Supermicro NVIDIA MGX Systems

CG1, CG2, C2 Optimized Systems

Josh Grossman, System
Yusuke Kondo, Marketing

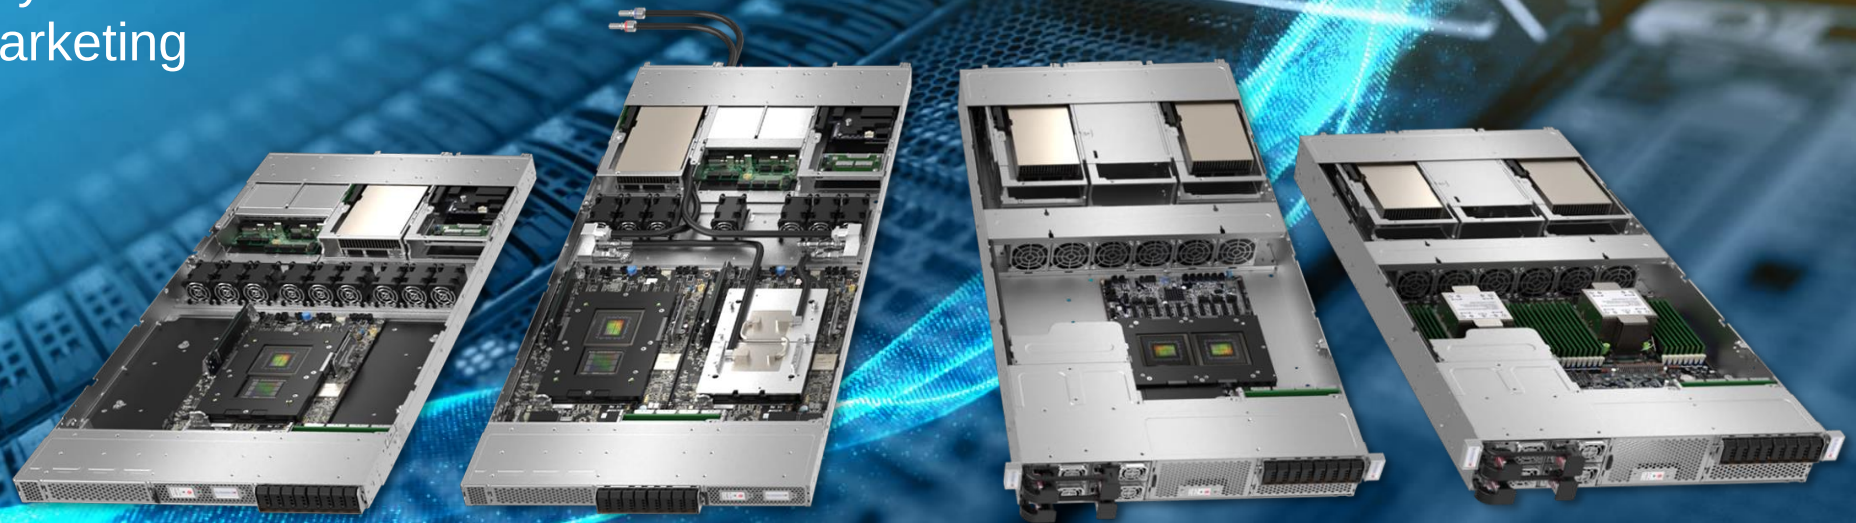

Supermicro GPU System Platforms

HGX Platforms

8U HGX H100 8-GPU

4U HGX H100 4-GPU

PCIe GPU Platforms

4U/5U-8/10 PCIe GPU System

4U PCIe GPU System

MGX Platforms

X86, Grace CPU Systems

Grace Hopper Systems

MGX 1U, 2U and 4U PCIe GPU Platforms

	1U 1 Node	1U 2 Node	2U 1 Node	2U 2 Node	4U 1 Node
Grace Hopper CG1	 AC/LC	 LC	 AC	 AC	
Grace Hopper CG1	 LC		 AC		
Grace Grace C2	 AC	 AC	 AC	 AC	
Intel X86	 AC		 AC		 AC

NVIDIA Grace Platform

Grace Hopper Superchip

Accelerated applications where CPU performance and system memory BW are critical since AI models continue to get bigger and GPUs get even faster.

Grace CPU Superchip

Applications that are not accelerated yet but where absolute performance, energy efficiency, and datacenter density matter, such as in scientific computing, data analytics, and hyperscale computing applications.

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

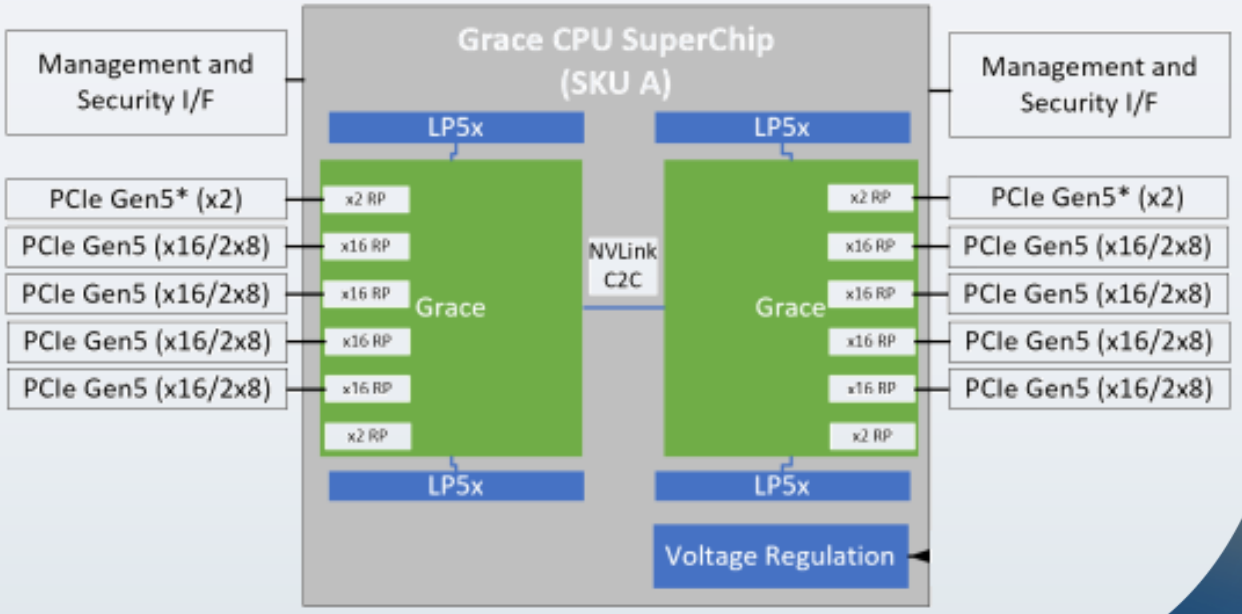

NVIDIA GH200 Grace Hopper Superchip

- 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

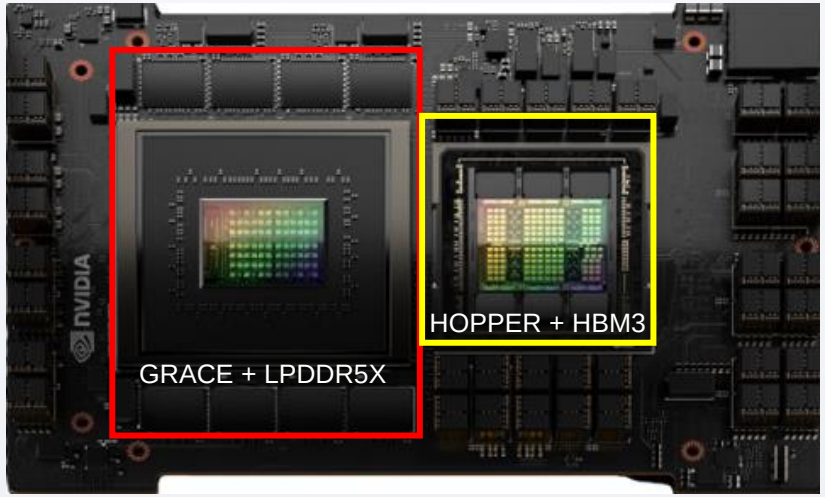

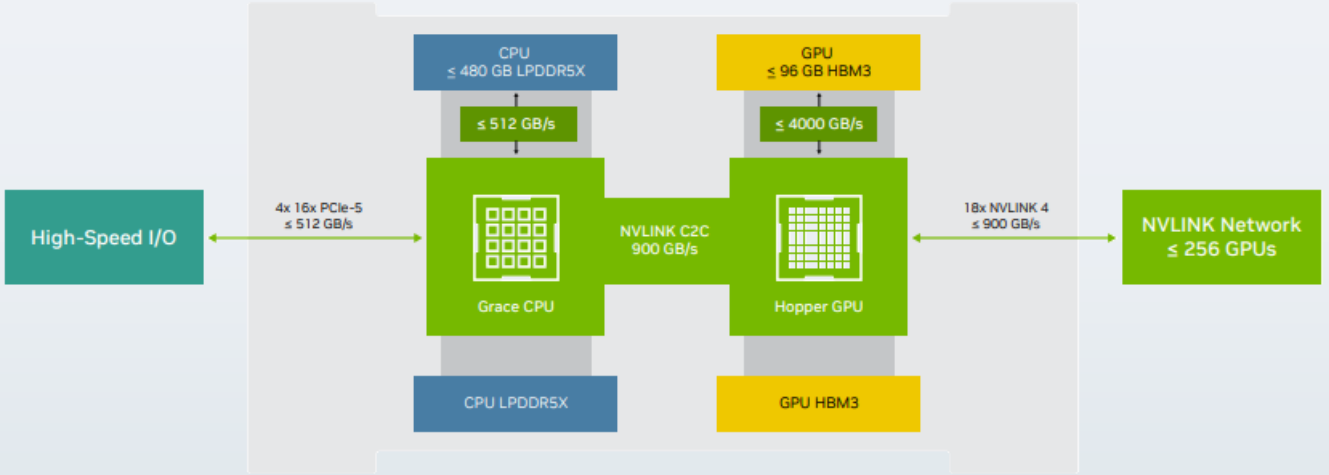

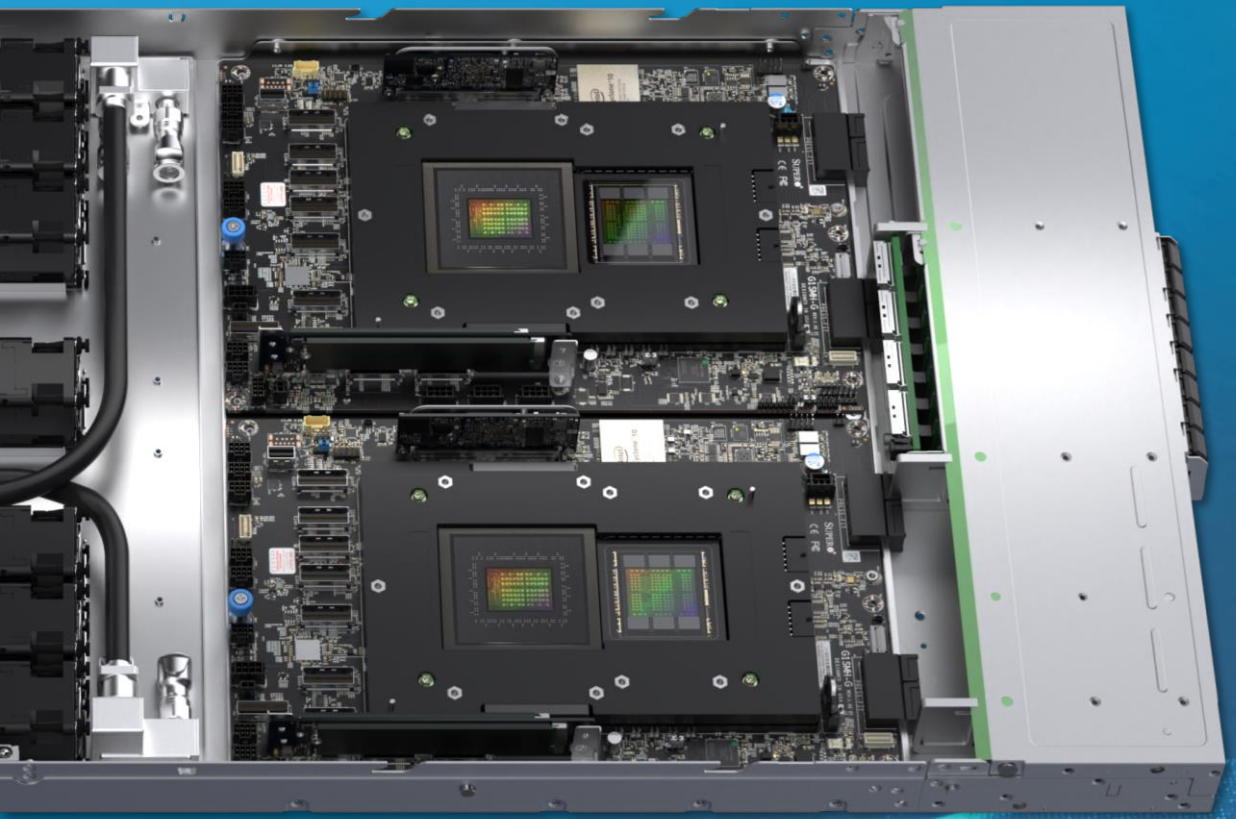

Supermicro NVIDIA MGX Systems

Product Overview

	1U with Grace Hopper	1U with Grace Hopper LC	1U 2-Node with Grace Hopper	1U 2-Node with Grace CPU	2U with Grace CPU	2U with x86 DP
Model	**ARS-111GL-NHR**	**ARS-111GL-NHR-LCC**	**ARS-111GL-DNHR-LCC**	**ARS-121L-DNR**	**ARS-221GL-NR**	**SYS-221GE-NR**
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	144-core Grace Arm Neoverse V2 CPU in a single chip per node (total of 288 cores in one system)	144-core Grace Arm Neoverse V2 CPU in a single chip	4th or 5th Generation Intel® Xeon® Scalable processors
Cooling	Air-cooled	Liquid-cooled	Liquid-cooled	Air-cooled	Air-cooled	Air-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	Please contact 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	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	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 per node, supporting NVIDIA BlueField-3 or ConnectX-7	3x PCIe 5.0 x16 slots per node, supporting NVIDIA BlueField-3 or ConnectX-7	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
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	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 2000W Titanium Level	2x 2000W Titanium Level	2x 2700W Titanium Level	2x 2700W Titanium Level	3x 2000W Titanium Level	3x 2000W Titanium Level

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	4th or 5th 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

1U Grace Hopper Superchip Systems

ARS-111GL-NHR, ARS-111GL-NHR-LCC

- 1U systems with NVIDIA GH200 Grace Hopper Superchip
- 72-core Grace Arm Neoverse V2 CPU + H100 Tensor Core GPU in a single chip
- Onboard memory:
  - CPU: 480GB LPDDR5X
  - GPU: 96GB HBM3
- 8x Hot-swap E1.S drives and 2x M.2
- NVMe drives
- Up to 3x PCIe 5.0 x16 slots supporting NVIDIA BlueField-3 or ConnectX-7
- Liquid-cooled and air-cooled systems available

1U 2-Node Grace Hopper Superchip System

ARS-111GL-DNHR-LCC

- Dual node 1U system with NVIDIA GH200 Grace Hopper Superchip per node
- 72-core Grace Arm Neoverse V2 CPU + H100 Tensor Core GPU in a single chip per node
- Onboard memory:
  - CPU: 480GB LPDDR5X per node
  - GPU: 96GB HBM3 per node
- 4x Hot-swap E1.S drives and 1x M.2
- NVMe drives per node
- 2x PCIe 5.0 x16 slots supporting NVIDIA BlueField-3 or ConnectX-7
- Liquid-cooled

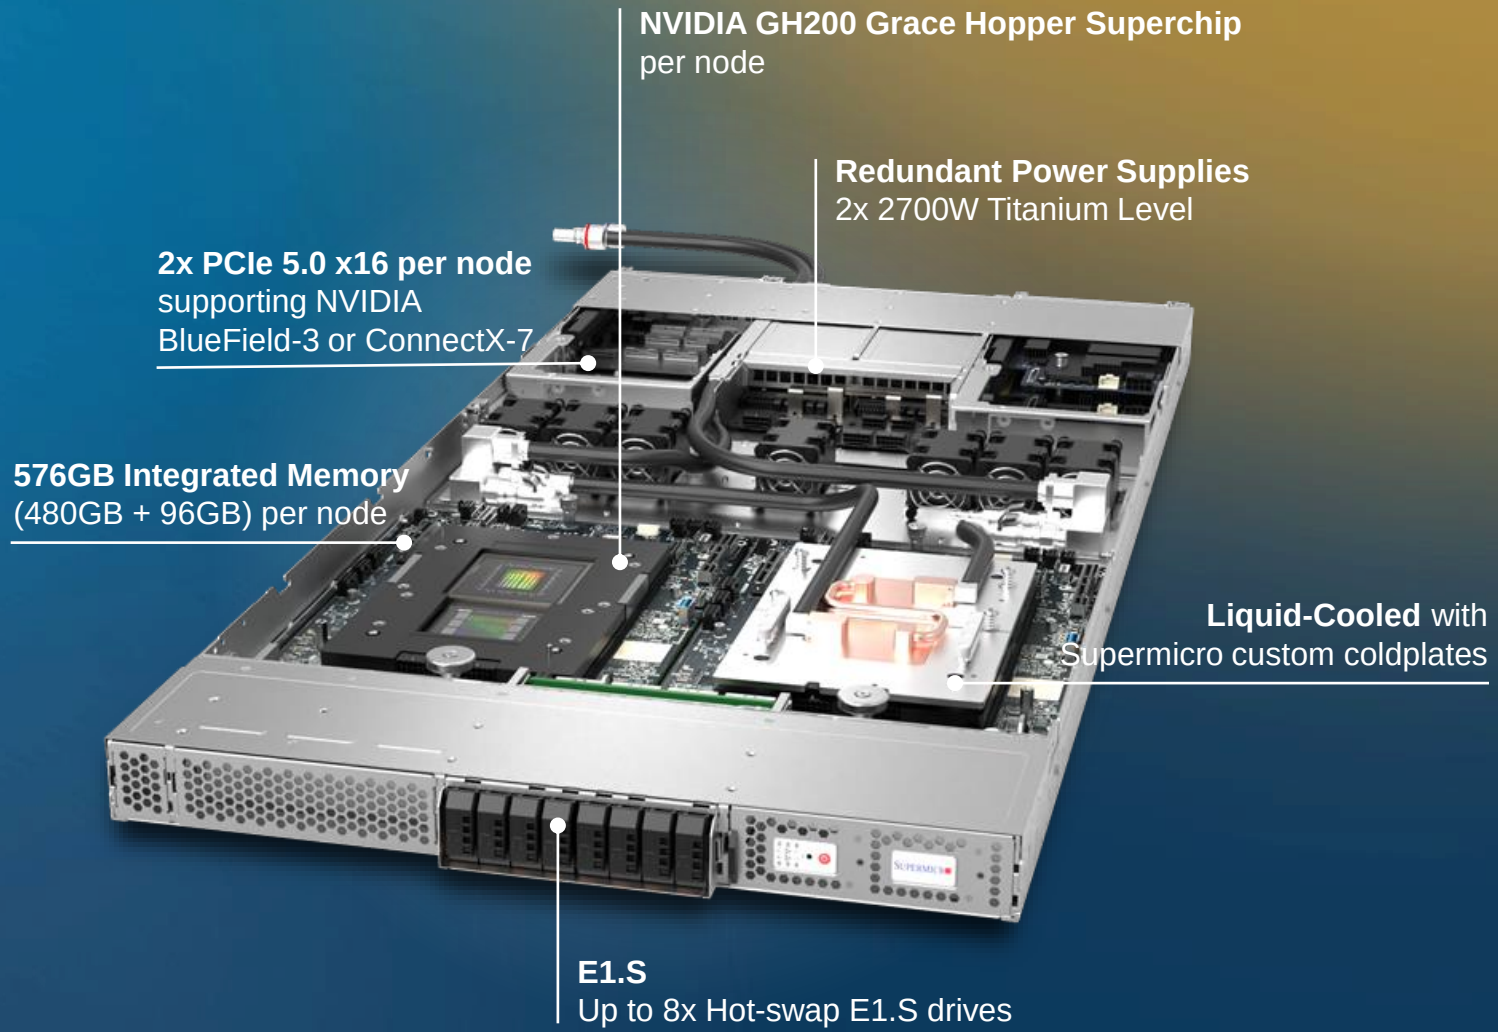

1U Grace Hopper (CG1) System

ARS-111GL-NHR

Key Features

- Single Grace Hopper Superchip up to 1000 Watts
- Up to 576 GB of total Shared Memory including 96GB HBM3 GPU Memory for LLM applications

Key Applications

- Versatile Solutions for HPC, ML and Training
- AI Inferencing
- Edge AI

Specifications

CPU – Grace Hopper Superchip Single Grace Hopper Superchip (up to 1000W TDP)	**Memory –** on Chip Up to 480GB LPDDR5 CPU; 96GB HBM3 GPU
Drives – 8 Hot-Swap Bays 8x E1.S 2x M.2	**Expansion –** 3 PCIe Slots 3x PCIe 5.0 x16
I/O ports (Per Server) 1x RJ45 1GbE IPMI 1x mini-DP (needed converter to VGA), 2x USB	**Power Supply –** Redundant 2x 2000W Titanium Level Efficiency Power Supplies

1U Grace Hopper (CG1) System Liquid Cooled

ARS-111GL-NHR-LCC

Key Features

- Single Grace Hopper Superchip up to 1000 Watts
- Up to 576 GB of total Shared Memory including 96GB HBM3 GPU Memory for LLM applications

Key Applications

- Versatile Solutions for HPC, ML and Training
- AI Inferencing

Specifications

CPU – Grace Hopper Superchip Single Grace Hopper Superchip (up to 1000W TDP)	**Memory –** on Chip Up to 480GB LPDDR5 CPU; 96GB HBM3 GPU
Drives – 8 Hot-Swap Bays 8x E1.S 2x M.2	**Expansion –** 3 PCIe Slots 3x PCIe 5.0 x16
I/O ports (Per Server) 1x RJ45 1GbE IPMI 1x mini-DP (needed converter to VGA), 2x USB	**Power Supply –** Redundant 2x 2000W Titanium Level Efficiency Power Supplies

1U2N Grace Hopper (CG1) System Liquid Cooled

ARS-111GL-DNHR-LCC

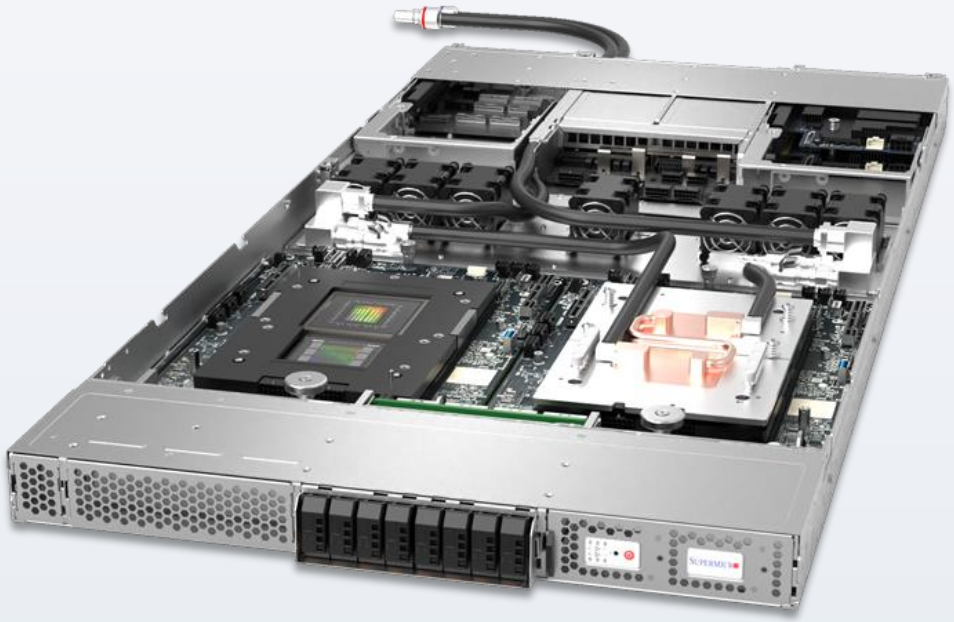

Key Features

- Dual Grace Hopper Superchips up to 1000 Watts per node
- Up to 576 GB of total Shared Memory including 96GB HBM3 GPU Memory per node for LLM applications

Key Applications

- Versatile Solutions for HPC, ML and Training
- Massive-scale AI training and inference

Specifications

CPU – Grace Hopper Superchip Dual Grace Hopper Superchips (1000W TDP per node, 2000W per system)	**Memory –** on Chip per Node Up to 480GB LPDDR5 CPU; 96GB HBM3 GPU
Drives – 8 Hot-Swap Bays 4x E1.S Per Node 2x M.2 Per Node	**Expansion –** 2 PCIe Slots per Node 2x PCIe 5.0 x16 Per Node
I/O ports (Per Node) 1x RJ45 1GbE IPMI 1x mini-DP (needed converter to VGA), 2x USB	**Power Supply –** Redundant 2x 2700W Titanium Level Efficiency Power Supplies

2U Grace CPU Superchip System

ARS-221GL-NR

- 2U systems with NVIDIA Grace CPU Superchip or x86 processors
- 144-core Grace Arm Neoverse V2 CPU in a single chip, or dual 4th and 5th Generation Intel® Xeon® Scalable processors (up to 60-core per socket)
- 480GB integrated LPDDR5X memory with ECC, up to 1TB/s memory bandwidth
- Up to 4 double-width GPUs including NVIDIA H100 PCIe, H100 NVL, L40S
- 8x Hot-swap E1.S drives and 2x M.2 NVMe drives
- 3x PCIe 5.0 x16 slots supporting NVIDIA BlueField-3 or ConnectX-7

2U Dual x86 Intel CPU System

SYS-221GE-NR

- 2U system with dual 4th or 5th Generation Intel® Xeon® Scalable processors (up to 60-core per socket)
- Up to 2TB, 32x DIMM slots, ECC DDR5-4800 DIMM
- Up to 4 double-width GPUs including NVIDIA H100 PCIe, H100 NVL, L40S
- 8x Hot-swap E1.S drives and 2x M.2 NVMe drives
- 3x PCIe 5.0 x16 slots supporting NVIDIA BlueField-3 or ConnectX-7

2U Grace (C2) 4GPU System

ARS-221GL-NR

Key Features

- Supports up to 4 Double Width GPUs (H100/H100 NVL and L40S/L40)
- Up to Two Grace CPUs on one Superchip up to 1000W per Super Chip TDP, Up to 144 Cores
- Up to 480GB Shared LPDDR5X Memory
- Front I/O and Rear I/O support

Key Applications

- AI Inference with H100 and L40S/L40
- High Performance Simulation of Complex 3D Graphics using L40S
- AI to high-performance computing (HPC) to data analytics, digital twins, and hyperscale cloud applications

Specifications

CPU – Grace Superchip One Grace Superchip (500W Per Superchip TDP)	**Memory –** Memory on Chip Per Server Up to 480GB LPDDR5X
Drives – 8 Hot-Swap Bays 8x E1.S 2x M.2	**Expansion –** 7PCIe Slots 7x PCIe 5.0 x16 (4FHFL/DW, 3 for I/O)
I/O ports 1x RJ45 1GbE IPMI 1x mini-DP (needed converter to VGA), 2x USB	**Power Supply –** Redundant 3x 2000W Titanium Level Efficiency Power Supplies

1U2N Grace (C2) 2GPU System

ARS-121L-DNR

Key Features

- Two Grace CPUs on one Superchip, Two Superchips on one System
- Up to 144 Core per one Grace Superchip
- Up to 480GB LPDDR5X Memory per one Grace Superchip
- Front I/O and Rear I/O support

Key Applications

- AI Inference with H100 Cards and L40S
- High Performance Simulation of Complex 3D Graphics using L40S
- AI to high-performance computing (HPC) to data analytics, digital twins, and hyperscale cloud applications

Specifications

CPU – Grace Superchip Two Grace Superchip (500W per node, 1000W per system)	**Memory –** on Chip Up to 480 GB LPDDR5X per Grace Superchips, 960 per system
Drives – 8 Hot-Swap Bays 4x E1.S Per Node 2x M.2 Per Node	**Expansion –** 2 PCIe Slots Per Node 2x PCIe 5.0 x16
I/O ports 1x RJ45 1GbE IPMI 1x mini-DP (needed converter to VGA), 2x USB	**Power Supply –** Redundant 2x 2700W Titanium Level Efficiency Power Supplies

2U MGX 4GPU System

SYS-221GE-NR

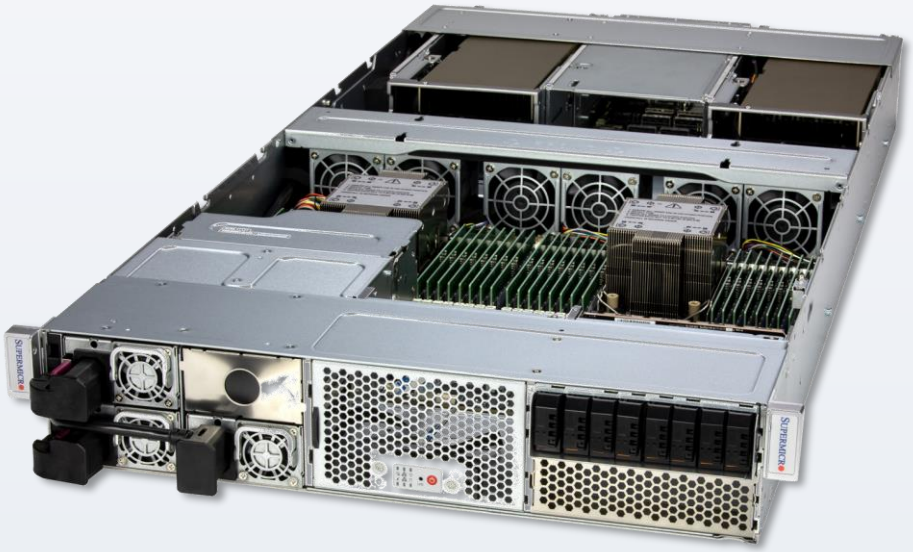

Key Features

- Supports up to 4 Double Width GPUs (NVIDIA L40S/L40/H100/H100 NVL PCIe)
- Supports Dual CPUs up to 350W TDP each

Key Applications

- AI Compute/Model Training/Deep Learning, HPC, DLRM (Deep Learning Recommendation Models)
- AI Inferencing and Digital Twin Creation/Omniverse
- High Performance Simulation of Complex 3D Graphics

Specifications

CPU – Dual Socket Dual Sapphire Rapids CPU (up to 350W TDP)	**Memory –** DIMM Slots 32x DIMM slots, ECC DDR5-4800MT/s
Drives – 8 Hot-Swap Bays 8x E1.S 2x M.2	**Expansion –** 7 PCIe Slots 7x PCIe 5.0 x16 (4 FHFL/DW, 3 for I/O)
I/O ports (Per Server) 1x RJ45 1GbE IPMI 1x mini-DP (needed converter to VGA), 2x USB	**Power Supply –** Redundant 3x 2000W Titanium Level Efficiency Power Supplies

Use Cases

ARS-111GL-NHR (Single Grace Hopper)

Industry: National Lab

Application: Climate Simulations, Genome Sequencing, Supernova Modeling

Territory: North America

Installation: PoC

- Optimized for density
- 1U with memory coherence between Grace Processor and H100 GPU
- Shared memory footprint of 576GB (480GB LPDDR5X + 96GB HBM3)

Use Cases

ARS-111GL-NHR (Single Grace Hopper)

AI and HPC

- The Grace Hopper Superchip is suitable for a broad array of tasks by utilizing a coherent pool of memory between the Grace CPU and H100 GPU.
- Delivers substantial speedups for AI workloads with high memory requirements

Diverse Workloads

- NVIDIA's full software stack is supported, including the NVIDIA HPC, NVIDIA AI, and NVIDIA Omniverse platforms.
- 1U form factor and air-cooling enables flexible deployments such as an Edge inference server

Use Cases

ARS-111GL-DNHR-LCC (Dual Grace Hopper)

Massive-Scale AI Training and Inference

- 96GB of HBM3 for GPU and 480GB of LPDDR5X for CPU totaling 576GB of memory (per node).
- Clustering through multiple BlueField-3 and ConnectX-7 at 400G for large language model (LLM) or recommender system training.

High-density HPC

- Balanced CPU and GPU performance in 1U
- Liquid-cooled for efficiency and reduced TCO for datacenters.

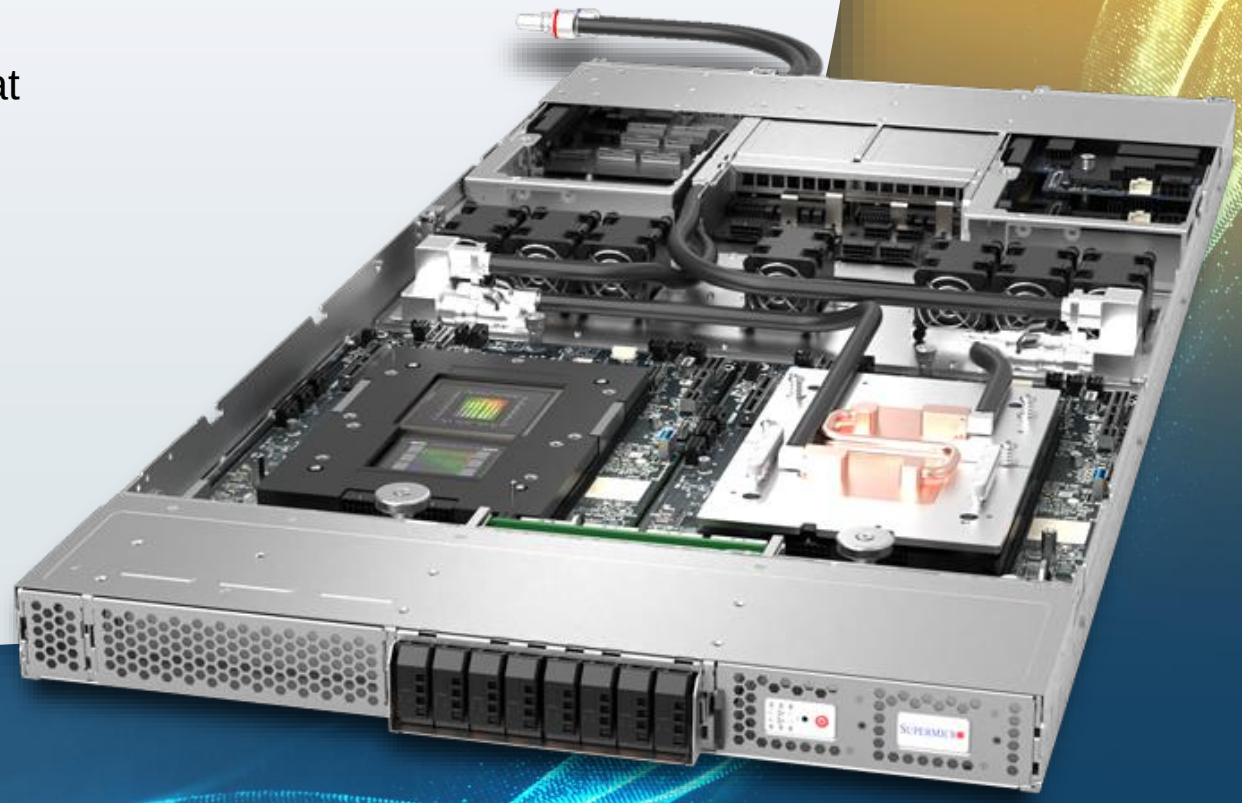

Use Cases

ARS-221L-DNR (2U Grace CPU)

Omniverse/Visualization

- Supports up to 4 double-width GPUs, such as NVIDIA L40S, for NVIDIA Omniverse OVX Server (remote visualization platform).
- Grace CPUs integrated high-density, high-bandwidth LPDDR5X accelerates a wide variety of complex 3D workflows and simulations with 2x the performance per watt vs. DIMM-based servers.

Sales Assets

MGX Landing Page, Datasheets, Sales Email Template, etc.

MGX Landing Page
www.supermicro.com/mgx

MGX Sales Training (US/EMEA/APAC)

1U Grace Hopper™ Superchip Systems

1U/2U NVIDIA Grace™ CPU Superchip and x86 Intel® Xeon® Systems

MGX Datasheets

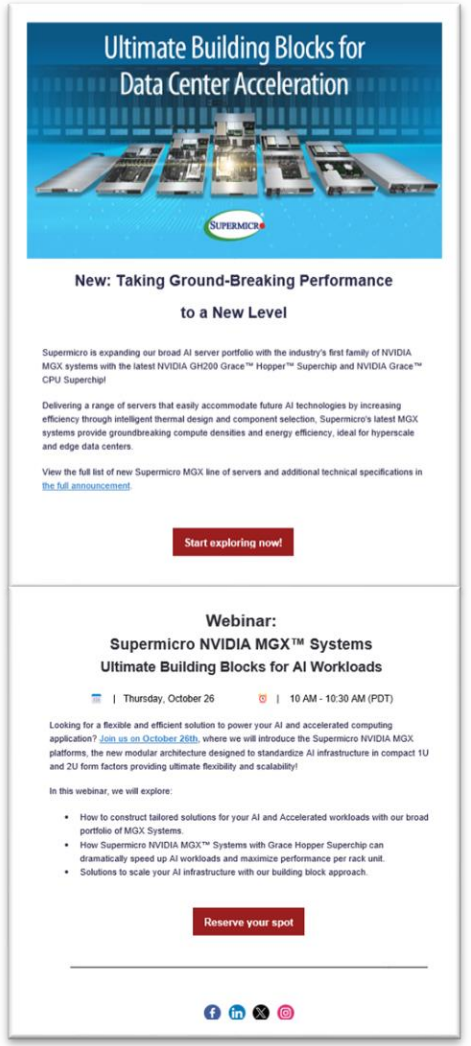

Sales Email Template

Promotions

Press Release, Webinar, eBlast, Social, etc.

Supermicro Starts Shipments of NVIDIA GH200 Grace Hopper Superchip-Based Servers, the Industry's First Family of NVIDIA MGX Systems

NEWS PROVIDED BY
Super Micro Computer, Inc. →
18 Oct, 2023, 16:58 ET

SHARE THIS ARTICLE

Supermicro's NVIDIA GH200 Superchip-Based Server Platform Increases AI Workload Performance Using a Tightly Integrated CPU and GPU and Incorporates the Latest DPU Networking and Communication Technologies

SAN JOSE, Calif., Oct. 18, 2023 /PRNewswire/ -- **Supermicro, Inc.** (NASDAQ: SMCI), a Total IT Solution manufacturer for AI, Cloud, Storage, and 5G/Edge, is announcing one of the industry's broadest portfolios of new GPU systems based on the NVIDIA reference architecture, featuring the latest NVIDIA GH200 Grace Hopper and NVIDIA Grace CPU Superchip. The new modular architecture is designed to standardize AI infrastructure and accelerated computing in compact 1U and 2U form factors while providing ultimate flexibility and expansion ability for current and future GPUs, DPUs, and CPUs. Supermicro's advanced liquid-cooling technology enables very high-density configurations, such as a 1U 2-node configuration with 2 NVIDIA GH200 Grace Hopper Superchips integrated with a high-speed interconnect. Supermicro can deliver thousands of rack-scale AI servers per month from facilities worldwide and ensures Plug-and-Play compatibility.

"Supermicro is a recognized leader in driving today's AI revolution, transforming data centers to deliver the promise of AI to many workloads," said Charles Liang, president and CEO of Supermicro. "It is crucial for us to bring systems that are highly modular, scalable, and universal for rapidly evolving AI technologies. Supermicro's NVIDIA MGX-based solutions show that our building-block strategy enables us to bring the latest systems to market quickly and are the most workload-optimized in the industry. By collaborating with NVIDIA, we are helping accelerate time to market for enterprises to develop new AI-enabled applications, simplifying deployment and reducing environmental impact. The range of new servers incorporates the latest industry technology optimized for AI, including NVIDIA GH200 Grace Hopper Superchips, BlueField, and PCIe 5.0 EDSFF slots."

Learn more about Supermicro's NVIDIA MGX systems at: www.supermicro.com/mgx

Press Release

Supermicro MGX Systems: Flexibility and Scalability for AI Workloads

Share Tweet Share Link

Webinar

Supermicro MGX Systems: Flexibility and Scalability for AI Workloads

Josh Grossman, Product Manager, Supermicro
Ivan Goldwasser, Director of Product Marketing, NVIDIA
Robert Moore, Marketing (Host), Supermicro

Josh Grossman, Supermicro; Ivan Goldwasser, NVIDIA; Robert Moore, Supermicro

About this talk

Looking for a flexible and efficient solution to power your AI and accelerated computing application? Join us for our webinar where we will introduce the Supermicro NVIDIA MGX platforms, the new modular architecture designed to standardize AI infrastructure in compact 1U and 2U form factors providing ultimate flexibility and scalability.

In this webinar, we will explore:

- How to construct tailored solutions for your AI and Accelerated workloads with our broad portfolio of MGX Systems
- How Supermicro NVIDIA MGX™ Systems with Grace Hopper™ Superchip can dramatically speed up AI workloads and maximize performance per rack unit
- Solutions to scale your AI infrastructure with our building block approach

Presenters:

- Josh Grossman, Product Manager, Supermicro
- Ivan Goldwasser, Director of Product Marketing, NVIDIA
- Robert Moore, Marketing (Host), Supermicro

Related topics:

Data Center | AI | Artificial Intelligence | Data Center Infrastructure | Data Center Management | NVIDIA

MGX Webinar

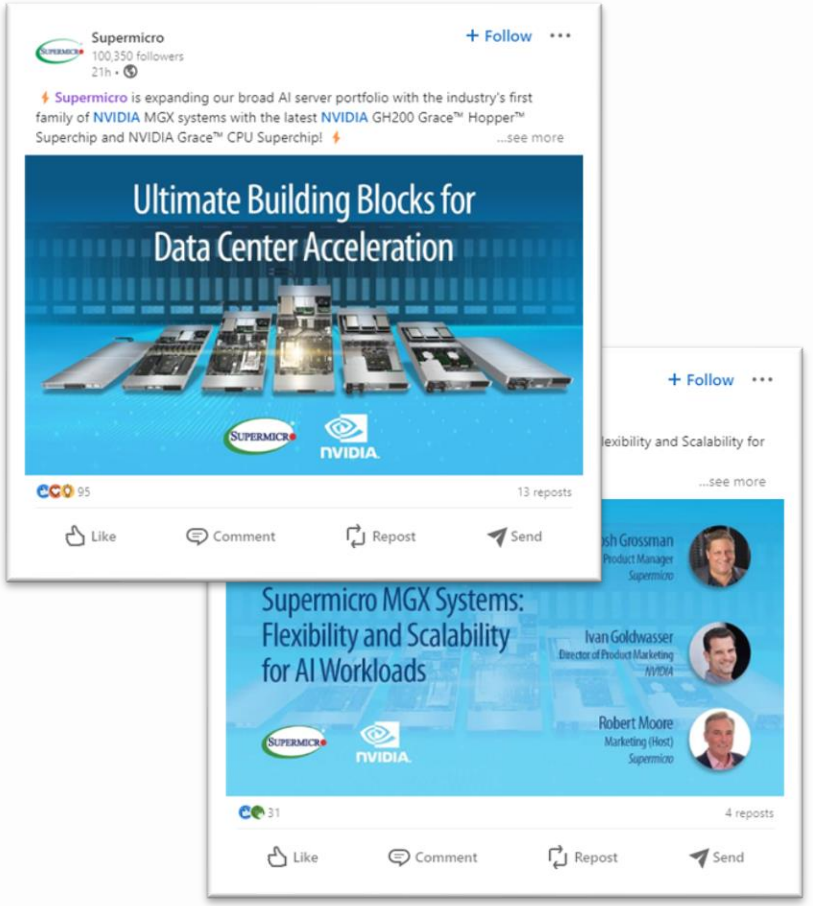

Supermicro
100,350 followers
21h •

+ Follow

Supermicro is expanding our broad AI server portfolio with the industry's first family of NVIDIA MGX systems with the latest NVIDIA GH200 Grace™ Hopper™ Superchip and NVIDIA Grace™ CPU Superchip! ...see more

Ultimate Building Blocks for Data Center Acceleration

95 13 reposts

Like Comment Repost Send

+ Follow

...see more

Supermicro MGX Systems: Flexibility and Scalability for AI Workloads

Josh Grossman, Product Manager, Supermicro
Ivan Goldwasser, Director of Product Marketing, NVIDIA
Robert Moore, Marketing (Host), Supermicro

31 4 reposts

Like Comment Repost Send

Organic/Paid Promo (eBlast, Social, Search Ads etc.)

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

Thank You

www.supermicro.com