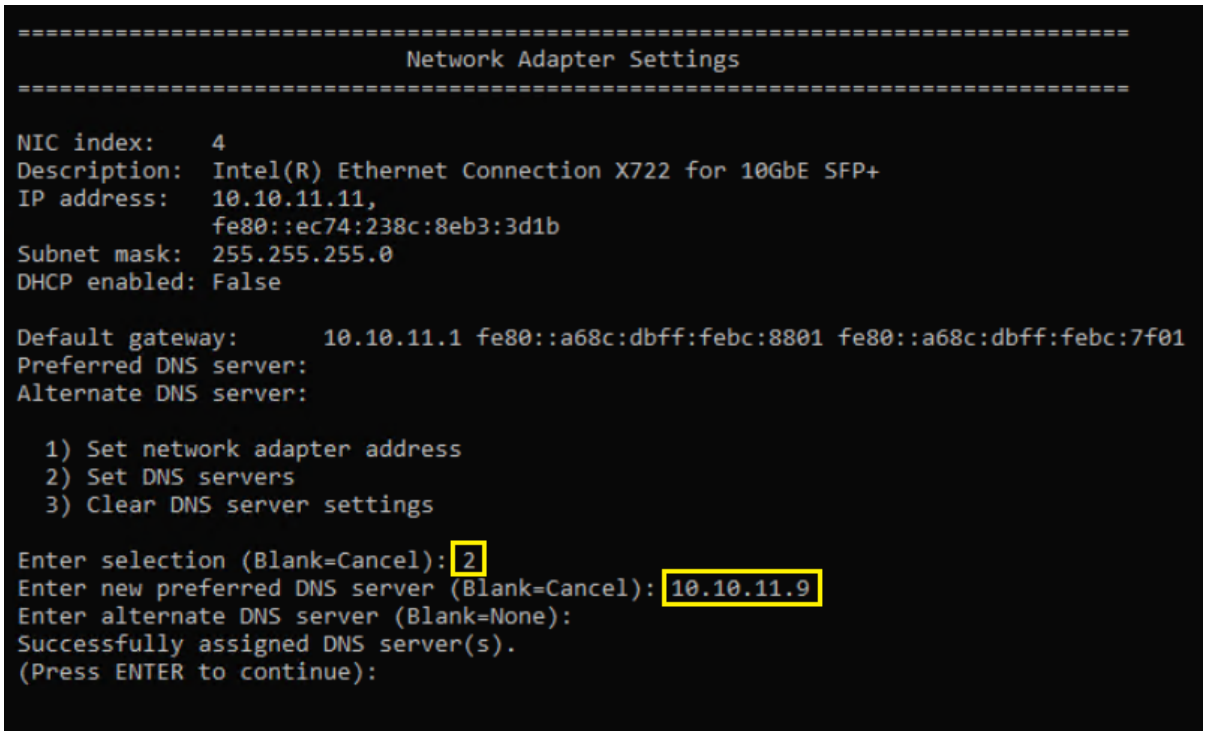

7. Once back in SConfig on the server that will become the first node in the HCI cluster, use Steps 1-6 above to configure the first management network interface on each of the other nodes.
8. Once the first management network interface has been configured on all systems that will become HCI cluster nodes, repeat Steps 1-7 above to configure the second management network interface on all nodes if desired. Our example deployment uses two LOM ports to carry management traffic.

With initial system configuration complete, we now turn our attention to WAC and its HCI failover cluster deployment wizard.

3.4 Windows Admin Center (WAC)

To get the most out of WAC and its HCI cluster deployment wizard, we highly recommend installing the Lenovo XClarity Integrator (LXCI) for WAC. With this WAC extension, firmware and device driver updates can be installed on all servers that will become nodes in the HCI cluster directly from the deployment wizard. Proceed to the next section to learn more about LXCI for WAC, including how to install it.

> **Note**: The examples and screenshots contained in this document were generated using the following software versions: HCI operating system “20H2”, WAC 1.3.2105, and LXCI v3.2.3.

3.4.1 Lenovo XClarity Integrator (LXCI) for Windows Admin Center

LXCI for WAC is a plug-in that integrates functions for managing, monitoring, and updating the Lenovo servers and their components with Windows OS or software application management system. It supports viewing Lenovo server hardware and firmware inventory, events, alerts, and health status, performing cluster-aware rolling updates of firmware and device drivers for Windows failover cluster nodes, displaying Lenovo ThinkAgile MX server topology views, and facilitating storage pool operations through wizards. Lenovo

XClarity Administrator (LXCA), which is optional, streamlines the Lenovo server management job, especially for large-scale deployments.

LXCI for WAC is installed as an extension to WAC using either the WAC feed, or a local shared folder containing the installation package. For additional details on LXCI for WAC, including how to install the extension via a local shared folder, see the following URL:

https://sysmgt.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.lxci_wac.doc%2Fwac_welcome.html

To install LXCI for WAC using the WAC feed, follow these steps:

1. Log in to WAC.
2. Click the settings icon  in the top right corner. The Settings page is displayed.
3. On the Settings page, click Extensions in the left navigation pane. The Extensions pane is displayed.
4. In the Extensions pane, click the Feeds tab.
5. On the Feeds tab, click Add. The Add package source pane is displayed on the right.
6. In the Add package source pane, either select the feed “https://aka.ms/sme-extension-feed” if it is already listed or click Add and enter the feed URL and then click Add.
7. Return to the Extensions pane and click the Available extensions tab.
8. On the Available extensions tab, select Lenovo XClarity Integrator from the list. License information will be displayed.
9. Read the license information. If you accept the license information, click Install.
10. When the "Install this extension?" window is displayed, click Confirm to continue. Installation confirmation will be displayed once the extension is installed.

11. Click the Installed extensions tab and scroll down to locate Lenovo XClarity Integrator to confirm that LXCI for WAC is now available for use.

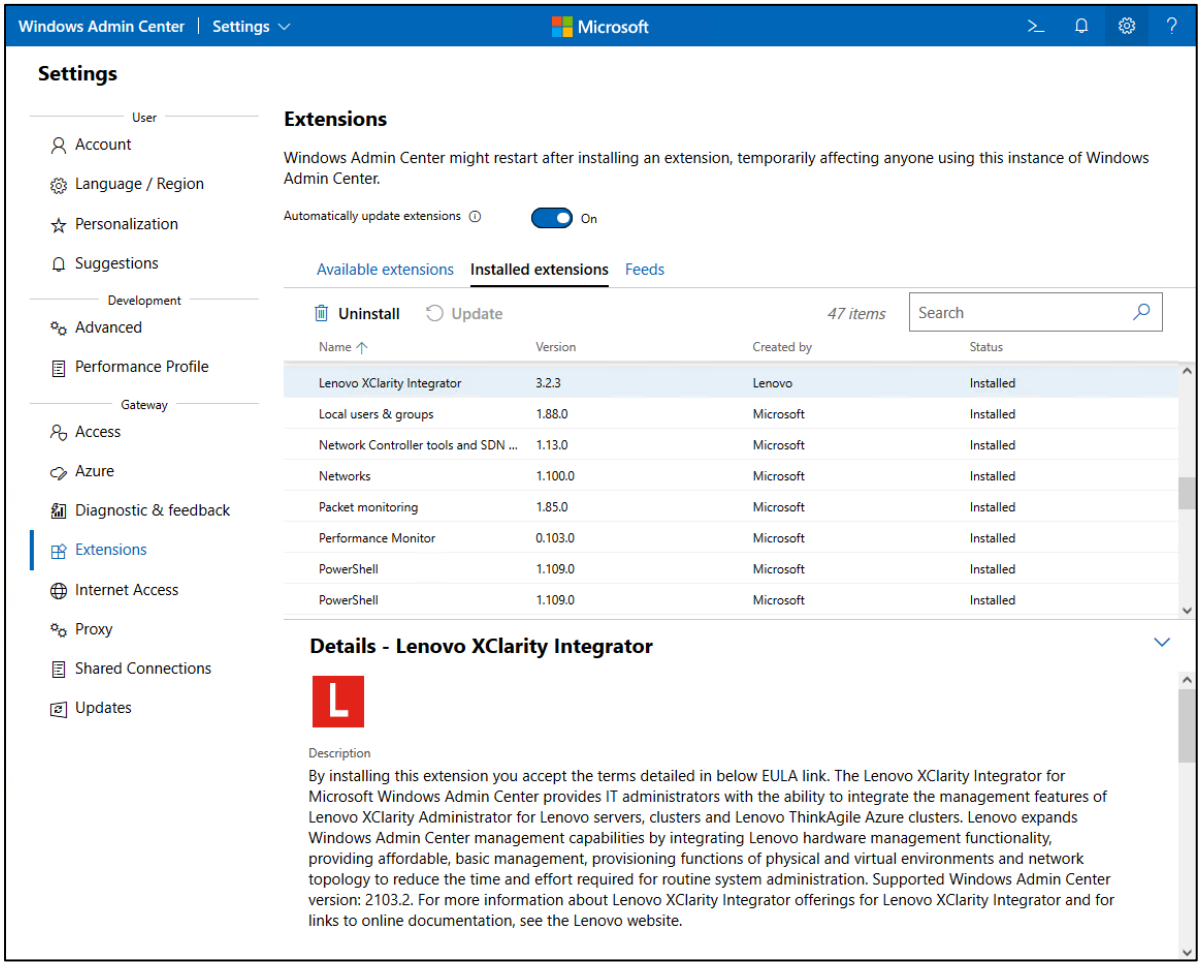

With the LXCI for WAC extension installed, we can proceed to the WAC cluster deployment wizard.

3.5 HCI cluster deployment

Once the above preparation has been completed, the process to deploy an Azure Stack HCI failover cluster is quite straight forward using the WAC deployment wizard. To deploy an HCI cluster via WAC, follow these steps:

1. Log in to WAC.
2. From the All connections view, click Add. The Add or create resources panel opens.

3. In the Server clusters box, click Create new.

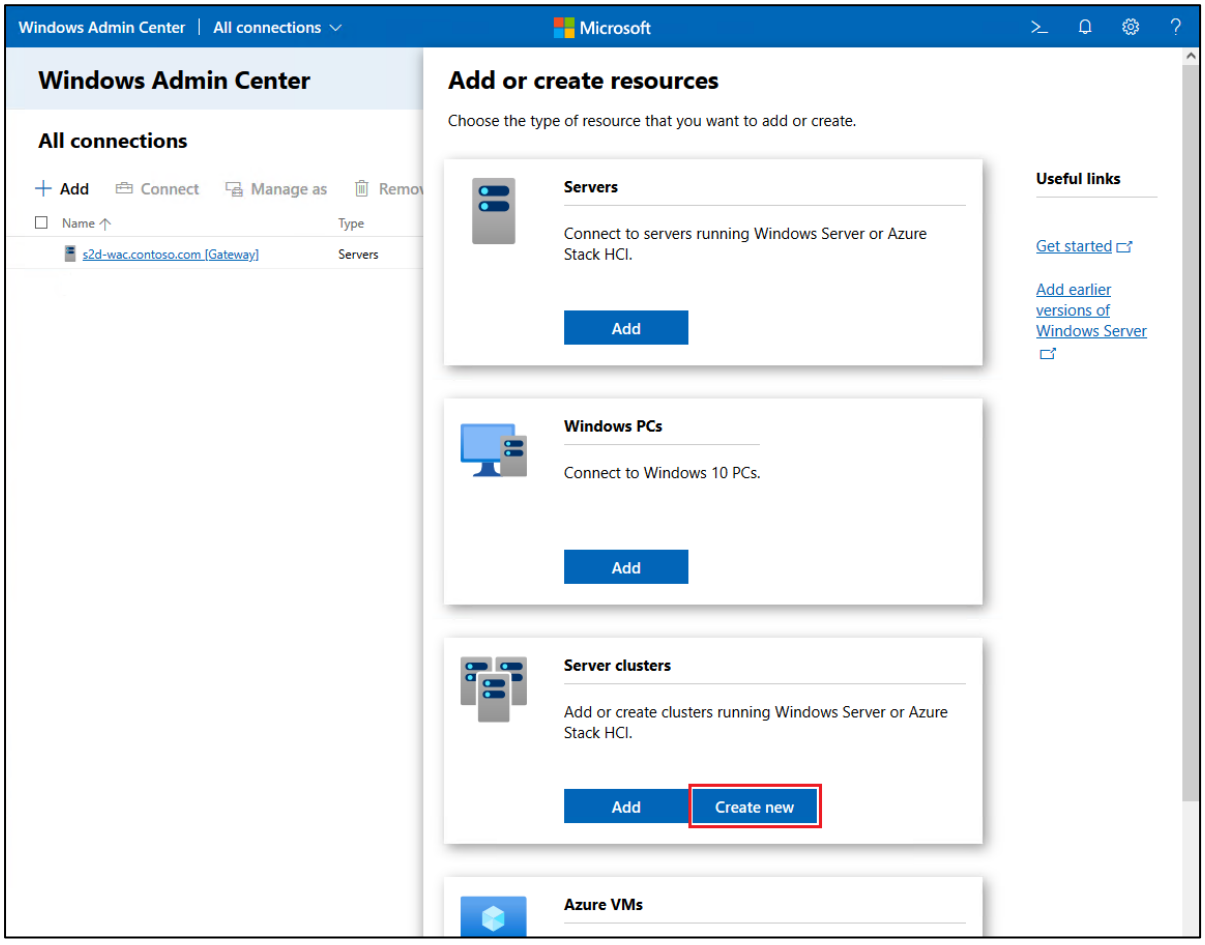

4. On the Cluster Creation page, select the Azure Stack HCI option. Our examples use the default All server in one site option. Once options have been selected, click Create. The Deploy an Azure Stack HCI cluster wizard opens.

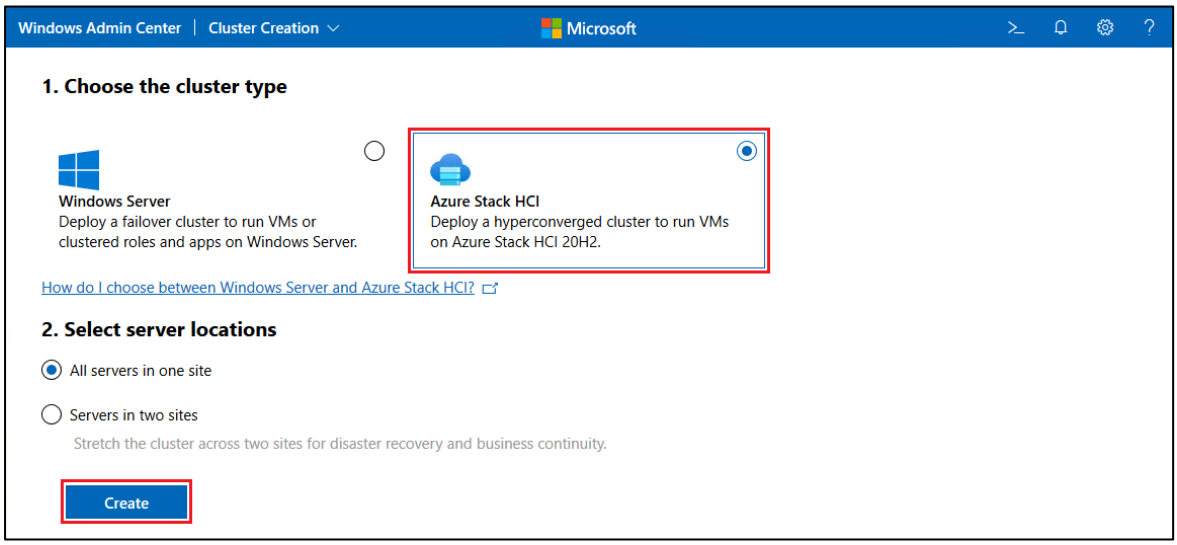

5. Review the Check the prerequisites panel to ensure that all prerequisites have been met and then click Next.

6. On the Add servers page, enter the credentials of an administrator account that will be used to connect to each of the servers that will become HCI cluster nodes.
7. Enter the IP address of the first server. Once the server is found, click Add. This will begin a short validation check and then add the server to the list.

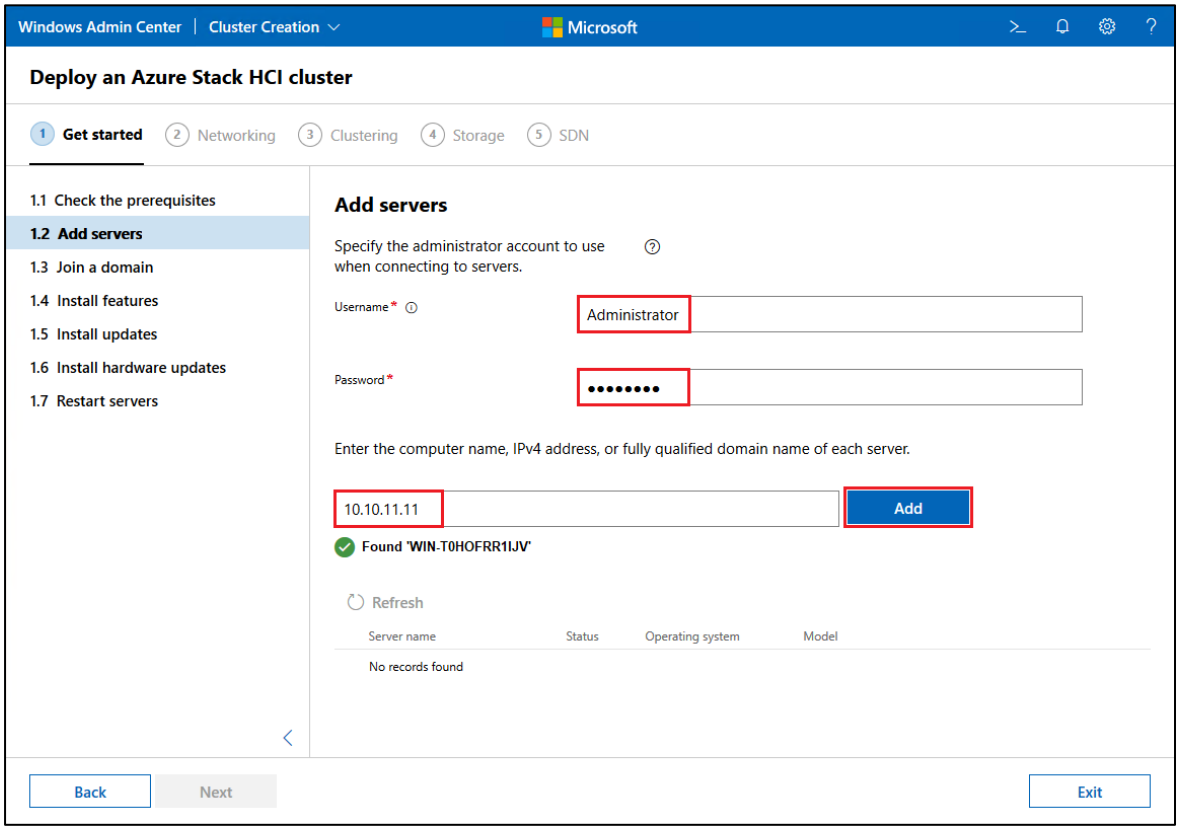

8. Repeat Step 7 for each of the remaining servers to be added to the cluster. Once all servers have been added to the list, click Next.

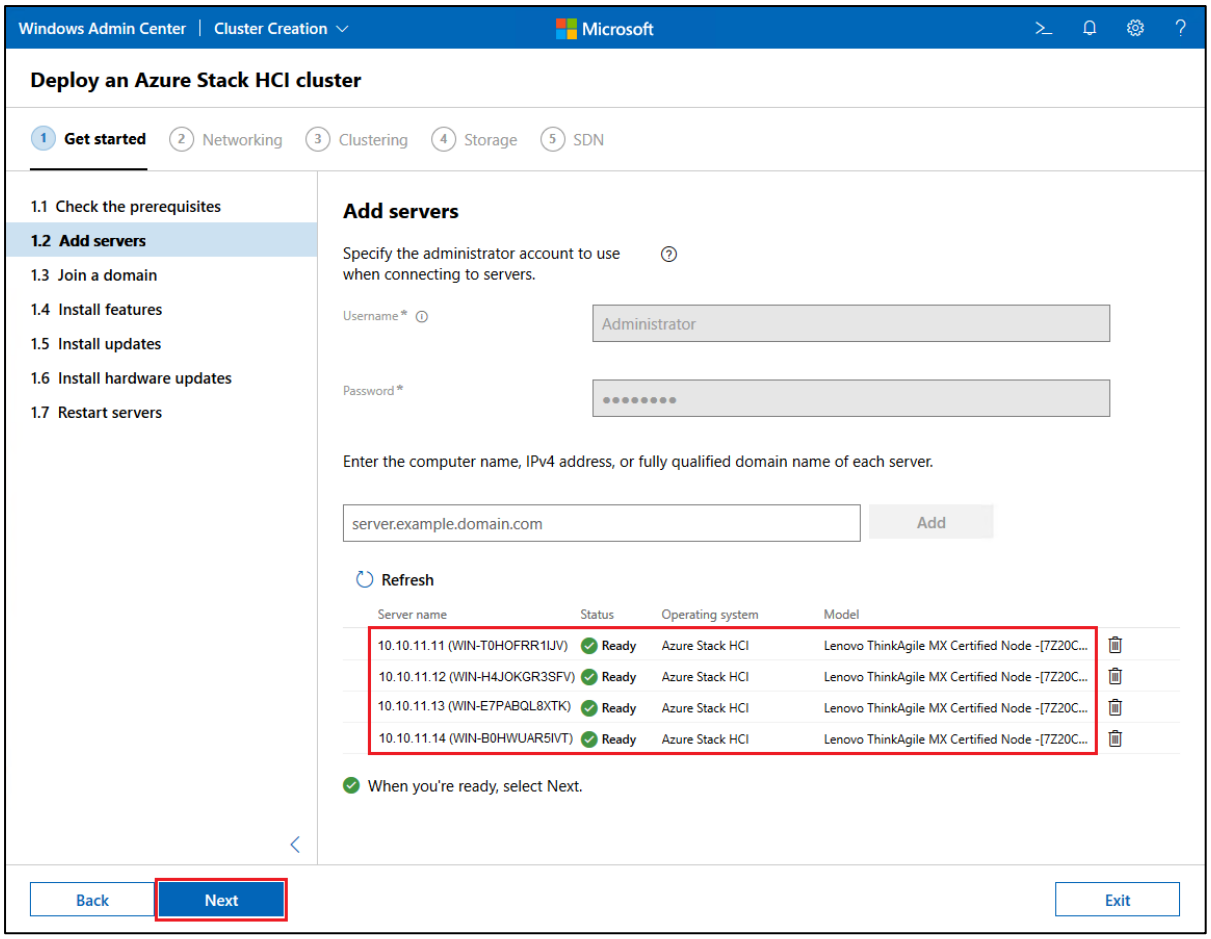

9. On the Join a domain page, enter the domain name and credentials for a domain account to use to join the servers to the domain specified. You can also change the computer name of each server if desired. Since we are working with servers that use the default random computer names generated

by Windows, we provide a new computer name for each server. Once all input is complete, click Apply changes. Progress will be shown as each server is added to the domain.

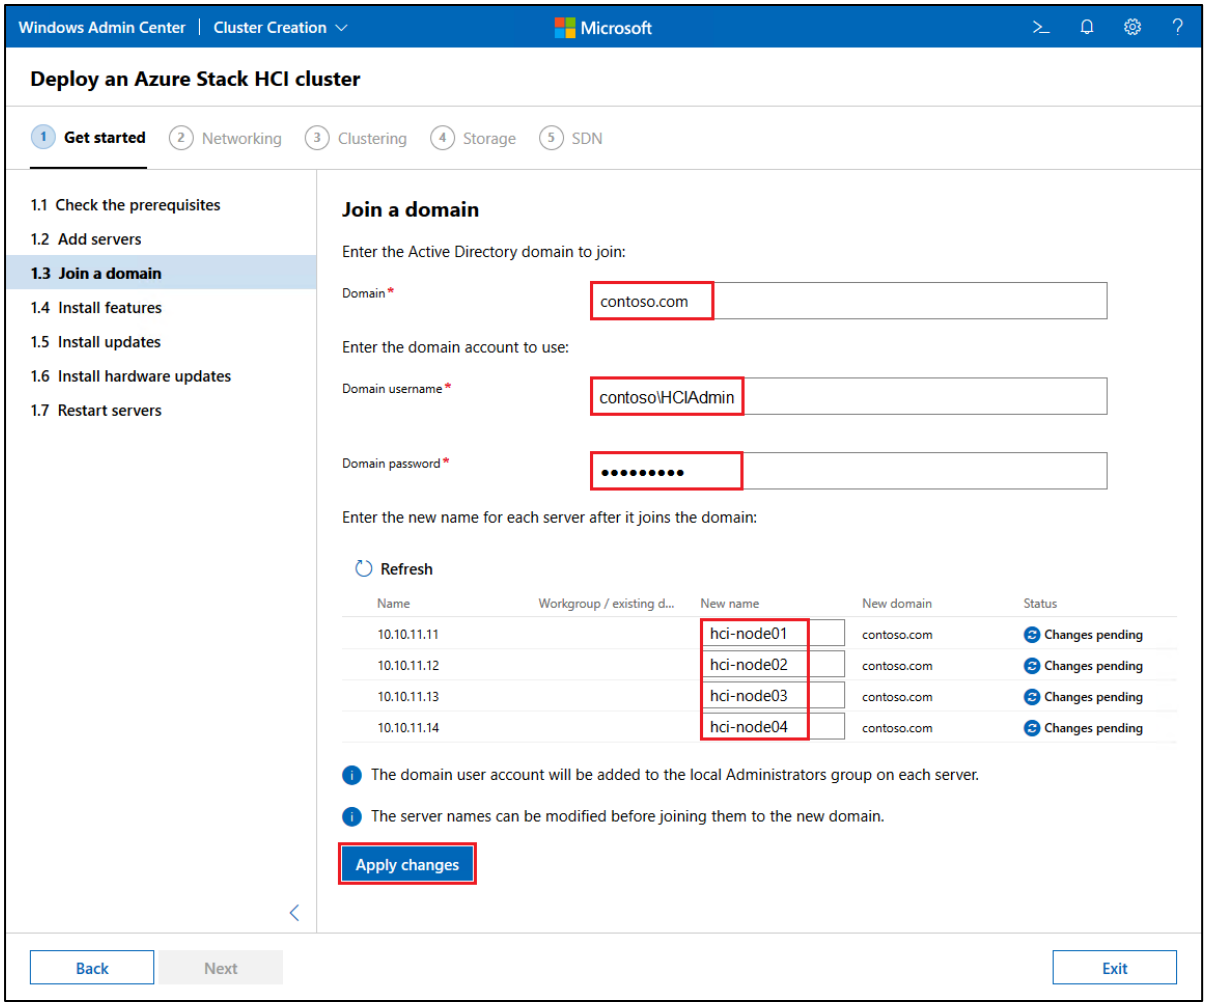

10. Once all servers have been added to the domain, click Next.

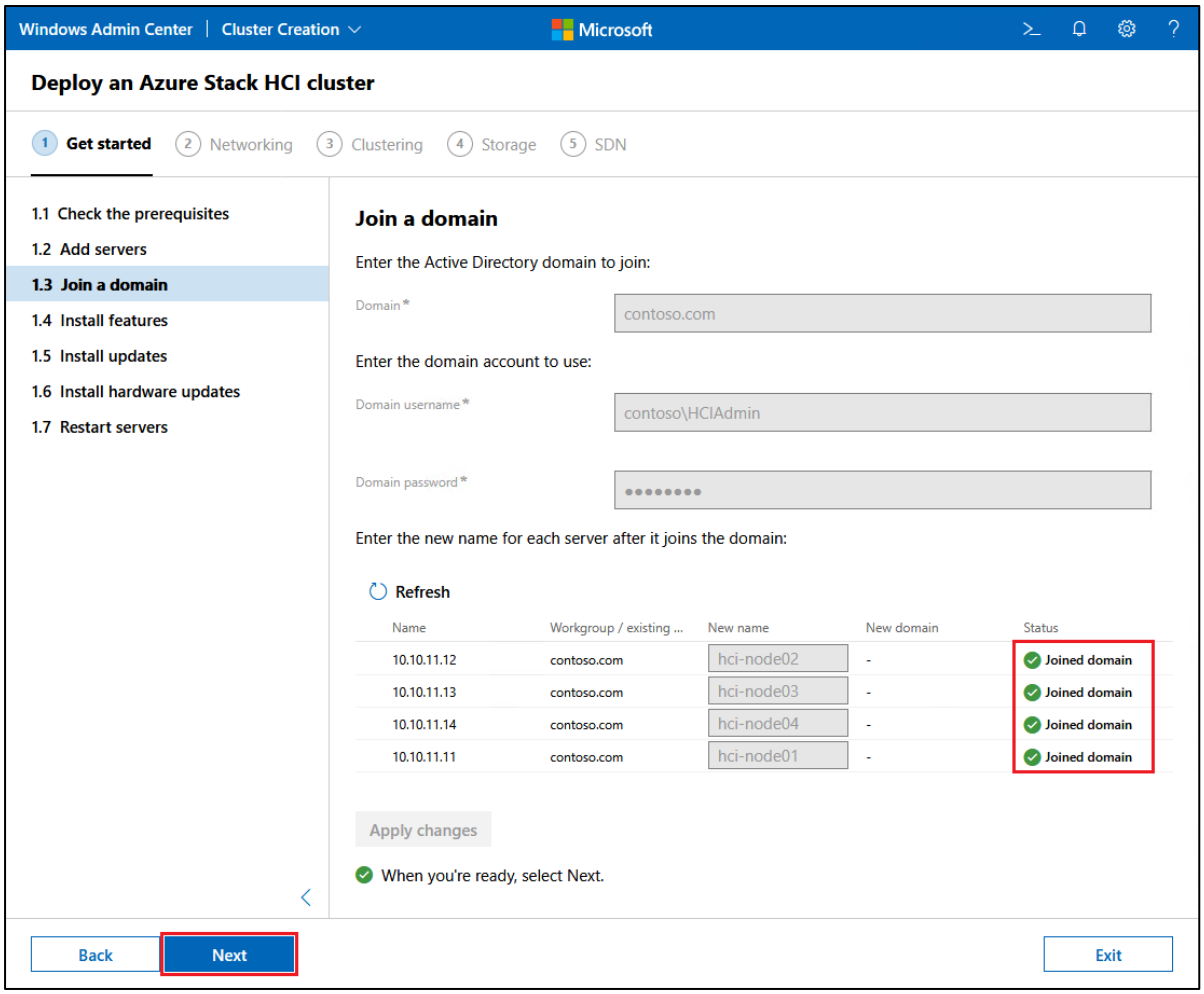

11. On the Install features page, if any servers show Status of “Not installed” click Install features.

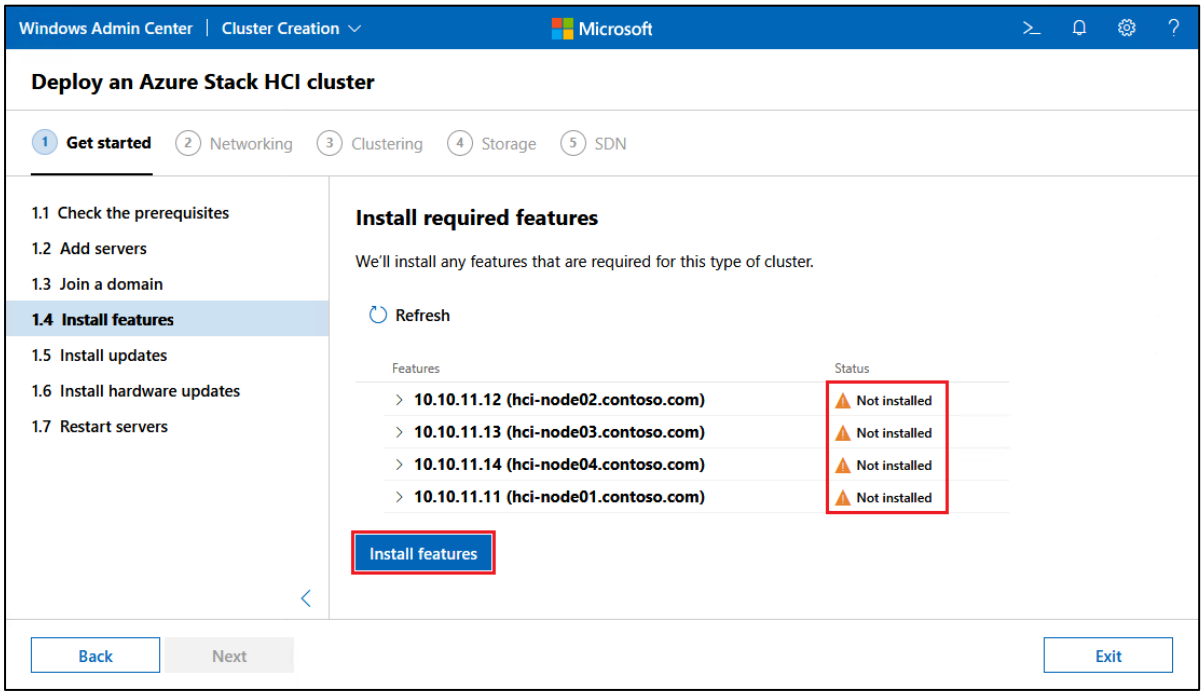

12. Once all required features have been installed on all servers, click Next.

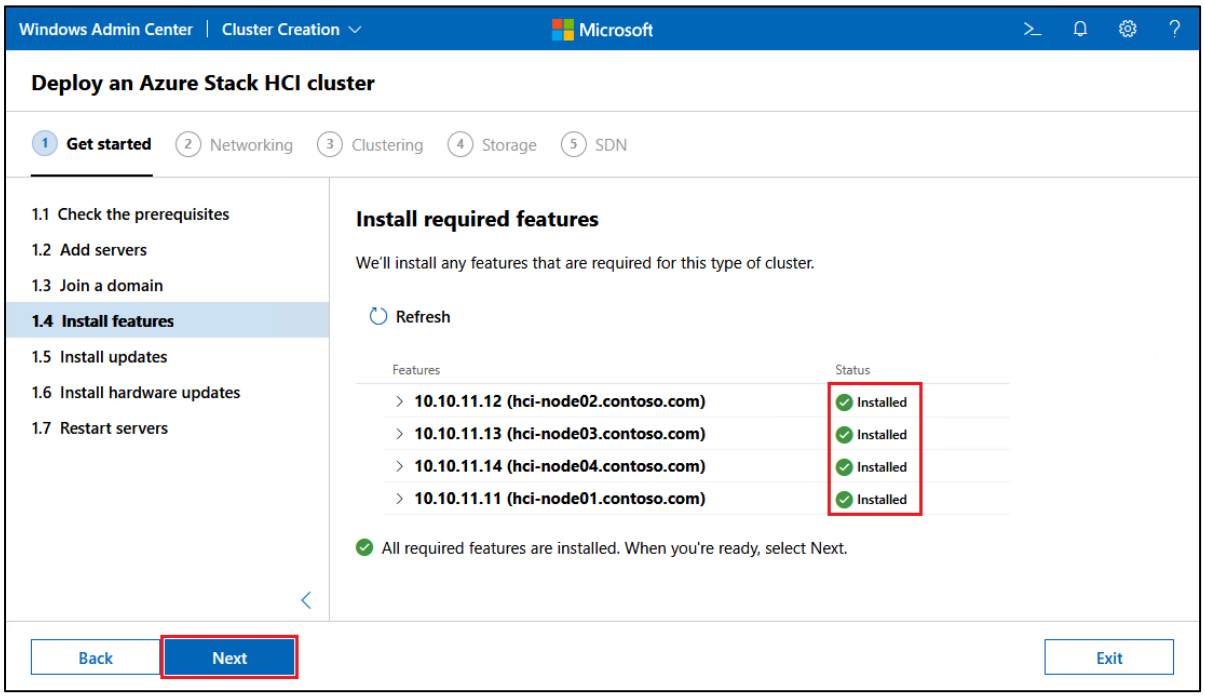

13. On the Install updates page, wait for the page to finish checking for updates. If the Server status column shows “Updates available” for any of the servers, click Install updates.

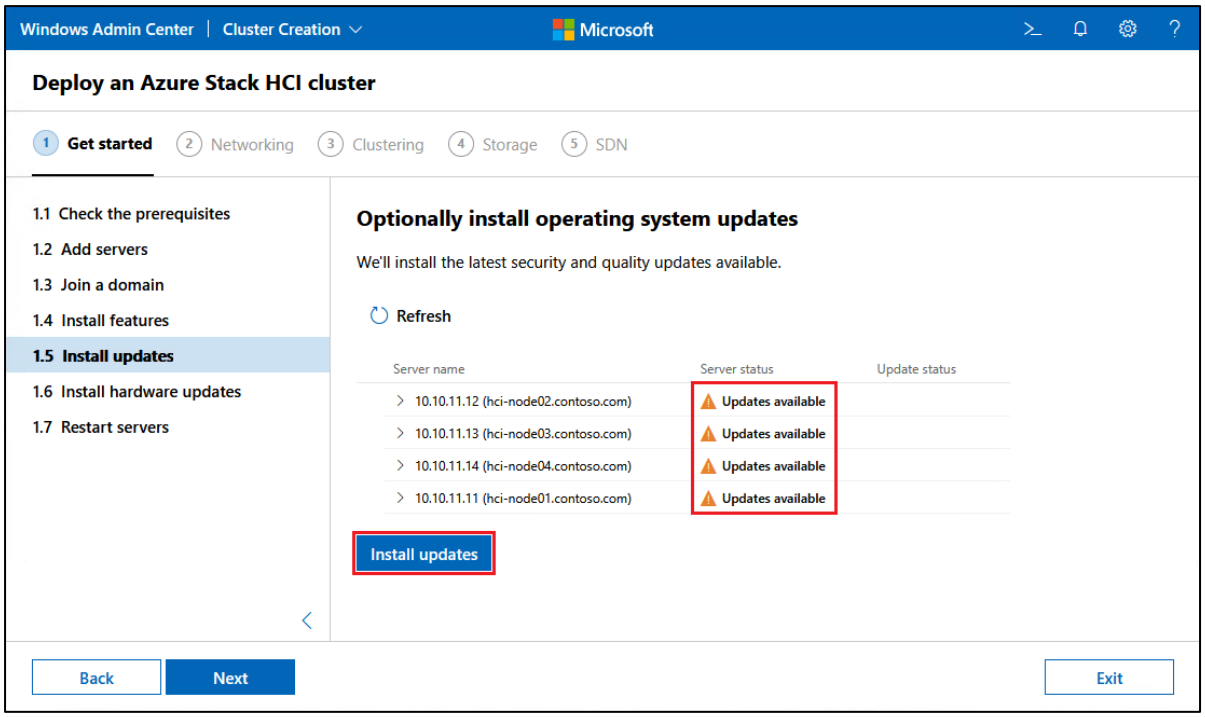

14. The wizard will provide status until all updates are complete. Once the Server status changes to “Ready” for all servers, click Next.

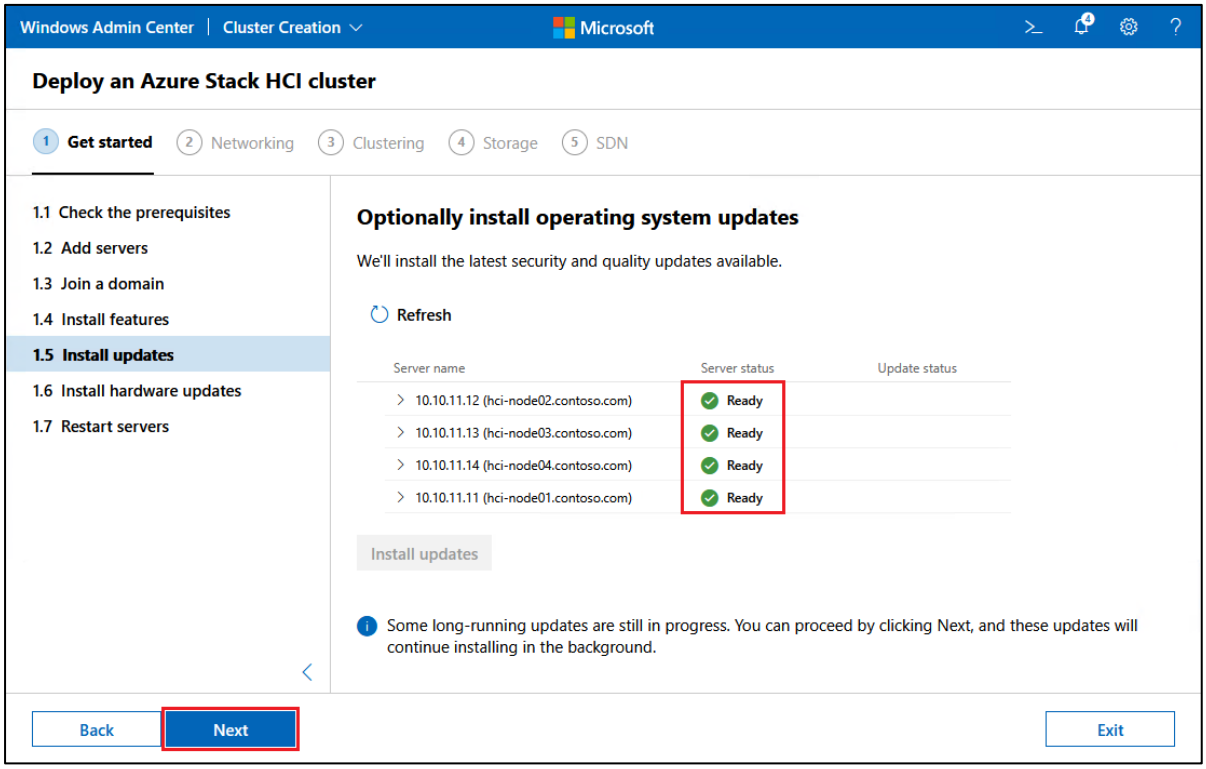

Opening the alert panel in WAC () will show actual update installation status, showing “Installing updates” for each node until each has completed, as shown in the following screen capture.

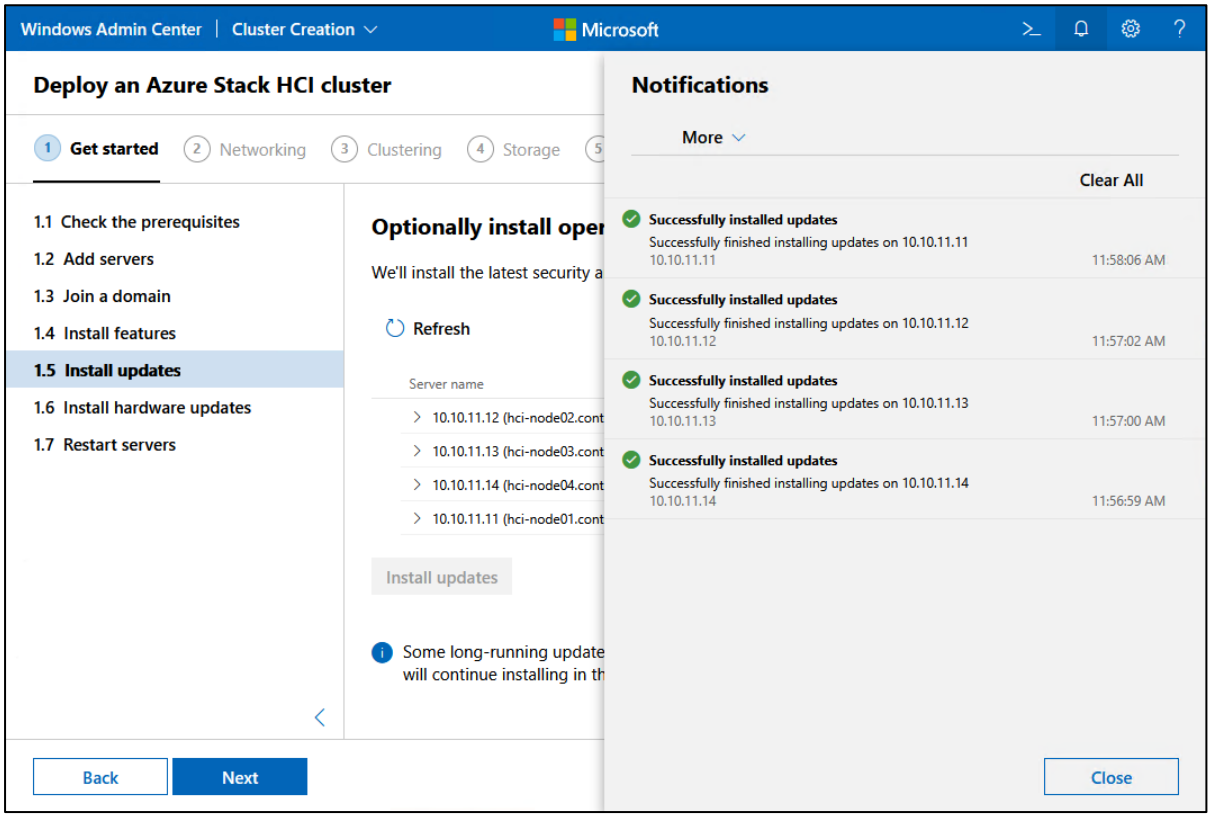

15. On the Install hardware updates page, wait for the “Checking a few things…” message to clear and then click Get updates.

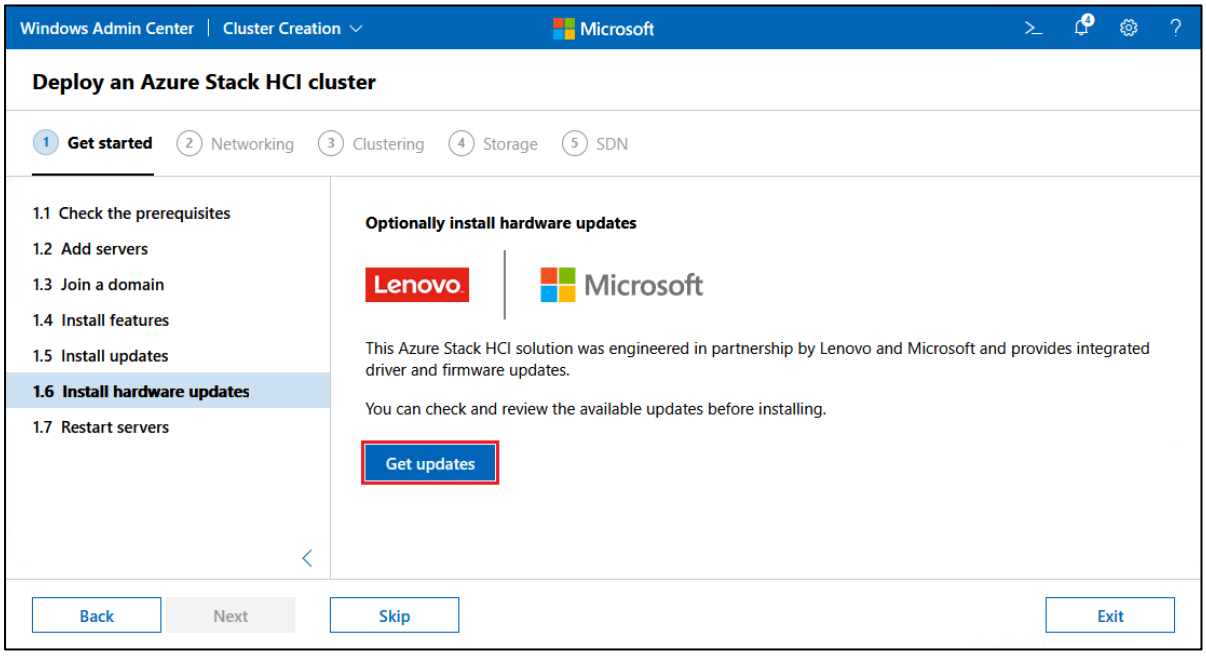

16. This portion of the process is handled by LXCI for WAC and will ensure that firmware and device drivers on all servers will be updated according to the current ThinkAgile MX Best Recipe. Before proceeding, Native OS Management must be enabled in LXCI for WAC. To do this, click on the ellipsis (…) in the upper right corner of the Lenovo XClarity Integrator panel.

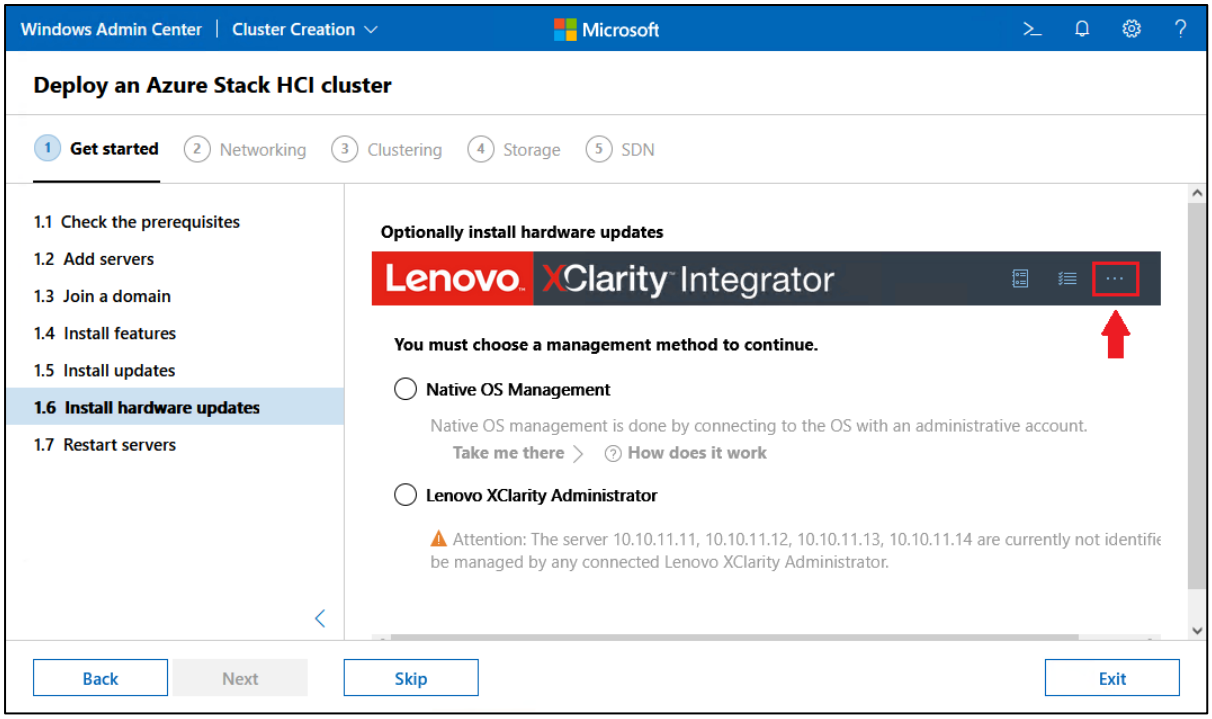

17. Next, click Settings.

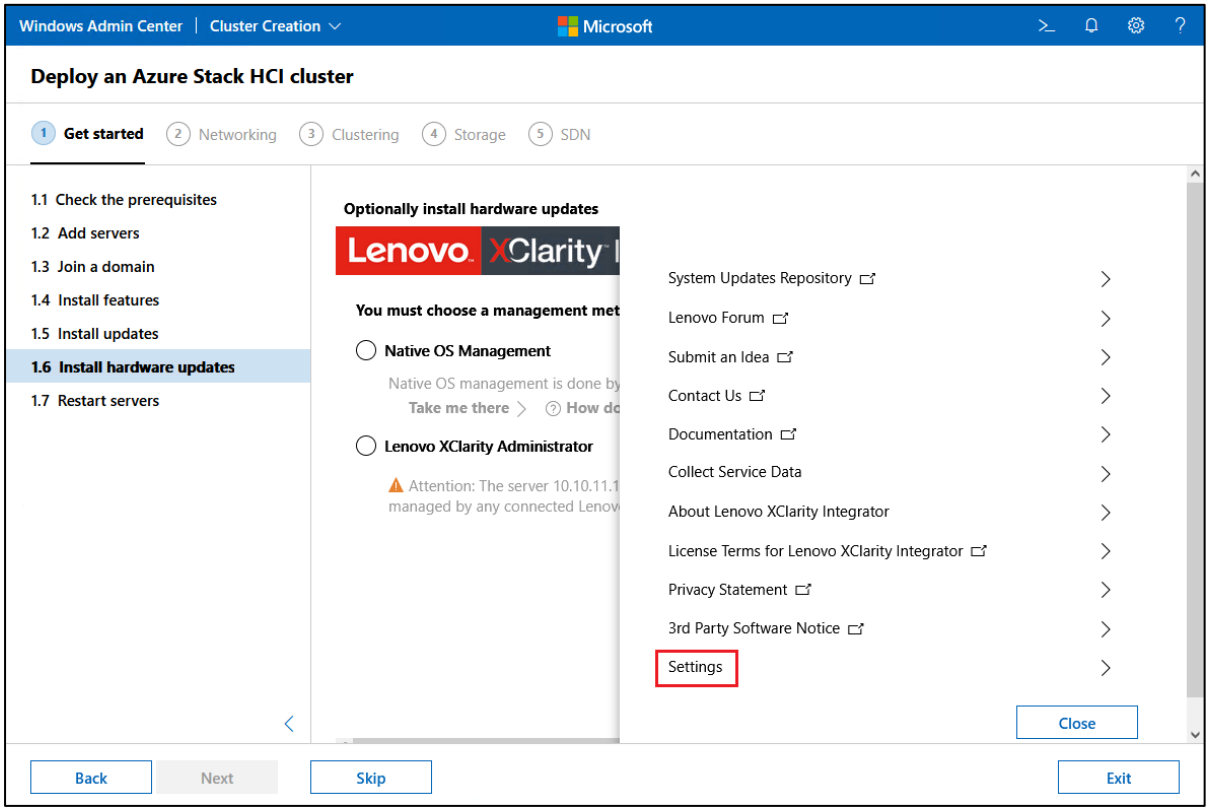

18. In the Settings panel, click Native OS Management in the left column, select the Yes radio button, click Apply, and then click Close.

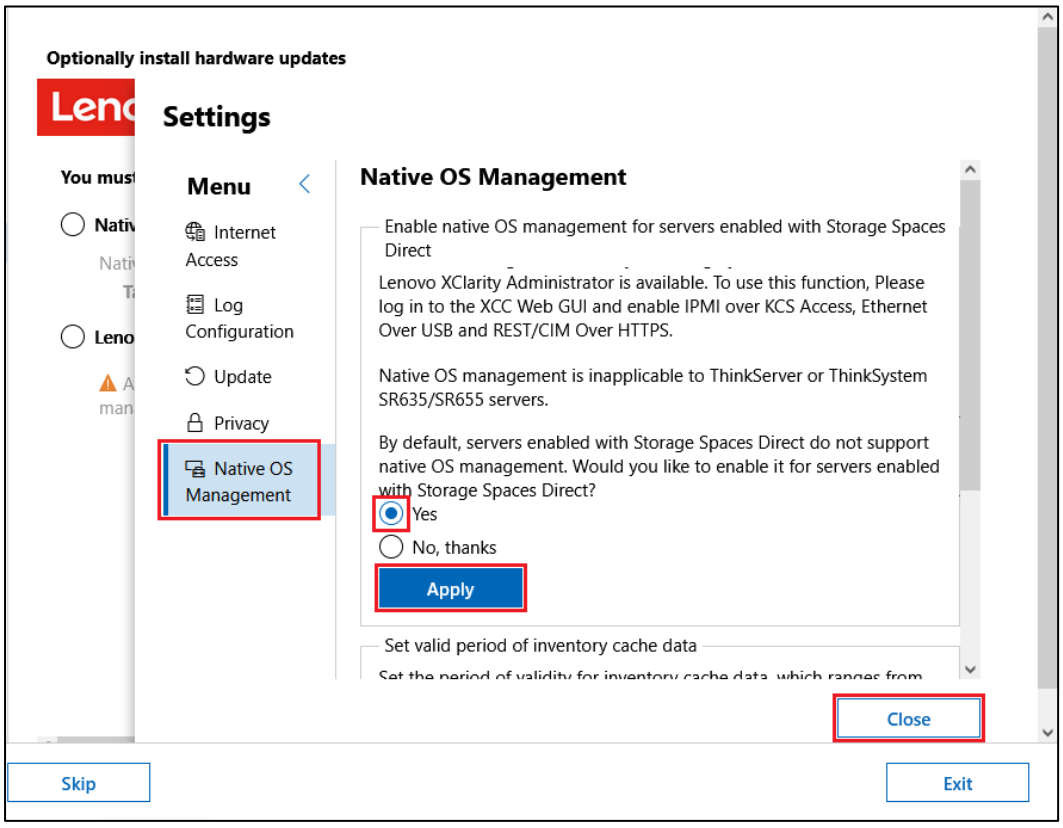

19. Back in the Lenovo XClarity panel, select the Native OS Management radio button and the click Take me there.

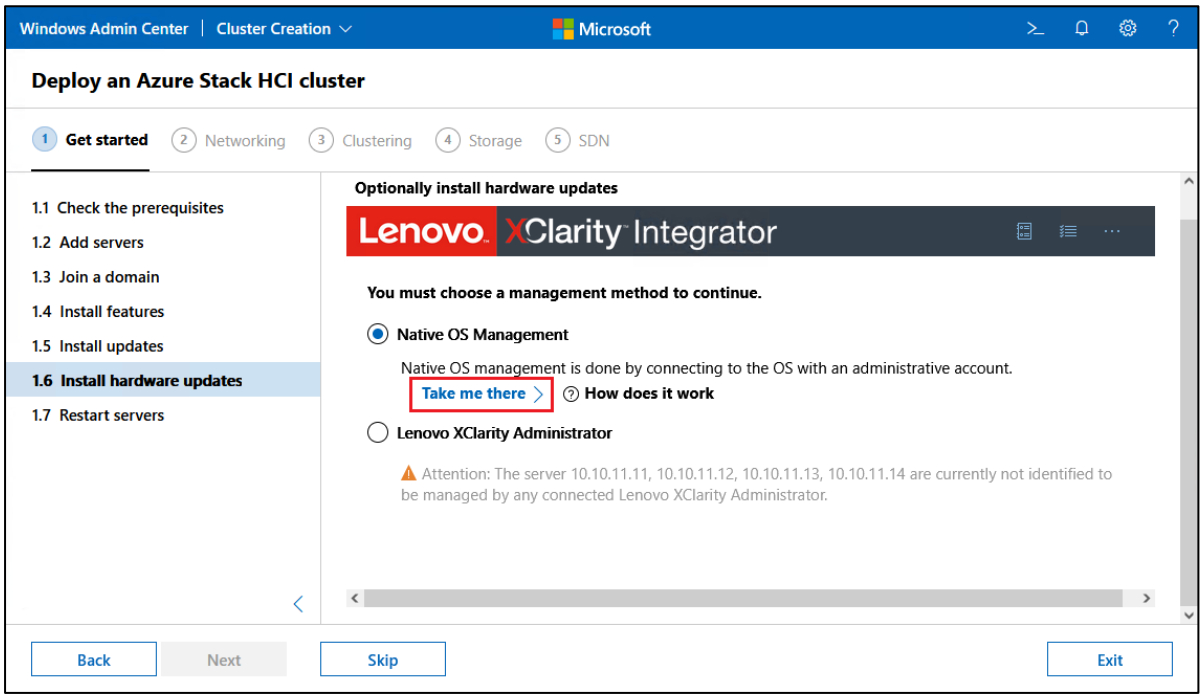

20. In the window that appears, select the Do not show this dialog in the future checkbox if desired, and then click OK.

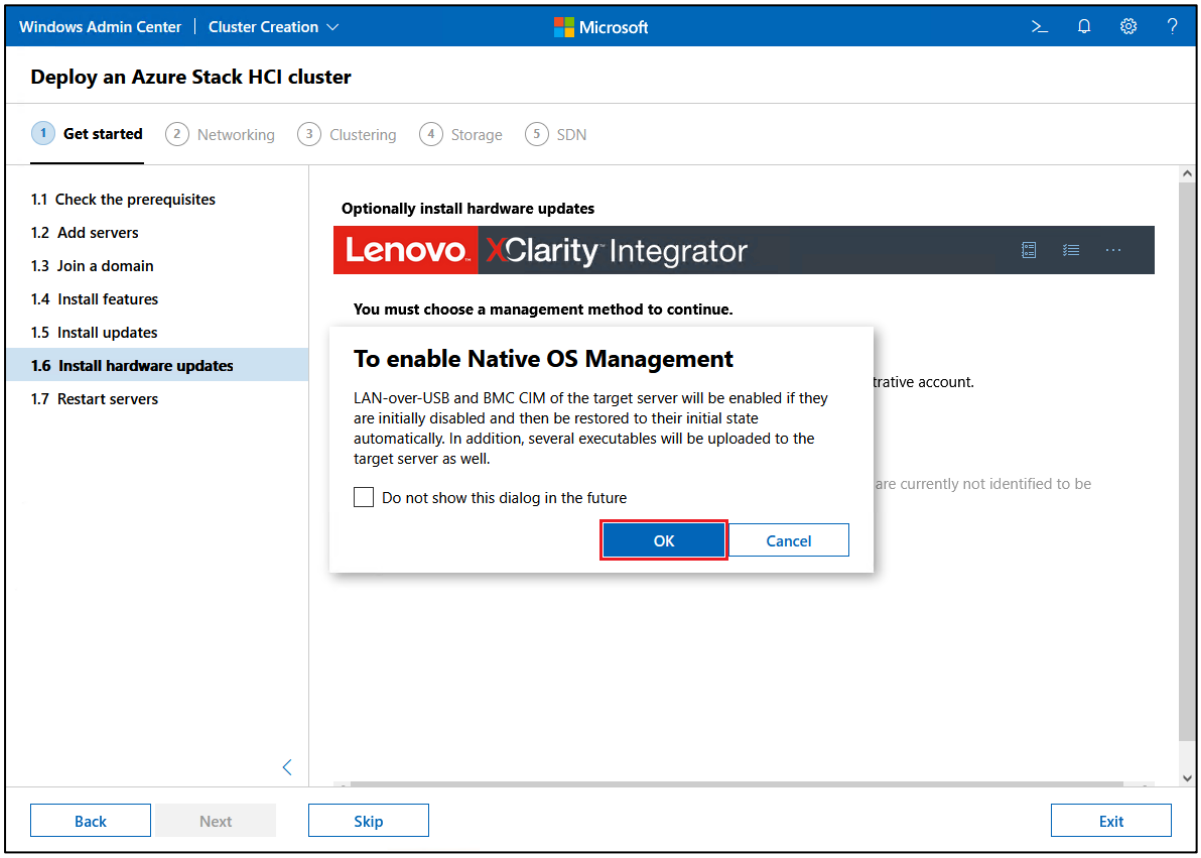

21. Wait for background activity to complete until the LXCI Method Selection panel is fully displayed and then select the Update with Best Recipe (firmware/driver) radio button.

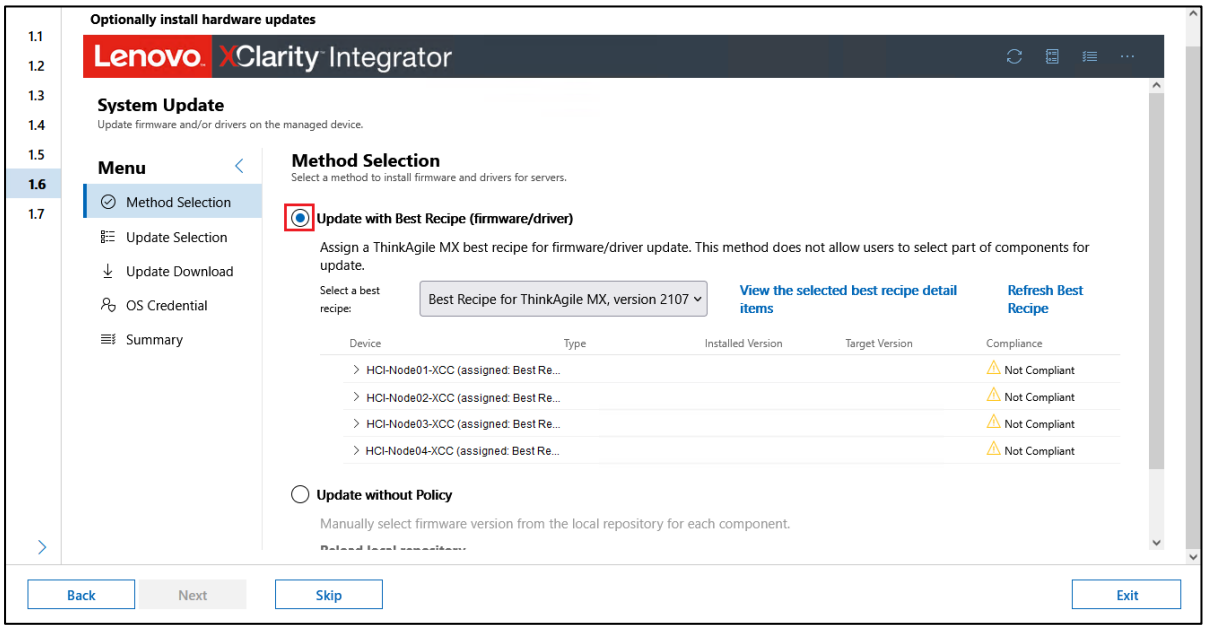

22. Wait for the panel to become fully available, which might take a few minutes, and then use the Select a best recipe dropdown list to select the latest Best Recipe shown in the list. Once the current Best Recipe has been selected, click Next (you might need to scroll down).

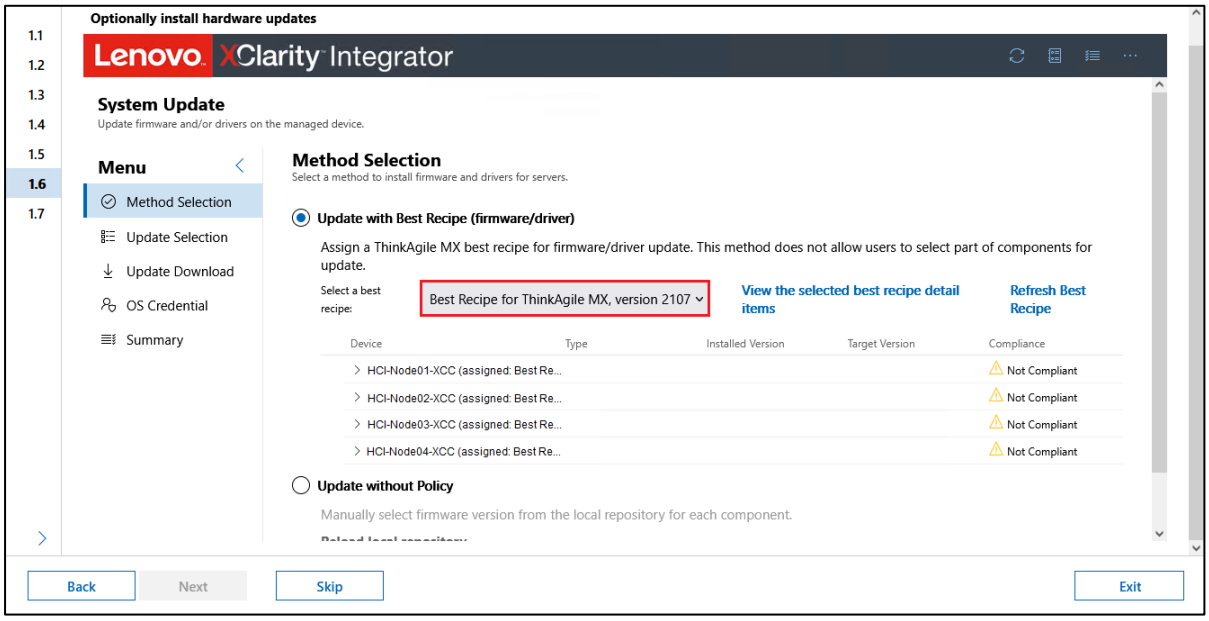

23. In the LXCI Update Selection panel, review the current and target firmware and driver versions if desired. You can click the “<” or “>” icons to show more or less of a panel or sub-panel. Once finished reviewing the pending updates, click Next.

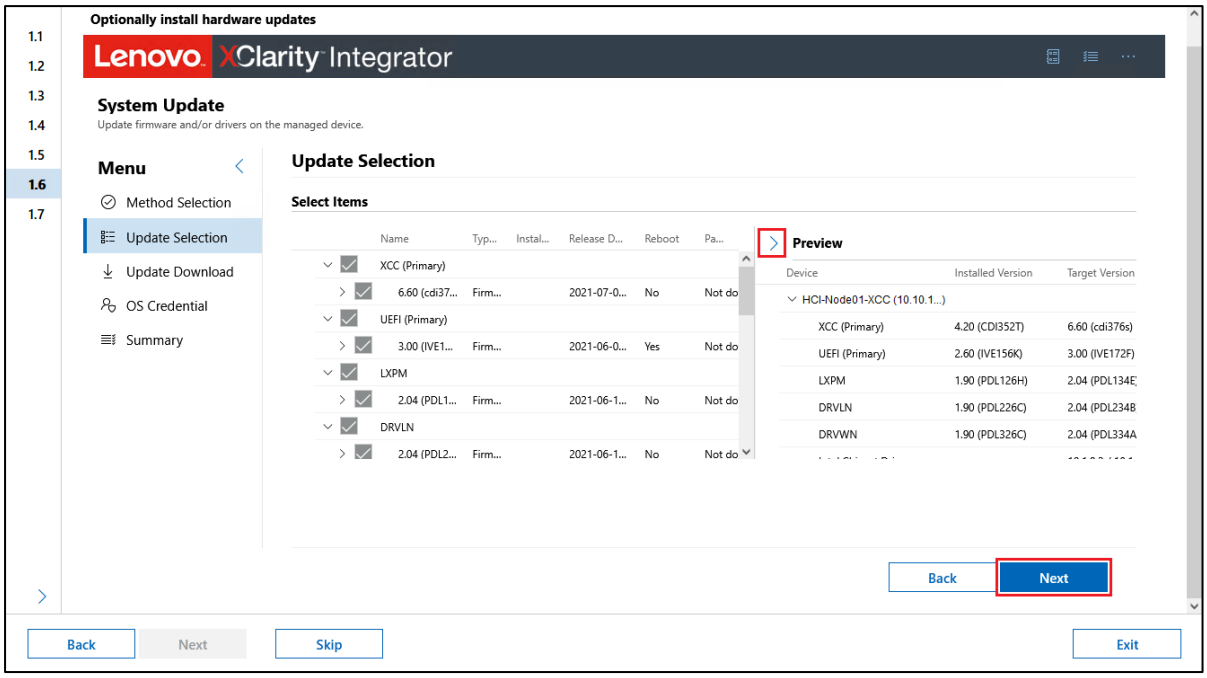

24. In the LXCI Update Download panel, click Download to begin downloading all firmware and driver update packages.

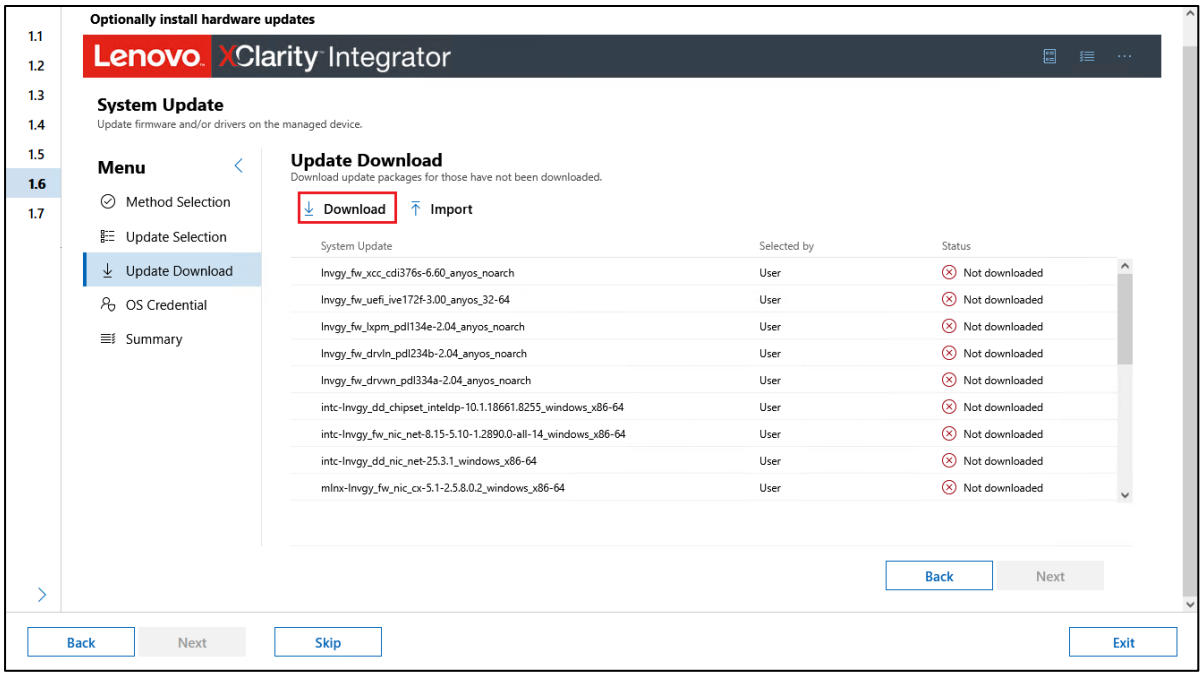

25. Depending on network bandwidth, it might take a few minutes to download all the update packages. Once the LXCI Update Download panel updates to indicate that all update packages have been downloaded, click Next.

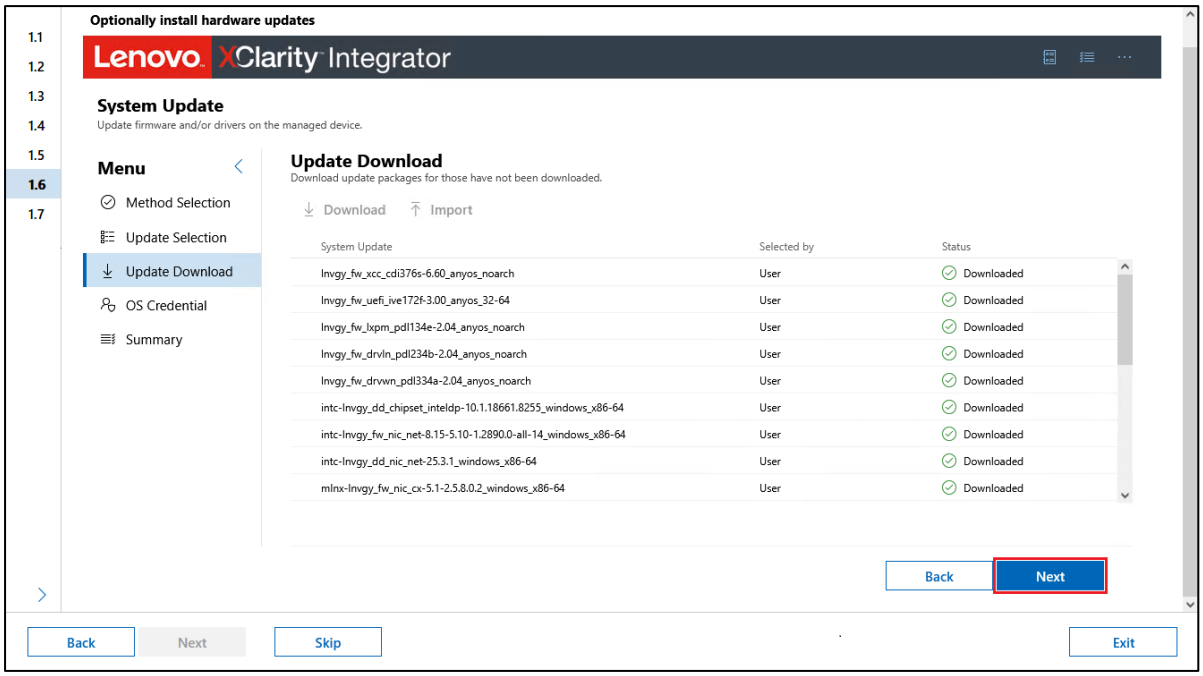

26. In the LXCI OS Credential panel, click on each Credential Needed link to enter appropriate credentials for each server.

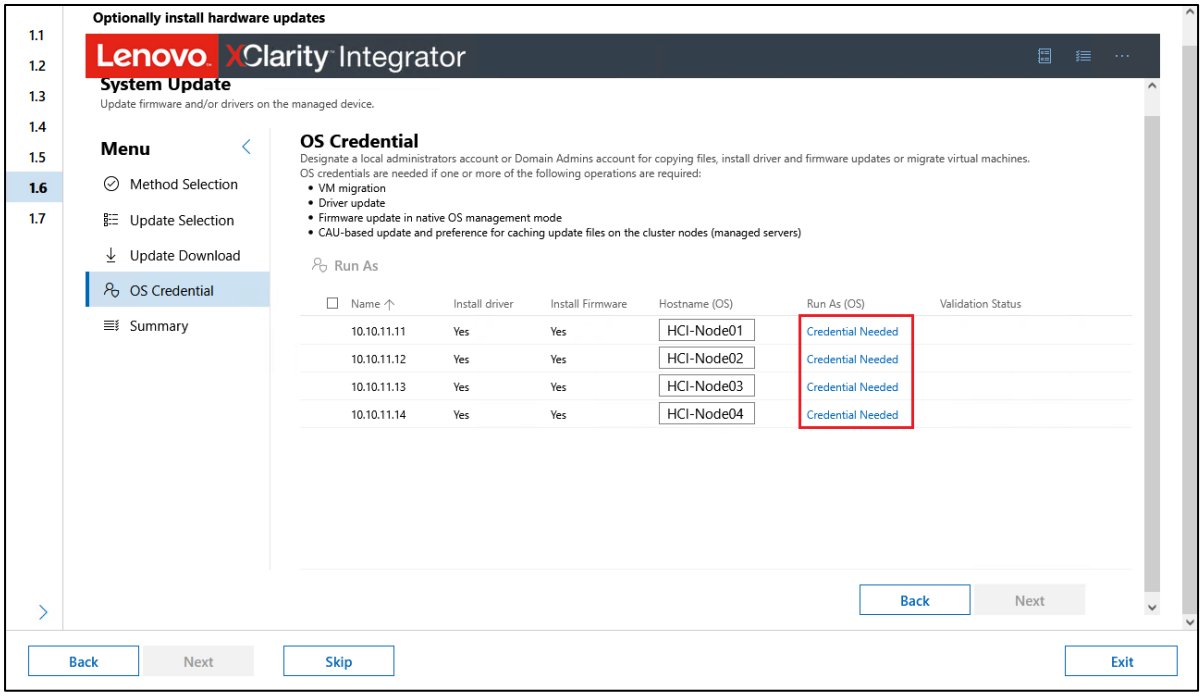

27. For each server that will be added to the cluster, enter an appropriate Username and Password and then click Submit.

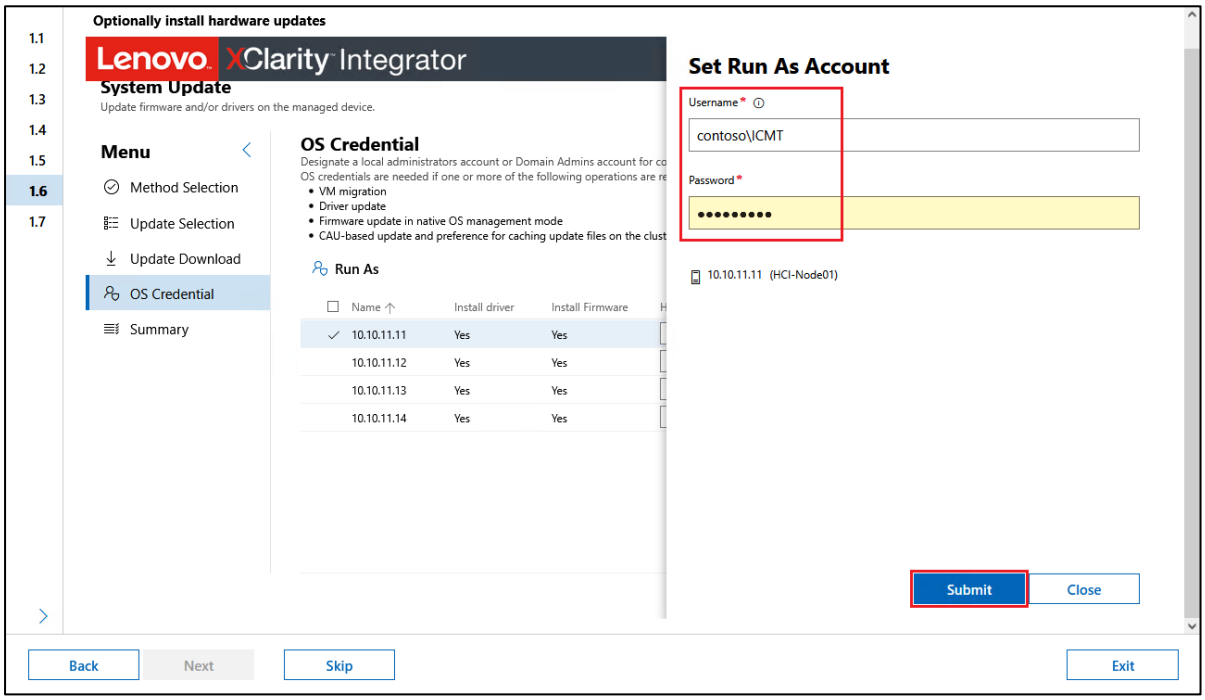

28. Once the credentials are tested against the system, click Submit and ignore errors.

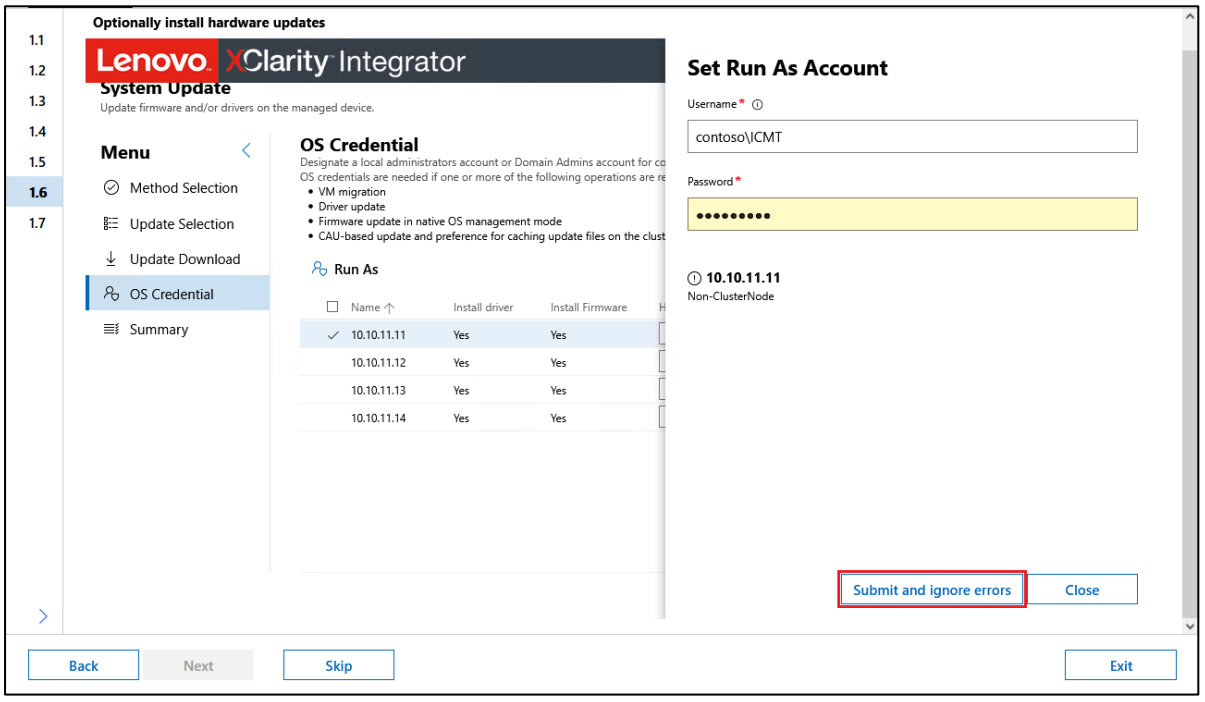

29. Once credentials have been entered for all systems, click Next.

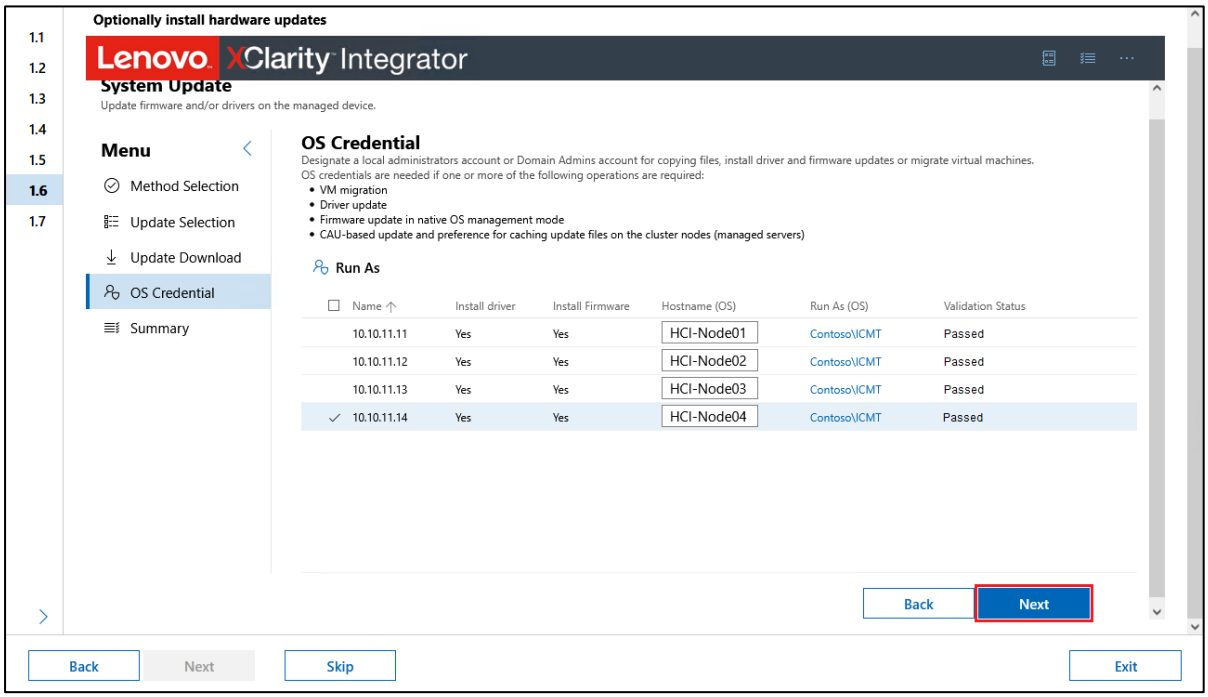

30. Review the Summary panel to ensure that all is as expected, and then click Submit.

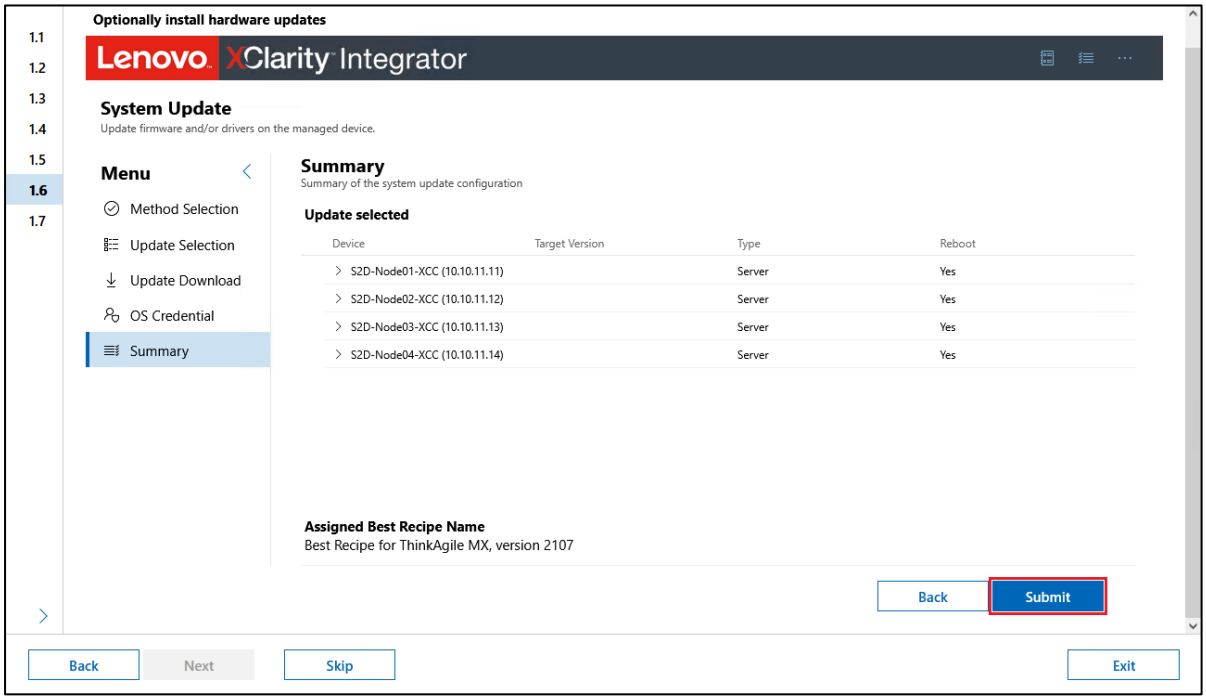

31. Click Yes in the Warning dialog that opens to begin the system update job.
32. In the System Update Progress panel you can watch the progress of the system update job. To get additional details, click the twisty next to any system in either column.

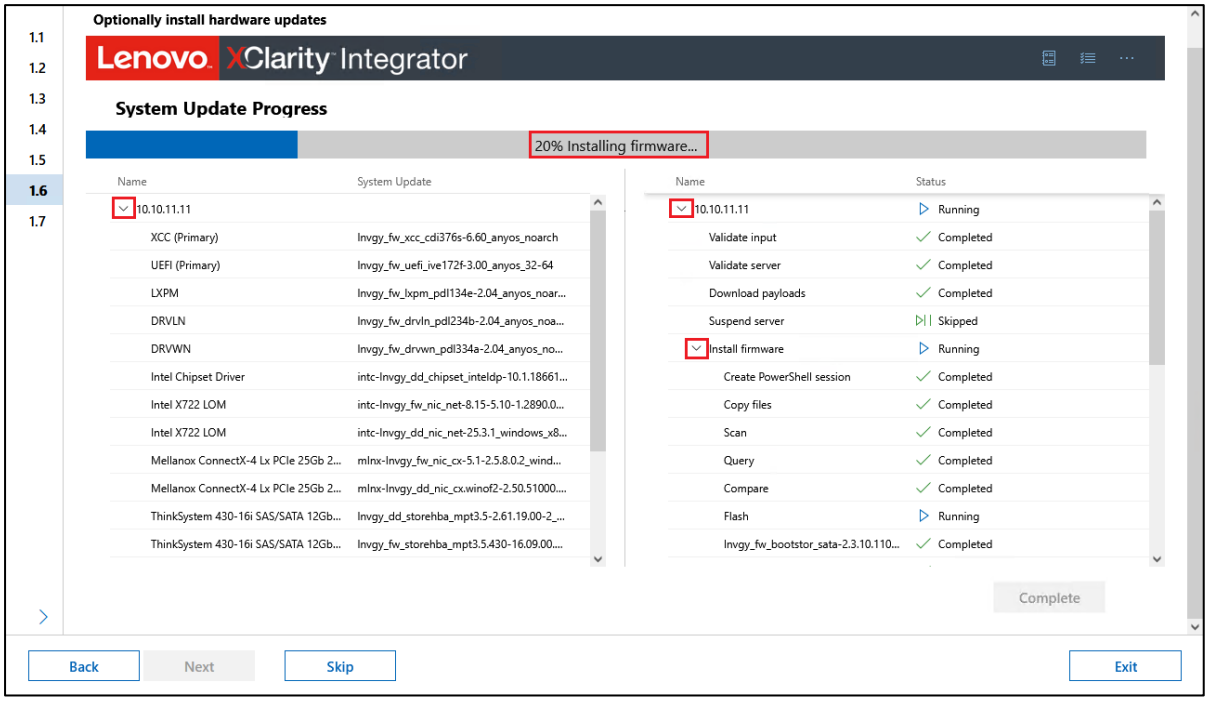

33. Once the system update process completes, click Next.

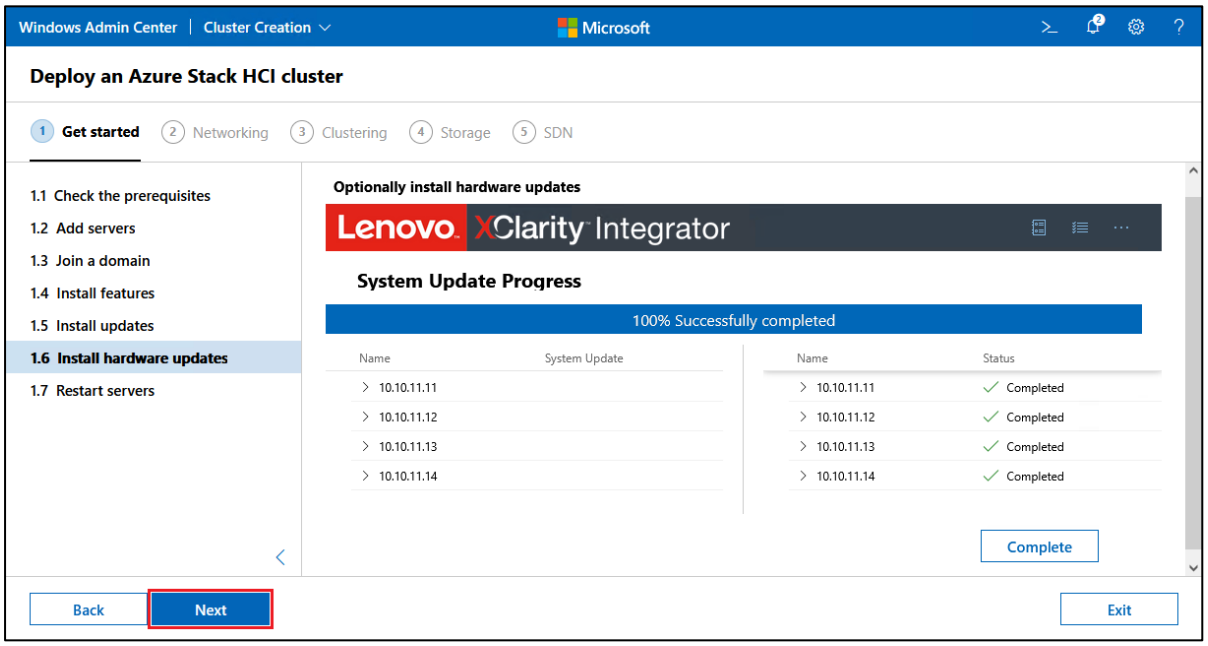

34. On the Restart servers page, click Restart servers.

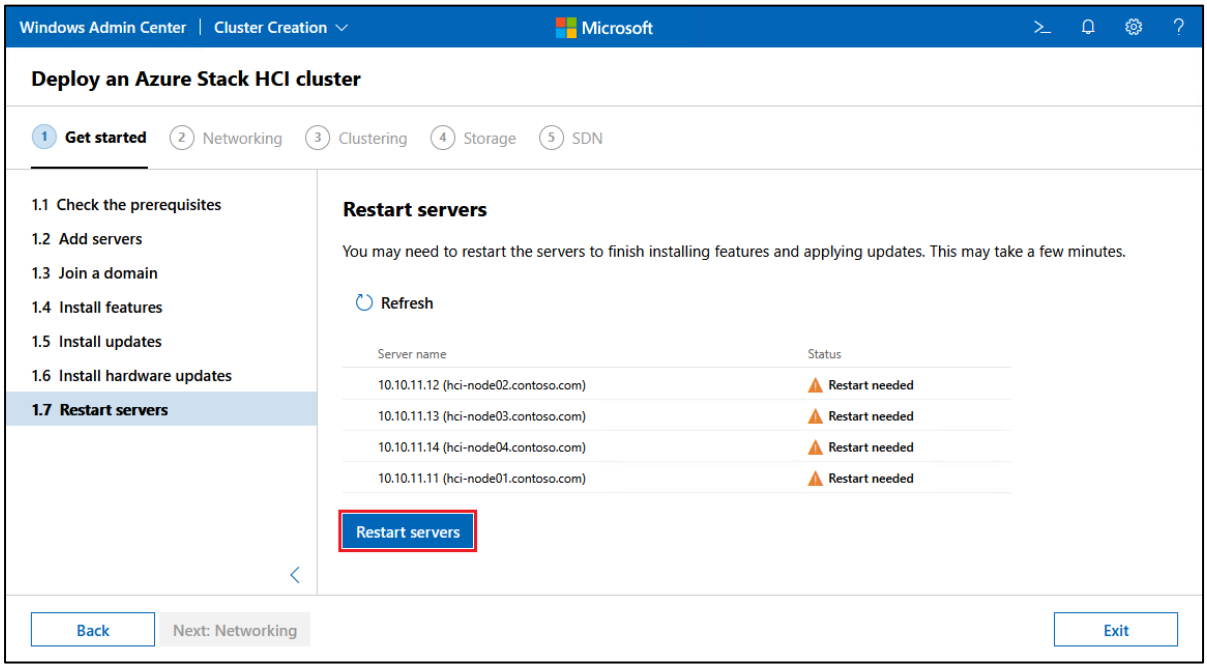

35. After the servers have been rebooted and their Status shows “Ready” click Next: Networking.

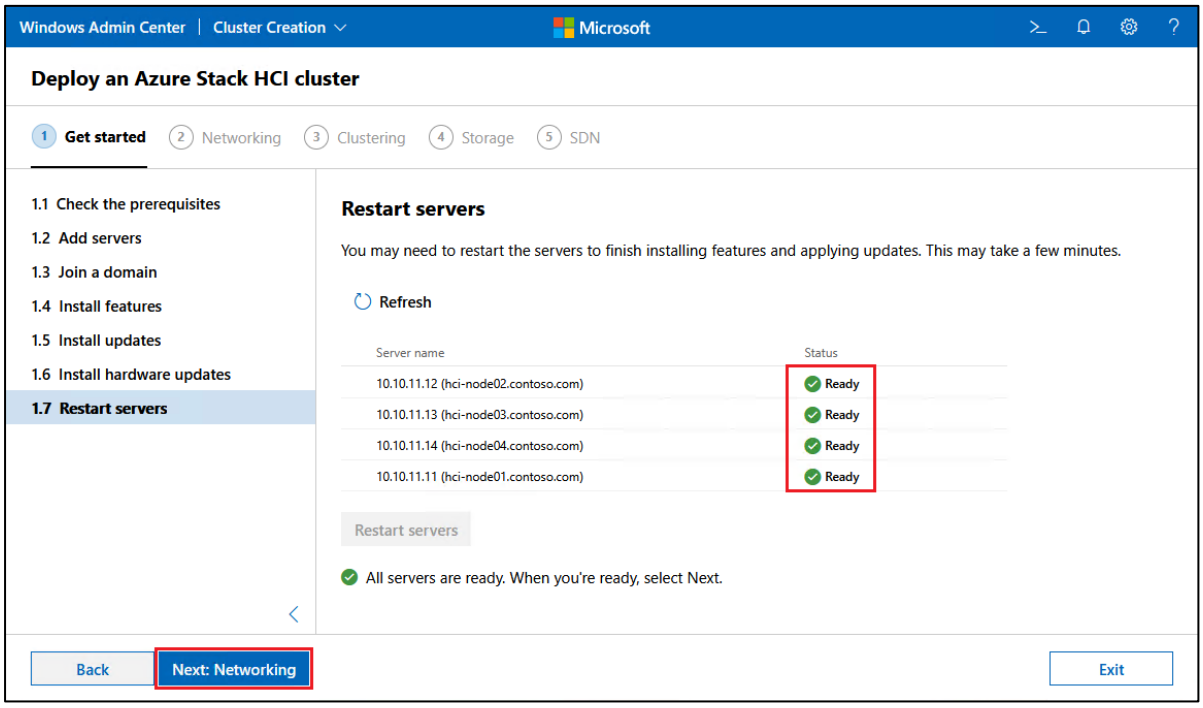

36. In the Check the network adapters panel, we disable the LAN Over USB interface on each node to avoid issues later (this interface will be automatically enabled when needed and then disabled when

no longer needed). For each node, select the “IBM USB Remote NDIS Network Device” and then click Disable. Once this interface has been disabled on all nodes, click Next.

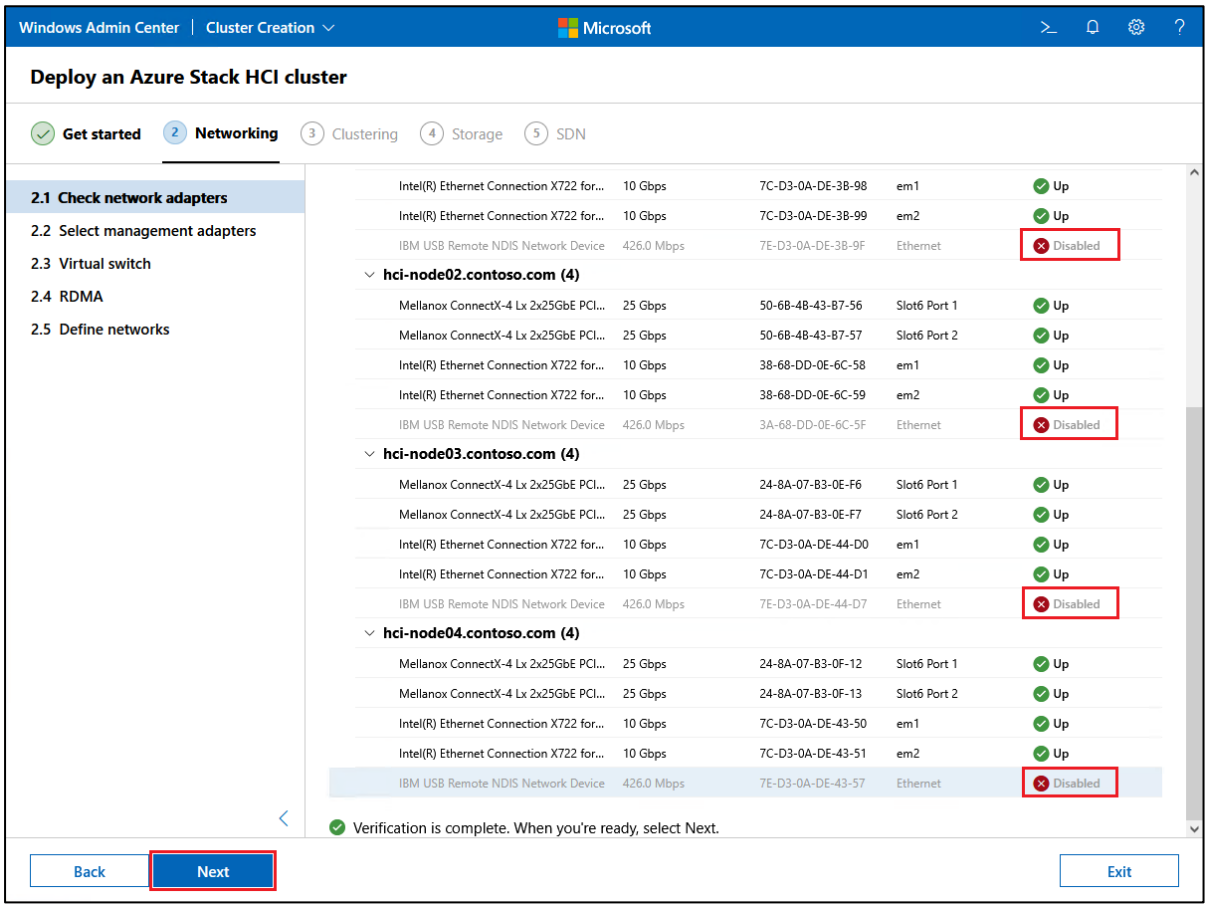

37. In the Select the adapters to use for management panel, choose the appropriate model at the top of the panel. As previously noted, we have configured two LOM ports in each server to carry management traffic in our example deployment. Once all management interfaces have been selected for all nodes, click Apply and test.

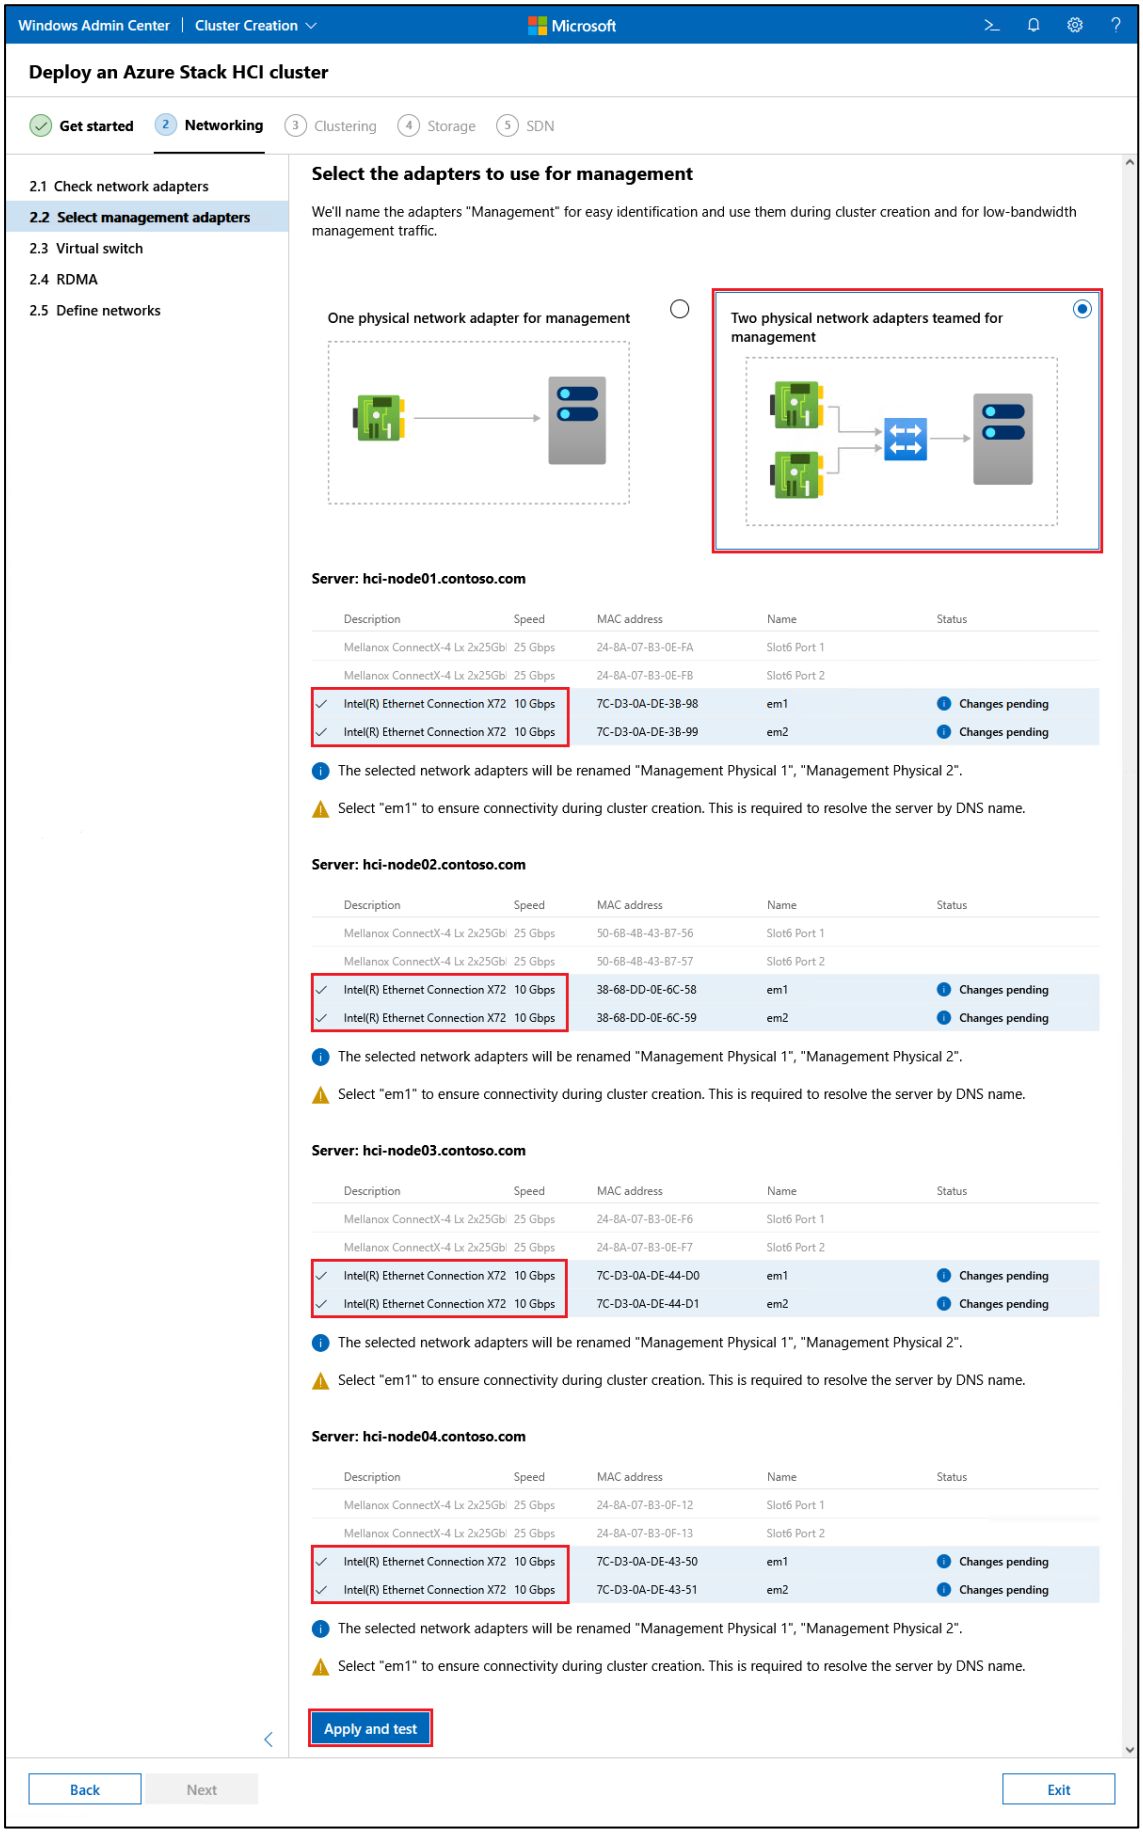
Windows Admin Center | Cluster Creation
Microsoft
Deploy an Azure Stack HCI cluster
Get started
2 Networking
3 Clustering
4 Storage
5 SDN
2.1 Check network adapters
2.2 Select management adapters
2.3 Virtual switch
2.4 RDMA
2.5 Define networks
Select the adapters to use for management
We'll name the adapters "Management" for easy identification and use them during cluster creation and for low-bandwidth management traffic.
One physical network adapter for management
Two physical network adapters teamed for management
Server: hci-node01.contoso.com
Description Speed MAC address Name Status
Mellanox ConnectX-4 Lx 2x25Gb 25 Gbps 24-8A-07-B3-0E-FA Slot6 Port 1
Mellanox ConnectX-4 Lx 2x25Gb 25 Gbps 24-8A-07-B3-0E-FB Slot6 Port 2
Intel(R) Ethernet Connection X72 10 Gbps 7C-D3-0A-DE-3B-98 em1 Changes pending
Intel(R) Ethernet Connection X72 10 Gbps 7C-D3-0A-DE-3B-99 em2 Changes pending
The selected network adapters will be renamed "Management Physical 1", "Management Physical 2".
Select "em1" to ensure connectivity during cluster creation. This is required to resolve the server by DNS name.
Server: hci-node02.contoso.com
Description Speed MAC address Name Status
Mellanox ConnectX-4 Lx 2x25Gb 25 Gbps 50-6B-4B-43-B7-56 Slot6 Port 1
Mellanox ConnectX-4 Lx 2x25Gb 25 Gbps 50-6B-4B-43-B7-57 Slot6 Port 2
Intel(R) Ethernet Connection X72 10 Gbps 38-68-DD-0E-6C-58 em1 Changes pending
Intel(R) Ethernet Connection X72 10 Gbps 38-68-DD-0E-6C-59 em2 Changes pending
The selected network adapters will be renamed "Management Physical 1", "Management Physical 2".
Select "em1" to ensure connectivity during cluster creation. This is required to resolve the server by DNS name.
Server: hci-node03.contoso.com
Description Speed MAC address Name Status
Mellanox ConnectX-4 Lx 2x25Gb 25 Gbps 24-8A-07-B3-0E-F6 Slot6 Port 1
Mellanox ConnectX-4 Lx 2x25Gb 25 Gbps 24-8A-07-B3-0E-F7 Slot6 Port 2
Intel(R) Ethernet Connection X72 10 Gbps 7C-D3-0A-DE-44-D0 em1 Changes pending
Intel(R) Ethernet Connection X72 10 Gbps 7C-D3-0A-DE-44-D1 em2 Changes pending
The selected network adapters will be renamed "Management Physical 1", "Management Physical 2".
Select "em1" to ensure connectivity during cluster creation. This is required to resolve the server by DNS name.
Server: hci-node04.contoso.com
Description Speed MAC address Name Status
Mellanox ConnectX-4 Lx 2x25Gb 25 Gbps 24-8A-07-B3-0F-12 Slot6 Port 1
Mellanox ConnectX-4 Lx 2x25Gb 25 Gbps 24-8A-07-B3-0F-13 Slot6 Port 2
Intel(R) Ethernet Connection X72 10 Gbps 7C-D3-0A-DE-43-50 em1 Changes pending
Intel(R) Ethernet Connection X72 10 Gbps 7C-D3-0A-DE-43-51 em2 Changes pending
The selected network adapters will be renamed "Management Physical 1", "Management Physical 2".
Select "em1" to ensure connectivity during cluster creation. This is required to resolve the server by DNS name.
Apply and test
Back
Next
Exit

38. Click Yes in the Confirm creation of teamed adapter dialog that opens to begin the process to create the team and test connectivity.
39. Status will be updated in WAC as the required changes are made and connectivity between the systems is verified.

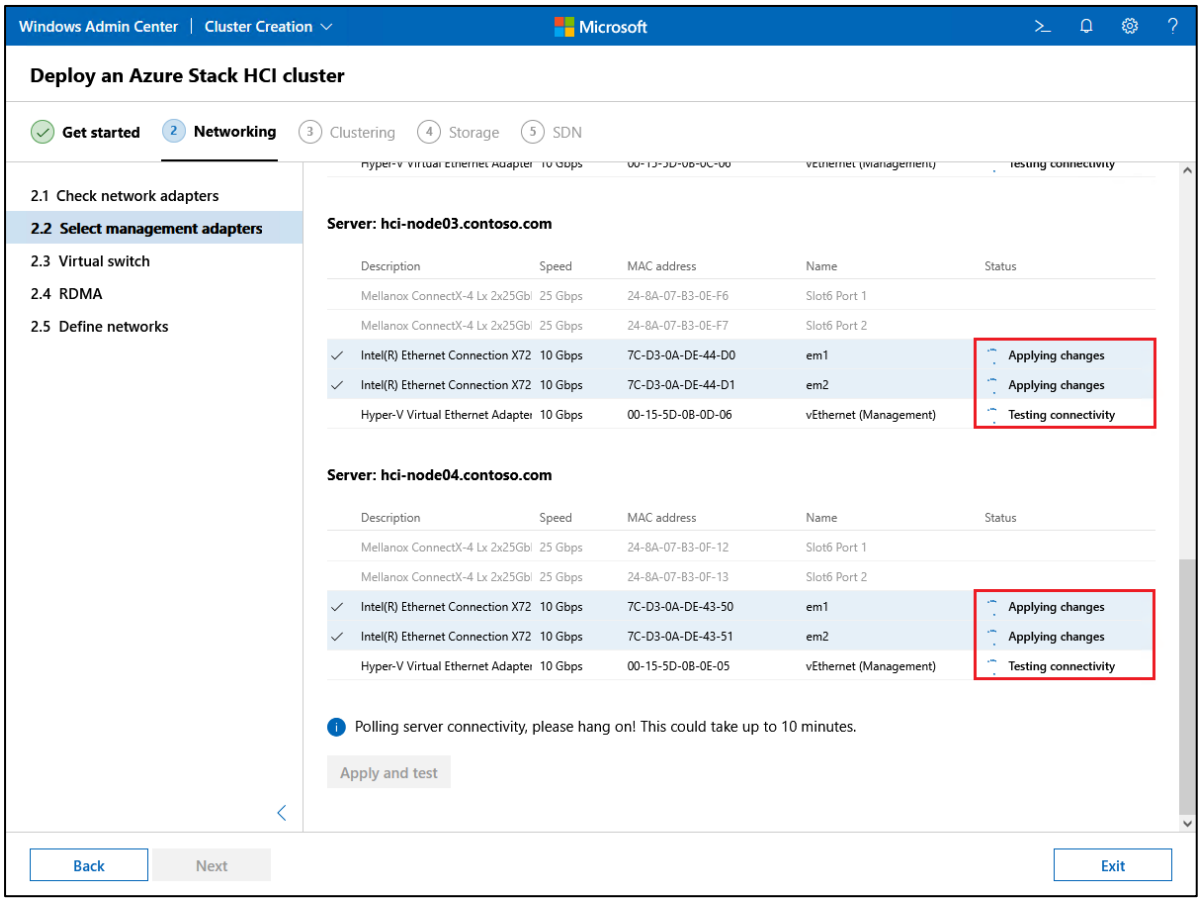

40. Once all systems show Status of “Changes applied” click Next.

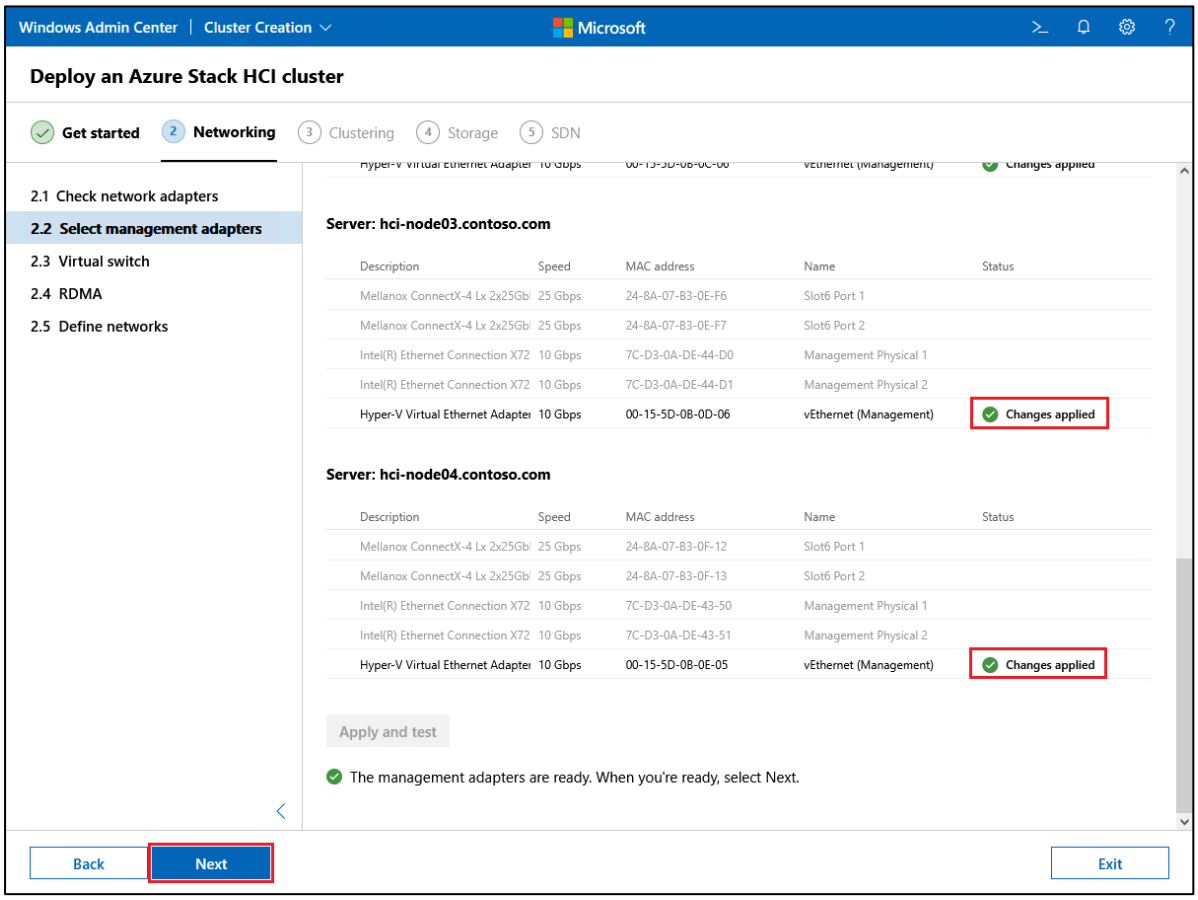

41. In the Virtual switch panel, select your preferred configuration from the available options. For our example, using a single dual-port Mellanox NIC in each server results in a single available configuration.
42. Once a configuration is selected, work through the panel, selecting each network interface that will carry East-West storage traffic. Near the bottom of the panel, open the Advanced twisty if desired to change additional settings. We recommend leaving all the defaults except for the Virtual switch name, which can be changed to something more recognizable in your environment (we leave the default switch name of “ConvergedSwitch” in our example deployment).

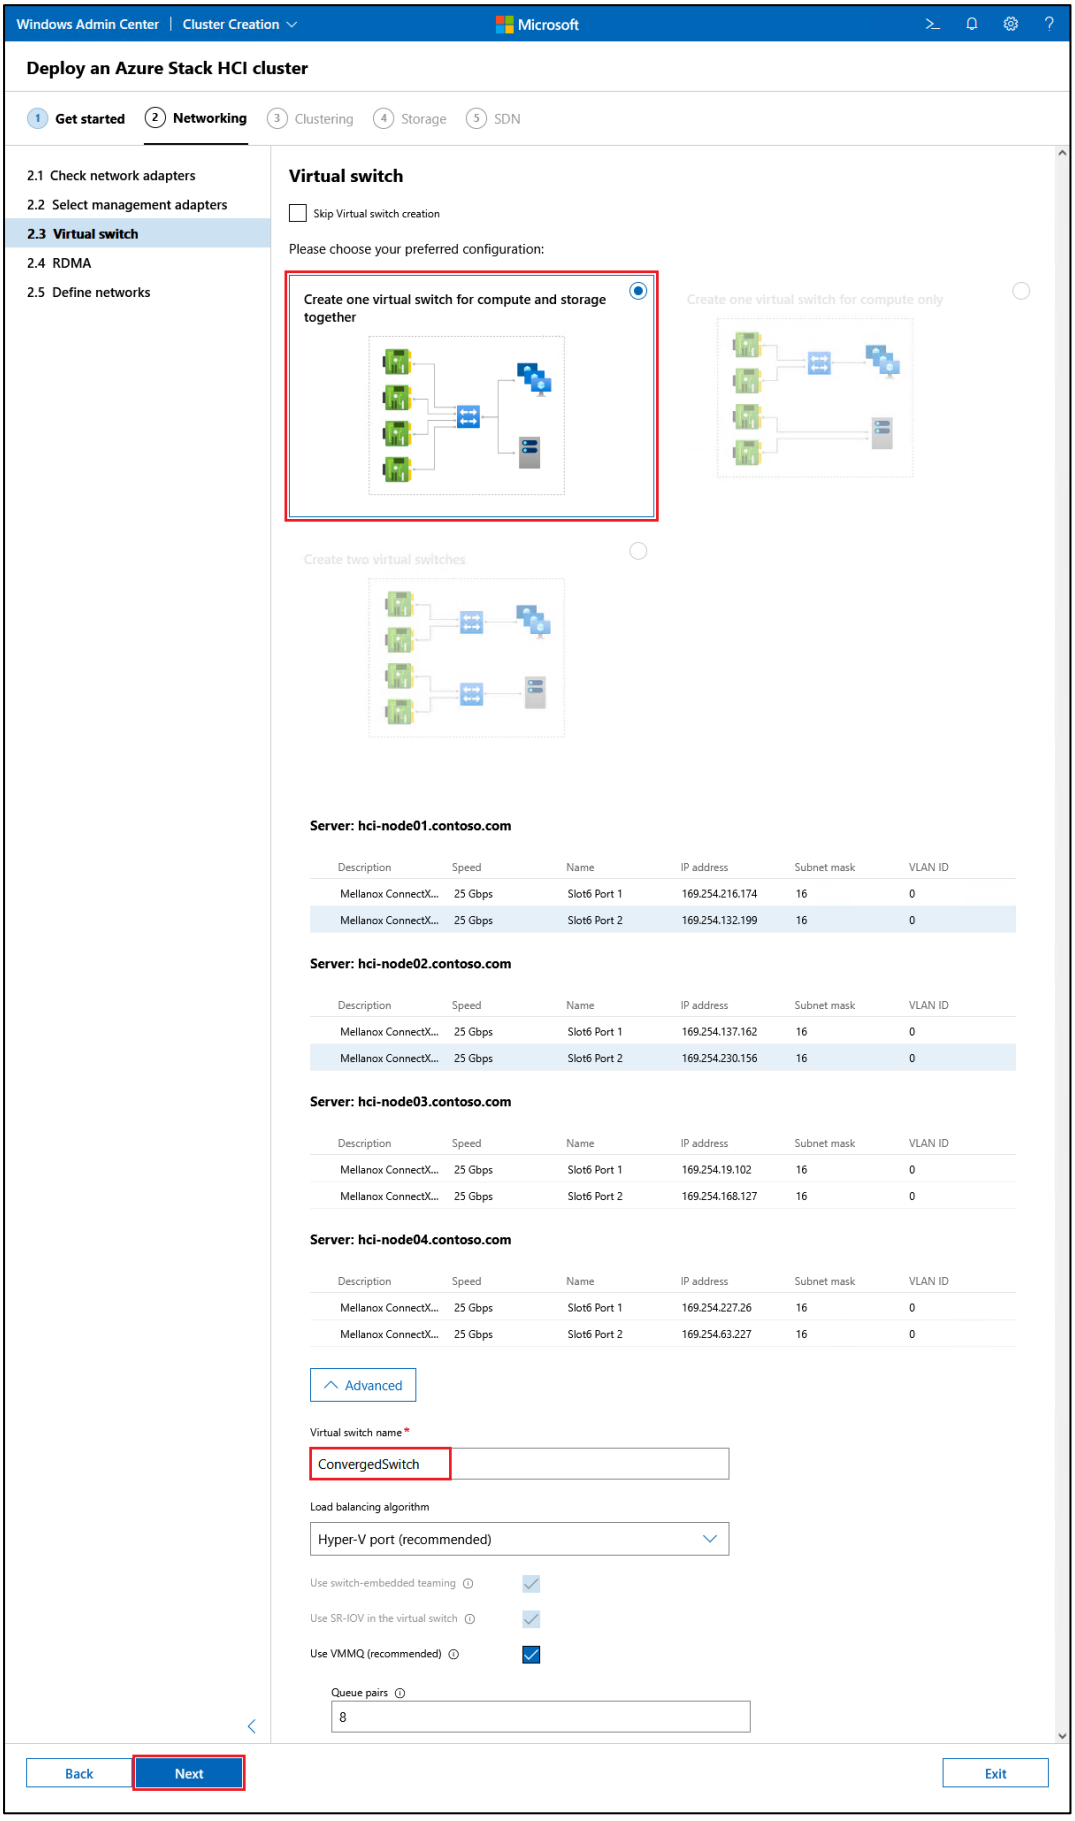
Windows Admin Center | Cluster Creation
Microsoft
Deploy an Azure Stack HCI cluster
1 Get started
2 Networking
3 Clustering
4 Storage
5 SDN
2.1 Check network adapters
2.2 Select management adapters
2.3 Virtual switch
2.4 RDMA
2.5 Define networks
Virtual switch
Skip Virtual switch creation
Please choose your preferred configuration:
Create one virtual switch for compute and storage together
Create one virtual switch for compute only
Create two virtual switches
Server: hci-node01.contoso.com
Description Speed Name IP address Subnet mask VLAN ID
Mellanox ConnectX... 25 Gbps Slot6 Port 1 169.254.216.174 16 0
Mellanox ConnectX... 25 Gbps Slot6 Port 2 169.254.132.199 16 0
Server: hci-node02.contoso.com
Description Speed Name IP address Subnet mask VLAN ID
Mellanox ConnectX... 25 Gbps Slot6 Port 1 169.254.137.162 16 0
Mellanox ConnectX... 25 Gbps Slot6 Port 2 169.254.230.156 16 0
Server: hci-node03.contoso.com
Description Speed Name IP address Subnet mask VLAN ID
Mellanox ConnectX... 25 Gbps Slot6 Port 1 169.254.19.102 16 0
Mellanox ConnectX... 25 Gbps Slot6 Port 2 169.254.168.127 16 0
Server: hci-node04.contoso.com
Description Speed Name IP address Subnet mask VLAN ID
Mellanox ConnectX... 25 Gbps Slot6 Port 1 169.254.227.26 16 0
Mellanox ConnectX... 25 Gbps Slot6 Port 2 169.254.63.227 16 0
Advanced
Virtual switch name *
ConvergedSwitch
Load balancing algorithm
Hyper-V port (recommended)
Use switch-embedded teaming
Use SR-IOV in the virtual switch
Use VMMQ (recommended)
Queue pairs
8
Back
Next
Exit

43. In the Optionally configure RDMA panel, select the Configure RDMA (Recommended) checkbox. If desired, open the Advanced twisty to expose Quality of Service (QoS) settings. Make any changes required for your environment and then click Apply changes.

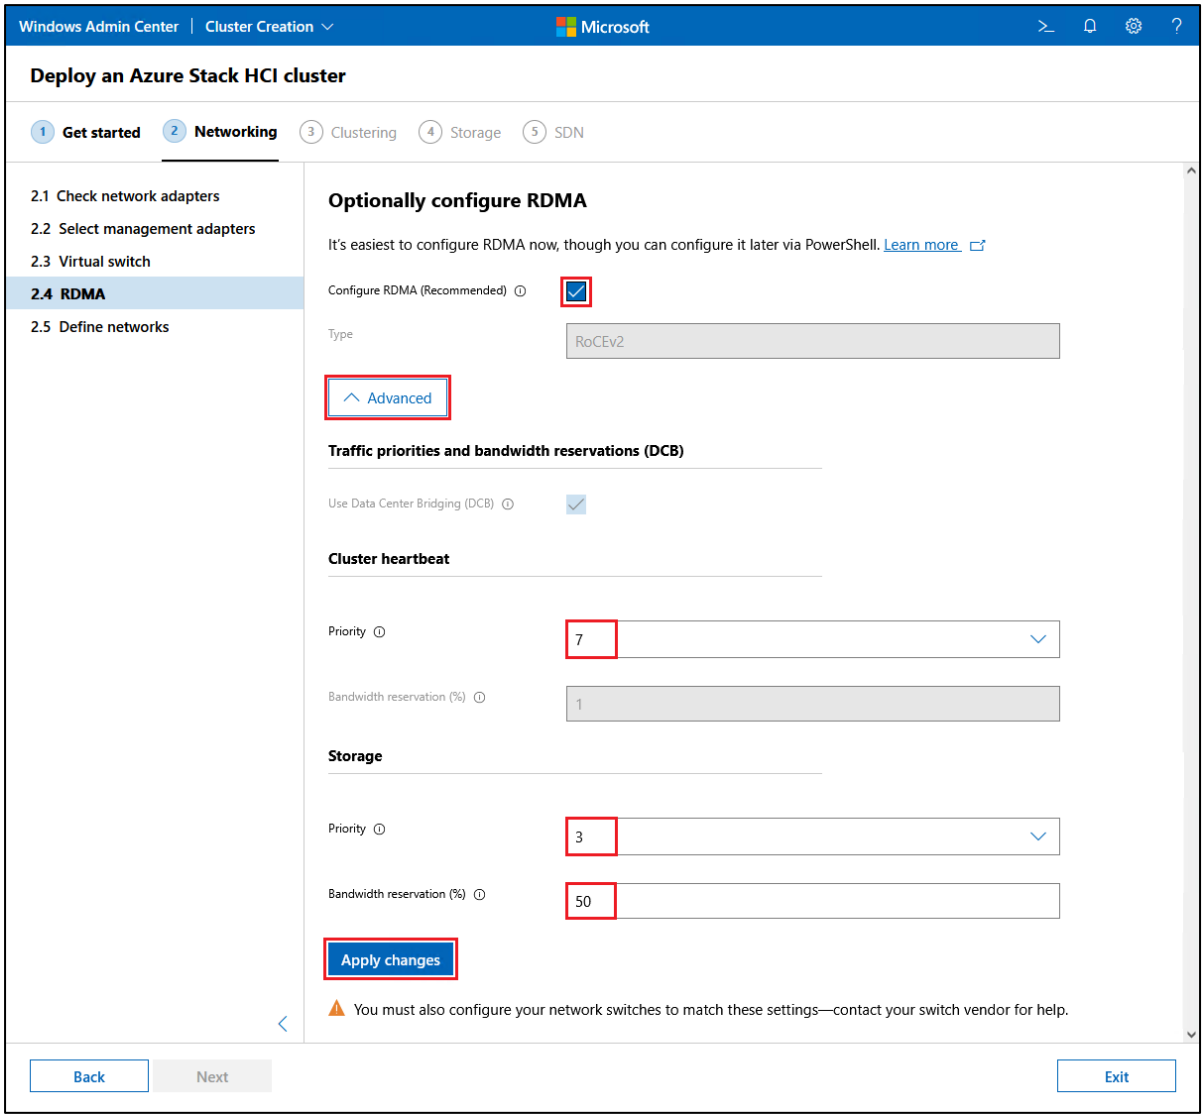

Note: If using the RoCEv2 implementation of RDMA, it is important to ensure that the QoS settings configured for RDMA on the network interfaces match the DCB settings on the network switches that will carry storage traffic. For more information, refer to the topic **RoCE: 2-16 nodes with network switches** in our companion document, *Microsoft Storage Spaces Direct (S2D) Deployment Guide*, found at the following URL:

https://lenovopress.com/lp0064

44. Once RDMA configuration is complete, click Next.

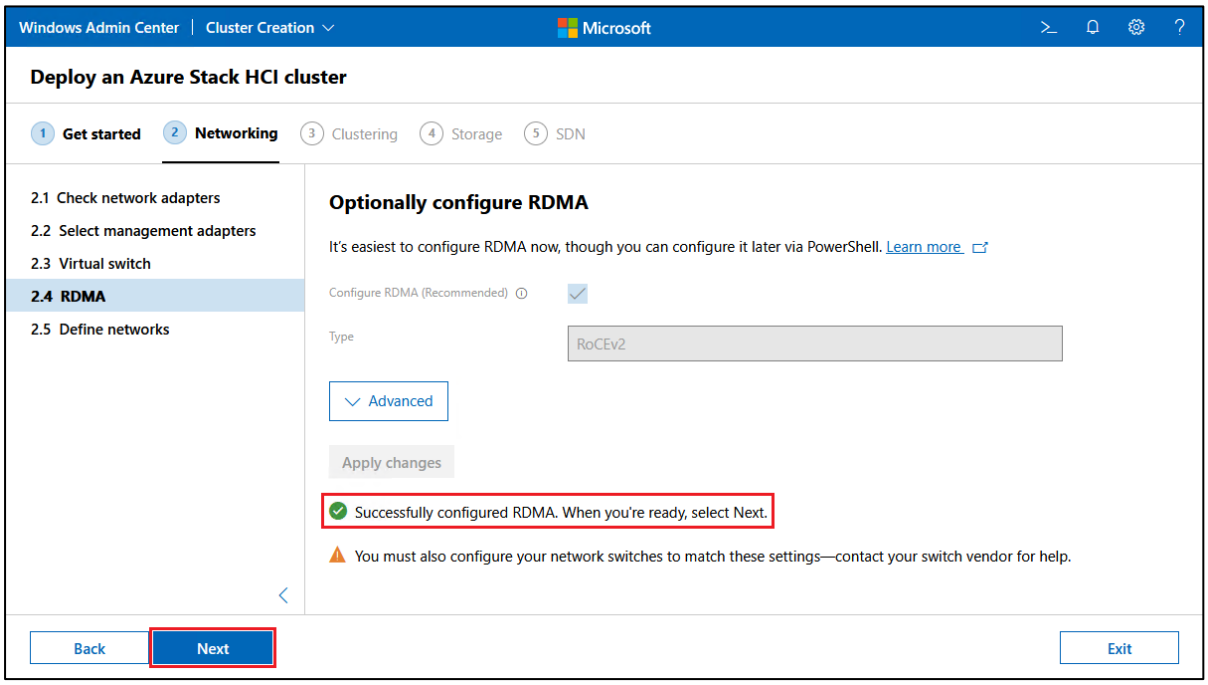

45. In the Define networks panel, Modify IP settings for IP address, subnet mask, and VLAN ID based on your environment. Note that for storage network interfaces, no default gateway should be specified. Settings will be checked briefly as they are specified, with obvious misconfigurations highlighted by the wizard, but full testing does not take place until Apply and test is clicked. Open the Advanced

twisty to expose settings for jumbo packet size and encapsulation overhead. Make any changes required for your environment and then click Apply and test.

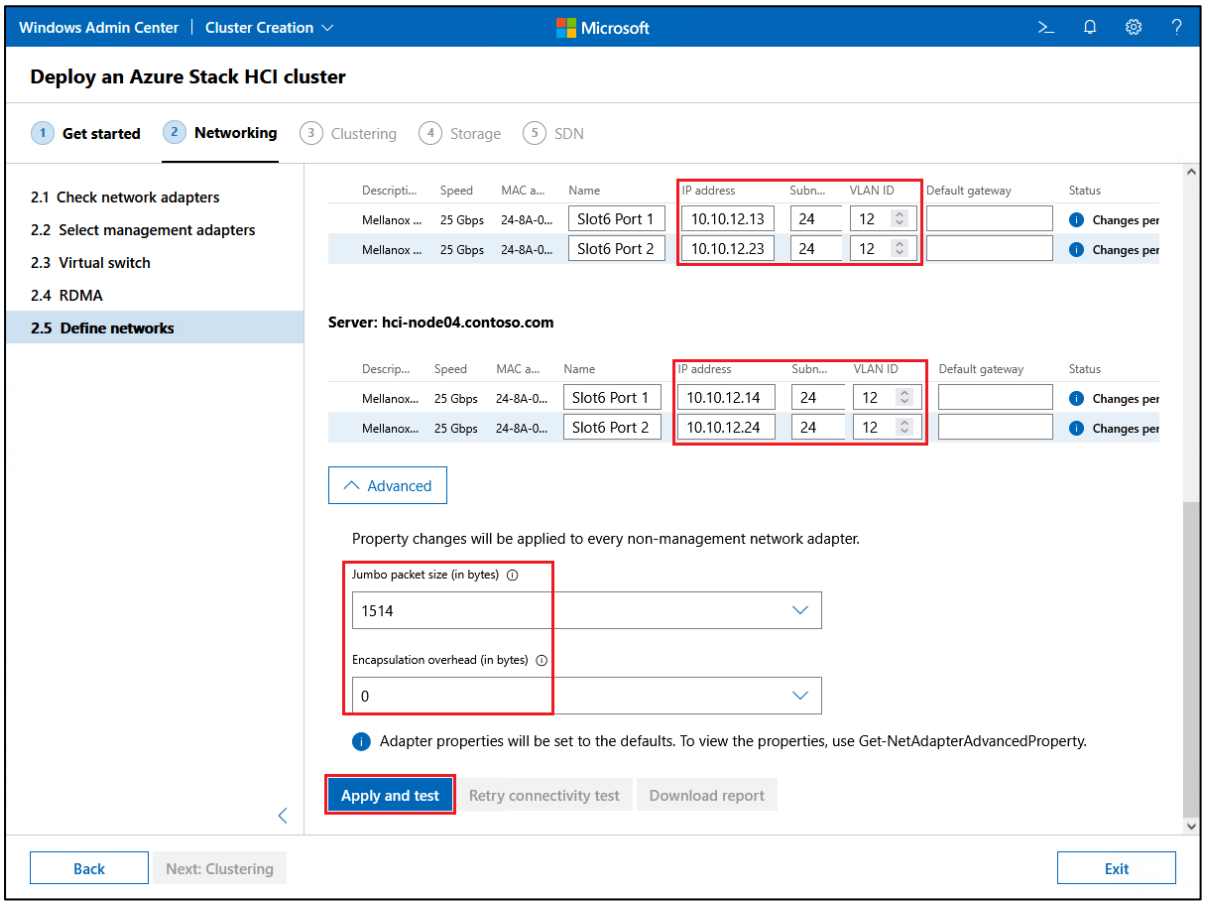

46. All specified network interface changes will be applied and connectivity between the interfaces is tested. You might see a prompt to Specify your credentials. If so, choose the appropriate radio button and enter appropriate credentials if necessary. If you choose to use an account with Domain Admin

privileges, check the Use these credentials for all connections checkbox. Once credentials have been entered, click Continue.

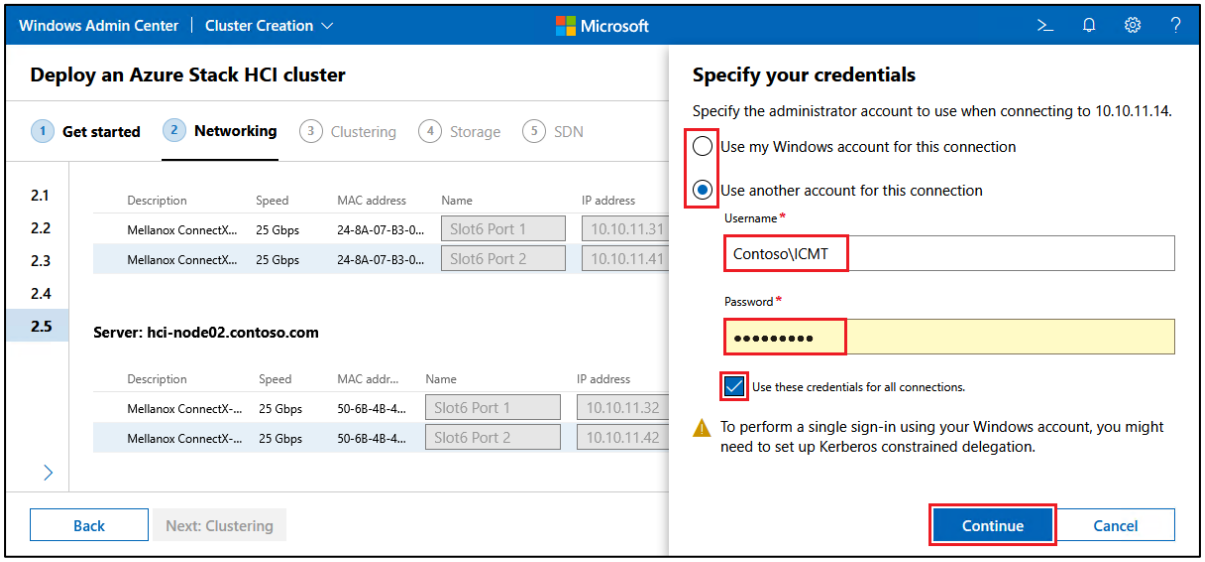

47. You can watch the configuration and connectivity test progress in the Status column. If you see a pop-up message asking to enable CredSSP, click Yes to proceed.

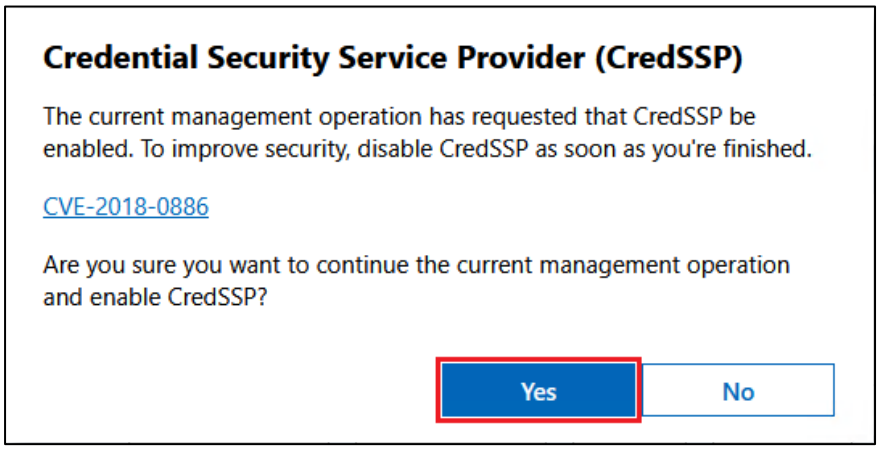

48. Once all testing shows a Status of Passed, choose Download report if desired before clicking Next: Clustering.

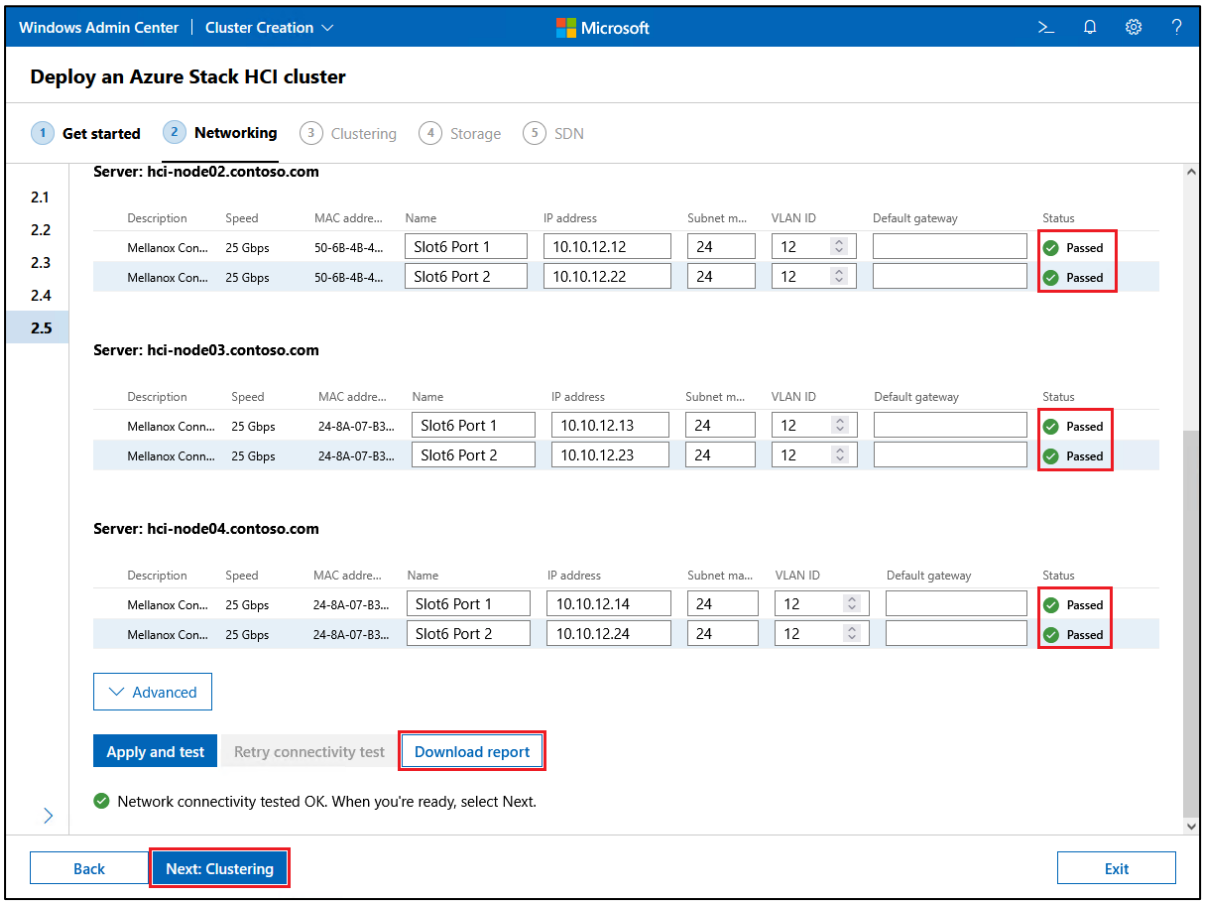

49. Once the Validate the cluster page becomes fully active, click Validate.

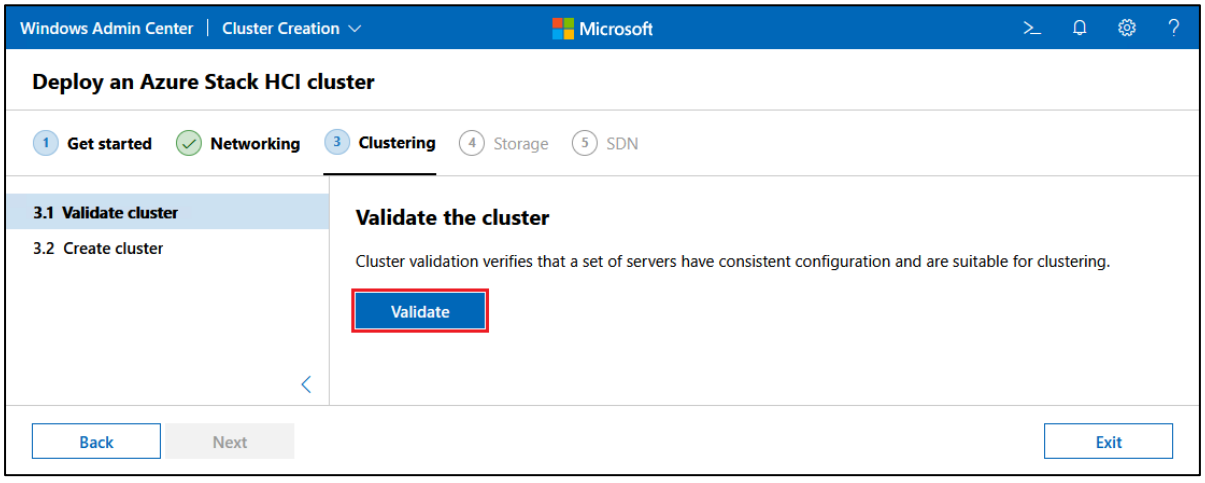

50. After completing the cluster validation process, a summary is shown as well as the option to download the Validation Report. If there are any Warnings or Errors, click Download report to gain additional information.

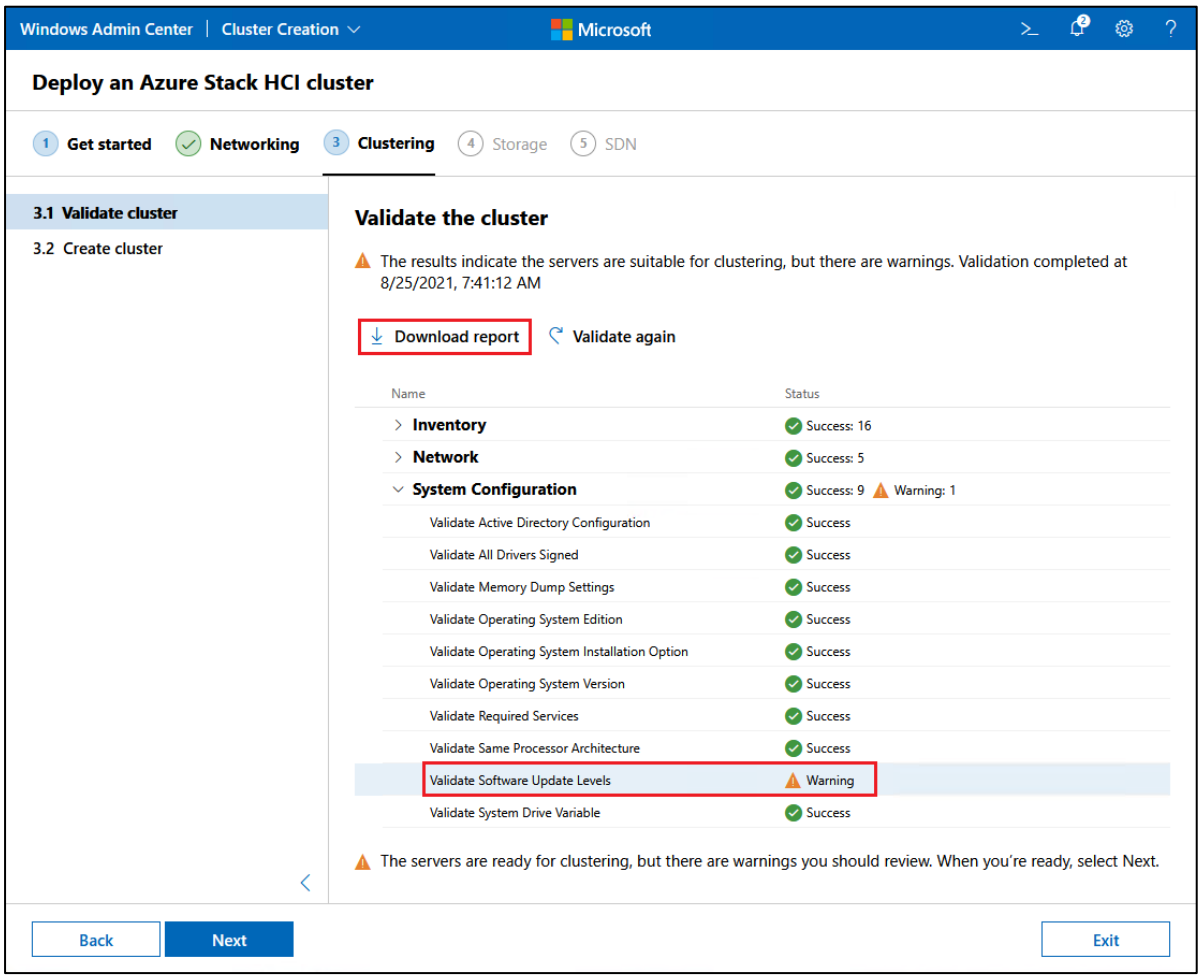

51. If necessary, review the cluster validation report to ensure that there are no issues that would impact the ability to create the failover cluster. In the example shown, the only warning comes from a missing Microsoft Windows Defender Antivirus definition file on Node 1. Since this will not cause serious

issues with the cluster and can be easily resolved after creating the cluster, we can ignore this particular warning.

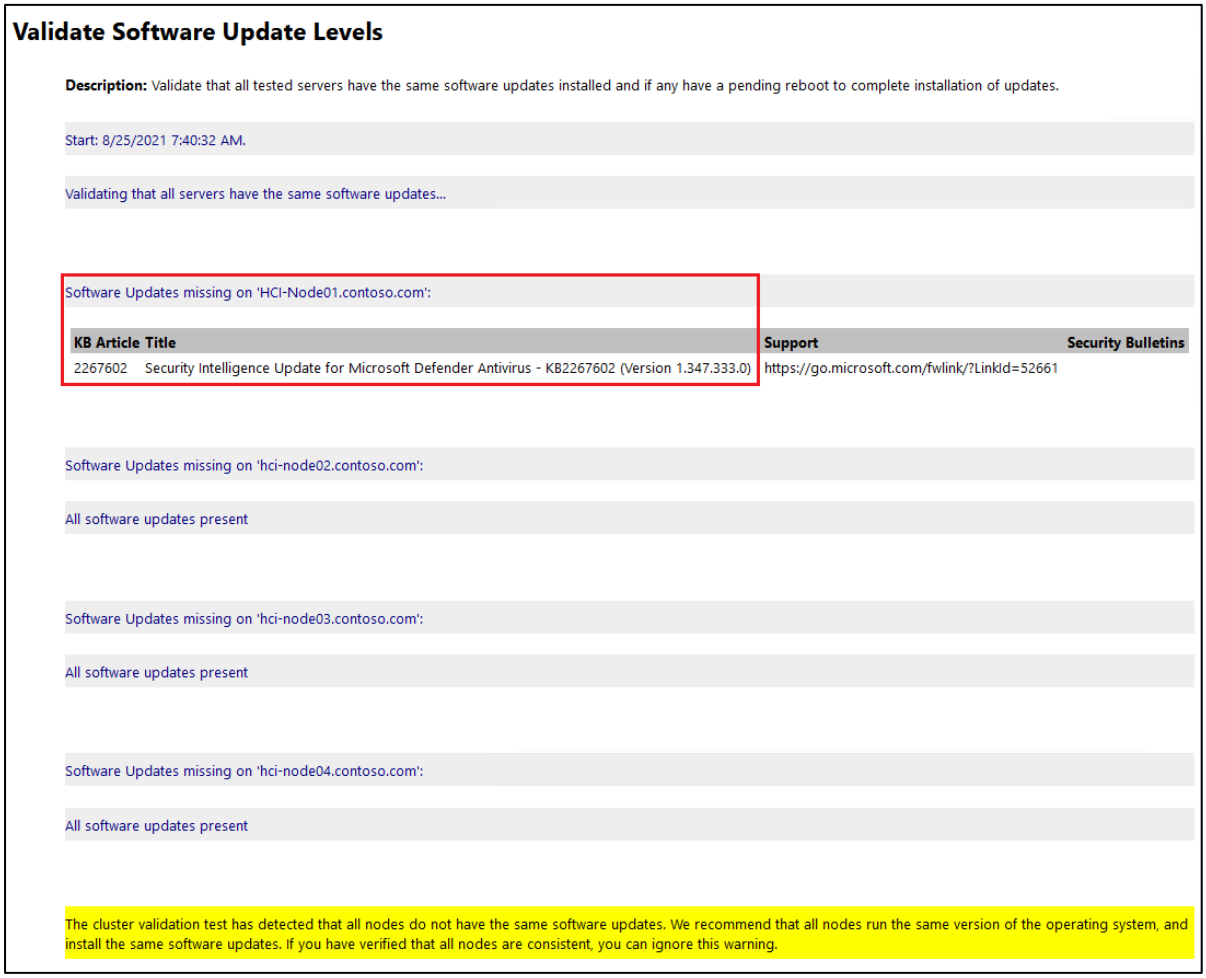

52. After resolving any serious issues, click Next to initiate creation of the failover cluster. When prompted, enter the cluster name and IP address that will be used for the Cluster Name Object (CNO) in Active Directory. Unselect the Use network checkbox for the **Storage** subnet if it is shown on this page, but leave the Management subnet selected and verify the IP address shown. For most

deployments, the Advanced configuration options should be left at their recommended settings. Click Create cluster when ready.

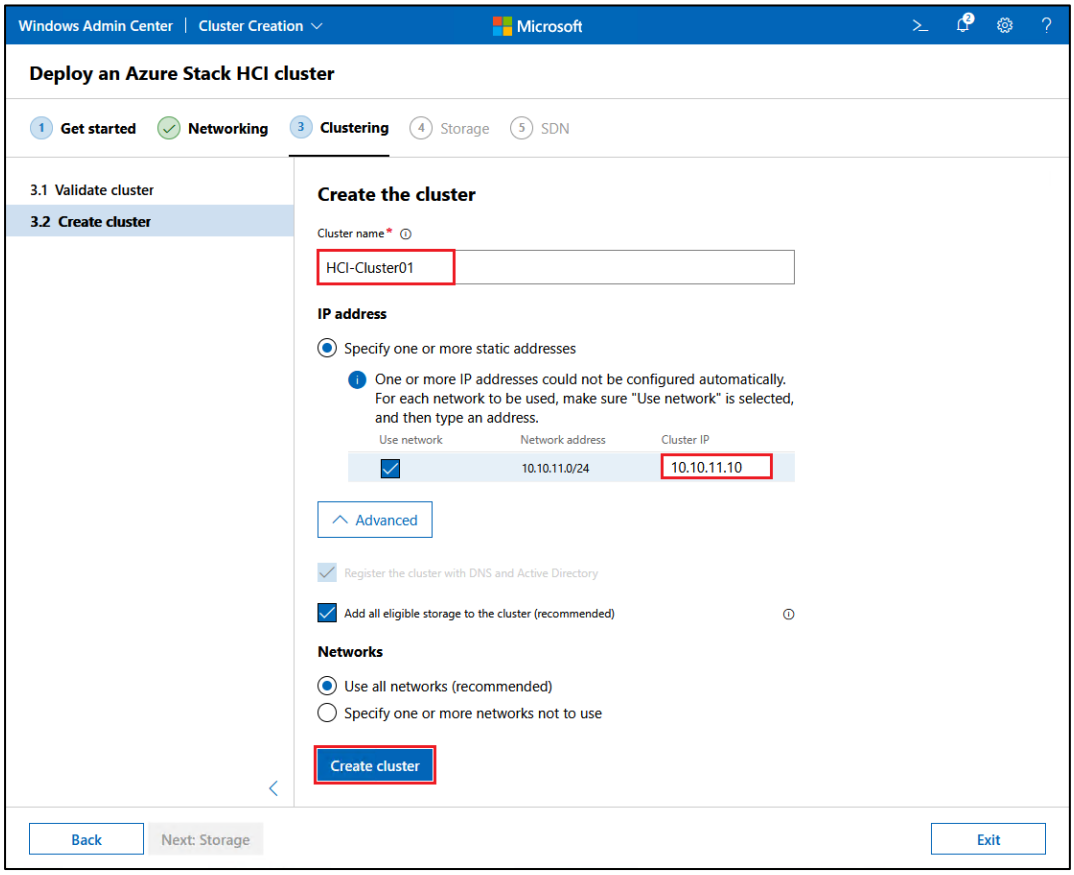

53. Once ready, click Next: Storage.

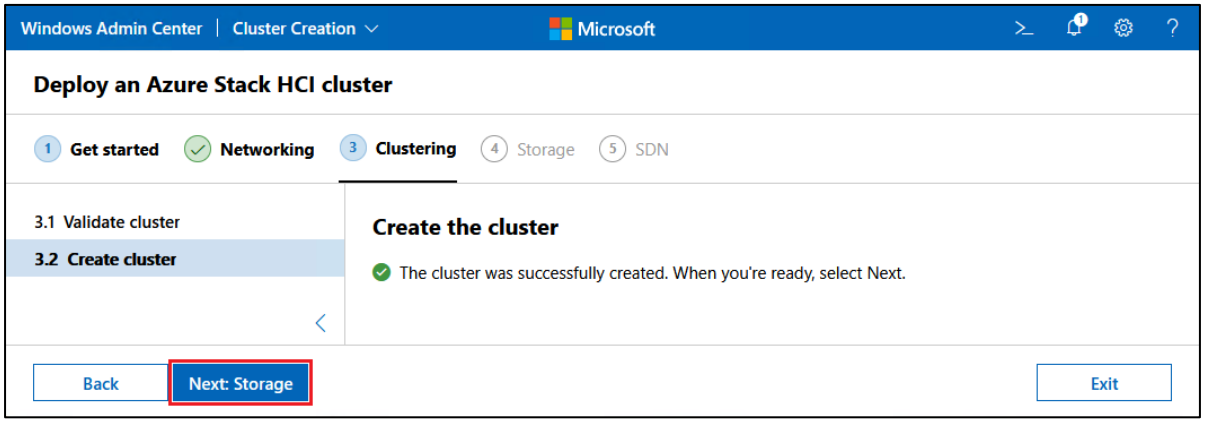

54. Even though this step is shown as optional, we recommend clicking Erase drives to ensure the drives contain no metadata that could interfere with adding the storage devices to the pool.

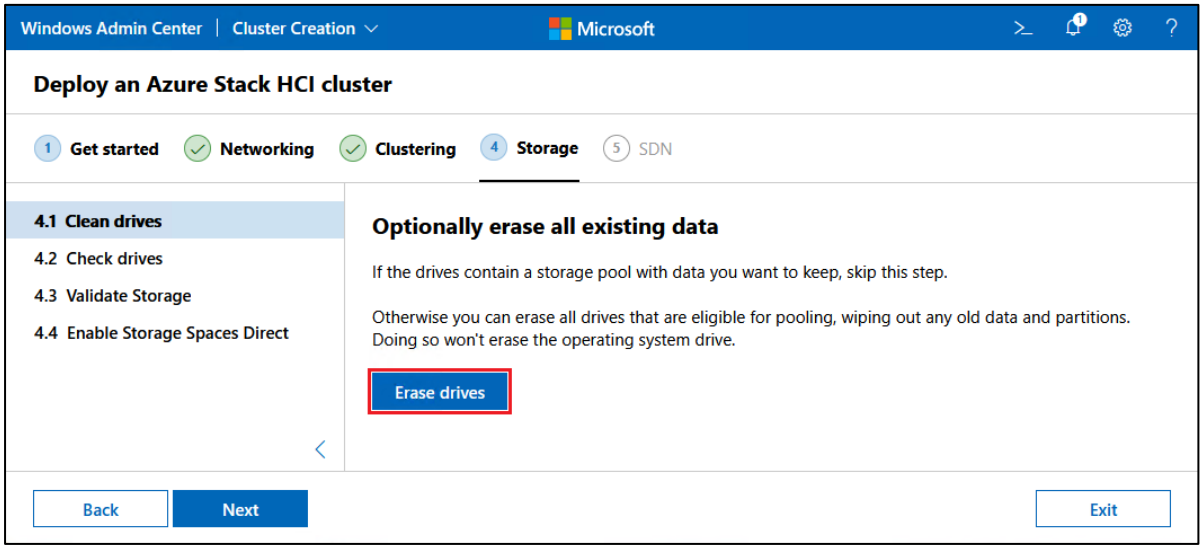

55. In the pop-up confirmation message that is displayed, click Erase drives.

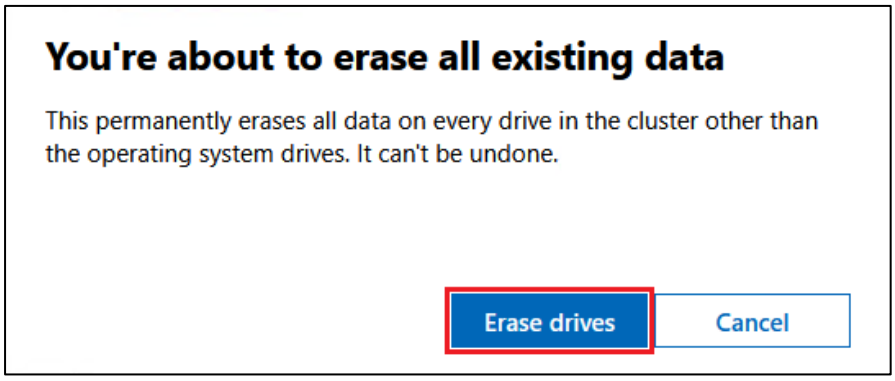

56. Once the process finishes and you see a message stating that all drives have been successfully erased, click Next.

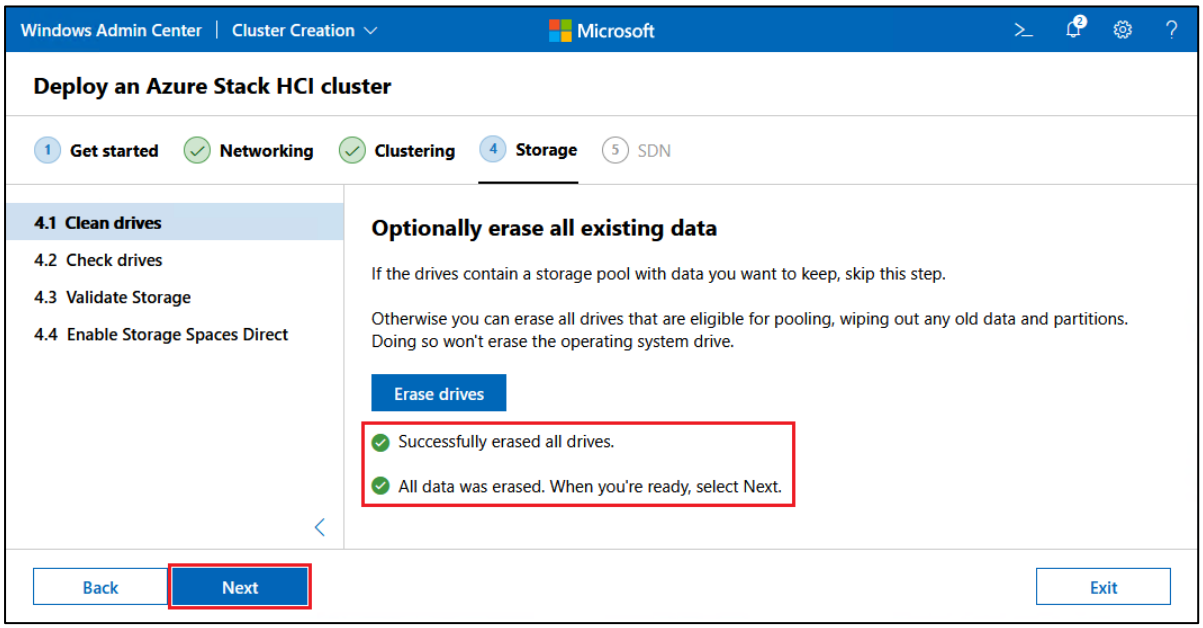

57. After verifying that all drives on all servers show “OK” in the Status column, click Next.

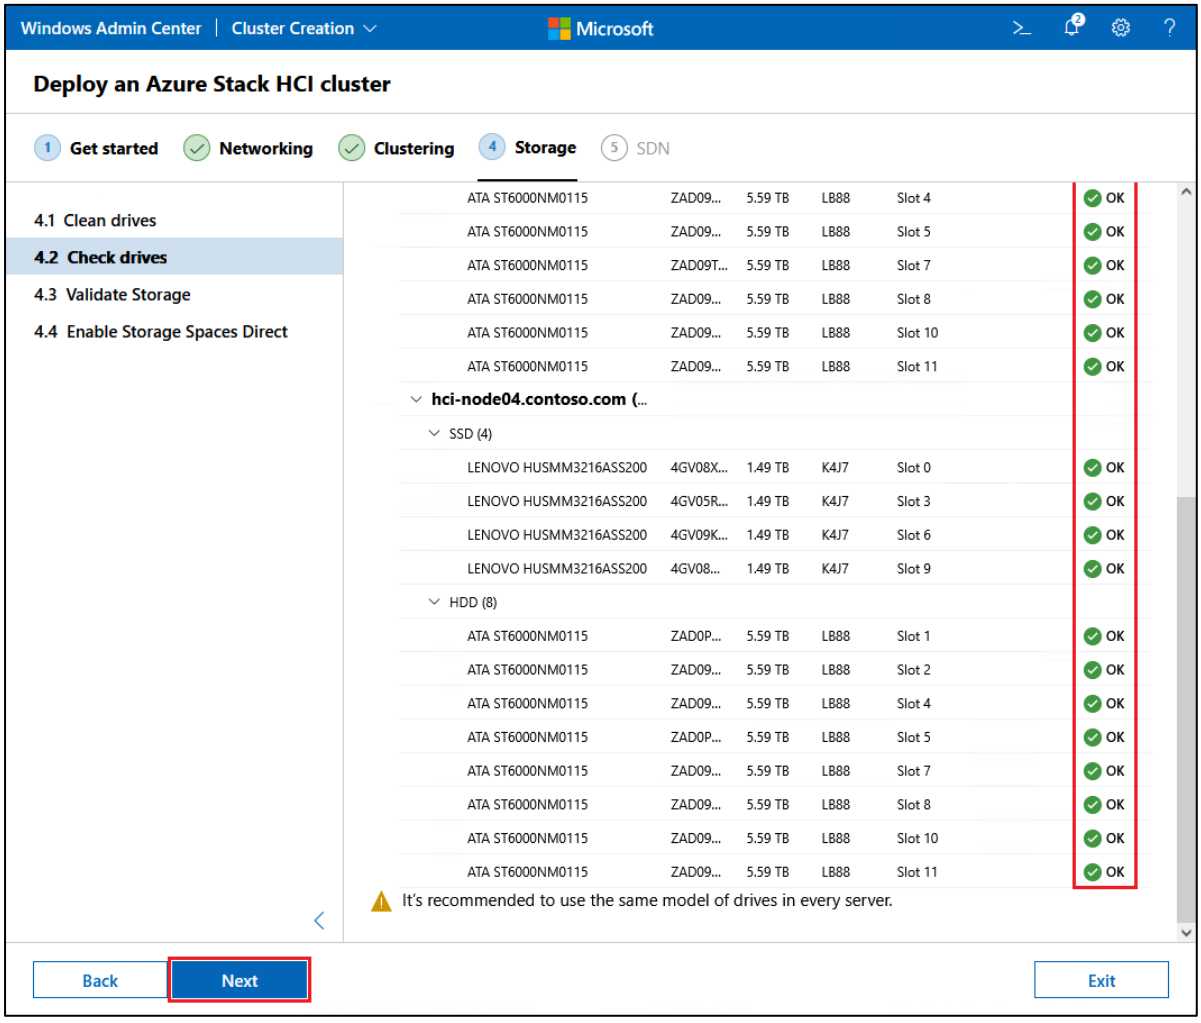

58. After verifying that all storage validation tests show “Success” in the Result column, click Download report if desired, and the click Next.

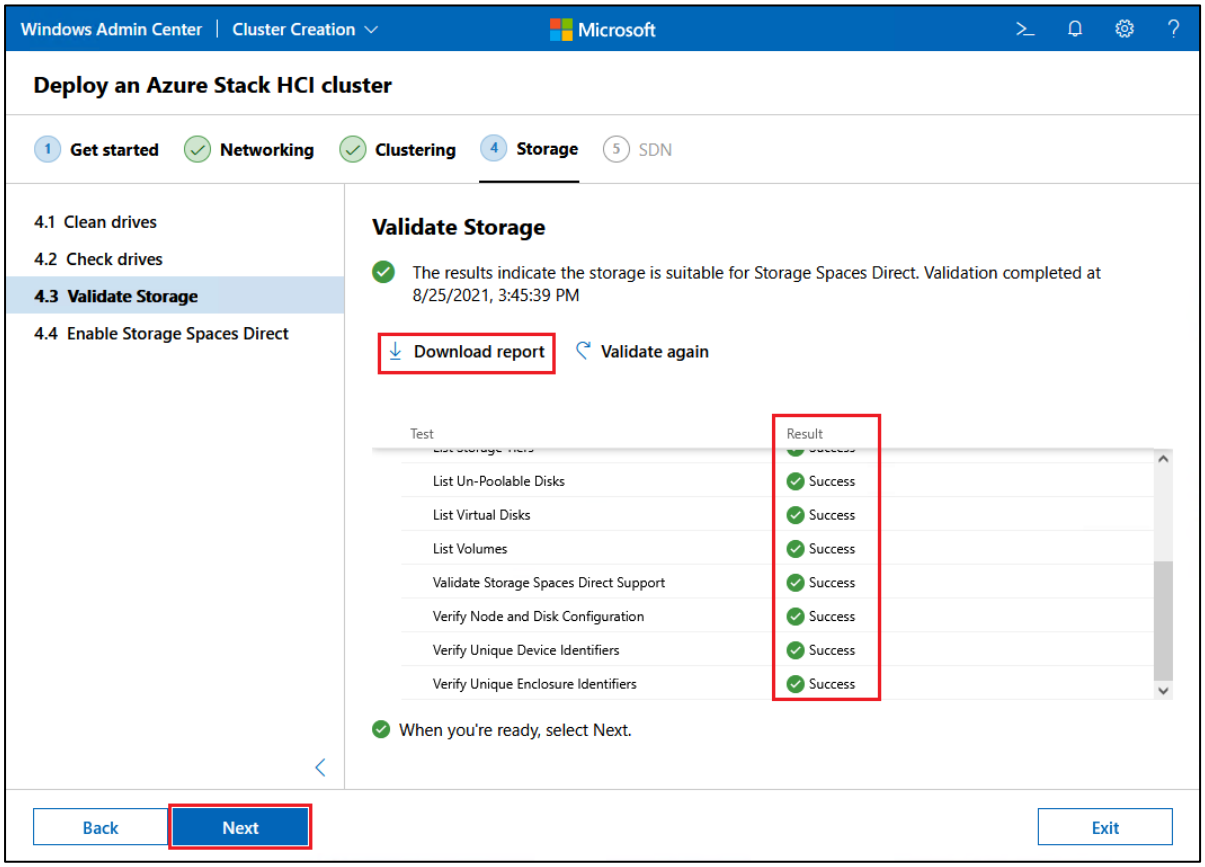

59. On the Enable Storage Spaces Direct panel, click Enable.

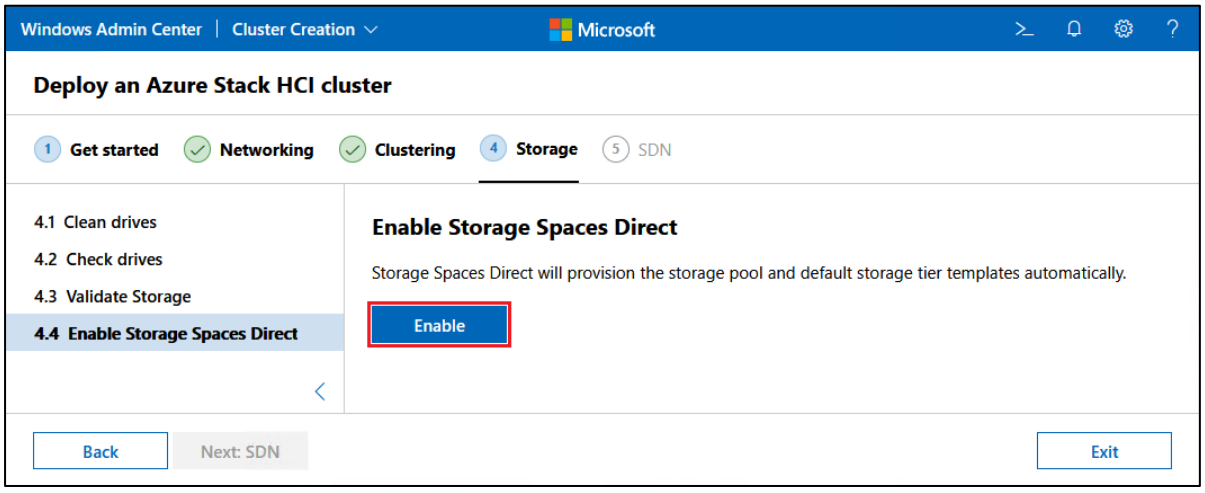

60. Progress will be shown until S2D has been enabled successfully. At this point, click Download report if desired, and then click Next: SDN.

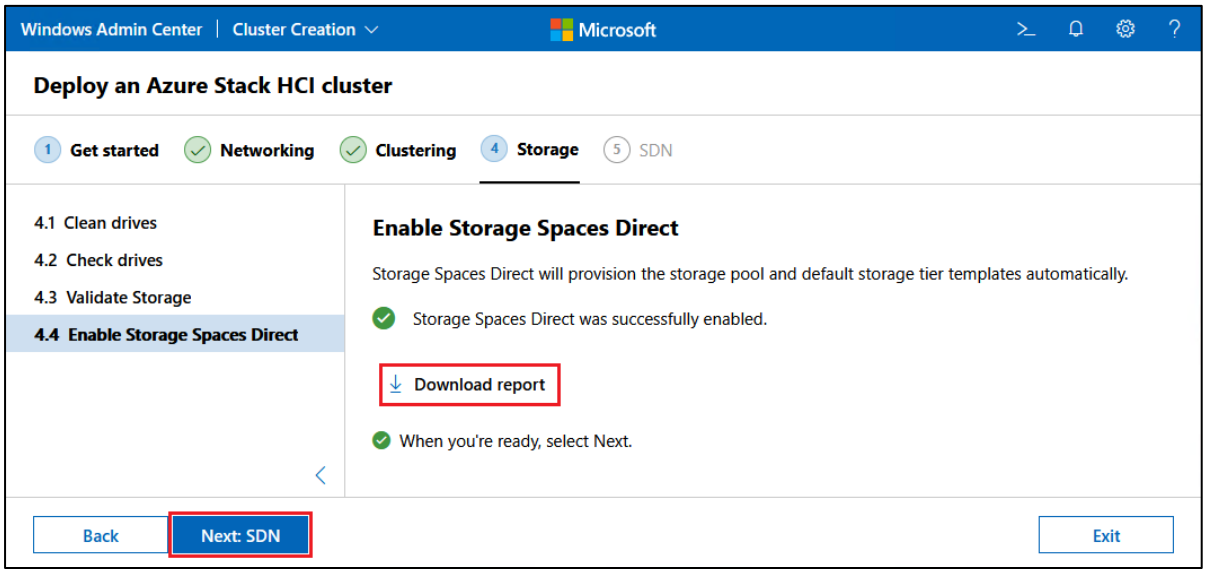

61. Configuring SDN is highly customer-dependent and often not required. For more information, refer to the following Microsoft article:

    https://docs.microsoft.com/en-us/azure-stack/hci/concepts/network-controller-overview

62. Determine whether to proceed with this portion of the deployment or skip it, based on your needs. In our example deployment, we click Skip to finish the deployment.

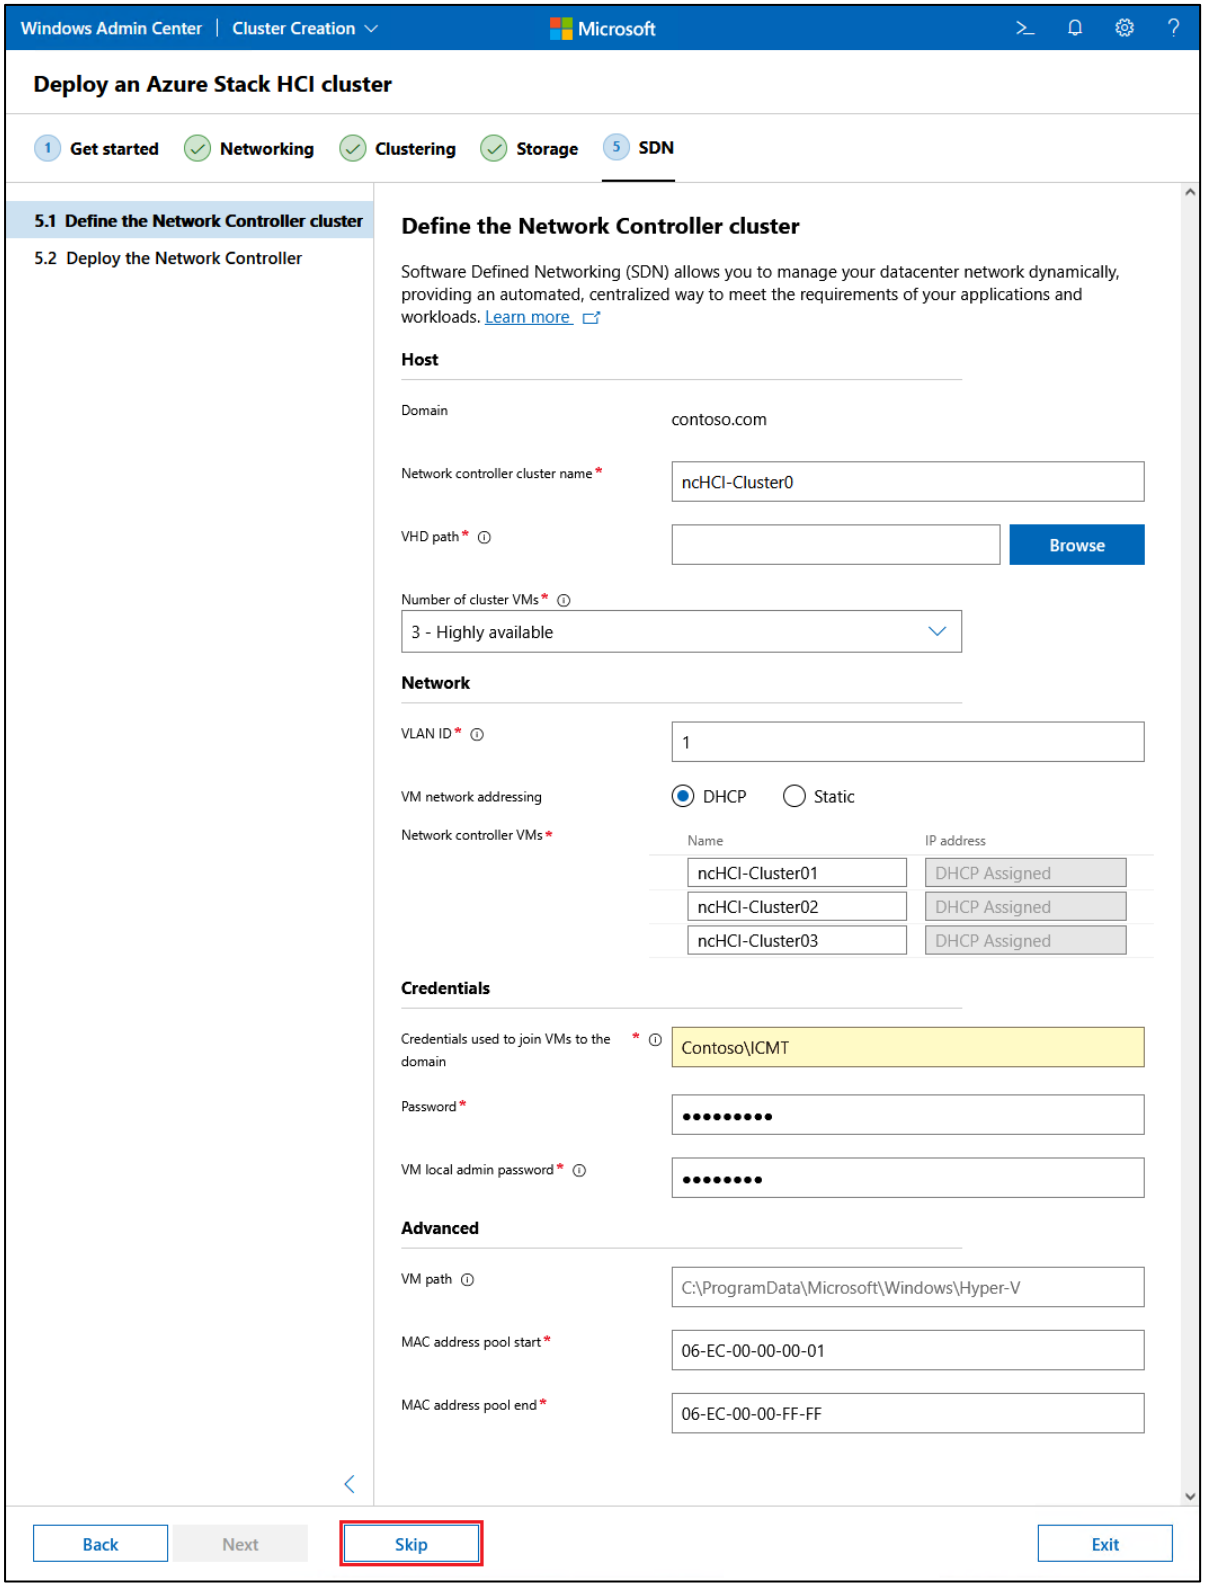

63. The Azure Stack HCI cluster has been successfully created. Click Go to connections list to return to the All connections pane.

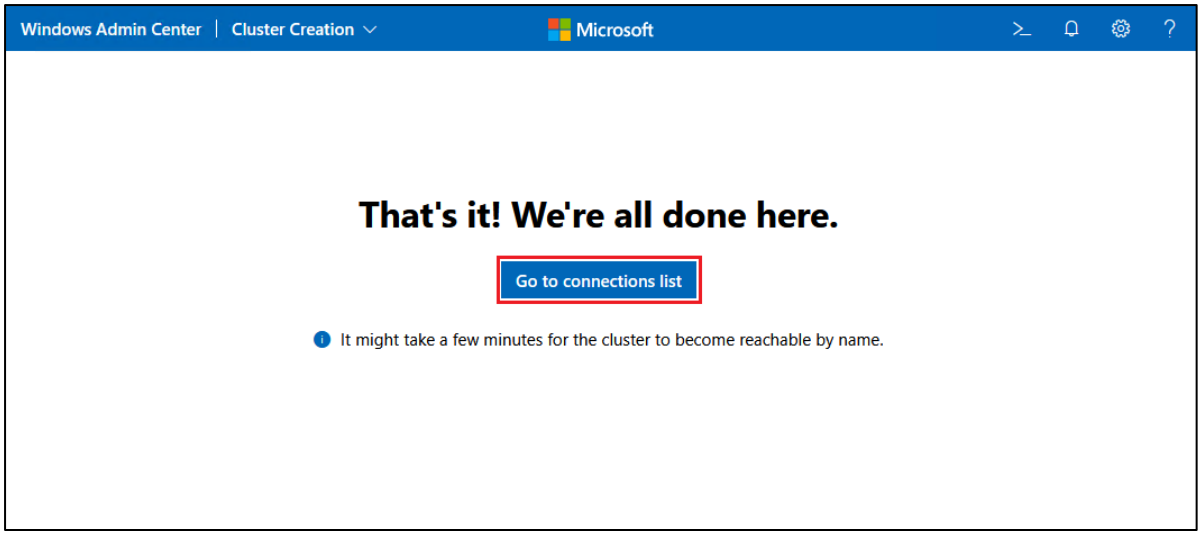

64. You should now see the new cluster in the All connections pane. If desired, click the cluster object to explore the capabilities of WAC for managing the new cluster.

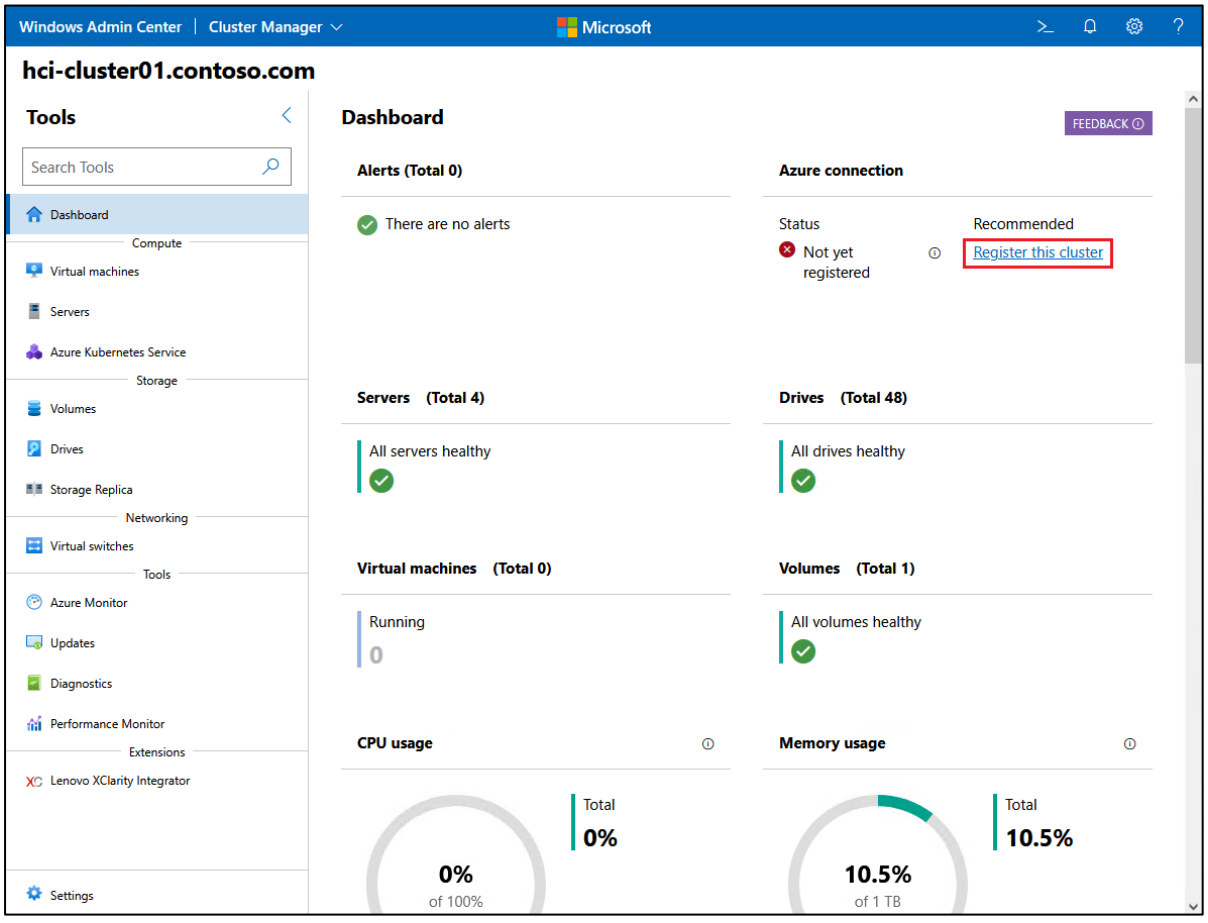

Now that the new HCI cluster is ready, we need to register it with Azure in order to run workloads on it. Click the Register this cluster link (seen in the screenshot above) to begin this process. Note that this step is

required by Microsoft, as follows:

> *"Registering with Azure is required, and your cluster is not fully supported until your registration is active. If you do not register your cluster with Azure upon deployment, or if your cluster is registered but has not connected to Azure for more than 30 days, the system will not allow new virtual machines (VMs) to be created or added."*

For additional details, refer to the Microsoft online documentation at the following URL:

https://docs.microsoft.com/en-us/azure-stack/hci/deploy/register-with-azure

3.6 Post-deployment configuration

There are a few configuration settings that must be made manually after completing the HCI cluster deployment wizard. Some of these steps are required to make Azure Stack HCI work as designed, while others are optional:

- Enable RDMA on vNICs (required)
- Disable LAN Over USB network interface (recommended)
- Disable Flow Control (global pause) on storage interfaces (recommended)
- Disable IPv6 on all network interfaces (optional)

3.6.1 Enable RDMA on vNICs (required)

Currently, the WAC Deployment Wizard fails to enable RDMA on the vNICs that it creates on each node. Since this is a mandatory step for any virtual network interfaces that will be used to carry East-West storage traffic, it must be done via PowerShell after completion of the wizard. To enable RDMA on the SMB vNICs, follow these steps *for each node* in the cluster:

1. In WAC, navigate to Server Manager and then select one of the HCI cluster nodes.
2. In the Tools pane on the left, select PowerShell. This will launch a remote PowerShell session against the node and request a password.
3. To check whether RDMA is enabled on the SMB vNICs, run the PowerShell command shown in Example 1.

Example 1 *PowerShell command to check RDMA enablement on vNIC*

```
Get-NetAdapterRdma | ? Name -Like *SMB* | Format-Table Name, Enabled
```

4. The command will likely return something similar to the following, where the Enabled column shows False for both SMB vNICs.

```
PS C:\> Get-NetAdapterRdma | ? Name -Like *SMB* | Format-Table Name, Enabled

Name                    Enabled
----                    -------
vSMB1                     False
vSMB2                     False
```

5. If this is the case, run the PowerShell command shown in Example 2 to enable RDMA on the SMB vNICs. If RDMA is already enabled on the SMB vNICs, proceed to the next section.

Example 2 *PowerShell command to enable RDMA on SMB vNICs*

```
Enable-NetAdapterRdma -Name "vSMB1","vSMB2"
```

6. Rerunning the command in Example 1 should now return the following:

```
PS C:\> Get-NetAdapterRdma | ? Name -Like *SMB* | Format-Table Name, Enabled

Name                    Enabled
----                    -------
vSMB1                      True
vSMB2                      True
```

7. Repeat Steps 1-6 above for each of the remaining HCL cluster nodes.

3.6.2 Disable LAN Over USB network interface (recommended)

Lenovo ThinkSystem rack servers provide a network interface that supports LAN Over USB for inband connections to the XCC management controller. Although this is a useful feature, WAC and the Failover Cluster Manager can become confused by the presence of this network interface. Fortunately, it can be safely disabled in the OS since it gets enabled automatically when needed (for example by Lenovo tools such as OneCLI) and then disabled again when no longer needed. To disable this interface follow these steps:

1. In WAC, navigate to Server Manager and then select one of the HCI cluster nodes.
2. In the Tools pane on the left, select PowerShell. This will launch a remote PowerShell session against the node and request a password.
3. Although this interface is typically identified by Windows Server operating systems as “Ethernet” it is a good idea to verify this before disabling the network interface. To do this, use the command shown in Example 3.

Example 3 *PowerShell command to verify network interface names*

```
Get-NetAdapter | Format-Table Name, InterfaceDescription, Status
```

4. The command should return something similar to the following. Find the network interface that shows “IBM USB Remote NDIS Network Device” in the InterfaceDescription column and make note of what is shown in the Name column for this interface (likely “Ethernet”).

```
PS C:\> Get-NetAdapter | Format-Table Name, InterfaceDescription, Status

Name                    InterfaceDescription                          Status
----                    --------------------                          ------
Management Physical 2   Intel(R) Ethernet Connection X722 for 10GbE SFP+     Up
Slot6 Port 2            Mellanox ConnectX-4 Lx 2x25GbE PCIe Adapter          Up
vEthernet (Management)  Hyper-V Virtual Ethernet Adapter                     Up
vSMB1                   Hyper-V Virtual Ethernet Adapter #2                  Up
vSMB2                   Hyper-V Virtual Ethernet Adapter #3                  Up
Ethernet                IBM USB Remote NDIS Network Adapter                  Up
Management Physical 1   Intel(R) Ethernet Connection X722 for 10GbE SFP+ #2  Up
Slot6 Port 1            Mellanox ConnectX-4 Lx 2x25GbE PCIe Adapter #2       Up
```

5. After verifying the name of the network interface used for LAN Over USB, use the command shown in Example 4 to disable this network interface (replace “Ethernet” with the appropriate name from your query results if necessary).

Example 4 *PowerShell command to disable a network interface*

```
Disable-NetAdapter -Name "Ethernet"
```

6. You will be asked to confirm this action. Press “y” and Enter to confirm.
7. Repeat the command shown in Example 3 to verify that the correct network interface has been disabled, as shown below.

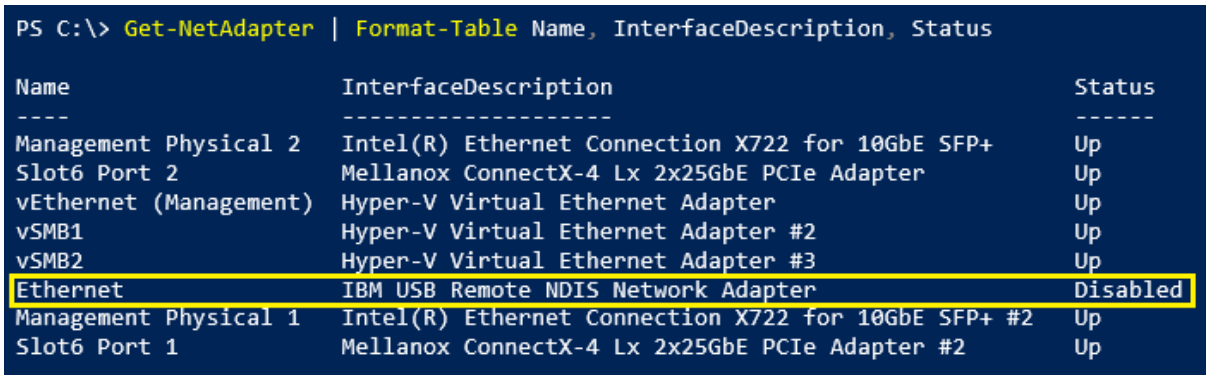

8. Repeat Steps 1-7 above for each of the remaining HCL cluster nodes.

3.6.3 Disable Flow Control (global pause) on storage interfaces (recommended)

Because of the Quality of Service (QOS) settings applied to the network interfaces in this solution, disabling Flow Control (global pause) is not strictly necessary. However, as a best practice and to avoid confusion during troubleshooting, we explicitly disable Flow Control on the physical NICs that carry storage traffic, since these pNICs should never issue a global pause command in any packet. For the Mellanox NIC used in our example, this is an Advanced Property on the NIC. To make this change, follow these steps:

1. In WAC, navigate to Server Manager and then select one of the HCI cluster nodes.
2. In the Tools pane on the left, select PowerShell. This will launch a remote PowerShell session against the node and request a password.
3. After verifying the appropriate pNIC names in your environment (“Slot6 Port 1” and “Slot6 Port 2” in our examples), use the PowerShell command shown in Example 5 to disable flow control on Mellanox pNIC ports. Make sure to provide the appropriate network interface names in your command.

Example 5 *PowerShell command to disable Flow Control advanced setting on Mellanox NIC ports*

```
Set-NetAdapterAdvancedProperty -Name "Slot6 Port 1","Slot6 Port 2" -RegistryKeyword
"*FlowControl" -RegistryValue 0
```

4. To verify that Flow Control has been disabled on the Mellanox pNIC ports, use the PowerShell command shown in Example 6. Again, make sure to provide the appropriate network interface names in the commands you execute.

Example 6 *PowerShell command to check advanced properties on Mellanox pNICs*

```
Get-NetAdapterAdvancedProperty -Name "Slot6 Port 1","Slot6 Port 2"
```

5. You should see in the output from the command shown in Example 6 that Flow Control has been disabled on both ports. The following screen capture from our lab environment shows partial output, including the Flow Control setting for Port 1. Make sure to check that both ports now show the Flow Control setting as Disabled.

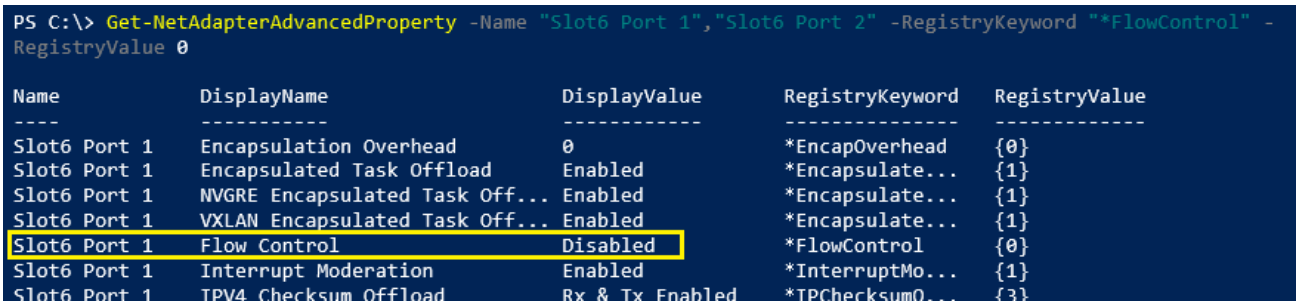

6. Repeat Steps 1-5 above for each of the remaining HCL cluster nodes.

3.6.4 Disable IPv6 (optional)

If you do not use IPv6 in your environment, you can disable it on all network interfaces without impacting the HCI cluster. Since there is currently no way to do this from WAC, we again turn to PowerShell. To disable IPv6 on all network interfaces, follow these steps:

1. In WAC, navigate to Server Manager and then select one of the HCI cluster nodes.
2. In the Tools pane on the left, select PowerShell. This will launch a remote PowerShell session against the node and request a password.
3. To check whether IPv6 is enabled any network interfaces, run the PowerShell command shown in Example 7.

Example 7 *PowerShell command to check IPv6 enablement on all network interfaces*

```
Get-NetAdapterBinding -ComponentID "ms_tcpip6"
```

4. The command will likely return something similar to the following, where the Enabled column shows True for at least some of the network interfaces.

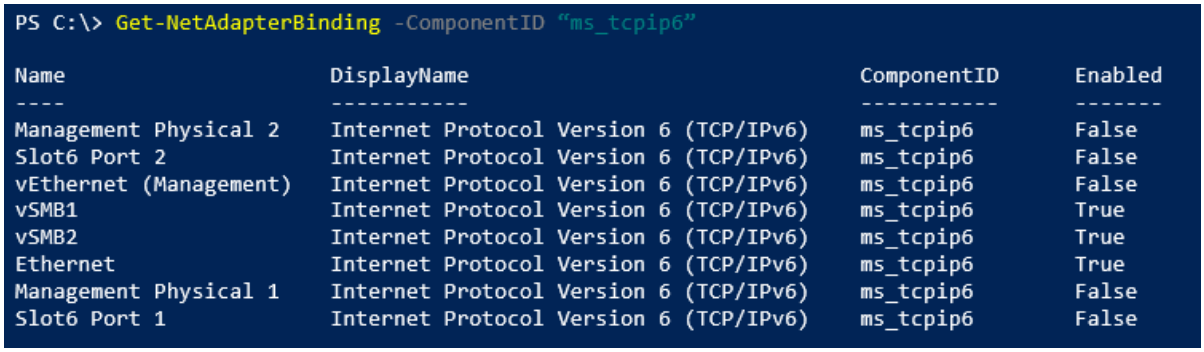

5. If this is the case, use the PowerShell command shown in Example 8 to disable IPv6 on all network interfaces.

Example 8 *PowerShell command to disable IPv6 on all network interfaces*

```
Get-NetAdapter | Disable-NetAdapterBinding -ComponentID "ms_tcpip6"
```

6. Rerunning the command in Example 7 should now return the following:

```
PS C:\> Get-NetAdapterBinding -ComponentID "ms_tcpip6"

Name                    DisplayName                           ComponentID      Enabled
----                    -----------                           -----------      -------
Management Physical 2   Internet Protocol Version 6 (TCP/IPv6)   ms_tcpip6        False
Slot6 Port 2            Internet Protocol Version 6 (TCP/IPv6)   ms_tcpip6        False
vEthernet (Management)  Internet Protocol Version 6 (TCP/IPv6)   ms_tcpip6        False
vSMB1                   Internet Protocol Version 6 (TCP/IPv6)   ms_tcpip6        False
vSMB2                   Internet Protocol Version 6 (TCP/IPv6)   ms_tcpip6        False
Ethernet                Internet Protocol Version 6 (TCP/IPv6)   ms_tcpip6        False
Management Physical 1   Internet Protocol Version 6 (TCP/IPv6)   ms_tcpip6        False
Slot6 Port 1            Internet Protocol Version 6 (TCP/IPv6)   ms_tcpip6        False
```

7. Repeat Steps 1-6 above for each of the remaining HCL cluster nodes.

Post-deployment configuration is now complete.

4 Summary

Lenovo has worked closely with Microsoft for many years to ensure our products perform smoothly and reliably with Microsoft operating systems and software. We have created Lenovo ThinkAgile MX Certified Node and Appliance solutions that contain only servers and server components that have been certified under the Microsoft Azure Stack HCI Program to run Microsoft Storage Spaces Direct (S2D) properly. These solutions provide a solid foundation for customers looking to consolidate both storage and compute capabilities on a single hardware platform. They provide outstanding performance, high availability protection and effortless scale-out growth potential to accommodate evolving business needs.

ThinkAgile MX Series platforms offer the choice of Azure Stack HCI Appliance or Azure Stack HCI Certified Node. These validated platforms help modernize on-premises infrastructure with pre-tested, pre-configured, and easy-to-order configurations, with seamless Azure integration. As a direct and indirect Microsoft Cloud Solution Provider (CSP), Lenovo offers cloud services and subscriptions through the Lenovo Cloud Marketplace, which enable Azure Stack HCI use cases with the ThinkAgile MX platforms.

This document has provided step-by-step instructions for deploying an Azure Stack HCI cluster on Lenovo ThinkAgile MX Certified Nodes and Appliances using the WAC deployment wizard. It has guided the reader through the wizard, explaining each step and providing real world examples from our labs. After completing the deployment wizard, additional configuration steps were presented to ensure readiness of this HCI solution for production use.

Additional resources

The following additional resources might help to provide background information related to Lenovo ThinkAgile MX solutions and Azure Stack HCI functionality.

Lenovo Professional Services

Lenovo offers an extensive range of solutions, from the simple OS-only laden product to much more complex solutions running cluster and cloud technologies. For customers looking for assistance in the form of design, deploy or migrate, Lenovo Professional Services is your go-to partner.

Our worldwide team of IT Specialists and IT Architects can help customers scope and size the right solutions to meet their requirements, and then accelerate the implementation of the solution with our on-site and remote services. For customers also looking to elevate their own skill sets, our Technology Trainers can craft services that encompass solution deployment plus skills transfer, all in a single affordable package.

To inquire about our extensive service offerings and solicit information on how we can assist in your new Storage Spaces Direct implementation, please contact us at x86svcs@lenovo.com.

For more information about our service portfolio, please see our website:

https://www3.lenovo.com/us/en/data-center/services/c/services?menu-id=services

Lenovo resources

Lenovo ThinkAgile MX Best Recipes

https://datacentersupport.lenovo.com/us/en/solutions/HT507406

ThinkAgile MX Best Recipe Updates Repository

https://thinkagile.lenovo.com/mx

ThinkAgile MX Information Center

https://thinkagile.lenovofiles.com/help/index.jsp?topic=%2Fcom.lenovo.thinkagile.7Z20.doc%2Fmx_intro.html

ThinkAgile MX Documents - Lenovo Press

https://lenovopress.com/servers/thinkagile/mx-series

Lenovo ThinkAgile MX Certified Configurations for Azure Stack HCI – V1 Servers

https://lenovopress.com/lp0866

Lenovo ThinkAgile MX Certified Configurations for Azure Stack HCI – V2 Servers

https://lenovopress.com/lp1520

Microsoft Storage Spaces Direct (S2D) Deployment Guide

https://lenovopress.com/lp0064

ThinkAgile MX1021 on SE350 Azure Stack HCI (S2D) Deployment Guide

https://lenovopress.com/lp1298

Lenovo XClarity Integrator for Windows Admin Center

https://support.lenovo.com/eg/en/solutions/HT507549

Lenovo XClarity Administrator Product Guide
https://lenovopress.com/tips1200-lenovo-xclarity-administrator

Lenovo ThinkAgile MX3520-H Appliance and MX Certified Node with Hybrid Storage for Microsoft Azure Stack HCI (Xeon SP Gen 2) Product Guide
https://lenovopress.com/lp1175

Lenovo ThinkAgile MX3520-F Appliance and MX Certified Node with All Flash Storage for Microsoft Azure Stack HCI (Xeon SP Gen 2) Product Guide
https://lenovopress.com/lp1176

Lenovo ThinkAgile MX1020 Appliance and MX1021 Certified Node for Microsoft Azure Stack HCI Product Guide
https://lenovopress.com/lp1296

ThinkAgile MX3330 and MX3331 1U Appliances and Certified Nodes (Intel Xeon SP Gen 3) Product Guide
https://lenovopress.com/lp1511

ThinkAgile MX3530 and MX3531 2U Appliances and Certified Nodes (Intel Xeon SP Gen 3) Product Guide
https://lenovopress.com/lp1512

Lenovo Professional Services
https://www.lenovo.com/us/en/data-center/services/c/services?menu-id=services

Lenovo Premier Support Overview, Resources and Tools
https://lenovopub.jiveon.com/docs/DOC-10010

Microsoft resources

Microsoft WSSD program
https://docs.microsoft.com/en-us/windows-server/sddc

HCI solutions from Microsoft
https://www.microsoft.com/en-us/cloud-platform/software-defined-datacenter

Microsoft blog “The technical value of WSSD validated HCI solutions”
https://cloudblogs.microsoft.com/windowsserver/2018/02/20/the-technical-value-of-wssd-validated-hci-solutions-part-1/
https://cloudblogs.microsoft.com/windowsserver/2018/02/21/the-technical-value-of-validated-hci-solutions-part-2/

Storage Spaces Direct overview
https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/storage-spaces-direct-overview

Taking a Storage Spaces Direct server offline for maintenance
https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/maintain-servers

Manage HCI with Windows Admin Center
https://docs.microsoft.com/en-us/windows-server/manage/windows-admin-center/use/manage-hyper-converged

Mirror-accelerated parity details
https://docs.microsoft.com/en-us/windows-server/storage/refs/mirror-accelerated-parity

What’s new in Windows Server 2022
https://docs.microsoft.com/en-us/windows-server/get-started/whats-new-in-windows-server-2022

What's new in Windows Server 2019

https://docs.microsoft.com/en-us/windows-server/get-started-19/whats-new-19

What's new in Storage in Windows Server

https://docs.microsoft.com/en-us/windows-server/storage/whats-new-in-storage

Delimit the allocation of volumes in Storage Spaces Direct

https://docs.microsoft.com/en-us/windows-server/storage/storage-spaces/delimit-volume-allocation

Change history

This is the initial release of this document.

Changes in the November 2021 update:

- Added steps to enable iWARP RDMA mode on Marvell network adapters before launching the WAC deployment wizard

Authors

This paper was produced by the following team of specialists:

Dave Feisthammel is a Senior Solutions Architect working at the Lenovo Bellevue Technology Center in Bellevue, Washington. He has over 25 years of experience in the IT field, including four years as an IBM client and over 18 years working for IBM and Lenovo. His areas of expertise include Windows Server and systems management, as well as virtualization, storage, and cloud technologies. He is currently a key contributor to Lenovo solutions related to Microsoft Azure Stack HCI and Azure Stack Hub.

Hussein Jammal is a Senior Solutions Architect Engineer and Microsoft Solution Lead in EMEA working in Bucharest, Romania. He has over 13 years of experience in the IT engineering field. He is focused on Microsoft Cloud on-premises solution development and enablement, including Azure Stack Hub, Azure Stack HCI and Azure Stack Edge solutions. He works regularly with customers on complex troubleshooting issues and proofs of concept.

Mike Miller is a Windows Engineer with the Lenovo Bellevue Technology Center in Bellevue, Washington. He has over 35 years in the IT industry, primarily in client/server support and development roles. The last 13 years have been focused on Windows Server operating systems and server-level hardware, particularly on operating system/hardware compatibility, advanced Windows features, and Windows test functions.

David Ye is a Principal Solutions Architect at Lenovo with over 25 years of experience in the IT field. He started his career at IBM as a Worldwide Windows Level 3 Support Engineer. In this role, he helped customers solve complex problems and critical issues. He is now working in the Lenovo Infrastructure Solutions Group, where he works with customers on Proof of Concept designs, solution sizing and reviews, and performance optimization. His areas of expertise are Windows Server, SAN Storage, Virtualization and Cloud, and Microsoft Exchange Server. He is currently leading the effort in Microsoft Azure Stack HCI and Azure Stack Hub solutions development.

A special thank you to the following Lenovo colleagues for their contributions to this project:

- Daniel Ghidali – Manager, Microsoft Technology and Enablement
- Vy Phan – Technical Program Manager, Microsoft OS and Solutions
- Oana Adelina Onofrei – Solutions Engineer, ISG Software Development
- Laurentiu Petre – Solutions Engineer, ISG Software Development
- Zhi (Paul) Wang – Senior Engineer, System Management Software Development
- David Watts – Senior IT Consultant, Lenovo Press

Trademarks and special notices

References in this document to Lenovo products or services do not imply that Lenovo intends to make them available in every country.

Lenovo, the Lenovo logo, ThinkSystem, ThinkCentre, ThinkVision, ThinkVantage, ThinkPlus and Rescue and Recovery are trademarks of Lenovo.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

Information is provided "AS IS" without warranty of any kind.

All customer examples described are presented as illustrations of how those customers have used Lenovo products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-Lenovo products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by Lenovo. Sources for non-Lenovo list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. Lenovo has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-Lenovo products. Questions on the capability of non-Lenovo products should be addressed to the supplier of those products.

All statements regarding Lenovo future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. Contact your local Lenovo office or Lenovo authorized reseller for the full text of the specific Statement of Direction.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in Lenovo product announcements. The information is presented here to communicate Lenovo’s current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard Lenovo benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Photographs shown are of engineering prototypes. Changes may be incorporated in production models.

Any references in this information to non-Lenovo websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this Lenovo product and use of those websites is at your own risk.