

Maximize Security and Workload Performance with Lenovo DPU-enabled Servers

Solution Brief

Tackling Performance, Latency and Security Issues in Hybrid Cloud Deployment

The latest server technology and architecture enables organizations to achieve a higher consolidation ratio for virtual machines and containers for different workloads. The dramatic increase of hybrid cloud adoption has led to an increase in east-west and north-south traffic which brings added complexity to achieve zero-trust security. Introducing an additional layer of security solutions brings challenges in terms of architecture, deployment and management - and it requires more compute power. Though servers have high-speed processors, memory and storage, the networking and security processing consumes considerable hardware resources, and it impacts achieving more density per server without compromising performance and latency. Businesses need to offload and accelerate workloads without introducing additional complexities for deployment and operations of their infrastructure.

Lenovo Solution

Select Lenovo ThinkAgile hyperconverged systems and ThinkSystem servers support NVIDIA BlueField-2 DPU (Data Processing Unit - also known as SmartNIC) adapters and VMware vSphere Distributed Services Engine to offload management and network services from the CPU. Lenovo ThinkAgile VX650 V3 DPU and ThinkSystem SR650 V3 DPU servers based on 4th Generation Intel® Xeon® Scalable processors are integrated with DPUs. With VMware vSphere 8.0u1 and NVIDIA BlueField-2 DPU adapters, these systems address performance, security, and latency challenges by offloading VMware NSX services to DPU. The DPU has up to 8 ARM cores and runs VMware vSphere Distributed Services Engine which isolates and offloads management and network services from workloads. These systems are validated, factory built and tested by Lenovo, VMware and NVIDIA engineering teams to simplify deployment and achieve efficiency through maximized workload performance.

Highlights

- Achieve higher throughput at lower latency and reduced CPU overhead
- Offload networking and NSX services to vSphere Distributed Services Engine
- Simplify cloud deployment with VCF and higher consolidation ratio
- Achieve zero trust security by hardening at the DPU layer instead of the application layer

Solution Components

These Lenovo solutions are co-engineered with VMware and NVIDIA to ease deployment and the management lifecycle. Lenovo DPU servers provide scalability, improved performance, and higher consolidation for different workloads. Applications and situations where the servers would excel include:

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



ThinkAgile VX650 V3
ThinkSystem SR650 V3

NVIDIA BlueField-2 DPU

NVIDIA BlueField-2 data processing unit have more processing power, memory and high speed interconnects to support offloading workloads from high density Lenovo servers equipped with the latest Intel Xeon Scalable processors. The DPU is fully programmable and supports a wide range of storage and networking features.

- 8 ARMv8 A72 cores (64-bit), 1MB L2 cache per 2 cores, 6MB L3 cache
- 32GB of on-board DDR4
- 1GbE out-of-band management port
- PCIe Gen 4.0
- 2x 25GbE SFP56 ports
- Hardware acceleration for Regular Expressions and encryption secure boot



VMware vSphere Distributed Services Engine

vSphere Distributed Services Engine is a new way to architect the core infrastructure by orchestrating and managing infrastructure services with Data Processing Units (DPUs). The following advantages are achieved with DSE enabled Lenovo servers:

- **Peak performance:** by offloading network processing to the SmartNIC, we can improve network bandwidth and reduce latency and free up core CPU cycles for top application performance.
- **Unified, consistent operations:** Consistent operations across all apps. This includes dramatically simplified lifecycle management across VMware Cloud Foundation (VCF) deployments. All of which is designed to dramatically reduce OpEx.
- **Zero-trust security model:** by offloading network security functions to the SmartNIC, we can provide comprehensive application security capabilities without compromising application performance

Lenovo ThinkAgile and ThinkSystem Servers with DPUs

Lenovo supports Data Processing Units on ThinkSystem SR650 V3 servers and on vSAN-based ThinkAgile VX650 V3 hyperconverged systems.

Table 1. Lenovo ThinkSystem and ThinkAgile Models with DPU Support

Model	SR650 V3 DPU	VX650 V3 DPU	VX650 V3 SAP DPU*
Part Numbers	7D75CTO4WW	7D6WCTO4WW 7D6WCTO5WW	7D6WCTO6WW
CPU	Intel Xeon 4th Gen SP	Intel Xeon 4th Gen SP	Intel Xeon 4th Gen SP
Form Factor	2U 2S	2U 2S	2U 2S
Memory	TruDDR5 32 DIMMs (8 TB Max)	TruDDR5 32 DIMMs (8 TB Max)	TruDDR5 32 DIMMs (8 TB Max)
GPU	8xSW 75W 6xSW 150W 3xDW 300W		
DPU	1x NVIDIA Bluefield-2 DPU with dual ports of 25GbE	1x NVIDIA Bluefield-2 DPU with dual ports of 25GbE	1x NVIDIA Bluefield-2 DPU with dual ports of 25GbE
Drives NVMe/SAS/SATA	32x2.5" 20x3.5"	32x2.5" 20x3.5"	32x2.5" 20x3.5"
PCIe 5.0	3 Slots 1/10/25/ 100 GbE	3 Slots 1/10/25/ 100 GbE	3 Slots 1/10/25/ 100/200 GbE

* The SAP systems are being certified

Lenovo XClarity Support for NVIDIA BlueField-2 DPU

Lenovo XClarity supports firmware upgrades for NVIDIA BlueField-2 adapters. XClarity support deployment of an OS image to the ARM system and configuration of the UEFI settings and security.

- DPU firmware and driver updates will be available in best recipe release cycles. Updates can be done through vLCM (VMware vSphere Lifecycle Management) also.
- ThinkAgile VX V3 BIOS settings are applied in factory, and are available in the [ThinkAgile VX Best Recipe](#)
- The DPU ESXi DCUI/shell can be accessed through out of band serial console for troubleshooting

The following parts are available for Lenovo ThinkAgile and ThinkSystem DPU systems:

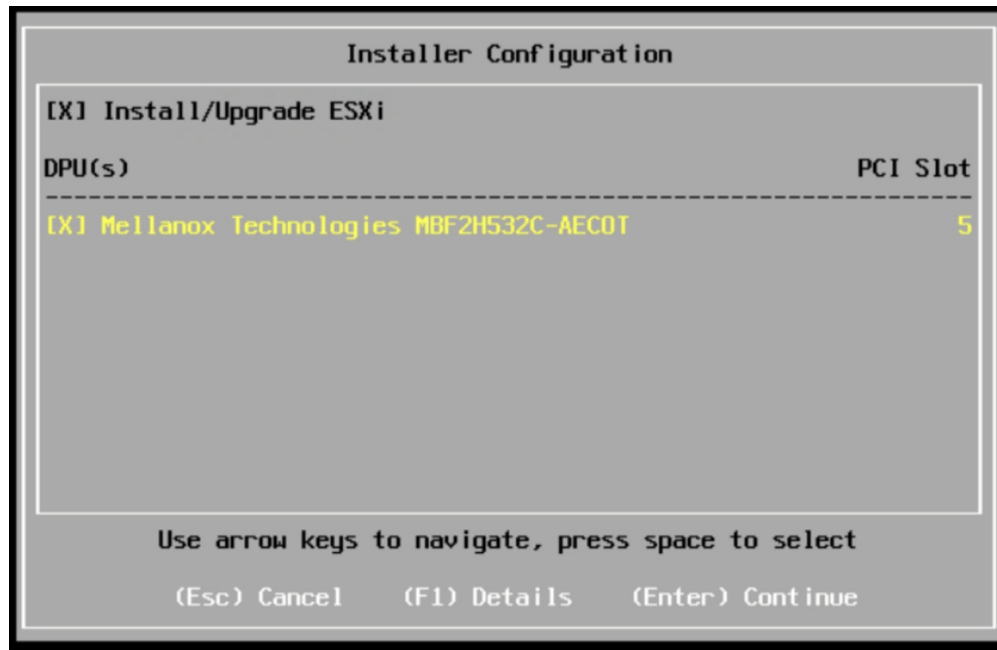
Table 2. Parts Available for ThinkSystem and ThinkAgile DPU Systems

Feature Code	Description
BRPK	ThinkSystem NVIDIA BlueField-2 25GbE SFP56 2-Port PCIe Ethernet DPU w/BMC & Crypto
BSH7	ThinkSystem OCP Enablement Kit for Distributed Services Engine

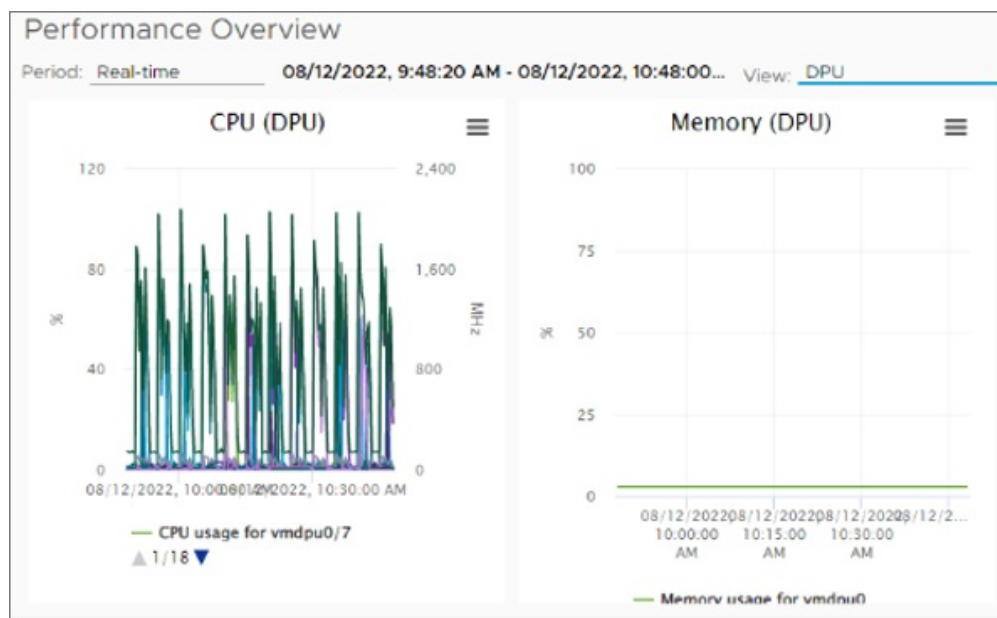
vSphere Distributed Services Engine (DSE)

The vSphere Distributed Services Engine is a feature in vSphere 8.0 and above, making it simple to include without adding new installation procedures. By integrating tightly with DPUs, vSphere Distributed Services Engine creates an alternate control fabric comprised of DPUs. The integration between ESXi on CPU and ESXi on DPU, NSX and vLCM are made easy for users to install and configure by leveraging existing tools and user experiences to preserve the existing Day 0, Day 1 and Day 2 operations.

ESXi Unified Installation

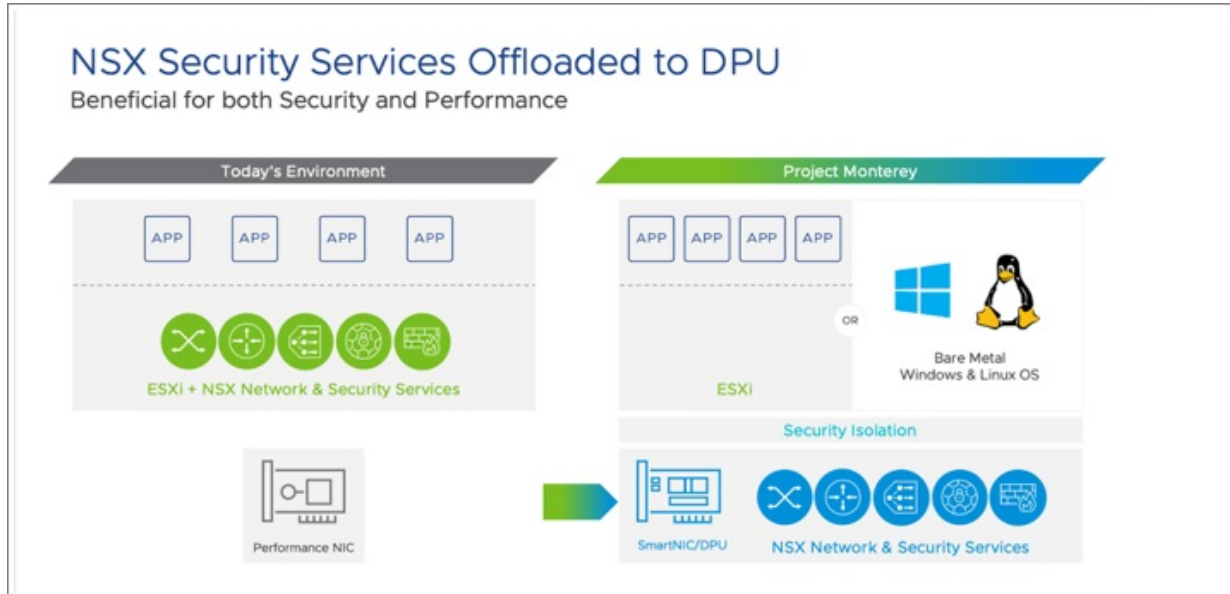


VMware vCenter Integration

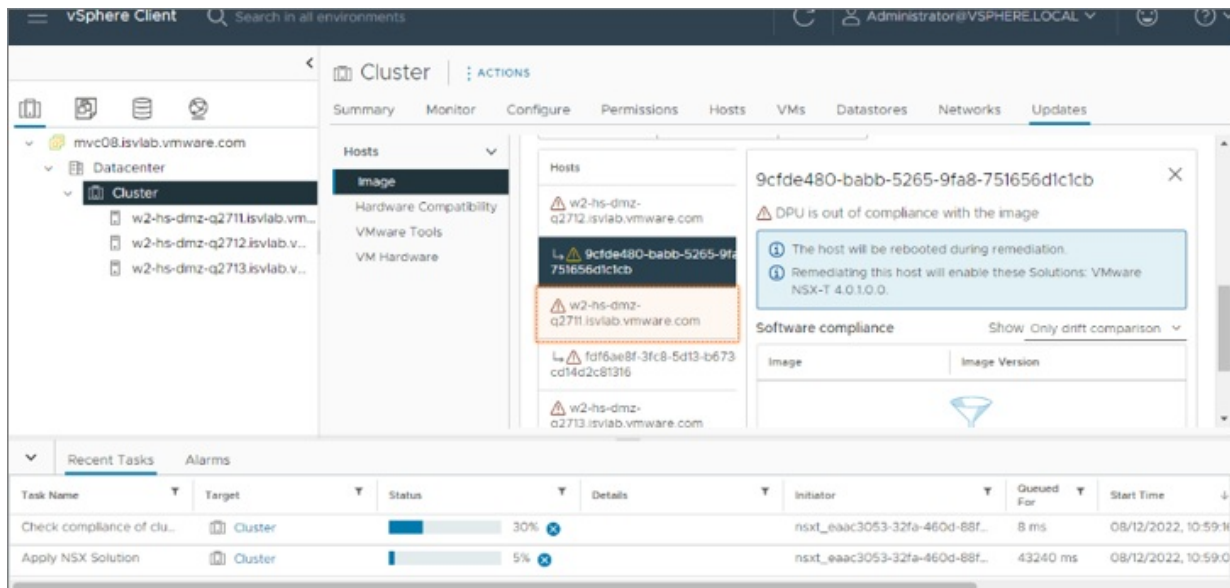


Unified NSX Services Deployment on DPU

NSX edge services processing requirements vary from 2-16 vCPU depending on the edge size. Refer to NSX component sizing in the [VMware Cloud Foundation Deployment Guide for ThinkAgile VX](#). With DPU integrated systems, the NSX edge services can be deployed on the DPU and it frees up CPU cycles to host more virtual machines and containers.

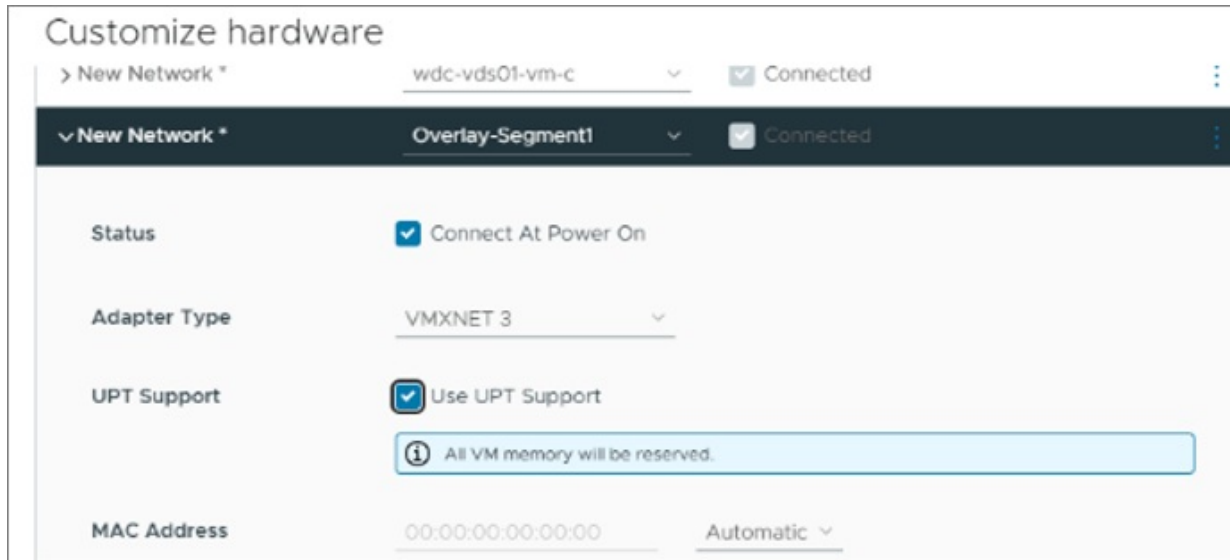


VMware Cloud Foundation supports installing NSX services directly on the DPU and this feature is available on SDDC Manager and can be finished in few clicks. Lenovo systems support NVIDIA BlueField-2 DPU adapters with 8 ARM cores. The adapters can provide enough capacity and performance to deploy many NSX services on cluster of nodes. When NSX services are running on DPU, it provides more isolation and security before the traffic reaches application workloads running on the CPU.




Universal Passthrough Version 2 (UPT v2)

vSphere Distributed Services engine is designed with an enhanced data path to access the NIC driver directly instead of using vSwitc. The new design adopts the best of SRIOV passthrough and Enhanced Data Path implementation on VMXNET3 to provide better performance without losing features such as DRS (disaster recovery services), HA (high availability) and vMotion. Virtual machines can be deployed by enabling UPT which is available only for hosts with DPU configured.



This performance test from VMware shows the Redis cache transactions were consuming 12 CPU cores for networking alone. When offloaded while using DPU with UPT v2, it freed up 12 cores for application workloads. Refer to this [VMware blog](#) for more information.



Bill of Materials

Table 3. Bill of Materials

Part number Feature code	Product Description	Qty
7D6WCTO4WW	Server : Lenovo ThinkAgile VX650 V3-DPU Integrated System	1
BRY9	ThinkAgile VX V3 2U 24x2.5" Chassis	1
B0W3	XClarity Pro	1
7S0XCTO5WW	XClarity Controller	1
BUE2	VMware vSAN 8 Enterprise	1
BQ6A	Intel Xeon Gold 6448H 32C 250W 2.4GHz Processor	2
BNF9	ThinkSystem 64GB TruDDR5 4800MHz (2Rx4) 10x4 RDIMM	16
B8P1	ThinkSystem 440-16i SAS/SATA PCIe Gen4 12Gb Internal HBA	1
B5MC	vSAN All Flash Config	1
BNW9	ThinkSystem 2.5" PM1655 1.6TB Mixed Use SAS 24Gb HS SSD	2
BP3K	ThinkSystem 2.5" PM1655 6.4TB Mixed Use SAS 24Gb HS SSD	6
B8LU	ThinkSystem 2U 8x2.5" SAS/SATA Backplane	1
BH8B	ThinkSystem 2U/4U 8x2.5" AnyBay Backplane	1
BS7A	Select Storage Devices - Configured M.2/7mm RAID	1
BM8X	ThinkSystem M.2 SATA/x4 NVMe 2-Bay Enablement Kit	1
BTTY	M.2 NVMe	1
BS7M	Intel VROC (VMD NVMe RAID) Standard for M.2	1
BS7F	M.2 NVMe Array 1 RAID 1	1
BKSR	ThinkSystem M.2 7450 PRO 960GB Read Intensive NVMe PCIe 4.0 x4 NHS SSD	2
BLA3	SW stack for ThinkAgile VX Appliance	1
BRPK	ThinkSystem NVIDIA BlueField-2 25GbE SFP56 2-Port PCIe Ethernet DPU w/BMC & Crypto	1
BPK9	ThinkSystem 1800W 230V Titanium Hot-Swap Gen2 Power Supply	2
BLL6	ThinkSystem 2U V3 Performance Fan Module	6
B8LA	ThinkSystem Toolless Slide Rail Kit v2	1
BS7J	M.2 NVMe Array 1 HDDs	2
BV2W	ThinkAgile VX650 V3-DPU IS	1
BQBP	ThinkSystem MCC CPU Clip	2
BM8T	ThinkSystem SR650 V3 Firmware and Root of Trust Security Module	1
BSH7	ThinkSystem OCP Enablement Kit for Distributed Services Engine	1
BP46	ThinkSystem 2U Main Air Duct	1
BLL3	ThinkSystem SR650 V3 PSU Duct	1
BPDR	ThinkSystem SR850 V3/SR650 V3 Standard Heatsink w/ Heatpipes	2
B8MM	ThinkSystem 2U MS 3FH Riser Filler	1
AVEQ	ThinkSystem 8x1 2.5" HDD Filler	1
AVEN	ThinkSystem 1x1 2.5" HDD Filler	8
AUTQ	ThinkSystem small Lenovo Label for 24x2.5"/12x3.5"/10x2.5"	1
BSLP	ThinkSystem Adapter NC-SI cable - 270mm	1
BMPF	ThinkSystem V3 2U Power Cable from MB to Front 2.5" BP v2	2

Part number Feature code	Product Description	Qty
BRPW	ThinkSystem SR650 V3,PCIe Gen4 CBL,SLx8-Swift,M.2-PCIe3	1
BACB	ThinkSystem V3 2U SAS/SATA Y Cable from CFF C0,C1/ C2,C3 to Front 8x2.5" BP	2
BSYM	ThinkSystem SR650 V3,PCIe4 Cable,Swift8x-SL8x,2in1,PCIe 6/5(MB) to BP1/BP2	1
BMP2	ThinkSystem V3 2U Power Cable from MB to CFF / Exp v2	1
BRPV	ThinkSystem SR650 V3,PCIe Gen4 CBL,SLx8-Swift,CFF IN-PCIe4	1
BPE3	ThinkSystem SR650 V3 MCIO8x to SL8x CBL, PCIe4, 8x2.5AnyBay, 200mm	2
BLKL	ThinkSystem V3 2U x16/x8/x8 PCIe Gen4 Riser1 or 2	1
7S06CTOSWW	VMware vSAN for VX	1
SAQT	VMware vSAN 8 Enterprise for 1 processor w/Lenovo 3Yr S&S	2
5PS7B23039	Essential ThinkAgile IS - 3Yr 24x7 4Hr Resp + YDYD VX650 V3 DPU	1

Conclusion

Lenovo ThinkAgile VX650 V3 DPU and ThinkSystem SR650 V3 DPU with Nvidia BlueField-2 DPU adapters and VMware vSphere Distributed Services Engine are empowering next generation infrastructure with accelerators and unified management, deployment and lifecycle operations. These next generation engineered hardware and software solutions from Lenovo, Intel, VMware and NVIDIA are one stop scalable solutions for modern application development and hybrid cloud scenarios with reduced TCO.

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 Lenovo workload solutions on ThinkAgile VX650 V3 and ThinkSystem SR650 V3 servers, contact your Lenovo Business Partner or visit: <https://www.lenovo.com/systems/solutions>

References:

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

Lenovo ThinkSystem SR650 V3 server: <https://lenovopress.lenovo.com/lp1601>

Related product families

Product families related to this document are the following:

- [ThinkAgile VX Series for VMware](#)
- [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, LP1768, was created or updated on June 30, 2023.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP1768>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP1768>.

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®

AnyBay®

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.