

A photograph of a server rack filled with multiple units, each with numerous fans. The image is overlaid with a blue digital aesthetic, including glowing lines, dots, and a grid pattern. The Gigabyte logo is in the top left. The bottom of the image features a large blue diagonal shape and a grey geometric pattern on the left.

GIGABYTE™

GIGAPOD

Advanced Rack-Scale Solutions

GIGABYTE GIGAPOD delivers scalable AI and HPC infrastructure, providing turnkey rack-level clusters with advanced cooling for extreme computing performance.

GIGAPOD AI


Unleash a Turnkey AI Data Center with High Performance GPU Compute

GIGABYTE has been pivotal in providing technology leaders with supercomputing infrastructure built on high-performance GPU servers. Our solutions are accelerated by the latest NVIDIA, AMD Instinct™, and Intel® Gaudi® 3 accelerator platforms. GIGAPOD is a service that provides professional support to create a cluster of racks all interconnected as a cohesive unit. An AI ecosystem platform thrives on a high degree of parallel processing, with each GPU server optimized for high-speed internal communication using technologies such as NVIDIA NVLink™ or AMD Infinity Fabric™, while scalable inter-rack connectivity enables coordinated performance across the entire cluster.

Complete GPU Ecosystem that Offers the Foundation for AI Infrastructure

GIGABYTE provides a comprehensive, rack-scale ecosystem that moves beyond individual servers to offer fully integrated solutions. Whether for massive AI factories or the burgeoning age of AI reasoning, GIGAPOD ensures that your choice of technology is backed by extreme performance and optimized thermal efficiency.


The GIGAPOD total solution is rounded out with specialized management servers for infrastructure oversight and control, as well as the proprietary GIGABYTE POD Manager (GPM) software suite for DCIM, workload orchestration, MLOps, and more.




AMD Instinct™
MI350/MI300 Series



Intel® Gaudi® 3



NVIDIA HGX™
B300/B200/H200



NVIDIA RTX PRO™ 6000
Blackwell Server Edition



NVIDIA GB300 NVL72

POD for AI Training & AI Factories



GIGAPOD AI with Liquid-Cooled Configurations

GPU Platform	GPU Server	GPU Servers per Rack	Power per Rack	Sockets per Rack	Rack Height	Racks per SU*	CDU
NVIDIA GB300 NVL72	XN15-CB0	18	140 kW	6 x 63A	48U	16+10	In-rack
							In-row
NVIDIA GB200 NVL72	XN14-CB0	18	130 kW	6 x 63A	48U	16+8	In-rack
							In-row
NVIDIA HGX™ Rubin NVL8	G2L4-SD4	8	215 kW	8 x 100A	48U	9+3	In-rack
		8	210 kW	8 x 100A	48U	9+3	In-row
		9	235 kW	8 x 100A	48U	8+3	In-row
		12	305 kW	8 x 100A	48U	6+4	In-row
NVIDIA HGX™ B300	G4L4-SD3 G4L4-ZD3	8	109 kW	6 x 63A	42U/48U	9+4/9+3	In-rack
				6 x 63A	42U/48U	9+3	In-row
	TO46-SD3	8	105 kW	6 x 63A	44OU	9+4	In-rack
				8 x 63A	44OU	9+4	In-row
NVIDIA HGX™ B200 NVIDIA HGX™ H200	G4L4/G4L3 Series	8	103 kW 80 kW	4 x 100A 4 x 63A	42U/48U	4+1	In-rack
							In-row
AMD Instinct™ MI355X AMD Instinct™ MI325X	G4L3-ZX1	8	120 kW 110 kW	4 x 100A	42U/48U	4+1	In-rack
							In-row

GIGAPOD AI with Air-Cooled Configurations

GPU Platform	GPU Server	GPU Servers per Rack	Power per Rack	Sockets per Rack	Rack Height	Racks per SU*
NVIDIA HGX™ H200 / B200 AMD Instinct™ MI300 / MI350	G893/G894 Series	4	58 - 70 kW	4 x 63A	42U/48U	8+1
NVIDIA HGX™ B300	G894 Series	4	70k W	4 x 63A	42U/48U	18+2
NVIDIA HGX™ B200	TO86-SD1	4	55 kW	4 or 8 x 63A	44OU	8+1
Intel® Gaudi® 3	G893-SG1 G894-SG1	4	62 kW	4 x 63A	42U/48U	8+1
NVIDIA RTX PRO™ 6000 Blackwell Server Edition	XL44/XL43 Series	8	80 kW	4 x 63A	42U/48U	4+1

*Compute Rack + Management Rack

GIGAPOD HPC


Scalable High-Density Supercomputing with Intelligent Scale-Out Management

GIGABYTE is at the forefront of enabling technology leaders with high-performance supercomputing infrastructure built on world-class CPU platforms from AMD and Intel. The GIGAPOD is a professional service designed to create a cluster of racks, all interconnected as a cohesive unit and optimized for the most demanding scientific and cloud computing environments.

To ensure sustained performance in high-density, multi-node setups, GIGAPOD HPC utilizes a full-system Direct Liquid Cooling (DLC) solution with integrated piping that achieves up to 91% system heat coverage. Validated for up to 105kW per rack, this thermal architecture removes heat directly at the source, eliminating performance bottlenecks while significantly improving energy efficiency in compute-intensive environments.

Leading CPU Platforms & Technical Foundation

Optimized for the most demanding scientific simulations and engineering analyses, this scalable architecture provides a stable foundation for national research centers and cloud infrastructures. Through multi-node configurations, GIGAPOD HPC leverages modular deployment to ensure rapid time-to-solution while maintaining peak operational efficiency.



The GIGAPOD ecosystem is complemented by specialized management servers and the proprietary GPM software suite, empowering IT personnel with centralized control over DCIM and workload orchestration for maximum operational efficiency.

POD for Engineering & Research Clusters

GIGAPOD HPC with Liquid-Cooled Configurations							
CPU Platform	Muti-node Server	Servers per Rack	Power per Rack	Sockets per Rack	Racks per SU*	Rack Height	CDU
Intel® Xeon® 6900	H174-A80-LAS1	30 (60-node)	105 kW	4 x 100A	4+1	42U	In-rack
							In-row
Intel® Xeon® 6900 Intel® Xeon® 6700/6500	H274-A81-LAZ1 H274-S61-LAW1	15 (60-node)	105 kW	4 x 100A	4+1	42U	In-rack
							In-row
Intel® Xeon® 6900 AMD EPYC™ 9005/9004	B684-A80-LAS1 B683-Z80-LAS1	5 (50-node)	90 kW	4 x 100A	4+1	42U	In-rack
							In-row


*Compute Rack + Management Rack

Why is GIGAPOD the rack scale service to deploy?



Industry Connections

Strategic AMD, Intel, and NVIDIA partnerships ensure rapid response to customer requirements.



Depth in Portfolio

A vast range of specialized server SKUs tailored for any enterprise application.



Scale Up or Out

Turnkey solutions designed for easy integration of new nodes and processors.



High Performance

Optimized server and rack designs ensuring peak performance with optional liquid cooling.



Proven Experience

Proven expertise in deploying large GPU clusters on accelerated delivery timelines.

GIGABYTE POD Manager (GPM)

Streamlined POD Monitoring & Automation

GIGABYTE POD Manager (GPM) is a comprehensive software suite designed to streamline the control, monitoring, and resource optimization of High-Performance Datacenters (PODs). Serving as the management core of the GIGAPOD solution, it provides an intuitive platform to simplify the complexities of large-scale AI and HPC infrastructure.




Key Capabilities:


- Datacenter Infrastructure Management (DCIM):** Provides centralized inventory tracking and real-time health/utilization visualization for servers, switches, and storage devices via GIGABYTE Server Management (GSM).
- Operating System Provisioning:** Automates network device discovery and enables batch OS deployment across multiple devices using customizable installation templates.
- Orchestration and Workload Deployment:** Manages clustered applications like Kubernetes and Hadoop with customizable resource allocation to scale for AI and HPC demands.
- Real-Time Monitoring and Alerting:** Features customizable dashboards with configurable alert thresholds, event logging, and proactive issue resolution to maintain high availability.
- Customizable Ecosystem:** Supports flexible cluster management by integrating with GIGABYTE's proprietary platform, NVIDIA Base Command, NVIDIA AI Enterprise, and other MLOps environments.

Applications & Cloud Services


Software Management - GIGABYTE POD Manager (GPM)



Infrastructure Management



Architecting Service



Compatible Server Platforms for GIGABYTE GIGAPOD

Liquid-Cooled Servers



XN16-CB0

1U Compute Tray for
NVIDIA Vera Rubin NVL72



XN15-CB0

1U Compute Tray for
NVIDIA GB300 NVL72




XN14-CB0

1U Compute Tray for
NVIDIA GB200 NVL72



G2L4-SD4

2U DP HPC/AI Server
NVIDIA HGX™ Rubin NVL8



TO46-SD3

40U DP HPC/AI Server
NVIDIA HGX™ B300



G4L4/G4L3 Series

4U DP HPC/AI Server
NVIDIA HGX™ / AMD Instinct™



H174-A80

1U 2-Node DP High-density Server
Intel® Xeon® 6



H274-A81/S61

2U 4-Node DP High-density Server
Intel® Xeon® 6



B684/B683 Series


6U 10-Node DP Blade Server
Intel® Xeon® 6 / AMD EPYC™ 9005

Air-Cooled Servers



G894/G893 Series

8U DP HPC/AI Server
NVIDIA HGX™ / AMD Instinct™ / Intel® Gaudi® 3



TO86-SD1

80U DP HPC/AI Server
NVIDIA HGX™ B200



XL44/XL43 Series

4U DP NVIDIA MGX Server
NVIDIA RTX PRO™ 6000 BSE



Learn more at <https://www.GigaComputing.com/en>

- * All specifications are subject to change without notice. Please visit our website for the latest information.
- * The entire materials provided herein are for reference only. Giga Computing reserves the right to modify or revise the content at anytime without prior notice.
- * All trademarks and logos are the properties of their respective holders.

Designed by

