

A photograph of several server racks in a data center. The racks are filled with server units, and red and blue liquid cooling tubes are connected to them. The image is partially obscured by a grey geometric pattern on the left and a purple diagonal shape on the right.

GIGABYTE™


GIGABYTE Direct Liquid Cooling Solution

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 Direct Liquid Cooling Solution

Productivity at Scale for the AI Factory Era: Servers × Liquid Cooling × G-REX

The rise of generative AI, large language models (LLMs), and high-performance computing (HPC) is fundamentally transforming data centers into specialized AI processing facilities. As rack power densities increasingly exceed 100 kW, scaling computation introduces severe energy, thermal, and operational challenges. The GIGABYTE Direct Liquid Cooling (DLC) architecture is systematically engineered to mitigate these constraints. By integrating purpose-built DLC servers, the GIGAPOD rack-scale infrastructure, and the G-REX management platform, GIGABYTE delivers a comprehensive solution ensuring high system availability, optimal thermal efficiency, and sustainable scalability.



Industry Challenges



Recent analyses by the U.S. Department of Energy (DOE) and the Lawrence Berkeley National Laboratory (LBNL) highlight escalating data center power consumption, straining existing electrical and thermal infrastructures.

- Limitations of Conventional Air Cooling:** Traditional air-cooling requires massive airflow and highly energy-intensive chillers, rendering it inadequate for dissipating the extreme thermal loads of high-density GPU clusters.
- Challenges of Scalability and Reliability:** Reliably operating hundreds of high-density racks in a live production environment is a major logistical undertaking, far exceeding the complexity of deploying isolated servers.
- IT and OT Integration Complexity:** Liquid cooling introduces operational technology (OT) components like Coolant Distribution Units (CDUs) and Rear-Door Heat Exchangers (RDHx). Without seamless integration, managing these alongside IT infrastructure creates operational friction.

Direct Liquid Cooling: Native Engineering over Retrofitting

Unlike aftermarket retrofits that merely adapt legacy architectures, GIGABYTE DLC servers are natively architected for liquid cooling prior to manufacturing. High-thermal-output components—specifically GPUs, CPUs, and high-bandwidth memory modules—interface directly with precision-engineered liquid conduits via micro-channel cold plates, extracting heat immediately at the source.

- **>92% Heat Capture Ratio (HCR):** The system extracts over 92% of total thermal output directly via liquid mediums, virtually eliminating the need for traditional Computer Room Air Conditioning (CRAC) overprovisioning.
- **Empirical Validation at Extreme Density:** The cold-plate technology is rigorously field-verified, demonstrating continuous stability under massive 140 kW rack-level IT loads.
- **Energy Efficacy:** The architecture eliminates high-RPM internal fans, drastically reduces chiller dependence, improves overall facility PUE, and creates a quieter working environment.



G-REX: Unified Administration for OT and IT

G-REX serves as the proprietary middleware of the GIGABYTE POD Manager (GPM), designed to dismantle silos between IT administration and facility OT by aggregating cooling and power metrics into a single pane of glass.

- **Cluster-Level Visibility and Control:** Delivers real-time monitoring of flow rates, temperatures, and power utilization via an IT-standardized interface, enabling proactive thermal management and predictive maintenance.
- **Open Standards and APIs:** Fully supporting Modbus, Redfish, and REST APIs, G-REX facilitates agnostic integration with third-party CDUs and rear-door heat exchangers, preventing vendor lock-in.

GIGAPOD: Modular and Scalable AI Infrastructure

GIGABYTE GIGAPOD extends DLC technology into a holistically integrated, liquid-cooled compute rack architecture. Subsystems—encompassing compute nodes, networking, power, and cooling—are pre-configured and factory-validated (L11/L12 paradigms), transforming months of onsite plumbing into days of rapid deployment.

- **GIGAPOD AI:** A purpose-built, GPU-centric architecture meticulously designed for large-scale AI training and LLM inference clusters. It provides structural and thermal support for flagship enterprise accelerators, including the NVIDIA HGX, GB200, and GB300 architectures, as well as the AMD Instinct series.
- **GIGAPOD HPC:** A highly scalable, CPU-centric architecture optimized for dense scientific simulations, complex engineering analyses, national laboratory deployments, and generalized cloud workloads.
- **Elevated Temperature Cooling (up to 45°C):** Engineered to operate stably with inlet water temperatures up to 45°C, this system minimizes chiller reliance and extends free-cooling windows, substantially mitigating operational energy expenditures.



48U GIGAPOD AI DLC Rack Configurations



	DLB5-CB0 (NVIDIA Vera Rubin NVL72)
GPU	72 x NVIDIA Rubin GPUs
CPU	36 x NVIDIA Vera CPUs
Compute Tray	18 x 1U GIGABYTE XN16-CB0-L01
NVLink Switch Tray	9 x 1U NVIDIA NVLink™ Switches
Management Switch	2 x NVIDIA Spectrum™ SN2201_M ToR Switches
Power Shelf	4 x 3U 110kW power shelves
Power Consumption	Up to 230 kW
Bus Bar	1 x Liquid cooled 54V DC 5000A+ Bus Bar
Cable Cartridge	4 x NVIDIA NVLink™ cable cartridges
CDU	Compatible with in-row CDU



	DLB2-CB3 (NVIDIA GB300 NVL72)
GPU	72 x NVIDIA Blackwell Ultra GPUs
CPU	36 x NVIDIA Grace™ CPUs
Compute Tray	18 x 1U GIGABYTE XN15-CB0-LA01
NVLink Switch Tray	9 x 1U NVIDIA NVLink™ Switches
Management Switch	2 x NVIDIA Spectrum™ SN2201 ToR Switches
Power Shelf	6 x 1U 33 kW power shelves
Power Consumption	Up to 140 kW
Bus Bar	1 x 54V DC 1400A Bus Bar
Cable Cartridge	4 x NVIDIA NVLink™ cable cartridges
CDU	Compatible with in-row or in-rack CDU



	DLB2-CB1 (NVIDIA GB200 NVL72)
GPU	72 x NVIDIA Blackwell GPUs
CPU	36 x NVIDIA Grace™ CPUs
Compute Tray	18 x 1U GIGABYTE XN14-CB0-LA01
NVLink Switch Tray	9 x 1U NVIDIA NVLink™ Switches
Management Switch	2 x NVIDIA Spectrum™ SN2201 ToR Switches
Power Shelf	6 x 1U 33 kW power shelves
Power Consumption	Up to 130 kW
Bus Bar	1 x 54V DC 1400 A Bus Bar
Cable Cartridge	4 x NVIDIA NVLink™ cable cartridges
CDU	Compatible with in-row or in-rack CDU

42U GIGAPOD AI DLC Rack Scale Solutions



	DL83-GP6	
Server Model	G4L4-SD3-LAX7	G4L4-ZD3-LAX7
Server Configuration	8 x 4U DLC GPU servers	
CPU	Intel® Xeon® 6700/6500-Series (16 x CPUs)	AMD EPYC™ 9005/9004 Series (16 x CPUs)
GPU	NVIDIA HGX™ B300 (64 x GPUs)	
Power Consumption	Up to 109 kW	
PDU	6 x PDUs; Input: 415 V, 100 A, 3-phase	
CDU	In-Rack Liquid-to-Liquid CDU; Cooling Capacity: 200 kW	



	DL83-GP5	
Server Model	G4L3-ZX1-LAT4	
Server Configuration	8 x 4U DLC GPU servers	
CPU	AMD EPYC™ 9005/9004 Series (16 x CPUs)	
GPU	AMD Instinct™ MI355X (64 x GPUs)	
Power Consumption	Up to 120 kW	
PDU	4 x PDUs; Input: 415 V, 125 A, 3-phase	
CDU	In-Rack Liquid-to-Liquid CDU; Cooling Capacity: 100 kW	



	DL83-GP2		
Server Model	G4L3-SD1-LAX5	G4L3-ZD1-LAX5	G4L3-ZX1-LAX2
Server Configuration	8 x 4U DLC GPU servers		
CPU	5th/4th Gen Intel® Xeon®	AMD EPYC™ 9005/9004 Series (16 x CPUs)	
GPU	NVIDIA HGX™ B200 (64 x GPUs)		AMD Instinct™ MI325X
Power Consumption	Up to 103 kW		Up to 110 kW
PDU	4 x PDUs; Input: 415 V, 100 A, 3-phase		
CDU	In-Rack Liquid-to-Liquid CDU; Cooling Capacity: 100 kW		



	DL83-GP1	
Server Model	G4L3-SD1-LAX3	G4L3-ZD1-LAX3
Server Configuration	8 x 4U DLC GPU servers	
CPU	5 th /4 th Gen Intel® Xeon® Scalable (16 x CPUs)	AMD EPYC™ 9005/9004 Series (16 x CPUs)
GPU	NVIDIA HGX™ H200 (64 x GPUs)	
Power Consumption	Up to 80 kW	
PDU	4 x PDUs; Input: 415 V, 100 A, 3-phase	
CDU	In-Rack Liquid-to-Liquid CDU; Cooling Capacity: 100 kW	

42U GIGAPOD AI DLC Rack Scale Solutions

6



	DL83-MG0
Server Model	XN24-VC0-LA61
Server Configuration	12 x 2U Liquid-cooled NVIDIA GB200 NVL4 servers
CPU	24 x NVIDIA Grace™ CPUs
GPU	48 x NVIDIA Blackwell GPUs
Power Consumption	Up to 80 kW
PDU	4 x PDUs; Input: 415 V, 100 A, 3-phase
CDU	In-Rack Liquid-to-Liquid CDU; Cooling Capacity: 100 kW



	DL83-MG1
Server Model	XLS4-SX2-LAS1
Server Configuration	5 x 6U Liquid-cooled NVIDIA RTX PRO™ servers
CPU	Intel® Xeon® 6700/6500-Series (10 x CPUs)
GPU	40 x NVIDIA RTX PRO™ 6000 Blackwell Server Edition GPUs
Power Consumption	Up to 39 kW
PDU	4 x PDUs; Input: 415 V, 100 A, 3-phase
CDU	In-Rack Liquid-to-Liquid CDU; Cooling Capacity: 100 kW




	DL83-GP0	
Server Model	G494-SB0-LAP1	G494-ZB0-LAP1
Server Configuration	4 x 4U DLC PCIe GPU servers	
CPU	Intel® Xeon® 6700/6500-Series (8 x CPUs)	AMD EPYC™ 9005/9004 Series (8 x CPUs)
GPU	32 x AMD Radeon™ AI PRO R9700S PCIe GPUs	
Power Consumption	Up to 30 kW	
PDU	2 x PDUs; Input: 415 V, 100 A, 3-phase	
CDU	In-Rack Liquid-to-Liquid CDU; Cooling Capacity: 100 kW	



	DL83-GP4
Server Model	G383-R80-LAP1
Server Configuration	8 x 3U DLC APU servers
APU	32 x AMD Instinct™ MI300A APUs
Power Consumption	Up to 50 kW
PDU	2 x PDUs; Input: 415 V, 100 A, 3-phase
CDU	In-Rack Liquid-to-Liquid CDU; Cooling Capacity: 100 kW

42U GIGAPOD HPC DLC Rack Scale Solutions



	DL83-BL0
Server Model	B683-Z80-LAS1
Server Configuration	5 x 6U10N DLC high density servers
CPU	AMD EPYC™ 9005/9004 Series Total 100 processors
Power Consumption	Up to 90 kW
PDU	4 x PDUs; Input: 415 V, 100 A, 3-phase
CDU	In-Rack Liquid-to-Liquid CDU; Cooling Capacity: 100 kW



	DL83-HD0		
Server Model	H274-S60-LAW1 H274-S61-LAW1	H273-Z80-LAW1/LAN1 H273-Z81-LAW1/LAN1	H263-S62-LAW1/LAN1 H263-S63-LAW1/LAN1
Server Configuration	16 x 2U4N DLC high density servers		
CPU	Intel® Xeon® 6700/6500 Total 128 processors	AMD EPYC™ 9005/9004 Total 128 processors	5 th /4 th Gen Intel® Xeon® Scalable Total 128 processors
Power Consumption	Up to 83 kW	Up to 90 kW	Up to 82 kW
PDU	2 x PDUs; Input: 415 V, 100 A, 3-phase		
CDU	In-Rack Liquid-to-Liquid CDU; Cooling Capacity: 100 kW		



	DL83-HD1	
Server Model	H273-Z85-LAZ1	H274-A81-LAZ1
Server Configuration	15 x 2U4N DLC high density servers	
CPU	AMD EPYC™ 9005/9004 Series Total 120 processors	Intel® Xeon® 6900-Series Total 120 processors
Power Consumption	Up to 105 kW	Up to 105 kW
PDU	2 x PDUs; Input: 415 V, 125 A, 3-phase	
CDU	In-Rack Liquid-to-Liquid CDU; Cooling Capacity: 100 kW	



	DL83-HD0
Server Model	H374-A80-LAW1 H374-A81-LAW1
Server Configuration	10 x 3U4N DLC high density servers
CPU	Intel® Xeon® 6900-Series Processors Total 80 processors
Power Consumption	Up to 76 kW
PDU	2 x PDUs; Input: 415 V, 100 A, 3-phase
CDU	In-Rack Liquid-to-Liquid CDU; Cooling Capacity: 100 kW

GIGABYTE Direct Liquid Cooling Servers



2U MGX Server
NVIDIA GB200 NVL4



2U HGX Server
NVIDIA HGX™ Rubin NVL8



3U APU Server
AMD Instinct™ MI300A



4U HGX / OAM Server
NVIDIA HGX™ / AMD Instinct™



4U PCIe GPU Server
AMD Radeon™ AI PRO R9700S



6U MGX Server
NVIDIA RTX PRO™ 6000 BSE



6U10N Blade Server
Intel® Xeon® / AMD EPYC™



3U4N High-density Server
Intel® Xeon® 6



2U4N High-density Server
Intel® Xeon® / AMD EPYC™



1U2N High-density Server
Intel® Xeon® 6




1U DP Rack Server
Intel® Xeon® / AMD EPYC™



G-REX
DLC Management Kit

Why GIGABYTE Direct Liquid Cooling Solution



Optimized All-in-One

A complete, high-performance rack solution ready for seamless deployment.



Flexible & Verified

Customizable configurations thoroughly tested with verified partner brands.



Fast Delivery, Cost-Effective

Shorter lead times and competitive pricing without compromising GIGABYTE's quality and service.



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

