



# Supernano Solutions

Q1, 2024

Igor Berin  
Senior BDM EMEA



# ABOUT SUPERMICRO



Revenue **\$14B+** (FY2024 guidance)  
\$7.1B (FY2023)  
\$5.2B (FY2022)

Worldwide Presence  
6M+ Sq ft. Facilities Worldwide  
1. Silicon Valley (HQ),  
2. Taiwan,  
3. The Netherlands,  
4. Malaysia and others

Production  
\$18B/yr Production Capacity (CY23)  
Top 5 Largest Server System Provider  
Worldwide (IDC & Gartner 2022), ~1.3M  
units annually

Human Resource in 4 Campuses  
6000+ headcount Worldwide,  
~50% Technical / R&D

Key Growth Matrix  
**#1** in Generative AI and LLM Platforms  
500%+ YoY Growth in Accel. Computing

# BUSINESS FOCUS

## Building Block Solutions

- Optimized subsystem design, including Motherboard, I/O, Chassis, Power supply, Thermal solution, Firmware, BMC, Security and Management Software, at scale
- Enable industry's broadest AI/Server/Storage/Edge/IoT product portfolio
- Customer time-to-market advantage




## Rack-Scale Solutions

- Plug-and-play design enables customers to plug in network & power, and ready to go on-line
- Onsite deployment and service
- Management and security software increase value and competitive advantages



## Green Computing

- High efficiency system designs, free air/liquid cooling
- Save up to 40% TCO for many of our partners
- Industry-wide green computing adoption could save \$10B per year of energy cost



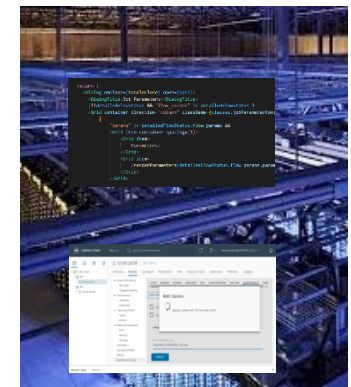
## Business Automation

- Operation and sales automation to efficiently serve enterprise and any size of business
- Command Center to provide 24/7 security and support



## Supermicro 4.0 (Coming Soon)

- Green computing everywhere
- Enhanced system/cloud software/application products
- New features and products for DC/Cloud total solutions
- More details later



# GLOBAL PRESENCE



## Production scale and cost optimization Economy of Scale and Cost

- **Silicon Valley Green Computing Park B20-B23**
  - Rack-Scale Integration (Liquid Cooling)
  - Command Center
  - B2B/C Programs
  - Cloud Services
- **APAC Science and Tech Center B62**
  - Increase 2X-3X APAC capacity in FY24, 25, 26
- **Supermicro Malaysia Campus**
  - High Volume Subsystem and Rack-Scale Production by Q3 2024

### Future Site Plans


- Additional Silicon Valley locations
- Additional Netherlands facility
- Mexico, Texas sites (in plan)




# Supermicro CPU Vendors



**X13**



**H13**



**R12**



NVIDIA Grace CPU

**R13**

# Industry's Most Comprehensive Portfolio



**Hyper-E and Hyper**  
Best-in-Class Performance and Flexibility Rackmount Servers



**Ultra and Ultra-E**  
High Performance & Flexibility Rackmount Systems for Enterprise Applications



**Cloud DC**  
All-in-one Rackmount Platforms for Cloud Data Centers



**WIO (UP)**  
Industry's Widest Variety of I/O Optimized Servers



**Mainstream**  
Versatile Entry Level and Volume Servers for Enterprise Applications



**BigTwin®**  
Highly Modular Multi-Node Systems with Tool-less Design



**TwinPro®**  
Cost-effective 2U Multi-node Platforms



**FatTwin®**  
Advanced Multi-node 4U Twin Architecture with 8 or 4 Nodes



**SuperBlade®**  
High Density x86 Multi-node Server for Enterprise Cloud, HPC



**SuperWorkstations**  
Workstations for High Performance Workloads



**MP 4-Way Server**  
Highest Performance and Flexibility for Enterprise Applications



**PCIe GPU Servers**  
High Density Systems for Double-width, Full Length PCIe GPUs



**HGX GPU Servers**  
High Performance and Flexibility with Advanced Architecture and Thermal Design



**SuperStorage®**  
Top-loading Server Optimized for Field Serviceability and Field Replacement




**IOT/Embedded**  
High-efficiency, High-performance Compact Form Factor for 5G and Edge computing

# Supermicro Rack Integration Services provides a “one-stop-shop” for your data center needs

## Optimized and Lab Tested Components for Superior Performance

- Turn-Key Data Center
- Accelerate Your Deployment
- Professional Rack Level Design
- Validation and Benchmarking



- Server
- Storage
- Network
- Software
- Cabling
- Power and Cooling
- Testing
- Benchmarking
- Full Rack Burn-in

# Direct Liquid Cooling (DLC)



1<sup>st</sup> Cycle

Facility to CDU

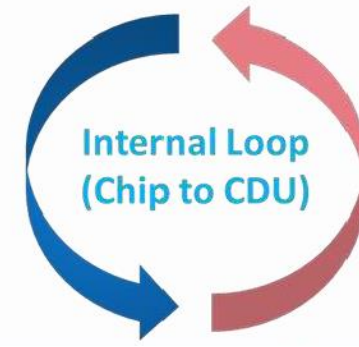


2<sup>nd</sup> Cycle

CDU to SERVER



CDU: Cooling Distribution Unit  
(45C facility water supported)






# Typical Industry's Workloads



GPU: AI, ML, Metaverse, Content Creation



SDS Storage



In Mem Data Base



IoT, O-RAN, Edge



Gaming, Video Surveillance/Analytics

# Hyperconverged



# Accelerate Everything

GPU Optimized Systems to Achieve 5X, 10X,... 100X Performance



## Large Scale AI Training Workloads

Large language models, Generative AI training, autonomous driving, robotics



## HPC/AI Workloads

Engineering simulation, scientific research, genomic sequencing, drug discovery



## Enterprise AI Inference & Training

AI-enabled services/applications, chatbots, business automation



## Visualization and Design

Graphical content development and automatic generation, digital twins, 3D collaboration



## Content Delivery and Virtualization

Content delivery networks (CDNs), video transcoding, live streaming, VDI



## AI Edge

Retail automation, manufacturing/logistics automation, medical diagnosis/predictive care, security, and many more

# GPU Platforms

*Highest Performance and Flexibility for AI/ML and HPC Applications*



Confidential

## HGX Platforms



8U-8GPU SYS-821GE-TNHR

Integrated Performance, HGX H100 8-GPU



4U-8GPU SYS-421GE-TNHR2 LCC

Integrated Performance, HGX H100 8-GPU



5U/4U-4GPU SYS-521GU-TNXR

Scalable Performance, HGX H100 4-GPU

## PCIe Gen5 Platforms



5U-10 GPU SYS-521GE-TNRT

Dual Root , PCIe GPU



4U-10 GPU SYS-421GE-TNRT

Dual Root, Direct Connect PCIe GPU



4U-4GPU SYS-741GE-TNRT

Flexible Solution, PCIe GPU

## MGX Platforms



CG1, CG2, C2 Systems

# AS -8125GS-TNMR2


## MI300X Solution



### Specifications




<b>CPU –</b> Dual EPYC 9004 Series processors, up to 256 cores/512 threads	<b>Memory –</b> 24x DIMM slots, ECC DDR5 supports up to 4800MHz
<b>Drives –</b> 2x onboard NVMe M.2 12x PCIe5.0 x4 NVMe U.2 (additional optional) 4x PCIe5.0 x4 NVMe U.2 2x SATA 2.5"	<b>Expansion –</b> 8x PCIe 5.0 x16 low profile 2x FHFL PCIe 5.0 x16 slots (Optional) 2x PCIe 5.0 x16 slots via additional PCIe switch
<b>System Cooling</b> 5x Front and 5x Rear counter-rotating fans with optimal fan speed control	<b>Power Supply –</b> 6x or 8x 3000-watt N+N Redundant Titanium Level Power Supplies
<b>Accelerator Support -</b> 8x MI300X	<b>Dimensions – W x H x D</b> 14.0" (H) x 17.2" (W) x 33.2" (D)



### Recommend System Configuration:

Item#	Description	Q'ty
AS -8125GS-TNMR2	[NR]H13DSG-OM, CSE-GP801TS-D2 for MI300X	1
PSE-GEN9534-0799	Genoa 9534 DP/UP 64C/128T 2.45G 256M 280W SP5	2
MEM-DR564L-CL01-ER48	64GB DDR5-4800 2RX4 (10X4) LP (16Gb) ECC RDIMM	24
HDS-MMN-MTFDKBA960TFR-15	Micron 7450 PRO 960GB NVMe PCIe 4.0 M.2 22x80mm TCG Opal	2
AOC-CX766003N-SQ0	Nvidia 900-9X766-003N-SQ0 PCIe 1-port IB 400GE OSFP Gen5	8
GPU-AMD-MI300X-OAM-0045H	[NR] AMD Instinct MI300X 192GB 8 OAM + UBB	1
AOC-STGS-I2T-O	Std LP 2-port 10G RJ45, Intel X550 (Retail Pack)	1



# Scalable AI Clusters

The Supermicro Scalable Rack Scale AI Solutions are based on a single rack as a building block, referred to as a Scalable Unit, which can then be expanded to four or eight racks.



Scalable Unit  
Base Package

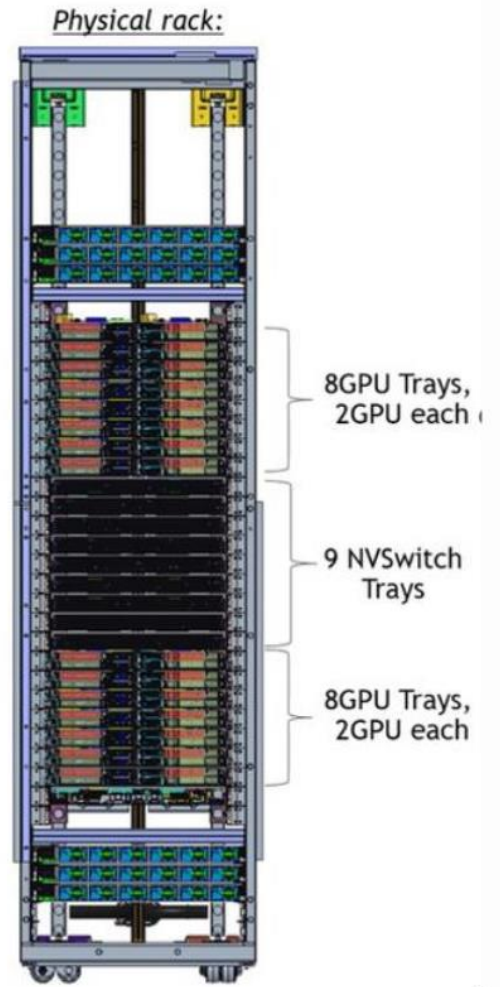
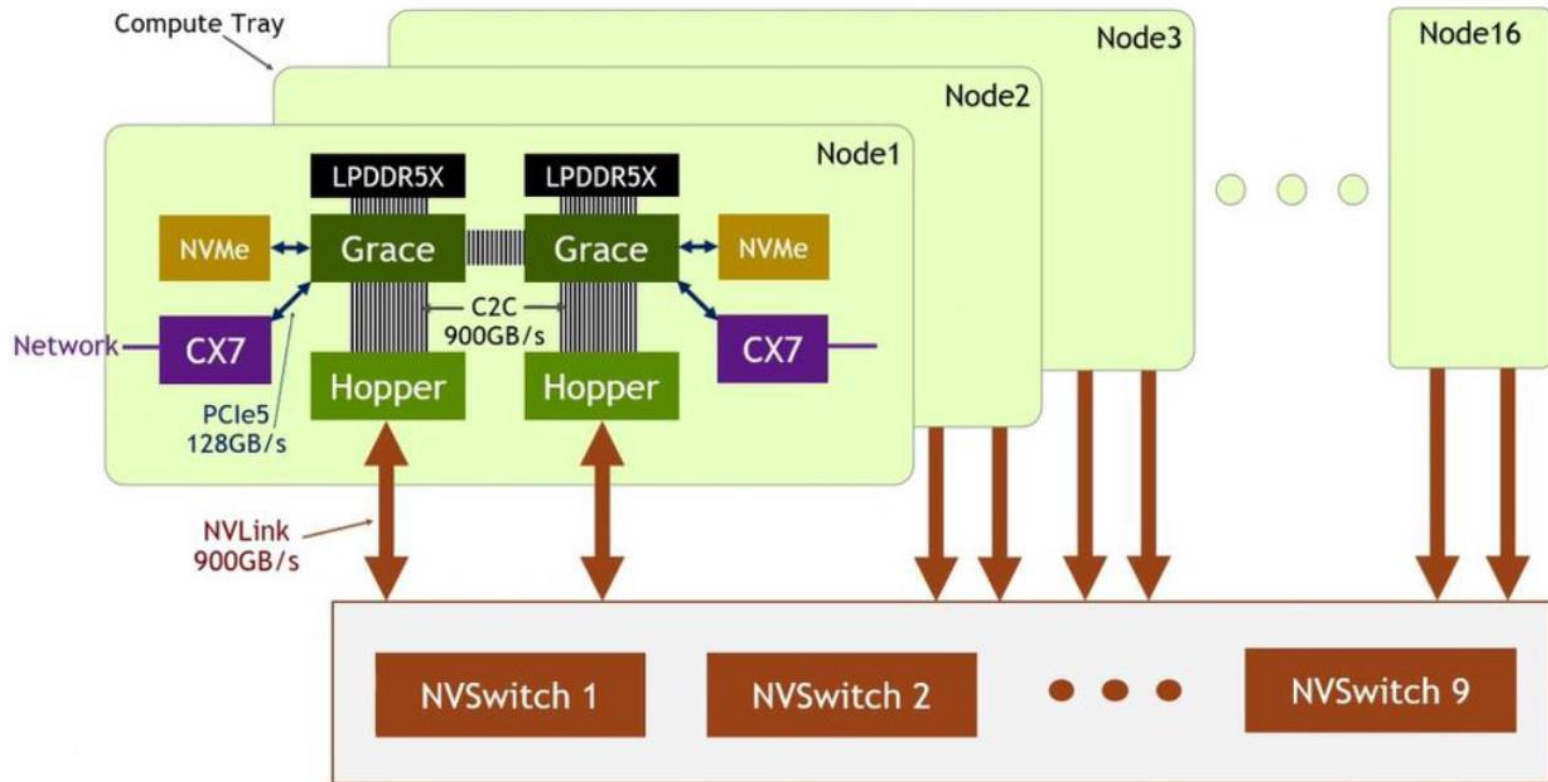
POD


SuperPOD

	<b>SU/Base Package SRS-42UGPU-AI-SU1</b>	<b>POD SRS-42UGPU-AI-SU2</b>	<b>SuperPOD SRS-42UGPU-AI-SU3</b>
<i>GPU Server (8U 8GPU)</i>	4x SYS-821GE-TNHR / 4x AS -8125GS-TNHR	16x SYS-821GE-TNHR / 16x AS -8125GS-TNHR	32x SYS-821GE-TNHR / 32x AS -8125GS-TNHR
<i>Total CPUs</i>	8x Intel® Xeon® Platinum 8480+ Processors or 8x AMD EPYC™ 9004 Processors	32x Intel® Xeon® Platinum 8480+ Processors or 16x AMD EPYC™ 9004 Processors	64x Intel® Xeon® Platinum 8480+ Processors or 64x AMD EPYC™ 9004 Processors
<i>Total GPUs</i>	32x NVIDIA HGX H100 SXM5	128x NVIDIA HGX H100 SXM5	256x NVIDIA HGX H100 SXM5
<i>Rack</i>	1x 42U (Optional 48U)	4x 42U (Optional 48U)	8x 42U (Optional 48U)
<i>Memory</i>	32TB DDR5 (X13) 24TB DDR5 (H13)	128TB DDR5 (X13) 96TB DDR5 (H13)	256TB DDR5 (X13) 192TB DDR5 (X13)
<i>Estimated Total Power Per Rack</i>	Max 45 kW	Max 180 kW	Max 360 kW
<i>Networking</i>	1x 400G 64-port NDR IB: SSE- MQM9700-NS2F	1x 400G 64-port NDR IB: SSE- MQM9700-NS2F	3x 400G 64-port NDR IB: SSE- MQM9700-NS2F
	1x SMC 100G Eth Switch (Storage)	1x SMC 100G Eth Switch (Storage)	1x SMC 100G Eth Switch (Storage)
	1x SMC 1G/25G MGT Switch	1x SMC 1G/25G MGT Switch	2x SMC 1G/25G MGT Switch

# OBERON GH200 System And Rack Architecture

32 GPU fully NVLink connected by NVSwitch



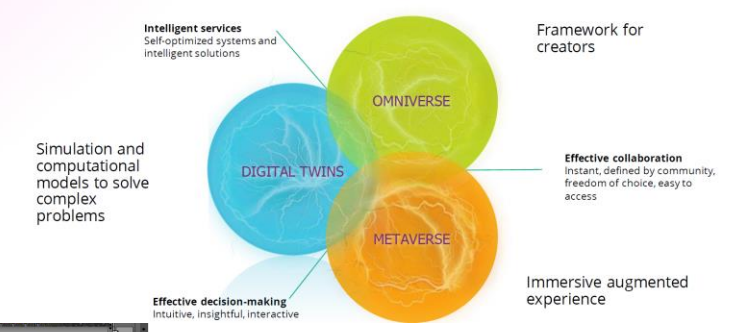


# Virtual Reality, Digital Twins, NVIDIA Omniverse

## Supermicro OVX Building Block Solution


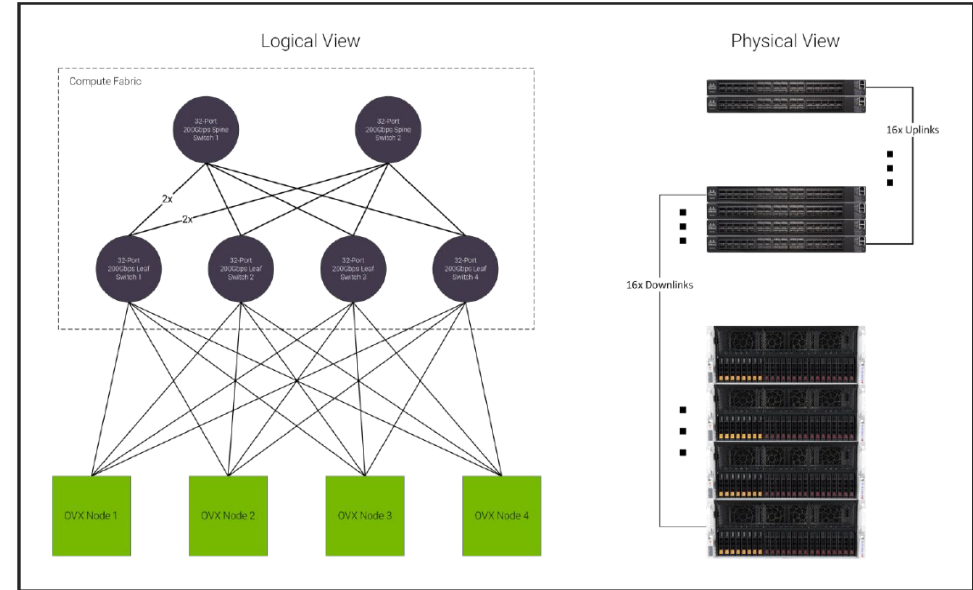
Rack-scale and a Full Turn-key Solution (L11/L12)

### REAL AND DIGITAL WORLDS CONVERGENCE




Scalable Unit / Base Package	
Systems	4x Supermicro OVX nodes (SYS-420GP-TNR) 1x Supermicro Nucleus server (SYS-120U-TNR)
Networking	6x 200Gbps 32-port Nvidia SN3700 Ethernet switches, Non-blocking
Graphics	32x Nvidia A40 GPUs
Power	Up to 18.56.kW
Other Notes	


Scales-out incrementally via 4-nodes after initial switch counts are met to 16 nodes for a Pod



Digital Twins




Simulations and Rendering



Design and Collaboration



## OLTP/OLAP, Data Analytics...



## MP SuperServer SYS-681E-TR

### Key Applications

In-Memory Database, Virtualization, ERP, CRM, Research Lab/National Lab, Scale-up HPC,

### Key Features

- Octa Socket E (LGA-4677) 4th Generation Intel® Xeon® Scalable processors;
- 128 DIMM Slots; Up to 32TB DDR5-4800MHz ECC RDIMM/RDIMM-3DS;
- Optional PCIe configurations up to 24 PCIe slots with support for 12 double-width GPU/Accelerator cards;
- 24x 2.5" Hot-swap NVMe/SAS3/SATA3 drive bays ; 2x internal M.2 NVMe/SATA; Optional RAID support via storage add-on card;
- 10 heavy duty fans w/ optimal fan speed control;



## MP SuperServer SYS-241E-TNRTTP

### Key Applications

Artificial Intelligence (AI), Business Intelligence, ERP, CRM, Scientific Virtualization, In-Memory Database, HCI, SAP HANA,

### Key Features


- Quad Socket E (LGA-4677) 4th Generation Intel® Xeon® Scalable processors;
- 64 DIMM Slots; Up to 16TB DDR5-4800MHz ECC RDIMM/RDIMM-3DS;
- Optional PCIe configurations up to 8 PCIe slots with support for two double-width GPU/Accelerator cards;
- Flexible networking options with 2 AIOM networking slot (OCP NIC 3.0 compatible);
- 24x 2.5" Hot-swap NVMe/SAS3/SATA3 Hybrid Drive bays (Default with 24 SAS3/SATA3 drive bays); 2x internal M.2 NVMe/SATA; Optional RAID support via storage add-on card;
- 6 heavy duty fans w/ optimal fan speed control;



## Features


- Supermicro multi-processor systems bring new levels of compute performance and flexibility with support for 4th Gen Intel® Xeon® Scalable processors to support mission-critical enterprise workloads
- 4- and 8-way systems with 4th Gen Intel® Xeon® Scalable processors up to 480 cores
- Large memory footprint with up to 64 DIMMs in 2U and 128 DIMMs in 6U up to 32TB of DDR5 Memory
- Next-generation PCIe 5.0 for GPU/accelerator and high-speed network interface cards
- Compute and hybrid storage-optimized configurations





# HPC: AL, ML, Inference, Simulation,

**NVIDIA H100 Tensor Core GPU**  
 Unprecedented performance, scalability, and security for every data center.




## Compute Nodes

# X13 SuperBlade®

Ultra High-Density Multi-Node Systems for Enterprise, Cloud, HPC, and AI Applications



### Optimized for Performance, Density and Advanced Networking

- Up to 20 nodes in 8U or 10 nodes in 6U with integrated switches
- Single or dual 4th Gen Intel® Xeon® Scalable processors with air-cooled support for up to 350W TDP CPUs
- Up to 32 DIMM slots per node supporting DDR5-4800MHz
- High-performance networking with 400G/200G InfiniBand and 100G/25G Ethernet support
- Up to 4 GPUs per node in a high-density, balanced architecture
- High-performance NVMe support in E1.S, U.2 and M.2 form factors

# X13 BigTwin®

Industry-leading Multi-node Architecture



### Highly configurable 2U 4-node and 2U 2-node systems optimized for density or storage

- Dual socket architecture featuring 4th Gen Intel® Xeon® Scalable processors
- Optimized thermal design with liquid cooling options
- All-hybrid hot-swappable NVMe/SAS/SATA drive bays - Up to 12 drives per node
- Resource Saving Architecture with shared power and cooling for increased efficiency
- Flexible networking with up to 200G Ethernet per node

## GPU Nodes

# X13 Universal GPU

Multi-Architecture Flexibility with Future-Proof Open-Standards-Based Design



### Ultimate modularity and customization options for AI and HPC environments

- Dual 4th Gen Intel® Xeon® Scalable processors
- Support for the latest industry standards including PCIe 5.0, DDR5 and Compute Express Link (CXL) 1.1
- Innovative modular architecture designed for flexibility and futureproofing with a range of form factors from 4U to 8U
- Supports next-generation GPUs including NVIDIA H100 and Intel Data Center GPU Max Series
- Optimized thermal capacity and airflow to support CPUs up to 350W and GPUs up to 700W with air cooling
- PCIe 5.0 x16 networking slots and up to 16 U.2 NVMe drive bays

# X13 PCIe GPU

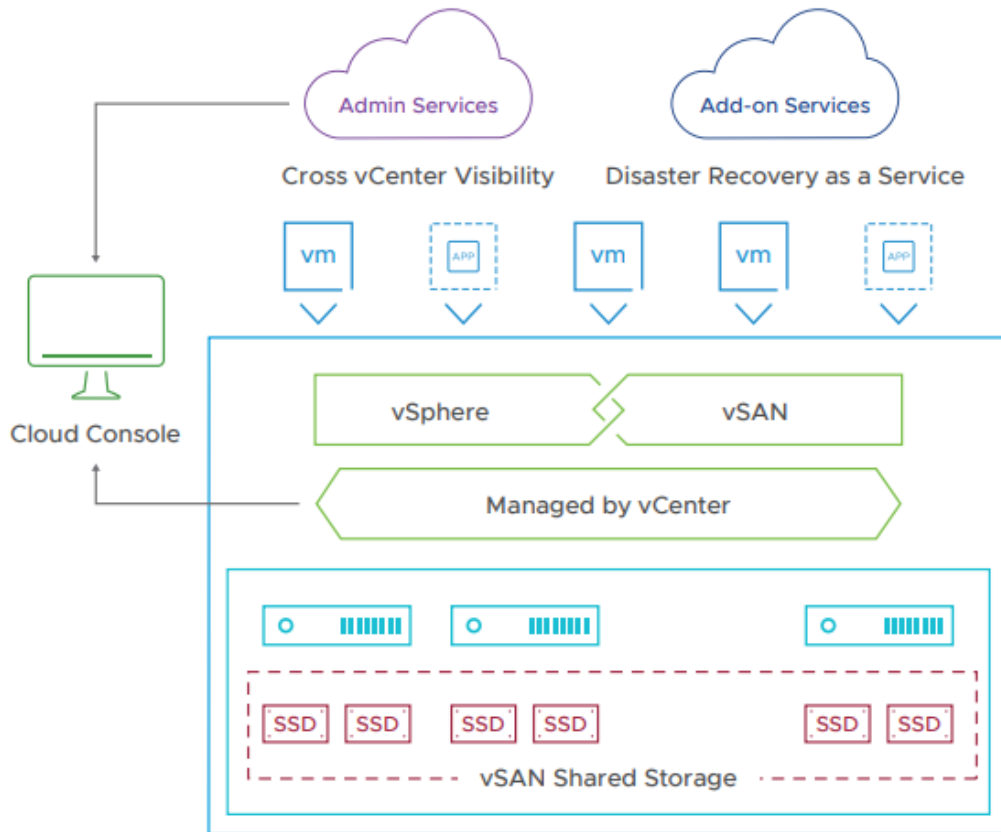

High Performance and Flexibility for AI, 3D Simulation and the Metaverse



### Maximum Flexibility with up to 10 PCIe GPUs in 4U/5U

- Dual 4th Gen Intel® Xeon® Scalable processors (formerly codenamed Sapphire Rapids)
- Support for the latest industry standards including PCIe 5.0, DDR5 and Compute Express Link (CXL) 1.1
- Supports NVIDIA H100, A100, Intel Data Center GPU Max Series and Intel Data Center GPU Flex Series PCIe GPUs
- Optional 1U expansion for enhanced thermal capacity
- Flexible storage with U.2 NVMe and optional direct-to-CPU storage configurations
- Dual root and direct-connect PCIe configurations available

# VMware: vSphere, vSAN



## X13 Hyper

Best-in-class Performance and Flexibility Rackmount Server



### Flagship Performance 1U and 2U Rackmount Systems

- Dual socket 4th Gen Intel® Xeon® Scalable processors with support for Intel Xeon CPU Max Series
- Free-air and liquid cooling options for maximum performance and efficiency
- 32 DIMM slots per node supporting DDR5-4800MHz
- NVMe SSD support with up to 24 drives in 2U
- Optional 2.5"/E1.S SSD hybrid configuration
- Up to 3 PCIe 5.0 slots in 1U or 8 PCI-E 5.0 slots in 2U
- PCIe 5.0 AIOM slots supporting up to 400G networking
- Tool-less system for simplified maintenance

## H13 Hyper Systems

Flexible and High Performance for Enterprise Data Centers




### Enterprise-focused platform designed for utmost performance and flexibility

Gain high performance, flexibility, scalability and serviceability to demanding IT environments, and to power mission-critical enterprise workloads.

- Dual AMD EPYC™ 9004 Series Processors
- Up to 6TB DDR5-4800MHz in 24 DIMMs
- Flexible NVMe, SAS, and SATA drive options
- Configurable PCIe 5.0 expansion capabilities with CXL 1.1+ memory expansion
- AIOM slots with OCP 3.0 support
- Titanium-Level efficiency power supplies

# All-Flash, Hybrid. Block, Object, File



## X13 Petascale All-Flash

Revolutionary Petascale NVMe for Unprecedented Density and Performance



### All-New 1U Storage Architecture

- Dual socket 4th Gen Intel® Xeon® Scalable processors
- 32 DIMM slots per node supporting DDR5-4800MHz
- 2x AIOM supporting PCIe 5.0 x16 and up to 2x PCIe 5.0 x16 slots
- Up to 24 EDSFF Short (E1.S) or 16 E3.S high-performance drives in a 1U chassis
- E1.S (9.5mm and 15mm) or E3.S (7.5mm) form factor support with E3.S 2T x8 CXL support on selected configurations



Object



Parallel File



Object/File

## Hybrid

### Storage SuperServer SSG-620P-E1CR24H



#### Key Features

- Dual socket 3rd Gen Intel® Xeon® Scalable processors, up to 72 Cores Per Node;
- 16 ECC DDR4-3200: LRDIMM/RDIMM;
- Dedicated PCIe 4.0 AIOM slot; 3 x PCIe 4.0 x16 Slots;
- Server remote management: IPMI 2.0 / KVM over LAN / Media over LAN per node;
- 24 3.5" Hot-swap SAS3/SATA3 drives, 4x Rear SATA/NVMe Slots, 2x SATA/NVMe M.2 (form factor: 2280);
- 5x 8cm hot-swap counter-rotate redundant PWM cooling fans;
- 1600W Redundant Power Supplies Titanium Level (96%);
- HW RAID support via Broadcom® 3908;

### Storage SuperServer SSG-640SP-E1CR90



#### Key Features

- 16 ECC DDR4-3200: LRDIMM/RDIMM;
- 3 x PCIe 4.0 x16 HHHH PCIe slots;
- 90 3.5"/2.5" Hot-swap SAS3/SATA3 drives, 2x Fixed slim SATA SSD, 2x NVMe M.2 (form factor: 2280 and 22110);
- 6 x 8cm hot-swap counter-rotate redundant PWM cooling fans;
- 2600W Redundant Power Supplies Titanium Level (96%);
- Drive Controller support via Broadcom® 3916 or 3616; Server remote management: IPMI 2.0 / KVM over LAN / Media over LAN;

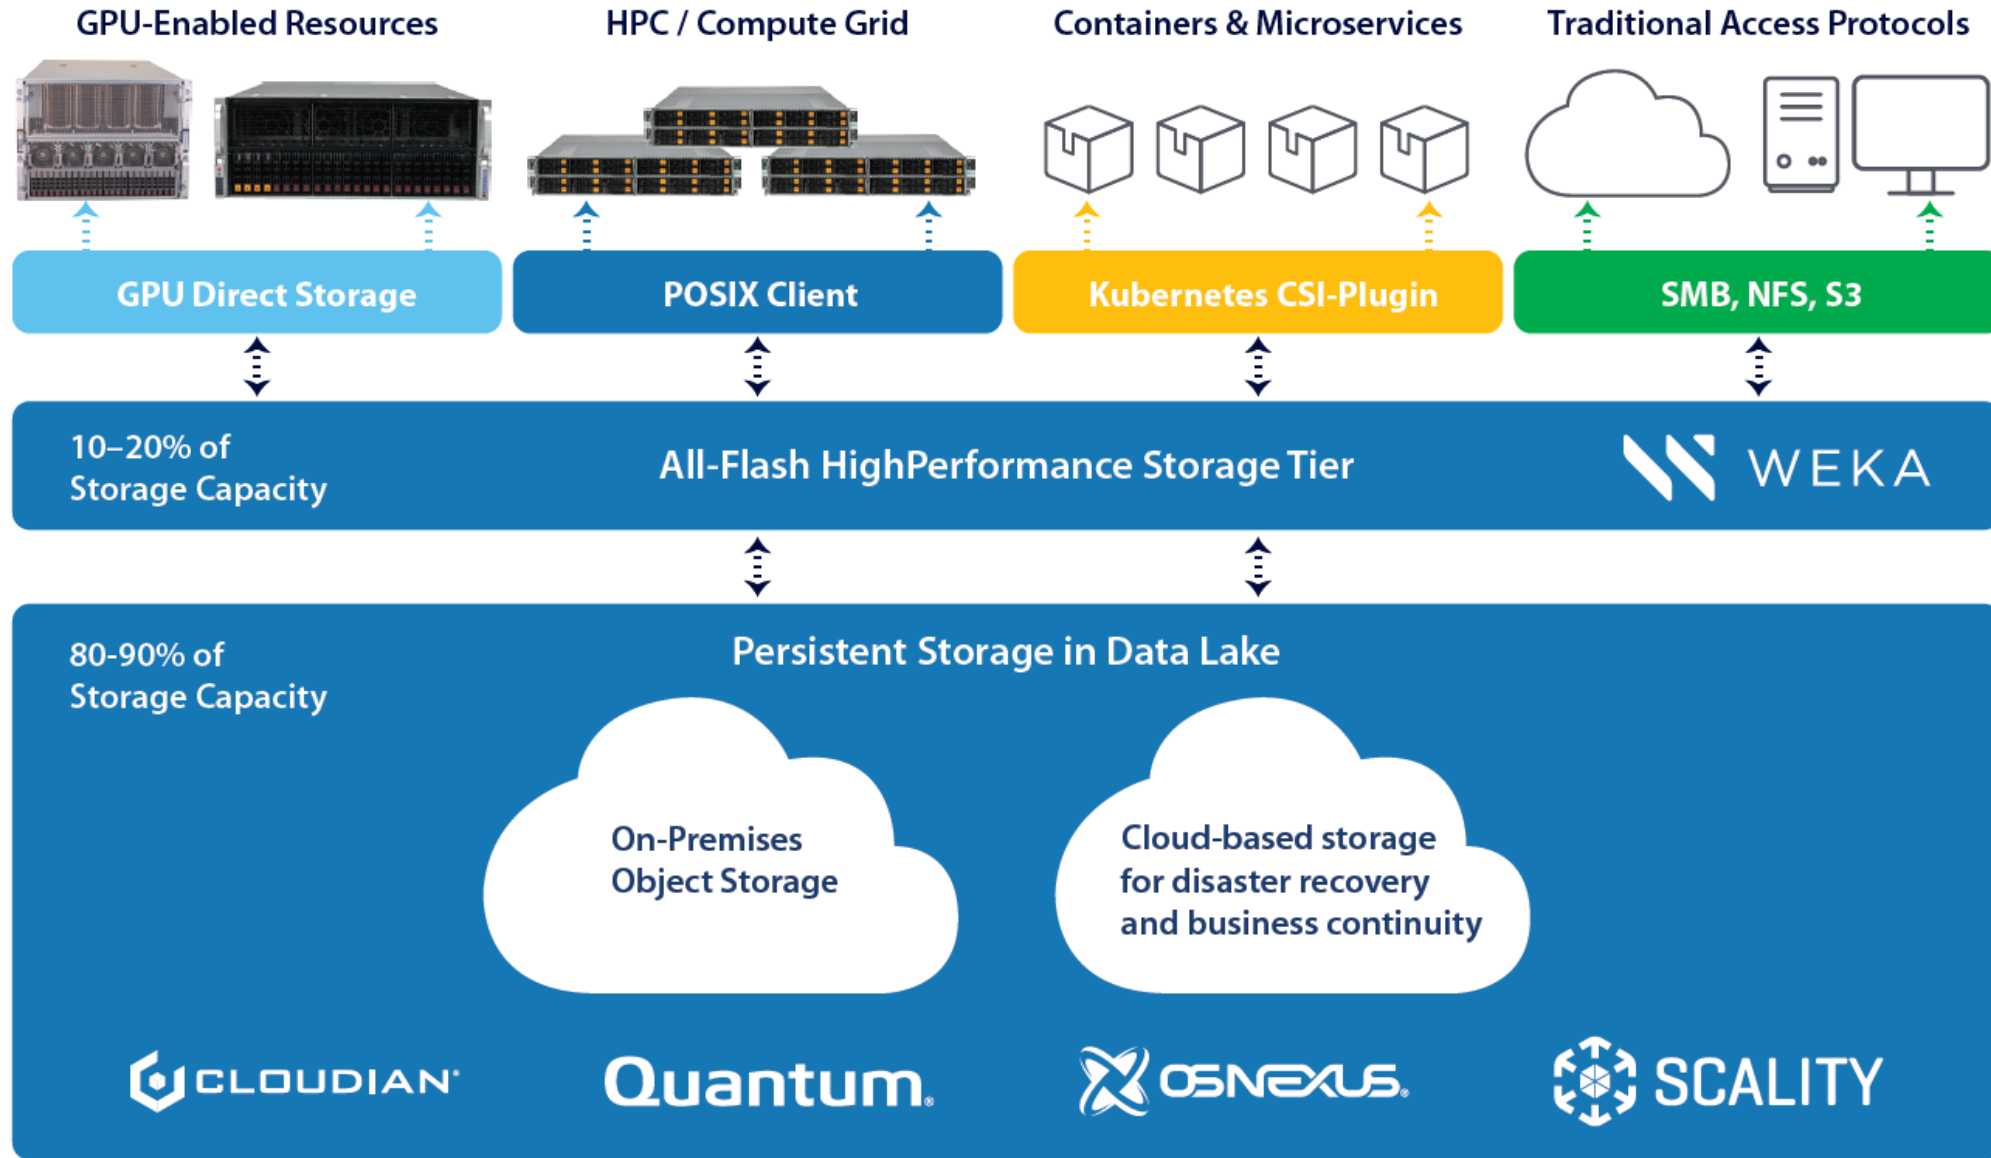


Object/File

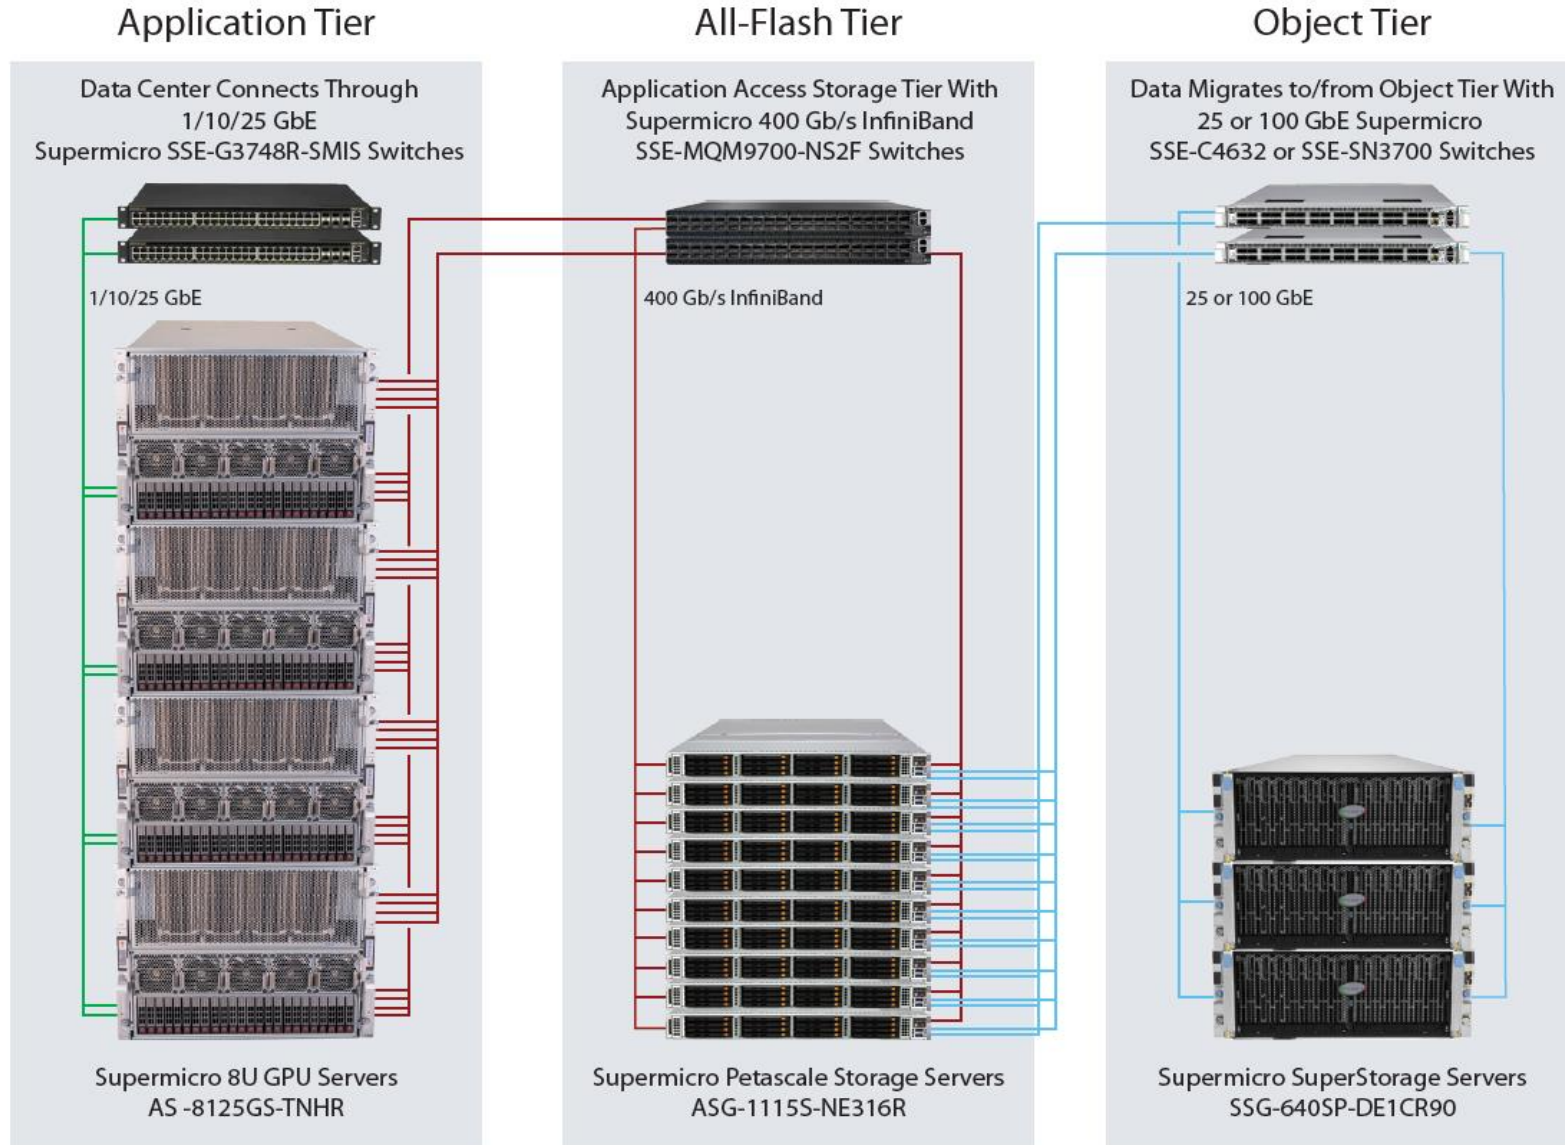


Block

# Tiered Storage Reference Architecture



# Tiered Storage in AI POD



## Optimized Platform for 5G/Telco Applications



### Highlight Systems

**NEW!**

**SuperEdge 2U 3-Node**  
SYS-211E-FRN2T  
SYS-211E-FRDN2T

**NEW!**

**SuperEdge 2U 4-Node**  
SYS-210TP-HPTR  
SYS-210TP-HPTRD

**NEW!**

**Hyper-E**  
SYS-221HE-FTNRD  
SYS-221HE-FTNR

**NEW!**

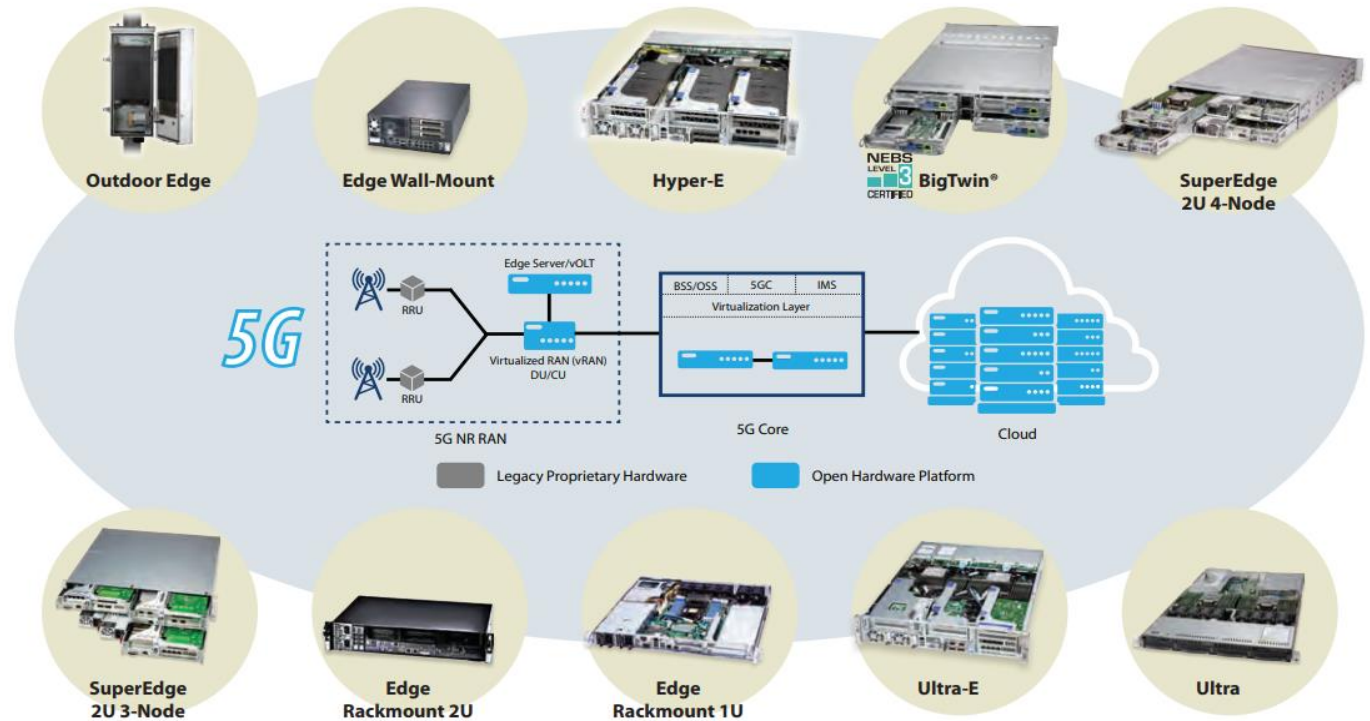
**SuperEdge 2U 3-Node**  
SYS-211E-FRN2T  
SYS-211E-FRDN2T

**NEW!**


**Edge Rackmount 1U**  
SYS-111E-FWTR  
SYS-111E-FDWTR

**NEW!**

**Outdoor Edge**  
SYS-E403-9D-16C-IPD2  
SYS-E403-9D-14CN-IPD2  
SYS-E403-9D-16C-IP



**NEC**    **MAVENIR**    **Rakuten Symphony**




- Supermicro  
NVIDIA MGX Systems  
Product Overview

# NVIDIA Grace CPU Superchip

- 2x Grace CPUs with attached LPDDR5X on mezzanine module
- 900GB/s NVLink Chip-to-Chip on the module
  - No off-the-module NVLink support



	Feature
Core count	144 Arm Neoverse V2 Cores with 4x128b SVE2
L1 cache	64KB i-cache + 64KB d-cache
L2 cache	1MB per core
L3 cache	234MB
LPDDR5X size	240GB, 480GB and 960GB on-module memory options
Memory bandwidth	Up to 1TB/s
NVIDIA NVLink-C2C bandwidth	900GB/s
PCIe links	Up to 8x PCIe Gen5 x16 option to bifurcate
Module thermal design power (TDP)	500W TDP with memory
Form factor	Superchip module
Thermal solution	Air cooled or liquid cooled

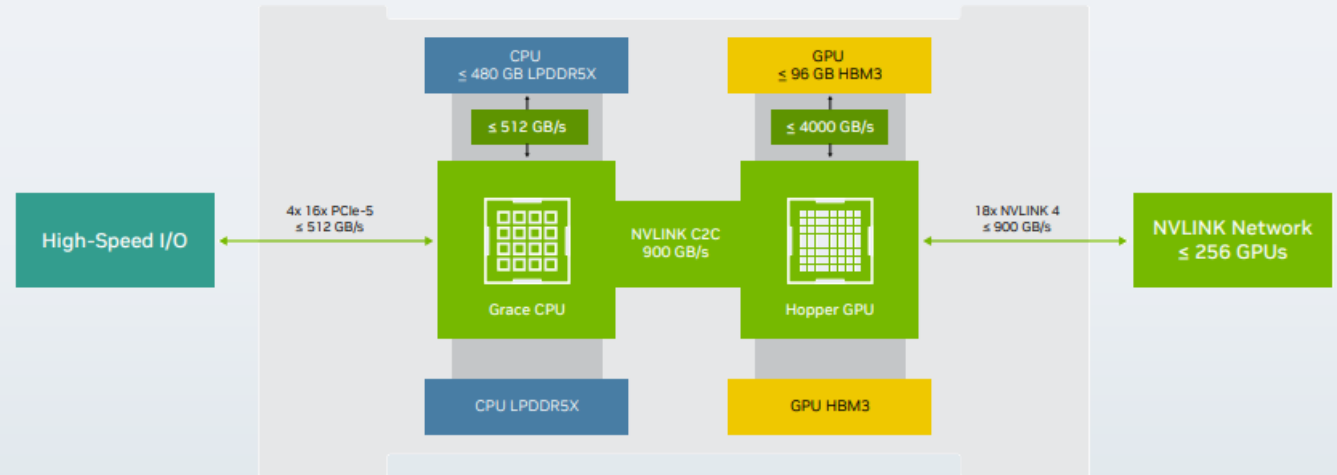
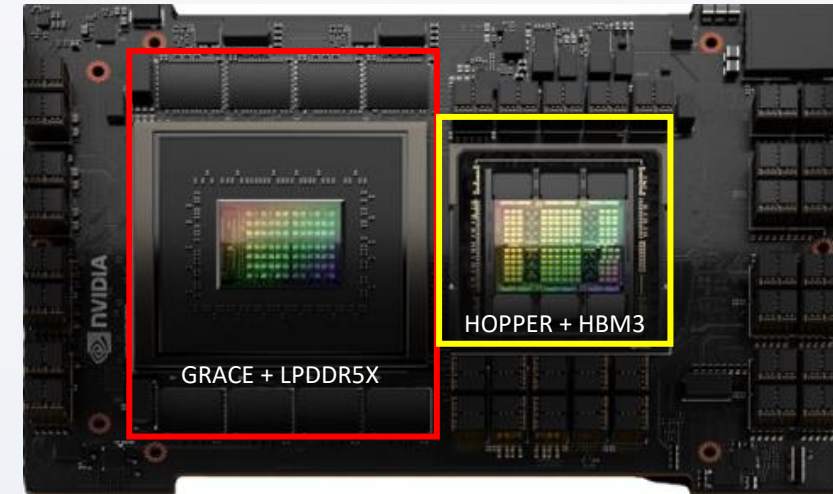




# NVIDIA GH200 Grace Hopper Superchip

- One Grace CPU with integrated LPDDR5X and one H100 Tensor Core GPU (Hopper) on mezzanine module
- Fast NVLink-C2C interface between CPU and GPU

Grace CPU	Feature
CPU core count	72 Arm Neoverse V2 cores
L1 cache	64KB i-cache + 64KB d-cache
L2 cache	1MB per core
L3 cache	117MB
LPDDR5X size	Up to 480GB
Memory bandwidth	Up to 512GB/s
PCIe links	Up to 4x PCIe x16 (Gen5)
Hopper H100 GPU	Feature
FP64	34 teraFLOPS
FP64 Tensor Core	67 teraFLOPS
FP32	67 teraFLOPS
TF32 Tensor Core	989 teraFLOPS*   494 teraFLOPS
BFLOAT16 Tensor Core	1,979 teraFLOPS*   990 teraFLOPS
FP16 Tensor Core	1,979 teraFLOPS*   990 teraFLOPS
FP8 Tensor Core	3,958 teraFLOPS*   1,979 teraFLOPS
INT8 Tensor Core	3,958 TOPS*   1,979 TOPS
HBM3 size	Up to 96GB
Memory bandwidth	Up to 4TB/s
NVIDIA NVLink-C2C CPU-to-GPU bandwidth	900 GB/s bidirectional
Module thermal design power (TDP)	Programmable from 450W to 1000W (CPU + GPU + memory)
Form factor	Superchip module
Thermal solution	Air cooled or liquid cooled



# 1U Grace Hopper Superchip Systems



1U with Grace Hopper



1U with Grace Hopper LC



1U 2-Node with Grace Hopper

Model	ARS-111GL-NHR	ARS-111GL-NHR-LCC	ARS-111GL-DNHR-LCC
CPU	72-core Grace Arm Neoverse V2 CPU + H100 Tensor Core GPU in a single chip	72-core Grace Arm Neoverse V2 CPU + H100 Tensor Core GPU in a single chip	72-core Grace Arm Neoverse V2 CPU + H100 Tensor Core GPU in a single chip per node
Cooling	Air-cooled	Liquid-cooled	Liquid-cooled
GPU Support	NVIDIA H100 Tensor Core GPU with 96GB of HBM3	NVIDIA H100 Tensor Core GPU with 96GB of HBM3	NVIDIA H100 Tensor Core GPU with 96GB of HBM3 per node
Memory	CPU: 480G integrated LPDDR5X with ECC GPU: 96GB HBM3	CPU: 480G integrated LPDDR5X with ECC GPU: 96GB HBM3	CPU: 480G integrated LPDDR5X with ECC per node GPU: 96GB HBM3 per node
Networking	3x PCIe 5.0 x16 slots supporting NVIDIA BlueField-3 or ConnectX-7	3x PCIe 5.0 x16 slots supporting NVIDIA BlueField-3 or ConnectX-7	2x PCIe 5.0 x16 slots supporting NVIDIA BlueField-3 or ConnectX-7
Storage	8x Hot-swap E1.S drives and 2x M.2 NVMe drives	8x Hot-swap E1.S drives and 2x M.2 NVMe drives	4x Hot-swap E1.S drives and 2x M.2 NVMe drives per node
Power Supplies	2x 2000W Titanium Level	2x 2000W Titanium Level	2x 2700W Titanium Level

# 1U/2U Grace CPU Superchip and x86 Systems



1U 2-Node with Grace CPU



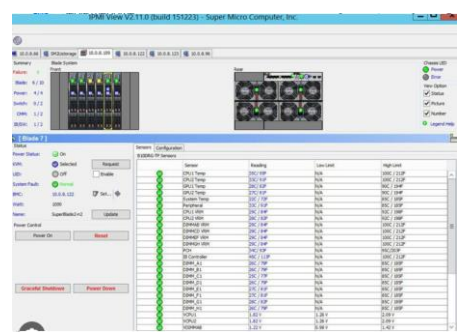
2U with Grace CPU



2U with x86 DP

Model	ARS-121L-DNR	ARS-221GL-NR	SYS-221GE-NR
CPU	144-core Grace Arm Neoverse V2 CPU in a single chip per node (total of 288 cores)	144-core Grace Arm Neoverse V2 CPU in a single chip	4 <sup>th</sup> or 5 <sup>th</sup> Generation Intel® Xeon® Scalable processors (up to 60-core per socket)
Cooling	Air-cooled	Air-cooled	Air-cooled
GPU Support	Please contact our sales for possible configurations	Up to 4 double-width GPUs including NVIDIA H100 PCIe, H100 NVL, L40S.	Up to 4 double-width GPUs including NVIDIA H100 PCIe, H100 NVL, L40S
Memory	Up to 480GB of integrated LPDDR5X with ECC and up to 1TB/s of memory bandwidth per node	Up to 480GB of integrated LPDDR5X with ECC and up to 1TB/s of memory bandwidth per node	Up to 2TB, 32x DIMM slots, ECC DDR5-4800 DIMM
Networking	3x PCIe 5.0 x16 slots supporting NVIDIA BlueField-3 or ConnectX-7	3x PCIe 5.0 x16 slots supporting NVIDIA BlueField-3 or ConnectX-7	2x PCIe 5.0 x16 slots supporting NVIDIA BlueField-3 or ConnectX-7
Storage	4x Hot-swap E1.S drives and 2x M.2 NVMe drives per node	8x Hot-swap E1.S drives and 2x M.2 NVMe drives	8x Hot-swap E1.S drives and 2x M.2 NVMe drives
Power Supplies	2x 2700W Titanium Level	3x 2000W Titanium Level	3x 2000W Titanium Level

# System Management Software



	Standard	Basic	Advanced	Enterprise
<b>Description</b>	Covers all core functionality to effectively set up, manage, and monitor your Supermicro systems. These features are available to all Supermicro users.	Extends the core functionality and makes system management easier with additional features, such as remote BIOS management and system updates.	Delivers a broad set of tools to help administrators improve the performance, up-time, and monitoring of Supermicro systems.	Offers an extensive platform to manage large data centers and coordinate automated lifecycle management, software-defined infrastructure, and more in a single pane of glass.
<b>License</b>	No license required	SFT-OOB-LIC	SFT-DCMS-SINGLE	SFT-DCMS-SINGLE + SFT-SDDC-SINGLE
<b>Key features<sup>(1)</sup></b>	<ul style="list-style-type: none"> <li>Secure remote console (KVM/HTML5)</li> <li>System temperature monitoring</li> <li>System power thresholds &amp; alerts</li> <li>Component monitoring</li> <li>Email alerting</li> <li>Remote configuration</li> <li>Offline diagnostics</li> <li>Crash dump</li> <li>License management</li> </ul>	<ul style="list-style-type: none"> <li>Remote BMC management</li> <li>Remote BIOS management</li> <li>Out-of-Band systems checks</li> <li>TPM Provisioning</li> <li>Mount/Unmount ISO images from Samba/HTTP</li> <li>Basic Redfish APIs</li> <li>CIM management</li> <li>SysLog</li> </ul>	<ul style="list-style-type: none"> <li>Remote OS deployment</li> <li>Auto-discovery</li> <li>Power capping</li> <li>RAID monitoring and configuration</li> <li>HDD monitoring</li> <li>Advanced Redfish APIs</li> <li>FW update policy</li> <li>System lock down</li> <li>Crash screen/video capture</li> </ul>	<ul style="list-style-type: none"> <li>3rd Party vendor support</li> <li>POD &amp; Rack-level management</li> <li>SDI Lifecycle management</li> <li>Manage Composable Dissaggregated Infrastructure</li> <li>Zero-touch provisioning for network configuration</li> <li>Single pane of glass for data center deployment</li> <li>Rich analytics &amp; telemetry</li> <li>User defined role-based access control</li> </ul>



## Service-Level Options

- **4-Hour Onsite Response**

A Supermicro authorized representative will arrive at the customer's site to begin hardware maintenance service within 4 hours after the service request has been received, the issue has been identified, and spare parts have been received onsite.

- **Next Business Day Response**

Service is available 8 hours per day within standard business hours, Monday to Friday, excluding local holidays. A Supermicro authorized representative will arrive at the customer's site to begin hardware maintenance service the next day after the service request has been received, the issue has been identified, and defective parts have been determined and shipped.

- **Integration Service**

Supermicro defines the requirements for the installation, upgrade or migration. We perform the planning, identify service requirements, create and execute the project plan, conduct verification testing, training, and provide technical documentation.

# The Bottomline!

- Supermicro is:
  - Unmatched Revolutionary Design Enabling You Fantastic Flexibility Configuring Your Sys!
  - Best PUE
  - Best TCO
  - Fantastic robustness and reliability (less than 0,3% failure rate)
  - Best price/performance Ratio
  - L12 Rack Design and Implementation
  - Ahead of Competition with new chipsets generations

**Don't trust the above?? Try us!!!**

Local and remote POC welcome!



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

- My Contact Details:

Igor Berin

Senior BDM EMEA

Tel: +48 728 233 205

E-mail: [igorb@supermicro.com](mailto:igorb@supermicro.com)



## DISCLAIMER

Super Micro Computer, Inc. may make changes to specifications and product descriptions at any time, without notice. The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors. Any performance tests and ratings are measured using systems that reflect the approximate performance of Super Micro Computer, Inc. products as measured by those tests. Any differences in software or hardware configuration may affect actual performance, and Super Micro Computer, Inc. does not control the design or implementation of third party benchmarks or websites referenced in this document. The information contained herein is subject to change and may be rendered inaccurate for many reasons, including but not limited to any changes in product and/or roadmap, component and hardware revision changes, new model and/or product releases, software changes, firmware changes, or the like. Super Micro Computer, Inc. assumes no obligation to update or otherwise correct or revise this information.

SUPER MICRO COMPUTER, INC. MAKES NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE CONTENTS HEREOF AND ASSUMES NO RESPONSIBILITY FOR ANY INACCURACIES, ERRORS OR OMISSIONS THAT MAY APPEAR IN THIS INFORMATION.

SUPER MICRO COMPUTER, INC. SPECIFICALLY DISCLAIMS ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE. IN NO EVENT WILL SUPER MICRO COMPUTER, INC. BE LIABLE TO ANY PERSON FOR ANY DIRECT, INDIRECT, SPECIAL OR OTHER CONSEQUENTIAL DAMAGES ARISING FROM THE USE OF ANY INFORMATION CONTAINED HEREIN, EVEN IF SUPER MICRO COMPUTER, Inc. IS EXPRESSLY ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

## ATTRIBUTION

© 2022 Super Micro Computer, Inc. All rights reserved.

**Thank You**



[www.supermicro.com](http://www.supermicro.com)