



SUPERMICRO + NVIDIA TOTAL SOLUTION FOR SCALABLE OVX DEPLOYMENTS

Effortlessly Build and Simulate Large Virtual 3D Worlds and Execute Compute-Intensive Workloads at Data Center Scale with NVIDIA Omniverse™ Enterprise

TABLE OF CONTENTS

Executive Summary 1

Designed for Operating Large-Scale Digital Twins and Running Simulations 1

Scalable Performance and Flexible Deployment Options 2

Deploy In Days, Not Weeks 2

Reasons To Deploy Supermicro Rack Scale OVX Solutions 2

Conclusion and Summary 4



Executive Summary

Supermicro is introducing a rack-scale implementation of NVIDIA Omniverse™ Enterprise powered by the Supermicro SYS-420GP-TNR GPU servers. The Supermicro solution is an NVIDIA-Certified OVX System purpose-built to power and operate large-scale digital twins for Omniverse Enterprise. This purpose-built turnkey solution features the best-in-class architecture to deliver the cutting-edge performance needed for data-intensive applications and real-time collaboration using NVIDIA Omniverse.

Designed for Operating Large-Scale Digital Twins and Running Simulations

With 256x NVIDIA L40 GPUs and 200 Gbps networking in a scalable 32-node SuperPOD, Supermicro OVX Rack solutions deliver unparallel performance to operate digital twins and run large simulations effortlessly.

- Designed and explicitly architected to meet the computing demands for digital twins and application-intensive workloads
- Validated solution enables data-center scale and reliability for real-time collaboration and high-fidelity simulations

Scalable Performance and Flexible Deployment Options

Deploy Supermicro Rack Scale Solutions of NVIDIA Omniverse™ Enterprise with 1 to 4 nodes as a proof of concept and quickly scale to hundreds of servers via SuperPODs to meet workload demands.

- Multi-rack plug-and-play design. Easily grow the cluster as an organization's workloads increase.
- SYS-420GP-TNR provides the best-in-class performance and flexible computing architecture with dual 3rd Gen Intel® Xeon® Scalable processors and NVIDIA GPUs.
- High-performance networking enables high bandwidth and low latency for workloads simultaneously utilizing multiple systems.

Deploy In Days, Not Weeks

Fully tested and validated design at the cluster level so customers can have Omniverse Enterprise operational in days. As a result, engineers, designers, artists, and scientists can focus on fewer system implementation complexities and more technological breakthroughs.

- Optimized rack layout and topology for power and cooling
- Validated and tuned OVX architecture that is designed for stability and future scalability

Reasons To Deploy Supermicro Rack Scale OVX Solutions

- Complete NVIDIA Omniverse™ Enterprise software stack included – makes it easy to get started with the NVIDIA Omniverse platform.
- Enabled by NVIDIA RTX™ and high-speed networking technologies – Enables fast responses for optimal collaboration.
- End-to-end thoroughly tested and integrated – Racks are plug-and-play design; Supermicro assembles and extensively tests the entire configuration prior to shipping to the customer.
- Purpose built for performance and acceleration – Selected components ensure optimal performance using the latest CPUs and GPUs.
- Highly scalable and easy to deploy from four-server clusters to multiple pods to handle the largest workloads.
- Four configurations are available, from single to multiple racks, to begin building metaverse applications, run large-scale simulations, and operate digital twins.

Supermicro OVX Solutions



	Small	Medium	Large (Pod)	SuperPOD
Appliance	SRS-48UOVX-SMAL-01	SRS-48UOVX-MED-01	SRS-48UOVX-POD-01	SRS-48UOVX-SPOD-01
OVX Nodes	4x SYS-420GP-TNR	8x SYS-420GP-TNR	16x SYS-420GP-TNR	32x SYS-420GP-TNR
Nucleus Servers	1x SYS-120U-TNR		2x SYS-120U-TNR	
Rack	1x 48U (Optional 42U)	2x 48U (Optional 42U)	3x 48U (Optional 42U)	6x 48U (Optional 42U)
Total CPUs	8x Intel® Xeon® Platinum 8362	16x Intel® Xeon® Platinum 8362	32x Intel® Xeon® Platinum 8362	64x Intel® Xeon® Platinum 8362
Total GPUs	32x NVIDIA L40	64x NVIDIA L40	128x NVIDIA L40	256x NVIDIA L40
Total System Memory	4TB	8TB	16TB	32TB
Networking	6x 200Gbps 32-port NVIDIA SN3700 Ethernet Switches		8x 200Gbps 32-port NVIDIA SN3700 Ethernet Switches	8x 200Gbps 64-port NVIDIA SN6400 Ethernet Switches
	1x 1Gbps 48-port NVIDIA SN2201 Ethernet Switch			
Total Storage	60.8TB NVMe (Raw)	121.6TB NVMe (Raw)	243.2TB NVMe (Raw)	486.4TB NVMe (Raw)
Estimated Power	18.61kW	34.10kW	66.64kW	133.21kW
Software	NVIDIA Omniverse™ Enterprise Starter Pack Subscription (1, 3, 4, and 5 yr. options)			



SYS-420GP-TNR OVX Specifications	
Overview	4U Dual Processor (3rd Gen Intel® Xeon®), Dual-Root GPU System, up to 10 PCIe GPUs
CPU	2x Intel® Xeon® Platinum 8362 CPUs (32C/64T, 2.8GHz, 265W TDP)
Memory (additional memory available)	1TB DDR4 (16x 64GB DIMMS)
Graphics	8x NVIDIA L40 GPUs (48GB, 300W TDP)
Networking	3x NVIDIA ConnectX-7 200Gbps NICs
Storage (additional storage available)	15.2TB NVMe (2x 7.6TB U.2) 1.92TB NVMe (2x 960GB M.2, Boot)
Estimated Power	3,874W (3.87kW)

Conclusion and Summary

NVIDIA Omniverse™ Enterprise revolutionizes design collaboration, simulation, and operation of digital twins. Globally dispersed teams can accelerate their workflows with one-click interoperability between leading software tools and seamlessly collaborate in a shared virtual world running from the data center. Supermicro's rack-scale OVX solution for Omniverse enables the development and operations of factory digital twins, harmonizes the flow of design data across the enterprise, and accelerates visualization for in-vehicle experiences, car configurators, and virtual showrooms.

Omniverse is driving digital transformation and innovation through the following:

- A Cutting-edge solution to help your business to beat the competition
- Seamless high-speed global connectivity allows for design and collaboration at scale
- Increase productivity and faster time to market
- Resource-saving and lower overhead costs

Digital Transformation Through Omniverse

Source: BMW Group



Digital Twins



Simulations and Rendering



Design and Collaboration



Global Connectivity
and Collaboration



Increased Productivity
and Maximize Creativity



Reduced Overhead Costs
and Inefficient Workflows

For More Information, please visit:

<https://www.supermicro.com/en/accelerators/nvidia/omniverse>

<https://www.supermicro.com/en/products/gpu>

[Supermicro X13 CPU Choices \(4th Gen Intel Xeon Scalable processors\)](#)

