

Lenovo EveryScale OVX Solution for NVIDIA Omniverse Solution Brief

Accelerated Adoption of Digital Twins with NVIDIA OVX

A digital twin is a highly virtual model that is an exact representation, or "twin", of a physical object or business process. Digital twins can be the virtual mirror of entire buildings, factories, distribution centers, cities, or even the whole planet.

Designing and conducting simulations within a digital twin delivers optimized business processes and improved results to companies that deploy them. The [Accenture Technology Trends 2022](#) report forecasts the worldwide market for digital twin platforms to grow from \$3.2 billion in 2020 to \$184.5 billion by 2030. This shows the competitive disadvantage to companies not taking this crucial next step in their Intelligent Transformation.

Before you can deploy any new device or autonomous agent in a real-world environment, organizations need to understand how they will behave in the physical world. This can be achieved through first simulating the change in a digital twin. Digital twins must be physically accurate representations of their real-world counterparts. True-to-life simulations require precision timing, ensuring synchronization across the physical and digital worlds with interactions occurring within the same real-time space for all actors in the simulation.

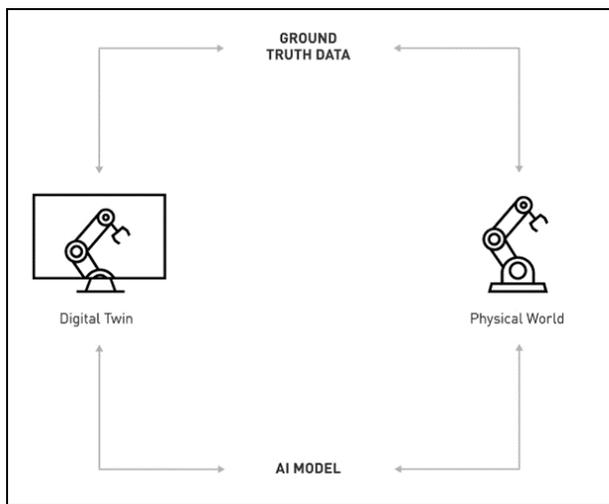


Figure 1. The relation of the digital twin to the physical world

Creating an industrial-scale digital twin simulation also involves capturing and modeling data from large numbers of machines, robots, sensors, autonomous vehicles, people, and processes. Building these sophisticated environments and accurately processing and modeling the data generated by the multitude of interactions of these components requires powerful computing solutions.

NVIDIA OVX is purpose-built to power large-scale industrial [digital twins](#) from the data center to create and operate massively complex models and true-to-reality simulation environments in real time leveraging [NVIDIA Omniverse](#) platform.

Removing OVX complexity for Lenovo customers

Lenovo partners with NVIDIA to bring our joint customers innovative solutions and infrastructures to support their Intelligent Transformations with smarter technology for all. This means equipping businesses across industries with the instruments to make data-driven decisions and gain faster insight from the edge to the core, achieving a competitive advantage.

Lenovo EveryScale OVX is a result of that deep partnership providing an architecture that is optimized to deliver extreme performance, scalability, and low latency to meet the demands of digital twins within the NVIDIA Omniverse. Lenovo's integrated solution approach removes the complexity in design, procurement, and testing of various infrastructure components providing predictable performance and solution-level interoperability support.

[Lenovo EveryScale](#) is the solution-framework Lenovo EveryScale OVX is built upon. Through solution-level interoperability testing Lenovo can warrant a fully end-to-end supported environment based on proven best practices while still tailoring it exactly to the customer's needs. That means that the infrastructure is not just supported on a component break and fix or "box"-level, but with a holistic perspective including software, firmware and even firmware-settings.

In addition to the interoperability solution support, Lenovo EveryScale is being deployed fully pre-integrated, pre-cabled, pre-loaded with the latest firmware "best recipe" to ensure a fast and reliable delivery and installation resulting in best-in-class time to solution.

Lenovo Services and NVIDIA Services provide the necessary experience to bring your Digital Twin environment to life partnering with your business lines and potential business consulting partners to fully incorporate the NVIDIA Omniverse into your business-verse.

OVX digital twin example use cases

The application of a digital twin is only limited by the reality it represents. The following two use cases give an idea of the power and value that a OVX digital twin can drive. Other real-life examples can be found on the [NVIDIA Digital Twins page](#).

Reducing planned and unplanned downtime for plant maintenance

Whether it is a turnaround in a refinery or a maintenance event in a power production plant (hydroelectric, nuclear, etc.) any interruption of production is extremely expensive. Leveraging a digital twin for planned and unplanned downtimes can massively reduce downtime duration resulting in \$ millions savings for the operator. An Omniverse digital twin could operate as a safe cowork space for engineers, contractors, and operators alike allowing lifts to be planned, measurements to be taken, and equipment to be designed.

Evolve and optimize manufacturing and supply chain

Operating a manufacturing environment does not end at the walls of the factory floor but requires a seamless engine from supply coming in over product being built to shipping out. A digital twin can realize the whole process from digital humans and trucks as well as autonomous robots that obey the laws of physics, it can create a perfect replication of reality. Through that digital twin the warehouse design and flow can be optimized, the intelligent robots can be trained, and overall productivity can be improved

OVX Software platform

NVIDIA Omniverse Software

NVIDIA Omniverse™ Enterprise is an end-to-end collaboration and simulation platform that fundamentally transforms complex design workflows, creating a more harmonious environment for creative teams.

A subscription gives access to all of the following Omniverse components

- **Omniverse Nucleus:** Collaboration engine which enables a variety of Omniverse Enterprise-enabled client applications (Apps, Connectors, and others) to share and modify authoritative representations of virtual worlds in Universal Scene Description (OpenUSD).
- **Omniverse Connectors:** Omniverse Connectors open the portals that allow popular content creation tools to connect to the Omniverse platform and save OpenUSD content. With Omniverse Connectors, users continue to work in their favorite industry 3D, CAD, and simulation software tools.
- **Omniverse Kit:** Toolkit and SDK for developing OpenUSD-based tools and applications.
- **Omniverse RTX Renderer:** An advanced, multi-GPU scalable, AI-enabled renderer based on NVIDIA RTX that enables both real-time ray tracing and referenced path tracing.
- **Omniverse Simulation:** Powerful suite of tools and SDKs that simulate a physically accurate world, including PhysX for rigid and soft body dynamics, combustible fluid, smoke and fire, and fracture and destruction.
- **Omniverse Launcher:** Easy-to-use client tool to download, install, and update Omniverse Apps, Extensions, and Connectors. Omniverse Launcher is provided in version for single users (Omniverse Launcher Workstation) and for IT-managed environments (Omniverse Enterprise Launcher).
- **Omniverse Extensions:** Building blocks of Omniverse. Extensions are built on Omniverse Kit SDK. Collections of extensions packaged together form Omniverse applications. Over 300 Omniverse Extensions are offered with complete source code to help developers easily create, add, and modify the tools and workflows they need to be productive.

The NVIDIA Omniverse Enterprise software licenses can be configured and ordered together with the Lenovo EveryScale OVX solution. For more information on NVIDIA Omniverse Enterprise licensing [click here](#).

Table 1. Features for NVIDIA Omniverse Enterprise

Feature	NVIDIA Omniverse Enterprise
COLLABORATION	Between multiple apps and licensed Creators
LICENSING	Per GPU Subscription Licenses, Annual and Multi-Year
SUPPORT	NVIDIA Enterprise Support
NUCLEUS	Enterprise Nucleus Server, Nucleus Workstation
CONNECTORS	Use of all production Connectors
ENTERPRISE SUPPORTED FOUNDATION APPLICATIONS	USD Composer, USD Presenter, USD Explorer Code

Lenovo EveryScale OVX solution also leverages RedHat Enterprise Linux and Ubuntu as OS and Lenovo XClarity family to simplify and automate the deployment and management of the solution.

To run the Lenovo EveryScale OVX solution as a fully virtualized environment Lenovo and NVIDIA partner with VMware leveraging vSphere.

OVX Infrastructure solution

The Lenovo EveryScale OVX solution is a purpose-built NVIDIA Omniverse computing offering based on Lenovo Scalable Infrastructure integrating Lenovo ThinkSystem servers and designed to meet the demands of large-scale digital twins.

Lenovo EveryScale 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.

The Lenovo EveryScale OVX computing solutions architecture is designed to scale leveraging NVIDIA Spectrum Ethernet switches and, when combined with NVIDIA Omniverse Enterprise, delivers a fully integrated solution that transforms workflows for digital twins.

As a result, despite the scale and complexity of an NVIDIA Omniverse digital twin solution, customers can focus their efforts on maximizing their business value, instead of consuming valuable resources to design, optimize, install, and support the infrastructure. Again, this is all thanks to Lenovo Scalable Infrastructure and the engineering partnership between Lenovo and NVIDIA.

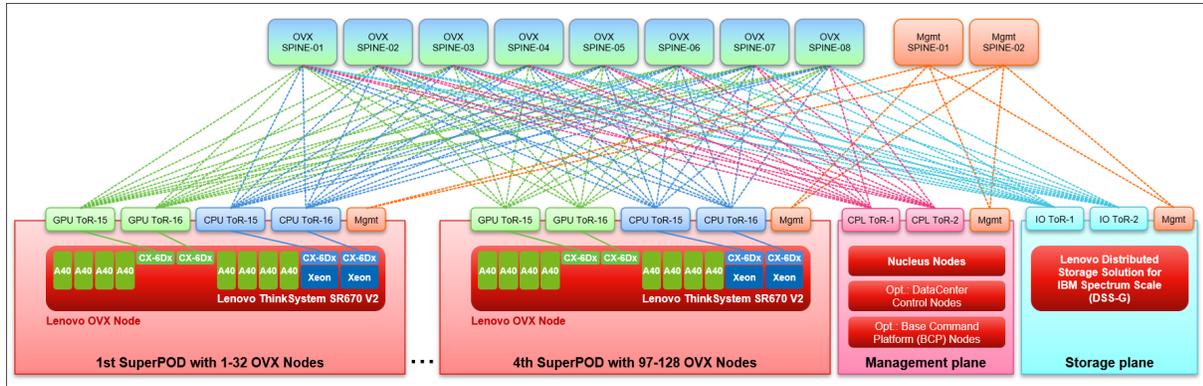


Figure 2. Lenovo EveryScale OVX Super-POD solution design

Scalable Performance

The Lenovo EveryScale OVX solution is designed for multi-node scalability with NVIDIA Spectrum Ethernet switch fabric. All POD and SuperPOD configurations are integrated with an optimized combination of network fabric, storage, and enterprise-grade software which combine to deliver unprecedented performance for the most demanding workloads. The OVX POD based architecture enables the deployment of one or more Scalable Units, providing the compute, low-latency networking, bandwidth, and performance required for massively complex simulations and real-time digital twins at the factory, city, or planetary-scale.

Table 2. Lenovo EveryScale OVX Size definitions

OVX Solution Size	Definition
Lenovo Everyscale OVX Node (CTO:7Z23OVX1WW, 7D9ROVX3WW)	<ul style="list-style-type: none"> • The foundational building block of the Lenovo EveryScale OVX solution. A Lenovo SR670 V2 or SR675 V3 configured as listed in the OVX Server platform section. • Capable of Initial POC. • Can be used for single-node Omniverse Enterprise Workloads.
Lenovo Everyscale OVX POD	<ul style="list-style-type: none"> • Lenovo EveryScale OVX PODs start at 4 node configurations and can scale up to 32 nodes. • Typical configurations include the Base-POD containing 4 OVX nodes and the Quad-POD containing 16 OVX nodes. • Customers can start with Base-POD and scale as needed. • Spectrum 2 or Spectrum 3 network fabric.
Lenovo Everyscale OVX SuperPOD	<ul style="list-style-type: none"> • Lenovo EveryScale OVX SuperPODs consist of at least 1 scalable unit (SU) of 32 Lenovo Everyscale OVX nodes. • Optimized infrastructure for large-scale, digital twin simulations. • Spectrum3 network fabric, storage, and orchestration software.

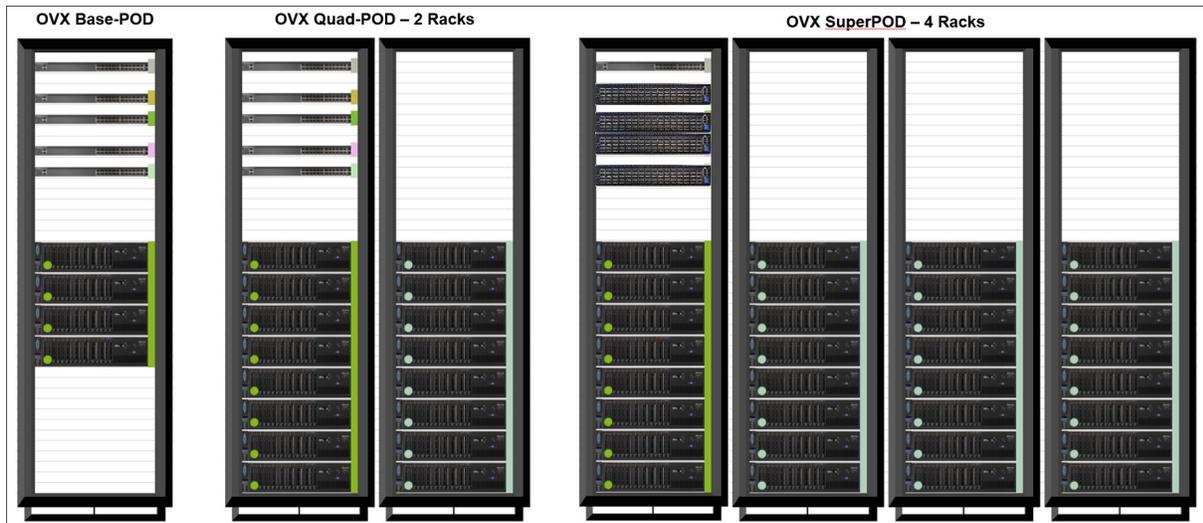


Figure 3. Lenovo EveryScale OVX scalable POD and SuperPOD rack view examples

OVX Server platform

Lenovo OVX node

The building block of the OVX computing node, the Lenovo ThinkSystem SR670 V2 server provides best-in-class cooling for the accelerators positioned in the front and allows both for front and rear IO connectivity for maximum graphic performance and IO throughput.

The OVX 1.0 node configuration combines eight [NVIDIA A40](#) GPUs, 3rd Gen Intel Xeon Scalable Processors, super-fast NVMe storage, and [NVIDIA ConnectX-6 Dx](#) network adapters with enterprise-grade management and orchestration software.

The OVX 2.0 node configuration replaces the eight NVIDIA A40s with the NVIDIA L40 or L40S GPUs and upgrades the networking card from a ConnectX-6 Dx to an NVIDIA ConnectX-7 adapter.

The new OVX 3.0 node configuration is built on the new SR675 V3 which leverages PCIe gen 5 and the latest AMD EPYC processors. It is outfitted with 4 or 8 NVIDIA L40S GPUs, CX7 or B3140 Bluefield-3 for East-West networking and finally a B3220 bluefield-3 for the North-South networking.

For more information regarding the Lenovo ThinkSystem SR670 V2 and SR675 V3 in general, see the [SR670 V2 product guide](#) and the [SR675 V3 product guide](#)

Below you find the detailed configuration of the Lenovo OVX node. The table lists the components for an OVX 1.0, OVX 2.0, and OVX 3.0 compute node.

Table 3. Lenovo OVX Node Configuration (CTO: 7Z23OVX1WW for OVX 1.0/2.0 and 7D9ROVX3WW for OVX3.0)

Configuration	OVX 1.0 (7Z23OVX1WW)	OVX 2.0 (7Z23OVX1WW)	OVX 2.0 "Cloud" (7Z23OVX1WW)	OVX 3.0 8GPU (7D9ROVX3WW)	OVX 3.0 4GPU (7D9ROVX3WW)
System	SR670 V2 8DW PCIe GPU Base, BFCW	SR670 V2 8DW PCIe GPU Base, BFCW	SR670 V2 8DW PCIe GPU Base, BFCW	SR675 V3 8DW PCIe GPU Base, BR7F	SR675 V3 4DW PCIe GPU Base, BR7G
Riser Config	GPUDirect RDMA Riser, BFGK	GPUDirect RDMA Riser, BFGK	GPUDirect RDMA Riser, BFGK	GPUDirect RDMA Riser, BR7S	GPUDirect RDMA Riser, BR7S
GPU	8x NVIDIA A40, BQZQ	8x NVIDIA L40, BT87	8x NVIDIA L40, BT87	8x NVIDIA L40, BT87 or L40S, BYFH	4x NVIDIA L40, BT87 or L40S, BYFH
Processor	2x Intel 8362 (32C 2.8GHz), BKDC	2x Intel 8362 (32C 2.8GHz), BKDC	2x Intel 8362 (32C 2.8GHz), BKDC	2x AMD 9354 (32C 3.25GHz), BR30 2x AMD 9554 (64C 3.1GHz), BPVJ 2x AMD 9654 (96C 2.4GHz), BPVK	2x AMD 9354 (32C 3.25GHz), BR30
Memory	32 x 32GB DDR4, B964	32 x 32GB DDR4, B964	16 x 64GB DDR4, B966	16 / 24x 32GB, BQ37 16 / 24x 64GB DDR5, BQ3D	16 / 24x 32GB, BQ37 16 / 24x 64GB DDR5, BQ3D
OS Storage	2x M.2 960GB PCIe 4.0, BKSR	2x M.2 960GB PCIe 4.0, BKSR			
Mgmt / IPMI	Optional	Optional	1x 4P 1GbE OCP (5719), B5T1	1x 4P 1GbE OCP (I350), B93E	1x 4P 1GbE OCP (I350), B93E
CPU Network	2x CX6 Dx 1P 200GbE, BQX9	1x CX7 2P VPI 200GbE, BQBN	1x BF2 2P 100GbE, VLH (4XC7A87899)	1x BF3 B3220 2P VPI 200G, BVBG	1x BF3 B3220 2P VPI 200G, BVBG
GPU Network (one to each PCIe switch)	2x CX6 Dx 1P 200GbE, BQX9	2x CX7 2P VPI 200GbE, BQBN	2x CX7 2P VPI NDR200, BQBN	4x CX7 2P VPI 200GbE/NDR200, BQBN or 1P NDR, BQ1N (future / 4x BF3 HPC B3140)	2x CX7 2P VPI 200GbE/NDR200, BQBN or 1P NDR, BQ1N (future / 2x BF3 HPC B3140)
PDUs	4x 2400W Plat., B962 / 2600W Tit., BKTJ	4x 2400W Plat., B962 / 2600W Tit., BKTJ	4x 2400W Plat., B962 / 2600W Tit., BKTJ	4x 2600W Titanium, BKTJ	4x 2600W Tit., BKTJ / (future / ITIC)

The following figure shows the architectural block diagram of the OVX 3.0 node.

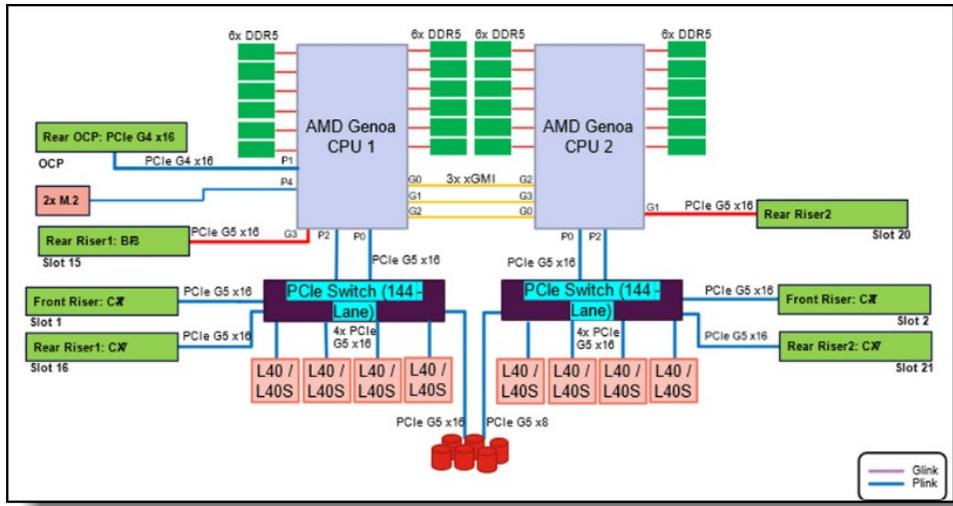


Figure 4. OVX 3.0 node – 8 GPU configuration block diagram (click to show larger version)

Lenovo Nucleus node

Nucleus is the collaborative user environment of NVIDIA Omniverse providing persistent data and scalable core microservices. Depending on the number of users in the Omniverse Environment the Nucleus Server configuration must fit those needs.

The Lenovo ThinkSystem SR655 V3 has been designed to combine highest performance and flexibility. It allows to support up to 500 users without live editing and be scaled out as requirements grow.



Figure 5. Nucleus node - Lenovo ThinkSystem SR655 V3

For more information regarding the Lenovo ThinkSystem SR655 V3 in general, [see the SR655 V3 product guide](#).

Below you can find the detailed example configuration of the Lenovo Nucleus node:

Table 4. Lenovo Nucleus node configuration

For Factor	Nucleus Server
System	Lenovo ThinkServer SR655 V3
CPU	1x AMD Genoa 16 Cores
System Memory	96GB ECC DDR5
Boot drive	1x 512GB M.2 NVMe SSD
Data drive	2x 1TB M.2 NVMe SSD
Networking	2x CX7 (2x200GB)
GPU	Not Applicable

Depending on concurrent users and targeted capabilities like live editing the optimal configuration choice can vary. A customer POC is recommended to optimize further for each unique deployment.

OVX Network platform

Digital twin simulations involve multiple autonomous systems interacting in the same space-time. To unlock the benefits of a digital twin and enable advanced visualization and simulation at scale, the architecture must be able to accurately simulate interactions with precise time synchronization across digital and physical worlds.

The Lenovo EveryScale OVX solution includes advanced networking technologies that enable that precise simulation of space and time. The NVIDIA Spectrum Ethernet Switch fabric and Nvidia Smart NICs deliver the necessary bandwidth, low latency, and precision timing to enable real-time accuracy and network synchronization, enhancing simulation capabilities.

NVIDIA Spectrum Ethernet Switches

The [NVIDIA Spectrum-2](#) SN3700 Ethernet switch offers 32 ports of 200GbE in a compact 1U form factor. It enables connectivity to endpoints at different speeds and carries a throughput of 12.8Tb/s, with an 8.33Bpps processing capacity and only 425ns latency. It is best used in small to medium sized OVX POD configurations.



Figure 6. OVX leaf switch - NVIDIA Spectrum-2 SN3700V Ethernet switch

The [NVIDIA Spectrum-3](#) SN4600 Ethernet switch offers 64 ports of 200GbE in a 2U form factor doubling the networking capacity over the SN3700. This switch is best used as central spine layer for the overall OVX solution or in large sized OVX SuperPOD configurations.



Figure 7. OVX spine switch - NVIDIA Spectrum-3 SN4600 Ethernet switch

ConnectX-6 Dx Network Adapter

The NVIDIA ConnectX-6 Dx SmartNIC is the industry's most secure and advanced datacenter network interface card to accelerate mission-critical applications, such as security, virtualization, SDN/NFV, big data, machine learning, and storage.

Its Advanced Timing and Synchronization capabilities with PTP hardware clock, PTP-based packet pacing, and nanosecond-level accuracy make it the ideal choice for OVX real-time accurate Digital Twin simulation. The ConnectX-6 Dx 1-port 200GbE PCIe Adapter provides a single-port of 200Gb/s Ethernet connectivity.

ConnectX-7 Network Adapter

The NVIDIA ConnectX-7 SmartNIC provides hardware-accelerated networking, storage, security, and manageability services at data center scale for cloud, telecommunications, AI, and enterprise workloads. ConnectX-7 empowers agile and high-performance networking solutions with features such as Accelerated Switching and Packet Processing (ASAP²), advanced RoCE, GPUDirect Storage, and in-line hardware acceleration for Transport Layer Security (TLS), IP Security (IPsec), and MAC Security (MACsec) encryption and decryption. ConnectX-7 enables organizations to meet their current and future networking needs in both high-bandwidth and high-density environments.

ThinkSystem NVIDIA BlueField-3 QSFP112 2-Port 200Gb DPU Adapter

The ThinkSystem NVIDIA BlueField-3 QSFP112 2-Port 200Gb DPU Adapter is an advanced compute platform that enables software-defined, hardware-accelerated IT infrastructures. With 16 Arm cores, 32GB of DDR5 memory, and two ports of 200Gb/s Ethernet or NDR InfiniBand network connectivity, BlueField-3 data processing unit (DPU) offloads, accelerates, and isolates software-defined networking, storage, security, and management functions.

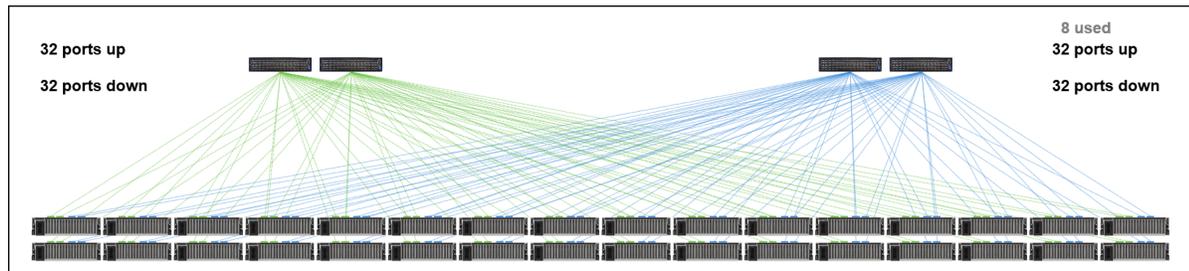


Figure 8. OVX Super-POD network topology diagram (click to show larger version)

NEW: Now quoting NVIDIA Spectrum-X networking platform for OVX

The NVIDIA Spectrum™-X Ethernet platform is designed specifically to improve the performance and efficiency of Ethernet-based AI clouds. This breakthrough technology achieves 1.6X better AI networking performance, along with consistent, predictable performance in multi-tenant environments. Spectrum-X is built on network innovations powered by the tight coupling of the NVIDIA Spectrum-4 Ethernet switch and NVIDIA® BlueField®-3 SuperNIC. Spectrum-X network optimizations reduce runtimes of massive transformer-based generative AI models and deliver faster time to insight.

NVIDIA Spectrum-X Networking Platform:

The Spectrum-X networking platform consists of the following components:

The Spectrum-4 Ethernet switch for smart-leaf, spine, and super-spine designs offers 64 ports of 800GbE in a dense 2U form factor, playing a vital role in NVIDIA Spectrum-X deployments. The SN5600 switch supports both standard leaf and spine designs with top-of-rack switches and rail-optimized end-of-row topologies. The SN5600 offers diverse connectivity in combinations of 10–800GbE and boasts an industry-leading total throughput of 51.2 terabits per second (Tb/s).

The BlueField-3 B3140 SuperNIC delivers up to 400GbE RoCE network connectivity between GPU servers and enables NVIDIA GPUDirect® RoCE for optimizing peak AI workload efficiency. The BlueField-3 SuperNIC accelerates AI networking and performs low-compute tasks. It is best suited for the B3140H model, providing an HHHH form-factor, and a sub-75 watt power envelope, making it an ideal choice for NVIDIA Spectrum-X deployments. NVIDIA integrates BlueField-3 SuperNICs across its data center computing systems, including the HGX™ H100 and OVX™ L40S-based systems.

LinkX transceivers and cables provide a robust selection of direct attach copper cables (DACs), active copper cables (ACCs), active optical cables (AOCs), and optics needed to connect octal small form-factor pluggable (OSFP)- and quad small form-factor pluggable (QSFP)-based fabrics. Spectrum-4 switches, combined with the LinkX portfolio, enable the most power-efficient AI fabric on the market.

Learn more about the NVIDIA Spectrum-X Networking Platform at: nvidia.com/en-us/networking/spectrumx

OVX Storage platform

NVIDIA Omniverse solutions are operating mostly with local storage on the OVX nodes. They do require a high performance and high reliable network file system though to read the initial data from and optionally execute for example continuous snapshots or backups to.

[Lenovo Distributed Storage Solution for IBM Spectrum Scale \(DSS-G\)](#) is the recommended Storage environment for the Lenovo EveryScale OVX solution. It is a software-defined storage (SDS) solution for dense scalable file and object storage suitable for high-performance and data-intensive environments. DSS-G combines the performance of Lenovo ThinkSystem servers, Lenovo storage enclosures, and industry leading IBM Spectrum Scale software, to offer a high performance, scalable building block approach to modern storage needs.

IBM Spectrum Scale, based on IBM General Parallel File System (GPFS) technology, is a high-performance and highly scalable parallel file system with an extensive suite of enterprise class data management features.

Lenovo is a strategic alliance partner of IBM, and combines IBM Spectrum Scale software with Lenovo servers, storage and networking components for integrated and customized solutions.

IBM Spectrum Scale offers access to a single file system or set of filesystems in a global namespace with shared file system access among IBM Spectrum Scale clusters, simultaneous file access from multiple nodes, high recoverability and data availability through replication, the ability to make changes while a file system is mounted, and simplified administration even in large environments.

OVX Customer Journey

The following steps are the OVX customer journey:

Step One – Proof Of Concept

Engage with Lenovo and NVIDIA on early-stage opportunities to qualify use cases and develop sizing recommendations:

- Initial POC on a single node (Single node CTO:7Z23OVX1WW or 7D9ROVX3WW)
- Virtual Machine or Bare-Metal
- Single-node Omniverse Enterprise Workloads

Step Two – Expansion

Seamlessly scale to POD and SuperPOD configurations:

- Add Networking, Infrastructure Management & Orchestration SW
- Multi-Node Omniverse Enterprise Workloads

Contact your Lenovo or NVIDIA sales representative to get started.

Lenovo Workstation and Displays

Enabling teams scattered across geographic locations to collaborate on 3D assets with various machines is a task that enterprise IT needs a solution for. Remote workforces demand efficiency, and this is a difficult balancing act for power-hungry workloads and a dispersed team. The demand for photorealistic, physically-accurate simulation begs higher compute power, accessible from anywhere.

Lenovo’s NVIDIA RTX powered ThinkPad, ThinkStation P Series Workstations & ThinkVision P series displays are tested and optimized to run NVIDIA’s Omniverse enterprise.

These workstations and displays are the perfect addition to an OVX deployment and can run the smaller than OVX workloads (known as OVE workloads) on their own.

Table 5. Lenovo OVE Workstations configurations

Form Factor	Mobile Workstation	Desktop Workstation
System	Lenovo ThinkStation P16, MTM: 21FA002EUS	Lenovo ThinkStation P7, MTM: 30F2/30F3xxCTOWW
CPU	Intel Core i9-13850HX	Intel Xeon W9-3475X 36C / 4.8GHz
System Memory	64GB (2x 32GB) DDR5	256GB (8x 32GB) ECC DDR5
Boot drive	1x 1TB M.2 NVMe SSD	1x 2TB M.2 NVMe SSD (or above)
Data drive	Not Required	2x 2TB M.2 NVMe SSDs
Networking	WiFi 6E	10GbE NIC + NVIDIA CX6 DX Active
GPU	1x NVIDIA RTX 5000 Ada Mobile	1x NVIDIA RTX 6000 Ada (Max 3x Supported)

ThinkVision P27h, P32p and P40 Displays

ThinkVision P series displays provide:

- 4K UHD (3840 x 2160) IPS
- Professional grade color with 99.5% Adobe RGB Gamut, 90% DCI-P3
- Near-edgeless display
- Pre-calibrated accuracy

ThinkPad Thunderbolt 4 Workstation Dock

- Supports 8K & up to 4x 4K Displays
- Intel vPro Remote Management
- Data Transfer Speeds up to 40Gbps
- Ethernet, USB, HDMI, DP, TB4 and more

For more information regarding Lenovo’s Omniverse Workstation offerings, see the following web page: <https://techtoday.lenovo.com/us/en/workstations/Omniverse>

Why Lenovo

Lenovo is a leading provider of data center infrastructure solutions and a premier NVIDIA OVX partner. The Lenovo EveryScale OVX solution is co-designed with NVIDIA. Lenovo will work with you and NVIDIA to identify, design, install and support the solution that best ensures your organization's needs.

Lenovo complements a portfolio of leading x86 infrastructure with a full range of storage, workstations, software, and comprehensive services that provides excellent performance, reliability, and security for your IT environment from the edge to the cloud.

One Stop Shop for OVX deployments

Lenovo is the one stop shop for your Omniverse and OVX deployments providing:

- End to end technology access from clients and workstations to edge devices and data center systems and solutions.
- Strategic partnership with NVIDIA resulting in co-designed Lenovo OVX solution.
- Only integrated NVIDIA certified OVX solution in the market - instead of plain boxes.
- 20 years track record on integrated HPC solutions with Lenovo Scalable Infrastructure
- Full solution from one provider with full control from supply chain to service and support, not assembling parts from a channel/distribution.
- Compact 3U form factor for OVX nodes.
- Global presence and availability of solution.

Integrated Software & Hardware

Optimized to power the most demanding workloads, Lenovo EveryScale OVX solution for NVIDIA Omniverse™ Enterprise provides an end-to-end collaboration and true-to-reality simulation platform that transforms complex digital twin workflows at any scale.

Enterprise Level Support

Lenovo EveryScale OVX solutions are based on Lenovo and NVIDIA-Certified Systems™, tested and validated to provide the necessary performance, manageability, security, and scalability. Comprehensive enterprise-grade support for Lenovo EveryScale OVX solutions will be provided by Lenovo and Omniverse software support will be provided directly by NVIDIA.

For details about Lenovo support, see the [Lenovo ISG Support Plan](#).

For More Information

To learn more about the Lenovo EveryScale OVX Solution for NVIDIA Omniverse, contact your Lenovo Sales Representative or Business Partner.

Products related to this document are the following:

- [Lenovo Scalable Infrastructure \(LeSI\)](#)
- [Lenovo ThinkSystem SR670 V2](#)
- [Lenovo ThinkSystem SR630 V2](#)
- [Lenovo ThinkSystem SR655 V3](#)
- [Lenovo ThinkSystem SR675 v3](#)
- [Lenovo Distributed Storage Solution for IBM Spectrum Scale \(DSS-G\)](#)
- [Lenovo High Performance Computing](#)
- [Lenovo TruScale Infrastructure Services](#)
- [Lenovo Workstations and Displays for Omniverse](#)

More information on NVIDIA Omniverse and related product families

- [NVIDIA Omniverse](#)
- [NVIDIA Omniverse Documentation](#)
- [NVIDIA A40](#)
- [NVIDIA L40](#)
- [NVIDIA ConnectX-6 Dx](#)
- [NVIDIA ConnectX-7](#)
- [NVIDIA Spectrum 3000 Series](#)
- [NVIDIA Spectrum 4000 Series](#)

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, LP1595, was created or updated on March 15, 2024.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/LP1595>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/LP1595>.

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®

Lenovo Services

ThinkPad®

ThinkServer®

ThinkStation®

ThinkSystem®

ThinkVision®

XClarity®

The following terms are trademarks of other companies:

Intel®, Intel Core™, Xeon®, Thunderbolt™, and vPro™ are trademarks of Intel Corporation or its subsidiaries.

Linux® is the trademark of Linus Torvalds in the U.S. and other countries.

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