SOLUTION BRIEF

SUPERMICRO AND CORNELIS ACCELERATE HPC WORKLOADS WITH FLEXTWIN™ AND CORNELIS CN5000 OMNI-PATH® 400G SOLUTION

Figure 1 - Supermicro FlexTwin and Cornelis CN5000 400G Solution

TABLE OF CONTENTS

Executive Summary 1
High Performance Computing Workloads 2
Joint Solution Overview 2
Customer Benefits 4
The Power of Liquid Cooling: Efficiency 5
Ideal Workloads 7
Results & Analysis 6
Summary 8
Appendix 9
For More Information 9

Executive Summary

Modern High-Performance Computing (HPC) deployments demand ever-greater compute density and energy efficiency — requirements that far exceed the limits of traditional air-cooled systems. Today, HPC success is not only measured by performance per dollar but increasingly by performance per dollar per watt, as data centers face tighter power and cooling constraints.

The Supermicro FlexTwin is a purpose-built HPC platform engineered not only for the best performance per dollar but also for energy efficiency. Designed for high performance and liquid-cooled efficiency, FlexTwin enables organizations to achieve higher compute density, lower total cost of ownership (TCO), and energy efficiency in demanding HPC environments.

The advanced liquid-cooling solution in FlexTwin enables Supermicro to achieve up to 95% heat-capture efficiency. This cold plate directly contacts and removes thermal energy from the highest thermal design power (TDP) components: the processors, memory, voltage regulator module (VRM), and the add-on-card (AOC) networking cards at the rear of the chassis.

A vital new addition to the cold plate is the direct cooling for the rear AOC networking cards. This addition is essential for managing the high heat generated by 400Gbps low-latency network fabrics, which are core to maximizing the performance of HPC environments and large-scale clusters. By collaborating with Cornelis® Omni-Path®, we were able to validate and showcase the performance-per-dollar you can achieve when combining FlexTwin advanced liquid cooling with the Cornelis® CN5000 solution.

High Performance Computing Workloads and Use Cases

Engineered for maximum performance-per-dollar-per-watt, the Supermicro FlexTwin system is the liquid-cooled, rack-scale foundation for a broad spectrum of the most demanding High-Performance Computing (HPC) workloads, including:

- Manufacturing: Computational Fluid Dynamics (CFD) for aerodynamic optimization; crash simulations; and Multiphysics modeling for materials refinement and production process optimization
- Scientific Research & Complex Modeling: Climate and weather modeling, computational physics, material and life sciences.
- Government & Defense: Mission-Critical Performance for National Security, Nuclear Science, military simulations, etc.
- Financial Services: High-frequency trading, risk analysis, and market simulation.
- Oil & Gas: Seismic data processing and reservoir simulation.

Joint Solution Overview

Supermicro FlexTwin

The FlexTwin is a masterclass in density-optimized design architecture and houses 4 powerful, hot-swappable dual-processor compute nodes within a compact 2U chassis. The FlexTwin architecture directly addresses four of the most critical challenges in today's HPC environments:

1. High Performance Computing

   Nodes can be configured with top-tier processors, such as AMD EPYC™ 9005 Series or Intel® Xeon® 6900 Series processors with P-cores. These processors are paired with DDR5 6400 MT/s memory, or, in the case of the Intel platform, which supports MRDIMM DDR5 8800MT/s for higher bandwidth memory.

2. Cost Optimized for HPC

   High performance always comes at a cost, but the solution ensures you get maximum value for your investment. Our modular design, combined with an energy-efficient liquid cooling system, delivers superior performance per dollar per watt for your HPC cluster. This approach significantly reduces both your Initial Capital Expenditure (CAPEX) and long-term Operational Expenditure (OPEX).

3. Energy Efficiency

The FlexTwin platform offers customers flexibility with two distinct liquid cooling solutions, allowing you to choose your ideal thermal efficiency:

- CPU-Only Cold Plate: This baseline solution provides focused cooling, capturing 61% of the total heat generated by the system.
- Advanced Cold Plate (Comprehensive): For maximum thermal density and efficiency, this solution expands the coverage to include additional high-TDP components, achieving a heat capture rate of up to 95%.
- Performance density and improved serviceability

With four nodes in a 2U form factor, FlexTwin enables massive core density, supporting up to 36,864 cores in a single 48U rack. All nodes, storage, and I/O are front-accessible and hot-swappable. This design simplifies installation and maintenance, eliminating the need to work in the hot aisle and disturbing complex rear-of-rack cabling.

Figure 2 - FlexTwin Server with Dual CPUs

Figure 3 - Supermicro FlexTwin Chassis with 4 Servers

Cornelis CN5000 Omni-Path 400G

The Cornelis CN5000 is an end-to-end HPC networking solution engineered for performance, efficiency, and scalability. At the heart of the platform is the CN5000 400Gbps SuperNIC, a purpose-built adapter that unlocks low-latency, congestion-free communication for GPU- and CPU-intensive workloads. Designed for high-throughput environments, the CN5000 SuperNIC is

now available in both air-cooled and liquid-cooled variants—offering unmatched thermal efficiency for next-generation HPC clusters with FlexTwin. The liquid-cooled Cornelis SuperNIC delivers superior performance per watt while enabling denser server configurations, ideal for high-density deployments where air cooling falls short. This fully integrated solution—from host to fabric to switch—ensures predictable performance at scale, making the CN5000 the intelligent interconnect for accelerating next-gen HPC workloads.

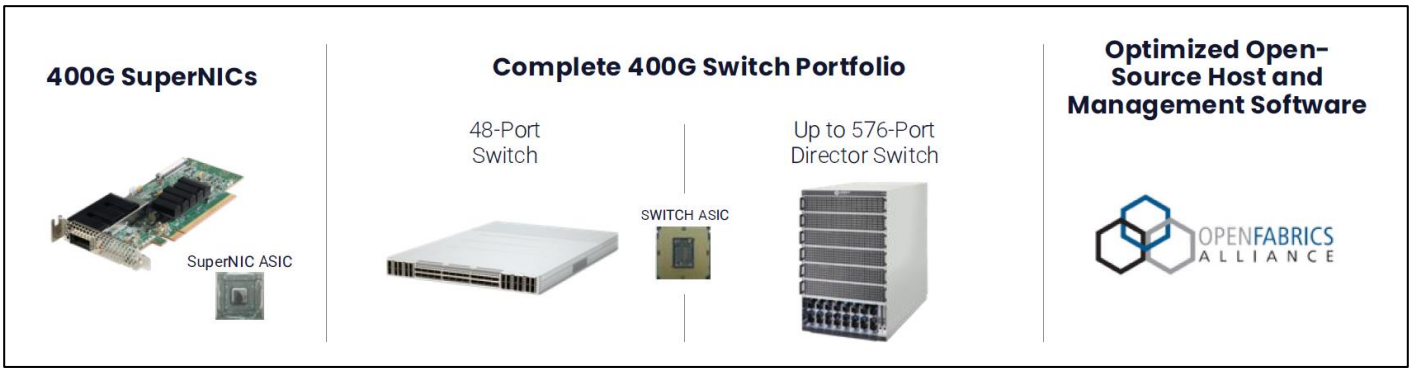

Figure 4 - Cornelis CN5000 Solution Portfolio

Key architectural pillars of Cornelis CN5000 Omni-Path that enable the acceleration of HPC workloads on FlexTwin are:

- Credit-Based Flow Control: Guarantees delivery and eliminates packet loss due to buffer overflow. This is essential for multi-iteration solvers where retransmissions break performance.
- Dynamic Lane Scaling & Link-Level Replay: If a lane fails mid-simulation, CN5000 downgrades gracefully without dropping the connection. At the same time, it locally detects bit errors via hardware replay — avoiding disruptive retransmissions that could invalidate solver iterations.
- Fine-Grained Adaptive Routing: Uses real-time telemetry to route around congestion, not through it. By dynamically selecting cleaner paths, it minimizes long-tail latency and accelerates operations that dominate CFD and structural solvers.
- Easy Integration: Drop-in compatible with Open MPI and standard applications — no vendor lock-in, no retraining required.

Customer Benefits

To validate the performance, cost efficiency, and energy efficiency of the FlexTwin + CN5000 platform, Supermicro and Cornelis jointly executed a comprehensive benchmark campaign using leading HPC solvers and representative workloads. Each test compared an 8-node cluster of FlexTwin systems connected with CN5000 networking against a competitive platform using an alternative 400 Gbps fabric. Tests were conducted under identical cluster configurations and thermal envelopes.

1. Application Performance and Performance per Dollar

Across manufacturing, physics, material life sciences, and climate workloads, the FlexTwin + CN5000 platform consistently delivers up to 1.5x higher performance in 8-node configurations compared to competitive systems. These gains are not isolated

peaks—they reflect a system that sustains throughput across all primary HPC domains, driven by FlexTwin system innovation and CN5000's ultra-low-latency fabric architecture.

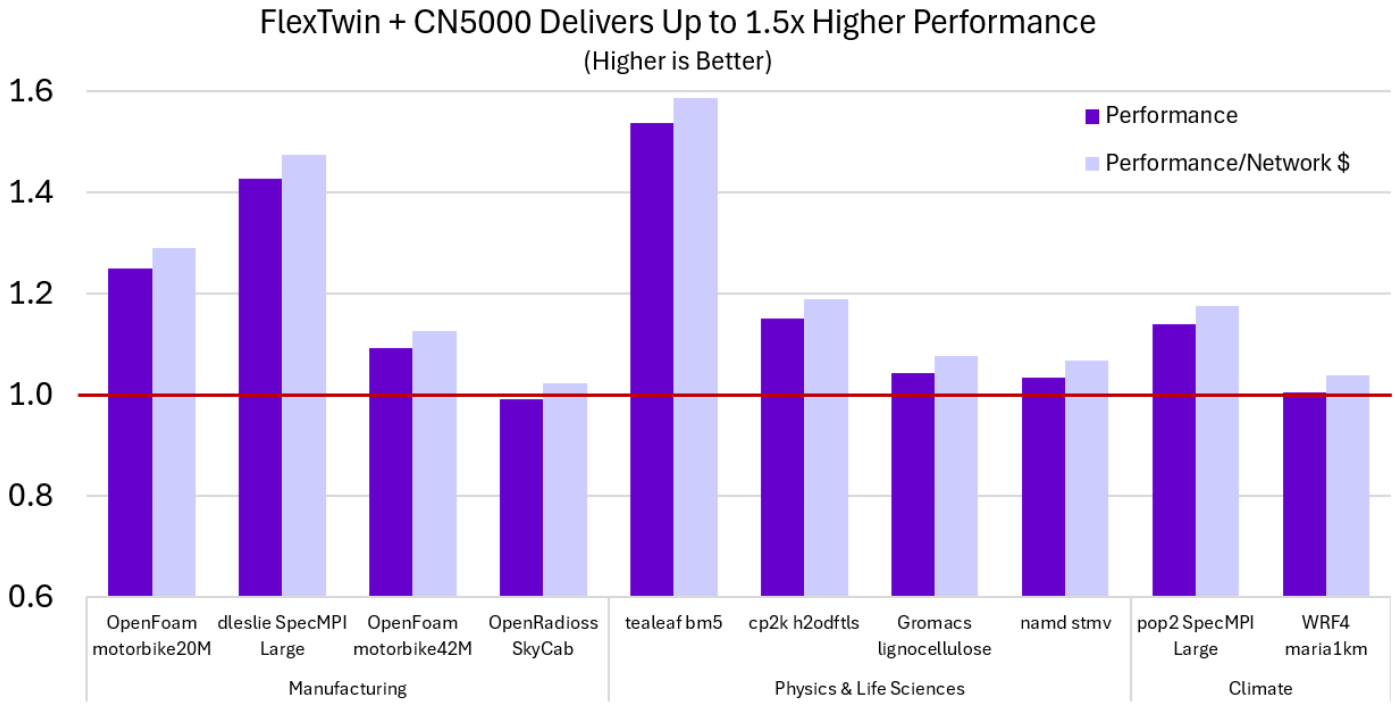

Figure 5 – Performance & Performance Per Dollar with Cornelis CN5000 vs Competitive Solution (1.0)

Beyond raw speed, the economics of performance matter. Using market-driven component cost data (switch, adapters, and cabling), the FlexTwin and CN5000 solution delivers an average 1.29x advantage in performance per network dollar, proving that high performance and cost efficiency can coexist. Customers deploying dozens or hundreds of nodes will see this advantage compound into lower total cost of ownership and reduced cost per simulation or design cycle.

2. Application Performance per Watt

When normalized for total fabric power draw, the FlexTwin and CN5000 solution delivers major energy efficiency gains relative to the competitive platform.

- 1.23x higher average performance per network power with CN5000's air-cooled fabric
- Up to 2.3x higher performance per network power with CN5000's liquid-cooled fabric.

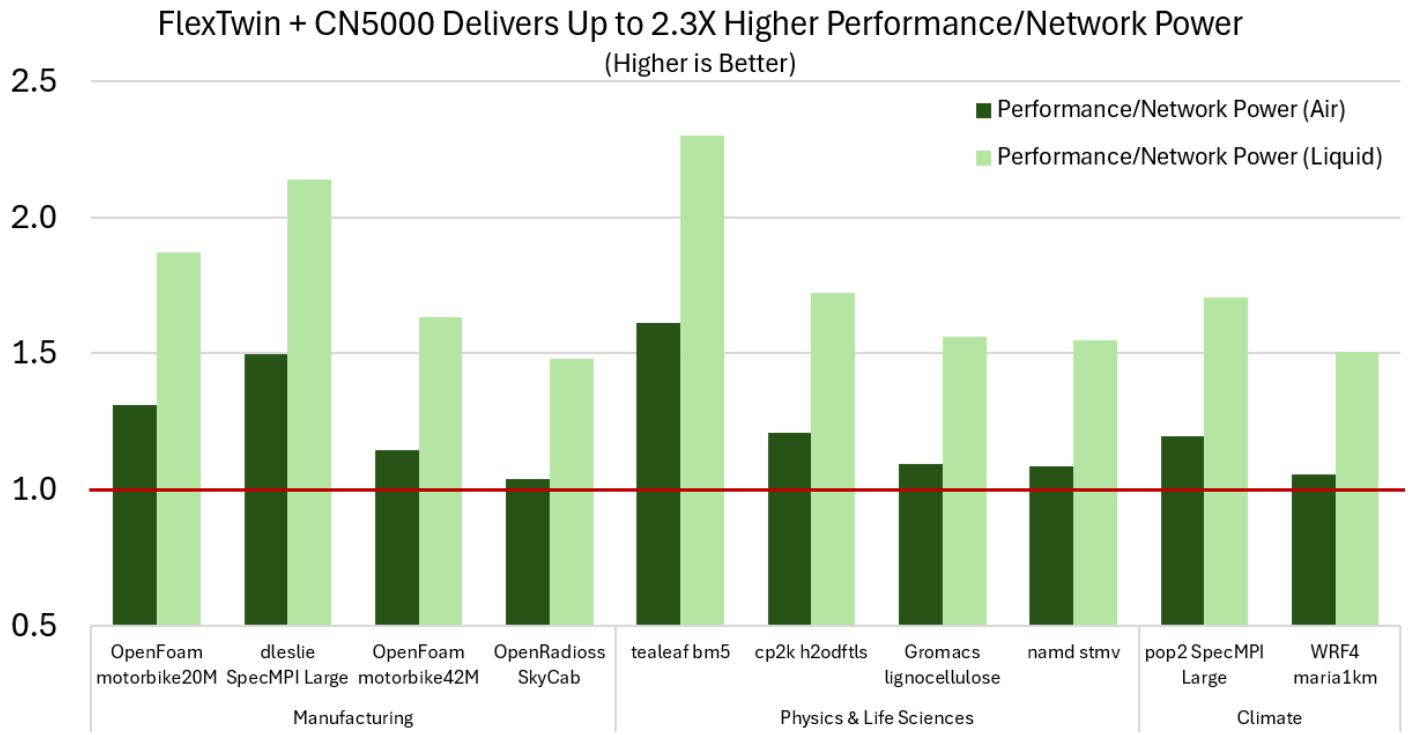

Figure 6 - Performance per Network Power with Cornelis CN5000 Air & Liquid Cooling vs Competition (1.0)

This means customers can double application throughput at the same power level or maintain identical performance while cutting network energy costs nearly in half. With the growing prevalence of power-capped data centers, that flexibility translates directly into more productive racks and a smaller carbon footprint.

3. Scaling Efficiency

Scaling efficiency measures how effectively a workload utilizes additional compute nodes as a cluster grows. In an ideal world, doubling the number of nodes would halve the solution time — a 100 % scaling efficiency. In practice, most HPC applications experience diminishing returns beyond a few nodes due to communication overhead and synchronization delays. Maintaining high scaling efficiency is therefore a key indicator of how well a system's network, compute architecture, and software stack work together under real-world load.

For HPC operators, this metric translates directly into productivity and economics: the closer efficiency stays to 100%, the more value each added node delivers. Clusters with high scaling efficiency waste less compute time in communication, solve problems faster, and require fewer total nodes — lowering both capital and operating costs.

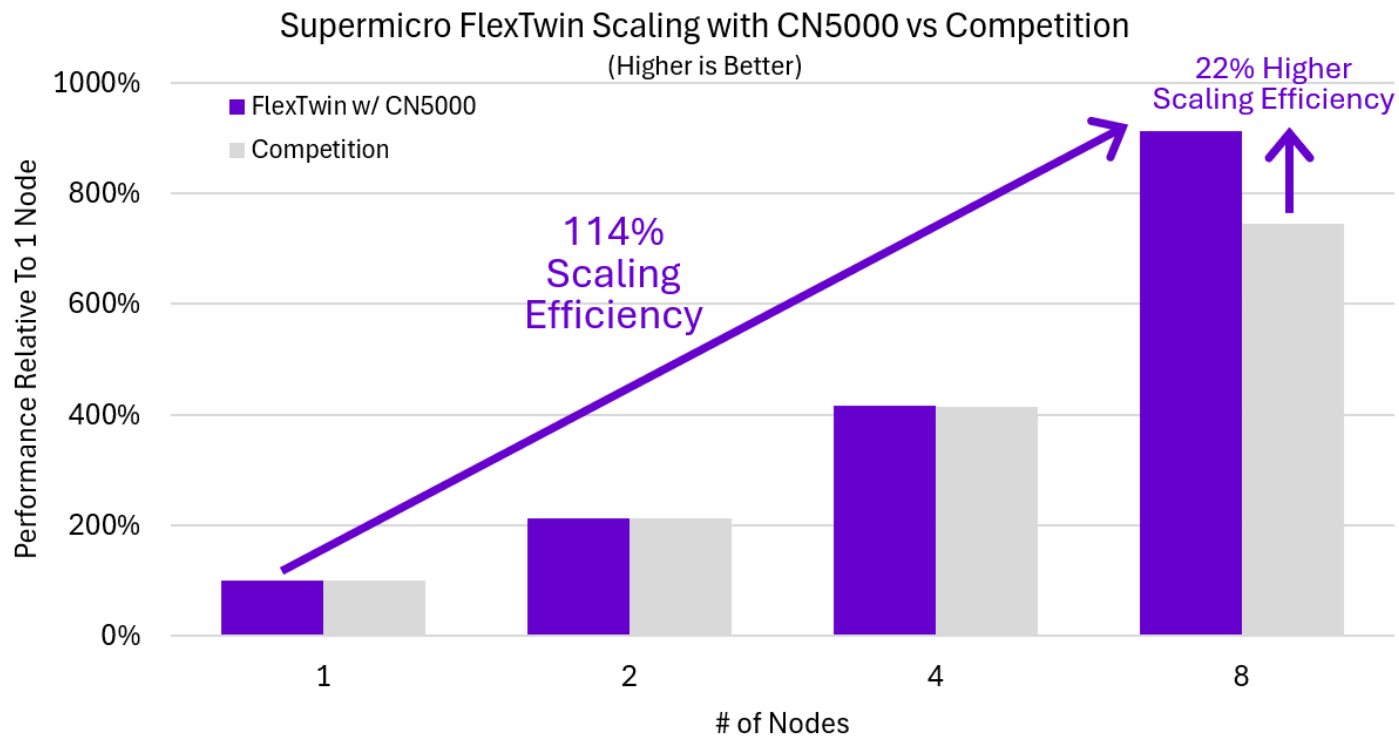

Figure 7 - Scaling performance with Cornelis CN5000

In testing with OpenFOAM motorbike20M, the Supermicro FlexTwin system with Cornelis CN5000 fabric achieved 114% scaling efficiency, representing a 1.22x improvement over a competitive interconnect. The superlinear efficiency observed at eight nodes shows that the FlexTwin and CN5000 platform not only avoided the usual communication overhead but actually accelerated under scale, as memory locality, communication overlap, and thermally stable CPU boost combined to deliver greater throughput than linear scaling alone would predict.

The combined liquid-cooled FlexTwin and CN5000 Omni-Path solution sets a new benchmark for performance per watt and per dollar. Customers can deploy denser racks, achieve faster time-to-solution, and reduce facility energy footprint—all within the same or lower power envelope.

The Power of Liquid Cooling: Efficiency at Scale

Supermicro integrated liquid-cooling technology from the ground up to meet these thermal demands, so customers can run computing-intensive workloads without compromising performance or energy efficiency.

This system provides a rack-level liquid-cooling solution by integrating both L10 and L11 components. This advanced cooling infrastructure ensures maximum performance and adaptability in any High-Performance Computing (HPC) environment, delivering an optimal performance-per-dollar-per-watt ratio.

To quantify this value, we partnered with Cornelis to demonstrate how the superior efficiency of our FlexTwin architecture—combined with the precise networking choice for each specific HPC workload—translates into a compelling performance-per-dollar advantage.

Figure 8 - Supermicro FlexTwin Rack with Liquid Cooling

Summary

The combination of Supermicro's advanced and highly optimized FlexTwin server, combined with the networking performance of the Cornelis CN5000 Omni-Path 400G Solution, will accelerate HPC workloads. To conquer the thermal and density challenges of next-generation HPC, organizations need a fully integrated, liquid-cooled foundation: the Supermicro FlexTwin with Cornelis CN5000 Omni-Path. This unique architecture is specifically engineered to manage the heat of 400G networking and top-tier processors, capturing up to 95% of system heat. The result is a demonstrable, tested advantage: our joint solution delivers up to 1.5x faster application performance, ensuring that every added node maximizes productivity. By choosing this validated platform, customers gain a powerful, cost-optimized, and energy-efficient solution that delivers a clear competitive edge across demanding scientific, defense, and financial workloads.

Appendix

Test Configuration

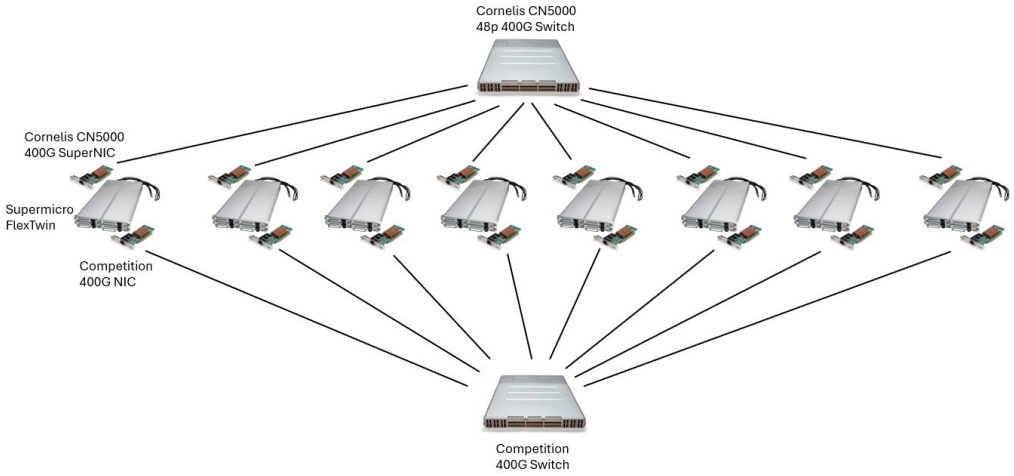

Network component cost data sourced from a list and estimated street pricing.

8-Node Cluster Platform
FlexTwin SYS-222FT-HEA-LCC using 2 socket, 2x Intel Xeon 6972P processors
1x 400G Cornelis CN5000 Omni-Path SuperNIC (at the back)
1x 400G Competition NIC (at the front)
Intel Hyper-Threading Technology enabled

OS
Rocky Linux 9.6 (Blue Onyx). 5.14.0-570.17.1.el9_6.x86_64 kernel

CN5000 Network Parameters and SW
Cornelis CN5000 Omni-Path SuperNIC, Cornelis CN5000 Omni-Path Switch
Cornelis OPX 12.0.2.0.14., 192ppn
Intel MPI 2021.15

For More Information

www.supermicro.com/flextwin

https://www.cornelisnetworks.com/products/cn5000

SUPERMICRO

As a global leader in high performance, high efficiency server technology and innovation, we develop and provide end-to-end green computing solutions to the data center, cloud computing, enterprise IT, big data, HPC, and embedded markets. Our Building Block Solutions® approach allows us to provide a broad range of SKUs, and enables us to build and deliver application-optimized solutions based upon your requirements.

Visit: www.supermicro.com

CORNELIS

Cornelis is a technology leader delivering purpose-built, high-performance fabrics accelerating high-performance computing, high performance data analytics, and artificial intelligence workloads in the cloud and in the datacenter. The company's products enable scientific, academic, governmental, and commercial customers to solve some of the world's toughest challenges by efficiently focusing the computational power of many processing devices at scale on a single problem, simultaneously improving both result accuracy and time-to-solution for their most complex application workloads.

Visit: www.cornelisnetworks.com