



Lenovo Weather and Climate Platform Guide

Planning / Implementation

Weather forecasting is incredibly important for society helping to prepare for and respond to the continuously shifting environmental conditions. Accurate weather predictions can impact agriculture, transportation, disaster management, and everyday activities, helping to protect lives and property.

Climate modeling, on the other hand, helps us understand long-term changes and trends, enabling policymakers and scientists to develop strategies to combat the climate crisis and take actions to mitigate its consequences.

To gain insights into weather and climate, numerical simulations are essential. These simulations use advanced computing technologies to provide detailed analysis, modeling, and projections. High-performance computing systems, like those used for example by Environment & Climate Change Canada (ECCC), allow researchers to run complex simulations that predict weather and climate phenomena more accurately.

This paper describes the supercomputer design at ECCC, intended as a platform guide and reference architecture for Weather and Climate environments. At the core of the design, beginning with the Intel Xeon 6 CPU which leverages a distinct memory subsystem, and extending to the server, compute infrastructure, rack layout, and network topology, this reference architecture aims to facilitate a foundation for performance and reliability required in the Weather & Climate vertical.

Case Study - Environment & Climate Change Canada

Canada uses the Global Environmental Multi-Scale Model (GEM), an advanced weather forecasting and data assimilation system developed by the Recherche en Prévision Numérique (RPN), Meteorological Research Branch (MRB), and the Canadian Meteorological Centre (CMC). The GEM Global Forecast System and the Global Deterministic Prediction System (GDPS) are being used for global data assimilation cycles and medium-range forecasting, as well as regional data assimilation spin-up cycles and short-range forecasting.

In addition to the operational GEM model, several experimental versions are being utilized to offer higher resolution insights and serve as test environments for new parameterizations, thereby continuously enhancing operations.

As is typical for mission-critical national weather forecasting, Canada is deploying two supercomputer environments. These are typically the two most powerful supercomputers in Canada and are represented on the Top500 ranking of the most powerful commercially available high-performance computing systems in the world.

Since 2021, these systems have utilized Lenovo Neptune technology. Lenovo Neptune supports energy-efficient supercomputing by improving operational and cooling efficiency. Consequently, the systems were ranked in the Green500 list of the most energy-efficient commercially available high-performance computing systems within the top three of general-purpose environments.

In addition to the supercomputers that execute large-scale numerical simulations, two pre-processing and post-processing environments are employed for data preparation and visualization of results. These tools facilitate the dissemination of information and reports to various stakeholders, like weather forecasts, alerts and warnings, and weather and climate projections.

Capacity computing refers to systems optimized to solve numerous small problems simultaneously. These systems maximize throughput, handling a high volume of tasks efficiently. Capacity computing is suitable for applications where many independent or loosely coupled tasks need to be processed in parallel, achieving high productivity by completing a large number of small-scale computations.

Capability computing, on the other hand, is designed to tackle large, complex problems that require substantial computational power. These systems focus on the ability to solve individual, computationally intensive tasks quickly. Capability computing is essential for applications where the priority is to execute major simulations or analyses that demand significant resources. The cost of delivering computational performance as capability is typically higher due to the need for advanced hardware and substantial processing power to handle demanding algorithms.

In high-performance computing (HPC), both capacity and capability computing play crucial roles. For organizations like Environment and Climate Change Canada (ECCC), capability computing is particularly important because it allows the rapid processing of large-scale weather and climate models, which are vital for accurate forecasts and risk assessments. Capacity computing is necessary complementing this by efficiently generating numerous products that depict the forecasted environmental state to various users.

The Lenovo ThinkSystem SC750 V4 Neptune, featuring Intel Xeon 6 processors, is engineered to provide the optimal balance between capacity and capability computing in numerical simulation for weather forecasting and climate research.

The Lenovo EveryScale Weather & Climate reference architecture is built on a Scalable Unit (SU) structure with 192 servers and 384 CPUs and scales up to 10 Scalable Units with 2,048 Servers and 4,096 CPUs within a standard FAT Tree Network Topology. Considering other scaling network topologies such as Dragonfly+, the reference architecture could grow exponentially.

A Scalable Unit consists of four compute racks and a central network rack that houses both high-speed InfiniBand leaf and spine switches, as well as Ethernet connectivity. The Lenovo Heavy Duty Rack Cabinets offer sufficient cable routing channels to efficiently direct InfiniBand connections to the adjacent network rack, while also accommodating all necessary Ethernet connections. These five-rack SUs are designed to scale conveniently, providing growth on demand for forecasting models of any size.




Figure 1. Lenovo EveryScale Weather & Climate Scalable Unit

Components

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

- Compute Infrastructure
- Network Infrastructure
- 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-bound workloads, positioning it as a preferred choice for meteorology and engineering applications like WRF, ICON, OpenFOAM, and Fluent.

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.

The Intel Xeon 6900 with P-core processors, boasting up to 128 Performance-cores, offers superior performance that make them ideal for vectorized workloads in fields like biology and chemistry, including applications such as NAMD, GROMACS, LAMMPS, CP2K, and Quantum ESPRESSO, and boost performance for machine and deep learning workloads as well.

With AI acceleration in every core, the Intel Advanced Matrix Extensions (Intel AMX) speeds up inferencing for INT8 and BF16, and it offers support for FP16-trained models, with up to 2,048 floating point operations per cycle per core for INT8 and 1,024 floating point operations per cycle per core for BF16/FP16. This can result in gains up to 2x higher Llama-13B (BF16) on gen to gen performance.

In the field of weather & climate prediction, performance improvements are also observed. Consider NEMO in the field of ocean modelling, there are many number of parameters most key being ocean temperature and velocity. Using Intel Xeon 6 processors ocean simulations run up to 2.35x faster than the previous generation Xeon processors, when comparing a timed step process from start to finish. For climate prediction codes such as GEM we also see better than a 2x improvement in performance due to the increased core count, additional memory channels and enhanced I/O capability when compared to the previous generation of processors.

MRDIMM technology

MRDIMM technology, or multiplexed rank DIMMs, represents a significant advancement in memory performance, particularly for memory-intensive workloads such as climate forecasting models like MOM6, NEMO, WRF, and GEM. Companies like Micron have been at the forefront of developing high-speed memory solutions, contributing to the reliability and efficiency of MRDIMMs.

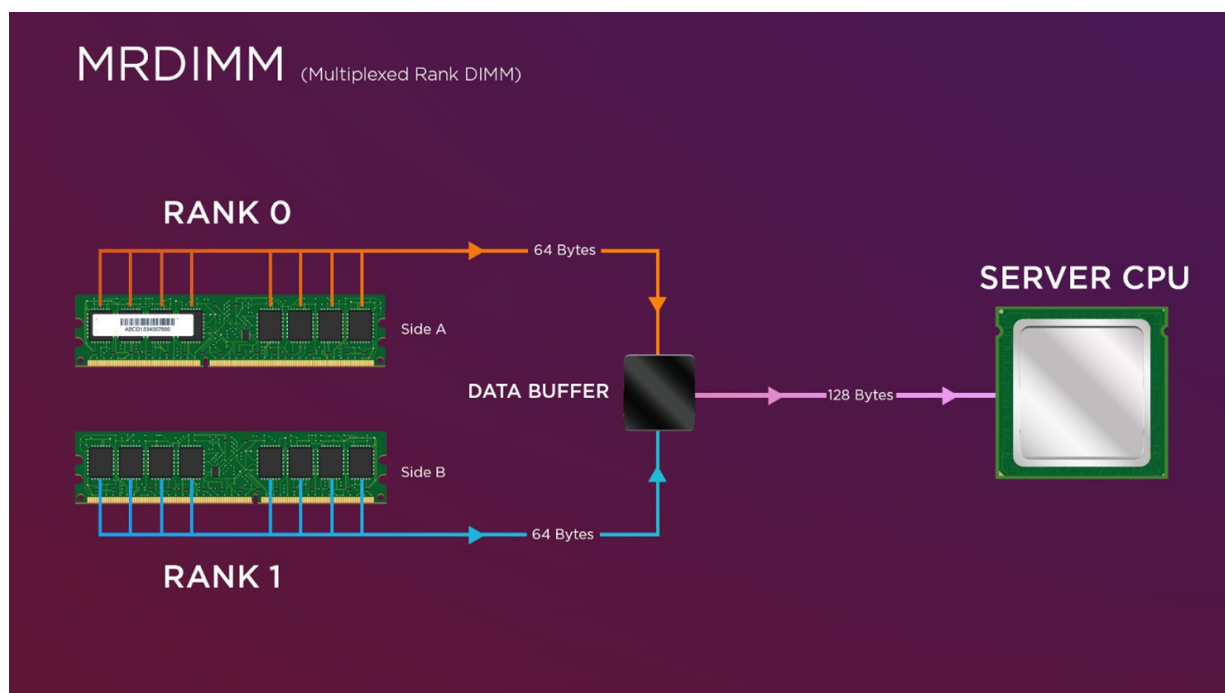


Figure 4. MRDIMM Multiplex Functionality

By operating at speeds of up to 8800 MT/s, these DIMMs drastically reduce the time required to access memory, in combination with Xeon6 leading to a 200% improvement in performance versus previous Xeon generation. This enhancement allows for lower power consumption or increased performance with limited increase in energy usage, making MRDIMMs an important choice for modern high-performance computing systems, such as those used for weather and climate research.

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

The scalable unit consists of four compute racks, each 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 compute nodes. This configuration allows both nodes within a single compute tray to access a single 400 Gbps InfiniBand connection, thereby reducing the cabling and switching requirements by up to 50% in the overall design.

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 weather and climate 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 Weather & Climate 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 weather and climate applications 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 connects to the spine network via 32 NDR uplinks. The spine layer consists of 16 QM9790 NDR InfiniBand spine switches.




Figure 7. Scale Out Network Topology

For larger weather and climate solutions requiring extensive scalability, the Dragonfly+ topology is a highly efficient option. This topology offers fewer diameter networks while maintaining high bandwidth and minimizing latency. Dragonfly+ is particularly beneficial for large node counts needing all-to-all or many-to-many communication, crucial for global weather and climate models. Adaptive routing within the Dragonfly+ ensures balanced global links between groups, optimizing performance. For smaller, regional tasks, jobs can be scheduled within a group of the Dragonfly+ topology, providing an optimal structure for localized or tree-friendly traffic. This design emphasizes balanced communication and robustness, making it suitable for high-complexity environments.




Figure 8. Four Group Dragonfly+ Network Topology

The Dragonfly+ group itself is structured in a Fat Tree topology. Fat Tree topologies are ideal for smaller clusters where workload is more localized or up/down friendly, and the complexities of managing multiple groups of a Dragonfly+ are not required. They ensure a flat communication hierarchy with balanced IO across the whole network. The simplicity of Fat Tree topology makes it easier to configure and maintain, providing a cost-effective solution for smaller-scale deployments. Additionally, Fat Tree networks offer predictable performance, which is crucial for applications with consistent and high throughput requirements.

Both topologies provide a solid network interconnect foundation for weather and climate solutions, addressing the needs for scalability, high bandwidth, and low latency, ensuring robust performance for complex simulations and models.

Performance

The performance characteristics of weather and climate workloads are generally balanced between compute-intensive and memory-intensive demands. *Compute-intensive* refers to workloads that benefit from a high number of processor cores, elevated CPU frequencies and increased instructions per cycle (IPC) capabilities. In contrast, *memory-intensive* workloads rely on high memory bandwidth and are particularly sensitive to the rate at which data is read from and written to memory. A large cache can also enhance performance for memory-bound workloads.

The degree to which a workload leans toward compute or memory intensity depends on factors such as the specific application, model resolution, and the physics being modeled. Higher-resolution models tend to be more memory-intensive, requiring increased memory bandwidth to run simulations efficiently. Conversely, lower-resolution models are often more compute-intensive and benefit from CPUs with higher clock speeds that accelerate scalar and vector operations.

Among the applications we tested, MOM6 and NEMO are ocean models, while WRF and GEM are mesoscale weather models. From a compute perspective, these four workloads presented in the plot below exhibit broadly similar behavior. However, NEMO is somewhat more memory-intensive, which contributes to its relatively stronger performance in memory bandwidth-constrained environments.

The figure below with the benchmark results highlights these distinctions.




Figure 9. Weather and Climate Application Performance Comparison

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:

- 5x Lenovo Heavy Duty 48U Rack Cabinets
- 12x Lenovo ThinkSystem N1380 Neptune Enclosures
- 96x Lenovo ThinkSystem SC750 V4 Compute Dual Node Trays
- 3x QM9790 Quantum NDR InfiniBand Leaf Switches
- 16x QM9790 Quantum NDR InfiniBand Spine Switches
- 4x 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 QM9790 Quantum NDR InfiniBand Spine 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	12
BYKR	ThinkSystem N1380 Neptune Enclosure Midplate Assembly	1	12
C4KW	0.95M, 63A 240-415V, 3-Phase Y-Splitter Floor Power Cable	2	24
BYKV	2.8M, 63A 240-415V, 3-Phase WYE IEC to Y-Splitter Rack Power Cable	2	24
BE0E	N+N Redundancy With Over-Subscription	1	12
BYKK	ThinkSystem N1380 Neptune EPDM Hose Connection	1	12
BYKJ	ThinkSystem N1380 Neptune System Management Module V3	1	12
BYJZ	ThinkSystem N1380 Neptune Enclosure	1	12
BYKH	ThinkSystem N1380 Neptune 15kW 3-Phase 200-480V Titanium Power Conversion Station	4	48
5WS7C20194	5Yr Premier 24x7 4Hr Resp N1380 Neptune Enclosure	1	12

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	96
BZ7E	ThinkSystem SC750 V4 ConnectX-7 Auxiliary Cable	1	96
5977	Select Storage devices - no configured RAID required	1	96
BPKR	TPM 2.0	1	96
C2WQ	Intel Xeon 6980P 128C 500W 2.0GHz Processor	4	384
BKSP	ThinkSystem NVIDIA ConnectX-7 NDR OSFP400 1-port PCIe Gen5 x16 InfiniBand Adapter (SharedIO) DWC	1	96
C0TY	ThinkSystem 32GB TruDDR5 8800MHz (2Rx8) MRDIMM	48	4608
BYK4	ThinkSystem SC750 V4 Neptune Tray	1	96
B7Y0	Enable IPMI-over-LAN	1	96
5WS7C20199	5Yr Premier 24x7 4Hr Resp SC750 V4 Neptune Tray	1	96

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	3
BP64	NVIDIA QM9790 64-Port Unmanaged Quantum NDR InfiniBand Switch (oPSE)	1	3
BRQ6	2.8m, 10A/100-250V, C15 to C14 Jumper Cord	2	6
BQJD	ThinkSystem NDRx2 OSFP800 IB Multi ModeTwin-Transceiver	16	48
BQJN	Lenovo 3M NVIDIA NDR Multi Mode Optical Cable	4	12
BQJR	Lenovo 10M NVIDIA NDR Multi Mode Optical Cable	14	42
BQJS	Lenovo 20M NVIDIA NDR Multi Mode Optical Cable	14	42
BQJX	Lenovo 2M NVIDIA NDRx2 OSFP800 to 2x NDR OSFP400 Passive Copper Splitter Cable	8	24
BQJY	Lenovo 3M NVIDIA NDRx2 OSFP800 to 2x NDR OSFP400 Passive Copper Splitter Cable	8	24
BP66	NVIDIA QM97xx Enterprise Rack Mount Kit	1	3
5WS7B96633	5Yr Premier 24x7 4Hr Resp NVID QM9790 oPSE	1	3

NVIDIA QM9790 Quantum NDR InfiniBand Spine Switches

Table 8. NVIDIA QM9790 Quantum NDR InfiniBand Spine Switches

Part number	Description	Quantity per system	Total quantity
0724HED	NVIDIA QM9790 64-Port Unmanaged Quantum NDR InfiniBand Switch (oPSE)	1	16
BP64	NVIDIA QM9790 64-Port Unmanaged Quantum NDR InfiniBand Switch (oPSE)	1	16
BRQ6	2.8m, 10A/100-250V, C15 to C14 Jumper Cord	2	32
BQJD	ThinkSystem NDRx2 OSFP800 IB Multi ModeTwin-Transceiver	32	512
BP66	NVIDIA QM97xx Enterprise Rack Mount Kit	1	16
5WS7B96633	5Yr Premier 24x7 4Hr Resp NVID QM9790 oPSE	1	16

NVIDIA SN2201 Gigabit Ethernet Leaf Switches

Table 9. NVIDIA SN2201 Gigabit Ethernet Leaf Switches

Part number	Description	Quantity per system	Total quantity
7D5FCTOGWW	Nvidia SN2201 1GbE Managed Switch with Cumulus (oPSE)	1	4
BPC8	Nvidia SN2201 1GbE Managed Switch with Cumulus (oPSE)	1	4
6201	1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	2	8
3793	3m Yellow Cat5e Cable	1	4
B306	Mellanox QSA 100G to 25G Cable Adapter	2	8
BFH2	Lenovo 25Gb SR SFP28 Ethernet Transceiver	2	8
BSNA	NVIDIA SN2201 Enterprise Rack Mount Kit for Recessed Mounting	1	4
5WS7B98278	5Yr Premier 24x7 4Hr Resp NVID SN2201 PSE	1	4

Lenovo Heavy Duty 48U Rack Cabinet

Table 10. Lenovo Heavy Duty 48U Rack Cabinet

Part number	Description	Quantity per system	Total quantity
1410P48	Lenovo EveryScale 48U Pearl Heavy Duty Rack Cabinet	1	5
BJ64	Lenovo EveryScale 48U Pearl Heavy Duty Rack Cabinet	1	5
BJPC	Side Panel Right Installation	1	5
BJ2N	Front Installation of 180mm Extension Kit	1	5
BJPB	Side Panel Left Installation	1	5
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	5
BJ68	ThinkSystem 48U Pearl Heavy Duty Rack Side Panel	2	10
BJ6A	ThinkSystem 48U Pearl Heavy Duty Rack Extension	2	10
BJ67	ThinkSystem 48U Pearl Heavy Duty Rack Rear Door	1	5
2304	Integration Prep	1	5
2310	Solution Specific Test	1	5
AU8K	LeROM Validation	1	5
B1EQ	Network Verification	1	5
5WS7B96703	5Yr Premier 24x7 4Hr Resp EveryScale 48U Rack	1	5

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, LP2220, was created or updated on May 16, 2025.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP2220>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP2220>.

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.

GDPS® is a trademark of IBM in the United States, other countries, or both.

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