

Lenovo CAE Platform Guide for Computational Fluid Dynamics

Planning / Implementation

In an era where innovation, speed, and precision define market leadership, Computer-Aided Engineering (CAE) has become an essential tool for modern product development. CAE enables organizations to simulate and analyze the physical behavior of products and systems in a virtual environment—long before physical prototypes are built. This approach helps reduce development costs, shorten time-to-market, and improve product quality and performance.

Across industries—from automotive and aerospace to energy, manufacturing, and consumer electronics—CAE supports critical engineering decisions. It allows teams to explore design alternatives, validate performance under real-world conditions, and ensure compliance with safety and regulatory standards. As products become more complex and performance expectations rise, the ability to simulate and optimize designs virtually is no longer optional—it's a competitive necessity.

However, the growing complexity of simulations—often involving millions of variables and intricate physical interactions—demands significant computational resources. This is where High-Performance Computing (HPC) becomes a strategic enabler. HPC platforms provide the computational power needed to run large-scale, high-fidelity simulations quickly and efficiently. With HPC, engineering teams can explore more design alternatives, run more detailed models, and make faster, data-driven decisions.

Importance of Computational Fluid Dynamics

One of the most powerful and widely used CAE tools is **Computational Fluid Dynamics (CFD)**. CFD is used to simulate the behavior of fluids—liquids and gases—as they interact with surfaces and environments. It plays a critical role in industries such as aerospace, automotive, energy, and electronics, where understanding airflow, heat transfer, and fluid behavior is essential to product success.

CFD enables engineers to:

- Analyze airflow over aircraft wings or vehicle bodies to reduce drag and improve fuel efficiency.
- Optimize cooling systems in electronics and power equipment.
- Simulate ventilation and air quality in buildings and industrial facilities.
- Improve the performance and safety of pumps, turbines, and other fluid-handling equipment.
- Support city planning by modeling urban heat effect.

However, CFD simulations are computationally intensive. They often involve solving millions of equations to capture the complex physics of fluid or air motion. CFD workloads are both compute-intensive and memory-intensive, requiring significant processing power and high memory bandwidth to solve complex physical models. These simulations typically scale efficiently across many CPU cores and nodes, making them ideal for HPC environments where the infrastructure enables large-scale execution—reducing turnaround time, improving model accuracy, and supporting more robust design optimization. The ability to increase complexity of the model at scale allows more factors to be considered taking into account macro factors affecting designs.

For decision-makers, the integration of CAE, CFD, and HPC represents a strategic investment in innovation. It empowers engineering teams to solve more complex problems, deliver better-performing products, and maintain a competitive edge in a rapidly evolving market.

CFD Software Platforms

The Lenovo EveryScale CAE Reference Architecture for CFD is designed to support a broad range of these tools, ensuring flexibility and performance across diverse CFD workloads.

Leading CFD Software Platforms

- **ANSYS® Fluent®**
One of the most widely adopted CFD tools in the industry, Fluent® offers robust capabilities for simulating complex fluid flow, turbulence, heat transfer, and chemical reactions. It is used extensively in aerospace, automotive, energy, and electronics sectors for high-fidelity simulations and design optimization.
- **Siemens™ Simcenter™ STAR-CCM+™**
STAR-CCM+™ is known for its integrated multiphysics capabilities, combining CFD with thermal, structural, and motion analysis. It is particularly valued for its automation, scalability, and ability to handle complex geometries and transient simulations, making it a strong choice for advanced engineering applications.
- **OpenFOAM®**
An open-source CFD toolbox, OpenFOAM® is widely used in academia and industry for its flexibility and extensibility. It supports a wide range of solvers and physical models and is ideal for organizations looking to customize their simulation workflows or reduce licensing costs.

Most commercial CFD software from CAE ISVs is traditionally licensed based on the number of CPU cores used, which can constrain scalability and increase costs for large simulations. However, newer licensing models—such as Simcenter STAR-CCM+'s Power Session and Ansys Fluent's HPC Ultimate—remove core count restrictions, allowing users to fully leverage high-core-count systems without incurring additional licensing fees. These models, along with open-source solvers like OpenFOAM, are reshaping compute strategies by encouraging the use of high-density CPUs that prioritize total throughput over per-core efficiency. This shift opens new opportunities for maximizing simulation performance and return in investment modern HPC environments. The Lenovo EveryScale architecture is optimized to support this evolution, delivering the compute power, memory bandwidth, and interconnect performance required to run these CFD workloads efficiently and at scale. Application performance improvements and scalability have been validated by Lenovo's expert performance engineers using production level hardware in the Lenovo HPC Innovation Center.

Overview

The Lenovo EveryScale CAE reference architecture for CFD is built on a Scalable Unit (SU) structure design, with 48Units with 1,920 servers and 3,840 CPUs within a standard FAT Tree Network Topology. With alternative network topologies such as Dragonfly+, the architecture has the potential to scale exponentially.

Each Scalable Unit comprises of a single compute rack equipped with a high-speed InfiniBand leaf switch and Ethernet connectivity. The Lenovo Heavy Duty Rack Cabinets offer sufficient cable routing channels to efficiently direct InfiniBand connections to the adjacent racks, while also accommodating all necessary Ethernet connections. This SU is purpose-built for seamless scalability, offering on-demand growth to support CFD simulations of any resolution or complexity.

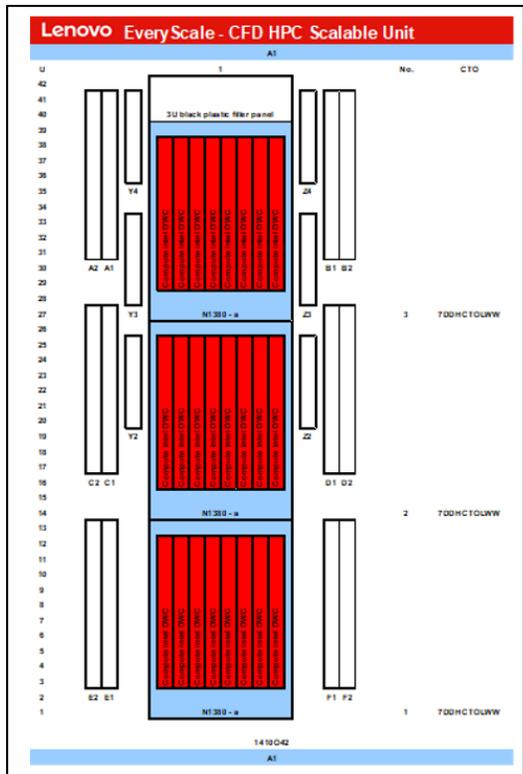


Figure 1. Single Rack Scalable Unit for CFD

Components

The main hardware components of Lenovo CAE RA for CFD are Compute nodes and the Networking infrastructure. As an integrated solution they come together in a Lenovo EveryScale Rack (Machine Type 1410).

- [Compute – N1380 and SC750 V4](#)
- [Network – Ethernet and InfiniBand](#)
- [Lenovo EveryScale Solution](#)

Compute Infrastructure

The Compute Infrastructure is built on the latest generation of Lenovo Neptune Supercomputing systems.

- [Lenovo ThinkSystem N1380](#)
- [Lenovo ThinkSystem SC750 V4](#)

Lenovo ThinkSystem N1380

The [ThinkSystem N1380 Neptune](#) chassis is the core building block, built to enable exascale-level performance while maintaining a standard 19-inch rack footprint. It uses liquid cooling to remove heat and increase performance and is engineered for the next decade of computational technology.



Figure 2. Lenovo ThinkSystem N1380 Enclosure

N1380 features an integrated manifold that offers a patented blind-mate mechanism with aerospace-grade drip-less connectors to the compute trays, ensuring safe and seamless operation. The unique design of the N1380 eliminates the need for internal airflow and power-consuming fans. As a result, it achieves a reduction in typical data center power consumption by up to 40% compared to similar air-cooled systems.

This newly developed enclosure incorporates up to four ThinkSystem 15kW Titanium Power Conversion Stations (PCS). These stations are directly fed with high current three-phase power and supply power to an internal 48V busbar, which in turn powers the compute trays. The PCS design is a game-changer, merging power conversion, rectification, and distribution into a single package. This is a significant transformation from traditional setups that require separate rack PDUs, additional cables and server power supplies, resulting in best-in-class efficiency.

Each 13U Lenovo ThinkSystem N1380 Neptune enclosure houses eight Lenovo ThinkSystem SC-series Neptune trays. Up to three N1380 enclosures fit into a standard 19" rack cabinet, packing 24 trays into just two 60x60 datacenter floor tiles.

The following table lists the configuration of the N1380 Enclosures.

Table 1. Configuration of the N1380 Enclosure

Part number	Description	Quantity
7DDHCTOLWW	Lenovo ThinkSystem N1380 Neptune Enclosure	1
BYKR	ThinkSystem N1380 Neptune Enclosure Midplate Assembly	1
C4KW	0.95M, 63A 240-415V, 3-Phase Y-Splitter Floor Power Cable	2
BYKV	2.8M, 63A 240-415V, 3-Phase WYE IEC to Y-Splitter Rack Power Cable	2
BE0E	N+N Redundancy With Over-Subscription	1
BYKK	ThinkSystem N1380 Neptune EPDM Hose Connection	1
BYKJ	ThinkSystem N1380 Neptune System Management Module V3	1
BYJZ	ThinkSystem N1380 Neptune Enclosure	1
BYKH	ThinkSystem N1380 Neptune 15kW 3-Phase 200-480V Titanium Power Conversion Station	4
5WS7C20194	5Yr Premier 24x7 4Hr Resp N1380 Neptune Enclosure	1

Lenovo ThinkSystem SC750 V4

The [ThinkSystem SC750 V4 Neptune](#) node is the next-generation high-performance server based on the sixth generation Lenovo Neptune direct water cooling platform.

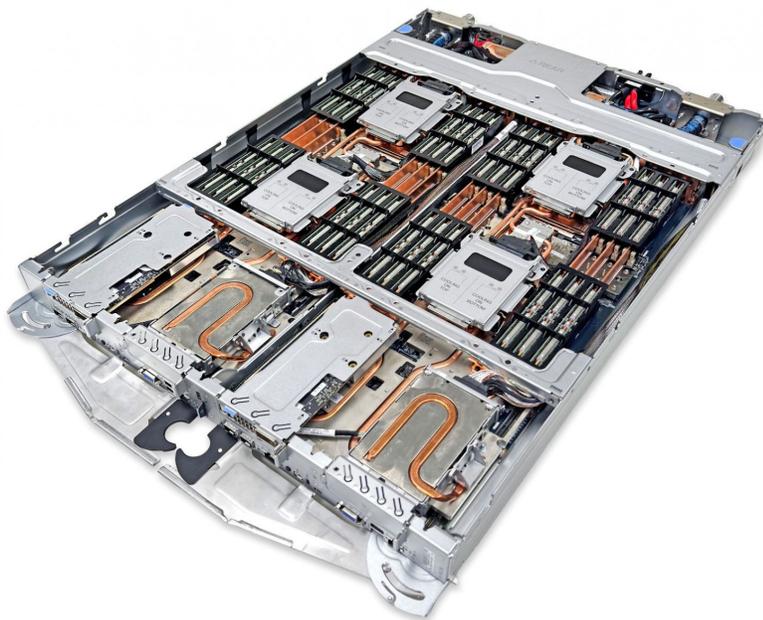


Figure 3. Lenovo ThinkSystem SC750 V4 Neptune Server Tray

Supporting the Intel Xeon 6900P-series, the ThinkSystem SC750 V4 Neptune stands as a powerhouse for demanding HPC workloads. Its industry-leading direct water-cooling system ensures steady heat dissipation, allowing CPUs to maintain accelerated operation and achieve up to a 10% performance enhancement.

With 12 channels of high-speed DDR5 RDIMM or an impressive 8800MHz high-bandwidth MRDIMM capability, it excels in memory bandwidth-intensive workloads, positioning it as a preferred choice for engineering and meteorology applications like Fluent, STAR-CCM+, OpenFOAM, WRF, and ICON.

Completing the package with support for high-performance NVMe and high-speed, low latency networking with the latest InfiniBand, Omnipath, and Ethernet choices, the SC750 V4 is your all-in-one solution for HPC workloads.

At its core, Lenovo Neptune applies 100% direct warm-water cooling, maximizing performance and energy efficiency without sacrificing accessibility or serviceability. The SC750 V4 is installed into the ThinkSystem N1380 Neptune enclosure which itself integrates seamlessly into a standard 19" rack cabinet. Featuring a patented blind-mate stainless steel dripless quick connection, SC750 V4 node trays can be added "hot" or removed for service without impacting other node trays in the enclosure.

This modular design ensures easy serviceability and extreme performance density, making the SC750 V4 the go-to choice for compute clusters of all sizes - from departmental/workgroup levels to the world's most powerful supercomputers – from Exascale to Everscale.

Intel Xeon 6900-Series processors with P-cores

Intel Xeon 6 processors with P-cores are optimized for high performance per core. With more cores, double the memory bandwidth, and AI acceleration in every core, Intel Xeon 6 processors with P-cores provide twice the performance for the widest range of workloads, including HPC and AI.

Domains such as CAE/CFD and weather/climate modeling present a more balanced performance profile—demanding both high compute throughput and substantial memory bandwidth. Simply increasing core counts can lead to diminishing returns unless accompanied by improvements in memory access speed, latency, and power delivery. The Intel Xeon 6900 series addresses these challenges by expanding the thermal design power (TDP) to 500 watts, which, when coupled with Lenovo Neptune cooling, helps to sustain or even boost CPU frequencies under heavy loads. Additionally, support for 12 memory channels and compatibility with DDR5-6400 MHz and MRDIMM-8800 MHz memory types significantly increases memory bandwidth, ensuring that high-core-count systems remain efficient and scalable for memory-intensive workloads. Reflecting these architectural advantages, CFD workloads from Fluent, STAR-CCM+, and OpenFOAM have observed over 2.4x faster computational times on average when running on Intel Xeon 6900 P-core processors compared to the previous generation—demonstrating the real-world impact of these enhancements on simulation throughput and productivity.

MRDIMM technology

MRDIMM technology, or multiplexed rank DIMMs, represents a significant advancement in memory performance, particularly for memory-intensive workloads such as CFD. Companies like Micron have been at the forefront of developing high-speed memory solutions, contributing to the reliability and efficiency of MRDIMMs.

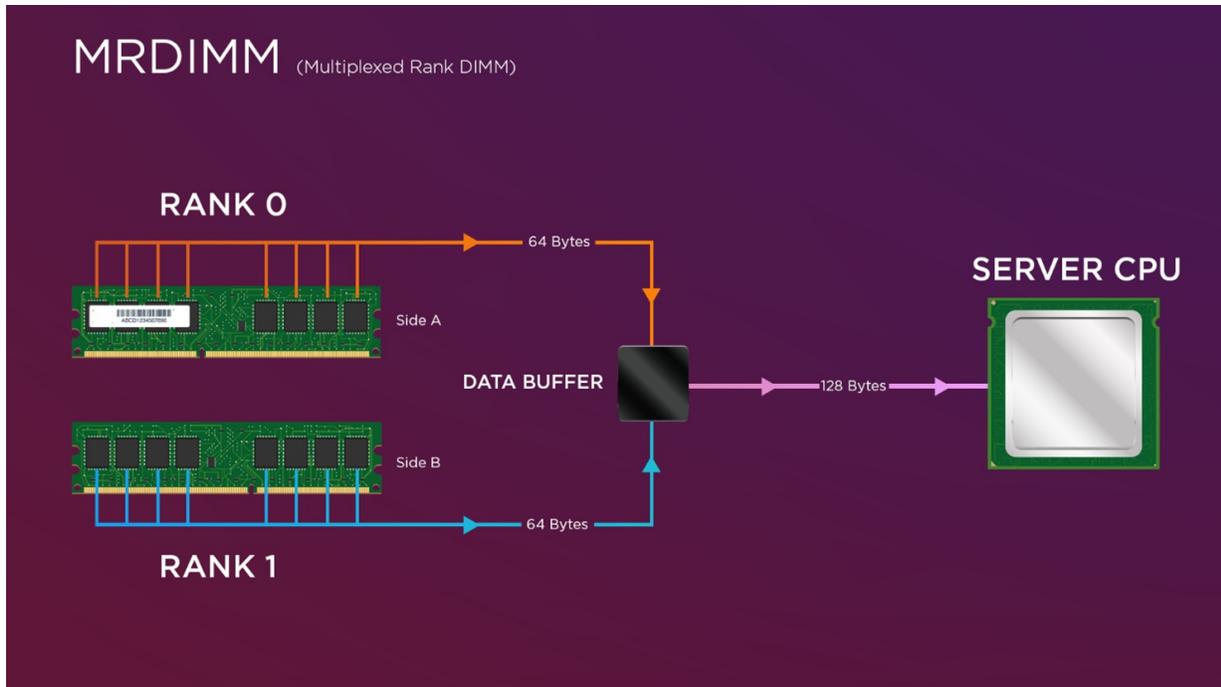


Figure 4. MRDIMM Multiplex Functionality

Operating at speeds of up to 8800 MT/s, MRDIMMs significantly reduce memory access latency and increase bandwidth, especially when paired with Intel Xeon 6th Gen processors. This combination has demonstrated up to a 200% performance improvement over the previous Xeon generation, making MRDIMMs a critical enabler for modern HPC environments. These benefits are particularly impactful for computationally intensive workloads like CFD, where memory speed and efficiency directly influence simulation performance. In real-world testing, CFD applications such as Fluent, STAR-CCM+, and OpenFOAM have shown an average of 1.2× faster computational times with MRDIMMs—highlighting the tangible value of high-bandwidth memory in accelerating engineering workflows.

Configuration

Each node is equipped with two Intel Xeon 6980P CPUs, each comprising 128 Xeon6 P-cores. This configuration provides 256 Xeon6 P-cores and 1.5TB of MRDIMM RAM per node, making the SC750 V4 highly suited for core and RAM-intensive tasks. Utilizing high-speed, low-latency network adapters at 200, 400, or 800Gbps, the SC750 V4, when paired with Intel Xeon6, offers exceptional scalability for the most demanding parallel MPI jobs.

The following table lists the configuration of the SC750 V4 Trays.

Table 2. Configuration of the SC750 V4 trays

Part number	Description	Quantity
7DDJCTOLWW	Lenovo ThinkSystem SC750 V4 Neptune Tray	1
BZ7E	ThinkSystem SC750 V4 ConnectX-7 Auxiliary Cable	1
5977	Select Storage devices - no configured RAID required	1
BPKR	TPM 2.0	1
C2WQ	Intel Xeon 6980P 128C 500W 2.0GHz Processor	4
BKSP	ThinkSystem NVIDIA ConnectX-7 NDR OSFP400 1-port PCIe Gen5 x16 InfiniBand Adapter (SharedIO) DWC	1
C0TY	ThinkSystem 32GB TruDDR5 8800MHz (2Rx8) MRDIMM	48
BYK4	ThinkSystem SC750 V4 Neptune Tray	1
B7Y0	Enable IPMI-over-LAN	1
5WS7C20199	5Yr Premier 24x7 4Hr Resp SC750 V4 Neptune Tray	1

Network Infrastructure

The Network Infrastructure is built on NVIDIA networking technology for both InfiniBand and Ethernet.

- [High Performance Network](#)
- [Management Network](#)

High Performance Network

Each Scalable Unit consist of a single compute rack housing 48 nodes shared across 24 Lenovo ThinkSystem SC750 V4 compute trays. Each compute tray is equipped with a high-speed InfiniBand adapter that supports up to 400 Gbps. The SC750 V4's InfiniBand adapters are Lenovo SharedIO capable and feature an internal high-speed PCIe link between two compute nodes. This configuration enables both nodes to share a single 4000 Gbps InfiniBand connection, effectively reducing cabling and switch port requirements by 50% across the system. The trade-off is that each compute node receives 200 Gbps of network bandwidth. However, performance testing has shown that this reduction has a minimal impact on CFD workloads, with less than a 10% performance drop observed even at large scales.

This streamlined approach offers a cost-effective solution for scaling CPU and memory-intensive workloads. By maintaining a balanced design, customers can accurately scale their workload when CPU and memory tasks heavily outweigh inter-node communication requirements. This optimized configuration results in significant cost savings while delivering optimal price and performance.

The following table lists the configuration of the NVIDIA QM9790 NDR InfiniBand Leaf Switch.

Table 3. Configuration of the NVIDIA QM9790 NDR InfiniBand Leaf Switch

Part number	Description	Quantity per system
0724HED	NVIDIA QM9790 64-Port Unmanaged Quantum NDR InfiniBand Switch (oPSE)	1
BP64	NVIDIA QM9790 64-Port Unmanaged Quantum NDR InfiniBand Switch (oPSE)	1
BRQ6	2.8m, 10A/100-250V, C15 to C14 Jumper Cord	2
BQJD	ThinkSystem NDRx2 OSFP800 IB Multi ModeTwin-Transceiver	16
BQJN	Lenovo 3M NVIDIA NDR Multi Mode Optical Cable	4
BQJR	Lenovo 10M NVIDIA NDR Multi Mode Optical Cable	14
BQJS	Lenovo 20M NVIDIA NDR Multi Mode Optical Cable	14
BQJX	Lenovo 2M NVIDIA NDRx2 OSFP800 to 2x NDR OSFP400 Passive Copper Splitter Cable	8
BQJY	Lenovo 3M NVIDIA NDRx2 OSFP800 to 2x NDR OSFP400 Passive Copper Splitter Cable	8
BP66	NVIDIA QM97xx Enterprise Rack Mount Kit	1
5WS7B96633	5Yr Premier 24x7 4Hr Resp NVID QM9790 oPSE	1

Management Network

Cluster management is usually done over Ethernet, and the SC750 V4 offers multiple options. It comes with 25GbE SFP28 Ethernet ports, a Gigabit Ethernet port, and a dedicated XClarity Controller (XCC) port. These can be customized based on cluster management and workload needs.

For stable environments with infrequent OS changes like CAE systems, the single Gigabit Ethernet port suffices. A CAT5e or CAT6 cable per node can use Network Controller Sideband Interface (NC-SI) for remote out-of-band and cluster management over one wire. For higher bandwidth needs or frequent updates, the 25Gb Ethernet interfaces offer additional capacity and support sideband communication to the XCC.

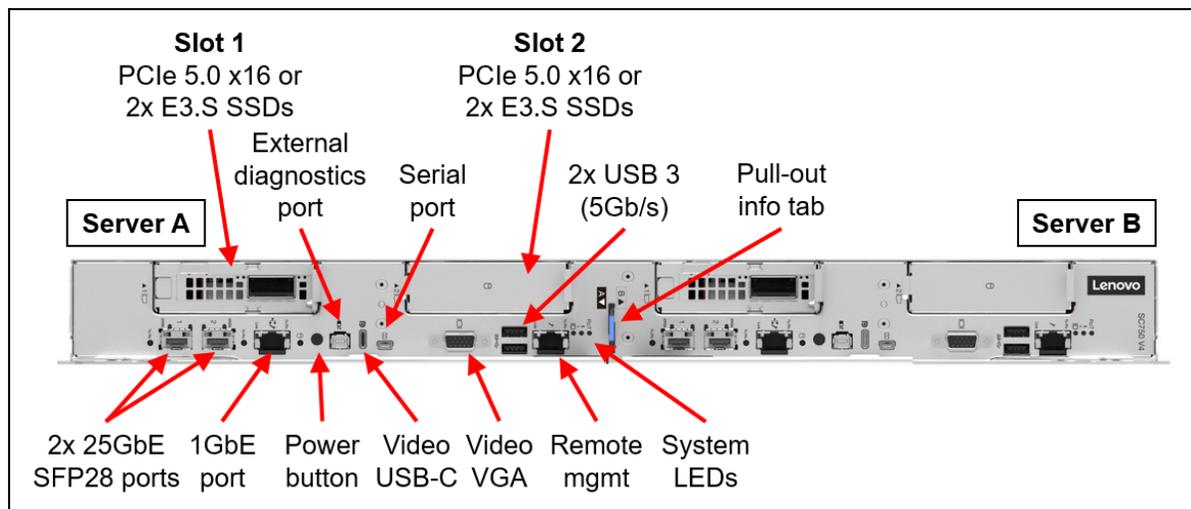


Figure 5. Front view the SC750 V4 with management ports

The SC750 V4 integrates the XCC through the Data Center Secure Control Module (DC-SCM) I/O board. This module also includes a Root of Trust module (NIST SP800-193 compliant), USB 3.2 ports, a VGA port, and MicroSD card capability for additional storage with the XCC, offering firmware storage options up to 4GB, including N-1 firmware history.

The N1380 enclosure features a System Management Module 3 (SMM) at the rear, managing both the enclosure and individual servers through a web browser or Redfish/IPMI 2.0 commands. The SMM provides remote connectivity to XCC controllers, node-level reporting, power control, enclosure power management, thermal management, and inventory tracking.

In the context of this CAE reference architecture, the gigabit Ethernet interfaces of each SC750 V4 compute tray are linked to an Nvidia SN2201 Management Leaf switch, ensuring connectivity for the compute nodes' cluster management, out-of-band management, and N1380 enclosure systems management modules (SMM).

The following table lists the configuration of the NVIDIA SN2201 1GbE Management Leaf Switches.

Table 4. Configuration of the NVIDIA SN2201 1GbE Management Leaf Switches

Part number	Description	Quantity per system
7D5FCTOGWW	NVIDIA SN2201 1GbE Managed Switch with Cumulus (oPSE)	1
BPC8	NVIDIA SN2201 1GbE Managed Switch with Cumulus (oPSE)	1
6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2
3793	3m Yellow Cat5e Cable	1
B306	Mellanox QSA 100G to 25G Cable Adapter	2
BFH2	Lenovo 25Gb SR SFP28 Ethernet Transceiver	2
BSNA	NVIDIA SN2201 Enterprise Rack Mount Kit for Recessed Mounting	1
5WS7B98278	5Yr Premier 24x7 4Hr Resp NVID SN2201 PSE	1

Lenovo EveryScale Solution

The Server and Networking components and Operating System can come together as a [Lenovo EveryScale Solution](#). It is a framework for designing, manufacturing, integrating and delivering data center solutions, with a focus on High Performance Computing (HPC), Technical Computing, and Artificial Intelligence (AI) environments.

Lenovo EveryScale provides [Best Recipe](#) guides to warrant [interoperability](#) of hardware, software and firmware among a variety of Lenovo and third-party components.

Addressing specific needs in the data center, while also optimizing the solution design for application performance, requires a significant level of effort and expertise. Customers need to choose the right hardware and software components, solve interoperability challenges across multiple vendors, and determine optimal firmware levels across the entire solution to ensure operational excellence, maximize performance, and drive best total cost of ownership.

Lenovo EveryScale reduces this burden on the customer by pre-testing and validating a large selection of Lenovo and third-party components, to create a “Best Recipe” of components and firmware levels that work seamlessly together as a solution. From this testing, customers can be confident that such a best practice solution will run optimally for their workloads, tailored to the client’s needs.

In addition to interoperability testing, Lenovo EveryScale hardware is pre-integrated, pre-cabled, pre-loaded with the best recipe and optionally an OS-image and tested at the rack level in manufacturing, to ensure a reliable delivery and minimize installation time in the customer data center.



Figure 6. Lenovo EveryScale Heavy Duty Enterprise Rack Cabinet

Scalability

As outlined in the component selection, the Lenovo reference design for CFD workloads employs a high-speed network utilizing NDR InfiniBand. This fifth-generation InfiniBand fabric is implemented in a Fat Tree topology, enabling scalability up to 2,048 compute nodes through Lenovo SharedIO and 32 QM9790 NDR InfiniBand leaf switches. Each leaf switch can connect to the spine network via 32 NDR uplinks. The spine layer would consist of 16 QM9790 NDR InfiniBand spine switches. The topology provides a solid network interconnect foundation for CFD solutions, addressing the needs for scalability, high bandwidth, and low latency, ensuring robust performance for complex simulations and models.

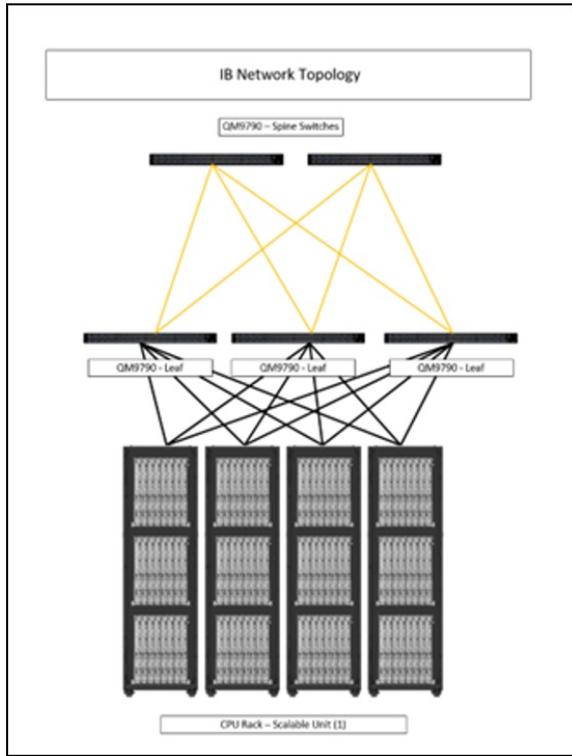


Figure 7. Scale Out Network Topology

Performance

The performance characteristics of CFD workloads are typically balanced between compute-intensive and memory-intensive demands. Compute-intensive workloads benefit from a high number of processor cores, elevated CPU frequencies, and increased instructions per cycle (IPC), which accelerate the execution of scalar and vector operations. In contrast, memory-intensive workloads depend heavily on high memory bandwidth and low latency, as performance is closely tied to the speed at which data can be read from and written to memory. Additionally, larger CPU caches can help mitigate memory bottlenecks by reducing the frequency of main memory accesses.

The degree to which a CFD workload leans toward compute or memory intensity depends on several factors, including the specific software application, the resolution of the simulation model, and the complexity of the physics being modeled. Higher-resolution models—which involve finer meshes and more detailed physics—tend to be more memory-bound, requiring greater memory bandwidth to maintain simulation efficiency. Conversely, lower-resolution models and simulations with chemical reactions such as combustion models are often more compute-bound and benefit from CPUs with higher clock speeds and strong single-threaded performance.

The following benchmark results highlight the performance impact of MRDIMMs across three leading CFD applications: ANSYS Fluent, Siemens Simcenter STAR-CCM+, and OpenFOAM. A diverse set of simulation models was selected, spanning from low-resolution models – such as Fluent’s Aircraft Wing model with 14 million cells (under 20 million cells) to high-resolution models exceeding 100 million cells, including STAR-CCM+’s VTM Bench model with 178 million cells. These benchmarks also included complex flow scenarios, such as thermal management simulations, to reflect real-world engineering challenges.

This range of model sizes and physics ensures a representative mix of workloads relevant to the broader engineering community. When averaged across all three applications and workload types, the Intel Xeon 6900-Series with MRDIMMs delivers a 2.4x overall performance improvement compared to the previous generation—highlighting its significant impact in accelerating CFD simulations across a broad range of use cases.

The chart below compares the performance of a 2-socket compute node featuring Intel 8592+ CPUs (Xeon 5) against a node equipped with Intel 6980P CPUs (Xeon 6) and MRDIMMs. It highlights relative performance across a range of Fluent, STAR-CCM+, and OpenFOAM workloads as model resolution increases. Performance is shown relative to the Intel 8592+ baseline, illustrating the gains achieved with the newer Xeon 6 architecture.

The figure below with the benchmark results highlights these distinctions.

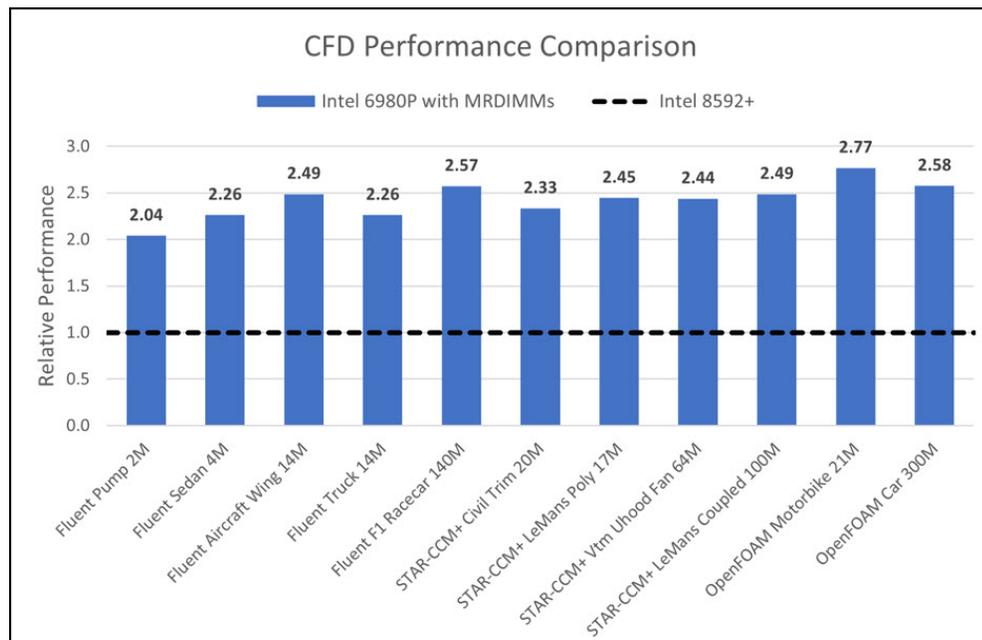


Figure 8. CFD Application Performance Comparison

Both memory bandwidth and compute power are key drivers of CFD workload performance. While it may be tempting to attribute the performance gains shown above solely to generational CPU improvements, further analysis reveals a more nuanced picture. When using the same Intel Xeon 6980P CPU configured with both MRDIMMs and standard DDR5 RDIMMs, the results indicate that the observed performance improvements are not solely due to CPU advancements. Servers with MRDIMMs deliver an additional performance improvement of 1.2x over DDR5 RDIMMs.

The chart below compares the performance of two 2-socket compute nodes, both featuring Intel 6980P CPUs, one configured with DDR5 RDIMMs and the other with MRDIMMs. It highlights relative performance across a range of Fluent, STAR-CCM+, and OpenFOAM workloads as model resolution increases. Performance is shown relative to the RDIMM-based node, illustrating the performance gains enabled by the MRDIMM memory technology.

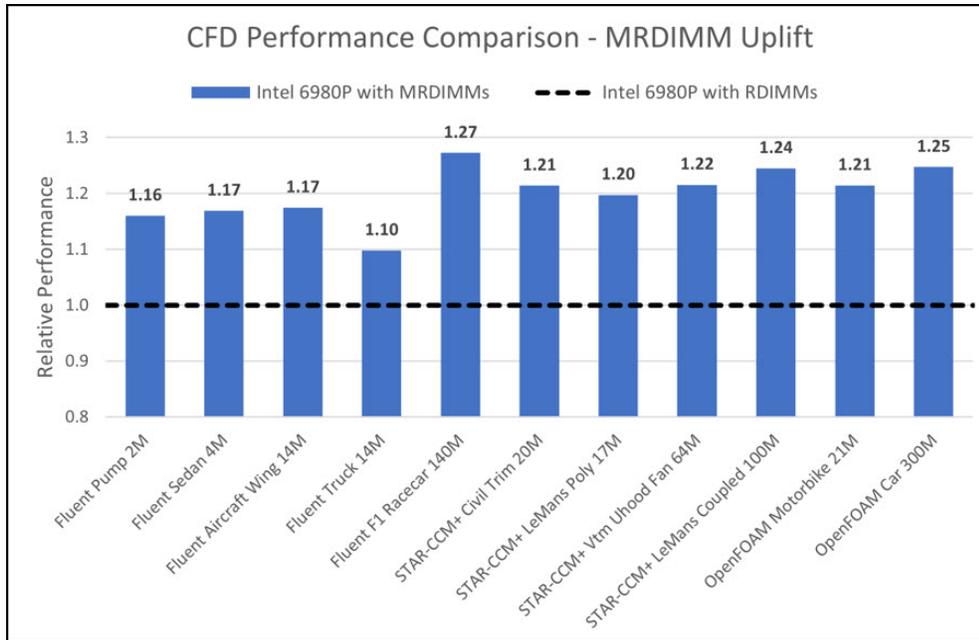


Figure 9. CFD Application MRDIMM Uplift

The results presented above are derived from benchmark runs by Lenovo HPC Innovation Center using the Lenovo ThinkSystem SC750 V5 with Intel Xeon 6 CPUs and Micron MRDIMM memory.

Lenovo TruScale

Lenovo TruScale XaaS is your set of flexible IT services that makes everything easier. Streamline IT procurement, simplify infrastructure and device management, and pay only for what you use – so your business is free to grow and go anywhere.

Lenovo TruScale is the unified solution that gives you simplified access to:

- The industry's broadest portfolio – from pocket to cloud – all delivered as a service
- A single-contract framework for full visibility and accountability
- The global scale to rapidly and securely build teams from anywhere
- Flexible fixed and metered pay-as-you-go models with minimal upfront cost
- The growth-driving combination of hardware, software, infrastructure, and solutions – all from one single provider with one point of accountability.

For information about Lenovo TruScale offerings that are available in your region, contact your local Lenovo sales representative or business partner.

Lenovo Financial Services

Why wait to obtain the technology you need now? No payments for 90 days and predictable, low monthly payments make it easy to budget for your Lenovo solution.

- **Flexible**

Our in-depth knowledge of the products, services and various market segments allows us to offer greater flexibility in structures, documentation and end of lease options.

- **100% Solution Financing**

Financing your entire solution including hardware, software, and services, ensures more predictability in your project planning with fixed, manageable payments and low monthly payments.

- **Device as a Service (DaaS)**

Leverage latest technology to advance your business. Customized solutions aligned to your needs. Flexibility to add equipment to support growth. Protect your technology with Lenovo's Premier Support service.

- **24/7 Asset management**

Manage your financed solutions with electronic access to your lease documents, payment histories, invoices and asset information.

- **Fair Market Value (FMV) and \$1 Purchase Option Leases**

Maximize your purchasing power with our lowest cost option. An FMV lease offers lower monthly payments than loans or lease-to-own financing. Think of an FMV lease as a rental. You have the flexibility at the end of the lease term to return the equipment, continue leasing it, or purchase it for the fair market value. In a \$1 Out Purchase Option lease, you own the equipment. It is a good option when you are confident you will use the equipment for an extended period beyond the finance term. Both lease types have merits depending on your needs. We can help you determine which option will best meet your technological and budgetary goals.

Ask your Lenovo Financial Services representative about this promotion and how to submit a credit application. For the majority of credit applicants, we have enough information to deliver an instant decision and send a notification within minutes.

Bill of materials – First Scalable Unit

This section provides an example Bill of Materials (BoM) of one Scaleable Unit (SU) deployment. This example BoM includes:

- 1x Lenovo Heavy Duty 48U Rack Cabinets
- 3x Lenovo ThinkSystem N1380 Neptune Enclosures
- 24x Lenovo ThinkSystem SC750 V4 Compute Dual Node Trays
- 1x QM9790 Quantum NDR InfiniBand Leaf Switches
- 1x SN2201 Spectrum Gigabit Ethernet Leaf Switches

Note: Storage is optional and not included in this BoM.

Tables in this section:

- [Lenovo ThinkSystem N1380 Neptune Enclosure](#)
- [ThinkSystem SC750 V4 Neptune Tray](#)
- [NVIDIA QM9790 Quantum NDR InfiniBand Leaf Switches](#)
- [NVIDIA SN2201 Gigabit Ethernet Leaf Switches](#)
- [Lenovo Heavy Duty 48U Rack Cabinet](#)

Lenovo ThinkSystem N1380 Neptune Enclosure

Table 5. Lenovo ThinkSystem N1380 Neptune Enclosure

Part number	Description	Quantity per system	Total quantity
7DDHCTOLWW	Lenovo ThinkSystem N1380 Neptune Enclosure	1	3
BYKR	ThinkSystem N1380 Neptune Enclosure Midplate Assembly	1	3
C4KW	0.95M, 63A 240-415V, 3-Phase Y-Splitter Floor Power Cable	2	6
BYKV	2.8M, 63A 240-415V, 3-Phase WYE IEC to Y-Splitter Rack Power Cable	2	6
BE0E	N+N Redundancy With Over-Subscription	1	3
BYKK	ThinkSystem N1380 Neptune EPDM Hose Connection	1	3
BYKJ	ThinkSystem N1380 Neptune System Management Module V3	1	3
BYJZ	ThinkSystem N1380 Neptune Enclosure	1	3
BYKH	ThinkSystem N1380 Neptune 15kW 3-Phase 200-480V Titanium Power Conversion Station	4	12
5WS7C20194	5Yr Premier 24x7 4Hr Resp N1380 Neptune Enclosure	1	3

ThinkSystem SC750 V4 Neptune Tray

Table 6. ThinkSystem SC750 V4 Neptune Tray

Part number	Description	Quantity per system	Total quantity
7DDJCTOLWW	Lenovo ThinkSystem SC750 V4 Neptune Tray	1	24
BZ7E	ThinkSystem SC750 V4 ConnectX-7 Auxiliary Cable	1	24
5977	Select Storage devices - no configured RAID required	1	24
BPKR	TPM 2.0	1	24
C2WQ	Intel Xeon 6980P 128C 500W 2.0GHz Processor	4	96
BKSP	ThinkSystem NVIDIA ConnectX-7 NDR OSFP400 1-port PCIe Gen5 x16 InfiniBand Adapter (SharedIO) DWC	1	24
C0TY	ThinkSystem 32GB TruDDR5 8800MHz (2Rx8) MRDIMM	48	1152
BYK4	ThinkSystem SC750 V4 Neptune Tray	1	24
B7Y0	Enable IPMI-over-LAN	1	24
5WS7C20199	5Yr Premier 24x7 4Hr Resp SC750 V4 Neptune Tray	1	24

NVIDIA QM9790 Quantum NDR InfiniBand Leaf Switches

Table 7. NVIDIA QM9790 Quantum NDR InfiniBand Leaf Switches

Part number	Description	Quantity per system	Total quantity
0724HED	NVIDIA QM9790 64-Port Unmanaged Quantum NDR InfiniBand Switch (oPSE)	1	2
BP64	NVIDIA QM9790 64-Port Unmanaged Quantum NDR InfiniBand Switch (oPSE)	1	2
BRQ6	2.8m, 10A/100-250V, C15 to C14 Jumper Cord	2	4
BQJD	ThinkSystem NDRx2 OSFP800 IB Multi ModeTwin-Transceiver	16	32
BQJN	Lenovo 3M NVIDIA NDR Multi Mode Optical Cable	4	8
BQJR	Lenovo 10M NVIDIA NDR Multi Mode Optical Cable	14	28
BQJS	Lenovo 20M NVIDIA NDR Multi Mode Optical Cable	14	28
BQJX	Lenovo 2M NVIDIA NDRx2 OSFP800 to 2x NDR OSFP400 Passive Copper Splitter Cable	6	12
BQJY	Lenovo 3M NVIDIA NDRx2 OSFP800 to 2x NDR OSFP400 Passive Copper Splitter Cable	6	12
BP66	NVIDIA QM97xx Enterprise Rack Mount Kit	1	2
5WS7B96633	5Yr Premier 24x7 4Hr Resp NVID QM9790 oPSE	1	2

NVIDIA SN2201 Gigabit Ethernet Leaf Switches

Table 8. NVIDIA SN2201 Gigabit Ethernet Leaf Switches

Part number	Description	Quantity per system	Total quantity
7D5FCTOGWW	Nvidia SN2201 1GbE Managed Switch with Cumulus (oPSE)	1	1
BPC8	Nvidia SN2201 1GbE Managed Switch with Cumulus (oPSE)	1	1
6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2	2
3793	3m Yellow Cat5e Cable	1	1
A1MT	10m Green Cat6 Cable	48	48
B306	Mellanox QSA 100G to 25G Cable Adapter	2	2
BFH2	Lenovo 25Gb SR SFP28 Ethernet Transceiver	2	2
BSNA	NVIDIA SN2201 Enterprise Rack Mount Kit for Recessed Mounting	1	1
5WS7B98278	5Yr Premier 24x7 4Hr Resp NVID SN2201 PSE	1	1

Lenovo Heavy Duty 48U Rack Cabinet

Table 9. Lenovo Heavy Duty 48U Rack Cabinet

Part number	Description	Quantity per system	Total quantity
1410P48	Lenovo EveryScale 48U Pearl Heavy Duty Rack Cabinet	1	1
BJ64	Lenovo EveryScale 48U Pearl Heavy Duty Rack Cabinet	1	1
BJPC	Side Panel Right Installation	1	1
BJ2N	Front Installation of 180mm Extension Kit	1	1
BJPB	Side Panel Left Installation	1	1
C0D6	0U 18 C13/C15 and 18 C13/C15/C19/C21 Switched and Monitored 32A 3 Phase WYE PDU	2	2
BJ2M	Rear Installation of 180mm Extension Kit	1	1
BJ68	ThinkSystem 48U Pearl Heavy Duty Rack Side Panel	2	2
BJ6A	ThinkSystem 48U Pearl Heavy Duty Rack Extension	2	2
BJ67	ThinkSystem 48U Pearl Heavy Duty Rack Rear Door	1	1
2304	Integration Prep	1	1
2310	Solution Specific Test	1	1
AU8K	LeROM Validation	1	1
B1EQ	Network Verification	1	1
5WS7B96703	5Yr Premier 24x7 4Hr Resp EveryScale 48U Rack	1	1

Related publications and links

For more information, see these resources:

- Lenovo EveryScale support page:
<https://datacentersupport.lenovo.com/us/en/solutions/ht505184>
- x-config configurator:
<https://lesc.lenovo.com/products/hardware/configurator/worldwide/bhui/asit/x-config.jnlp>

Authors

Martin W Hiegl is the Executive Director of Advanced Solutions at Lenovo, responsible for the global High-Performance Computing (HPC) and Enterprise Artificial Intelligence (EAI) solution business. He oversees the global EAI and HPC functions, including Sales, Product, Development, Service, and Support, and leads a team of subject matter specialists in Solution Management, Solution Architecture, and Solution Engineering. This team applies their extensive expertise in associated technologies, Supercomputer solution design, Neptune water-cooling infrastructure, Data Science and application performance to support Lenovo's role as the most trusted partner in Enterprise AI and HPC Infrastructure Solutions. Martin holds a Diplom (DH) from DHBW, Stuttgart, and a Bachelor of Arts (Hons) from Open University, London, both in Business Informatics. Additionally, he holds a United States patent pertaining to serial computer expansion bus connection.

David DeCastro is the Director of HPC Solutions Architecture at Lenovo leading a team of subject matter experts and application performance engineers of various disciplines across the HPC portfolio including Weather & Climate. His role consists of developing HPC solution architecture, planning roadmaps, and optimizing go to market strategies for the Lenovo HPC community. David has over 15 years of HPC experience across a variety of industries including, Higher Education, Government, Media & Entertainment, Oil & Gas and Aerospace. He has been heavily engaged optimizing HPC solutions to provide customers the best value for their research investments in a manner that delivers on time and within budget.

Kevin Dean is the HPC Performance and CAE Segment Architect on the HPC Worldwide Customer Engagement Team within the Infrastructure Solutions Group at Lenovo. The role consists of leading the HPC performance engineering process and strategy as well as providing CAE application performance support and leading the customer support for the manufacturing vertical. Kevin has 8 years of experience in HPC and AI application performance support at Lenovo plus 12 years of aerodynamic design and computational fluid dynamics experience in the US defense and automotive racing industries. Kevin holds an MS degree in Aerospace Engineering from the University of Florida and a BS degree in Aerospace Engineering from Virginia Polytechnic Institute and State University.

Related product families

Product families related to this document are the following:

- [High Performance Computing](#)
- [ThinkSystem SC750 V4 Server](#)

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 2025. All rights reserved.

This document, LP2243, was created or updated on June 26, 2025.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP2243>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP2243>.

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®

from Exascale to Everyscale®

Neptune®

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.