

Scalable VMware vSAN Storage Architectures on Lenovo ThinkAgile VX

Solution Brief

Business Trends

Software Defined Datacenter Architecture: The growth of technology, modern applications, cloud, artificial intelligence and data-driven ecosystems bring the need for on-premises infrastructure to meet high performance and high-density workload requirements across datacenters. Servers need to have software defined solutions to build robust storage systems and integrated software stack for cloud and cloud native technologies and third-party software to build flexible and scalable architecture to meet any workload. This trend means businesses need faster servers and storage to achieve the performance and low latency required.

Infrastructure and Workloads Modernization: Adoption of hybrid cloud, artificial intelligence and machine learning and open source technologies drives enablement for tens of software and applications to coexist on the shared infrastructure and need increased processing and data storage capabilities. Ever growing data and AI/ML workloads require storage system with low latency and high capacity drives, high speed network adapters and interconnects that are pre-tested and pre-validated engineered solution to reduce deployment complexities. Business needs for consolidation enterprise workloads and rapid provision end-end infrastructure and software stack without compromising performance and integration capabilities.

Lenovo ThinkAgile VX Solutions for VMware vSAN

Lenovo ThinkAgile VX systems are the perfect choice for hyperconverged infrastructure and provide an outstanding platform to support the different VMware vSAN™ architectures. Lenovo and VMware's over 20 year partnership and collaboration continues to strongly drive innovation and technical enablement for vSAN-based storage solutions. This includes validation, certification, configuration and support for ThinkAgile VX systems.

Lenovo ThinkAgile VX V3 hyperconverged systems are equipped with 4th generation Intel® Xeon® Scalable processors and VMware vSAN 8. They are **Accelerated by Intel** offerings that drive greater performance through a number of enhancements, including higher cores, embedded accelerators, GPU, DPU, DDR5 and PCIe Gen 5 components.

Lenovo ThinkAgile VX servers are available as Integrated Systems and Certified Nodes. Both are factory integrated, pre-configured systems with Lenovo hardware, VMware software, and deployment services. VX Integrated Systems provide a quick and convenient path to implement a hyperconverged solution powered by VMware vSAN and a single point of contact provided by Lenovo for purchasing, deploying, and supporting the solution. VX Certified Nodes come with optional VMware software and services.

ThinkAgile VX Integrated Systems can also be up and running quickly with a web-based deployment wizard. The installer can install and configure VMware ESXi, vCenter Server and Lenovo XClarity Integrator and either create or join a cluster.

Lenovo ThinkAgile VX systems support all vSAN architectures:

Table 1. VMware vSAN Storage Architectures

Original Storage Architecture (OSA)	<ul style="list-style-type: none"> • Tiered architecture with flash cache tier and capacity tier • Support NVMe, SSD and HDD drives • Minimum 2 drives and 10GbE • Disk group uses one cache and multiple capacity drives
Express Storage Architecture (ESA)	<ul style="list-style-type: none"> • Single tier storage pool architecture with all NVMe drives • Minimum 4 drives and 25GbE • Minimum 3 nodes
vSAN Max Architecture	<ul style="list-style-type: none"> • Disaggregated storage based on Express Storage Architecture (ESA) • Compute nodes are not part of the cluster • Minimum 6 nodes and 100GbE is recommended • Minimum 75 TB per node • Ready Node configurations available

Comparison of vSAN OSA and ESA Solutions

The table below shows a basic comparison of vSAN OSA and ESA solutions.

Table 2. Comparison of VMware vSAN OSA and ESA

Disk Type Support	Cache tier - NVMe, SSD Capacity tier - NVMe, SSD, HDD	All NVMe Capacity drives (Minimum 4 drives per host)
Network (Minimum)	10 GbE	25 GbE
Compression Encryption	Datastore level	Policy driven virtual machine level
Deduplication	Supported	Not supported
Data transfer over network	Not compressed	Compressed
vSAN File Services	Supported	Supported (vSphere 8.0 U2)
Default Storage Policy	RAID1	RAID5
Maximum VMs per host	200	500 (vSphere ESXi 8.0 U2)

ThinkAgile VX and vSAN Storage Architectures



ThinkAgile VX650 V3



ThinkAgile VX630 V3

Applications where the servers would excel include:

- Mission Critical applications
- SAP HANA and ERP applications
- CRM
- Business Intelligence (BI)
- Databases and Data Warehouse
- Analytics
- Virtual Desktops
- Artificial Intelligence
- Server Consolidation
- Virtualization

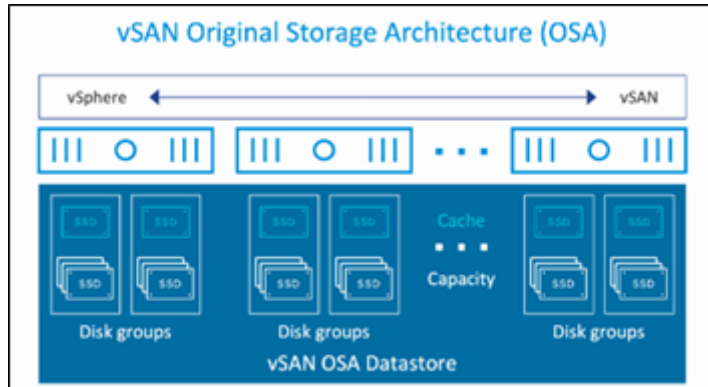


Figure 1. vSAN Original Storage Architecture (OSA)

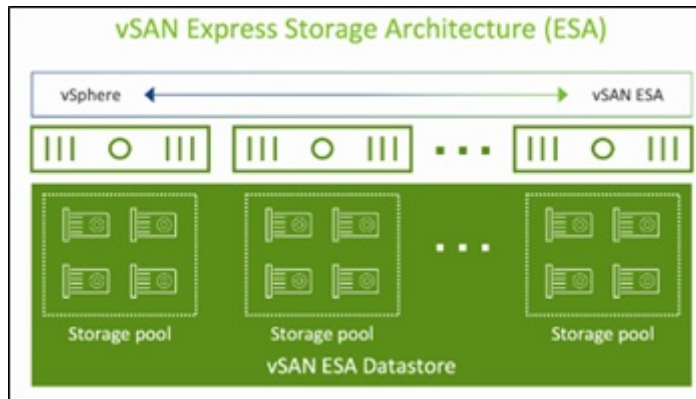


Figure 2. vSAN Express Storage Architecture (ESA)

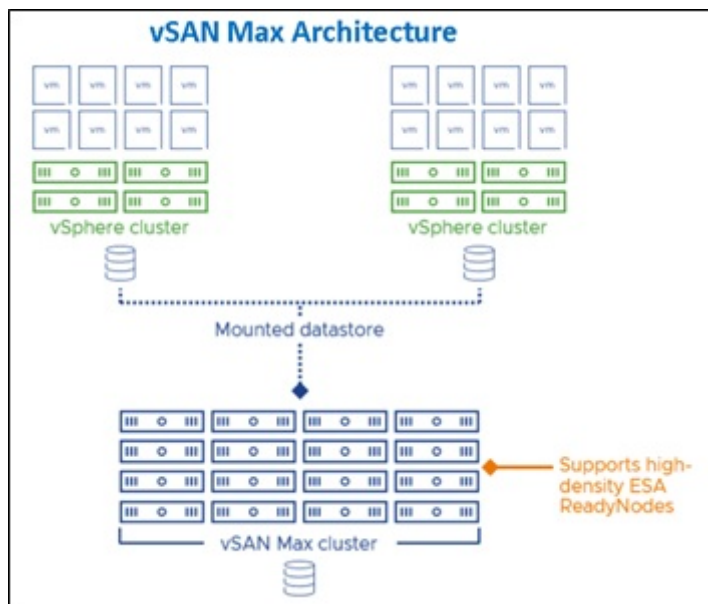


Figure 3. vSAN Max Architecture

Use Case: VMware Horizon VDI on vSAN ESA vs. OSA

The Login Enterprise VDI benchmark was used to measure performance for 1400 knowledge users on a 4-node ThinkAgile VX650 V3 integrated system cluster with 2x Intel Xeon Platinum 8468 48C processors for both vSAN OSA and vSAN ESA. The performance is similar on both architectures and the login performance looks better on OSA which could be due to the cache tier. The application performance in ESA during steady state looks better than OSA.

The ESA testing was done with 8x ThinkSystem P5620 3.2TB Mixed Use NVMe PCIe 4.0 drives per node. The figure below shows ESA login performance with response rate increases above 50 seconds after 800 users.




Figure 4. Login Enterprise Test Results with vSAN ESA

The OSA testing was done with ThinkSystem P5620 1.6TB Mixed Use NVMe PCIe 4.0 drives for the cache tier and ThinkSystem S4520 1.92TB Read Intensive SATA 6Gb HS SSD drives for the capacity tier. The figure below shows OSA login performance with response rate increases above 50 seconds after 1200 users. The OSA cache architecture can provide benefit for workloads which can leverage cache.

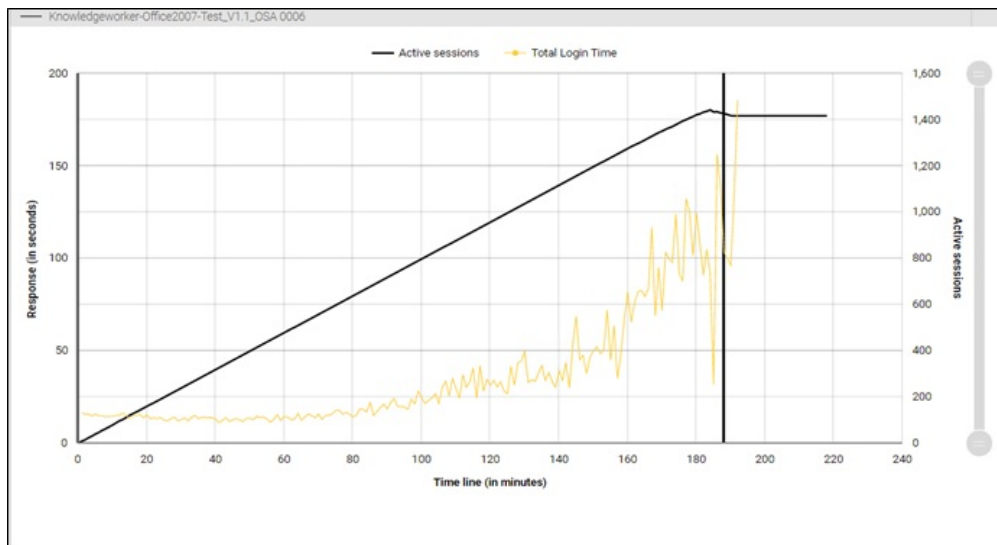


Figure 5. Login Enterprise Test Results with vSAN OSA

HCIBench FIO Benchmark on vSAN OSA and ESA

The HCIBench tool was used to measure performance for ESA and OSA in the different profiles listed in the table below. A total of 16 virtual machines are configured with 4x vCPU, 8GB memory and 8x 50GB data disks with data distributed evenly across 4 nodes. The disk configurations are the same as used for the VDI benchmark and all the testing was performed with 4th Gen Intel Xeon Scalable Processors.

Scenario	Block size	vmdk disks per vm	Threads	Read %	Write %	Random %	Application
4K-Read 100%	4K	8	32	100%	0%	100%	Generic Workload
8K-Read 70%	8K	8	16	70%	30%	100%	Database
16K-Read 70%	16K	8	8	70%	30%	70%	Data Warehouse
32K-Read 50%	32K	8	4	50%	50%	50%	Video/Image/3D
128K-Read 50%	128K	8	1	0%	100%	0%	Video/Image/3D

Throughput on 25Gbe NICs

In this testing, we compared throughput of OSA RAID1 vs. OSA RAID5 vs. ESA RAID5 across the different scenarios with 2x Mellanox ConnectX-6 Lx 10/25GbE SFP28 2-port PCIe Ethernet Adapters. The vSAN OSA RAID1 testing with 3 disk groups performed better than ESA for scenarios with different read and write combination. Also, vSAN ESA consumes ~5-10% more CPU resources than vSAN OSA. The vSAN OSA RAID5 performance is lower than ESA RAID5 and it is due to log based file system architecture specifically addressed for storage efficiency in ESA. ESA RAID5 is default and recommended option and it provided better throughput than ESA RAID1.

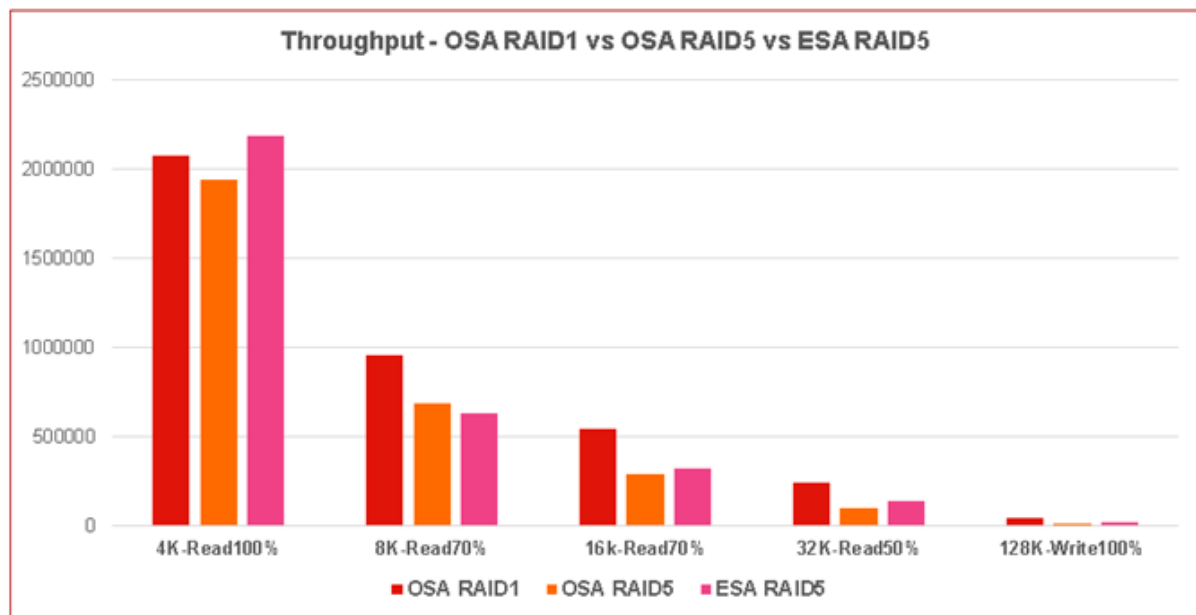


Figure 6. Throughput - OSA RAID1 vs. OSA RAID5 vs. ESA RAID5

Throughput on 100GbE NICs with 32 and 48 Core 4th Gen Intel Xeon Processors

The figure below shows the throughput for both OSA RAID1 and ESA RAID5 on 2x Intel Xeon Gold 6448Y CPU 32C 2.1 GHz and 2x Intel Xeon Platinum 8468 48C 2.1 GHz with 100 GbE bandwidth.

- The 100 GbE throughput outperforms 25GbE results for OSA and ESA
- Both OSA and ESA performance looks comparable on 100 GbE. The OSA requires more disks to get more throughput by increasing number of disk groups.
- An ESA configuration with today's software can scale up to and even saturate 32 cores. While the performance of vSAN itself is better on medium core count (MCC) CPUs, e.g., the 32 core Intel Xeon Gold 6448Y processor, the additional cores available on extreme core count (XCC) CPUs, e.g., the 48 core Intel Xeon Platinum 8486 processor, remain available to support other compute-intensive workloads, e.g., virtual desktop or AI use cases, potentially providing a superior overall solution.

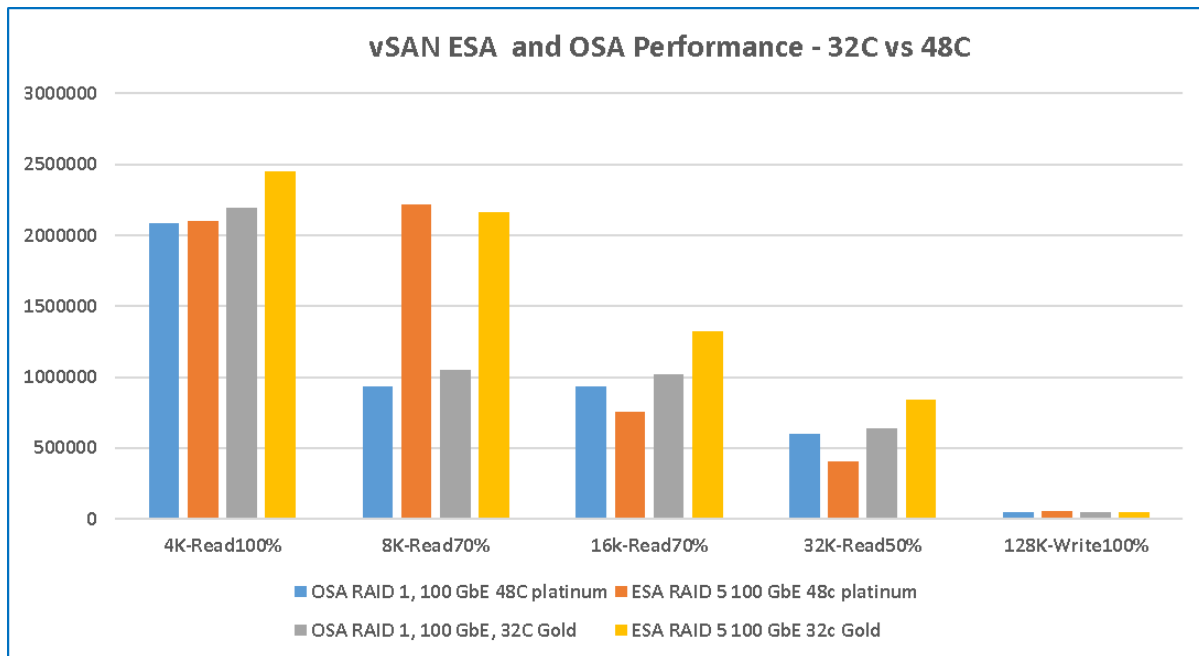


Figure 7. vSAN ESA and OSA Performance - 32C vs. 48C

vSAN OSA Scale Up Performance

The figure below shows throughput results for different numbers of disk groups for vSAN OSA. The throughput increases with the greater number of disk groups and the write intensive scenarios do not show considerable benefits. Adding more disk groups requires more cache drives.




Figure 8. OSA Disk Group Performance - RAID1

Design Guidelines for OSA and ESA

Both vSAN OSA and ESA are scalable solutions and ESA reduces storage footprint up to 40% and also it can accommodate 60% more virtual machines. vSAN OSA supports vSphere 7 & 8 and provides support for a wide number of options with NVMe, SSD and HDD and the cache tier provides better performance for many workloads. The vSAN cluster configuration should be balanced for CPU, memory and storage requirements based on the workloads running on them and the cluster must be either scaled up or scaled out appropriately. The table below provides different aspects to be considered during initial assessment when designing different vSAN solutions.

Feature	OSA	ESA
vSphere Support	ESXi 7.x, 8.x	ESXi 8.x
Drive Options	More drives required	Less drives required
Maximum Capacity per Node	35 drives (5 disk groups, 7 capacity drive per disk group)	24 drives
Scale Up	<ul style="list-style-type: none"> • Increase Disk Groups (Cache and Capacity) • Increase Capacity disks and ensure enough cache capacity 	Capacity can be increased without any dependencies
Scale Out	<ul style="list-style-type: none"> • Add nodes and disks together • If total disk capacity reached 70% threshold and if no more drives can be added, then add more hosts 	Add nodes
Changing Storage Policy	Requires reformatting disks	Applied at storage policy level
Disk Maintenance	<ul style="list-style-type: none"> • Removing Cache drive requires removal of disk group and creation of new disk group. The larger the disk group the more time is required to evacuate data and rebalance after maintenance. • Adding/removing capacity drives are easy and data migration is at the disk level • Need a free host with enough drives to migrate data from the failed node 	Drives can be added/removed easily, and data migration is at the disk level
Compression	At vSAN datastore level	At virtual machine level
Optimal Performance	RAID1	RAID5 is comparable to RAID1 OSA

Accelerated by Intel

To deliver the best experience possible, Lenovo and Intel have optimized this solution to leverage Intel capabilities like processor accelerators not available in other systems. Accelerated by Intel means enhanced performance to help you achieve new innovations and insight that can give your company an edge.



Lenovo and VMware

With co-located engineering organizations and a history of technical collaboration, VMware and Lenovo consistently deliver innovative joint solutions for the data center. Lenovo's leadership in reliability, customer satisfaction, and performance, combined with VMware's leadership in software and cloud services, continues to deliver innovative data center solutions and lower TCO for our joint customers.

Why Lenovo

Lenovo is a US\$70 billion revenue Fortune Global 500 company serving customers in 180 markets around the world. Focused on a bold vision to deliver smarter technology for all, we are developing world-changing technologies that power (through devices and infrastructure) and empower (through solutions, services and software) millions of customers every day.

For More Information

To learn more about this Lenovo solution contact your Lenovo Business Partner or visit:

<https://www.lenovo.com/us/en/servers-storage/sdi/thinkagile-vx-series/>

References:

Lenovo ThinkAgile VX630 V3 Integrated System and Certified Node: <https://lenovopress.lenovo.com/lp1672>

Lenovo ThinkAgile VX650 V3 Integrated System and Certified Node: <https://lenovopress.lenovo.com/lp1673>

VMware vSAN Design Guide: <https://core.vmware.com/resource/vmware-vsan-design-guide>

VMware ESA Ready Node Hardware Guidance:

https://www.vmware.com/resources/compatibility/vsanesa_profile.php

Related product families

Product families related to this document are the following:

- [ThinkAgile VX Series for VMware](#)
- [ThinkSystem SR630 V3 Server](#)
- [ThinkSystem SR650 V3 Server](#)
- [VMware Alliance](#)

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.
8001 Development Drive
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.

© Copyright Lenovo 2024. All rights reserved.

This document, LP1872, was created or updated on December 15, 2023.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP1872>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP1872>.

Trademarks

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®

XClarity®

The following terms are trademarks of other companies:

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

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