

ITPOD servers

A New Era for Complex AI Workloads



ITPOD Servers with NVIDIA GPUs

ITPOD delivers next-generation computing power by combining traditional server components—processors, memory, and storage—with cutting-edge NVIDIA GPU accelerators. Engineered for large-scale generative AI, predictive analytics, and intelligent classification systems, these servers support up to 8 GPU accelerators, including the latest NVIDIA H200 with NVLink bridge . Seamlessly scale performance by linking two or four GPUs into a unified computing environment—ideal for mission-critical AI/ML systems that drive business transformation.

Already trusted within the ITG Corporation, ITPOD servers power cloud platforms, data centers, telecom networks, and enterprise and government IT infrastructures.

Operating principles

Unlike CPUs, optimized for sequential processing across a few powerful cores, GPUs feature thousands of efficient cores designed to handle massive parallel workloads. This architectural advantage makes GPUs the engine of choice for AI training, deep learning, and high-performance computing.



ITPOD-SY4108G-D12R-G4



ITPOD-SYR4108G-D12R-G5

Perfect For

Generative Artificial Intelligence (Generative AI)

High-Performance Computing (HPC) & Scientific Research

Big Data Analytics

Advanced Performance Features

Direct GPU-to-CPU connectivity for maximum throughput

Support up to 8 full-height, double-width GPUs (H200, H100, L40S, A100 tested)

NVLink interconnect for ultra-fast GPU-to-GPU data sharing

Standard Support

3 years

Extended Service Package

Up to 5 years in 24x7 mode

Replacement of faulty components

Next-business-day

Remote technical support

24/7 access to self-service web portal

All supplied equipment is covered by warranty and supported with spare parts from local warehouses in UAE and all our worldwide locations

ITPOD-SY4108G-D12R-G4



Form factor:
4U – Optimized for dense, high-capacity deployments in data centers and AI clusters

Processor:
Dual 5th Gen Intel® Xeon® Scalable processors (up to 350W TDP)

Memory:
32 DDR5 RDIMM slots running at 5600 MT/s

Storage subsystem:
12 front-access bays for 2.5" and 3.5" SATA/SAS/ NVMe drives with hot-swap support
2 M.2 NVMe slots (2280 / 22110)

Networking:
2 × 10GbE RJ45 ports (Intel® X710, optional)

Expansion & GPU Support:
Up to **8 full-height, double-width, full-length GPUs** (PCIe 5.0 x16)
- Direct GPU-to-CPU architecture minimizes latency and maximizes throughput
- Fully tested with NVIDIA H200, H100, L40S and A100 – perfect for generative AI, LLM training, and HPC
- Additional expansion: 2 × PCIe 5.0 x16 and 1 × PCIe 4.0 x8 full-height slots

Power Supply:
4 × 2200W / 3200W hot-swap PSUs
Supports 3+1 or 2+2 redundancy configurations

ITPOD-SYR4108G-D12R-G5



Form factor:
4U – Designed for maximum scalability and efficient integration into high-density data centers and AI clusters

Processor:
Dual AMD EPYC™ 9005 Series ("Turin") processors (up to 350W TDP)

Memory:
24 DDR5 RDIMM slots supporting speeds up to 4800/6400 MT/s

Storage subsystem:
12 front-access universal bays for 2.5" and 3.5" SATA/SAS/NVMe drives with hot-swap support
2 M.2 NVMe slots (2280 form factor)

Networking:
2 × 1GbE RJ45 ports (Intel® i350, integrated)

Expansion & GPU Support:
Supports up to **8 full-height, double-width, full-length GPUs** (PCIe 5.0 x16)
- Direct GPU-to-CPU architecture ensures maximum bandwidth and minimal latency
- Proven compatibility with NVIDIA H200, H100, L40S, and A100 – ideal for generative AI, large language models, and scientific computing
- Additional expansion: 4 × full-height PCIe 5.0 x16 slots for high-performance add-ons

Power Supply:
4 × 2200W / 3200W hot-swap power supplies
Configurable in 3+1 or 2+2 redundancy modes

ITPOD specialists will help design a custom architecture for your needs, select optimal GPU server configurations, and calculate requirements for energy efficiency and fault tolerance.

info@itpod.com +971 43 583 090



itpod.com



LinkedIn