



1U 4 GPU Solution Extends Supermicro's Leadership in Supercomputing Servers

GPU / Xeon Phi™ SuperServers Offer Superior Performance, Density and Efficiency

Supermicro Offers the Best Supercomputing Servers

Supermicro continues to enable a highly scalable, energy efficient future for parallel computing with our latest Green Supercomputing solutions. Highlighting our expertise in maximizing compute density, performance and power efficiency, our latest 1U SuperServer® supports up to four NVIDIA GPU or Intel® Xeon Phi™ accelerator cards and dual CPUs and is powered by our highest efficiency redundant Titanium Level digital power supplies.

The new Supermicro SYS-1028GQ-TRT supports up to 4 NVIDIA® Tesla® K80 dual-GPU accelerators (up to 300W) or Intel Xeon Phi™ with a streamlined layout architecture that enables PCI-E direct connect for best signal integrity as well as elimination of complex cabling, repeaters, and GPU pre-heat for maximum airflow and cooling. The system also supports dual Intel® Xeon® E5-2600 v3 processors (up to 145W), up to 1TB ECC DDR4-2133MHz in 16 DIMM slots, front 2x 2.5" hot-swap SATA drive bays plus 2x 2.5" internal drives, dual 10GBase-T ports and intelligent, redundant 2000W (1+1) Titanium Level high-efficiency power supplies. This architecture allows the SYS-1028GQ-TRT to operate at high ambient temperatures, allowing customers to save significant TCO (Total Cost of Ownership) and establishing a best-in-class solution for the widest range of customers' supercomputing challenges.



Figure 1: SYS-1028GQ-TR(T) with 4 NVIDIA GPU cards

Supermicro Is the #1 Customer Choice for Supercomputing Servers

The Supermicro SYS-1028GQ is the clear winner when compared with the Dell PowerEdge C4130, its primary rival in the 1U 4 GPU supercomputing server market space. Supermicro's superior CPU direct connect design optimizes signal integrity and airflow to eliminate the GPU pre-heating, PCI-E extension cables, and re-drivers that add to the Dell system's complexity, cost, extra power consumption, and latency. Power efficiency is further enhanced by fully redundant Titanium Level (96%+) power supplies, while the Dell system employs lower wattage power supplies with the potential for under-powering a fully loaded system and degrading performance when one power supply is unavailable. The Supermicro SYS-1028GQ also supports up to 4x 2.5" drive bays with fully redundant power supplies, while the Dell system must eliminate one power supply in order to add 2 additional 2.5" drive bays. The Supermicro system is unique in supporting NVIDIA GeForce GPU cards which may be required for certain customer applications.

Feature	Supermicro SYS-1028GQ-TR(T)	Dell PowerEdge C4130
Direct CPU connection for lowest latency and power consumption	Yes	Requires extra re-driver chips and cables
Front 2x 2.5" Hot-swap and internal 2x 2.5" drive bays	Yes	No (Sacrifices 1 PWS for 2x 2.5" bays)
Active GeForce GPU support	Yes	No
High Energy Efficiency Power Supplies	2000W fully redundant	1600W redundant when no hot-swap HDDs installed
35°C Ambient Temperature Operation	Yes	No (Conditional 25°C ambient)

Table 1: Competitive Comparison

Supermicro Supercomputing Servers are Optimized for Parallel Applications

Supermicro SYS-1028GQ supercomputing servers are optimized for many processing-intensive, low latency, and TCO sensitive environments.

The ability to scale multiple systems into large high-density deployments while optimizing TCO is a big advantage in the competitive and rapidly evolving *Oil and Gas* industry.

Engineering and Scientific Research fields can dramatically accelerate application performance with minimal investment in development by exploiting the flexibility of Supermicro's supercomputing servers that support up to 4x 2.5" drive bays to store large user data sets.

The SYS-1028GQ high density compute platform is also well-suited for *High-Performance Computing Clusters* with its onboard 10GbE LAN support for high speed networking and energy efficiency to support large, highly scalable deployments.

Many other additional applications can benefit from the scalability, efficiency, and performance of the

Supermicro SYS-1028GQ including Cloud and Visualization, 3D CAD/CAM/CAE, Medical Image Processing, Computational Finance, and Deep Learning.



Figure 2: SYS-1028GQ-TR(T) with 4 Intel Xeon Phi™ cards

Supermicro Has the Widest and Deepest Supercomputing Server Product Line

With an extensive range of SuperServer platforms in 2U 6x, 4U 8x, 4U 4-node FatTwin™ 12x, and 7U SuperBlade® supporting 30x GPU cards, Supermicro offers an unrivaled range of flexible configurations to meet any scale of supercomputing challenge.



Chart 1: Supermicro Supercomputing Server Product Line