Lenovo ThinkAgile MX Certified Configurations for Azure Local - V1 Servers

Last Update: January 2025

Provides details of certified configurations for Lenovo ThinkAgile MX solutions that are based on SR650 rack servers

Applies to both Azure Local and Windows Server S2D

Provides guidance for properly configuring nodes for an Azure Local or Windows Server S2D cluster

Lists supported options that can be used when configuring Azure Local or Windows Server S2D cluster nodes

Dave Feisthammel

Guy Fusman

Mike Miller

David Ye

LENOVO PRESS

Abstract

This document provides background information regarding the Microsoft Windows Server Software-Defined (WSSD) certification program for Azure Local and Windows Server S2D, 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 Local and Windows Server S2D that are based on ThinkSystem™ SR650 servers, including processor, memory, network, and storage components available for each cluster node. This includes the following solutions:

- ThinkAgile MX3520-H Hybrid Appliance
- ThinkAgile MX3520-F All-Flash Appliance
- ThinkAgile MX Certified Node on SR650

Looking for Lenovo ThinkAgile MX solutions that are based on our V2, V3, or Edge servers? Check our companion documents.

- For ThinkSystem V2 rack servers: http://lenovopress.com/lp1520
- For ThinkSystem V3 rack servers: http://lenovopress.com/lp1741
- For ThinkSystem Edge servers: http://lenovopress.com/lp1984

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

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.

Host operating system Azure Local 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 Local operating system or take advantage of Azure Local 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 scale-on-demand compute/storage integrated solutions
- Easily provisioned new IT services that reduce deployment time
- Better performance and lower Total Cost of Ownership (TCO)
- Flexible infrastructure and data centers

For decades Lenovo has worked closely with Microsoft 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 Local certification program.

Deploying Lenovo certified configurations for Microsoft HCI solutions takes the guesswork out of system configuration. Whether you intend to build an Azure Local instance or a Windows Server S2D cluster (disaggregated or hyperconverged), you can rest assured that purchasing a certified configuration will provide a rock solid foundation with minimal obstacles along the way. All node configurations shown in this document are certified by Lenovo and validated by Microsoft for out-of-the-box optimization. Using Lenovo ThinkAgile MX solution 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 presents the Lenovo certified configurations for ThinkSystem SR650 servers that have been validated for use in a Microsoft HCI solution under the Microsoft Windows Server-Software Defined (WSSD) certification program. Details of each node are specified, including all key components. Since there is latitude for customization in these configurations, the rules for customization are also described.

For an overview of the Microsoft Azure Local solution, visit the following URL:

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

To help understand the technical value of the WSSD program and learn why deploying a certified configuration for Azure Local or Windows Server S2D provides an optimal path to success, read this Microsoft blog post:

https://cloudblogs.microsoft.com/windowsserver/2018/02/20/the-technical-value-of-wssd-validated-hci-solutions

To see all Lenovo certified solutions in the Microsoft Azure Local Catalog, visit:

https://azurestackhcisolutions.azure.microsoft.com/#/catalog?vendorName=Lenovo

ThinkAgile MX Series solutions

The Microsoft WSSD program allows OEM partners to deliver pre-engineered, validated HCI solutions. Whether your preference is for a Certified Node, an Integrated System, or a Premier Solution, Lenovo has designed, tested and validated the ThinkAgile MX Series offerings to quickly and easily provide the solutions you need. This results in a quickly deployable, robust, and high-performance hyperconverged solution that can help to rapidly solve your IT challenges.

ThinkAgile MX Certified Node (CN)

Lenovo ThinkAgile MX Certified Nodes map to Microsoft “Azure Local Validated Nodes” in the Microsoft Azure Local Catalog (see link above). These solutions package Microsoft-certified HCI capabilities into easy-to-use machine types to provide the following:

- Easy to order
- Enforced configuration rules to ensure a valid configuration
- Best recipe firmware and device drivers
- Premier Support (where available)
- Optional services such as deployment, management, etc.

ThinkAgile MX Integrated System (IS)

Lenovo ThinkAgile MX Integrated Systems (some models known as Appliances) map to Microsoft “Azure Local Integrated Systems” in the Microsoft Azure Local Catalog (see link above). These solutions are based on exactly the same hardware as ThinkAgile MX Certified Nodes. The only differences between a ThinkAgile MX Certified Node and Integrated System that are based on the same server (for example, the ThinkSystem SR650 rack server) is that the Integrated System configuration includes the following items:

- Azure Local operating system, along with the latest Best Recipe device drivers, is preloaded before shipping to the customer
- ThinkAgile Premier Support for 3 years (can be uplifted to a longer term or quicker response time)

The remainder of this document focuses on describing the existing Lenovo configurations that have been certified under the Microsoft WSSD program 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. This document assumes the reader has prior knowledge of Microsoft HCI technologies, including Azure Local and Windows Server S2D.

Lenovo certified configurations for Microsoft Azure Local

The Microsoft WSSD certification program allows for solution certification using a min/max paradigm. The configurations presented in this document represent examples of what has been certified, rather than an exhaustive list of the certified configurations that are available. Refer to “Component selection” on page 15 for additional information regarding the components that have been certified.

Table 1 lists the key components of the example configurations for Azure Local or Windows Server S2D that have been certified under the Microsoft WSSD program. Depending on the

configuration, the number of nodes can range from 1 to 16 in a single cluster. Note that we regularly certify 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 NH80T**1** and NH80T**2**.

Table 1 Example configuration highlights for Lenovo ThinkAgile MX solutions using SR650 rack servers[1]

Config	Server/CPU/RAM	Cache	Capacity	Storage Controller	Storage Network
SH40T1 (hybrid)	MX3520-H (SR650) 2 CPUs 192GB-1.5TB	4 x 800GB SSD FC: B170	10 x 4TB FC: AUU8	430-16i FC: AUNM	Mellanox CX-4 2-port 25GbE FC: AUAJ[2]
SH60T1 (hybrid)	MX3520-H (SR650) 2 CPUs 192GB-1.5TB	4 x 1.6TB SSD FC: B171	10 x 6TB FC: AUUA	430-16i FC: AUNM	Mellanox CX-4 2-port 25GbE FC: AUAJ[2]
NH80T1 (hybrid)	MX3520-H (SR650) 2 CPUs 192GB-1.5TB	4 x 3.2TB NVMe (U.2) FC: B2XG	10 x 8TB FC: AUU9	430-16i FC: AUNM	Mellanox CX-4 2-port 25GbE FC: AUAJ[2]
NH120T1 (hybrid)	MX3520-H (SR650) 2 CPUs 192GB-1.5TB	4 x 3.2TB NVMe (U.2) FC: B2XG	10 x 12TB FC: B118	430-16i FC: AUNM	Mellanox CX-4 2-port 25GbE FC: AUAJ[2]
NS61T1 (all-flash)	MX3520-F (SR650) 2 CPUs 192GB-1.5TB	4 x 750GB Optane NVMe (U.2) FC: B2ZJ	16 x 3.84TB SSD FC: B49C	3 x 430-8i FC: AUNL	Mellanox CX-4 2-port 100GbE FC: ATRP[3]
NS77T1 (all-flash)	MX3520-F (SR650) 2 CPUs 192GB-1.5TB	4 x 3.2TB NVMe (U.2) FC: B11K	20 x 3.84TB SSD FC: B49C	3 x 430-8i FC: AUNL	Mellanox CX-4 2-port 100GbE FC: ATRP[3]
NN38T1 (all-flash)	MX3520-F (SR650) 2 CPUs 192GB-1.5TB	12 x 3.2TB NVMe (U.2) FC: B11K (all-NVMe)		3 x 430-8i FC: AUNL	Mellanox CX-4 2-port 100GbE FC: ATRP[3]
SS92T1 (all-flash)	MX3520-F (SR650) 2 CPUs 192GB-1.5TB	24 x 3.84TB SSD FC: B49C (all-SSD)		3 x 430-8i FC: AUNL	Mellanox CX-4 2-port 100GbE FC: ATRP[3]

[1] This list is not exhaustive and can be customized. Refer to the “Component selection” on page 15 for information about customizing these configurations. All configurations use a dual 480GB M.2 SSD configured as a RAID-1 mirrored volume for OS boot.

[2] Mellanox CX-4 1-port 40GbE (FC ATRN) and ThinkSystem Intel E810-DA2 10/25GbE SFP28 2-Port PCIe Ethernet Adapter (FC BCD6) are also certified for this configuration.

[3] Mellanox CX-4 2-port 25GbE (FC AUAJ) and ThinkSystem Intel E810-DA2 10/25GbE SFP28 2-Port PCIe Ethernet Adapter (FC BCD6) are also certified for this configuration.

Lenovo certified configuration details

This section includes details for each example configuration contained in Table 1 that have been certified under the Microsoft WSSD certification program. Each configuration lists the

Lenovo ThinkAgile MX Certified Node or ThinkSystem rack server that is used for the Azure Local or Windows Server 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 15 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.

SH40T1 hybrid configuration

This configuration uses the Lenovo ThinkAgile MX Certified Node configured with SSDs for the cache tier and HDDs for the capacity tier. Total raw capacity of this configuration is 40TB per node.

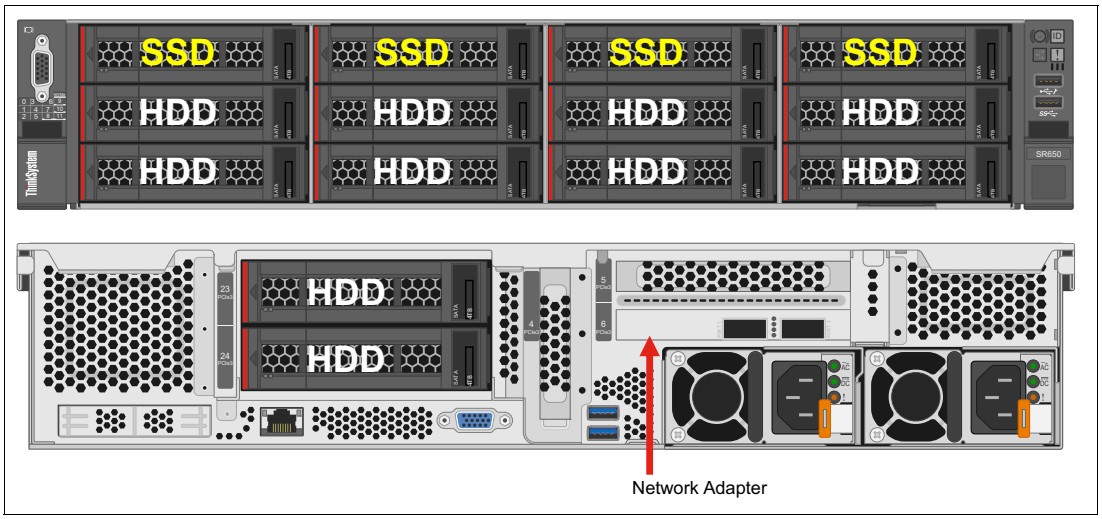

Figure 1 Lenovo ThinkAgile MX Certified Node configuration SH40T1

Additional details include the following:

- Network adapter: The following network adapters have been certified:
  - Mellanox ConnectX-4 2-port 10/25GbE Ethernet Adapter (FC AUAJ)
  - 2 x Mellanox ConnectX-4 1-port 40GbE Ethernet Adapter (FC ATRN)
    - 2 x Mellanox QSA 100G to 25G Cable Adapter (FC B306) are required if network switches do not support 40GbE
  - ThinkSystem Intel E810-DA2 10/25GbE SFP28 2-Port PCIe Ethernet Adapter (FC BCD6)
- Storage: The following storage devices have been certified:
  - Dual 480GB M.2 adapter configured for RAID-1 for OS boot (FC B919)
  - 430-16i SAS HBA (RAID not supported, FC AUNM)
  - 4 x 800GB LFF High Performance SAS SSD for cache (FC B170)
  - 10 x 4TB LFF 6Gbps NL SATA HDD for capacity (FC AUU8)

This is a general purpose configuration that uses SSD and HDD storage devices. It is recommended when raw capacity requirements are less than 40TB per node. Network bandwidth of 10GbE is generally adequate for this configuration. This is one of the configurations that has been certified for use in a 2-node Microsoft HCI solution.

SH60T1 hybrid configuration

This configuration uses the Lenovo ThinkAgile MX Certified Node configured with SSDs for the cache tier and HDDs for the capacity tier. Total raw capacity of this configuration is 60TB per node.

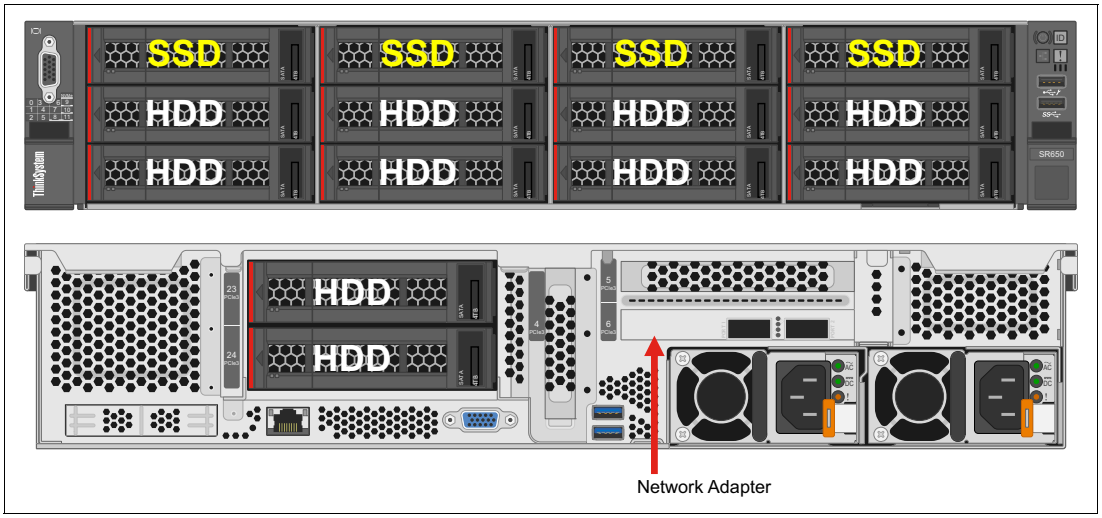

Figure 2 Lenovo ThinkAgile MX Certified Node configuration SH60T1

Additional details include the following:

- CPU: 2 x Intel Gold or Platinum family processors
- Memory: 192GB - 1.5TB
- Network adapter: The following network adapters have been certified:
  - Mellanox ConnectX-4 2-port 10/25GbE Ethernet Adapter (FC AUAJ)
  - 2 x Mellanox ConnectX-4 1-port 40GbE Ethernet Adapter (FC ATRN)
    - 2 x Mellanox QSA 100G to 25G Cable Adapter (FC B306) are required if network switches do not support 40GbE
  - ThinkSystem Intel E810-DA2 10/25GbE SFP28 2-Port PCIe Ethernet Adapter (FC BCD6)
- Storage: The following storage devices have been certified:
  - Dual 480GB M.2 adapter configured for RAID-1 for OS boot (FC B919)
  - 430-16i SAS HBA (RAID not supported, FC AUNM)
  - 4 x 1.6TB LFF High Performance SAS SSD for cache (FC B171)
  - 10 x 6TB LFF 6Gbps NL SATA HDD for capacity (FC AUUA)

This is a general purpose configuration that uses SSD and HDD storage devices, with increased raw capacity of 60TB per node. It is recommended when a bit more storage capacity is required. A 16-node Microsoft HCI solution built using this configuration will provide a total raw storage capacity of nearly a petabyte.

NH80T1 hybrid configuration

This configuration uses the Lenovo ThinkAgile MX Certified Node configured with hot-swap NVMe U.2 devices for the cache tier and HDDs for the capacity tier. Total raw capacity of this configuration is 80TB per node.

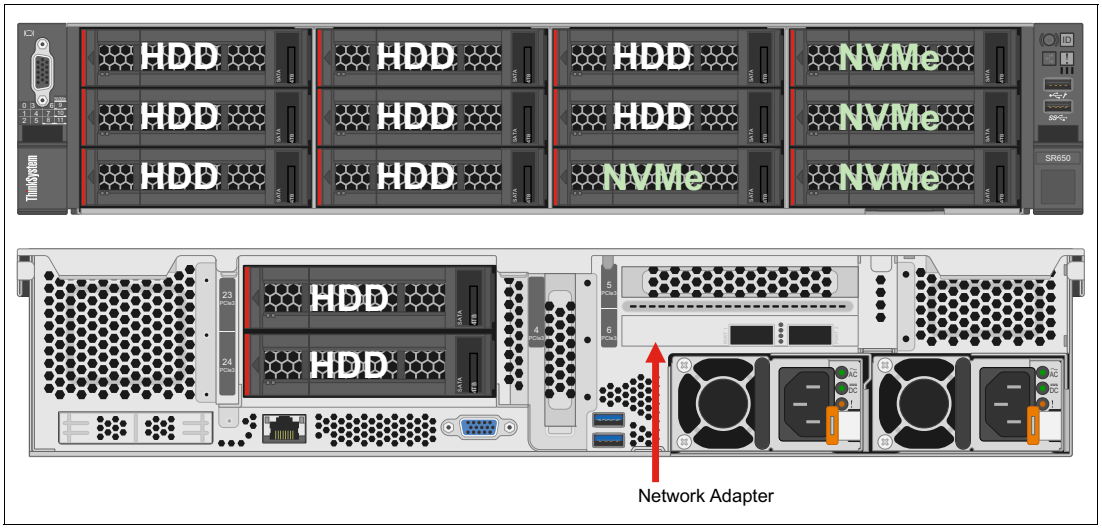

Figure 3 Lenovo ThinkAgile MX Certified Node configuration NH80T1

Additional details include the following:

- CPU: 2 x Intel Gold or Platinum family processors
- Memory: 192GB - 1.5TB
- Network adapter: The following network adapters have been certified:
  - Mellanox ConnectX-4 2-port 10/25GbE Ethernet Adapter (FC AUAJ)
  - 2 x Mellanox ConnectX-4 1-port 40GbE Ethernet Adapter (FC ATRN)
    - 2 x Mellanox QSA 100G to 25G Cable Adapter (FC B306) are required if network switches do not support 40GbE
  - ThinkSystem Intel E810-DA2 10/25GbE SFP28 2-Port PCIe Ethernet Adapter (FC BCD6)
- Storage: The following storage devices have been certified:
  - Dual 480GB M.2 adapter configured for RAID-1 for OS boot (FC B919)
  - 430-16i SAS HBA (RAID not supported, FC AUNM)
  - 3 or 4 x 3.2TB LFF HS NVMe U.2 for cache (FC B2XG)
  - NVMe U.2 devices require AnyBay™ drive bays
  - 10 x 8TB LFF 6Gbps NL SATA HDD for capacity (FC AUU9)

Note: It is recommended to use a minimum of 25GbE network bandwidth for better HDD rebuild times for HDDs with a capacity of 8TB or more.

This is a high performance configuration that uses hot-swap NVMe U.2 devices inserted into the AnyBay drive bays as cache for the HDD capacity tier and has the same raw capacity of 80TB per node as Configuration NH80T1a. It is highly recommended to use a minimum network bandwidth of 25GbE in order to keep up with NVMe storage performance and also to potentially reduce HDD rebuild times. This is one of the configurations that has been certified for use in a 2-node Microsoft HCI solution.

NH120T1 hybrid configuration

This configuration uses the Lenovo ThinkAgile MX Certified Node configured with hot-swap NVMe U.2 devices for the cache tier and HDDs for the capacity tier. Total raw capacity of this configuration is 120TB per node.

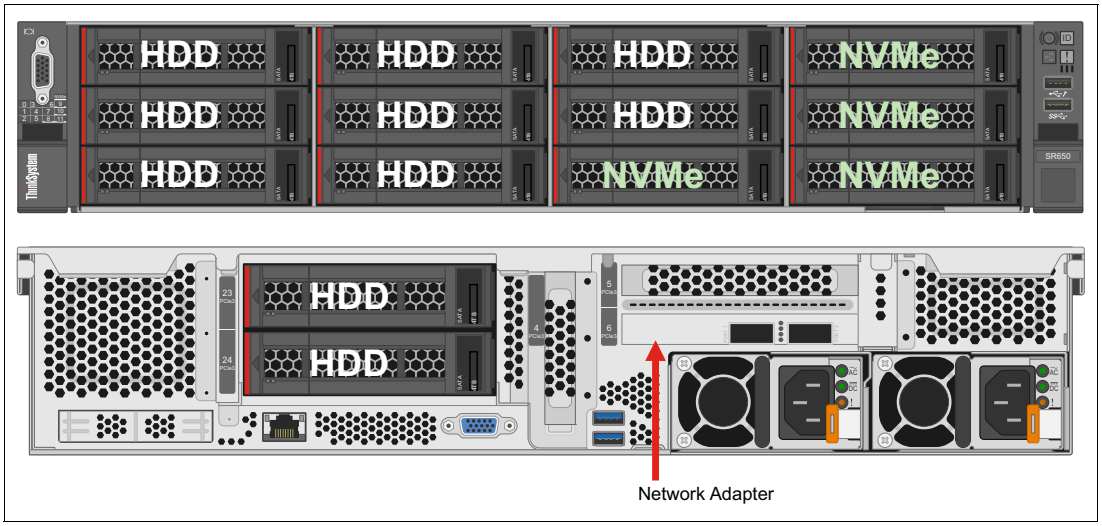

Figure 4 Lenovo ThinkAgile MX Certified Node configuration NH120T1

Additional details include the following:

- CPU: 2 x Intel Gold or Platinum family processors
- Memory: 192GB - 1.5TB
- Network adapter: The following network adapters have been certified:
  - Mellanox ConnectX-4 2-port 10/25GbE Ethernet Adapter (FC AUAJ)
  - 2 x Mellanox ConnectX-4 1-port 40GbE Ethernet Adapter (FC ATRN)
    - 2 x Mellanox QSA 100G to 25G Cable Adapter (FC B306) are required if network switches do not support 40GbE
  - ThinkSystem Intel E810-DA2 10/25GbE SFP28 2-Port PCIe Ethernet Adapter (FC BCD6)
- Storage: The following storage devices have been certified:
  - Dual 480GB M.2 adapter configured for RAID-1 for OS boot (FC B919)
  - 430-16i SAS HBA (RAID not supported, FC AUNM)
  - 4 x 3.2TB LFF HS NVMe U.2 for cache (FC B2XG)
  - NVMe U.2 devices require AnyBay drive bays
  - 10 x 12TB LFF 6Gbps NL SATA HDD for capacity (FC B118)

Note: It is recommended to use a minimum of 25GbE network bandwidth for better HDD rebuild times for HDDs with a capacity of 8TB or more.

This is a high performance configuration that uses hot-swap NVMe U.2 devices inserted into the AnyBay drive bays as cache for the HDD capacity tier and has a total raw capacity of 120TB per node. It is highly recommended to use a minimum network bandwidth of 25GbE in order to keep up with NVMe storage performance and also to potentially reduce HDD rebuild times. A 16-node Microsoft HCI solution built using this configuration will provide a total raw storage capacity of nearly 2 Petabytes.

NS61T1 all-flash configuration

This configuration uses the Lenovo ThinkAgile MX Certified Node with 24 2.5” drive bays configured with U.2 NVMe devices for the cache tier and SSDs for the capacity tier. Total raw capacity of this configuration is approximately 61TB per node. The focus of this configuration is performance rather than large capacity.

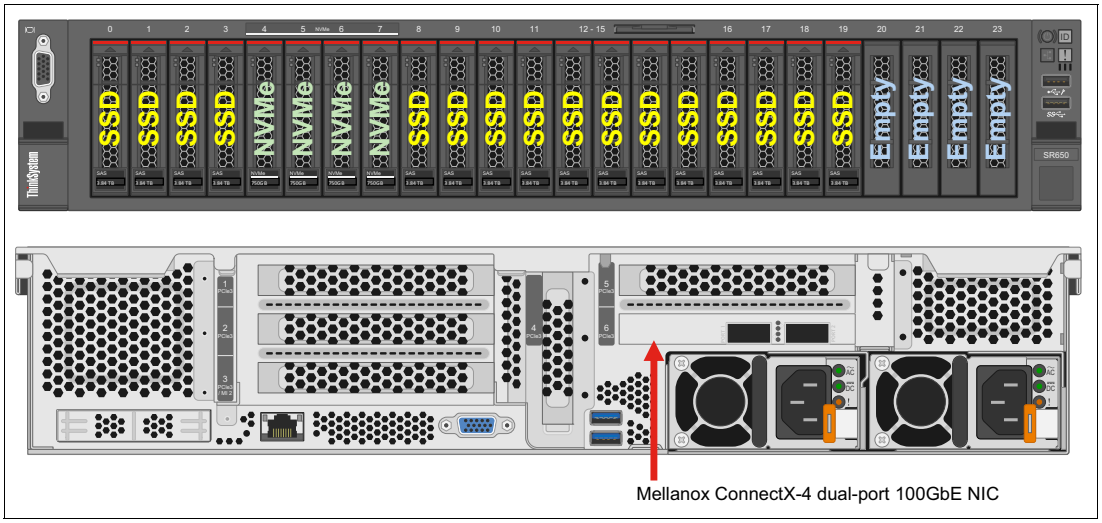

Figure 5 Lenovo ThinkAgile MX Certified Node configuration NS61T1

Additional details include the following:

- CPU: 2 x Intel Gold or Platinum family processors
- Memory: 192GB - 1.5TB
- Network adapter: The following network adapters have been certified:
  - Mellanox ConnectX-4 2-port 100GbE Ethernet Adapter (FC ATRP)
  - 2 x Mellanox ConnectX-6 HDR100 QSFP56 1-port PCIe InfiniBand Adapter (FC B4R9)
  - Mellanox ConnectX-6 HDR100 QSFP56 2-port PCIe InfiniBand Adapter (FC B4RA)
- Storage: The following storage devices have been certified:
  - Dual 480GB M.2 adapter configured for RAID-1 for OS boot (FC B919)
  - 3 x 430-8i SAS HBA (RAID not supported, FC AUNL)
  - 4 x 750GB High Performance Optane U.2 NVMe for cache (FC B2ZJ)
  - 16 x 3.84TB 6Gbps SATA SSD for capacity (FC B49C)

This is an ultra-high performance all-flash configuration that uses NVMe devices as cache for the SSD capacity tier, but has a smaller raw capacity of approximately 61TB per node. In order to achieve maximum performance, this configuration includes a 2-port 100GbE Mellanox network adapter in each node. The Mellanox ConnectX-6 adapters shown above support Ethernet, including RoCEv2, and have been certified for Azure Local.

NS77T1 all-flash configuration

This configuration uses the Lenovo ThinkAgile MX Certified Node with 24 2.5” drive bays configured with U.2 NVMe devices for the cache tier and SSDs for the capacity tier. Total raw capacity of this configuration is approximately 77TB per node. The focus of this configuration is performance rather than large capacity.

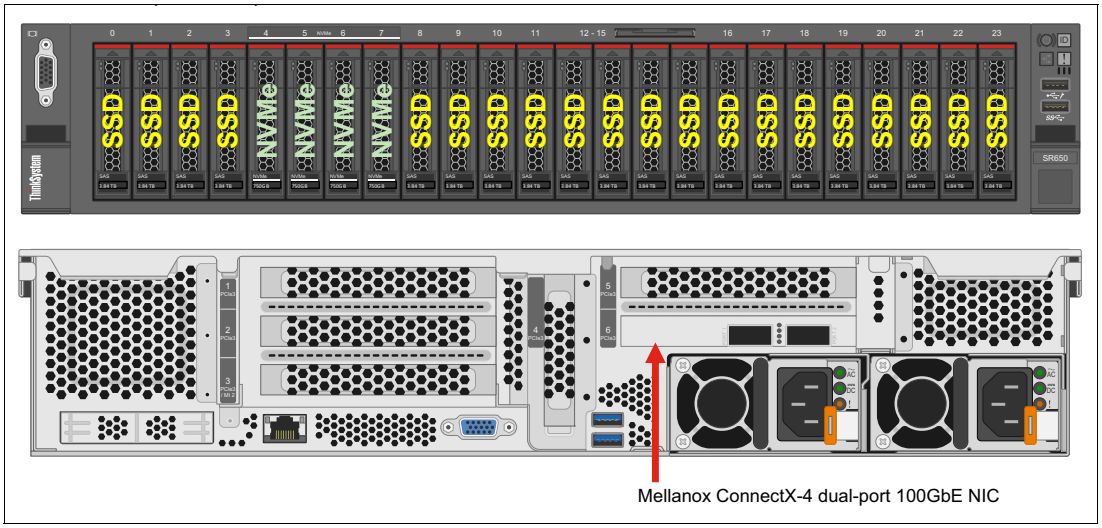

Figure 6 Lenovo ThinkAgile MX Certified Node configuration NS77T1

Additional details include the following:

- CPU: 2 x Intel Gold or Platinum family processors
- Memory: 192GB - 1.5TB
- Network adapter: The following network adapters have been certified:
  - Mellanox ConnectX-4 2-port 100GbE Ethernet Adapter (FC ATRP)
  - 2 x Mellanox ConnectX-6 HDR100 QSFP56 1-port PCIe InfiniBand Adapter (FC B4R9)
  - Mellanox ConnectX-6 HDR100 QSFP56 2-port PCIe InfiniBand Adapter (FC B4RA)
- Storage: The following storage devices have been certified:
  - Dual 480GB M.2 adapter configured for RAID-1 for OS boot (FC B919)
  - 3 x 430-8i SAS HBA (RAID not supported, FC AUNL)
  - 4 x 3.2TB High Performance U.2 NVMe for cache (FC B11K)
  - 20 x 3.84TB 6Gbps SATA SSD for capacity (FC B0Z2)

This is an ultra-high performance all-flash configuration that uses NVMe devices as cache for the SSD capacity tier, but has a smaller raw capacity of approximately 77TB per node. In order to achieve maximum performance, this configuration includes a 2-port 100GbE Mellanox network adapter in each node. The Mellanox ConnectX-6 adapters shown above support Ethernet, including RoCEv2, and have been certified for Azure Local.

NN38T1 all-flash configuration (all-NVMe)

This configuration uses the Lenovo ThinkAgile MX Certified Node with 24 2.5” drive bays configured with U.2 NVMe devices as the only storage devices. Total raw capacity of this configuration is approximately 38TB per node. The focus of this configuration is performance rather than large capacity.

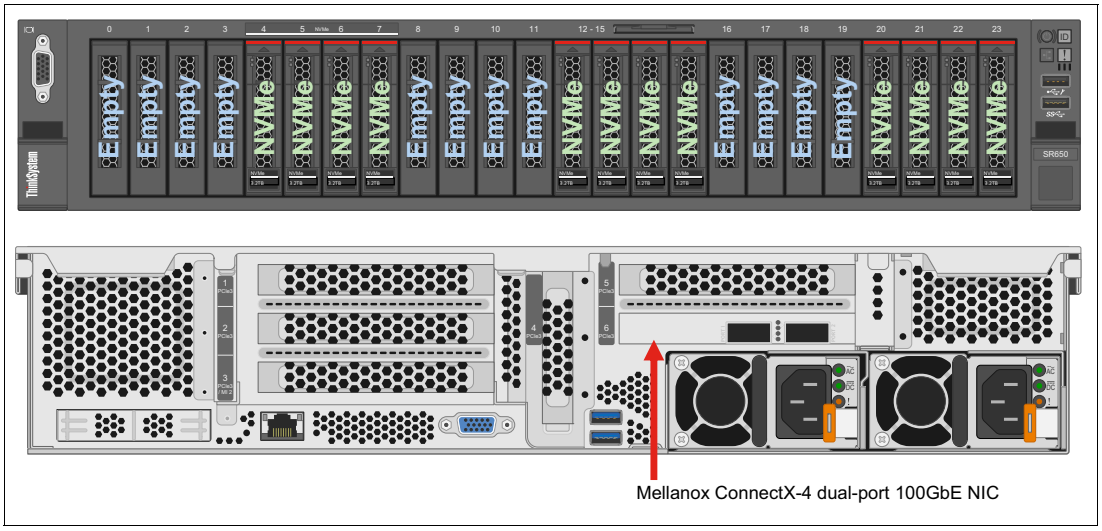

Figure 7 Lenovo ThinkAgile MX Certified Node configuration NN38T1

Additional details include the following:

- CPU: 2 x Intel Gold or Platinum family processors
- Memory: 192GB - 1.5TB
- Network adapter: The following network adapters have been certified:
  - Mellanox ConnectX-4 2-port 100GbE Ethernet Adapter (FC ATRP)
  - 2 x Mellanox ConnectX-6 HDR100 QSFP56 1-port PCIe InfiniBand Adapter (FC B4R9)
  - Mellanox ConnectX-6 HDR100 QSFP56 2-port PCIe InfiniBand Adapter (FC B4RA)
- Storage: The following storage devices have been certified:
  - Dual 480GB M.2 adapter configured for RAID-1 for OS boot (FC B919)
  - 3 x 430-8i SAS HBA (RAID not supported, FC AUNL)
  - 2 x ThinkSystem 1610-4p NVMe Switch adapter (FC AUV2)
  - 12 x 3.2TB High Performance U.2 NVMe (FC B11K)

This is an ultra-high performance all-NVMe configuration that uses only NVMe devices for storage, but has a smaller raw capacity of approximately 38TB per node. In order to achieve maximum performance, this configuration includes a 2-port 100GbE Mellanox network adapter in each node. The Mellanox ConnectX-6 adapters shown above support Ethernet, including RoCEv2, and have been certified for Azure Local.

SS92T1 all-flash configuration (all-SSD)

This configuration uses the Lenovo ThinkAgile MX Certified Node with 24 2.5” drive bays configured with SSDs as the only storage devices. Total raw capacity of this configuration is approximately 92TB per node. The focus of this configuration is performance rather than large capacity.

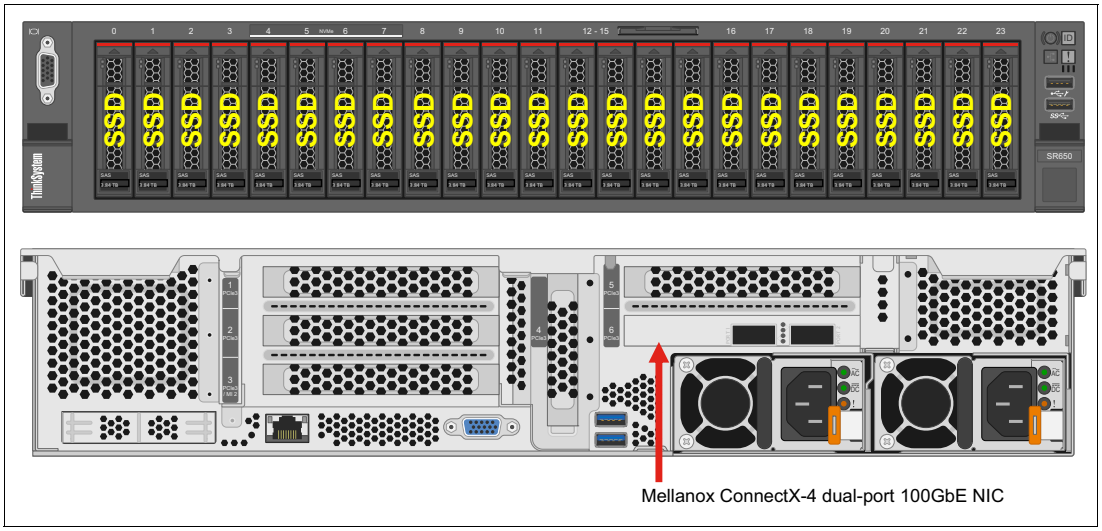

Figure 8 Lenovo ThinkAgile MX Certified Node configuration SS92T1

Additional details include the following:

- CPU: 2 x Intel Gold or Platinum family processors
- Memory: 192GB - 1.5TB
- Network adapter: The following network adapters have been certified:
  - Mellanox ConnectX-4 2-port 100GbE Ethernet Adapter (FC ATRP)
  - 2 x Mellanox ConnectX-6 HDR100 QSFP56 1-port PCIe InfiniBand Adapter (FC B4R9)
  - Mellanox ConnectX-6 HDR100 QSFP56 2-port PCIe InfiniBand Adapter (FC B4RA)
- Storage: The following storage devices have been certified:
  - Dual 480GB M.2 adapter configured for RAID-1 for OS boot (FC B919)
  - 3 x 430-8i SAS HBA (RAID not supported, FC AUNL)
  - 24 x 3.84TB 6Gbps SATA SSD for capacity (FC B0Z2)

This is an ultra-high performance all-SSD configuration that uses only SSD devices for storage, but has a smaller raw capacity of approximately 92TB per node. In order to achieve maximum performance, this configuration includes a 2-port 100GbE Mellanox network adapter in each node. The Mellanox ConnectX-6 adapters shown above support Ethernet, including RoCEv2, and have been certified for Azure Local.

Small cluster configurations

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

Direct-connect networking

For a 2-node 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-4 10/25GbE 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 9 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.

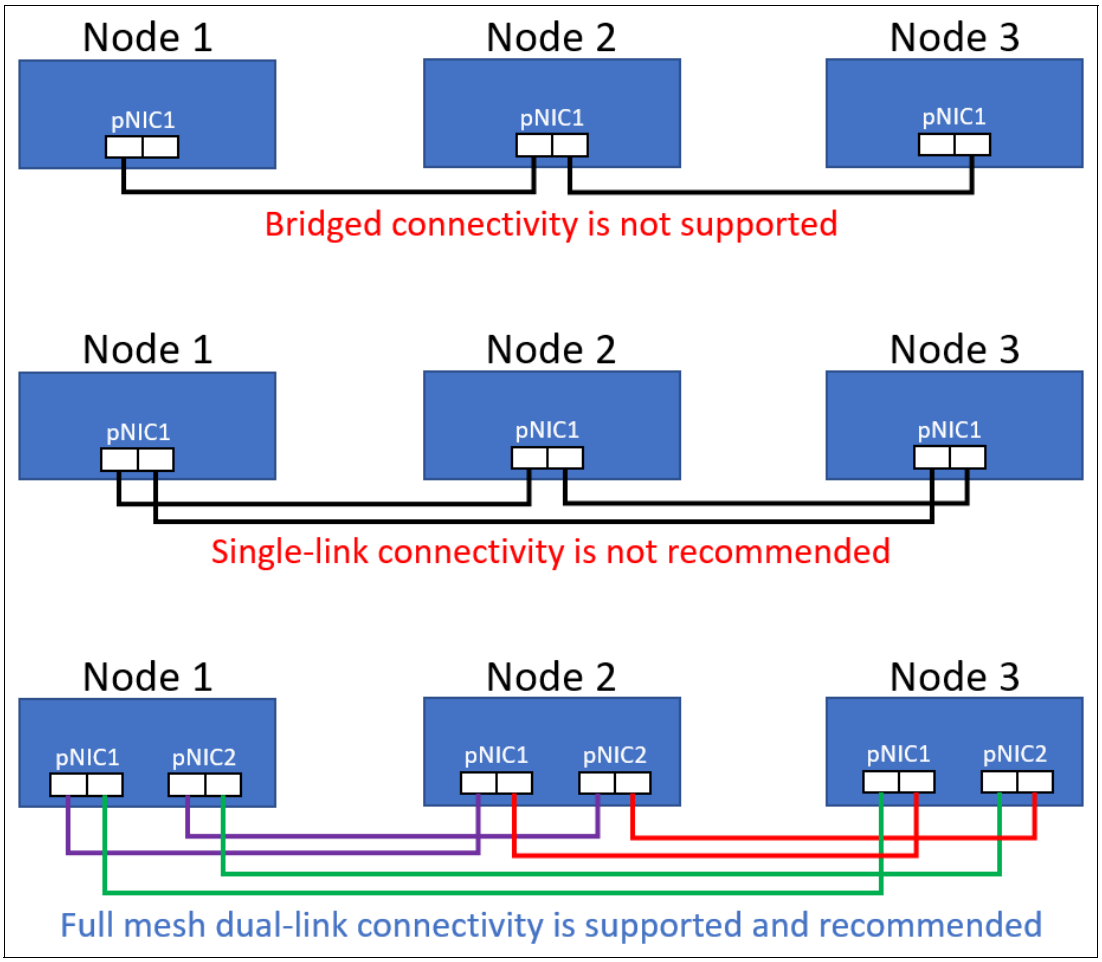

Figure 9 Various node-to-node network connectivity models

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, a separate network connection is still required from the customer network to the cluster (aka "north-south" traffic). A low-cost option for this additional network interface is to use one of the LAN On Motherboard (LOM) cards available for the ThinkAgile MX Certified Node.

Important: LOM ports are not certified to carry RDMA storage traffic inside the solution. These cards are offered only to connect the cluster to the customer network.

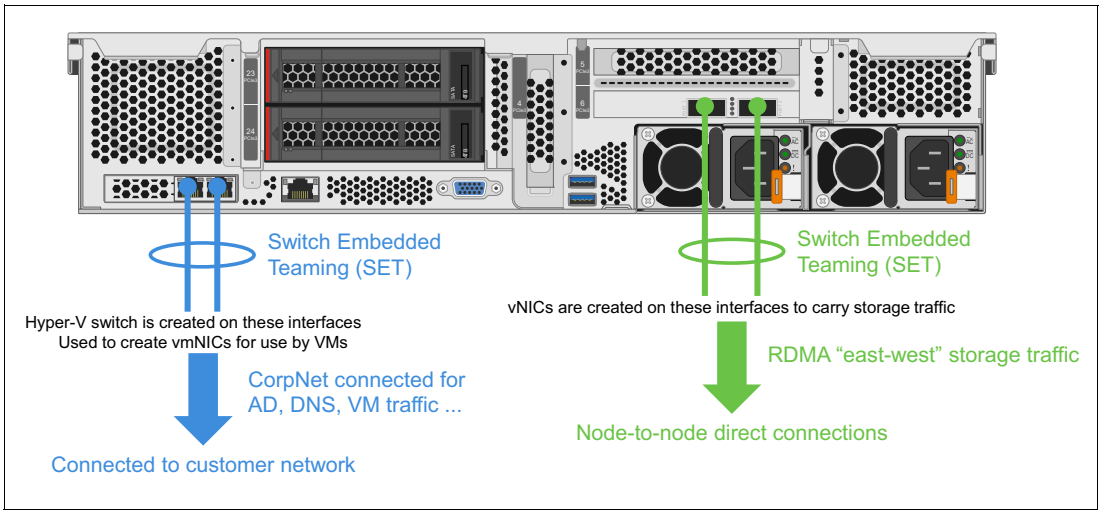

Figure 10 Diagram showing network connectivity for a ThinkAgile MX Certified Node that is part of a 2-node direct-connected HCI cluster.

Component selection

The Lenovo certified configurations listed above 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 ThinkAgile MX Certified Nodes are based on the ThinkSystem SR650 server, so the SR650 server itself has been certified.
- The Lenovo ThinkSystem SR630 has also been certified for use as an Azure Local or Windows Server S2D node even though there is no ThinkAgile MX solution based on this server.
- In general, the number of nodes can range from 1 to 16.

Processors

- Two Intel processors with a recommended minimum of 8 cores per CPU in the Silver (4100 series), Gold (5100 or 6100 series), or Platinum (8100 series) processor families
- 205 watt processors are not supported

Memory

- Minimum of 192GB is required for converged and 384GB for hyperconverged
- Maximum of 1.5TB per node

- A “balanced memory configuration” is strongly recommended - for details, see the following URL: http://lenovopress.com/lp0742.pdf

OS Boot

- Minimum requirement is 200GB OS boot volume
- M.2 Mirroring Kit with dual 480GB M.2 SSD configured as RAID-1 for resilience are specified in all Lenovo ThinkAgile MX Certified Node configurations
- If hot swap storage is preferred for boot, we recommend a RAID card configured for RAID-1 with two SSDs or two HDDs

GPUs

The following GPUs are supported for ThinkAgile MX solutions running on Lenovo SR650 rack servers:

- ThinkSystem NVIDIA A2 16GB PCIe Gen4 Passive GPU (FC BQTZ)
- ThinkSystem NVIDIA A10 24GB PCIe Gen4 Passive GPU (FC BFTZ)
- ThinkSystem NVIDIA A16 64GB PCIe Gen4 Passive GPU (FC BNFE)
- ThinkSystem NVIDIA A30 24GB PCIe Gen4 Passive GPU (FC BJHG)
- ThinkSystem NVIDIA A100 40GB PCIe Gen4 Passive GPU (FC BEL5)
- ThinkSystem NVIDIA T4 16GB PCIe Passive GPU (FC B4YB)

Storage Network Adapters

We recommend 25GbE or 100GbE, depending on storage configuration, for optimal performance. Storage network adapters are shown below by RDMA protocol and apply to solutions based on the SR650 rack server, as noted.

- For RoCE v2 on SR650 solutions:
  - Mellanox ConnectX-4 Lx 2-port 10/25GbE Ethernet Adapter (FC AUAJ)
  - Mellanox ConnectX-6 HDR100 QSFP56 2-port PCIe InfiniBand Adapter (FC B4RA)
  - Mellanox ConnectX-6 HDR100 QSFP56 1-port PCIe InfiniBand Adapter (FC B4R9)
  - Mellanox ConnectX-6 Lx 10/25GbE SFP28 2-Port PCIe Ethernet Adapter (FC BE4U)
  - Mellanox ConnectX-6 Dx 100GbE QSFP56 2-port PCIe Ethernet Adapter (FC B8PP)
  - Mellanox ConnectX-4 Lx 1-port 40GbE Ethernet Adapter. Previously used with Mellanox QSA 100G to 25G Cable Adapter if two single-port NICs is preferred.

Note: Network switches must support the RoCE v2 feature set for best storage performance. See “Network switches” on page 21 for more information regarding Lenovo and NVIDIA/Mellanox network switches that have been tested with ThinkAgile MX solutions.

- For iWARP on SR650 solutions:
  - ThinkSystem Intel E810-DA2 10/25GbE SFP28 2-Port PCIe Ethernet Adapter (FC BCD6)

Non-storage Network Adapters

For ThinkAgile MX solutions that use the SR650 rack servers, LAN On Motherboard (LOM) ports are supported for north-south (Management and Compute) traffic, but not for east-west (Storage) traffic, as discussed in “Direct-connect networking” on page 14.

- Management and Compute traffic:
  - ThinkSystem 10Gb 2-port Base-T LOM (FC AUKL)

- ThinkSystem 10Gb 2-port SFP+ LOM (FC AUKJ)
- ThinkSystem 10Gb 4-port SFP+ LOM (FC AUKK)

Storage HBAs

- Hybrid storage configurations
  - ThinkSystem 430-16i SAS/SATA 12Gb HBA
  - ThinkSystem 440-16i SAS/SATA 12Gb HBA
  - ThinkSystem 4350-16i SAS/SATA 12Gb HBA
- All-flash storage configuarions
  - ThinkSystem 430-8i SAS/SATA 12Gb HBA
  - ThinkSystem 440-8i SAS/SATA 12Gb HBA
  - ThinkSystem 4350-8i SAS/SATA 12Gb HBA

NVMe switch adapters

- ThinkSystem 1610-4p NVMe Switch adapter
  - NVMe switches are used for configurations that include more than 4 NVMe devices

Storage devices

- For configurations with two storage device types, the number of devices can be reduced to a minimum of two cache and four capacity devices
- For configurations with a single storage device type (all-SSD or all-NVMe), the number of devices can be reduced to a total of 4 SSD or 4 NVMe devices
- NVMe U.2 devices require AnyBay option

A minimum 10% cache to capacity ratio (e.g. 2x 800GB SSD and 4x 4TB HDD) is strongly recommend. 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 may be significantly impacted.

Table 2 provides a list of all certified Lenovo storage devices that can be used to configure a Hybrid storage HCI solution based on Lenovo ThinkSystem SR650 rack server. This table does not include OS boot devices.

Table 2 Lenovo storage devices certified for ThinkAgile MX hybrid storage solutions on SR650

Storage Devices Used for Hybrid Solutions based on SR650	Feature Code	Type	Usage
ThinkSystem 3.5" Intel P4610 1.6TB Mainstream NVMe PCIe 3.0 x4 HS SSD	B58C	NVMe	Cache
ThinkSystem 3.5" Intel P5600 1.6TB Mainstream NVMe PCIe 4.0 x4 HS SSD	BCFM	SSD	Cache
ThinkSystem 3.5" Intel P5600 3.2TB Mainstream NVMe PCIe 4.0 x4 HS SSD	BCFJ	SSD	Cache
ThinkSystem 3.5" Intel P5600 6.4TB Mainstream NVMe PCIe 4.0 x4 HS SSD	BCFQ	SSD	Cache
ThinkSystem 3.5" Intel P5620 1.6TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEK	SSD	Cache
ThinkSystem 3.5" Intel P5620 3.2TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEM	SSD	Cache
ThinkSystem 3.5" Intel P5620 6.4TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEN	SSD	Cache
ThinkSystem 3.5" Intel P5620 12.8TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	BNEP	SSD	Cache
ThinkSystem 3.5” SS530 800GB Performance SAS 12Gb HS SSD	B4Y8	SSD	Cache
ThinkSystem 3.5” SS530 1.6TB Performance SAS 12Gb HS SSD	B4Y9	SSD	Cache
ThinkSystem 3.5” SS530 3.2TB Performance SAS 12Gb HS SSD	B4YA	SSD	Cache

Storage Devices Used for Hybrid Solutions based on SR650	Feature Code	Type	Usage
ThinkSystem 3.5" PM1645a 800GB Mainstream SAS 12Gb HS SSD	B8HT	SSD	Cache
ThinkSystem 3.5" PM1645a 1.6TB Mainstream SAS 12Gb HS SSD	B8JN	SSD	Cache
ThinkSystem 3.5" PM1645a 3.2TB Mainstream SAS 12Gb HS SSD	B8JK	SSD	Cache
ThinkSystem 3.5" PM1655 800GB Mixed Use SAS 24Gb HS SSD	BNW7	SSD	Cache
ThinkSystem 3.5" PM1655 1.6TB Mixed Use SAS 24Gb HS SSD	BNWA	SSD	Cache
ThinkSystem 3.5" PM1655 3.2TB Mixed Use SAS 24Gb HS SSD	BNWB	SSD	Cache
ThinkSystem 3.5" PM1655 6.4TB Mixed Use SAS 24Gb HS SSD	BP3G	SSD	Cache
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" 4TB 7.2K NL SAS 12Gb HS 512n HDD	AUU6	HDD	Capacity
ThinkSystem 3.5" 6TB 7.2K NL SAS 12Gb HS 512e HDD	AUU7	HDD	Capacity
ThinkSystem 3.5" 8TB 7.2K NL SAS 12Gb HS 512e HDD	B0YR	HDD	Capacity
ThinkSystem 3.5" 10TB 7.2K NL SAS 12Gb HS 512e HDD	AUUG	HDD	Capacity
ThinkSystem 3.5" 12TB 7.2K NL SAS 12Gb HS 512e HDD	B117	HDD	Capacity
ThinkSystem 3.5" 14TB 7.2K NL SAS 12Gb HS 512e HDD	B496	HDD	Capacity

Table 3 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 SR650 rack server. This table does not include OS boot devices.

Table 3 Lenovo storage devices certified for ThinkAgile MX all-flash storage solutions on SR630 and SR650

Storage Devices Used for All-Flash Solutions based on SR650	Feature Code	Type	Usage
ThinkSystem U.2 Intel P4800X 750GB Performance NVMe PCIe 3.0 x4 HS SSD	B2ZJ	NVMe	Cache
ThinkSystem 2.5” U.2 Intel P5600 1.6TB Mainstream NVMe PCIe 3.0 x4 HS SSD	B589	NVMe	Cache
ThinkSystem 2.5” U.2 Intel P5600 3.2TB Mainstream NVMe PCIe 3.0 x4 HS SSD	B58A	NVMe	Cache
ThinkSystem 2.5” U.2 Intel P5600 6.4TB Mainstream NVMe PCIe 3.0 x4 HS SSD	B58B	NVMe	Cache
ThinkSystem 2.5” U.2 Intel P5620 1.6TB Mixed Use NVMe PCIe 3.0 x4 HS SSD	BNEG	NVMe	Cache
ThinkSystem 2.5” U.2 Intel P5620 3.2TB Mixed Use NVMe PCIe 3.0 x4 HS SSD	BNEH	NVMe	Cache
ThinkSystem 2.5” U.2 Intel P5620 6.4TB Mixed Use NVMe PCIe 3.0 x4 HS SSD	BNEZ	NVMe	Cache
ThinkSystem 2.5” U.2 Intel P5620 12.8TB Mixed Use NVMe PCIe 3.0 x4 HS SSD	BA4V	NVMe	Cache
ThinkSystem 2.5” SS530 400GB Performance SAS 12Gb HS SSD	B4Y4	SSD	Cache

Storage Devices Used for All-Flash Solutions based on SR650	Feature Code	Type	Usage
ThinkSystem 2.5” SS530 800GB Performance SAS 12Gb HS SSD	B4Y5	SSD	Cache
ThinkSystem 2.5” SS530 1.6TB Performance SAS 12Gb HS SSD	B4Y6	SSD	Cache
ThinkSystem 2.5” SS530 3.2TB Performance SAS 12Gb HS SSD	B4Y7	SSD	Cache
ThinkSystem 2.5" PM1645a 800GB Mainstream SAS 12Gb HS SSD	B8HU	SSD	Cache
ThinkSystem 2.5" PM1645a 1.6TB Mainstream SAS 12Gb HS SSD	B8J4	SSD	Cache
ThinkSystem 2.5" PM1645a 3.2TB Mainstream SAS 12Gb HS SSD	B8JD	SSD	Cache
ThinkSystem 2.5" PM1655 800GB Mixed Use SAS 24Gb HS SSD	BNW8	SSD	Cache
ThinkSystem 2.5" PM1655 1.6TB Mixed Use SAS 24Gb HS SSD	BNW9	SSD	Cache
ThinkSystem 2.5" PM1655 3.2TB Mixed Use SAS 24Gb HS SSD	BNW6	SSD	Cache
ThinkSystem 2.5" PM1655 6.4TB Mixed Use SAS 24Gb HS SSD	BP3K	SSD	Cache
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" Intel S4510 1.92TB Entry SATA 6Gb HS SSD	B49B	SSD	Capacity
ThinkSystem 2.5" Intel S4510 3.84TB Entry SATA 6Gb HS SSD	B49C	SSD	Capacity
ThinkSystem 2.5" Intel S4520 480GB Read Intensive SATA 6Gb HS SSD	BA7G	SSD	Capacity
ThinkSystem 2.5" Intel S4520 960GB Read Intensive SATA 6Gb HS SSD	BA7H	SSD	Capacity
ThinkSystem 2.5" Intel S4520 1.92TB Read Intensive SATA 6Gb HS SSD	BA7J	SSD	Capacity
ThinkSystem 2.5" Intel S4520 3.84TB Read Intensive SATA 6Gb HS SSD	BK77	SSD	Capacity
ThinkSystem 2.5" Intel S4520 7.68TB Read Intensive SATA 6Gb HS SSD	BK78	SSD	Capacity
ThinkSystem 2.5" 5300 1.92TB Entry SATA 6Gb HS SSD	B8J5	SSD	Capacity
ThinkSystem 2.5" 5300 3.84TB Entry SATA 6Gb HS SSD	B8JP	SSD	Capacity
ThinkSystem 2.5" 5300 7.68TB Entry SATA 6Gb HS SSD	B8J2	SSD	Capacity
ThinkSystem 2.5" 5300 480GB Mainstream SATA 6Gb HS SSD	B8HY	SSD	Capacity
ThinkSystem 2.5" 5300 960GB Mainstream SATA 6Gb HS SSD	B8J6	SSD	Capacity
ThinkSystem 2.5" 5300 1.92TB Mainstream SATA 6Gb HS SSD	B8JE	SSD	Capacity
ThinkSystem 2.5" 5300 3.84TB Mainstream SATA 6Gb HS SSD	B8J7	SSD	Capacity
ThinkSystem 2.5" 5400 PRO 480GB Read Intensive SATA 6Gb HS SSD	BQ1P	SSD	Capacity

Storage Devices Used for All-Flash Solutions based on SR650	Feature Code	Type	Usage
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" PM1645a 800GB Mainstream SAS 12Gb HS SSD	B8HU	SSD	Capacity
ThinkSystem 2.5" PM1645a 1.6TB Mainstream SAS 12Gb HS SSD	B8J4	SSD	Capacity
ThinkSystem 2.5" PM1645a 3.2TB Mainstream SAS 12Gb HS SSD	B8JD	SSD	Capacity
ThinkSystem 2.5" PM1643a 1.92TB Entry SAS 12Gb HS SSD	B91B	SSD	Capacity
ThinkSystem 2.5" PM1643a 3.84TB Entry SAS 12Gb HS SSD	B91C	SSD	Capacity
ThinkSystem 2.5" PM1643a 7.68TB Entry SAS 12Gb HS SSD	B91D	SSD	Capacity
ThinkSystem 2.5" PM1653 960GB Read Intensive SAS 24Gb HS SSD	BNWC	SSD	Capacity
ThinkSystem 2.5" PM1653 1.92TB Read Intensive SAS 24Gb HS SSD	BNWE	SSD	Capacity
ThinkSystem 2.5" PM1653 3.84TB Read Intensive SAS 24Gb HS SSD	BNWF	SSD	Capacity
ThinkSystem 2.5" PM1653 7.68TB Read Intensive SAS 24Gb HS SSD	BP3E	SSD	Capacity
ThinkSystem 2.5" PM1653 15.36TB Read Intensive SAS 24Gb HS SSD	BP3J	SSD	Capacity
ThinkSystem 2.5" PM1653 30.72TB Read Intensive SAS 24Gb HS SSD	BP3D	SSD	Capacity

Note: Do not use a storage device for a purpose other than listed in the Usage column. For example, the Intel S4500 and S4510 SSDs have been certified for use only as a capacity device, so should not be used as a cache device.

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, previous generations reach their end of life. Table 4 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 4 Storage device end of life summary

End of Life Device	Date	Replacement Device
ThinkSystem U.2 PX04PMB 960GB Mainstream 3.5" NVMe PCIe 3.0 x4 HS SSD	June 2018	None
PX04PMC 1.6TB Performance NVMe PCIe 3.0 x4 Flash Adapter (AIC)	June 2018	None
PX04PMC 3.2TB Performance NVMe PCIe 3.0 x4 Flash Adapter (AIC)	June 2018	None

End of Life Device	Date	Replacement Device
PX04PMC 1.92TB Mainstream NVMe PCIe 3.0 x4 Flash Adapter (AIC)	June 2018	None
Intel S4500 SSD devices	September 2018	Intel S4510 SSD devices
Intel P4600 NVMe devices	March 2019	Intel P4610 NVMe devices
ThinkSystem 3.5" HUSMM32 SSD devices	June 2019	ThinkSystem 3.5” SS530 SSD devices
ThinkSystem 2.5" 5200 SSD devices	August 2020	ThinkSystem 2.5" 5300 SSD devices
ThinkSystem 3.5” SS530 SSD devices	November 2020	ThinkSystem 3.5" PM1645a SSD devices
ThinkSystem 2.5” SS530 SSD devices	November 2020	ThinkSystem 2.5" PM1645a SSD devices
ThinkSystem M.2 5100 SSD devices	February 2021	ThinkSystem M.2 5300 SSD devices
ThinkSystem U.2 Intel P4610 SSD devices	July 2021	ThinkSystem U.2 Intel P5600 SSD devices
Intel S4510 SSD devices	June 2022	Intel S4520 SSD devices
M.2 Intel P4511 NVMe devices	June 2022	M.2 7450M NVMe devices
ThinkSystem 2.5" U.2 P5500 NVMe devices	March 2023	ThinkSystem 2.5" U.2 P5520 NVMe devices
ThinkSystem 2.5" U.2 P5600 NVMe devices	March 2023	ThinkSystem 2.5" U.2 P5620 NVMe devices
ThinkSystem 2.5" PM1643a SSD devices	March 2023	ThinkSystem 2.5" PM1653 SSD devices
ThinkSystem 2.5" PM1645a SSD devices	March 2023	ThinkSystem 2.5" PM1655 SSD devices
ThinkSystem 2.5" 5300 Entry SSD devices	March 2023	ThinkSystem 2.5" 5400 PRO Read Intensive SSD devices
ThinkSystem 2.5" 5300 Mainstream SSD devices	March 2023	ThinkSystem 2.5" 5400 PRO Read Intensive SSD devices

Network switches

Network switches that have been tested in our labs include Lenovo and NVIDIA (Mellanox) switches. Although Lenovo no longer sells network switches, information about them is included here for customers who already own them. Mellanox switches must be ordered directly from NVIDIA.

Lenovo network switches

Lenovo network switches are no longer being offered. The information contained in this section is provided in case customers want to verify that their existing Lenovo switches can be used for an Azure Local or Windows Server S2D solution.

Although network switches are not specifically certified under the Microsoft HCI certification programs, all of the Lenovo certified configurations for Microsoft HCI discussed in this document have undergone rigorous end-to-end solution validation using Lenovo network switches to carry all solution traffic.

Table 5 lists the recommended Lenovo networking switches for Azure Local or Windows Server S2D deployment. These switches support the Remote Direct Memory Access (RDMA) feature of Microsoft SMB 3.x, which is used extensively by Microsoft HCI solutions and are fully compatible with the Mellanox ConnectX-4 Lx network adapters used in these solutions to provide the highest storage performance.

Table 5 Recommended Lenovo network switches for Azure Local or Windows Server S2D

Lenovo Switch	Speed	Part Number	Feature Code
RackSwitch™ G8272	10GbE	7159CRW/7159CFV	ASRD/ASRE
ThinkSystem NE1032 RackSwitch	10GbE	7159A1X/7159A2X	AU3A/AU39
ThinkSystem NE2572 RackSwitch	10/25GbE	7159E1X/7159E2X	AV19/AV1A
ThinkSystem NE10032 RackSwitch	100GbE	7159D1X/7159D2X	AV17/AV18

Note: The first part number and feature code listed in Table 5 is for a switch with rear to front airflow. The second part number and feature code is for front to rear airflow.

RackSwitch G8272

The Lenovo RackSwitch G8272 uses 10 Gb SFP+ and 40 Gb QSFP+ Ethernet technology and is specifically designed for the data center. It is ideal for today's big data, cloud, and optimized workload solutions. It is an enterprise class Layer 2 and Layer 3 full featured switch that delivers line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data. Large data center-grade buffers help keep traffic moving, while the hot-swap redundant power supplies and fans (along with numerous high-availability features) help provide high availability for business sensitive traffic. In addition to the 10GbE and 40GbE connections, the G8272 can use 1GbE connections.

ThinkSystem NE1032 RackSwitch

The Lenovo ThinkSystem NE1032 RackSwitch is a 1U rack-mount 10 GbE switch that delivers lossless, low-latency performance with feature-rich design that supports virtualization, Converged Enhanced Ethernet (CEE), high availability, and enterprise class Layer 2 and Layer 3 functionality. The switch delivers line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data. The NE1032 RackSwitch has 32x SFP+ ports that support 1 GbE and 10 GbE optical transceivers, active optical cables (AOCs), and DACs. The switch helps consolidate server and storage networks into a single fabric, and it is an ideal choice for virtualization, cloud, and enterprise workload solutions.

ThinkSystem NE2572 RackSwitch

The Lenovo ThinkSystem NE2572 RackSwitch is designed for the data center and provides 10/25 GbE connectivity with 40/100 GbE upstream links. It is ideal for big data, cloud, and enterprise workload solutions. It is an enterprise class Layer 2 and Layer 3 full featured switch that delivers line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data. Large data center-grade buffers help keep traffic moving, while the hot-swap redundant power supplies and fans (along with numerous high-availability software features) help provide high availability for business sensitive traffic. The NE2572 RackSwitch has 48x SFP28/SFP+ ports that support 10 GbE SFP+ and 25 GbE SFP28 optical transceivers,

AOCs, and DACs. The switch also offers 6x QSFP28/QSFP+ ports that support 40 GbE QSFP+ and 100 GbE QSFP28 optical transceivers, AOCs, and DACs. These ports can also be split out into four 10 GbE (for 40 GbE QSFP+) or 25 GbE (for 100 GbE QSFP28) connections by using breakout cables.

ThinkSystem NE10032 RackSwitch

The Lenovo ThinkSystem NE10032 RackSwitch uses 100 GbE QSFP28 and 40 GbE QSFP+ Ethernet technology and is specifically designed for the data center. It is ideal for today's big data, cloud, and enterprise workload solutions. It is an enterprise class Layer 2 and Layer 3 full featured switch that delivers line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data. Large data center-grade buffers help keep traffic moving, while the hot-swap redundant power supplies and fans (along with numerous high-availability features) help provide high availability for business sensitive traffic. The NE10032 RackSwitch has 32x QSFP+/QSFP28 ports that support 40 GbE and 100 GbE optical transceivers, AOCs, and DACs. These ports can also be split out into four 10 GbE (for 40 GbE ports) or 25 GbE (for 100 GbE ports) connections by using breakout cables.

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 Onyx 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 Onyx 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 Onyx 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 Azure Local or Windows Server 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 ThinkSystem SR650 server, order Feature Code B0MK to enable TPM 2.0 or Feature Code AUK7 to enable TPM 2.0 and Secure Boot.

Note that Secure Boot is required by Microsoft for all ThinkAgile MX Appliance solutions, so will be selected by default for all Appliance 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. For the SR650, order Feature Codes B173.

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 provides important benefits for managing a Microsoft Azure Local or Windows Server S2D 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 Local instances as well as Lenovo ThinkAgile MX Appliances 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 for 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://pubs.lenovo.com/lxci-wac

Summary

Lenovo is a key partner in the Microsoft WSSD and Azure Local 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, Lenovo's 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 Local 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 Azure Local or

Windows Server S2D solution will save time, money, and effort associated with designing and building a do-it-yourself solution.

Change History

Changes in the January 2025 update:

- Updated terminology based on Microsoft's change of the Azure Stack HCI operating system name to "Azure Local"

Changes in the December 2024 update:

- Removed information for Lenovo Edge servers, which can now be found in a separate document specifically written for solutions built using these servers
- Made several general updates to content related to Windows Server and Azure Local operating systems

Changes in the December 2023 update:

- Updated "Storage Network Adapters" on page 16 and "Non-storage Network Adapters" on page 16 to clarify which network adapters have been certified for each traffic type (storage, management, and compute).

Changes in the April 2023 update:

- Updated certified storage device lists (Table 2 on page 17 and Table 3 on page 18)
- Updated storage device EOL table (Table 4 on page 20)

Changes in the February 2023 update:

- Updated list of supported GPUs (page 22)
- Updated certified storage device lists (Table 2 on page 17, Table 3 on page 18, and Table on page 20)
- Separated the certified storage devices for SE350-based solutions into a new table (Table on page 20)
- Updated storage device EOL table (Table 4 on page 20)

Changes in the June 2022 update:

- Added NVIDIA A2 and A30 GPUs to list of supported GPUs (page 22)
- Added 4350 HBAs to list of supported Storage HBAs (page 23)
- Added Intel E810 NIC to list of supported iWARP network adapters (page 23)
- Updated certified storage device lists (Tables 3 and 4)
- Updated storage device EOL table (Table 5)

Changes in the January 2022 update:

- Updated certified network adapters (page 23)
- Added ThinkSystem 440-8i and 440-16i SAS/SATA PCIe Gen4 12Gb HBAs (page 23)
- Updated certified storage device lists (Tables 3 and 4)
- Updated storage device EOL table (Table 5)
- Clarified the use of NationZ TCM instead of TPM in China (page 29)

Changes in the October 2021 update:

- Updated the document title to help differentiate it from its companion document for ThinkSystem SR630 V2 and SR650 V2 servers (https://lenovopress.com/lp1520)
- Added details regarding TPM 2.0 and Secure Boot security options
- Added details regarding Lenovo XClarity Integrator for WAC

Changes in the July 2021 update:

- Updated the list of GPUs supported in ThinkAgile MX solutions in “Network switches” on page 21
- Updated storage device EOL table (Table 5)

Changes in the April 2021 update:

- Added notes regarding support for NVIDIA T4 GPU in ThinkAgile MX1020 and MX1021 solutions
- Added section “Lenovo network switches” on page 21 that provides details regarding NVIDIA (Mellanox) network switches that have been tested with ThinkAgile MX solutions
- Added the ThinkAgile MX1020 on SE350 Appliance to the notes

Changes in the February 2021 update:

- Updated all-flash storage device table (Table 4) with new supported drives for SE350-based ThinkAgile MX solutions
- Updated storage device EOL table (Table 5)

Changes in the January 2021 update:

- Corrected GPUs shown in Table 6

Changes in the December 2020 update:

- Added a brief description of ThinkAgile MX Appliance offerings in “ThinkAgile MX Integrated System (IS)” on page 4
- Added “Network switches” on page 21 to provide information regarding GPU adapter support for ThinkAgile MX solutions
- Updated storage device tables (Table 2 on page 17 and Table 4) with new supported drives
- Removed supported OS boot devices from storage device tables
- Added supported OS boot devices to “Component selection” on page 15

Changes in the August 2020 update:

- Added reference to Lenovo Press document *ThinkAgile MX1021 on SE350 Azure Stack HCI (S2D) Deployment Guide*
- Updated configuration details for some all-flash configurations, since the Cavium/QLogic 25GbE network adapter has now been certified for these configurations
- Added Mellanox ConnectX-6 HDR100 QSFP56 PCIe InfiniBand Adapters (1-port and 2-port models) as supported and certified network adapters for high performance all-flash configurations
- Added the SE350 Wireless Network Module as a supported and certified network module for ThinkAgile MX1021 on SE350

- Updated storage device tables (Tables 3 and 4) with new supported drives and removed devices that are no longer available
- Updated storage device tables (Tables 3 and 4) to include boot device
- Updated Table 5 with additional storage devices that have or will soon reach their end of life

Changes in the May 2020 update:

- Added ThinkAgile MX1021 on SE350 to the ThinkAgile MX Certified Nodes family, including example configurations NN16T1a (single-tier all-NVMe), NN12T1a (two-tier all-NVMe), and SS08T1a (single-tier all-SSD)
- Added a section that discusses special configuration details for the ThinkSystem SE350 Edge Server in ThinkAgile MX1021 on SE350 solutions
- Added supported data storage devices to Table 4 for ThinkAgile MX1021 on SE350

Changes in the October 2019 update:

- Changed the document title to accurately reflect Microsoft's change of "Storage Spaces Direct" to "Azure Stack HCI"
- Added multiple comments, mainly in "Component selection" on page 15, regarding the Lenovo ThinkSystem SE350, which has been certified for Azure Stack HCI, but is not yet offered as a ThinkAgile MX Certified Node

Changes in the July 2019 update:

- Replaced NS58T1a all-flash configuration with NS61T1a to ensure the number of capacity devices (16) is an equal multiple of cache devices (4)
- Added an All-SSD configuration example (SS92T1a)
- Added an All-NVMe configuration example (NN38T1a)
- Updated storage device tables
- Corrected typos and updated graphics for accuracy

Changes in the March 2019 update:

- Added information about the Microsoft Azure Stack HCI program
- Added Lenovo ThinkSystem SR630 as a certified general purpose server
- Added ThinkSystem 3.5" Intel P4610 NVMe devices as replacements for P4600 devices
- Added ThinkSystem U.2 Intel P4610 NVMe devices as replacements for P4600 devices

Changes in the December 2018 update:

- Clarified the availability of ThinkAgile MX Certified Nodes configuration in the Lenovo Data Center Solution Configurator
- Removed configuration SH32T1a, since it is exactly the same as SH40T1a with fewer HDD capacity devices

Changes in the November 2018 update:

- Added ThinkAgile MX Certified Node description and details
- Updated the layout for Table 1 to improve readability
- Added 1-port Mellanox NIC as an option if two 1-port NICs are preferred
- Added section "Small cluster configurations" on page 14 to provide additional detail for these configurations

- Added ThinkSystem 430-8i SAS/SATA 12Gb HBA for all-flash configurations
- Split storage device table into two tables, one for Hybrid Storage configs (Table 2 on page 17) and one for All-Flash configs (Table 4)

Changes in the August 2018 update:

- Added Cavium/QLogic network adapter to the list of certified NICs
- Updated Nodes column in Table 1, including 2-node and 16-node configurations
- Updated certified storage devices shown in Table 3
- Updated processor selection criteria in the “Component selection” on page 15

Changes in the May 2018 update:

- Added configuration SR650-NH120T1a
- Updated document title and content regarding certified configurations

Authors

This paper was produced by the following specialists:

Dave Feisthammel is a Senior Solutions Architect working at the Lenovo Center for Microsoft Technologies 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.

Guy Fusman is a Windows Engineer working at the Lenovo Center for Microsoft Technologies in Bellevue, Washington. He has over 25 years in the IT industry, focused on client, server, and cloud solutions support and development. For the last 8 years he has primarily worked with Microsoft Azure Stack Hub and Azure Stack HCI solutions running on Lenovo servers.

Mike Miller is a Windows Engineer working at the Lenovo Center for Microsoft Technologies 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
- Hussein Jammal - Advisory Engineer, Microsoft Solutions Lead EMEA
- Vinay Kulkarni, Principal Technical Consultant - Microsoft Solutions and Enablement
- Oana Adelina Onofrei, Solutions Engineer - ISG Software Development
- Laurentiu Petre, Solutions Engineer - ISG Software Development

- Vy Phan, Technical Program Manager - Microsoft OS and Solutions
- 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 January 9, 2025.

Send us your comments via the **Rate & Provide Feedback** form found at
http://lenovopress.com/lp0866

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:

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

The following terms are trademarks of other companies:

Intel, 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.