




QCT Next Gen Server Solutions Powered by AMD EPYC™ 9005 Series Processors



QCT Next Gen Server Solutions powered by the latest AMD EPYC™ 9005 Series processors and AMD Instinct™ MI325X accelerators feature the AMD “Zen 5” and AMD CDNA™ 3 architectures to offer leading AI performance and energy efficiency for training and inference tasks.

AMD
EPYC

AMD Instinct™ MI325X GPU Accelerators



Memory:

Up to 256GB of HBM3E memory

Memory bandwidth:

6TB/s peak memory bandwidth

Architecture:

AMD CDNA™ 3

Chiplets:

Eight compute, four I/O, and eight memory chiplets using advanced 2.5D and 3D packaging technology

Built for enterprise AI, offering seamless integration for fast, cost-effective upgrades

Drop-in compatible with existing AMD Instinct™ MI300X infrastructures

Infinity Fabric™ technology for seamless data sharing

AMD ROCm™ software to accelerate AI inferencing and training



Elevated Data Center Performance with AMD EPYC™ Series Processors

Max core count per socket

128 cores/ 256 threads

96 cores/ 192 threads

Max TDP

500W

400W

L3 cache

Up to 512 MB

Up to 384 MB

Memory channels / socket

12 channels, 2DPC

12 channels, 2DPC

Memory frequency

Up to 6000 MHz 1DPC

Up to 4800 MHz 1DPC

CPU interlinks

GMI-3, 32GT/s

GMI-3, 32GT/s

PCIe Lane

PCIe 5.0 128 lanes

PCIe 5.0 128 lanes

QCT Product Info

AMD
EPYC

Enterprise

HPC

AI/Acceleration




SDS/HCI



Virtualization



Design & Simulation



IC design & EDA



Machine Learning



AI Training & Inferencing



QuantaGrid S44NL-1U



QuantaGrid D44N-1U



QuantaGrid D74A-7U



QuantaGrid D75T-7U



QuantaGrid D75M-5U

QuantaGrid D74A-7U



Advanced Performance for the Most Extreme AI-HPC Workloads



- Powered by dual AMD EPYC™ 9004 or 9005 Series Processors with up to 500W TDP and AMD Instinct™ MI300X or MI325X GPU accelerators
- UBB architecture to support AI & HPC workloads
- PCIe 5.0 OCP expansion slots for RDMA networking
- Modularized system design for optimal flexibility and easy serviceability
- High NVMe storage density, supporting up to 18x 2.5" NVMe drives
- Accelerated parallel computing performance for the most extreme AI-HPC workloads
- Enhanced serviceability with tool-less, hot-pluggable designs

Processor	(2) AMD EPYC™ 9004 or 9005 Series Processors, up to 500W TDP
Memory	Up to 6TB memory capacity with (24) DDR5 RDIMM up to 6000 MHz
Networking	Dedicated (1) GbE management port
Storage	Front (18) 2.5" hot-pluggable NVMe drives
Expansion Slot	(2) PCIe 5.0 x16 OCP 3.0 SFF slots (10) PCIe 5.0 x16 OCP 3.0 TSFF slots
GPU Expansion Slot	(8) Accelerator modules on one single baseboard (AMD Instinct™ MI300X/ MI325X and NVIDIA HGX™ H100 8-GPU baseboard)
Dimensions	447.8 mm x 307.85 mm x 886 mm
Form Factor	7U rackmount server

QuantaGrid D75T-7U



Transform Computational Capabilities and Achieve High Performance



- Powered by dual AMD EPYC™ 9005 Series Processors with up to 500W TDP and AMD Instinct™ MI300X/MI325X accelerators
- UBB architecture to support AI & HPC workloads
- PCIe 5.0 expansion slots for RDMA networking
- Modularized system design for optimal flexibility and easy serviceability
- High NVMe storage density, supporting up to 18x 2.5" NVMe drives.
- Accelerated parallel computing performance for the most extreme AI-HPC workloads
- Enhanced serviceability with tool-less, hot-pluggable designs

Processor	(2) AMD EPYC™ 9005 Series Processors, up to 500W TDP
Memory	Up to 6TB memory capacity with (24) DDR5 RDIMM up to 6000 MHz
Networking	Dedicated (1) GbE management port
Storage Expansion Slot	Front (18) 2.5" hot-pluggable NVMe drives (2) PCIe 5.0 x16 FHHL slots (8) PCIe 5.0 x16 HHHL slots
GPU Expansion Slot	(8) Accelerator modules on one single baseboard (AMD Instinct™ MI300X/ MI325X 8-GPU baseboard)
Dimensions	447.8 mm x 307.85 mm x 950 mm
Form Factor	7U rackmount server

QuantaGrid D75M-5U



Advanced Performance for the Most Extreme AI-HPC Workloads



- Powered by AMD EPYC™ 9004 or 9005 Processors with up to 500W TDP
- Liquid cooling solution designed for AMD CPU platforms
- All-NVMe drive bays for GPUDirect® Storage
- Flexible with NVIDIA topology for BlueField®-3 DPU and SuperNIC
- PCIe 5.0 expansion slots for GPUDirect® RDMA
- Accelerated parallel computing performance for the most extreme AI-HPC workloads
- Enhanced serviceability with tool-less, hot-pluggable designs

Processor	(2) AMD EPYC™ 9004 or 9005 Series Processors, up to 500W TDP
Memory	Up to 6TB memory capacity with (24) DDR5 RDIMM up to 6000 MHz
Networking	Dedicated (1) GbE management port
Storage Expansion Slot	Front (18) 2.5" hot-pluggable NVMe drives (2) PCIe 5.0 x16 FHHL slots (8) PCIe 5.0 x16 HHHL slots
GPU Expansion Slot	(8) Accelerator modules on one single baseboard (NVIDIA HGX™ H200 8-GPU baseboard)
Dimensions	447.8 mm x 219.5 mm x 950 mm
Form Factor	5U rackmount server

QuantaGrid S44NL-1U



Cost-effective Server with Scalable Expansion Options and Optimized Performance



- Supports 1x AMD EPYC™ 9004 or 9005 Processor, up to 500W TDP
- Advanced air-cooling architecture supporting top-bin CPUs
- Liquid cooling-ready by design
- Up to 5x PCIe 5.0 expansion slots and the DC-SCM architecture to meet different configuration needs
- Featuring up to 2x single-width accelerators to support AI inference workloads
- System optimized for cloud, enterprise, AI, HPC, networking, security and IoT workloads

Processor	(1) AMD EPYC™ 9004 or 9005 Series processors, up to 500W TDP
Memory	Up to 6TB memory capacity with (24) DDR5 RDIMM up to 5200 MHz
Networking	Dedicated (1) GbE management port
Storage	All U.2 Front (12) 2.5" hot-pluggable NVMe/SATA/SAS drives All E1.S Front (12) or (16) E1.S hot-pluggable NVMe drives
Expansion Slot	Option 1-1 (2) PCIe 5.0 x16 FHHL slots (2) PCIe 5.0 x16 OCP 3.0 slots Option 1-2, 2-1 (2) PCIe 5.0 x16 HHHL slots (2) PCIe 5.0 x16 OCP 3.0 slots Option 1-3, 2-2 (3) PCIe 5.0 x16 HHHL slots (2) PCIe 5.0 x16 OCP 3.0 slots
Dimensions	440 mm x 43.2 mm x 780 mm
Form Factor	1U rackmount server

QuantaGrid D44N-1U



Ultimate 1U Server with AMD EPYC Performance Breakthrough



- Powered by dual AMD EPYC™ 9004 or 9005 Series Processors, up to 500W TDP
- Advanced air-cooling architecture supporting top-bin CPUs
- Liquid cooling-ready by design
- Up to 5x PCIe 5.0 expansion slots and the DC-SCM architecture to meet different configuration needs
- Featuring up to 2x single-width accelerators to support AI inference workloads
- System optimized for cloud, enterprise, AI, HPC, networking, security and IoT workloads

Processor	(2) AMD EPYC™ 9004 or 9005 Series processors, up to 500W TDP
Memory	Up to 6TB memory capacity with (24) DDR5 RDIMM up to 6000 MHz
Networking	Dedicated (1) GbE management port
Storage	All U.2 Front (12) 2.5" hot-pluggable NVMe/SATA/SAS drives All E1.S Front (12) or (16) E1.S hot-pluggable NVMe drives
Expansion Slot	Option 1-1 (2) PCIe 5.0 x16 FHHL slots (2) PCIe 5.0 x16 OCP 3.0 slots Option 1-2, 2-1 (2) PCIe 5.0 x16 HHHL slots (2) PCIe 5.0 x16 OCP 3.0 slots Option 1-3, 2-2 (3) PCIe 5.0 x16 HHHL slots (2) PCIe 5.0 x16 OCP 3.0 slots
Dimensions	440 mm x 43.2 mm x 780 mm
Form Factor	1U rackmount server

About QCT

Quanta Cloud Technology (QCT) designs, manufactures, integrates and services cutting edge offerings for 5G Telco/Edge, AI/HPC, Cloud, and Enterprise infrastructure via its own global network. Product lines include hyper-converged and software-defined data center solutions as well as servers, storage, and network switches from 1U to entire racks with a diverse ecosystem of hardware components and software partners to fit a variety of business verticals and workload parameters.

<https://www.qct.io/>

www.QCT.io



Contact Us



© 2024 Quanta Computer Inc. All rights reserved. Specifications and figures are subject to change without prior notice.

All other brand trademarks, logos, and names are the property of their respective owners. All campaign statements and product images contained herein are copyrighted and may not be reprinted and/or reproduced, in whole or in part without the written consent of Quanta Computer Inc.

QCT, the QCT logo, Rackgo, Quanta, and the Quanta logo are trademarks or registered trademarks of Quanta Computer Inc. All trademarks and logos are the properties of their respective holders.

AMD, the AMD arrow logo, EPYC, and combinations thereof are trademarks of Advanced Micro Devices, Inc.

ver Oct., 2024