Moderní infrastruktura pro AI

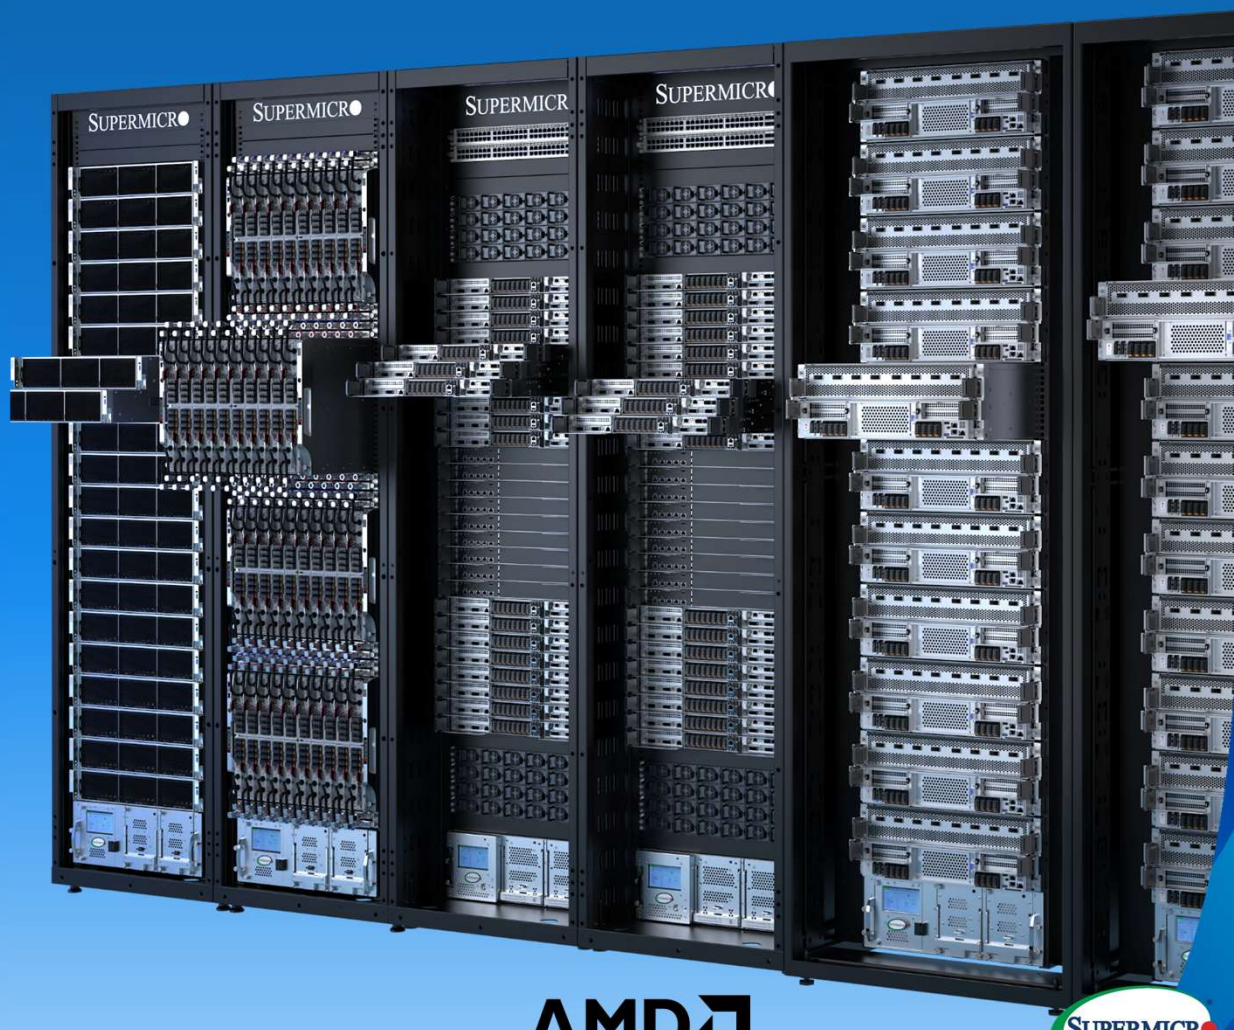

Petr Karbus
Senior Sales Manager | EER

May 2026

AMD

SUPERMICRO

About SUPERMICRO

Revenue	**$33B+** (FY2026) $22B (FY2025) $14.9B (FY2024) Build 1993
Worldwide Presence	6M+ Sq ft. Facilities Worldwide 1. Silicon Valley (HQ), 2. Taiwan, 3. The Netherlands, 4. Malaysia and others
Production	$50B/yr Production Capacity (CY26) ~5000 Air-cooled Server Racks/mo. ~2000 Liquid-cooled Server Rack/mo.
Human Resource in 4 Campuses	~7000 headcount Worldwide, ~50% Technical / R&D
Key Growth Matrix	**#1** in Generative AI and LLM Platforms 200%+ YoY Growth in Accel. Computing

AMD

SUPERMICRO

Business Progression

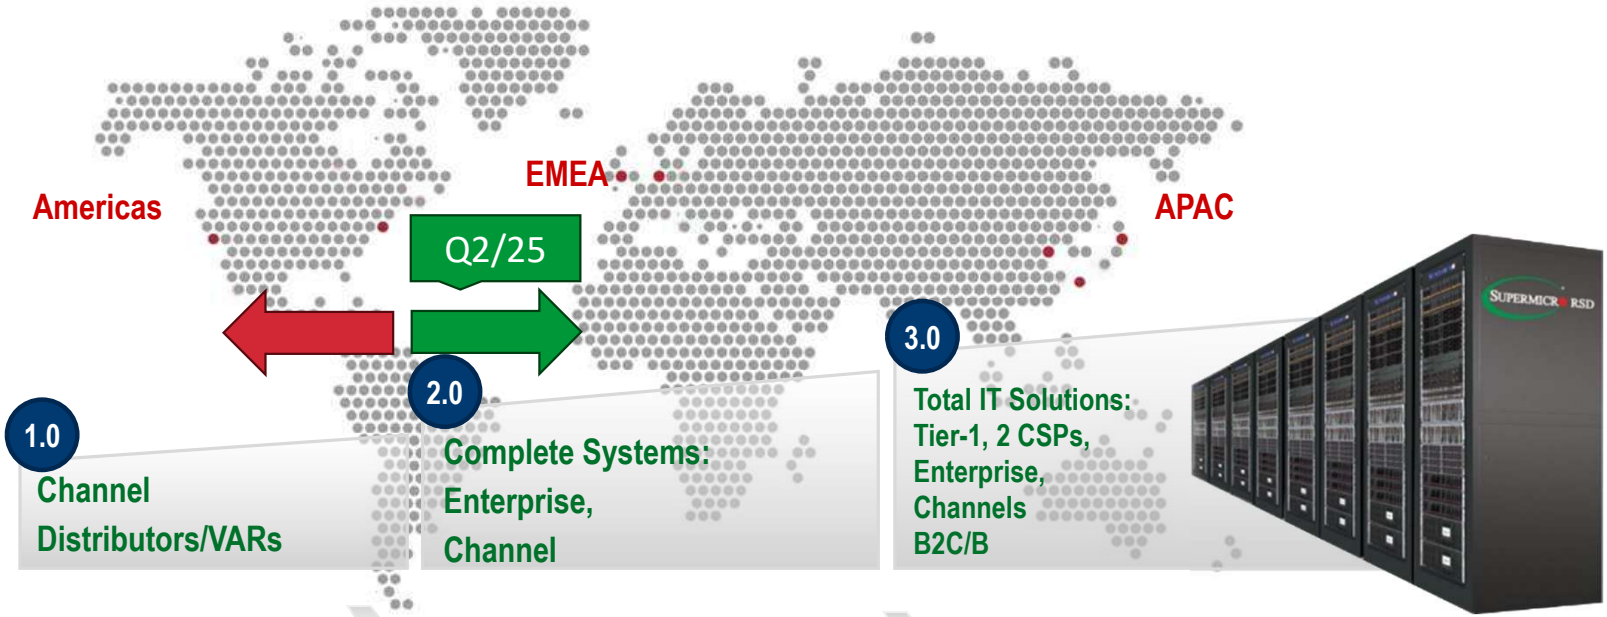

Market

Innovation

Subsystems

Motherboards
Chassis
Server Building Blocks

Server Systems

AI/ML Solutions
Twin Architectures
SuperBlade
Storage
Cloud Optimized
5G / Edge / IoT

Total Solutions

Rack-Scale/ Cloud
Management Software
Global Services
Resource-Savings
Comprehensive Software & Security Suites
Business Automation B2B/B2C Business Model

AMD

SUPERMICRO

Business Focus

Building Block Solutions

- Optimized subsystem design, including Motherboard, Power Supply, Chassis, Thermal solution, I/O, Firmware, BMC, Security and 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, direct liquid cooling (DLC) and free air cooling
- Save up to 40% TCO for many of our partners
- Industry-wide green computing adoption could save $10B per year of energy cost
- Save billions of trees

Business Automation

- Operation and sales automation to efficiently serve enterprise and any size of business
- Global manufacturing footprint and onsite deployments and technical support
- Command Center to provide 24/7 security and real-time monitoring

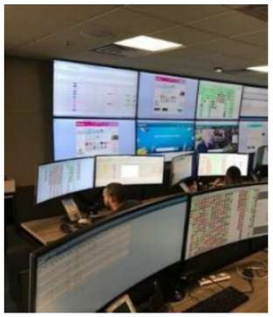

Supermicro 4.0 (DCBBS)

- Improve datacenter TTM/TTO 30% faster, Save $, Quality and Service
- Enhanced system/cloud software/application products
- DLC, networking, cabling and pre/post-deployment services for DC/Cloud total solutions

AMD

SUPERMICRO

Industry's Most Comprehensive Portfolio

Rack Mounted

Multi Processor
8 x CPU Sockets

Multi Processor
4 x CPU Sockets

Hyper
2 or 1 x CPU Sockets

CloudDC
2 or 1 x CPU Sockets

WIO
1 x CPU Socket

GPU

HGX
8 x GPU SXM

HGX DLC
8 x GPU SXM

PCIe GPU

MGX
Grace Hooper

Multi Node

Superblade

MicroCloud

FlexTwin

GrandTwin

BigTwin

Storage

All Flash

Front Load

Simply Double

Top Load

AMD

SUPERMICRO

Supermicro Solutions

GPU Acceleration (AI/ML, HPC, Omniverse)

Data Base & ERP

Cloud & Virtualization

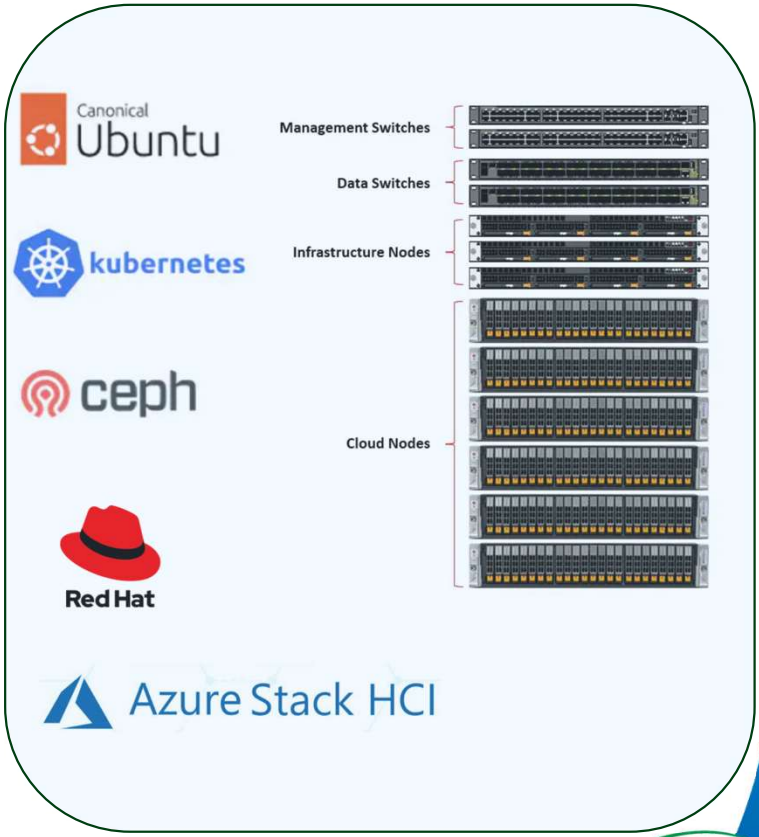

AMD

SUPERMICRO

Supermicro Solutions - Storage

All Flash

SSG-122B-NE316R

1U front-loading all-flash storage server with 16 E3.S NVMe drives and PCIe 5.0

MINIO

VAST

WEKA

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 HHHL 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;

OSNEXUS Quantum

Qumulo AMD

SUPERMICRO

Supermicro CPU Vendors

X14

H14

R13

NVDIA Grace CPU

G1

AMD

SUPERMICRO

A+ Server Solutions

World's Most Versatile Portfolio of AMD Processor-Based Systems
Supporting the Latest AMD EPYC™ series processors

Supermicro H14 Generation Servers

The Most Comprehensive Portfolio of AMD Processor-Based Systems,
now with AMD EPYC™ 9005 series processors, Including Servers, Storage, GPU-Optimized, and Multi-Node Solutions to Exactly Match System Requirements to Your Workloads

AMD

SUPERMICRO

Introducing H14 Generation A+ Servers

EPYC™ 9005 SERIES PROCESSORS

- Up to 192 cores 384 threads per socket
- Up to 6TB of memory of DDR5 with ECC 6400MT/s and Advanced Memory Device Correction (AMDC) and supports 2 DIMMs per channel (2DPC) with single socket
- PCIe 5.0 up to 160 lanes
- Next Generation Reliability, Availability, and Serviceability (RAS)

WORKLOAD OPTIMIZED SYSTEMS WITH OPEN ARCHITECTURES

- Vast I/O, storage, networking and expansion slot options for maximum versatility
- Flexibile networking options with Advanced I/O Modules (AIOMs), up to 400Gbps throughput per card and OCP 3.0 support
- Market-leading GPU optimized servers for large scale AI/ML and HPC workloads
- Compute Express Link (CXL 2.0) peripheral support including memory expansion through PCIe 5.0 lanes

INCREASED OPERATIONAL EFFICIENCY

- Tool-less chassis design
- Rear and Front I/O options
- Hot-swappable nodes with shared power for multi-node system
- Titanium level redundant power supplies
- Efficient resource-saving multi-node designs with shared power and cooling

AMD

SUPERMICRO

Introducing H14 Generation A+ Servers

H14 FlexTwin™ Liquid-Cooled System

2U Liquid-cooled multi-node Architecture for high-density and efficiency

H14 GPU Optimized PCIe System

Universal GPU optimized for AI/ Deep Learning and HPC

H14 Accelerated GPU Servers

Maximum Acceleration for AI Training and LLM

H14 Hyper System

Industry Leading IOPS Servers with Energy Efficiency and Flexibility

H14 CloudDC System

All-in-One Servers with Flexible I/O Options for Cloud Scale Data Centers

H14 GrandTwin® System

Leading Multi-Node Architecture with Front I/O

AMD

SUPERMICRO

Introducing H14 Generation A+ Servers

H14 Flextwin™ SYSTEM

Liquid-Cooled Multi-Node Server for Maximum Density and Power Efficiency

Per Node:

Dual AMD EPYC™ 9005/9004 series processors

Up to 24 DIMMs 6TB DDR5-6400 in 1DPC

1 OCP 3.0 compatible AIOM slot

Optional 2 E1.S drive bays

Up to 4 redundant 3200W Titanium Level power supplies

AS -2126FT-HE-LCC

Key Applications • HPC • Oil and Gas • Scientific Research

AMD

SUPERMICRO

Introducing H14 Generation A+ Servers

H14 5U PCIe GPU SYSTEM

Flexible, High Density GPU Systems for AI and HPC

- Dual AMD EPYC™ 9005/9004 series processors
- Support up to 10 double-width PCIe accelerators
- Up to 24 DIMMs 6TB DDR5-6400 in 1DPC
- Up to 13 PCIe 5.0 x16 FHFL slots
- Up to 8 front hot-swap 2.5" NVMe drive bays

AS -5126GS-TNRT

Key Applications • AI • HPC • 3D Rendering Farm • Virtualization • Research

Introducing H14 Generation A+ Servers

H14 8-GPU SYSTEMS

Next-Gen Large Scale AI Training Platform

Industry standard OCP Accelerator Module (OAM) with 8 accelerators interconnected on an AMD Instinct™ and NVIDIA HGX™

Industry-leading up to 2TB HBM3e GPU memory in a single server node

1:1 400G networking dedicated for each GPU designed for large scale AI and supercomputing clusters

Dual AMD EPYC™ 9005/9004 series processors

Up to 24 DIMMs for up to 6 TB of DDR5-6400 memory

AS -8126GS-TNMR

Key Applications • Large Scale Deep Learning • Generative AI and Large Language Model Training • AI-fused HPC applications • Industrial Automation • Business Intelligence & Analytics

AMD

SUPERMICRO

Introducing H14 Generation A+ Servers

H14 Hyper SYSTEM

Industry Leading IOPS Rackmount Servers with Energy Efficiency and Flexibility

Dual AMD EPYC™ 9005/9004 series processors

Up to 24 DIMMs 6TB DDR5-6400 in 1DPC

3 PCIe 5.0x 16 (1U),
up to 4 PCIe 5.0 x16 or 8 PCIe 5.0 x8 (2U)

Capable to support GPU and CXL 2.0

Up to 12 front hot-swap 2.5" NVMe/ SAS/SATA drives (1U)

Up to 24 front hot-swap 2.5" NVMe/ SAS/SATA drives (2U)

Flexible networking options with
1 AIOM/OCP 3.0 support

AS -1126HS-TN

Key Applications •Virtualization • Software Defined Storage • Enterprise Server • AI Inference • HPC • Cloud Computing

AMD

SUPERMICRO

Introducing H14 Generation A+ Servers

H14 cloud dc SYSTEM

All-in-One Servers with Flexible I/O Options for Cloud-Scale Data Centers

Single AMD EPYC™ 9005/9004 series processor

Up to 12 DIMMs 3TB DDR5-6400 in 1DPC

Designed with DC-MHS compliance

2 PCIe 5.0 x16 FHHL slots

Flexible networking option with
1 AIOM/ OCP3 support

Up to 12 front hot-swap 2.5"
NVMe/SAS/SATA drives

AS -1116CS-TN

Key Applications • Cloud Computing • Web Server • Hyper-converged Storage • Virtualization, File Servers • Head-node Computing • Telcom Security Server • CDN

AMD

SUPERMICRO

Introducing H14 Generation A+ Servers

H14 Grandtwin® SYSTEM

Leading Multi-Node Architecture with Front I/O

Per Node:

Single AMD EPYC™ 9005/9004 series processor

Up to 16 DIMMs 4TB DDR5-4400 in 2DPC

Up to 4 front hot-swap 2.5" NVMe/SATA drives

Flexible networking with optional AIOM slots

Redundant 2200W Titanium Level power supplies

AS -2116GT-HNTF

Key Applications • HPC • Telco Edge • Cloud Computing • EDA • CDN • Web Hosting Applications

AMD

SUPERMICRO

Supported GPU By Supermicro

NVIDIA | AMD | intel

Multi socket

HGX
H100
H200
B100
B200
B300

CDNA3
MI300X
MI350X

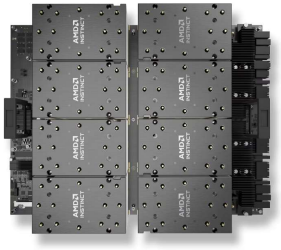

Gaudi3 UBB

PCIe

H100 NVL
L40S
L4
RTX 6000 Ada

MI210

Gaudi3 PCIe

CPU+GPU

Grace Hooper
GH200
GB200
GB300

CDNA3
MI300A
MI350

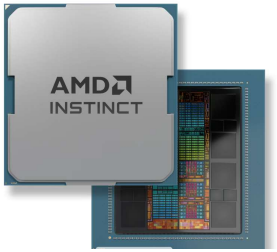

AMD

SUPERMICRO

GPU Accelerated Workloads

AMD

SUPERMICRO

Building AI Infrastructure

What is AI Infrastructure?

- *Compute GPU nodes*
- *Fast Interconnect (Network)*
- *Supporting Sys (Storage, MNGM)*
- *Orchestration Tools (Cluster Management, Cloud & Virtualization)*

What do we need for AI Infrastructure?

- *Concept*
- *Planning*
- *Data Center*

AMD

SUPERMICRO

Why Use GPU for Workloads?

1.Parallel Processing Power: GPUs are designed to handle multiple tasks simultaneously, making them highly efficient for parallel computations. In deep learning, many operations (like matrix multiplications) can be parallelized, which GPUs excel at due to their architecture with numerous cores.

2.High Performance: GPUs are optimized for handling large amounts of data and performing complex calculations quickly. They can process thousands of arithmetic operations in parallel, significantly speeding up model training compared to CPUs.

3.Deep Learning Framework Support: Most deep learning frameworks (like TensorFlow, PyTorch, and MXNet) are designed to leverage GPU acceleration. They have libraries that automatically distribute computations across multiple GPU cores, maximizing performance.

4.Memory Bandwidth: GPUs have high memory bandwidth, allowing them to efficiently handle the large amounts of data involved in deep learning tasks. This helps prevent bottlenecks that can slow down training on CPUs.

5.Specialized Architectures: Modern GPUs often include specialized cores and features specifically tailored for deep learning tasks, such as Tensor Cores for accelerated matrix operations (e.g., in NVIDIA GPUs).

6.Cost-Effectiveness: GPUs can offer significant speedups in model training time compared to CPUs. This means that training large models or processing extensive datasets can be done more quickly, potentially reducing overall training costs in terms of time and resources.

AMD

SUPERMICRO

Supermicro Networking - Ethernet

SSE-T7132S/SR (32 ports)

Key Features

- 32x 400Gbps Ethernet ports (QSFP-DD)
- SONiC Networking Operating System
- Fully shared packet buffering
- Redundant hot-pluggable power supplies
- 1U form factor ideal for spine/super-spine
- Regular and reverse airflow models

SSE-C4632SB/SRB (32 ports)

Key Features

- 32 x100Gbps Ethernet ports (QSFP28)
- 1:1 Non-blocking connectivity
- 1U form factor for flexible installation
- Data-Center friendly - regular and reverse airflow models
- Hot-pluggable power supplies
- Broadcom Advanced Enterprise SONiC Switch Software pre-installed

SSE-T8032S

Key Features

- 64x 400Gbps Ethernet ports
- Broadcom Advanced Enterprise SONiC Networking Operating System
- Fully shared packet buffering
- Redundant hot-pluggable power supplies
- 1U form factor ideal for leaf/spine/super-spine
- Regular airflow model

SSE-SN3700-VS2

Key Features

- 32 x 200 Gbps Ethernet ports (QSFP56)
- Connectivity at different speeds with throughput of 12.8Tb/s
- Cumulus Linux Networking Operating System
- Fully shared packet buffering
- Best-in-class VXLAN scale
- Redundant hot-pluggable power supplies
- 1U form factor ideal for ToR super spine

SMC networking

- One time investment, no additional costs (ports, data transfer..)
- Standard & Reverse airflow models,
  OS Sonic management

AMD

SUPERMICRO

Supermicro Networking - Infiniband

System Specifications	
Performance	400Gb/s per port
Switch radix	64 400Gb/s non-blocking ports with aggregate data throughput up to 51.2Tb/s
Connectors and cabling	32 octal small form-factor pluggable (OSFP) connectors; passive or active copper or active fiber cable; optical module

System Specifications			
	Q3200-RA	**Q3400-LD**	**Q3400-RA**
Performance	Two switches, each of 28.8Tb/s throughput	115.2Tb/s throughput	115.2Tb/s throughput
Switch radix	Two switches, each of 36 800Gb/s non-blocking ports	144 800Gb/s non-blocking ports	144 800Gb/s non-blocking ports
Connectors and cabling	Two groups of 18 OSFP connectors	72 OSFP connectors	72 OSFP connectors

HPC usage

- High throughput, low latency

AMD

SUPERMICRO

AI Infrastructure Complete Solution from Supermicro

Supermicro Rack-Scale Advantage:

Leverage Proven Building Blocks

Simplify Cluster-Scale Architecture

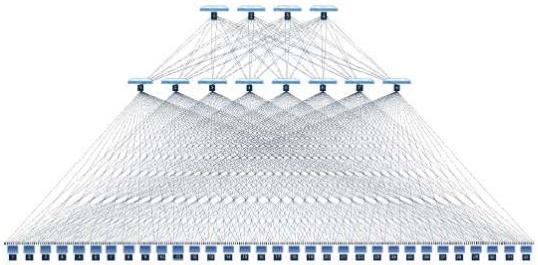

Deploy Plug & Play Racks

Rack Solution Design & Deployment Steps:

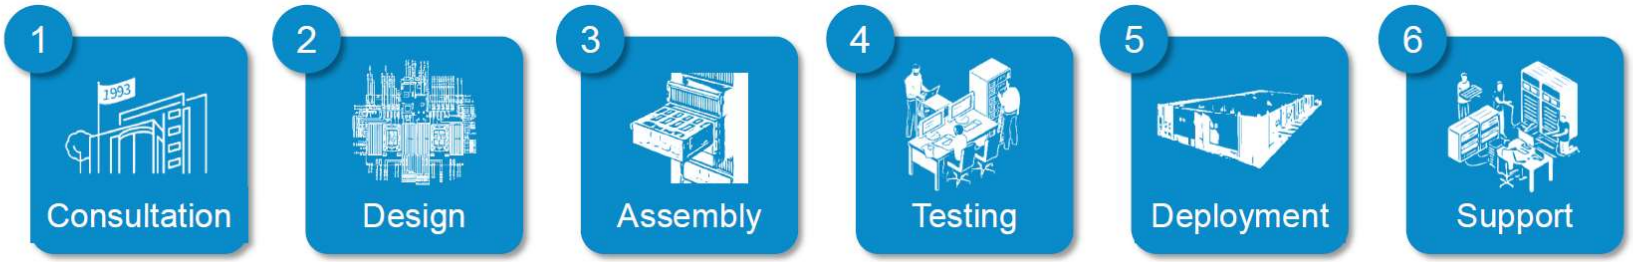

AMD

SUPERMICRO

Security Management Feature Overview

AMD

Product Security Landscape

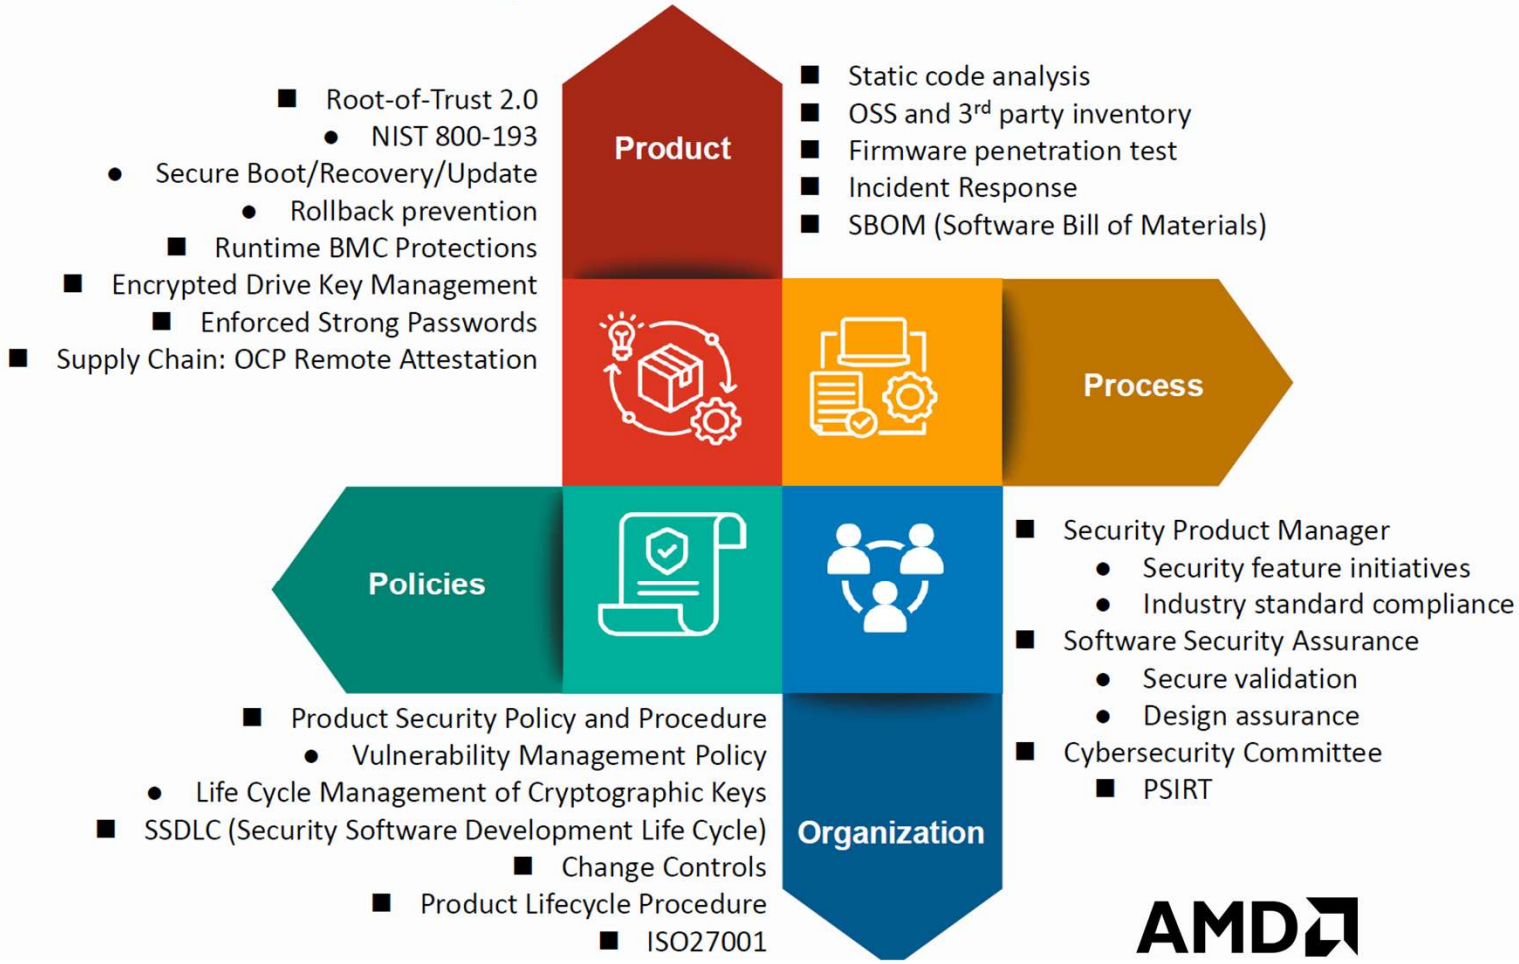

- Root-of-Trust 2.0
  - NIST 800-193
  - Secure Boot/Recovery/Update
  - Rollback prevention
- Runtime BMC Protections
- Encrypted Drive Key Management
- Enforced Strong Passwords
- Supply Chain: OCP Remote Attestation

- Static code analysis
- OSS and 3rd party inventory
- Firmware penetration test
- Incident Response
- SBOM (Software Bill of Materials)

- Security Product Manager
  - Security feature initiatives
  - Industry standard compliance
- Software Security Assurance
  - Secure validation
  - Design assurance
- Cybersecurity Committee
  - PSIRT

- Product Security Policy and Procedure
  - Vulnerability Management Policy
  - Life Cycle Management of Cryptographic Keys
- SSDLC (Security Software Development Life Cycle)
- Change Controls
- Product Lifecycle Procedure
- ISO27001

AMD

Product Security Feature Overview

Confidential

Category	Security Specifications	Intel X14 Products	AMD H14 Products	Intel X13 Products	AMD H13 Products	Intel X12 Products	AMD H12 Products
Hardware	Silicon Root of Trust	✓	✓	✓	✓	✓	✓
	Chassis Intrusion Protection	✓	✓	✓	✓	✓	✓
	Trusted Platform Module (TPM) 2.0	✓	✓	✓	✓	✓	✓
BIOS/BMC	Secure Boot	✓	✓	✓	✓	✓	✓
	Secure Drive Erase *	✓	✓	✓	✓	✓	✓
	Secure Flash	✓	✓	✓	✓	✓	✓
	Secure Firmware Updates	✓	✓	✓	✓	✓	✓
	Cryptographically signed firmware	✓	✓	✓	✓	✓	✓
	Secure Redfish API[3]	✓	✓	✓	✓	✓	✓
	Password Security	✓	✓	✓	✓	✓	✓
	USB dynamic enable/disable			✓		✓	
	HDD Password *	✓	✓	✓	✓	✓	✓
	BMC Unique Password	✓	✓	✓	✓	✓	✓
	Automatic Firmware Recovery	✓	✓	✓	✓	✓	✓
	Anti-rollback	✓	✓	✓	✓	✓	✓
	Runtime BMC Protections	✓	✓	✓	✓		
	System Lockdown	✓	✓	✓	✓	✓	✓
	Supply Chain Security: Remote Attestation	✓	✓	✓	✓		
	Drive Key Management (Super-Guardians) *	✓	✓	✓	✓		
	IPMI 2.0 User Locking	✓	✓	✓	✓		
	Security State Monitoring	✓	✓	✓	✓		
	Security Protocol and Data Model (SPDM) Management[2]	✓	✓	✓	✓		
Standards	NIST SP 800-193	✓	✓	✓	✓	✓	✓
	NIST SP 800-147b	✓	✓	✓	✓	✓	✓
	NIST SP 800-88	✓	✓	✓	✓	✓	✓

Certain features may not apply to all products. [1] On select models and configurations. [2] Planned on Q4 CY2024. [3] TLS v1.2 / v1.3 supported. RMCP+ Cipher Suite3 / 17 supported. * BIOS Feature for Host/Server

https://www.supermicro.com/en/support/security_center#!product-security

AMD

Supermicro Service & Support

Supermicro Global Onsite Services

Warranty
No SLA
Standard Inventory Pool
Supermicro Regional HQ
One-Way shipping In-Region
Technical support provided during business hours at Supermicro Regional HQ
Technical support CRM
Depot based Repair at Supermicro Regional HQ
No Technical Account Manager

Onsite Service		Bronze SMSADx	Silver OSNBDx	Gold OS4HRx
SLA		3 Business Day	Next Business Day	4-hour
Parts	Separate Stock to meet SLA	●	●	●
Local Depot	In-Country or Region to meet SLA	●	●	●
Advanced Logistics	2-Way Shipping Import/Export, & VAT.	●	●	●
24x7 Service Desk	Entitlement checking, Guided Troubleshooting Phone Support Partner escalation	●	●	●
Web Portal	Ticket submission, Case tracking & Routing	●	●	●
Field Engineer	Global Availability Component replacement	■	●	●
Technical Account Manager	Account Management Single point of contact Escalation Support QBR Review		●	●

Additional Onsite Services

*requires Onsite Service code

Liquid Cooling & Preventive Maintenance	Service Renewals	Managed Services	Multi-Vendor Product Service
Buyback Program	Digital Media Retention	Onsite Integration	Supermicro University

● *Included with Service*
■ *leverages Customer Staff*
x *describes service term in years*

AMD

SUPERMICRO

Supermicro Ready-to-Deploy Gold Series

Supermicro Gold Series systems take the guesswork out of enterprise server acquisition

Purchase with Confidence

Based on our most popular configurations, Supermicro Gold Series are pre-configured for specific workloads. These configuration have been pre-tested and are ready to go from day one.

Short Lead Times

Gold Series produce are pre-configured and ready to ship. No need to wait for parts and assembly. In most cases, Gold Series systems will ship from Supermicro's warehouse on the next business day.

Effortless Deployment

Systems are delivered to the customer with components already installed and pre-tested.
Unpack. Rack. Powerup.
Done.

AMD

SUPERMICRO

Gold Series Products

Enterprise Compute

SYS-112H-TN-01-G2
SYS-112H-TN-02-G2

SYS-212H-TN-01-G2

AS -2115HS-TNR-01-G2

AS -1115CS-TNR-01-G2

AS -2115GT-HNTR-01-G2

SBI-612B-1NE34-01-G2

Standard form factors designed for enterprise applications, cloud data centers, and general compute workloads

Enterprise Storage

AS -2015CS-TNR-01-G2

ASG-2015S-E1CR24L-01-G2

SSG-542B-E1CR60-01-G2

SSG-542B-DE1CR90-01-G2

Storage solutions for a range of performance and density requirements

Enterprise AI

SYS-822GS-NB3RT-01-G2

AS -8126GS-NB3RT-01-G2

SYS-212GB-FNR-01-G2

SYS-422GA-NRT-01-G2

Pre-configured with GPUs to accelerate AI inference and training workloads

Intelligent Edge

SYS-111AD-HN2-01-G2

ARS-E103-JONX-H2-01-G2

AS -E300-14GR-01-G2

SYS-E300-14AR-01-G2

Compact systems for the intelligent edge, ready to be deployed in remote environments or integrated into specialized equipment

AMD

SUPERMICRO

Contact details:

Petr Karbus
Senior Sales Manager EER
+420 602 225 313
petrk@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

AMD

SUPERMICRO

www.supermicro.com

AMD