



Vince Chen  
Sr. Director, Solution Architecture  
Supermicro  
August 8, 2024

# Optimized Storage for Scalable AI Applications



# Agenda

- AI Storage Requirements
- All Flash / Multi-tier Storage Solution
- Supermicro's Storage Products
- Summary

# Supermicro Rack-Scale AI Total Solutions



**5,000 Racks** capacity per month expected by the end of FY 24 of which....

- **2,000+ Racks will be Direct Liquid Cooled**
- **100KW+ Racks** now ready to deploy

## Total IT Solutions for Better Customer Value Hardware + Software + Services

- Rack-scale plug-and-play solutions with Building Block Architecture
- Scalable optimized AI (Generative AI) & general compute, storage, switch, infrastructure cooling, software & service
- Free-air, liquid-cooling and traditional air-cooling technologies provide customers the optimized datacenter solution
- Leading energy efficiency and lower TCO
- One-stop shop

Server software management solutions drive optimization and higher infrastructure security



# AI Storage for AI Factory or AI Cloud

- Large scale AI/GPU cluster infrastructures. Scale by Scalable Unit (SU, aka Pod, e.g. 256 GPUs)
- Highest performance AI training and inference.
- Storage requirements
  - All NVMe
  - Per SU performance requirement:  
Read 125 GBps / Write 62 GBps
  - InfiniBand or Ethernet/RDMA at 200Gb/s or 400Gb/s.
  - Parallel filesystem, High performance Object storage



Table 6. Guidelines for storage performance

Performance Characteristic	Good (GBps)	Better (GBps)	Best (GBps)
Single-node read	4	8	40
Single-node write	2	4	20
Single SU aggregate system read	15	40	125
Single SU aggregate system write	7	20	62
4 SU aggregate system read	60	160	500
4 SU aggregate system write	30	80	250

NVIDIA DGX SuperPOD storage requirements reference

<https://docs.nvidia.com/dgx-super-pod-reference-architecture-dgx-h100.pdf>

# AI Storage for Enterprises

- Pod level deployments
- Enterprise use cases, inference vs training
- Storage requirements
  - All NVMe or Tiered in PB scale.
  - Balance between Performance and TCO.
  - > 100 Gbps, Ethernet preferred.
  - RDMA optional
  - Parallel filesystem, High performance Object storage



Where AIOps and MLOps are performed

- GPU-dense servers accelerate AI training and inferencing
- GPUDirect storage transfers directly from/to GPU memory and eliminates CPU or main memory overhead
- Supermicro offers a range of GPU-dense servers to meet every level of application need

## Supermicro-Powered Application Tier

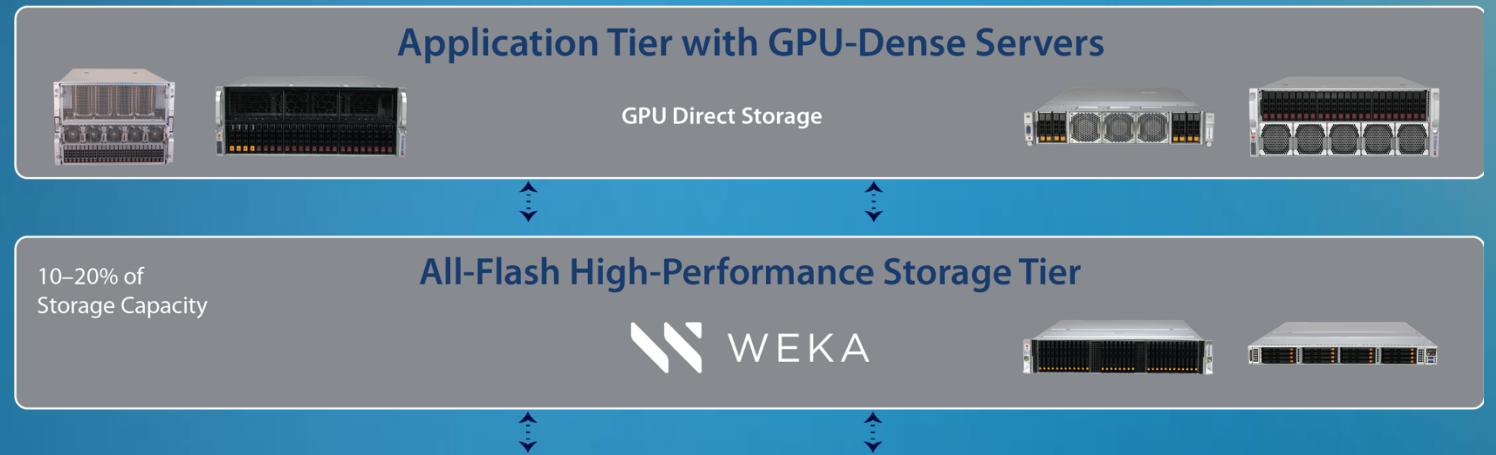


No vendor offers more choices for GPU-accelerated computing

Active data is stored here

- Optimized for performance delivers data to applications as fast as they can consume it
  - More utilization -> Better ROI
- Supermicro Peta-scale storage systems
  - 1U and 2U servers using the latest E3.S NVMe storage
- WEKA Data Platform
  - Scale-out, hierarchical storage solution
  - Clustered storage solution
  - Data protection and performance

## High-Performance All-Flash Tier

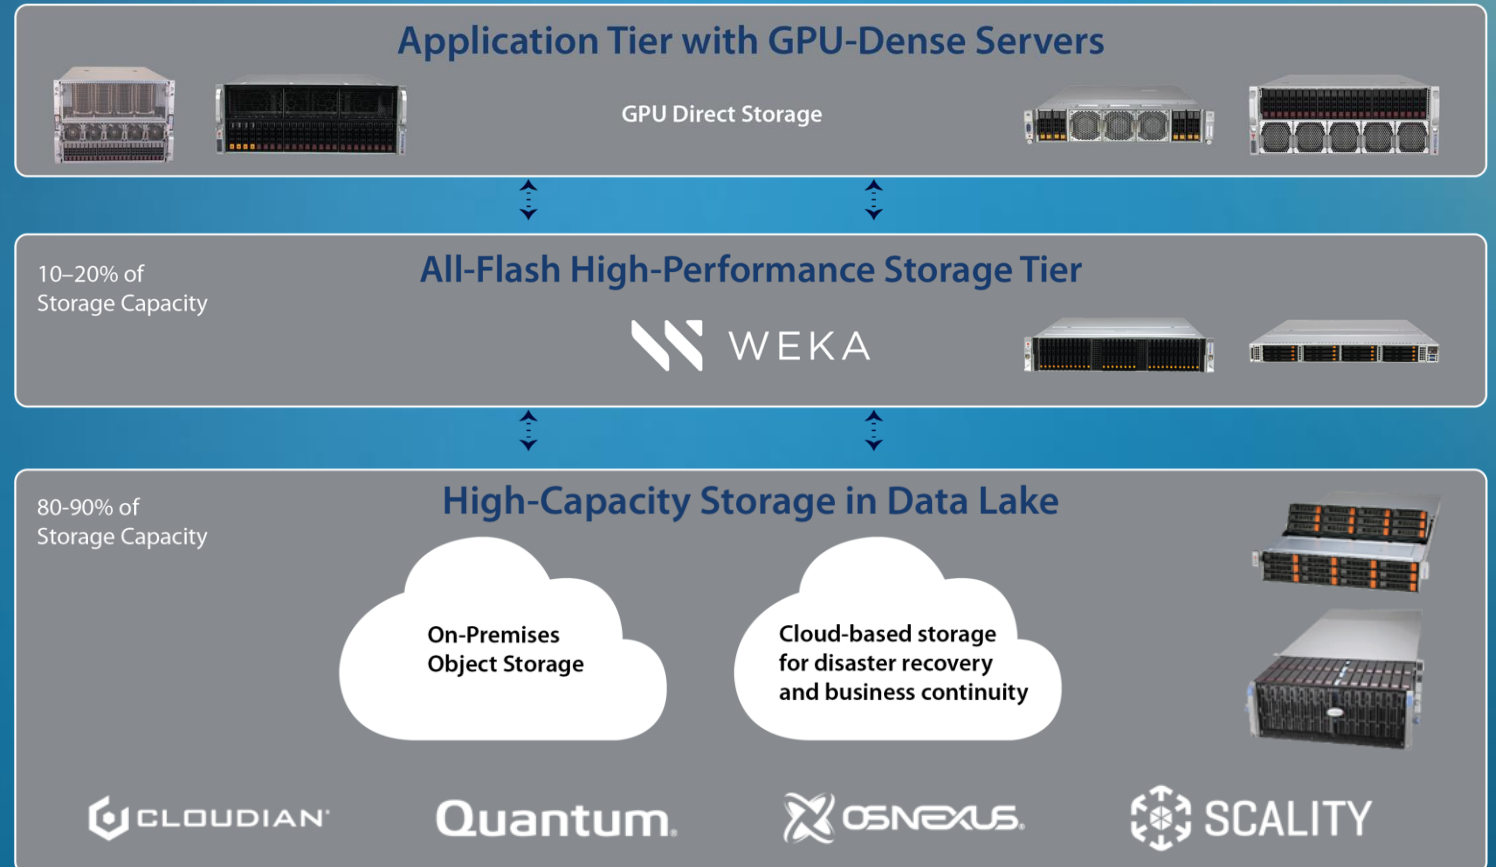


We worked with WEKA engineers to optimize for Supermicro storage servers

# High-Capacity Data Lake

All training data sets and models stored on-premises

- Data lake uses capacity-optimized storage
- High-capacity spinning-disk storage at lower cost per TB
- Capable back up & Tiering to the private cloud
- Supermicro servers join a scale-out cluster supported by 3<sup>rd</sup> party software partners



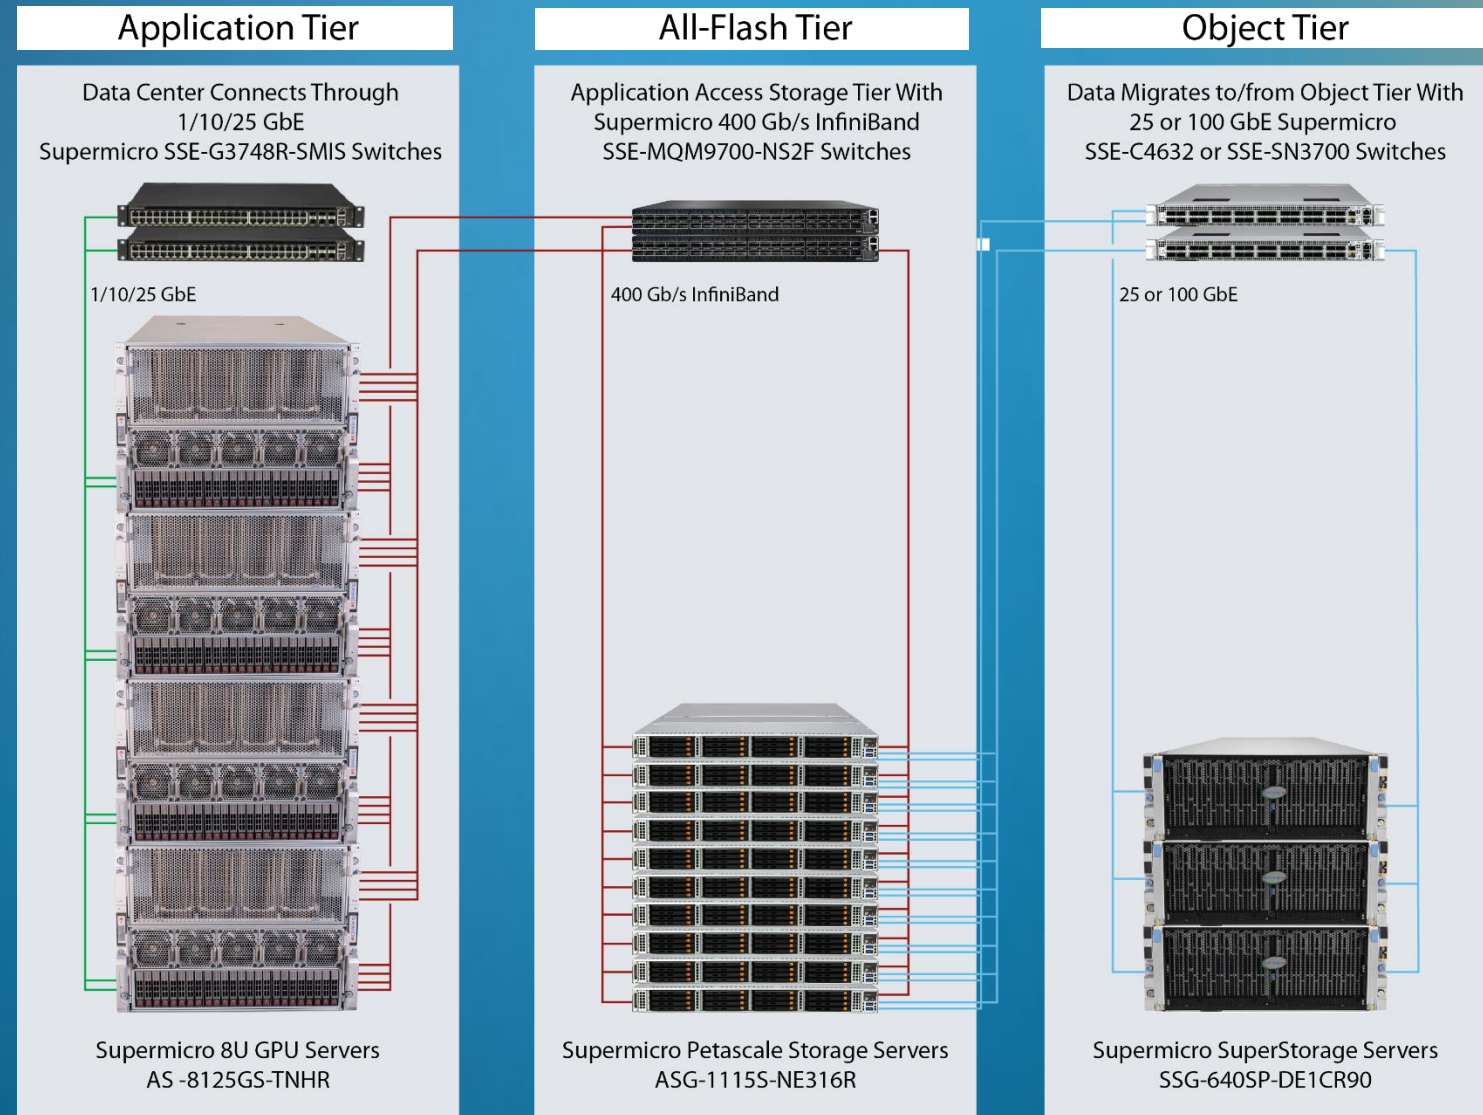
The key to cost-effectively storing all your data, safely, on premises



# Tested and Validated Solution

- Flexible application tier
  - Range of GPU server choices
- Flexible all-flash tier
  - Petascale storage servers
- Capacity-optimized object tier
  - SuperStorage servers
- Interconnections
  - 100–400-Gbps InfiniBand or Ethernet to application tier
  - 25 or 100 GbE for object tier
- Designed for Ease-of-Use for AIOps / MLOps
- Tailored design, integration, testing and deployment to meet exact customer requirements.

Engineered for AI-powered manufacturing defect analysis



# Supermicro Storage Server Portfolio Overview

From High-Capacity 90-Drive Systems to High-Performance All-Flash NVMe Systems

Capacity

## Enterprise



Featuring up to 36 LFF drives, front/rear drive access for enterprise storage building block.

## Top Loading



Featuring up to 90 LFF drives and single or twin controller for cloud scale object and file storage.

## Simply Double



Featuring up to 24 LFF and 4 NVMe drives for hybrid performance storage.

## High Availability SBB



Featuring up to 24 dual port NVMe and dual canister design for high availability storage.

## Petascale All-Flash

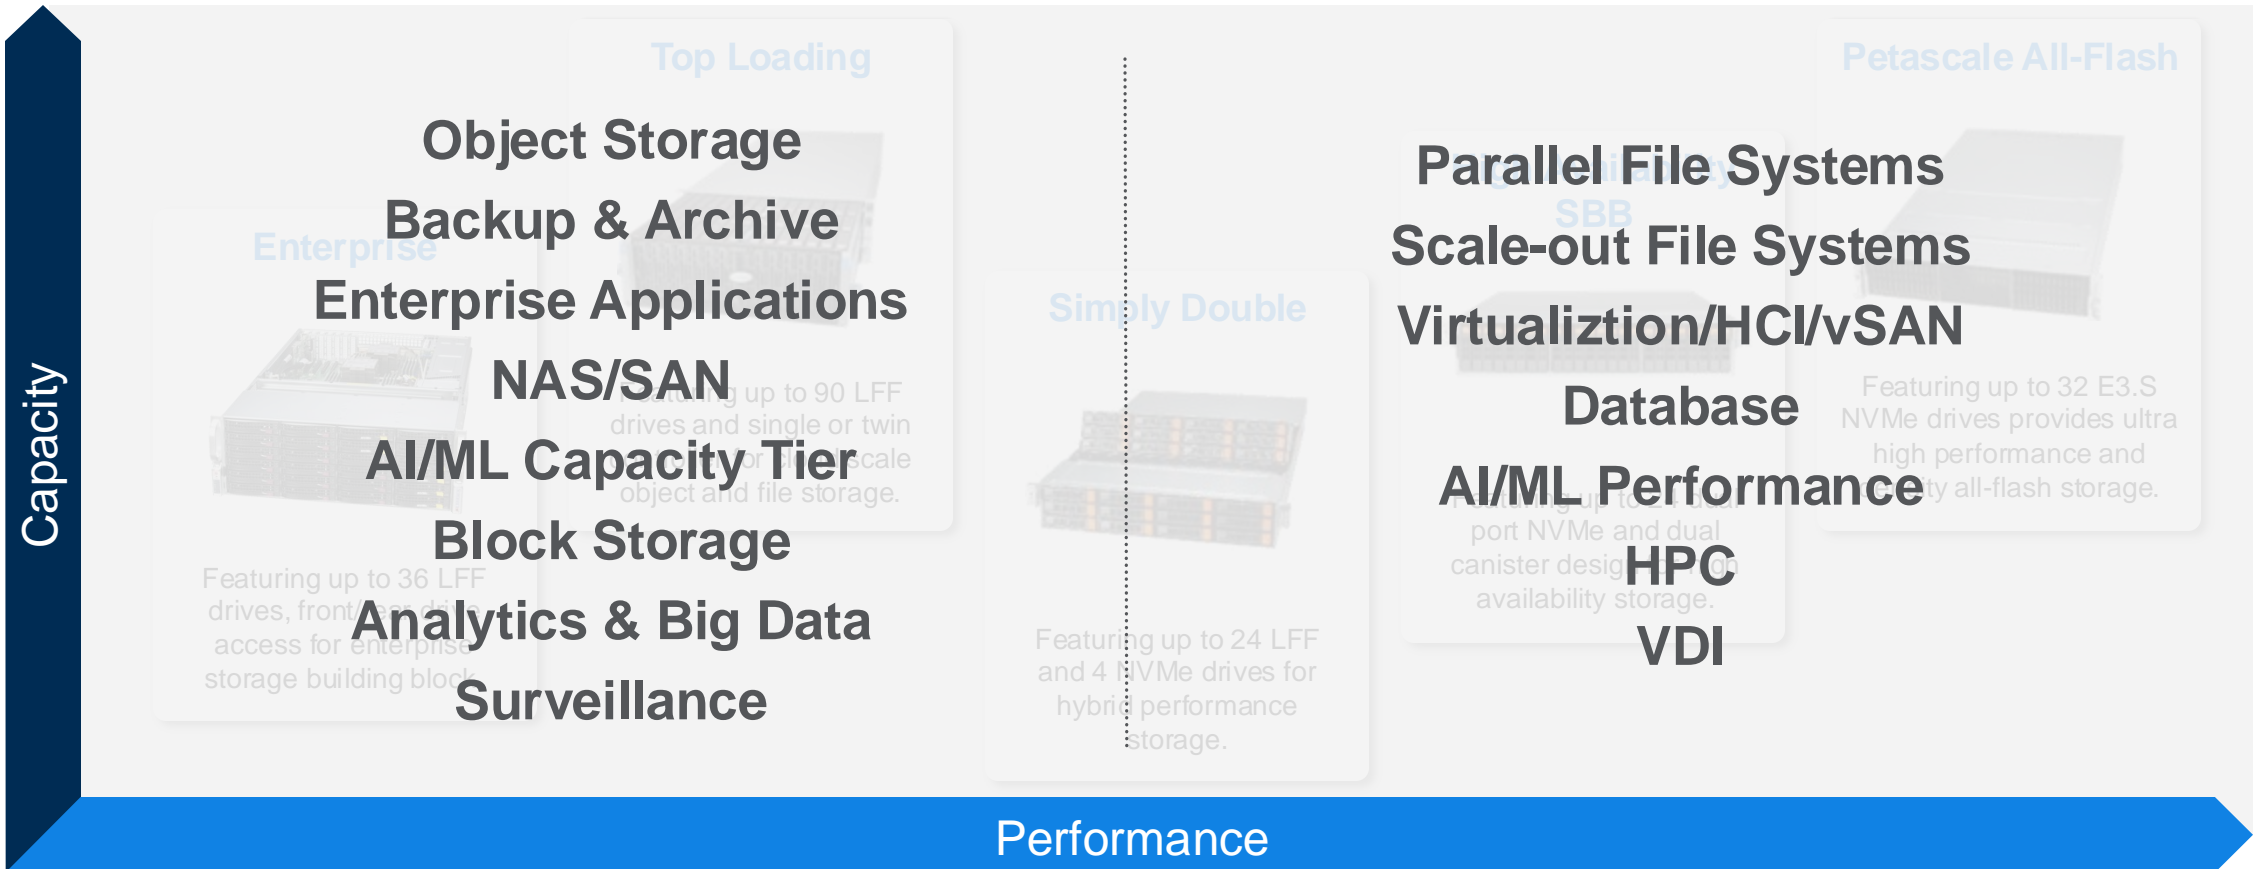


Featuring up to 32 E3.S NVMe drives provides ultra high performance and density all-flash storage.

Performance

# Supermicro Storage Server Portfolio Overview

From High-Capacity 90-Drive Systems to High-Performance All-Flash NVMe Systems



# Supermicro PetaScale All-Flash Storage Systems



## • Key Features


- Support dual Intel and single AMD single PCIe Gen5 Processor
- Support 1U E1.S and E3.S and 2U E3.3 TLC, QLC and CXL device and up to 1PB all flash in 2U
- Optimized thermal design with EDSFF design
- Balanced PCIe lanes for front SSD and rear IO design
- Up to 30M 4KB RR IOPS and 230+ GB/s 128KB SR BW.

## • Target Solutions


- High performance AI SDS building blocks
- Virtualization & dense VDI
- High performance object storage
- Hyperconverged infrastructure
- High performance CDN and video streaming



# AI Storage White Paper



- This paper is available at [www.supermicro.com/en/solutions/ai-storage](http://www.supermicro.com/en/solutions/ai-storage)



**Thank You**



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

© 2021 Super Micro Computer, Inc. All rights reserved.