Reference Architecture: Microsoft SQL Server on Azure Local with ThinkAgile MX650 V3

Last update: **24 June 2025**

Version 1.1

Technical overview of Azure Local Architecture on MX Series servers

Contains SQL Server performance test data and sizing guidance

Covers reliability and performance benefits of hyper-converged solutions

Includes detailed configuration and best practices for optimized performance

David West
Vinay Kulkarni
Laurentiu Petre

Click here to check for updates

Table of Contents

1 Introduction 1

2 Business problem and business value 2

2.1 Business problem 2

2.2 Business value 2

3 Architectural overview 4

4 Component model 5

4.1 Microsoft Azure Stack HCI OS 5

4.2 Windows Admin Center 5

4.3 Microsoft Storage Spaces Direct 5

4.4 Lenovo ThinkAgile MX Certified Nodes 6

4.5 Microsoft SQL Server 2022 6

5 Deployment considerations 8

5.1 Hardware resources 8

5.2 Software resources 9

5.3 Network configuration 9

5.4 Storage Configuration 9

5.5 SQL Server HA considerations 10

5.6 Best practices 10

6 Solution validation 12

6.1 Monitoring tools 12

6.2 Workload generation tool 12

6.3 Performance Test Results 12

6.4 Processor generation comparisons 13

7 Conclusion 14

8 Appendix: Lenovo Bill of materials 15

Resources 16

Document history 17

Trademarks and special notices 18

1 Introduction

This document describes the reference architecture for Microsoft SQL Server on Microsoft Azure Local and ThinkAgile MX series systems. The Lenovo workloads team worked together on the architectural design, testing and engineering effort to create this Reference Architecture. This paper in intended to provide the planning, design considerations, and best practices for implementing Microsoft SQL Server on Azure Local with Lenovo server products.

Deploying hyperconverged infrastructure (HCI) has become the new standard for organizations looking to modernize their infrastructure. Large storage deployments are being replaced by HCI-based solutions for most general-purpose workloads. HCI delivers 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.

Microsoft SQL Server is one of the most widely adopted business critical database solutions on the virtualized platform that requires an infrastructure layer with consistent performance, high availability, and simplified management. With the evolution of the modern hardware platforms—Server CPU cores count increasing dramatically, 25/100Gbps networking growing rapidly, and NVMe SSD drives of NAND flash technology becoming dominant, it is time to redefine how SQL Server environments can leverage the new hardware platform to take a big leap forward.

In this solution, we provide design and virtualized sizing guidance, solution validation reports, and best practices for enterprise infrastructure administrators and SQL Server application workload owners to run virtualized SQL Server workloads on Azure Local with Lenovo ThinkAgile MX Series.

The architecture described herein has been designed and validated by Lenovo SQL and HCI engineers, while following Microsoft best practices for Hyper-Converged Infrastructure deployments.

The intended audience of this document is IT professionals, technical architects, sales engineers, and consultants to assist in planning, designing and implementing Microsoft SQL Server on Azure Local instances. It is assumed that the reader is familiar with the concepts, administration and operations of Microsoft SQL Server, Windows Admin Center, Windows command line and Lenovo ThinkAgile MX Series, and related components.

2 Business problem and business value

This section will cover some of the business problems modern IT organizations are facing, followed by how this solution solves the problem and provides business value.

2.1 Business problem

In many of today's enterprise data centers, IT administrators are inundated with excessive amounts of data. Traditional servers no longer provide the return on investment nor meet the required performance efficiency to manage all of this data. Companies are also faced with increasing data storage costs and inadequately tuned hardware bandwidth.

Additionally, organizations are struggling with complex processes and manual operating procedures that introduce potential for human error while managing the companies most critical data. Configuring high availability is often a complex task that ends up misconfigured resulting in data loss or business down time. IT managers would rather spend time optimizing and running the business applications, not worrying about the underlying infrastructure.

2.2 Business value

Microsoft Azure Local provides a high level of performance and efficiency from next-generation hardware devices with optimized throughput, superior efficiency, enhanced resilience, and agile operation. Business critical applications like SQL Server are one of the top use cases for virtualized workloads on Azure Local that offers expedited deployment with highly automated processes and reduces storage costs compared to legacy storage solutions. This reference architecture is also a showcase of Windows Storage Spaces Direct (S2D) which virtualizes internal storage into pools. All of this is enabled on the Lenovo ThinkAgile MX Series for operating and managing Microsoft SQL Server database workloads in a fully integrated SDDC environment.

Lenovo ThinkAgile MX Series with Azure Local enables customers to focus on the outcome of mission-critical applications rather than on building your infrastructure by providing simple, agile, and cost-effective solutions to support the business needs. ThinkAgile MX Series helps to modernize the data center or implement a hybrid cloud with Azure cloud integration options. Available in several models, the MX Series enables you to start with as few as two nodes and scale capacity and performance easily, without disruption to day-to-day operations.

Here are the benefits of deploying SQL Server in an Azure Local environment:

- Next-generation HCI: Azure Local unlocks the capabilities of modern hardware by adding optimization for high-performance, NVMe-based flash devices with Storage Spaces Direct.
- Reliable resilience: Azure Local improves data availability with fail-over upon failure and cloud-based data protection workflows for backup/restore.
- Intuitive operations: Azure Local provides simplified operations for storage device provisioning and servicing, and proactive insights to detect anomalies and prevent potential issues with Windows Admin Center.

- Azure Local delivers scalable and consistent performance for mission-critical SQL Server database workloads both vertically (scale-up) and horizontally (scale-out).

3 Architectural overview

In this solution, we validated different numbers of VMs to determine the best performance scalability for SQL Server 2022 virtualized workloads running on Azure Local.

From a scale up perspective, we increased the number of SQL Server VMs from 8 to 16 and showcased the performance capability of Azure Local instance by running multiple database workloads simultaneously.

Azure Local is designed for performance scalability, highly available data resiliency and rich data services by taking advantage of modern hardware. High performance can be achieved using Hyper-V and Storage Spaces Direct technology which are built into Windows Server. Several technologies like NVMe storage and Remote Direct Memory Access (RDMA) networking are natively supported in Windows Server to enable the highest levels of performance.

The diagram below shows the overall HCI and S2D architecture. It starts with a foundation of servers running Azure Stack HCI OS with physical hard drives that are placed into storage pools. From these pools volumes are provisioned to host virtual disks (VHD). Hyper-V virtual machines and their data volumes run on the virtual disks. The drawing shows 4 hosts for illustration purposes, for our testing we used a 2-node instance.

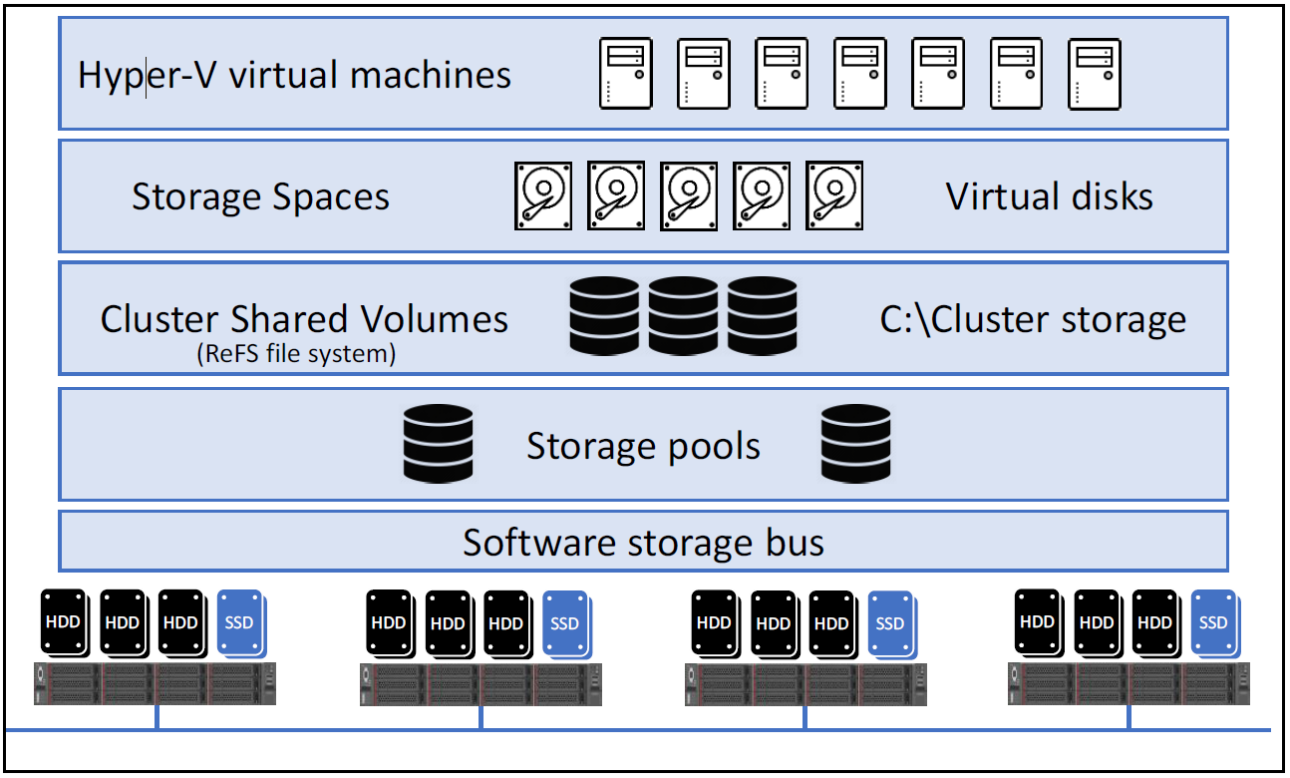

Figure 1 Architecture of Azure Local and Storage Spaces Direct

4 Component model

The solution technology components covered in this section include the following:

- Microsoft Azure Stack HCI Operating System
- Windows Admin Center (WAC)
- Microsoft Storage Spaces Direct (S2D) software defined storage
- Lenovo ThinkAgile MX Series and MX Certified Nodes
- Microsoft SQL Server 2022

4.1 Microsoft Azure Stack HCI OS

Azure Stack HCI OS is a Microsoft operating system which includes Hyper-V and comprises Microsoft's virtualization platform that transforms data centers into aggregated computing infrastructures that include CPU, memory, storage, and networking resources. This 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.

4.2 Windows Admin Center

Windows Admin Center (WAC) is the main tool for deploying and managing these infrastructures as a unified operating environment and provides operators with the tools to administer the servers and instances in the environment. It is a graphical interface with wizard driven deployment guidance during setup, and ease of administration after deployment. We also have available the Lenovo XClarity Integrator (LXCI) for WAC. With this WAC extension, firmware and device driver updates can be installed on all servers that will become nodes in the instance directly from the deployment wizard.

4.3 Microsoft Storage Spaces Direct

S2D is illustrated in the architecture drawing above. It uses the server's internal physical hard drives organized into virtualized storage pools. From these pools, volumes are provisioned and used to host virtual disks. S2D is setup and configured during the WAC deployment method. However, the servers and network switches need to be prepared ahead of time, to enable and optimize RDMA for east-west storage network traffic. The preparation steps are covered in another Lenovo document - *Microsoft Storage Spaces Direct (S2D) Deployment Guide*, found at the following URL: https://lenovopress.com/lp0064

4.4 Lenovo ThinkAgile MX Certified Nodes

Lenovo ThinkAgile MX is a fully integrated system built on the most reliable and secure servers that are tested and validated for Azure Local and S2D compliance. ThinkAgile MX Series helps you modernize your core data center, edge locations, and deploy a hybrid cloud environment using Azure cloud and services. With all-NVMe devices, ThinkAgile MX Series for Azure Local scales compute or storage capacity easily and accelerates performance for your vital business-critical application deployments in both on-prem and cloud environments.

Lenovo ThinkAgile MX Certified Nodes and Appliance solutions contain only servers and server components that have been certified under the Microsoft Azure Local Program to run Microsoft Storage Spaces Direct (S2D) properly. These solutions provide a solid foundation for customers looking to consolidate both storage and compute capabilities on a single hardware platform. They provide outstanding performance, high availability protection and effortless scale-out growth potential to accommodate evolving business needs.

ThinkAgile MX supports the latest 4th Gen Intel® Xeon® processor family CPU´s or AMD 4th Generation EPYC™ Series Processors that boast enhancements in performance, security, and energy efficiency. These new generations support new DDR memory, PCIe Gen5 PCIe I/O. All-NVMe SSD models support inline de-duplication, compression, and encryption to give you an optimized, secure, high-performance platform with maximum usable capacity.

4.5 Microsoft SQL Server 2022

Microsoft SQL Server 2022 is the new data platform to help ensure uptime with fully managed disaster recovery, high availability and integration with Azure cloud services and management. It takes advantage of performance and availability for faster queries and to help ensure business continuity and accelerate query performance and tuning with no code changes required.

SQL Server 2022 includes updates to existing features like Intelligent Query Processing in addition to management, platform or language.

Starting with SQL 2022, runtimes for R, Python, and Java are no longer installed with SQL Setup. Instead, install any desired custom runtime(s) and packages.

Performance enhancements in SQL Server 2022 include:

- Improvements have been made to all columnstore indexes that benefit from enhanced segment elimination by data type.
- Concurrent updates to global allocation map pages reduce page latch contention
- Improvements in buffer pool scan operations on large-memory systems by using multiple CPU cores for parallel scans
- Improvements to Clustered ColumnStore Indices to sort existing data in memory before index builder compresses the data
- Support for Intel QuickAssist Technology (QAT) backup compression with software or hardware acceleration

- TempDB performance enhancements for scalability
- Shrink database uses low priority processing to minimize impact on performance
- In-memory OLTP enhancements

SQL Server management improvements include:

- Additional Azure integration
- Link to Azure SQL Managed Instance
- Accelerated Database Recovery (ADR)
- Always On Availability Group enhancements

In this solution, we adopted SQL Server 2022 workloads on Lenovo ThinkAgile MX Series using all NVMe SSD drives with Azure Local to demonstrate unprecedented performance capability, scalability, and resiliency.

5 Deployment considerations

This section covers the solution deployment resources and configuration, including the following.

- Hardware resources
- Software resources
- Network configuration
- Storage configuration
- SQL Server HA considerations

5.1 Hardware resources

The solution was composed of the following hardware, to create a 2 node Azure Local instance.

Table 1 Hardware configuration (per server)

PROPERTY	SPECIFICATION
Server model name	2 x Lenovo ThinkAgile MX650 V3 Certified Node
CPU	2 x Intel Xeon Platinum 8480+ 56C 350W 2.0GHz 4th Gen Scalable Processors
RAM	4 TB DDR5 4800
Network adapter	Mellanox ConnectX-6, 10/25 Gb/s
Disks	8 x 1.6TB Intel 5620 NVMe SSD

5.2 Software resources

The following software and levels were used in the solution.

Table 2 Software resources

Software	Version	Purpose
Azure Stack HCI OS	22H2	Microsoft HCI Operating System
Storage Spaces Direct	22H2, feature in OS	Virtualized storage pools and volumes to host SQL server Hyper-V VMs
Hyper-V	22H2, feature in OS	Microsoft's Hypervisor and VM management tools
Windows Admin Center	2211	Web-based tool for Azure Local deployment and administration
Microsoft SQL Server 2022 Enterprise	RTM build – 16.0.1000.6	Microsoft SQL Server database server platform, running on VMs
HammerDB	4.6	Generate OLTP workload derived from TPC-C like benchmark

5.3 Network configuration

Azure Local typically requires a 25GbE network as minimum, to support the high-speed storage network. Remote Direct Memory Access (RDMA) is used as well. RDMA has lower CPU utilization and less I/O latency for database workloads.

In this solution, we used 10/25 GbE network with RDMA to fully exploit the capabilities of high-performing NVMe devices used by the S2D network.

For more details about the physical switch and NIC configuration, please reference the companion document for S2D configuration, *Microsoft Storage Spaces Direct (S2D) Deployment Guide*, at the following URL: https://lenovopress.com/lp0064

5.4 Storage Configuration

In this solution, we used S2D with mirroring as the best mix of resiliency and performance. Each server included 8 NVMe drives which were all used for the storage pool, as the underlying storage for the SQL Server data volumes – including the Databases, Logs, TempDB. Each server's boot drive used M.2 SSDs, to

free up all the NVMe drives for the data storage pool.

5.5 SQL Server HA considerations

This section covers SQL Server high availability options on Azure Local.

SQL Server supports two types of high availability and disaster-recovery technologies that are both fully supported on Azure Local: Always on Availability Groups (AG) and Failover Cluster Instances (FCI).

Always On Availability Group

SQL Server AG maximizes the availability of a set of user databases that fail over together—primary database and up to 8 secondary replicas. By running AG on Azure Local, it further provides database-level HA protection in addition to the HCI host level instance availability.

Failover Cluster Instance

Azure Local supports persistent reservation on a shared virtual disk which is required by SQL Server FCI to create a highly scalable, available, and high-performance HCI solution. Windows Cluster Shared Volumes (CSV) provide access to shared storage with resiliency and failover. By taking advantage of CSVs, shared disks that are required by SQL Server FCI are easy to manage, and simple to migrate.

5.6 Best practices

The following list of recommendations provide the best practices and sizing guidance to run SQL Server on Azure Local.

Size the SQL Server VM based on the performance and capacity requirements of the desired workload in terms of the database customer count, IOPS numbers, and others for typical OLTP workloads.

Avoid over-committing physical CPU resources even if hyper-threading is enabled, since over-committing CPU resources does not guarantee further performance.

For mission-critical workloads to achieve better performance, configure the Operating Mode for Maximum Performance in the BIOS settings.

Enable Hyper-Threading in the BIOS settings.

Configure the power profile in Windows Server to "High performance".

Set the Windows page file to fixed size of 4 GB, and ensure it is on the C drive, not on and data drives.

Used fixed size virtual disks (VHD) for the SQL data drives not dynamic thin provisioned

SQL Server database and log drives are recommended to be formatted with 64KB NTFS cluster size.

The OS and SQL Server binary drives are recommended to be formatted with standard default 4KB NTFS cluster size.

TempDB is shared by many processes and users as a temporary working area. Default configuration will be suitable for most workloads. Use the SQL install experience for guided configuration. More info in Microsoft TempDB Database documentation. Experiment to determine whether more or larger TempDB files are

beneficial for your test environment.

If the server is dedicated to the SQL Server workload, use the default dynamic memory management model or follow Microsoft SQL documentation guidelines for manually configuring memory options if finer grain control is desired. In general, the recommendation is to reserve 80% of available memory for SQL.

Use 25Gbps network as a minimum requirement. For better performance and throughput, consider 40Gbps/100Gbps network for mission-critical SQL Server workloads.

Use RDMA capable network adapters to offload host CPU utilization and boost the SQL Server workloads as needed.

Use multiple database files and separate data, log, and TempDB on different virtual disks.

For mission-critical workload, configure the max degree of parallelism (MAXDOP) depending on the number of logical processors per NUMA node. For OLTP workloads, this is usually set to 1 which disables parallelism. However, consider experimenting with different MAXDOP settings depending on your test environment. See *Configure the max degree of parallelism*.

6 Solution validation

This section covers the solution testing. It includes the monitoring and benchmark tools, configuration and final results.

6.1 Monitoring tools

The following tools were used to monitor and log the performance testing.

Windows Performance Monitor

Windows Performance Monitor is a Windows tool that enables users to capture statistics about CPU, memory, and disk utilization from operating system levels. It also provides counters for monitoring SQL Server performance and status.

Here are the major performance counters we collected during the tests:

- Logical disk
- Processor information
- SQL Server: Database TPS

HammerDB Transaction Counter

HammerDB provides a real-time graph view of the transactions rates occurring. This can help gauge whether the test is performing as expected or is experiencing performance bottlenecks.

6.2 Workload generation tool

The OLTP workload tested in this solution is derived from Microsoft TPC-C using HammerDB. This tool can run industry-standard benchmark and scalability testing. With HammerDB, you can verify your database environment size and loads.

HammerDB is an open-source load testing and benchmarking tool for databases, available from http:///www.hammerdb.com. It offers tools for testing performance of OLTP and Analytics workloads. The OLTP workload is based on TPC-C benchmark from http://www.tpc.org. Hammerdb was run on separate load servers to avoid impacting performance of the system under test (SUT). Network connectivity between the load server and SUT should allow for 3 Gb/s bandwidth per instance.

6.3 Performance Test Results

This section covers the test server configuration details and results of performance testing.

Important note: The OLTP workload tested in this solution is derived from the TPC-C Benchmark and is not comparable to published TPC-C Benchmark results, as the implementation does not comply with all requirements of the TPC-C Benchmark. The TPM results published in this solution is a normalized value.

Benchmarks simulated: TPC-C on 8 VMs and 16 VMs

Database size: Each VM has a 100 GB database, comprised of 800 warehouses. VMs were distributed equally between the two nodes.

Processor: Intel 4th Generation 8280+ 64core

Run time parameters:
Virtual users: 400
User delay: 1 ms

The following table lists the benchmark results.

Table 3 Benchmark results

Metric	**8 VMs Total**	**16 VMs Total**
Transactions Per Minute (TPM)	7.5 million	6.6 million

6.4 Processor generation comparisons

This section looks at the previous CPU generation and compares the sizing and performance results.

These comparisons were done over several years, and by different teams, so the test methods and parameters may vary slightly.

The following table summarizes the comparisons of TPC-C results from multiple VMs.

Table 4 Benchmark details and comparison

Database tested	**SQL 2022 Enterprise**	**SQL 2019 Enterprise**
Processor generation	SR650 V3, 4th Gen Intel Sapphire Rapids 8480+ 64C 2.0 GHz	SR650 V2, 3rd Gen Intel Ice Lake 6330 28C 2.2 GHz
Hardware configuration (per node)	2 TB RAM 12 1.6TB NVMe All Flash	512 GB RAM Cache: 2x 3.2 NVMe Capacity: 4x 14TB SATA SSD
Warehouses	800	800
Virtual Users	400	400
Number of nodes	2	2 and 4
TPM (million) 4 VMs		4.0 (2 node)
TPM (million) 8 VMs	7.5	8.0 (4 node)
TPM (million) 16 VMs	6.6	

7 Conclusion

Running SQL Server on Azure Local on Lenovo ThinkAgile MX Series is a fast way to unlock capabilities, improve operation efficiency, and achieve high performance with industry-leading HCI on a next-generation hardware platform for mission-critical database workloads.

In this solution, we demonstrated the performance of consolidating SQL Server VMs running on Azure Local. By combining the latest hardware advancements with the rich feature set of Azure Local, Hyper-V, and Storage Spaces Direct, resiliency and high performance can be achieved with technology that comes built into Windows Server. Several technologies like NVMe storage and Remote Direct Memory Access (RDMA) networking are natively supported in Windows Server to enable the highest levels of performance.

8 Appendix: Lenovo Bill of materials

This appendix contains the bill of materials (BOMs) for the solution.

The BOM lists in this appendix are not meant to be exhaustive and must always be double-checked with the configuration tools. Any discussion of pricing, support, and maintenance options is outside the scope of this document.

7D76CTO1WW	Server: ThinkSystem SR650 V3	1
		1
BNOM	Intel Xeon Platinum 8480+ 64C 350W 2.0GHz Processor	2
BNFC	ThinkSystem 128GB TruDDR5 4800 MHz (4Rx4) 3DS RDIMM	32
B8NY	ThinkSystem RAID 940-8i 4GB Flash PCIe Gen4 12Gb Adapter	1
BNEG	ThinkSystem 2.5" U.2 P5620 1.6TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	8
B8LU	ThinkSystem 2U 8 x 2.5" SAS/SATA Backplane	1
BH8D	ThinkSystem 2U/4U 8 x 2.5" NVMe Backplane	1
BM8X	ThinkSystem M.2 SATA/x4 NVMe 2-Bay Enablement Kit	1
AUUV	ThinkSystem M.2 128GB SATA 6Gbps Non-Hot Swap SSD	2
B93E	ThinkSystem Intel I350 1GbE RJ45 4-port OCP Ethernet Adapter	1
BLKM	ThinkSystem V3 2U x16/x16/E PCIe Gen4 Riser1 or 2	2
BMUF	ThinkSystem 1800W 230V Platinum Hot-Swap Gen2 Power Supply	2
BLL6	ThinkSystem 2U V3 Performance Fan Module	6
BRQ1	ThinkSystem SR650 V3, SATA CBL, SLx8-SLx4, M.2-M.2(MB),150mm	1
BSYM	ThinkSystem SR650 V3, PCIe4 Cable, Swift8x-SL8x,2in1, PCIe 6/5(MB) to BP1/BP2/td>	1
BETS	ThinkSystem V3 2U SFF C0 (RAID) to Front 8x2.5" BP1	1
BPE3	ThinkSystem SR650 V3 MCIO8x to SL8x CBL, PCIe4, 8x2.5 AnyBay, 200mm	2

Resources

To learn more about this Lenovo solution contact your Lenovo Business Partner or visit the following pages.

- *Microsoft SQL Server 2022*
- *Lenovo ThinkAgile MX650 V3 Certified Nodes*
- *Lenovo guide to Storage Spaces Direct deployment*
- *Lenovo guide to Azure Stack HCI deployment with WAC*

Document history

Version 1.0	December 2023	First version of document, based on Lenovo RA template
Version 1.1	June 2025	Changed Azure Stack HCI to Azure Local Changed Azure Stack HCI cluster to Azure Local instance Changed cluster to instance

Trademarks and special notices

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

The following terms are trademarks of Lenovo in the United States, other countries, or both:
Lenovo®
AnyBay®
ThinkAgile®
ThinkSystem®
XClarity®

The following terms are trademarks of other companies:

Intel® and Xeon® are trademarks of Intel Corporation or its subsidiaries.

Microsoft®, Azure®, Hyper-V®, SQL Server®, Windows Server®, and Windows® are trademarks of Microsoft Corporation in the United States, other countries, or both.

TPC and TPC-C are trademarks of Transaction Processing Performance Council.

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

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

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

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

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

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

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

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

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