



## P9000G6(LC)



### Advanced Liquid-Cooled Platform for AI and HPC Workloads

#### Highlights

- Support dual 4th and 5th Gen Intel® Xeon® Scalable processor
- Support NVIDIA HGX™ B200 system with liquid cooling solution (GPU+CPU)
- Best GPU Communication with NVIDIA NVLink™ Bridge
- Support DDR5 DIMM, 4800/5600 MT/s @ 1DPC, 4400 MT/s @ 2DPC
- Support up to 12 U.2 NVMe SSD
- Inventec-designed switchboard provides maximum bandwidth for NVIDIA GPUDirect® RDMA
- Increase power efficiency and reliability by decoupling 12V and 54V power source
- Modular architecture to provide customer the best performance platform

With AI applications expanding rapidly and HPC workloads becoming increasingly complex, the P9000G6 (LC) is designed to deliver the performance, scalability, and energy efficiency required by today's data centers. Featuring dual Intel® Xeon® Scalable processors and the latest NVIDIA HGX™ B200 platform, it provides a robust foundation for AI training, scientific computing, and other data-intensive tasks. The system uses liquid cooling across both GPU and CPU modules, ensuring consistent thermal performance under high compute loads.

#### An Extremely Flexible AI Server Ready for the Next Wave of AI and HPC


The P9000G6 (LC) features a flexible 8U design that combines dual Intel® Xeon® Scalable processors with the NVIDIA HGX™ B200 8-GPU platform, delivering high compute density for demanding AI workloads. It supports up to 12 U.2 NVMe SSDs, located under the same PCIe switch as the GPUs to enable fast local data access and enhance NVIDIA GPUDirect® Storage (GDS) performance. To maintain performance under sustained loads, the system integrates a liquid cooling solution that covers both the GPU and CPU zones. This ensures stable operation in data-intensive environments such as model training, inference acceleration, and AI simulation. With its balanced architecture and thermal efficiency, the P9000G6 (LC) provides a reliable foundation for next-generation AI computing.

## Engineered for Fluid AI Performance

The P9000G6 (LC) is built to handle the data-heavy nature of modern AI workloads. With DDR5 memory running at speeds up to 5600 MT/s at 1DPC and 4400 MT/s at 2DPC, the system ensures fast and efficient data access. NVIDIA NVLink™ Bridge connects GPUs with high-bandwidth interconnects, enabling smooth coordination across parallel tasks. To further optimize performance, Inventec's custom switchboard architecture supports NVIDIA GPUDirect® RDMA, streamlining data transfer between CPUs and GPUs. This reduces latency and helps maintain consistent throughput under intensive compute conditions.

## Modular Architecture for Simplified Maintenance

In line with Inventec's focus on serviceability, the P9000G6 (LC) features a modular design that enables quick, tool-free access to key components. Fans, storage, NICs, and power modules are easy to maintain and replace. The clear separation between host and GPU modules enhances thermal zoning and improves accessibility, making maintenance and system management more efficient.



### About Inventec Data Center Solutions (Inventec EBG)

Inventec Data Center Solutions (Inventec EBG) was established in 1998 and has been focusing on the design and manufacturing of server systems in Inventec Corporation. Over decades, Inventec EBG has been the key server system supplier of the global branding clients.

Inventec logos are trademarks or registered trademarks of Inventec Corporation. Inventec reserves the right to modify this document, the Specifications and photos from time to time without notifying the Party. The entire materials provided herein are for reference only.

# P9000G6(LC)

Advanced Liquid-Cooled Platform  
for AI and HPC Workloads

Form Factor	8U Rackmount W x H x D: 448x 352x 850 mm (17.63x 13.85x 33.46inch)
Processor	4th and 5th Gen Intel Xeon Scalable Processor up to 385W
GPU Module	NVIDIA HGX™ B200 (LC)
PCIe Slot (from MLB)	Option(1) 1x FHHLDW Option(2) 2x FHHLSW Option(3) 1x OCP 3.0 SFF + 1x FHHLSW
PCIe Slot (from Switch)	8x Low Profile card (PCIeG5x16)
Storage (from Switch)	8x U.2 NVMe SSD bays
Storage (from MLB)	4x U.2 NVMe SSD bays from CPU 2x M.2 SATA SSD or NVMe SSD
Management Port	1x RJ45 for BMC(AST2600) remote management
Cooling	Liquid cooling
Power Supply	3300W 54V PSU, support 3+3 redundancy 2400W 12V PSU, support 1+1 redundancy
Fan	5x 8038 Fan for GPU cooling 5x 8056 Fan for CPU cooling

### Inventec Corporation (TAO)

No.88, Dazhi Rd., Taoyuan Dist., Taoyuan City 33068, Taiwan  
Tel: 886-3-390-0000 Fax: 886-3-376-2370  
Email: TAOproductsupport@inventec.com

Learn more at <https://ebg.inventec.com/>



All title and intellectual property rights in and to this document, the Specifications and photos contained therein, remain the exclusive property of Inventec or its suppliers. Intel, the Intel logo, Xeon, the Xeon inside, are trademarks or registered trademarks of Intel Corporation in the U.S. and/or other countries. All trademarks and logos are the properties of their representative holders