

The background of the slide features a close-up, slightly angled view of a server rack filled with multiple server units. The units are densely packed, showing various components like fans, ports, and circuit boards. The lighting is somewhat dim, emphasizing the industrial and technological nature of the scene. In the bottom left corner, there is a decorative geometric pattern of overlapping triangles in shades of gray and white. A large, solid orange triangle points upwards from the bottom right corner, partially overlapping the server rack image and the main title area.

GIGABYTE™

GIGABYTE OCP ORV3 Compliant Solutions


GIGABYTE is an active member of the Open Compute Project (OCP), regularly attending the OCP's annual summits and continuously designing and releasing new server nodes and racks based on the OCP's Open Rack Standard specifications as well as providing the best performing add-on cards for your OCP solution.

GIGABYTE One Stop OCP Solutions

Empowering Innovation, Built for Open Compute.

The Open Compute Project (OCP) is a global, collaborative community focused on open-sourcing and redesigning hardware technology for scalable computing, primarily for data centers.

GIGABYTE is an active member of the OCP, regularly attending the OCP's annual summits and continuously designing and releasing new compute, storage and GPU server hardware based on the OCP Open Rack Standard specifications and providing the best performing mezzanine cards for your OCP solution. GIGABYTE's latest OCP server product line is based on OCP Open Rack V3 (ORV3) specification requirements. The products are designed for a OCP rack and feature a separate PSU system, with power supplied to each server node by a bus-bar system running along the rear of the rack.



GIGABYTE OCP Solutions Advantages



Efficient Rack Density

Optimized node configurations balance compute density with power consumption needs.



Thermal Optimization

Best-in-class thermal design, utilizing the cold aisle/hot aisle concept to effectively manage heat.



Greater Power Efficiency

Centralized power shelf reduces operating expenses with a low Power Usage Effectiveness (PUE).



Simplified Maintenance

Front-access, tool-less design for easier serviceability and component replacement.



Higher Reliability (MTBF)

Centralized power and fewer components increase system Mean Time Between Failures (MTBF).

GIGAPOD: Scalable AI Supercomputing Cluster

Unleash a Turnkey AI Data Center with High Throughput and an Incredible Level of Compute

GIGABYTE GIGAPOD is a scaling supercomputing solution designed for groundbreaking AI workloads, from large language model (LLM) training to real-time inference. Made up of 32 topline GIGABYTE GPU Servers incorporating 256 accelerators, GIGAPOD can be deployed across 9 racks if cooled by air, or a highly dense configuration of 5 racks if utilizing direct liquid cooling (DLC).



32 x GPU Nodes

2 x Management Nodes

22 x Storage Nodes

In AI data center planning, the GPU/Accelerator is the foundational hardware choice. GIGABYTE's industry partners (AMD, Intel, NVIDIA) each offer uniquely advanced products. The selection is complex, relying on factors like performance (AI training or inference), cost, availability, and ecosystem. GIGABYTE simplifies this by offering the choices, customization, and expertise needed to create data centers ideal for demanding AI/ML models.

In addition to hardware, GIGAPOD features GIGABYTE POD Manager (GPM) software to provide complete central control and foster an ideal AI environment. Used by hyperscalers the world over, GIGAPOD is the starting point of modern AI triumphs.

Applications & Cloud Services

Software Management - GIGABYTE POD Manager (GPM)

Workload Management

AI

NVIDIA AI Enterprise

MLOps Platform

Bare-metal, HPC, Virtualization, Containers, NVIDIA Omniverse™, etc.

Cluster Management

NVIDIA Base Command™

GIGABYTE Cluster Manager

Infrastructure Management

Compute Nodes

Management Nodes

Storage Systems

Cooling Systems

Network Fabrics

OS Provisioning

Power Resource

Monitoring

Architecting Service

Project Consultation and Evaluation

Site Planning and Floor Plan Design

Hardware and Software Integration

Cabling Design and Installation

Power and Thermal Management

Data Center Deployment and Validation

Open Rack

DOA3-ST0

Open Rack Version	Open Rack Version 3
Dimensions (WxHxD)	21" 44OU Rack 600 x 2286 x 1068 mm *19" EIA Compatible
Available Node Spac	220U + 210U *23rd OU occupied by reinforced bar
Bus Bar	1 x 48V DC bus bar
Power Shelf	Optional 10U 33 kW power shelf *Support N+N redundancy
Option Parts	EIA conversion shelves, L-shaped shelves, Structural add-ons, Cable management...




Front View



Rear View

GPU Node



TO86-SD1-AA05

Dimensions (WxHxD)	80U GPU Server 537 x 380.5 x 853 mm
GPU	Air-cooled NVIDIA HGX™ B200
CPU	Dual Intel® Xeon® 6700/6500 series
Memory	32 x DDR5 DIMMs 8-Channel RDIMM/MRDIMM
Networking	2 x 10GbE LAN 1 x MLAN
Storage	8 x 2.5" Gen5 NVMe drives 1 x M.2 slot (PCIe 5.0 x4)
Expansion Slots	12 x FHHL PCIe 5.0 x16 slots
Power Supply	48V to 54V DC bus bar

TO86-ZX1-AA03

Dimensions (WxHxD)	80U GPU Server 537 x 380.5 x 853 mm
GPU	Air-cooled AMD Instinct™ MI350X
CPU	Dual AMD EPYC™ 9005/9004 series
Memory	24 x DDR5 DIMMs 12-Channel RDIMM
Networking	2 x 10GbE LAN 1 x MLAN
Storage	8 x 2.5" Gen5 NVMe drives 2 x M.2 slots (PCIe 3.0 x4 & x1)
Expansion Slots	12 x FHHL PCIe 5.0 x16 slots
Power Supply	48V to 54V DC bus bar

Storage Node

TO24-JD0

TO24-JD1





TO24-JD2






TO24-JD3

Dimensions (WxHxD)	20U JBOD 537 x 93 x 807.7 mm	20U JBOD 537 x 93 x 807.7 mm	20U JBOD 537 x 93 x 807.7 mm	20U JBOD 537 x 93 x 807.7 mm
Storage	32 x 3.5" SATA/SAS4 *1 Host 32 drives	32 x 3.5" SATA/SAS4 *2 Host 32 drives	32 x 3.5" SATA/SAS3 *1 Host 32 drives	32 x 3.5" SATA/SAS3 *2 Host, each host 16 drives
Expander	3 x SAS4x32 Single-path	6 x SAS4x32 Dual-path, HA support	1 x SAS4x48 Single-path	2 x SAS4x32 Single-path
IO Port	4 x Mini-SAS HD ports 1 x MLAN	8 x Mini-SAS HD ports 1 x MLAN	4 x Mini-SAS HD ports 1 x MLAN	4 x Mini-SAS HD ports 1 x MLAN
BMC	ASPEED AST2520	ASPEED AST2520	ASPEED AST2520	ASPEED AST2520
Power Supply	48V DC bus bar	48V DC bus bar	48V DC bus bar	48V DC bus bar



Compute Node

	TO25-ZU4-AA01	TO25-ZU5-AA01	TO25-Z10-AA01	TO25-Z11-AA01
   				
Dimensions (WxHxD)	20U 2-node 258 x 86 x 771 mm	20U 2-node 258 x 86 x 771 mm	20U 2-node 265.1 x 90.7 x 742.9 mm	20U 2-node 265.1 x 90.7 x 742.9 mm
CPU	Single AMD EPYC™ 9005/9004	Single AMD EPYC™ 9005/9004	Single AMD EPYC™ 9005/9004	Single AMD EPYC™ 9005/9004
Memory	12 x DDR5 RDIMM	12 x DDR5 RDIMM	12 x DDR5 RDIMM	12 x DDR5 RDIMM
Networking	1 x MLAN, 2 x 10GbE LAN	1 x MLAN, 2 x 10GbE LAN	1 x MLAN	1 x MLAN
Storage	2 x 9.5mm E1.S NVMe 4 x 2.5"NVMe/SATA/SAS4	2 x 9.5mm E1.S NVMe 8 x 2.5"NVMe/SATA/SAS4	4 x 9.5mm E1.S NVMe Optional 1 x M.2 NVMe	12 x 9.5mm E1.S NVMe Optional 1 x M.2 NVMe
Expansion Slots	2 x PCIe Gen5 x16 FHFL 2 x OCP 3.0 Gen5 x16	2 x PCIe Gen5 x16 LP	2 x PCIe Gen5 x16 LP 1 x OCP 3.0 Gen5 x16	1 x PCIe Gen5 x16 LP 1 x OCP 3.0 Gen5 x16
Power Supply	48V to 54V DC bus bar	48V to 54V DC bus bar	48V DC bus bar	48V DC bus bar

	TO25-Z12-AA01	TO25-S12-AA01	TO25-S10-AA01	TO25-S11-AA01	TO23-H60
    					
Dimensions (WxHxD)	20U 2-node 262.7 x 90 x 740 mm	20U 2-node 265.1 x 90.7 x 742.9 mm	20U 2-node 262.7 x 90 x 740 mm	20U 2-node 262.7 x 90 x 740 mm	20U 3-node 180 x 90 x 740 mm
CPU	AMD EPYC™ 9005/9004 Single processor	5th/4th Gen Intel® Xeon® Single processor	5th/4th Gen Intel® Xeon® Single processor	5th/4th Gen Intel® Xeon® Single processor	3rd Gen Intel® Xeon® Dual processors
Memory	12 x DDR5 RDIMM	8 x DDR5 RDIMM	8 x DDR5 RDIMM	8 x DDR5 RDIMM	16 x DDR4 RDIMM/LRDIMM
Networking	1 x MLAN	1 x MLAN	1 x MLAN	1 x MLAN	1 x MLAN
Storage	4 x 9.5mm E1.S NVMe Optional 1 x M.2 NVMe	2 x 9.5mm E1.S NVMe 1 x M.2 SATA	2 x 9.5mm E1.S NVMe 1 x M.2 SATA	10 x 9.5mm E1.S NVMe 1 x M.2 SATA	4 x 2.5" NVMe/SATA/SAS
Expansion Slots	2 x PCIe Gen5 x16 FHFL - Supports 1 x Dual slot GPU 1 x LP PCIe Gen5 x16 1 x OCP 3.0 Gen5 x16	2 x PCIe Gen5 x16 FHFL - Supports 1 x Dual slot GPU 2 x PCIe Gen5 x16 FHFL 1 x PCIe Gen5 x8 LP	2 x PCIe Gen5 x16 FHFL 1 x PCIe Gen5 x16 LP 1 x PCIe Gen5 x8 LP	2 x PCIe Gen5 x16 FHFL 1 x PCIe Gen5 x8 LP	2 x PCIe Gen4 x16 LP 2 x OCP 3.0 Gen4 x16
Power Supply	48V DC bus bar	48V DC bus bar	48V DC bus bar	48V DC bus bar	48V DC bus bar

Node Tray

	TO25-BT1-AA01	TO25-BT0-AA01	TO23-BT0
  			
Dimensions (WxHxD)	20U 2-node tray 537 x 93.5 x 883 mm	20U 2-node tray 537 x 93.5 x 918.5 mm	20U 3-node tray 537 x 93.2 x 803 mm
No. of Bus Bars	1 x 48V to 54V bus bar connector	1 x 48V bus bar connector	1 x 48V bus bar connector

OCP Immersion Cooling

Taking another step forward, GIGABYTE has also introduced an immersion cooling total solution. Similar to its air cooling and direct liquid cooling counterparts, the single-phase immersion cooling solution covers everything, from add-on cards to the immersion tank. This next-generation cooling solution is perfect for companies interested in adopting more environmentally friendly and cost-effective alternatives while wanting to avoid the complicated and time-consuming verification process.

A103-CC0 Immersion Tank	
Hardware Capacity	180U + 20U (power shelf) + 1U x 2 (EIA)
Heat Dissipation Capacity	80 kW with 25°C inlet water 60 kW with 35°C inlet water
Dimensions (L x W x H)	Tank: 0.91m x 1.16m x 1.49m CDU: 0.86m x 0.48m x 1.62m
Weight	Tank: 450 kg (w/o coolant); CDU: 325 kg
CDU Power Consumption	0.75 kW
Coolant Volume	750 L
Electrical Connection	PDU: IEC 60309 32A-3P+N+E, 6H, Plug x1 Power Shelf: IEC 60309 63A-3P+N+E, 6H, Plug x1 CDU: IEC 60320 C19, Plug x1
Footprint	2.1m x 0.9m
Cooling Water Inlet	20-35°C
Water Flow Rate	240 LPM
PUE	1.02




	TO15-S40-IA01	TO15-S41-IA01	TO15-Z40-IA01	TO15-Z20-IA01
Dimensions (WxHxD)	10U GPU Server 537 x 45 x 805.7 mm	10U GPU Server 537 x 45 x 805.7 mm	10U GPU Server 537 x 45 x 805.7 mm	10U GPU Server 537 x 45 x 805.7 mm
CPU	5th/4th Gen Intel® Xeon® Dual processors	5th/4th Gen Intel® Xeon® Dual processors	AMD EPYC™ 9005/9004 Single processor	AMD EPYC™ 9005/9004 Single processor
Memory	24 x DDR5 DIMMs 8-Channel RDIMM	16 x DDR5 DIMMs 8-Channel RDIMM	24 x DDR5 DIMMs 12-Channel RDIMM	12 x DDR5 DIMMs 12-Channel RDIMM
Networking	2 x 10GbE LAN 1 x MLAN	2 x 10GbE LAN 1 x MLAN	2 x 10GbE LAN 1 x MLAN	2 x 10GbE LAN 1 x MLAN
Storage	4 x 15mm E1.S NVMe Optional 2 x M.2 slots (SATA)	4 x 15mm E1.S NVMe Optional 2 x M.2 slots (SATA)	4 x 15mm E1.S NVMe Optional 2 x M.2 slots (SATA)	4 x 15mm E1.S NVMe 2 x M.2 slots (PCIe Gen5 x4) Optional 2 x M.2 slots (SATA)
Expansion Slots	4 x Dual-slot Gen5 GPUs 2 x LP PCIe Gen5 x16 slots	8 x Single-slot Gen5 GPUs 2 x LP PCIe Gen5 x16 slots	4 x Dual-slot Gen5 GPUs 2 x LP PCIe Gen5 x16 slots	4 x Dual slot-Gen5 GPUs 2 x LP PCIe Gen5 x16 slots
Power Supply	48V DC bus bar	48V DC bus bar	48V DC bus bar	48V DC bus bar

OCP Direct Liquid Cooling

GIGABYTE Direct Liquid Cooling solution focuses on innovative breakthroughs in AI, HPC and cloud computing, delivering outstanding efficiency in heat dissipation while achieving high system availability and stability. GIGABYTE provides comprehensive DLC solutions, including optimized compute nodes, GPU servers, racks, and a full liquid cooling ecosystem with components like cold plate, coolant distribution unit (CDU), rack manifold, and rear door heat exchanger (RDHx). Tailored solutions are available for POC, small-scale adoption, and full-scale data center deployment.

DCA3-ST0 DLC Rack

Open Rack Version	Open Rack Version 3
Dimensions (W x H x D)	21" 44OU Rack 800 x 2286 x 1200 mm 800 x 2286 x 1450 mm (with RDHx)
Available Node Spac	22OU + 21OU *23rd OU occupied by reinforced bar
Bus Bar	1 x 48V DC bus bar
Power Shelf	Optional 10U 33 kW power shelf Support N+N redundancy
Manifold	Single vertical manifold with UQD06 quick connectors
Option Parts	CDU, RDHx, DLC management server, EIA conversion shelves, L-shaped shelves, Structural add-ons, Cable management...



TO46-SD3-LA07

TO46-ZX1-LA04



Dimensions (WxHxD)	40U GPU Server 537.0 x 191 x 812.2 mm	40U GPU Server 537.0 x 191 x 812.2 mm
GPU	Liquid-cooled NVIDIA HGX™ B300	Liquid-cooled AMD Instinct™ MI355X
CPU	Dual Intel® Xeon® 6700/6500 series	Dual AMD EPYC™ 9005/9004 series
Memory	32 x DDR5 DIMMs 8-Channel RDIMM/MRDIMM	24 x DDR5 DIMMs 12-Channel RDIMM
Networking	8 x 800Gb/S OSFP 2 x 10GbE LAN 1 x MLAN	2 x 10GbE LAN 1 x MLAN
Storage	8 x 2.5" Gen5 NVMe drives 1 x M.2 slot (PCIe 5.0 x4) 1 x M.2 slot (PCIe 5.0 x2)	8 x 2.5" Gen5 NVMe drives 1 x M.2 slot (PCIe 3.0 x4) 1 x M.2 slot (PCIe 3.0 x1)
Expansion Slots	4 x FHHL dual-slot PCIe 5.0 x16 slots	4 x FHHL dual-slot PCIe 5.0 x16 slots 8 x FHHL single-slot PCIe 5.0 x16 slots
Power Supply	48V to 54V DC bus bar	48V to 54V DC bus bar



GigaComputing.com/en



Giga Computing



Giga Computing



Giga Computing



GigaComputing



Designed by

**GIGA
COMPUTING**

* All intellectual property rights, including without limitation to copyright and trademark of this work and its derivative works are the property of, or are licensed to, Giga Computing Technology Co., Ltd. Any unauthorized use is strictly prohibited.

* 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.