



# Reference Architecture: Microsoft SQL Server on VMware vSAN Express Storage Architecture with ThinkAgile VX Series

Last update: 22 November 2023

Version 1.0

---

Technical overview of VMware Express Storage Architecture on VX 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

Mark Xu

Deji Akomolafe

Catherine Xu

Brian Faleiro

David West

Chandrakandh Mouleeswaran



# Table of Contents

---

<b>1</b>	<b>Introduction.....</b>	<b>1</b>
<b>2</b>	<b>Business problem and business value .....</b>	<b>2</b>
2.1	Business problem .....	2
2.2	Business value.....	2
<b>3</b>	<b>Architectural overview .....</b>	<b>4</b>
<b>4</b>	<b>Component model .....</b>	<b>6</b>
4.1	VMware vSphere .....	6
4.2	VMware vSAN Express Storage Architecture .....	7
4.3	Lenovo ThinkAgile VX Series.....	8
4.4	Microsoft SQL Server 2022.....	8
<b>5</b>	<b>Deployment considerations.....</b>	<b>9</b>
5.1	Hardware resources.....	9
5.2	Software resources .....	10
5.3	Network configuration .....	10
5.4	Storage Configuration .....	10
5.5	SQL Server HA considerations .....	11
5.6	Best practices .....	12
<b>6</b>	<b>Solution validation.....</b>	<b>14</b>
6.1	Monitoring tools.....	14
6.2	Database workload generation tool.....	14
6.3	Scale-up OLTP workload test.....	15
6.4	Scale-out OLTP workload test.....	17
6.5	SQL Server data protection and snapshot test .....	18
<b>7</b>	<b>Conclusion.....</b>	<b>20</b>

<b>8 Appendix: Lenovo Bill of materials .....</b>	<b>21</b>
<b>Resources.....</b>	<b>22</b>
<b>Document history.....</b>	<b>23</b>
<b>Trademarks and special notices .....</b>	<b>24</b>

# 1 Introduction

---

This document describes the reference architecture for Microsoft SQL Server on vSAN Express Storage Architecture from VMware. The Lenovo and VMware teams worked together on the architectural vision, testing and joint engineering effort to create this Reference Architecture. This paper is intended to provide the planning, design considerations, and best practices for implementing Microsoft SQL Server on VMware vSAN Express Storage Architecture with Lenovo products.

Microsoft SQL Server is one of the most widely adopted business critical database solutions on the virtualized platform that requires the infrastructure layer with consistent performance, high availability, and simplified management. With the evolution of the modern hardware platforms—server CPU core 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 can leverage the new hardware platform to take a big leap forward.

In this solution, we provide design and virtualized sizing guidance, solution validation report, and best practices for enterprise infrastructure administrators and SQL Server application workload owners to run virtualized SQL Server workloads on VMware vSAN Express Storage Architecture with Lenovo ThinkAgile VX Series.

The architecture described herein has been validated by both Lenovo and VMware.

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 VMware vSAN Express Storage Architecture. It is assumed that the reader is familiar with the concepts, administration and operations of Microsoft SQL Server, VMware vSphere®, VMware vSAN and Lenovo ThinkAgile VX 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 company's 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

VMware vSAN™ 8 with [Express Storage Architecture](#)™ delivers a high level of performance and efficiency from your 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 vSAN Express Storage Architecture that offers expedited deployment with highly automated processes and reduces storage costs compared to legacy storage solutions. This reference architecture is a showcase of VMware vSAN Express Storage Architecture on Lenovo ThinkAgile VX Series for operating and managing Microsoft SQL Server database workloads in a fully integrated SDDC environment.

Lenovo ThinkAgile VX Series with vSAN Express Storage Architecture 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 VX Series helps to modernize data center or implement a hybrid cloud with VMware cloud platforms by arriving from the factory as a complete engineered system, preconfigured, tested, validated, software preloaded and optimized to provide the simplest way to bring a new VMware Hyperconverged Infrastructure (HCI) environment online. Available in several models, VX Series enables you to start with as few as three nodes and scale capacity and performance easily, without disruption to day-to-day operations.

Here are the benefits of deploying SQL Server in a VMware vSAN Express Storage Architecture environment:

- Next-generation HCI: vSAN Express Storage Architecture unlocks the capabilities of modern hardware by adding optimization for high-performance, NVMe-based TLC flash devices with vSAN.
- Performance without tradeoff: vSAN Express Storage Architecture delivers up to 4x higher performance, consistent and scalable with the new high performant RAID-5/6 technology at the performance of RAID-1. It also provides enterprise-grade snapshots with no trade-offs, and a negligible performance impact even with deep chains.

- Ready-for-anything resilience: vSAN Express Storage Architecture improves data availability with reduced failure domains upon failure and accelerates data protection workflows for backup/restore with scalable native snapshots.
- Superior resource and space efficiency: vSAN Express Storage Architecture saves up to 70% more usable capacity, and up to 40% lower TCO. With the new default RAID-5/6 technology, Express Storage Architecture has less capacity overhead over similarly protected RAID-1 VMs. Express Storage Architecture also enables data services on a per-VM/object basis.
- Intuitive and agile operations: vSAN Express Storage Architecture provides storage policies with simplified operations for storage device provisioning and servicing, and proactive insights to detect anomalies and prevent potential issues.
- vSAN Express Storage Architecture delivers scalable and consistent performance for mission-critical SQL Server database workloads both vertically (scale-up) and horizontally (scale-out).
- vSAN Express Storage Architecture provides high-available data protection and rich data services for SQL Server databases with the new snapshot technology. The snapshot has minimal performance impact for database workloads with simplified and fast snapshot management experience.
- An 8-node vSAN Express Storage Architecture cluster on Lenovo ThinkAgile VX Series can service up to 720K aggregated IOPS generated by VMware OLTP workloads derived from an industry standard benchmark toolkit. The results demonstrated in this solution showcase linear scalability and predictable database performance in terms of transaction per second (TPS): 14.x TPS improvement for the scale-up test and 7.x TPS improvement for the scale-out test.

## 3 Architectural overview

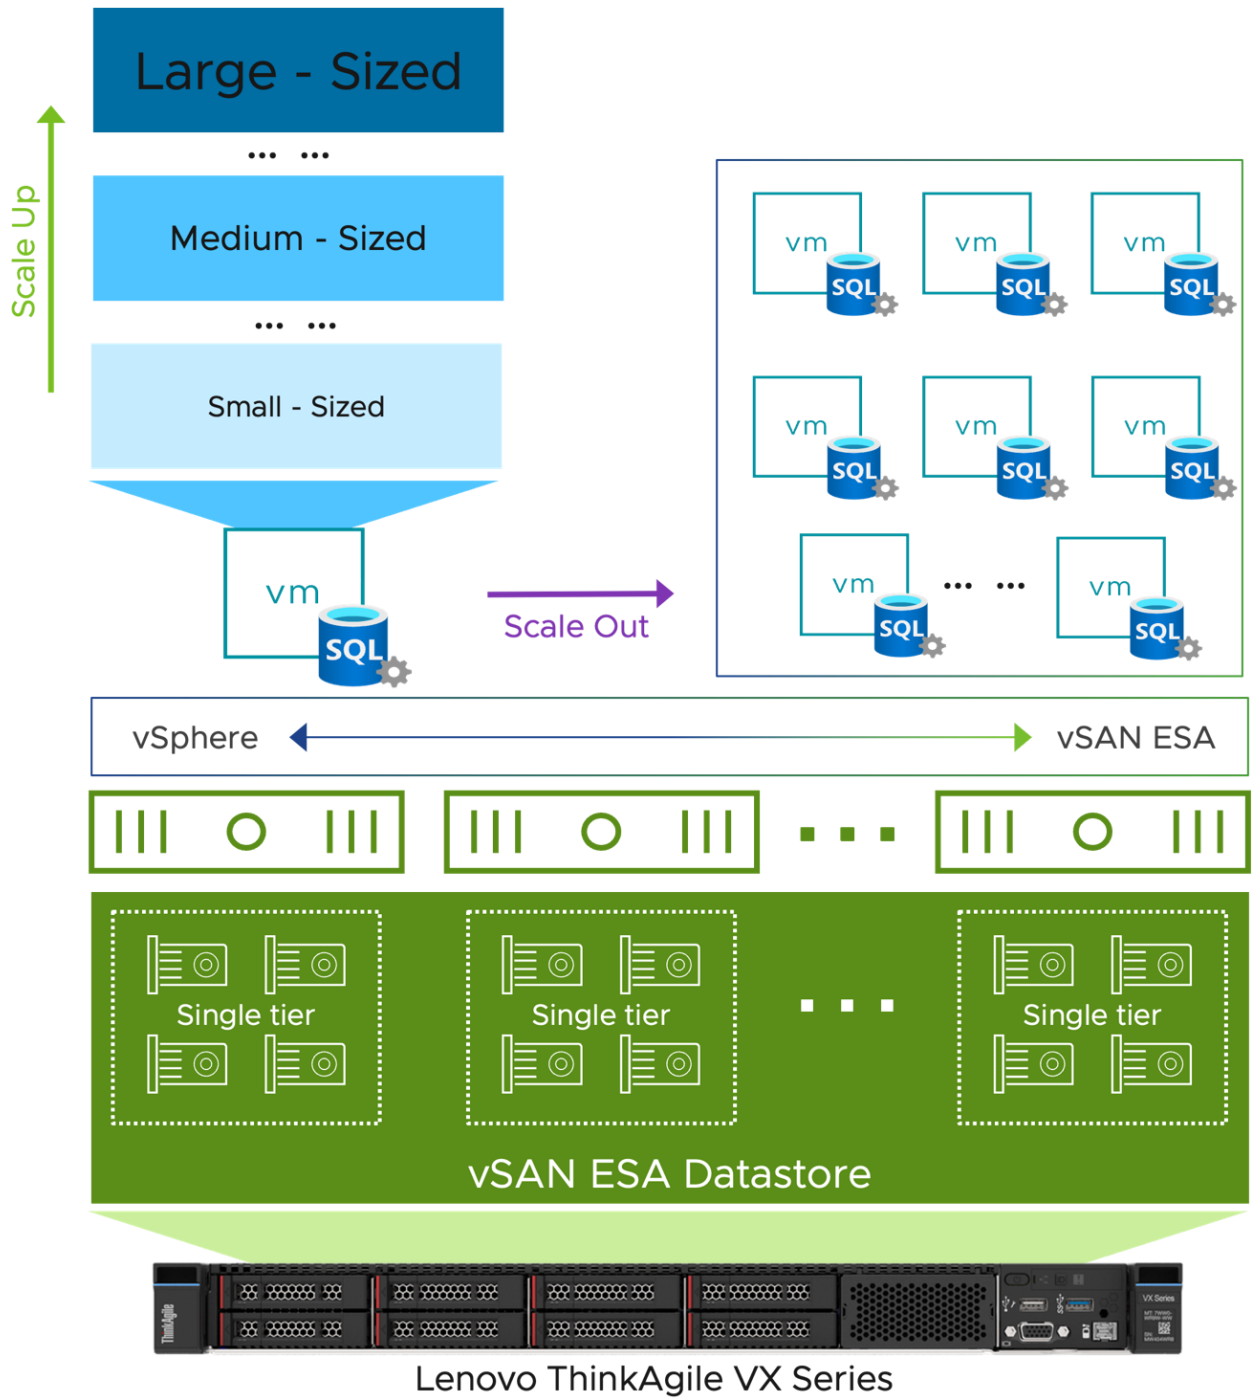
---

In this solution, we designed and validated different VM T-Shirt sizing to accommodate the best performance scalability and capability for SQL Server 2022 workloads running on vSAN Express Storage Architecture.

From the vertical perspective, we scaled up the SQL Server VM from small-sized, medium-sized to large-sized databases in terms of the number of target customers for each database. The scale-up test demonstrated the performance scalability and flexibility of vSAN Express Storage Architecture servicing business critical database workloads.

From the horizontal perspective, we scaled up the number of SQL Server VMs to up to 8 and showcased the overall performance capability of the vSAN Express Storage Architecture cluster by running the heavy database workloads simultaneously.

vSAN Express Storage Architecture is designed for excellent performance, scalability, highly available data resiliency and rich data services by taking advantage of modern hardware such as fast/efficient NVMe drives and 25/100Gbe networks. In this solution, we validated the above goals with the following solution architecture diagram that runs SQL Server 2022 on vSAN Express Storage Architecture with Lenovo ThinkAgile VX Series.



**Figure 1 Architecture Diagram**



# 4 Component model

The solution technology components include the following:

- VMware vSphere
- VMware vSAN Express Storage Architecture
- Lenovo ThinkAgile VX Series
- Microsoft SQL Server 2022

## 4.1 VMware vSphere

VMware vSphere is VMware's virtualization platform, which transforms data centers into aggregated computing infrastructures that include CPU, memory, storage, and networking resources. vSphere manages these infrastructures as a unified operating environment and provides operators with the tools to administer the data centers that participate in that environment. The two core components of vSphere are ESXi™ and vCenter Server®. ESXi is the hypervisor platform used to create and run virtualized workloads. vCenter Server is the management plane for the hosts and workloads running on the ESXi hosts.

VMware vSphere 8 is the new enterprise workload platform that brings the benefits of cloud to on-premises workloads. It supercharges performance with DPU and GPU based acceleration, enhances operational efficiency through the VMware Cloud Console, seamlessly integrates with add-on hybrid cloud services, and accelerates innovation with an enterprise-ready integrated Kubernetes runtime that runs containers alongside VMs.

**INTRODUCING**

**VMware vSphere 8**  
The Enterprise Workload Platform

- Get Cloud Benefits On-Premises**  
Enhance existing on-premises workloads in place with cloud services
- Supercharge Workload Performance**  
Meet the throughput and latency needs of modern distributed workloads
- Enhance Operational Efficiency**  
Efficiently reduce IT maintenance windows
- Accelerate Innovation for DevOps**  
Easily discover, access and deploy developer services

**Figure 2 vSphere 8 - The Enterprise Workload Platform**

## 4.2 VMware vSAN Express Storage Architecture

vSAN 8 introduced Express Storage Architecture as an optional, alternative storage architecture to the vSAN original storage architecture. When running on qualified hardware in approved vSAN ReadyNode, the vSAN Express Storage Architecture offers supreme levels of performance, scalability, resilience, and data services without compromising performance. The vSAN Express Storage Architecture unlocks the capabilities of modern hardware to allow the workloads of today and tomorrow perform as the hardware allows.

vSAN Express Storage Architecture evolves beyond the concept of disk groups, discrete caching, and capacity tiers, enabling users to claim storage devices for vSAN into a “storage pool” where all devices are added to a host’s storage pool to contribute to the capacity of vSAN. This improves the serviceability of the drives and the data availability management and helps drive down costs. A new fast and efficient data path paired with an efficient and resilient data structure give our users the ability to store and access data fast and efficiently. It also provides an all-new snapshot engine that meets high-level performance and scalability while remaining fully compatible for use with our ISV partners who provide backup solutions. And finally, vSAN Express Storage Architecture makes administering an environment simpler by improving how devices contribute as a storage resource to the cluster.

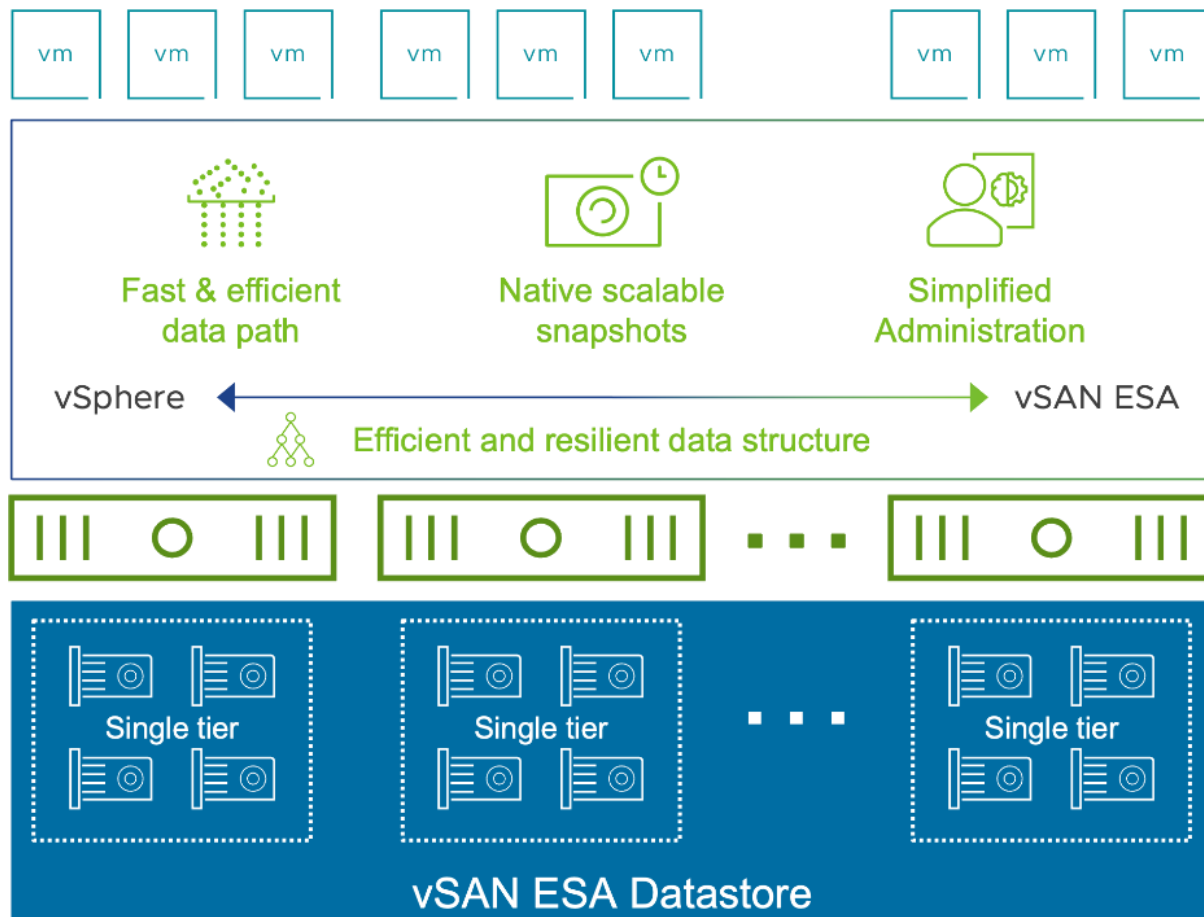


Figure 3 VMware vSAN Express Storage Architecture

## 4.3 Lenovo ThinkAgile VX Series

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

ThinkAgile VX 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 and our Intel platforms also enable hardware accelerator (Data Processing Unit – DPU) support. All-NVMe SSD models support inline deduplication, compression, and encryption to give you an optimized, secure, high-performance platform with maximum usable capacity.

Lenovo ThinkAgile VX Series streamlines the deployment, configuration, and management of your HCI environments using VX Deployer tool with full integration directly into vSphere Lifecycle Manager (vLCM) for single pane of glass management via Lenovo XClarity Integrator for VMware (LXCI) Hardware Support Manager (HSM). vLCM simplifies and speeds up software and hardware firmware updating with task automation via Best Recipes. These are released frequently throughout the year to ensure you are compliant and secure. This enables you to free up valuable personnel for more productive tasks.

## 4.4 Microsoft SQL Server 2022

Microsoft SQL Server 2022 is the latest version of the data platform and helps ensure uptime with fully managed disaster recovery by continuously replicating data to and from the cloud. SQL Server 2022 adopts an immutable ledger to help protect data from tampering to achieve database security and compliance goals with the least vulnerability over the past few years. 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 also helps manage and govern the entire data estate to overcome data silos and drive insights in near real time by breaking the wall between operational and analytical stores.

In this solution, we adopted SQL Server 2022 workloads on Lenovo ThinkAgile VX Series using all NVMe SSD drives with VMware vSAN Express Storage Architecture 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.

**Table 1 Hardware configuration**

PROPERTY	SPECIFICATION
Server model name	8 x Lenovo ThinkAgile VX7531 Node
CPU	2 x Intel(R) Xeon(R) Gold 6348 CPU @ 2.60GHz, 28 cores each
RAM	1024 GB
Network adapter	Mellanox Technologies MT2892 Family [ConnectX-6 Dx], 100 Gbits/s full duplex
Storage adapter	Intel Volume Management Device
Disks	6 x 6.4TB Micron 7450 NVMe SSD storage with max endurance on each host

## 5.2 Software resources

The following software and levels were used in the solution.

**Table 2 Software resources**

Software	Version	Purpose
VMware vSphere	VMware ESXi, 8.0.1, 21495797 VMware vCenter, 8.0.1, 21560480	VMware vSphere is the leading server virtualization software for containerized and business critical enterprise applications. VMware vCenter Server provides a centralized platform for managing VMware vSphere environments.
VMware vSAN Express Storage Architecture	Built in with ESXi	VMware vSAN is a software-defined, enterprise storage solution that supports HCI systems.
Microsoft SQL Server 2022	RTM build – 16.0.1000.6	Microsoft SQL Server database server platform.
Windows Server 2022	Datacenter edition	Windows Operating System.
MSTPCE Toolkit	1.12.0	Generate OLTP workload derived from TPC-E like benchmark.

## 5.3 Network configuration

vSAN Express Storage Architecture typically requires a 25GbE network as minimum. vSAN supports Remote Direct Memory Access (RDMA) since vSAN 7.0 Update 2 and later releases. vSAN over RDMA has lower CPU utilization and less I/O latency for database workloads.

In this solution, we used 100GbE network with RDMA and Jumbo Frame enabled for vSAN Express Storage Architecture to fully exploit the capabilities of high-performing NVMe devices used by Express Storage Architecture. To enable lossless traffic for vSAN RDMA network, you need careful design and configuration on the physical switch side. For example, the Data Center Bridging (DCB) mode must be configured as IEEE and the Priority Flow Control (PFC) value must be set to 3. [vSAN Health Check](#) gives an overall evaluation under the “RDMA configuration health” section. Make sure the physical NIC is working properly with the appropriate firmware version in the ESXi Server. For more details about physical switch and NIC configuration, contact the corresponding vendors for more details.

## 5.4 Storage Configuration

Storage policy with erasure coding is the new default policy for workloads running on vSAN Express Storage Architecture. In this solution, we used “vSAN Express Storage Architecture default policy—RAID-6” as the storage policy for SQL Server VMs deployed on vSAN Express Storage Architecture. The testbed is configured with 8 physical hosts and the Express Storage Architecture RAID-6 is the best storage policy that

provides higher resiliency with FTT=2 for database workloads, equivalent performance with no compromise as compared to RAID-1, and better space-efficiency (1.5x) for large database size.

As a best practice, we separated the SQL Server database into multiple files and virtual disks. In this solution, we provisioned 16 data disks and 1 log disks and spread them among different VMware Paravirtual SCSI adapters.

**Table 3 Virtual disk configuration**

Virtual Disk usage	Size	PVSCSI Controller
Operating System	1 x 100GB	0
Data disks (1-8)	8 x 1TB	1
Data disks (9-16)	8 x 1TB	2
Log disk	1 x 1TB	3
Tempdb disk	1 x 1TB	3

## 5.5 SQL Server HA considerations

This section covers SQL Server high availability options on Express Storage Architecture.

SQL Server supports two types of high availability and disaster-recovery technologies that are both fully supported to run natively on vSAN Express Storage Architecture: Always on Availability Group (AG) and Failover Cluster Instance (FCI). For both HA technologies, VMware recommend implementing with VMware vSphere High Availability (vSphere HA) and VMware vSphere Distributed Resource Scheduler™ (DRS) to fully automated VM placement based on a proper configuration of VM rules. For more details, refer to [Planning Highly Available and Mission Critical Microsoft SQL Server on Windows Deployments with VMware vSphere](#).

### 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 Express Storage Architecture, it further provides database-level HA protection in addition to the storage level replicas provides by vSAN. Since both level of data replicas is present on database and storage layer, storage planning is a key in this situation.

VMware recommended the Express Storage Architecture storage policy of Failures to Tolerate (FTT) equal or greater than 1 for virtual disks that the AG database resides on. This ensures dual protection level of mission-critical production database with AG and vSAN. vSAN Express Storage Architecture also enables in-place compression by default with minimal performance overhead whereas the capacity saving is evident for multiple AG replicas.

### Failover Cluster Instance

Same as vSAN Original Storage Architecture, vSAN Express Storage Architecture supports SCSI-3 persistent reservation on a shared virtual disk that required by SQL Server FCI with highly scalable, available, and high-

performance HCI solution. Shared disk on vSAN further eliminates the management difficulty of traditional physical Raw Device Mapping (pRDM). By taking advantage of vSAN Express Storage Architecture, shared disks that required by SQL Server FCI are easy to manage, fast to operate, and simple to migrate. Live vMotion is enabled inside the vSAN cluster without workload interference.

For more details of configuration and migration of shared disk on vSAN, refer to [Configuring a shared disk resource for Windows Server Failover Cluster \(WSFC\) and migrating SQL Server Failover Cluster Instance \(FCI\) from SAN \(RDMs\) to vSAN](#).

## 5.6 Best practices

In this solution, we validated the VMware OLTP workload derived by MSTPCE Toolkit and measured the performance scalability and capability of running SQL Server 2022 on vSAN Express Storage Architecture with Lenovo ThinkAgile VX Series.

The following recommendations provide the best practices and sizing guidance to run SQL Server on vSphere.

- Hardware consideration:
  - Follow [vSAN ESA ReadyNode Configurator](#) to make sure the hardware components meet the desired compatibility requirements.
  - Follow [ThinkAgile VX Series for VMware vSAN](#) for general hardware best practices.
- Compute consideration:
  - For general best practices, follow [Architecting Microsoft SQL Server on VMware vSphere](#).
  - Size the SQL Server VM based on the performance and capacity requirements of the desired workload. You may refer to numbers described in the scale-up result table to size your SQL Server VM in terms of the database customer count, IOPS numbers, and others for typical OLTP workloads. VMware also does not recommend over-committing physical CPU resources even if hyper-threading is enabled, since over-committing CPU resources do not guarantee a further scalable performance as expected.
  - For mission-critical workloads to achieve better performance, configure the Operating Mode for Maximum Performance in the BIOS settings.
  - Enable MWAIT in the system BIOS settings. For more details, see [VMware KB83076](#).
- Network consideration:
  - Use 25Gbps network as a minimum requirement for vSAN Express Storage Architecture. For better performance and throughput, consider 40Gbps/100Gbps network for mission-critical SQL Server workloads.
  - Use RDMA capable network to offload host CPU utilization and boost the SQL Server workloads as needed.
- Express Storage Architecture Storage consideration:

- Use erasure coding RAID-5/6 as the default storage policy for vSAN Express Storage Architecture as it eliminates the trade-off of performance and deterministic space efficiency. Select FTT=1 using RAID-5 and FTT=2 using RAID-6 depending on the number of hosts presented in the Express Storage Architecture cluster and your data availability requirement.
- Express Storage Architecture can support up to 32 maximum snapshots in a chain. As a best practice, always maintain a short snapshot chain for management and storage consideration. For heavy and latency sensitive workloads, take snapshot when the workload is idle or light-loaded to ensure minimum impact. Refer to [Best practices for using VMware snapshots in the vSphere environment](#) for more details.
- SQL Server consideration:
  - Reserve enough memory resource for SQL Server—60% to 80% available memory size is a recommended value.
  - 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. 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

### vSAN Performance Service

[vSAN Performance Service](#) is used to monitor the performance of the vSAN environment, using the vSphere web client. The performance service collects and analyzes performance statistics and displays the data in a graphical format. You can use the performance charts to manage your workload and determine the root cause of problems.

### vSAN Health Check

[vSAN Health Check](#) delivers a simplified troubleshooting and monitoring experience of all things related to vSAN. Through the vSphere web client, it offers multiple health checks specifically for vSAN including cluster, hardware compatibility, data, limits, physical disks. It is used to check the vSAN health before the mixed-workload environment deployment.

### 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 and Exchange performance and status.

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

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

## 6.2 Database workload generation tool

### MSTPCE Toolkit

The VMware OLTP workload tested in this solution is derived from Microsoft TPC-E toolkit that can conduct industry-standard benchmark testing and scalability testing. With the MSTPCE Toolkit, you can update your database environment, and mitigate the risks of planned database changes. We used this tool to generate VMware OLTP workloads to validate SQL Server on VMware vSAN Express Storage Architecture platform.

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

## 6.3 Scale-up OLTP workload test


The scale-up OLTP workload test is designed to help sizing SQL Server VMs running on vSAN Express Storage Architecture cluster with different customer count and database size. Table 4 listed the tested VM sizing with corresponding customer count for each of the test database. The database is generated by MSTPCE Toolkit, and a rough estimation is 10GB for each 1K customer count in terms of database size. The SQL Server memory size is configured at 75 percent of the available memory size of the VM.

**Table 4 Scale-up test configuration**

VM Sizing	Customer Count	CPU sizing	Memory (GB)	SQL Server Memory (GB)
Tiny	40K	4 vCPU	32	24
Small	80K	8 vCPU	64	48
Medium	160K	16 vCPU	128	96
Large	320K	32 vCPU	256	192
xLarge	480K	48 vCPU	384	288
Monster	640K	64 vCPU	512	384

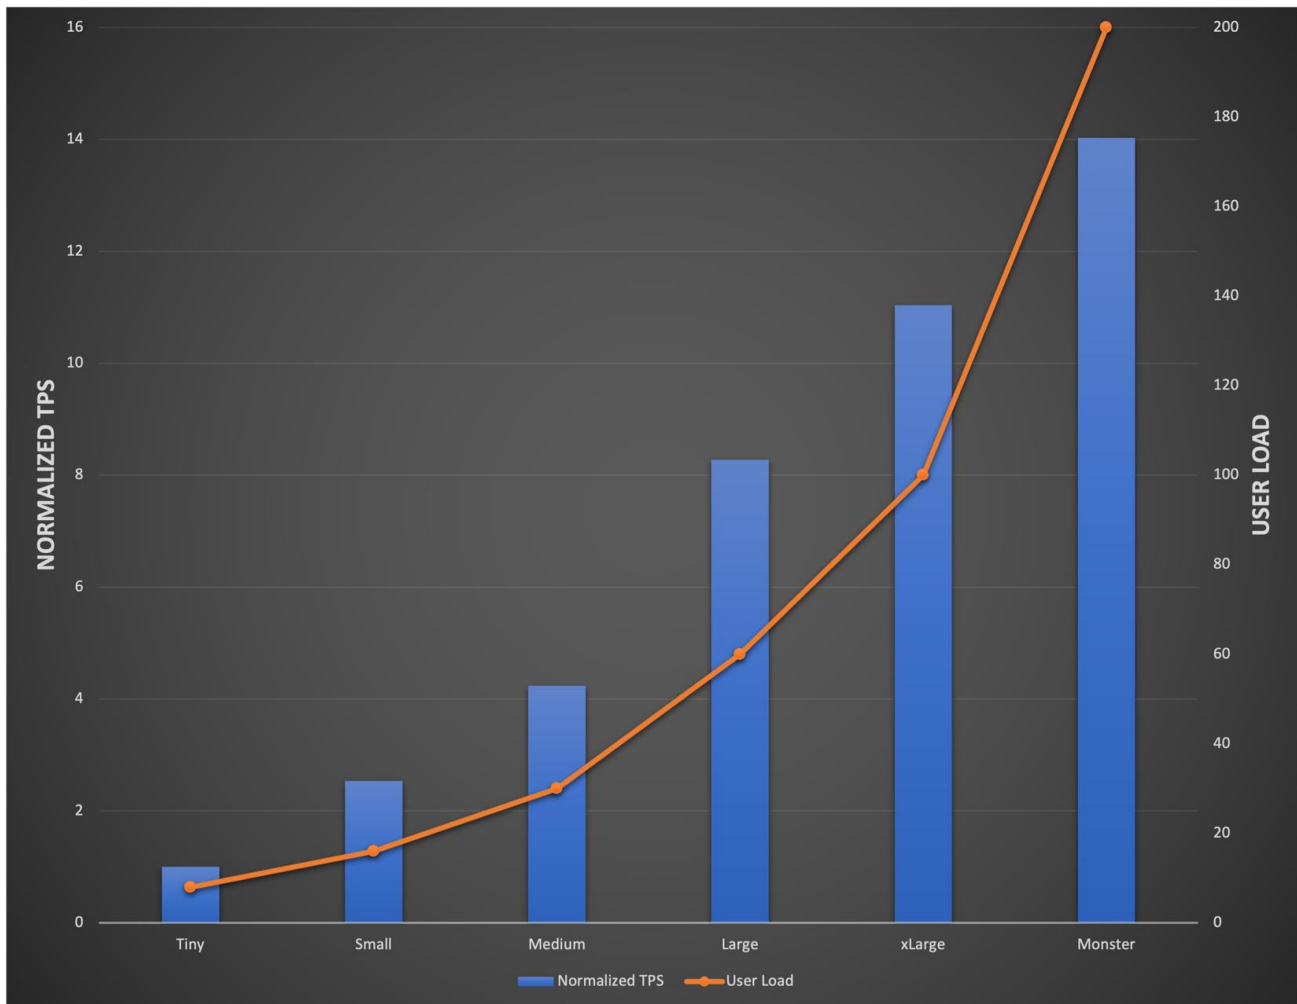
The test objective is to reach an average of 60 to 70 percent of VM CPU utilization measured by the performance counters collected by Windows performance monitor within each test run. The overall test duration was 2 hours, and we measured the performance after the OLTP workload entered a steady state for 1 hour (started at the 30 minutes of the test run until the 90 minutes).

Figure 4 shows the test results of total IOPS (read+write) and VM CPU utilization for each SQL Server VM size that reached 60~70 percent in the performance counter. With the monster VM (64 vCPU+512GB memory), we achieved 114K read IOPS and 15K write IOPS with a total IOPS value of 129K. All the data and log disks response time were around 1 millisecond with the help of single-tier all NVMe disks configured by vSAN Express Storage Architecture. With the T-shirt VM sizing of different customer count, it demonstrated a linear performance scalability for SQL Server OLTP workloads running on vSAN Express Storage Architecture.



**Figure 4 Scale-up test results - IOPS and CPU usage**

Figure 5 shows the normalized TPS result and the corresponding user load for each test run. We took the TPS result of “Tiny” VM configuration monitored in Windows Performance Monitor as baseline (TPS result shown as “1”). As we scaled up the VM configuration from “Tiny” (4vCPU and 32GB memory) to “Monster” (64vCPU and 512GB memory), the TPS result boosted up by 14 times, which showcased a well-designed and predicible database performance for SQL Server running on vSAN Express Storage Architecture.



**Figure 5 Scale-up test results - Normalized TPS and user load**

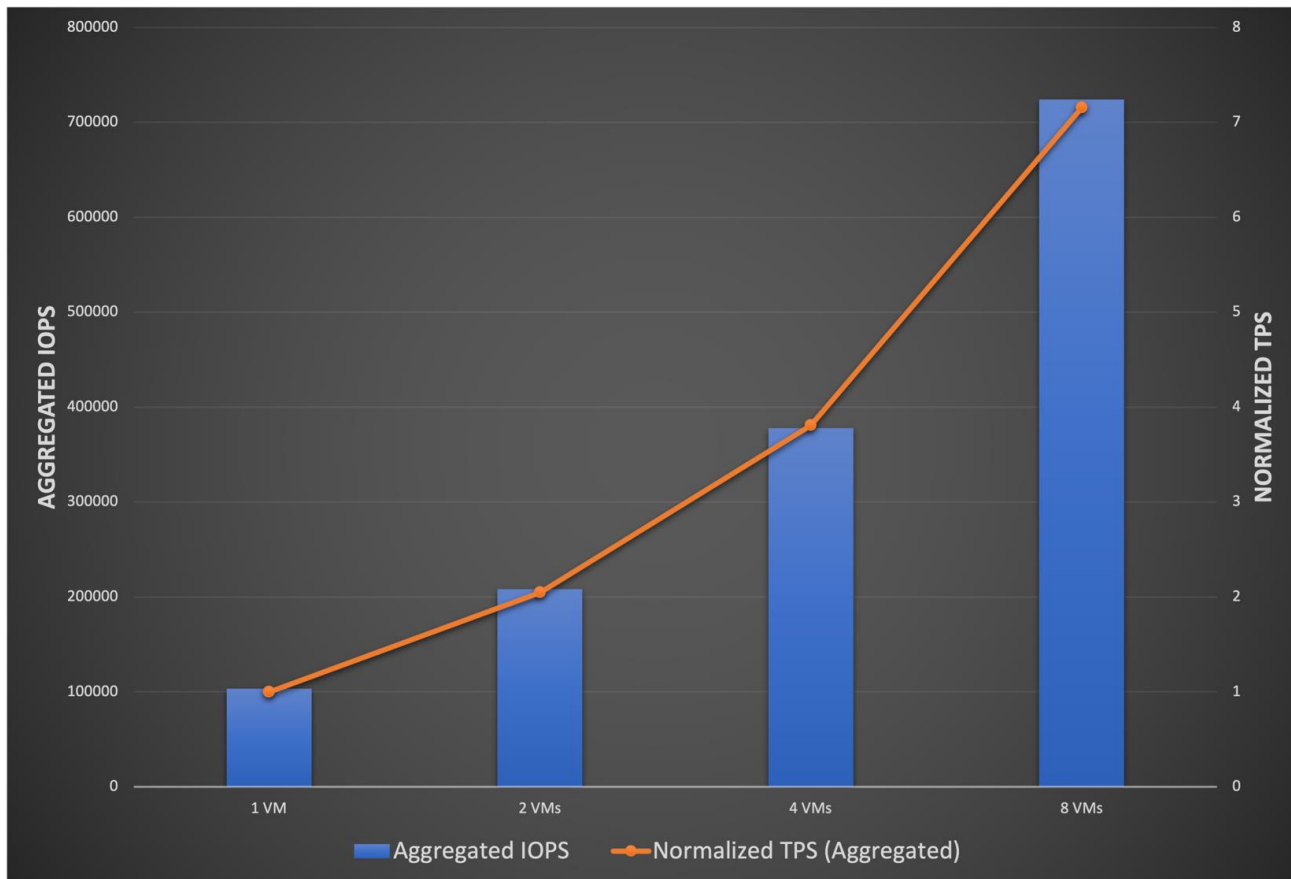
## 6.4 Scale-out OLTP workload test

The scale-out OLTP workload test is designed to demonstrate the performance capability and consistency of vSAN Express Storage Architecture cluster with multiple heavy mission-critical SQL Server OLTP workloads. We ran VMware OLTP workloads generated by MSTPCE Toolkit against up to 8 SQL Server VMs at the same time. The databases were generated by MSTPCE toolkit with 480K customer count with a rough size of 5TB. The SQL Server memory size is configured at 75 percent of the available memory size of the VM.

The test objective is to test the overall performance of the vSAN Express Storage Architecture cluster on the 8-node Lenovo ThinkAgile VX 7531 Series with SQL Server OLTP workloads. We started the test with single VM on the Express Storage Architecture cluster, and as we scale out the number of SQL Server VMs, we monitored the aggregated IOPS and normalized TPS results to evaluate the scale-out capability of the entire system.

Figure 6 shows the result of single VM, double VMs, 4 VMs and 8 VMs scale-out test results. We achieved an aggregated read and write IOPS over 720K for 8 SQL Server VMs running the VMware OLTP workloads derived from MSTPCE toolkit. The TPS result shows a linear scalability in terms of the number of SQL Server

VMs—with 7.x times TPS improvement of 8 VMs compared to single VM.



**Figure 6 Scale-out test results - Aggregated IOPS and Normalized TPS**

## 6.5 SQL Server data protection and snapshot test

With the growing demand for data protection solutions of mission-critical database workloads, vSAN Express Storage Architecture provides high performance and efficient snapshot management for SQL Server.

Leveraging the Express Storage Architecture snapshot technology, the data protection software can back up and restore SQL Server VM residing on Express Storage Architecture using the vSphere Storage APIs (known as VADP) faster and more efficient. Express Storage Architecture provides higher bandwidth and flexible snapshots for data protection workflows. In this solution, we tested the snapshot performance impact of running VMware OLTP workload.


### vSAN Express Storage Architecture Snapshot Test

vSAN Express Storage Architecture enhanced the snapshot feature by using a highly efficient lookup table with a B-Tree to provide scalable and high-performance native snapshots. Snapshot creation and deletion are extremely fast and performance impact is minimized.

Figure 7 shows the vSAN Express Storage Architecture snapshot test result by running the VMware OLTP workload. The picture was captured in Windows Performance monitor with the counter “Disk Transfer/sec\_Total”, which showcased the IOPS number of the SQL Server test VM.

We generated the OLTP workload as shown in previous test with 480K customer size database. The snapshots were created every 15 minutes for the SQL Server VM, and we monitored the performance impact. The first snapshot was created after the VMware OLTP workload entered a steady state. There was a slight performance dip right after the snapshot took place with a short stun time of the virtual disks. The performance got back right after the first snapshot happened and we could see the line remained flat afterwards. Same observation was achieved for the following snapshots, and we did 8 snapshots in a row with a 15-minute interval. From the chart we could conclude there was zero performance impact of Express Storage Architecture snapshot on the VMware OLTP workload on the SQL Server VM.

Express Storage Architecture also improves snapshots deletion and consolidation of snapshot chain. While the snapshot creation can be achieved just on a simple click, the snapshot deletion of the virtual disks of the SQL Server test database was completed quite fast within several seconds, no matter how long the snapshot chain was.



**Figure 7 Express Storage Architecture snapshot test results – Windows Performance Monitor**

## 7 Conclusion

---

Running SQL Server on VMware vSAN Express Storage Architecture on Lenovo ThinkAgile VX Series is a fast way to unlock capabilities, improve operation efficiency, and achieve high performance with industry-leading HCI on the next-generation hardware platform for mission-critical database workloads.

In this solution, we demonstrated the performance scalability for SQL Server VMs running on vSAN Express Storage Architecture from both vertical direction with scale-up capability and horizontal direction with scale-out capability. The results demonstrate a linear scalability with predictable and consistent performance for SQL Server workloads. vSAN Express Storage Architecture also provides rich data services and higher data availability for business-critical workloads with various storage policy options such as erasure coding RAID-5/6 and data compression technologies. The Express Storage Architecture snapshot technology has minimal performance impact and simplified snapshot management. Snapshot creation, deletion, and consolidation are much faster and database backup software can take advantage of the Express Storage Architecture snapshot for efficient and effective data protection of SQL Server workloads.

## 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.

Part number	Product Description	Qty
7Z63CTO5 WW	Server: Lenovo ThinkAgile VX7531 Certified Node	1
BJLK	ThinkAgile VX 2U 2.5" Chassis with 8, 16 or 24 Bays	1
BB2L	Intel Xeon Gold 6348 28C 235W 2.6GHz Processor	2
B964	ThinkSystem 32GB TruDDR4 3200 MHz (2Rx4 1.2V) RDIMM	32
B8P1	ThinkSystem 440-16i SAS/SATA PCIe Gen4 12Gb Internal HBA	1
BT2G	vSAN ESA	1
BYRN	AF-2	1
BNEL	ThinkSystem 2.5" U.3 7450 MAX 6.4TB Mixed Use NVMe PCIe 4.0 x4 HS SSD	6
B5XH	ThinkSystem M.2 SATA 2-Bay RAID Enablement Kit	1
BQ1Y	ThinkSystem M.2 5400 PRO 480GB Read Intensive SATA 6Gb NHS SSD	2
BQ8S	VMware ESXi 8.0 U1 (Factory Installed)	1
B93E	ThinkSystem Intel I350 1GbE RJ45 4-port OCP Ethernet Adapter	1
B4RA	ThinkSystem Mellanox ConnectX-6 HDR100/100GbE QSFP56 2-port PCIe VPI Adapter	1
BQ0W	ThinkSystem V2 1100W (230Vac/115Vac) Platinum Hot Swap Power Supply	2
BH8E	ThinkSystem 2U Performance Fan	6



# Resources

---

- [VMware vSphere](#)
- [VMware vSAN](#)
- [VMware Express Storage Architecture](#)
- [VMware Express Storage Architecture Frequently Asked Questions](#)
- [Microsoft SQL Server 2022](#)
- [Architecting Microsoft SQL Server on VMware vSphere](#)
- [Lenovo ThinkAgile VX Series for VMware vSAN](#)

# Document history

---

Version 1.0	November 2023	First version of document, based on original VMware RA document
-------------	---------------	---

# Trademarks and special notices

---

© Copyright Lenovo 2023.

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

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

Lenovo®

ThinkAgile®

ThinkSystem®

TruDDR4

XClarity®

The following terms are trademarks of other companies:

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

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

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

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