Cloud Computing Infrastructure Hardware Requirement - Official Technical Overview & Hardware Datasheet

EXECUTIVE SUMMARY

As the foundational layer of modern cloud computing, physical infrastructure hardware directly dictates the performance, scalability, and reliability of virtualized services. This document provides a comprehensive technical overview of the reference hardware requirements for cloud computing infrastructure, including compute nodes, storage arrays, and fabric interconnect switches. The specifications outlined herein address multi-tenant isolation, east-west traffic patterns, and near-zero downtime operational models.

ARCHITECTURE & CHASSIS DESIGN

The cloud-optimized hardware platform is organized into three discrete plane domains: the Compute Plane (x86/ARM-based servers with hardware offload engines), the Storage Plane (NVMe-oF and distributed block storage enclosures), and the Fabric Plane (Ethernet spine-leaf switches with congestion control). The modular chassis architecture supports front-serviceable compute sleds, redundant power distribution units (PDUs), and independent fan trays with N+2 redundancy. Backplane signaling supports PCIe Gen5 and CXL 2.0 interconnect fabrics for memory pooling and accelerator attachment.

HARDWARE FEATURES

- Processor Sockets: Dual-socket configuration supporting 64 cores per socket (AMD EPYC 4th Gen / Intel Xeon Scalable 4th Gen)
- Memory Capacity: 32 DIMM slots per chassis, supporting DDR5-4800 ECC RDIMMs up to 8TB total
- Storage Controllers: Tri-mode RAID-on-chip (RoC) supporting NVMe, SAS4, and SATA drives
- Network Interface: Dual-port 100GbE/200GbE per compute node with RDMA over Converged Ethernet (RoCEv2) offload
- Management: Dedicated 1GbE out-of-band management port with Redfish API compliance

COMPLIANCE & STANDARDS

The hardware platform complies with: OCP Open Compute Project v3.0, ODCC (Open Data Center Committee) mechanical specifications, PCI-SIG Gen5 certification, NVMe Express 2.0, IEEE 802.1Qau (Congestion Notification), ENERGY STAR for Servers v3.0, and ASHRAE Class A4 environmental guidelines.

TECHNICAL SPECIFICATIONS

Parameter	**Specification**
Form Factor	1RU / 2RU / 4RU Chassis, 19-inch rack mountable
Switching Capacity	3.2 Tbps per 1RU compute node; 12.8 Tbps per spine switch (non-blocking)
Power Supply	1+1 Redundant, 2300W AC (200-240V) or DC (-48V), Titanium efficiency
Operating Temperature	5°C to 40°C (41°F to 104°F), up to 45°C for short duration
Cooling	N+2 redundant hot-swappable fan trays, front-to-back airflow
MTBF	280,000 hours at 40°C ambient

Weight (Fully Loaded)	18 kg (1RU), 32 kg (2RU), 55 kg (4RU)
Security	TPM 2.0, Secure Boot, hardware root of trust, side-band authentication

ORDERING OPTIONS

The cloud infrastructure hardware is available in three base configurations:

- CLOUD-BASE-1: 1RU compute node, single socket, 16 DIMMs, 4x NVMe drives, dual 25GbE

- CLOUD-STD-2: 2RU compute/storage hybrid, dual socket, 32 DIMMs, 12x SAS/NVMe bays, quad 100GbE

- CLOUD-HPC-3: 4RU accelerator-optimized, dual socket, 32 DIMMs, 8x GPU PCIe slots, 8x 200GbE

All configurations include 1+1 redundant hot-swap power supplies (2300W Titanium efficiency) and N+2 redundant fan assemblies. Extended warranty and on-site support contracts are available in 3-year and 5-year increments.