



Lenovo ThinkAgile MX Certified Configurations for Azure Stack HCI – V3 Servers

Last Update: June 2024

Provides details of Lenovo certified configurations for SR630 V3 and SR650 V3 based servers

Describes the Microsoft Azure Stack HCI Program

Provides guidance for properly configuring nodes for an Azure Stack HCI cluster

Lists supported options that can be used when configuring Azure Stack HCI cluster nodes

Dave Feisthammel

Hussein Jammal

Mike Miller

David Ye



Abstract

This document provides background information regarding the Microsoft Azure Stack HCI program for Windows Server 2022, as well as the benefits of deploying certified configurations based on Lenovo ThinkAgile™ MX Certified Nodes and Appliances. We focus on details of current Lenovo certified configurations for Azure Stack HCI that are based on ThinkSystem™ SR630 V3 and SR650 V3 servers, including processor, memory, network, and storage components available for each cluster node, including the following solutions:

- ▶ ThinkAgile MX630 V3 IS (Integrated System/Appliance)
- ▶ ThinkAgile MX630 V3 CN (Certified Node)
- ▶ ThinkAgile MX650 V3 IS (Integrated System/Appliance)
- ▶ ThinkAgile MX650 V3 CN (Certified Node)

Looking for Lenovo ThinkAgile MX solutions that are based on our V1 servers? Check our companion document at <http://lenovopress.com/lp0866>.

Looking for Lenovo ThinkAgile MX solutions that are based on our V2 servers? Check our companion document at <http://lenovopress.com/lp1520>.

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.

Contents

Introduction	3
Microsoft HCI certification overview	3
ThinkAgile MX Series solutions	4
Lenovo certified configurations for Microsoft Azure Stack HCI	5
Component selection	14
Network switches	21
Other recommendations	23
Summary	24
Change History	25
Authors	25
Notices	27
Trademarks	28

Introduction

Deploying hyperconverged infrastructure (HCI) 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, a host operating system from Microsoft, 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 take advantage of Azure Stack HCI functional capabilities that are built into 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. Our customers can reap the benefits of our partnership with Microsoft by taking advantage of HCI solutions that have been certified under the Microsoft Azure Stack HCI program for Windows Server 2019.

Deploying Lenovo certified configurations for Microsoft HCI solutions takes the guesswork out of system configuration. Whether you intend to build a converged or hyperconverged S2D cluster, you can rest assured that purchasing a certified configuration will provide a rock solid foundation with minimal obstacles along the way. These node configurations are certified by Lenovo and validated by Microsoft for out-of-the-box optimization. Using the Lenovo ThinkAgile MX Certified Node and Integrated System (Appliance) configurations presented in this document, you can get up and running without lengthy design and build time, knowing that the solution will work as intended.

This document briefly discusses the Microsoft HCI certification programs, and then presents the Lenovo certified configurations for ThinkSystem SR630 V3 and ThinkSystem SR650 V3 servers that have been validated for use in a Microsoft HCI solution under these programs. Details of each node configuration are specified, including all key components. Since there is some latitude for component customization in these configurations, the rules for customization are also described.

Microsoft HCI certification overview

To provide the best experience and support to HCI customers in production, Microsoft introduced the WSSD certification program, which includes Windows Server operating systems through Windows Server 2016. For Windows Server 2019 and beyond, Microsoft has rebranded their HCI certification program as Microsoft "Azure Stack HCI."

Microsoft Azure Stack HCI program

Beginning with Windows Server 2019, Microsoft has rebranded their HCI certification program as the "Azure Stack HCI program." According to Microsoft, "Azure Stack HCI is a

hyper-converged Windows Server 2019 cluster that uses validated hardware to run virtualized workloads on-premises, optionally connecting 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.”

Many of the certification requirements from the WSSD program have been carried over into the Azure Stack HCI program, which begins with Windows Server 2019 logo certification. Each key hardware component must pass rigorous testing procedures and be certified as an Azure Stack HCI component before it can be included in an Azure Stack HCI solution. In addition to the specific certification requirements that must be met by the individual hardware components, Microsoft requires end-to-end solution validation for each configuration to be certified. This involves running the fully configured HCI solution for many hours, while putting it through various usage and potential failure scenarios.

Key components must have SDDC “Additional Qualifiers” certification (SDDC-AQ). These components include the following:

- Network adapters
- Storage adapters (SAS/SATA HBAs)
- Storage devices (NVMe, SSD, and HDD)

Perhaps the greatest value to be derived from this program from a customer perspective is to reduce the risks and unknowns associated with deploying an HCI solution using “off the shelf” components. To earn certification in the Azure Stack HCI program, Lenovo has met or exceeded multiple criteria set by Microsoft for quality, accelerated time to value, out-of-the-box optimization, and expedited problem resolution.

What is unique about Lenovo certified configurations for Microsoft HCI solutions is our rigorous evaluation process to select the best components from our existing Lenovo product portfolio. The main objective is to ensure our customers will have great confidence in our HCI solutions for a production environment.

For more information about the Microsoft Azure Stack HCI program, visit the following URL:

<https://docs.microsoft.com/en-us/windows-server/azure-stack-hci>

To learn why deploying a certified configuration for S2D is an optimal path to success for S2D deployment, read the two-part Microsoft blog post at the following URLs:

<https://cloudblogs.microsoft.com/windowsserver/2018/02/20/the-technical-value-of-validated-hci-solutions-part-1>

<https://cloudblogs.microsoft.com/windowsserver/2018/02/21/the-technical-value-of-validated-hci-solutions-part-2>

To browse all Lenovo certified solutions for Azure Stack HCI in the Microsoft Catalog, visit the following URL:

<https://azurestackhcisolutions.azure.microsoft.com/#/catalog?Manufacturer=Lenovo>

ThinkAgile MX Series solutions

As previously discussed, the Microsoft HCI certification programs allow OEM partners to deliver pre-engineered, validated HCI solutions. Whether your preference is for a Certified Node or an Integrated System (Appliance), Lenovo has designed, tested and validated the

ThinkAgile MX Series offerings to quickly and easily provide the solutions you need, with the confidence required to exceed the stringent requirements of today's IT. The result is that you can quickly deploy a robust, high-performance storage solution and rapidly solve your IT challenges.

ThinkAgile MX Certified Node

The Lenovo ThinkAgile MX Certified Node Series of solutions maps to Microsoft "Azure Stack HCI Validated Nodes." These solutions package Microsoft-certified HCI solutions into easy-to-use machine types to provide the following:

- ▶ Easy to order
- ▶ Enforced configuration rules to ensure a valid configuration
- ▶ Best recipe firmware
- ▶ ThinkAgile Advantage (where available)
- ▶ Optional services such as deployment, management, etc.

ThinkAgile MX Integrated System

Lenovo ThinkAgile MX Integrated Systems (previously know as Appliances) are based on exactly the same hardware as ThinkAgile MX Certified Nodes. The only differences between a ThinkAgile MX Certified Node and Integrated System (Appliance) that are based on the same server (for example, the ThinkSystem SR630 V3 rack server) is that the Integrated System configuration includes the following items:

- ▶ Azure Stack HCI operating system is preloaded before shipping to the customer
- ▶ ThinkAgile Advantage Support for 3 years (can be uplifted to a longer term, quicker response time, or both via Lenovo Premier Support offerings)

The remainder of this document focuses on describing the existing Lenovo configurations that have been certified under the Microsoft HCI certification programs and the details of key components contained in each configuration. The purpose of this document is to provide guidance for Lenovo customers and technical pre-sales personnel during the process of configuring a Microsoft certified HCI solution for production usage. This document assumes the reader has prior knowledge of Microsoft HCI technologies, including S2D.

Lenovo certified configurations for Microsoft Azure Stack HCI

The Microsoft HCI certification programs allow for solution certification using a min/max paradigm. Under the program, OEM partners are allowed to certify a minimum configuration and a maximum configuration in order to receive certification of all configurations that lie between these extremes. Therefore, the configurations presented in this document represent examples of what has been certified, rather than an exhaustive list of the only certified configurations that are available. Refer to "Component selection" on page 14 for additional information regarding the components that have been certified.

Table 1 lists the key components of example configurations that are certified for use by Azure Stack HCI. This list of configurations is not exhaustive and can be customized. Refer to "Component selection" on page 14 for information about customizing these configurations. The Feature Code (FC) shown for each option is shown in the table to make it easy to find each device in Lenovo's Data Center Solution Configurator.

The number of nodes in an Azure Stack HCI cluster can range from a minimum of 2 to a maximum of 16. Where the cache and capacity columns show different storage devices,

these configurations represent a two-tier storage structure in which the devices shown in the cache column are used to provide a high-performance cache layer in front of the devices shown in the capacity column, which make up the total raw storage capacity of the node. Where a single storage device type spans the cache and capacity columns, these configurations represent a single tier storage structure. Note that Lenovo regularly certifies additional configurations as time and resources allow.

The format of the configuration name follows a specific pattern. The first two alphabetic characters define the storage types included in the configuration (“N” for NVMe, “S” for SSD, and “H” for HDD). The next three or four alphanumeric characters define the total raw storage capacity of the node (e.g. “80T” indicates a total capacity of 80TB per node). The last numeric character defines the configuration sequence for the given component parameters. For example, if there are two certified configurations that contain NVMe and HDD storage devices with a total raw capacity of 80TB per node, they would be referred to as NH80T1 and NH80T2.

Table 1 Configuration highlights for Lenovo ThinkAgile MX630 V3 and MX650 V3 solution nodes

Config	CPU/RAM	Cache	Capacity	Storage Controller	Storage Network
NS15T1 (All-flash)	MX630 V3 2 CPUs 256GB	2 x 1.6TB ThinkSystem U.2 P5800X NVMe FC: BMM8	8 x 1.92TB ThinkSystem 5400 MAX SSD FC: BQ22	440-16i HBA FC: BM50	Intel E810-DA2 10/25GbE PCIe FC: BCD6
NH60T1 (Hybrid)	MX650 V3 2 CPUs 256GB	2 x 3.2TB ThinkSystem U.2 P5620 NVMe FC: BNEM	10 x 6TB ThinkSystem 7.2K SATA HDD FC: AUUA	440-16i Internal HBA FC: B8P1	Mellanox CX-6 Lx PCIe 10/25GbE FC: BE4U
NS64T1 (All-flash)	MX650 V3 2 CPUs 512GB	4 x 3.2TB ThinkSystem U.2 P5620 NVMe FC: BNEM	20 x 3.2TB ThinkSystem PM1655 SSD FC: BP3K	440-16i Internal HBA FC: B8P1	Mellanox CX-6 Dx PCIe 100GbE FC: B8PP
NN61T1 (All-flash)	MX630 V3 2CPUs 512GB	16 x ThinkSystem E1.S 7450 PRO 3.84TB Read Intensive NVMe PCIe 4.0 x4 HS SSD (all-NVMe single tier) FC: BP3L		NVMe Retimer FC: B98C	Intel E810-DA4 10/25GbE OCP FC: BP8L
NN92T1 (All-flash)	MX650 V3 2CPUs 512GB	24 x ThinkSystem 2.5" U.2 P5520 3.84TB Read Intensive NVMe PCIe 4.0 x4 HS SSD (all-NVMe single tier) FC: BCFR		NVMe Retimer FC: B98C	Mellanox CX-6 Dx PCIe 100GbE FC: B8PP
SH16T1 (Hybrid)	MX630 V3 2CPUs 256GB	2 x ThinkSystem 2.5" PM1655 800GB SSD FC: BNW8	4 x ThinkSystem 3.5" 4TB 7.2K SATA 6Gb HDD FC:AUU8	430-8i HBA FC: AUNL	Intel E810-DA2 10/25GbE PCIe FC: BCD6

Lenovo certified configuration details

This section includes details of each of the example Lenovo configurations contained in Table 1 that have been certified under the Microsoft HCI certification programs. Each configuration lists the Lenovo ThinkAgile MX Certified Node or ThinkSystem rack server that is used for the S2D cluster node, as well as the storage and network devices that have been certified for the configuration.

Again, the configurations shown are example configurations and are not meant to provide an exhaustive list of all available certified configurations. Refer to “Component selection” on page 14 for additional information regarding components that have been certified. If you have questions about the validity of a configuration you would like to purchase, check with your account team.

NS15T1 All-flash configuration

This configuration uses the Lenovo ThinkAgile MX630 V3 CN or IS configured with NVMe storage devices for the cache tier and SSD devices for the capacity tier. Total raw capacity of this configuration is roughly 15TB per node. As seen in Figure 1, this configuration can be ordered with either three low profile PCIe slots or one full height and one low profile PCIe slot at the rear of the server.




Figure 1 Lenovo ThinkAgile MX630 V3 Certified Node configuration NS15T1

Additional details include the following:

- ▶ CPU: 2 x Intel Xeon Gold 5420+ 28C 205W 2.0GHz Processor (FC BQ65)
- ▶ Memory: 16 x ThinkSystem 16GB TruDDR5 4800 MHz RDIMM (FC BKTL)
- ▶ Storage:
 - ThinkSystem 440-16i SAS/SATA PCIe Gen4 12Gb HBA (FC BM50)
 - 2 x ThinkSystem 2.5" U.2 P5800X 1.6TB NVMe PCIe 3.0 x4 HS SSD for the cache tier (FC BMM8)
 - 8 x ThinkSystem 2.5" 5400 MAX 1.92TB SATA 6Gb HS SSD for the capacity tier (FC BQ22)
- ▶ Storage (RDMA) network adapter: ThinkSystem Intel E810-DA2 10/25GbE SFP28 2-Port PCIe Ethernet Adapter (FC BCD6)

This is an all-flash configuration that uses NVMe storage devices for the cache tier and SSD devices for the capacity tier. This configuration is recommended when raw capacity requirements are less than 20TB per node.

NH60T1 Hybrid configuration

This configuration uses the Lenovo ThinkAgile MX650 V3 CN or IS configured with NVMe storage devices for the cache tier and HDD devices for the capacity tier. Total raw capacity of this configuration is roughly 60TB per node.




Figure 2 Lenovo ThinkAgile MX650 V3 Certified Node configuration NH60T1a

Additional details include the following:

- ▶ CPU: 2 x Intel Xeon Gold 6430 32C 270W 2.2GHz Processor (FC BPPC)
- ▶ Memory: 16 x ThinkSystem 16GB TruDDR5 4800 MHz RDIMM (FC BKTL)
- ▶ Storage:
 - ThinkSystem 440-16i SAS/SATA PCIe Gen4 12Gb Internal HBA (FC B8P1)
 - 2 x ThinkSystem 3.5" U.2 P5620 3.2TB Mixed Use NVMe PCIe 4.0 x4 HS SSD for the cache tier (FC BNEM)
 - 10 x ThinkSystem 3.5" 6TB 7.2K SATA 6Gb HS 512e HDD for the capacity tier (FC AUUA)
- ▶ Storage (RDMA) network adapter: Mellanox ConnectX-6 Lx 10/25GbE SFP28 2-Port PCIe Ethernet Adapter (FC BE4U)
- ▶ Non-storage Network adapter: ThinkSystem Intel X710-T2L 10GBASE-T 2-port OCP Ethernet Adapter (FC BCD5)

This is a hybrid configuration that uses NVMe storage devices for the cache tier and HDD devices for the capacity tier, with increased raw capacity of 60TB per node. This configuration is recommended when more storage capacity is required. A 16-node Microsoft HCI solution built using this configuration will provide a total raw storage capacity of nearly 1 petabyte, which can be increased further by choosing larger HDDs.

NS64T1 All-flash configuration

This configuration uses the Lenovo ThinkAgile MX650 V3 CN or IS configured with NVMe storage devices for the cache tier and SSD devices for the capacity tier. Total raw capacity of this configuration is roughly 64TB per node.




Figure 3 Lenovo ThinkAgile MX650 V3 Certified Node configuration NS64T1a

Additional details include the following:

- ▶ CPU: 2 x Intel Xeon Gold 5420+ 28C 205W 2.0GHz Processor (FC BQ65)
- ▶ Memory: 32 x ThinkSystem 16GB TruDDR5 4800 MHz RDIMM (FC BKTL)
- ▶ Storage:
 - ThinkSystem 440-16i SAS/SATA PCIe Gen4 12Gb Internal HBA (FC B8P1)
 - 4 x ThinkSystem 2.5" U.2 P5620 3.2TB Mixed Use NVMe PCIe 4.0 x4 HS SSD for the cache tier (FC BNEH)
 - 20 x ThinkSystem 2.5" PM1655 3.2TB Mixed Use SAS 12Gb HS SSD for the capacity tier (FC BNW6)
- ▶ Storage (RDMA) network adapter: ThinkSystem Mellanox ConnectX-6 Dx 100GbE QSFP56 2-port PCIe Ethernet Adapter (FC B8PP)

This is an all-flash configuration that uses NVMe storage devices for the cache tier and SSD devices for the capacity tier, with increased raw capacity of 64TB per node. This configuration is recommended when more storage capacity is required. Network bandwidth of 100GbE is recommended for this configuration in order to keep up with the storage performance. A 16-node Microsoft HCI solution built using this configuration will provide a total raw storage capacity of roughly 1 petabyte, which can be increased further by choosing larger SSDs for the capacity tier.

NN61T1 All-flash configuration

This configuration uses the Lenovo ThinkAgile MX630 V3 CN or IS configured with E1.S NVMe storage devices in a single tier. Total raw capacity of this configuration is roughly 61TB per node since one or two (two required for software RAID) of the 4TB E1.S NVMe devices is used for boot.




Figure 4 Lenovo ThinkAgile MX630 V3 Certified Node configuration NN61T1a

Additional details include the following:

- ▶ CPU: 2 x Intel Xeon Gold 6430 32C 270W 2.2GHz Processor (FC BPPC)
- ▶ Memory: 32 x ThinkSystem 16GB TruDDR5 4800 MHz RDIMM (FC BKTL)
- ▶ Storage:
 - ThinkSystem 4-Port PCIe Gen4 NVMe Retimer Adapter (FC B98C)
 - 16 x ThinkSystem E1.S 7450 PRO 3.84TB Read Intensive NVMe PCIe 4.0 x4 HS SSD (FC BP3L)
- ▶ Storage (RDMA) network adapter: ThinkSystem Intel E810-DA4 10/25GbE SFP28 4-Port OCP Ethernet Adapter (FC BP8L)
 - An OCP device is required here since no PCIe slots are available

This is an all-NVMe all-flash configuration that uses NVMe storage devices in a single storage tier. This configuration is recommended when a small datacenter footprint and extreme storage performance are requirements. Network bandwidth of 25GbE is recommended for this configuration.

NN92T1 All-flash configuration

This configuration uses the Lenovo ThinkAgile MX650 V3 CN or IS configured with NVMe storage devices in a single storage tier. Total raw capacity of this configuration is roughly 92TB per node.




Figure 5 Lenovo ThinkAgile MX650 V3 Certified Node configuration NN92T1a

Additional details include the following:

- ▶ CPU: 2 x Intel Xeon Gold 5420+ 28C 205W 2.0GHz Processor (FC BQ65)
- ▶ Memory: 32 x ThinkSystem 16GB TruDDR5 4800 MHz RDIMM (FC BKTL)
- ▶ Storage:
 - ThinkSystem 4-Port PCIe Gen4 NVMe Retimer Adapter (FC B98C)
 - 24 x ThinkSystem 2.5" U.2 P5520 3.84TB Read Intensive NVMe PCIe 4.0 x4 HS SSD (FC BCFR)
- ▶ Storage (RDMA) network adapter: ThinkSystem Mellanox ConnectX-6 Dx 100GbE QSFP56 2-port PCIe Ethernet Adapter (FC B8PP)

This is an all-NVMe all-flash configuration that uses NVMe storage devices in a single storage tier. This configuration is recommended when extreme storage performance is a requirement. Network bandwidth of 100GbE is recommended for this configuration in order to keep up with the NVMe storage performance. A 16-node Microsoft HCI solution built using this configuration will provide a total raw storage capacity of just under 1.5 petabytes.

SH16T1 Hybrid configuration

This configuration uses the Lenovo ThinkAgile MX630 V3 CN or IS configured with SSD storage devices at the rear of the server for the cache tier and HDD devices at the front of the server for the capacity tier. Total raw capacity of this configuration is roughly 16TB per node. Note that for this configuration an OCP version of the storage network adapter must be used, since all 3 PCIe slots are consumed by the HBA and 2 SSD devices.




Figure 6 Lenovo ThinkAgile MX630 V3 Certified Node configuration SH16T1a

Additional details include the following:

- ▶ CPU: 2 x Intel Xeon Gold 6426Y 16C 185W 2.5GHz Processor (FC BPQF)
- ▶ Memory: 16 x ThinkSystem 16GB TruDDR5 4800 MHz RDIMM (FC BKTL)
- ▶ Storage:
 - ThinkSystem 440-8i SAS/SATA PCIe Gen4 12Gb HBA (FC BM51)
 - 2 x ThinkSystem 2.5" PM1655 800GB SAS 12Gb HS SSD for the cache tier (FC BNW8)
 - 4 x ThinkSystem 3.5" 4TB 7.2K SATA 6Gb HS 512n HDD for the capacity tier (FC AUU8)
- ▶ Storage (RDMA) network adapter: ThinkSystem Intel E810-DA4 10/25GbE SFP28 4-Port OCP Ethernet Adapter (FC BP8L)
 - An OCP device is required here since all 3 PCIe slots are consumed

This is a hybrid configuration that uses SSD storage devices for the cache tier and HDD devices for the capacity tier. This configuration is recommended for small clusters (2-4 nodes) in which the total raw storage capacity requirement is relatively low. Network bandwidth of 25GbE is recommended for this configuration.

Small cluster configurations

There are a few special factors that might come into play when considering a 2- or 3-node HCI configuration. This section outlines the details that are specific to these small HCI clusters.

Direct-connect networking

For a 2- or 3-node HCI cluster, it is possible to connect the network adapters directly to each other without placing a network switch between the nodes. For a 2-node cluster using the 2-port Mellanox ConnectX-6 Lx 10/25GbE PCIe network adapter as an example, this means that Port 1 of the adapter on one node can be cabled directly into Port 1 of the second node and Port 2 from each node can be direct-connected as well. In this example, the network cables are standard SFP28 Direct Attach Cables (DACs). There is no need for a “crossover” cable.

Figure 7 shows diagrams of various network connectivity models between cluster nodes. Microsoft does not support bridged connectivity between cluster nodes and does not recommend single-link connectivity. The only recommended approach is to provide full mesh

dual-link connectivity between all nodes for east-west storage traffic. For a 3-node cluster, the only way to provide multiple network connections to each of the other two nodes without using a switch between them is by using two dual-port network adapters.




Figure 7 Various node-to-node network connectivity models

One of the most significant benefits associated with the direct-connect method is that high-speed network switches are not required for Storage traffic inside the cluster (aka “east-west” traffic). However, since the network adapters are connected to each other, a separate network connection is required from the customer network to the cluster for Management and Compute traffic (aka “north-south” traffic).

Figure 8 shows how an optional low-speed OCP network adapter (on the left) and a high-speed PCIe network adapter (on the right) are used in a 2-node direct-connected solution. Note that a high-speed OCP network adapter could be used for Storage traffic and low-speed PCIe network adapter could be used for Management and Compute traffic. In all cases, the network adapter being used must be certified to carry the traffic type intended. For details, see “Storage network adapters” on page 16 and “Non-storage network adapters” on page 17.

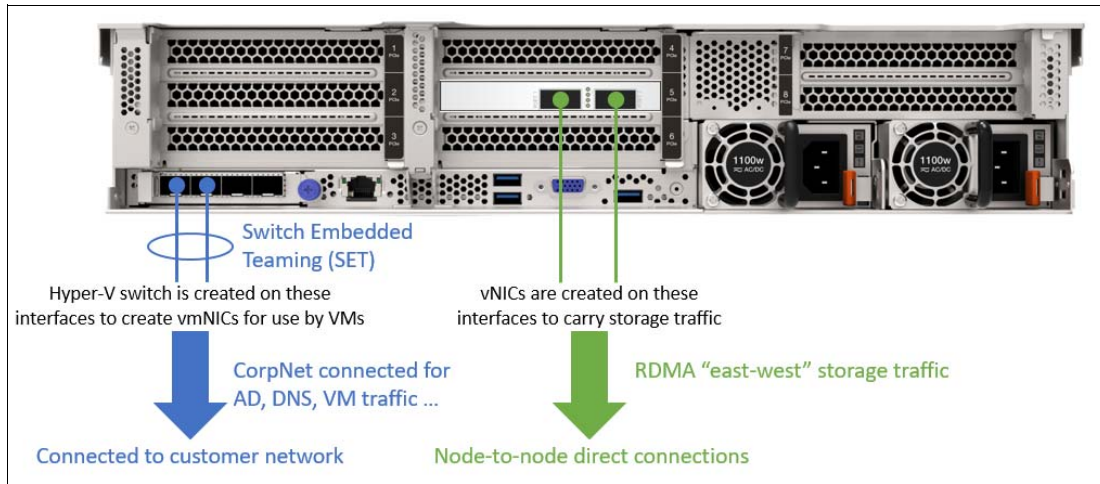


Figure 8 Diagram showing network connectivity for a ThinkSystem SR650 V3 that is part of a 2-node direct-connected HCI cluster.

USB file share witness

A feature for Microsoft 2-node HCI clusters that was introduced in Windows Server 2019 is the ability to use a "USB witness." This capability allows the requirement for a cluster witness to be satisfied by a file share configured on a USB thumb drive inserted into a network router. This reduces the complexity of cluster setup in small environments, such as branch office locations. Figure 9 illustrates the USB witness capability for a direct-connected 2-node HCI cluster.

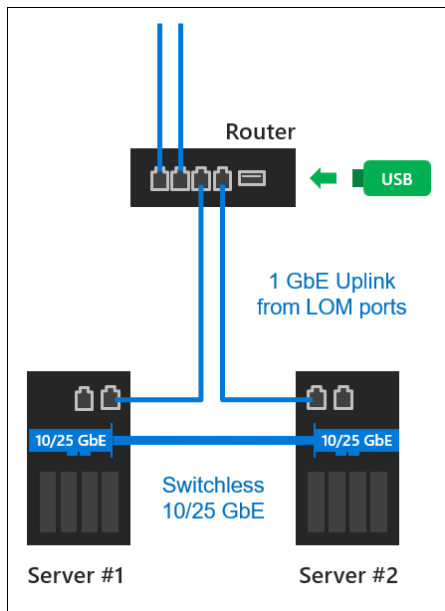


Figure 9 Illustration of USB file share witness for a direct-connected 2-node HCI cluster.

Component selection

The Lenovo certified configurations listed above are based on ThinkSystem SR630 V3 and ThinkSystem SR650 V3 rack servers and include several common hardware components. Depending on workloads and other requirements, there is some flexibility in customization of

each configuration to meet a large range of customer needs. However, the following configuration guidelines *must* be followed:

Nodes

- ▶ Lenovo ThinkSystem SR630 V3 and ThinkSystem SR650 V3 rack servers have been certified for Azure Stack HCI
 - Lenovo ThinkAgile MX630 V3 IS and MX630 V3 CN solutions are based on the SR630 V3 rack server
 - Lenovo ThinkAgile MX650 V3 IS and MX650 V3 CN solutions are based on the SR650 V3 rack server
 - See Table 2 for additional details of ThinkAgile MX solutions
- ▶ The number of nodes can range from 2 to 16

Table 2 ThinkAgile MX V3 solution details

ThinkAgile MX solution name	Rack Server	Machine Type	CTO
ThinkAgile MX630 V3 IS (Integrated System)	SR630 V3	7D6U	7D6UCTO1WW
ThinkAgile MX630 V3 CN (Certified Node)	SR630 V3	7D6U	7D6UCTO2WW
ThinkAgile MX650 V3 IS (Integrated System)	SR650 V3	7D6S	7D6SCTO1WW
ThinkAgile MX650 V3 CN (Certified Node)	SR650 V3	7D6S	7D6SCTO2WW

Processors

- ▶ Two Intel processors with a minimum of 8 cores per CPU in the Silver (4400 series) Gold (5400 or 6400 series) or Platinum (8400 series) processor families

Memory

Lenovo ThinkSystem SR630 V3 and SR650 V3 rack servers include 32 DIMMs sockets. How these sockets are populated can impact memory performance dramatically. We recommend populating in groups of 16 DIMMs to maximize performance. Table 3 shows performance characteristics of various memory configurations, including the only Unbalanced configurations that are supported (but not recommended) for use by Azure Stack HCI.

Table 3 Memory performance based on DIMM socket population

Configuration Type	Total Number of DIMMs	Relative Performance	Example
Balanced	16	100%	16 x 16GB DIMMs
Balanced	32	98%	32 x 16GB DIMMs
Near-Balanced	32 (mixed size)	97%	16 x 16GB DIMMs + 16 x 32GB DIMMs
Unbalanced	12	75%	12 x 16GB DIMMs
Unbalanced	24	72%	24 x 16GB DIMMs
Unbalanced	28	71%	28 x 16GB DIMMs

General memory requirements are as follows:

- ▶ Minimum of 256GB per node (16 x 16GB DIMMs)
- ▶ Maximum of 4TB per node (32 x 128GB DIMMs)
- ▶ Persistent memory is not currently supported

OS Boot

The following RAID controllers are supported for the OS boot drive. There are several drives to choose from for OS boot. The solution requires a pair of drives configured in a RAID-1 mirror set.

- ▶ ThinkSystem M.2 SATA/x4 NVMe 2-Bay Enablement Kit (FC BM8X)
- ▶ ThinkSystem 7mm SATA/NVMe 2-Bay Rear Enablement Kit v2 (FC BU0N)
- ▶ ThinkSystem RAID 540-8i/16i PCIe Gen4 12Gb Adapter (FC BMFT)
- ▶ M.2 SATA Array 1 RAID 1 (FC BS7C)
- ▶ ThinkSystem M.2 RAID B540i-2i SATA/NVMe Adapter (FC BYFF)

GPUs

The following two tables show supported GPUs and whether they will support GPU-P functionality upon release of an updated device driver from NVIDIA. Table 4 shows the GPUs that are supported for ThinkAgile MX solutions running on Lenovo SR630 V3 rack servers.

Table 4 Supported GPUs for SR630 V3 servers

Description	GPU-P Support	Feature Code
ThinkSystem NVIDIA A2 16GB PCIe Gen4 Passive GPU w/o CEC	Yes	BQZT

Table 5 shows the GPUs that are supported for ThinkAgile MX solutions running on Lenovo SR650 V3 rack servers.

Table 5 Supported GPUs for SR650 V3 servers

Description	GPU-P Support	Feature Code
ThinkSystem NVIDIA A2 16GB PCIe Gen4 Passive GPU w/o CEC	Yes	BQZT
ThinkSystem NVIDIA A16 64GB PCIe Gen4 Passive GPU	Yes	BNFE
ThinkSystem NVIDIA A16 64GB PCIe Gen4 Passive GPU w/o CEC	Yes	BQZU
ThinkSystem NVIDIA A30 24GB PCIe Gen4 Passive GPU	No	BJHG
ThinkSystem NVIDIA A40 48GB PCIe Gen4 Passive GPU w/o CEC	Yes	BQZQ
ThinkSystem NVIDIA A100 40GB PCIe Gen4 Passive GPU w/o CEC	No	BQZP

Storage network adapters

We recommend 25GbE or 100GbE, depending on storage configuration, for optimal performance. Storage network adapters are shown below by RDMA protocol.

- ▶ For RoCE v2:

- ThinkSystem Mellanox ConnectX-6 Lx 10/25GbE SFP28 2-Port OCP Ethernet Adapter (FC BE4T)
- ThinkSystem Mellanox ConnectX-6 Lx 10/25GbE SFP28 2-Port PCIe Ethernet Adapter (FC BE4U)
- ThinkSystem Mellanox ConnectX-6 Dx 100GbE QSFP56 2-Port PCIe Ethernet Adapter (FC B8PP) Use this NIC for all-flash storage configurations when the additional throughput is required.

Note: Network switches must support the RoCE v2 feature set for best storage performance. See “Network switches” on page 22 for more information regarding Lenovo and NVIDIA/Mellanox network switches that have been tested with ThinkAgile MX solutions.

- ▶ For iWARP:
 - ThinkSystem Intel E810-DA2 10/25GbE SFP28 2-Port OCP Ethernet Adapter (FC BCD4)
 - ThinkSystem Intel E810-DA4 10/25GbE SFP28 4-Port OCP Ethernet Adapter (FC BP8L)
 - ThinkSystem Intel E810-DA2 10/25GbE SFP28 2-Port PCIe Ethernet Adapter (FC BCD6)
 - ThinkSystem Intel E810-DA4 10/25GbE SFP28 4-Port PCIe Ethernet Adapter (FC BP8M)

Non-storage network adapters

Network adapters that support Management and Compute traffic (but not storage traffic) are shown below.

- ▶ OCP Network Adapters
 - ThinkSystem Intel X710-T2L 10GBASE-T 2-Port OCP Ethernet Adapter (FC BCD5)
 - ThinkSystem Intel X710-T4L 10GBASE-T 4-Port OCP Ethernet Adapter (FC BPPY)
 - ThinkSystem Broadcom 57414 10/25GbE SFP28 2-Port OCP Ethernet Adapter (FC BN2T)
 - ThinkSystem Broadcom 57416 10GBASE-T 2-Port OCP Ethernet Adapter (FC B5ST)
 - ThinkSystem Broadcom 57454 10GBASE-T 4-Port OCP Ethernet Adapter (FC B5T4)
 - ThinkSystem Broadcom 57504 10/25GbE SFP28 4-Port OCP Ethernet Adapter (FC BPPW)
- ▶ PCIe Network Adapters
 - ThinkSystem Intel X710-T2L 10GBASE-T 2-Port PCIe Ethernet Adapter (FC BNWL)
 - ThinkSystem Intel X710-T4L 10GBASE-T 4-Port PCIe Ethernet Adapter (FC BMXB)
 - ThinkSystem Broadcom 57454 10GBASE-T 4-Port PCIe Ethernet Adapter (FC B5SU)

Storage HBAs

- ▶ ThinkSystem 440-8i SAS/SATA PCIe Gen4 12Gb HBA (FC BM51)
- ▶ ThinkSystem 440-16i SAS/SATA PCIe Gen4 12Gb HBA (FC BM50)
- ▶ ThinkSystem 440-16i SAS/SATA PCIe Gen4 12Gb Internal HBA (FC B8P1) For MX650 V3 solutions only.
- ▶ ThinkSystem 4350-8i SAS/SATA 12Gb HBA (FC BJHH)
- ▶ ThinkSystem 4350-16i SAS/SATA 12Gb HBA (FC BJHJ)
- ▶ ThinkSystem 4-Port PCIe Gen4 NVMe Retimer Adapter (FC B98C) Technically, this is not an HBA, but serves the same purpose for NVMe devices.

Storage devices

- ▶ Hybrid solutions can only use two-tier storage configurations
- ▶ All-flash solutions can use single-tier or two-tier storage configurations
- ▶ Drive rules for single-tier storage configurations are as follows:
 - Drives can be either SAS/SATA SSD or NVMe SSD (includes EDSFF)
 - All drives must be the same feature code
 - A minimum of 4 drives is required
- ▶ Drive rules for two-tier storage configurations are as follows:
 - A minimum of 2 cache drives and 4 capacity drives is required
 - All cache drives must be the same feature code
 - All capacity drives must be the same feature code
 - Storage device types must match one of the following:
 - SAS SSD devices for cache (high endurance devices) and SAS/SATA SSD devices for capacity (low endurance devices)
 - NVMe SSD devices for both cache and capacity
 - NVMe SSD devices for cache and SAS/SATA SSD devices for capacity

We strongly recommend a minimum 10% cache to capacity ratio (e.g. 2x 800GB SSD and 4x 4TB HDD). Although this is not a requirement, care should be taken to provide enough cache space for the amount of capacity available in the solution or performance can be impacted significantly.

Table 6 provides a list of all certified Lenovo storage devices that can be used to configure a Hybrid Storage HCI solution based on the Lenovo ThinkSystem SR630 V3 and SR650 V3 rack servers. This table does not include OS boot devices, which are covered in “Component selection” on page 14.

Table 6 *Lenovo storage devices certified for Microsoft Azure Stack HCI in ThinkAgile MX Hybrid solutions*

Storage Devices Used for Lenovo ThinkAgile MX Hybrid Solutions	FC	Type	Usage
2.5" Small Form Factor devices			
ThinkSystem 2.5" U.2 P5620 1.6TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEG	NVMe	Cache
ThinkSystem 2.5" U.2 P5620 3.2TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEH	NVMe	Cache
ThinkSystem 2.5" U.2 P5620 6.4TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEZ	NVMe	Cache
ThinkSystem 2.5" U.2 P5620 12.8TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BA4V	NVMe	Cache
ThinkSystem 2.5" U.3 7450M 800GB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNF1	NVMe	Cache
ThinkSystem 2.5" U.3 7450M 1.6TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEJ	NVMe	Cache
ThinkSystem 2.5" U.3 7450M 3.2TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEY	NVMe	Cache
ThinkSystem 2.5" U.3 7450M 6.4TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEL	NVMe	Cache
ThinkSystem 2.5" PM1655 800GB Mixed Use SAS 12Gb HS SSD	BNW8	SSD	Cache
ThinkSystem 2.5" PM1655 1.6TB Mixed Use SAS 12Gb HS SSD	BNW9	SSD	Cache
ThinkSystem 2.5" PM1655 3.2TB Mixed Use SAS 12Gb HS SSD	BNW6	SSD	Cache
ThinkSystem 2.5" PM1655 6.4TB Mixed Use SAS 12Gb HS SSD	BP3K	SSD	Cache
ThinkSystem 2.5" 600GB 10K SAS 12Gb HS 512n HDD	AULZ	HDD	Capacity
ThinkSystem 2.5" 1.2TB 10K SAS 12Gb HS 512n HDD	AUM1	HDD	Capacity
ThinkSystem 2.5" 1.8TB 10K SAS 12Gb HS 512e HDD	AUM2	HDD	Capacity

Storage Devices Used for Lenovo ThinkAgile MX Hybrid Solutions	FC	Type	Usage
ThinkSystem 2.5" 2.4TB 10K SAS 12Gb HS 512e HDD	BOYS	HDD	Capacity
ThinkSystem 2.5" 600GB 15K SAS 12Gb HS 512n HDD	AULW	HDD	Capacity
ThinkSystem 2.5" 900GB 15K SAS 12Gb HS 512e HDD	AULX	HDD	Capacity
3.5" Large Form Factor devices			
ThinkSystem 3.5" U.2 P5800X 800GB Write Intensive NVMe PCIe 4.0 x4 HS SSD	BMM5	NVMe	Cache
ThinkSystem 3.5" U.2 P5800X 1.6TB Write Intensive NVMe PCIe 4.0 x4 HS SSD	BMM6	NVMe	Cache
ThinkSystem 3.5" U.2 P5620 1.6TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEK	NVMe	Cache
ThinkSystem 3.5" U.2 P5620 3.2TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEM	NVMe	Cache
ThinkSystem 3.5" U.2 P5620 6.4TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEN	NVMe	Cache
ThinkSystem 3.5" U.2 P5620 12.8TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEP	NVMe	Cache
ThinkSystem 3.5" PM1655 800GB Mixed Use SAS 12Gb HS SSD	BNW7	SSD	Cache
ThinkSystem 3.5" PM1655 1.6TB Mixed Use SAS 12Gb HS SSD	BNWA	SSD	Cache
ThinkSystem 3.5" PM1655 3.2TB Mixed Use SAS 12Gb HS SSD	BNWB	SSD	Cache
ThinkSystem 3.5" PM1655 6.4TB Mixed Use SAS 12Gb HS SSD	BP3G	SSD	Cache
ThinkSystem 3.5" 1TB 7.2K SATA 6Gb HS 512n HDD	AUUF	HDD	Capacity
ThinkSystem 3.5" 4TB 7.2K SATA 6Gb HS 512n HDD	AUU8	HDD	Capacity
ThinkSystem 3.5" 6TB 7.2K SATA 6Gb HS 512e HDD	AUUA	HDD	Capacity
ThinkSystem 3.5" 8TB 7.2K SATA 6Gb HS 512e HDD	AUU9	HDD	Capacity
ThinkSystem 3.5" 10TB 7.2K SATA 6Gb HS 512e HDD	AUUB	HDD	Capacity
ThinkSystem 3.5" 12TB 7.2K SATA 6Gb HS 512e HDD	B118	HDD	Capacity
ThinkSystem 3.5" 14TB 7.2K SATA 6Gb HS 512e HDD	B497	HDD	Capacity
ThinkSystem 3.5" 16TB 7.2K SATA 6Gb HS 512e HDD	B7F0	HDD	Capacity
ThinkSystem 3.5" 4TB 7.2K SAS 12Gb HS 512n HDD	AUU6	HDD	Capacity
ThinkSystem 3.5" 6TB 7.2K SAS 12Gb HS 512n HDD	AUU7	HDD	Capacity
ThinkSystem 3.5" 8TB 7.2K SAS 12Gb HS 512n HDD	BOYR	HDD	Capacity
ThinkSystem 3.5" 10TB 7.2K SAS 12Gb HS 512n HDD	AUUG	HDD	Capacity
ThinkSystem 3.5" 12TB 7.2K SAS 12Gb HS 512n HDD	B117	HDD	Capacity
ThinkSystem 3.5" 14TB 7.2K SAS 12Gb HS 512n HDD	B496	HDD	Capacity
ThinkSystem 3.5" 16TB 7.2K SAS 12Gb Hot Swap 512e HDD	B7EZ	HDD	Capacity

Table 7 provides a list of all certified Lenovo storage devices that can be used to configure an All-Flash HCI solution based on the Lenovo ThinkSystem SR630 V3 and SR650 V3 rack servers. Note that 3.5" Large Form Factor drives are not supported for All-Flash solutions. This table does not include OS boot devices, which are covered in "Component selection" on page 14.

Table 7 Lenovo storage devices certified for Microsoft Azure Stack HCI in ThinkAgile MX All-Flash solutions

Storage Devices Used for Lenovo ThinkAgile MX All-Flash Solutions	FC	Type	Usage ¹
ThinkSystem 2.5" U.2 P5800X 800GB Write Intensive NVMe PCIe 4.0 x4 HS SSD	BKKZ	NVMe	Cache ²
ThinkSystem 2.5" U.2 P5800X 1.6TB Write Intensive NVMe PCIe 4.0 x4 HS SSD	BMM8	NVMe	Cache ²
ThinkSystem 2.5" U.2 P5620 1.6TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEG	NVMe	Cache ²
ThinkSystem 2.5" U.2 P5620 3.2TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEH	NVMe	Cache ²
ThinkSystem 2.5" U.2 P5620 6.4TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEZ	NVMe	Cache ²
ThinkSystem 2.5" U.2 P5620 12.8TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BA4V	NVMe	Cache ²
ThinkSystem 2.5" U.3 7450M 800GB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNF1	NVMe	Cache ²
ThinkSystem 2.5" U.3 7450M 1.6TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEJ	NVMe	Cache ²
ThinkSystem 2.5" U.3 7450M 3.2TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEY	NVMe	Cache ²
ThinkSystem 2.5" U.3 7450M 6.4TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEL	NVMe	Cache ²
ThinkSystem 2.5" PM1655 800GB Mixed Use SAS 12Gb HS SSD	BNW8	SSD	Cache ²
ThinkSystem 2.5" PM1655 1.6TB Mixed Use SAS 12Gb HS SSD	BNW9	SSD	Cache ²
ThinkSystem 2.5" PM1655 3.2TB Mixed Use SAS 12Gb HS SSD	BNW6	SSD	Cache ²
ThinkSystem 2.5" PM1655 6.4TB Mixed Use SAS 12Gb HS SSD	BP3K	SSD	Cache ²
ThinkSystem E1.S 5.9mm 7450P 3.84TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BP3L	NVMe	Capacity
ThinkSystem E1.S 5.9mm 7450P 7.68TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BPKW	NVMe	Capacity
ThinkSystem 2.5" U.3 7450P 960GB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BNF3	NVMe	Capacity
ThinkSystem 2.5" U.3 7450P 1.92TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BNF2	NVMe	Capacity
ThinkSystem 2.5" U.3 7450P 3.84TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BNF5	NVMe	Capacity
ThinkSystem 2.5" U.3 7450P 7.68TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BNF4	NVMe	Capacity
ThinkSystem 7mm U.3 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BPZ4	NVMe	Capacity
ThinkSystem 7mm U.3 7450 PRO 1.92TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BPZ5	NVMe	Capacity
ThinkSystem 7mm U.3 7450 PRO 3.84TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BPZ6	NVMe	Capacity
ThinkSystem 2.5" U.2 P5520 1.92TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BMGD	NVMe	Capacity
ThinkSystem 2.5" U.2 P5520 3.84TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BMGE	NVMe	Capacity
ThinkSystem 2.5" U.2 P5520 7.68TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BNEF	NVMe	Capacity
ThinkSystem 2.5" U.2 P5520 15.36TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BNEQ	NVMe	Capacity
ThinkSystem 2.5" Intel S4620 480GB Mixed Use SATA 6Gb HS SSD	BA7Q	SSD	Capacity
ThinkSystem 2.5" Intel S4620 960GB Mixed Use SATA 6Gb HS SSD	BA4T	SSD	Capacity
ThinkSystem 2.5" Intel S4620 1.92TB Mixed Use SATA 6Gb HS SSD	BA4U	SSD	Capacity
ThinkSystem 2.5" Intel S4620 3.84TB Mixed Use SATA 6Gb HS SSD	BK7L	SSD	Capacity
ThinkSystem 2.5" U.2 P5520 1.92TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BMGD	NVMe	Capacity
ThinkSystem 2.5" U.2 P5520 3.84TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BMGE	NVMe	Capacity

Storage Devices Used for Lenovo ThinkAgile MX All-Flash Solutions	FC	Type	Usage¹
ThinkSystem 2.5" U.2 P5520 7.68TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BNEF	NVMe	Capacity
ThinkSystem 2.5" U.2 P5520 15.36TB Read Intensive NVMe PCIe 4.0 x4 HS SSD	BNEQ	NVMe	Capacity
ThinkSystem 2.5" PM1653 960GB Read Intensive SAS 12Gb HS SSD	BNWC	SSD	Capacity
ThinkSystem 2.5" PM1653 1.92TB Read Intensive SAS 12Gb HS SSD	BNWE	SSD	Capacity
ThinkSystem 2.5" PM1653 3.84TB Read Intensive SAS 12Gb HS SSD	BNWF	SSD	Capacity
ThinkSystem 2.5" PM1653 7.68TB Read Intensive SAS 12Gb HS SSD	BP3E	SSD	Capacity
ThinkSystem 2.5" PM1653 15.36TB Read Intensive SAS 12Gb HS SSD	BP3J	SSD	Capacity
ThinkSystem 2.5" PM1653 30.72TB Read Intensive SAS 12Gb HS SSD	BP3D	SSD	Capacity
ThinkSystem 2.5" S4520 480GB Read Intensive SATA 6Gb HS SSD	BA7G	SSD	Capacity
ThinkSystem 2.5" S4520 960GB Read Intensive SATA 6Gb HS SSD	BA7H	SSD	Capacity
ThinkSystem 2.5" S4520 1.92TB Read Intensive SATA 6Gb HS SSD	BA7J	SSD	Capacity
ThinkSystem 2.5" S4520 3.84TB Read Intensive SATA 6Gb HS SSD	BK77	SSD	Capacity
ThinkSystem 2.5" S4520 7.68TB Read Intensive SATA 6Gb HS SSD	BK78	SSD	Capacity
ThinkSystem 2.5" 5400 MAX 480GB Mixed Use SATA 6Gb HS SSD	BQ21	SSD	Capacity
ThinkSystem 2.5" 5400 MAX 960GB Mixed Use SATA 6Gb HS SSD	BQ24	SSD	Capacity
ThinkSystem 2.5" 5400 MAX 1.92TB Mixed Use SATA 6Gb HS SSD	BQ22	SSD	Capacity
ThinkSystem 2.5" 5400 MAX 3.84TB Mixed Use SATA 6Gb HS SSD	BQ23	SSD	Capacity
ThinkSystem 7mm 5400 PRO 480GB Read Intensive SATA 6Gb HS SSD	BQ1V	SSD	Capacity
ThinkSystem 7mm 5400 PRO 960GB Read Intensive SATA 6Gb HS SSD	BQ1W	SSD	Capacity
ThinkSystem 7mm 5400 PRO 1.92TB Read Intensive SATA 6Gb HS SSD	BR13	SSD	Capacity
ThinkSystem 7mm 5400 PRO 3.84TB Read Intensive SATA 6Gb HS SSD	BR12	SSD	Capacity
ThinkSystem 7mm 5400 PRO 7.68TB Read Intensive SATA 6Gb HS SSD	BR11	SSD	Capacity
ThinkSystem 2.5" 5400 PRO 480GB Read Intensive SATA 6Gb HS SSD	BQ1P	SSD	Capacity
ThinkSystem 2.5" 5400 PRO 960GB Read Intensive SATA 6Gb HS SSD	BQ1R	SSD	Capacity
ThinkSystem 2.5" 5400 PRO 1.92TB Read Intensive SATA 6Gb HS SSD	BQ1X	SSD	Capacity
ThinkSystem 2.5" 5400 PRO 3.84TB Read Intensive SATA 6Gb HS SSD	BQ1S	SSD	Capacity
ThinkSystem 2.5" 5400 PRO 7.68TB Read Intensive SATA 6Gb HS SSD	BQ1T	SSD	Capacity
ThinkSystem 2.5" PM893 480GB Read Intensive SATA 6Gb HS SSD	BM8B	SSD	Capacity
ThinkSystem 2.5" PM893 960GB Read Intensive SATA 6Gb HS SSD	BM8A	SSD	Capacity
ThinkSystem 2.5" PM893 1.92TB Read Intensive SATA 6Gb HS SSD	BM89	SSD	Capacity
ThinkSystem 2.5" PM893 3.84TB Read Intensive SATA 6Gb HS SSD	BM88	SSD	Capacity
ThinkSystem 2.5" PM893 7.68TB Read Intensive SATA 6Gb HS SSD	BM87	SSD	Capacity

¹ Do not use a storage device for a purpose other than listed in the Usage column.

² This device is supported as a Cache device in two-tier all-flash configurations and as a Capacity device in single tier all-flash configurations (i.e. all-NVMe or all-SSD).

Storage device end of life

More than any other component in a certified solution, the storage devices available are constantly changing as new, faster, larger devices are brought to market and previous generations reach their end of life. Table 8 provides details on which devices have reached or are nearing their projected end of life, including estimated last availability date and replacement device (if one is available).

Table 8 Storage device end of life summary

End of Life Device	Date	Replacement Device
No devices approaching end of life		

Network switches

Since Lenovo no longer sells network switches, we have tested NVIDIA (Mellanox) network switches in our labs. Mellanox switches can be ordered directly from NVIDIA or through your Lenovo business partner/reseller.

NVIDIA/Mellanox network switches

Although NVIDIA/Mellanox switches are not orderable from Lenovo, the following Mellanox network switches have been tested with ThinkAgile MX solutions and proper switch functionality has been verified:

NVIDIA MSN2010-CB2F Spectrum Based 25GbE/100GbE with Cumulus Linux OS
 1U, Half-Width Open Ethernet switch with 18 SFP28 and 4 QSFP28 Ports
<https://www.mellanox.com/sites/default/files/doc-2020/br-sn2000-series.pdf>
<https://www.mellanox.com/sites/default/files/doc-2020/pb-sn2010.pdf>

NVIDIA MSN2410-CB2F Spectrum Based 25GbE/100GbE with Cumulus Linux OS
 1U, Full-Width Open Ethernet switch with 48 SFP28 Ports 8 QSFP28 Ports
<https://www.mellanox.com/sites/default/files/doc-2020/br-sn2000-series.pdf>
<https://www.mellanox.com/sites/default/files/doc-2020/pb-sn2410.pdf>

NVIDIA MSN3700-CS2F Spectrum-2 Based 100GbE with Cumulus Linux OS
 1U, Full-Width Open Ethernet switch with 32 QSFP28 Ports
<https://www.mellanox.com/files/doc-2020/br-sn3000-series.pdf>

Other recommendations

We also recommend the features and upgrades in this section to maximize the security and manageability of the S2D solution built using the Lenovo certified configurations discussed in this document.

TPM 2.0 and Secure Boot

Trusted Platform Module (TPM) is an international standard for a secure cryptoprocessor, a dedicated microcontroller designed to secure hardware through integrated cryptographic

keys. TPM technology is designed to provide hardware-based, security-related functions and is used extensively by Microsoft in Windows Server technologies including BitLocker, Device Guard, Credential Guard, UEFI Secure Boot, and others. There is no additional cost to enable TPM 2.0 on Lenovo ThinkSystem servers.

For ThinkAgile MX solutions that are based on the SR630 V3 and SR650 V3 rack servers, order Feature Code BPKR to enable TPM 2.0 or Feature Code BPKQ to enable TPM 2.0 and Secure Boot. Note that Secure Boot is required by Microsoft for all ThinkAgile MX Integrated System (Appliance) solutions, so will be selected by default for all Integrated System configurations.

Note: TPM is not supported in PRC. For systems shipped to China, NationZ TCM is used and supported.

ThinkSystem XClarity Controller Standard to Enterprise Level

The Lenovo XClarity™ Controller is the next generation management controller that replaces the baseboard management controller (BMC) for Lenovo ThinkSystem servers. Although the XCC Standard Level includes many important manageability features, we recommend upgrading to the XCC Enterprise Level of functionality. This enhanced set of features includes Virtual Console (out of band browser-based remote control), Virtual Media mounting, and other remote management capabilities.

Lenovo XClarity Pro

Lenovo XClarity Administrator (LXCA) is a centralized resource management solution that is aimed at reducing complexity, speeding response, and enhancing the availability of Lenovo server systems and solutions. LXCA provides agent-free hardware management for our servers, storage, network switches, hyperconverged and ThinkAgile solutions. Lenovo XClarity Pro offers additional functionality that provide important benefits to managing a Microsoft Azure Stack HCI cluster solution. For more information, see the LXCA Product Guide at the following URL:

<https://lenovopress.com/tips1200-lenovo-xclarity-administrator>

Lenovo XClarity Integrator for Microsoft Windows Admin Center

Lenovo XClarity Integrator for Microsoft Windows Admin Center (LXCI for WAC) provides IT administrators with a smooth and seamless experience in managing Lenovo servers. Using WAC's Server Manager or Cluster Manager extension, IT administrators can manage Lenovo servers as single hosts or directly manage them as Microsoft Windows Failover clusters. In addition, they are able to manage Azure Stack HCI clusters as well as Lenovo ThinkAgile MX Integrated Systems and Certified Nodes through the LXCI snap-ins integrated into WAC's cluster creation and Cluster-Aware Updating (CAU) functions. The LXCI for WAC extension simplifies server management of IT administrators, making it possible to remotely manage servers throughout their life cycle in a single unified UI. For more information, see the LXCI for WAC Information Center at the following URL:

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

Summary

Lenovo is a key partner in the Microsoft WSSD and Azure Stack HCI programs for certification of HCI solutions. Based on Lenovo's investment in these programs and the

tremendous amount of time, resources, and effort dedicated to certification and validation testing for each certified configuration discussed in this document, our customers can rest assured that the configurations presented will perform smoothly and reliably right out of the box.

This document has provided some background information related to the Microsoft WSSD and Azure Stack HCI programs, as well as details of Lenovo certified configurations that have been certified and validated under the program to run Storage Spaces Direct. Selecting from the list of Lenovo certified configurations found in this document to build an S2D HCI solution will save time, money, and effort associated with designing and building a do-it-yourself solution that could be riddled with issues.

We will add more examples of Lenovo certified configurations for Microsoft HCI solutions to this document as additional certifications are completed.

Change History

Changes in the June 2024 update:

- ▶ Updated “OS Boot” on page 16, adding the ThinkSystem M.2 RAID B540i-2i SATA/NVMe Adapter as a supported RAID adapter for the boot volume.

Changes in the December 2023 update:

- ▶ Updated “Storage network adapters” on page 16 and “Non-storage network adapters” on page 17 to clarify which network adapters have been certified for each traffic type (storage, management, and compute).

Authors

This paper was produced by the following 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:

- ▶ Emanuel Andreescu, Microsoft Solutions Engineer
- ▶ Guy Fusman, Microsoft Cloud Solution Engineer
- ▶ Daniel Ghidali, Manager - Microsoft Technology and Enablement
- ▶ Vinay Kulkarni, Principal Technical Consultant - Microsoft Solutions and Enablement
- ▶ Laurentiu Petre, Solutions Engineer - ISG Software Development
- ▶ Vy Phan, Technical Program Manager - Microsoft OS and Solutions
- ▶ Garrett Tashiro, Microsoft Solutions Intern
- ▶ David Watts, Senior IT Consultant - Lenovo Press

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service.

Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
1009 Think Place - Building One
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary.

Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

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

Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

This document was created or updated on May 6, 2024.

Send us your comments via the **Rate & Provide Feedback** form found at <http://lenovopress.com/lp1741>

Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at <http://www.lenovo.com/legal/copytrade.html>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®	Lenovo(logo)®	TruDDR4™
Lenovo XClarity™	ThinkAgile™	
RackSwitch™	ThinkSystem™	

The following terms are trademarks of other companies:

Intel, Xeon, and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

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