

Next-Generation Enterprise Server Infrastructure: Vantageo 2260 & 22T1 with GPU-Accelerated RAID

Technical White Paper

March 2026

Vantageo India Pvt Ltd

Enterprise Server Solutions Division

Executive Summary

Modern enterprise workloads artificial intelligence, high-performance computing, cloud infrastructure, and data analytics demand unprecedented compute density, memory bandwidth, and storage throughput. Traditional server architectures with conventional RAID controllers increasingly struggle to deliver the performance required for NVMe-dense configurations, creating bottlenecks that undermine expensive CPU and GPU investments.

Vantageo addresses these challenges with two flagship platforms: the **2260 series** powered by Intel Xeon 6 processors and the **22T1** featuring AMD EPYC 9005/9004 processors. Both platforms integrate Graid Technology's SupremeRAID SR-1000 (model SR-BUN-1000-FD12), delivering GPU-accelerated RAID that eliminates storage bottlenecks while maintaining enterprise-grade data protection[1][2].

This white paper examines the architectural advantages, performance characteristics, and competitive positioning of these platforms against established solutions from HPE and Dell.

Market Context and Competitive Landscape

The NVMe Performance Gap

Enterprise servers from HPE (ProLiant DL380/DL360) and Dell (PowerEdge R650/R750) continue to dominate data center deployments. However, their traditional RAID architectures face fundamental limitations when scaling NVMe storage:

- **HPE Smart Array controllers** deliver up to 46 GB/s bidirectional throughput and 3.5 million IOPS across 32 lanes, creating bottlenecks in dense NVMe configurations[3]
- **Dell PERC11 NVMe RAID** achieves 15.5× bandwidth improvement over previous generations but still caps at "tens of GB/s" aggregate throughput due to inline controller architecture[4]
- Both solutions force all storage traffic through PCIe-connected RAID controller ASICs, fundamentally limiting scalability as NVMe drive counts increase[5]

The GPU-Accelerated RAID Alternative

Grain SupremeRAID SR-1000 employs a fundamentally different architecture: data flows directly between CPU and NVMe drives via native PCIe fabric, while an Nvidia T1000 GPU handles RAID parity calculations out-of-path[6]. This approach delivers:

- Up to 220 GB/s sequential throughput and 16 million IOPS in optimal configurations[7]
- Near-linear performance scaling from 4 to 32 NVMe drives without controller bottlenecks[8]
- Minimal CPU overhead (<2%) compared to software RAID solutions that can consume 20-40% of CPU cycles[9]

For AI workloads with 16-24 NVMe drives, real-world testing shows 100-200+ GB/s protected throughput—2-4× higher than comparable HPE/Dell RAID configurations[8][10].

Vantageo 2260 Series: Intel Xeon 6 Platform

Platform Overview

The Vantageo 2260 is a 2U dual-socket server supporting Intel Xeon 6 processors (6th generation, Birchstream architecture). Available in multiple variants optimized for different workload profiles, the platform targets compute-intensive AI inference, high-performance computing, virtualization, and data-intensive storage arrays[11].

Technical Specifications

Processor Architecture

- Sockets: 2× LGA4710 (Socket E2)
- Processor Support: Intel Xeon 6700E-series, 6500P-series, 6700P-series
- TDP Support: Up to 350W per socket with air cooling
- Core Options: Up to 288 cores per socket (P-core variants)[12]
- AI Acceleration: Intel AVX2, VNNI (Vector Neural Network Instructions), fast-convert for BF16/FP16[12]

Memory Subsystem

- Capacity: 32 DDR5 DIMM slots (16 per CPU, 2DPC configuration)
- Maximum Memory: Up to 8,192 GB (8 TB) using 256 GB RDIMMs
- Memory Technology: DDR5 RDIMM, RDIMM-3DS, MRDIMM support
- Memory Speed:
 - RDIMM: Up to 6400 MT/s (1DPC), 5200 MT/s (2DPC)
 - MRDIMM: Up to 8000 MT/s (1DPC only, P-core processors only)

- Architecture: 8-channel DDR5 per processor (2× memory bandwidth vs. previous generation)[11][13]

Storage Configuration (2260-RM variant)

- Front Bays: 24× hot-swap 2.5" U.2 NVMe drive bays (PCIe 5.0 ×4 per drive)
 - 12 drives from CPU0, 12 drives from CPU1
- Rear Storage: 2× hot-swap E1.S 9.5mm NVMe drive bays (PCIe 5.0 ×4, from CPU1)
- M.2 Slots: 2× M.2 2280/22110 NVMe slots (PCIe 5.0 ×2, from CPU0)
- Total NVMe Capacity: Up to 26 NVMe drives per chassis
- RAID Support: Intel VROC (Virtual RAID on CPU) with optional key, or Graid SupremeRAID SR-1000[11][14]

Expansion and I/O

- PCIe Slots:
 - 2× PCIe 5.0 ×16 FHFL double-wide slots (supports up to 400W PCIe cards, 1 per CPU)
 - 1× PCIe 5.0 ×16 FHFL double-wide slot (CPU1)
- OCP Network: 1× PCIe 5.0 ×16 OCP 3.0 NIC mezzanine slot (CPU0, NCSI supported)
- Management: 1× 1000Base-T dedicated management port
- Front I/O: 1× USB 3.0, 1× USB 2.0, power/UID/reset buttons, status LEDs
- Rear I/O: 1× Mini DisplayPort, 1× COM port, 1× USB 2.0[11]

Server Management

- BMC: ASPEED AST2600 with AMI MegaRAC firmware
- Standards: IPMI 2.0, DMTF Redfish API
- Features: Dual BIOS, dual BMC, eMMC for local BMC storage
- Security: TPM 2.0 module support (optional), ASPEED AST1060 Hardware Root of Trust (optional)
- Remote Management: KVM-over-IP, virtual media, remote firmware updates[11]

Power and Cooling

- Power Supply: 1+1 redundant 2700W CRPS (Common Redundant Power Supply), 80 PLUS Titanium efficiency
- Input: AC 100-127V/12A or 200-240V/16A, 50-60Hz; DC 240V/16A option
- Cooling: 2× EVAC air cooling modules (350W CPU support), 6× hot-swap 6056 system fans
- Power Consumption: Optimized for <1.3 PUE in modern data centers[11]

Physical Specifications

- Form Factor: 2U rackmount (19" standard)
- Dimensions: 438mm (W) × 87mm (H) × 770mm (D) / 17.2" × 3.4" × 30.3"

- Weight: Approximately 26.25 kg (57.87 lbs) configured
- Operating Temperature: 0°C to 35°C (50°F to 95°F)
- Non-operating Temperature: -20°C to 70°C (-4°F to 158°F)[11]

Intel Xeon 6 Processor Advantages

Architectural Improvements

Intel Xeon 6 represents a significant generational leap over 5th Gen Xeon Scalable processors:

1. Dual Architecture Strategy
 - P-cores (Performance): Up to 288 cores per socket, optimized for high single-threaded performance and compute-intensive workloads including AI inference and HPC[12]
 - E-cores (Efficiency): Up to 144 cores per socket, optimized for containerized microservices, multi-tenant virtualization, and networking/edge workloads[15]
2. Memory Performance
 - 2× memory performance over 5th Gen Xeon with MRDIMM support
 - 37% greater memory bandwidth with MRDIMMs vs. standard DDR5 RDIMMs
 - Critical for memory-intensive HPC applications and large-scale AI model inference[13][16]
3. AI Acceleration
 - Built-in Intel AVX2 and VNNI instructions accelerate AI inference without discrete accelerators
 - Fast-convert instructions for BF16 and FP16 data types common in AI workloads
 - Up to 216 MB L3 cache with exceptionally low latency for large working sets[12][17]
4. I/O Throughput
 - Up to 192 PCIe Gen 5 lanes in dual-socket configurations (136 lanes single-socket)
 - Intel UPI 2.0 provides 20% higher inter-socket bandwidth vs. previous generation
 - Up to 64 lanes of CXL 2.0 at 32 GT/s per lane for memory expansion and accelerator connectivity[13]

Performance Positioning

Intel's benchmarking shows Xeon 6 delivering up to 40% average higher performance across broad workloads vs. previous generation, with particular strength in:

- AI inference and machine learning (better performance than any other general-purpose CPU per Intel claims)[12][15]
- High-performance computing with parallel workloads
- Virtualized environments requiring high single-thread performance

- Host CPU role for GPU-accelerated servers, where single-threaded performance drives faster data transfers to GPU accelerators[18]

Graid SupremeRAID Integration

SR-BUN-1000-FD12 Configuration

The 2260 series integrates Graid SupremeRAID SR-1000 as an alternative to traditional hardware RAID:

- GPU: Nvidia T1000 (PCIe Gen3 ×16) handles RAID parity calculations
- Drive Support: Up to 12 NVMe drives in FD12 license bundle (SR-1000 card supports up to 32 drives total)
- RAID Levels: 0, 1, 5, 6, 10 with GPU-accelerated parity
- Data Path: Direct CPU-to-NVMe via native PCIe fabric; GPU operates out-of-path for parity only[6][7]

Performance Advantages vs. HPE/Dell RAID

Metric	Graid SR-1000	HPE SmartRAID	Dell PERC11
Peak Sequential Throughput	100-220 GB/s	~46 GB/s	Tens of GB/s
Peak IOPS (Random 4K)	6M-16M+	Up to 3.5M	Up to 10M
Scaling Behavior	Linear, near-JBOD	Bottlenecks at high density	PCIe/ASIC limits
CPU Overhead	<2%	Low (HW offload)	Low (HW offload)
Architecture	Out-of-path GPU	Inline HW controller	Inline HW controller

Table 1: RAID performance comparison for 16-24 NVMe configurations

For AI workloads with 16-24 NVMe drives (Samsung PM1743 Gen5, Kioxia/Micron Gen4/5, WD/Intel enterprise NVMe), real-world measurements show:

- RAID5 with 8 NVMe: 30.7 GB/s read, 11.3 GB/s write; 8.88M read IOPS, 863K write IOPS[8]
- Scaling to 24 NVMe: 100-200+ GB/s aggregate throughput, maintaining 60-70% of raw JBOD performance with full RAID protection[8][10]
- AI/HPC workloads: RAID5 delivers 64 GB/s read, 44 GB/s write vs. direct JBOD 102 GB/s—acceptable efficiency with data protection[10]

By comparison, HPE SmartRAID caps at ~46 GB/s total bidirectional bandwidth even with 32 lanes, while Dell PERC11 hits only ~50% of NVMe potential at 32 drives due to inline controller architecture[3][4][5].

Target Workloads and Use Cases

AI Inference and Fine-Tuning

- High core count (up to 288 P-cores) and built-in VNNI instructions accelerate small-to-medium language models (up to 13B parameters) without discrete GPUs
- High memory bandwidth (DDR5-6400/8000) supports large model weights in-memory
- SupremeRAID enables 100+ GB/s training data pipeline feeding GPU accelerators in expansion slots[11][12]

High-Performance Computing

- Dual 288-core processors deliver 576 cores per 2U chassis for parallel scientific computing
- Large L3 cache (up to 216 MB) and low-latency memory access benefit HPC applications
- NVMe-dense storage supports high-throughput checkpoint/restart and scratch space[12][13]

Cloud Computing and Virtualization

- High core density consolidates more VM instances per chassis
- Intel VT-x and VT-d virtualization extensions with IOMMU support
- Fast storage fabric supports virtual machine storage requirements without bottlenecks[11][15]

Content Delivery Networks and Caching

- Large NVMe capacity (up to 26 drives × multi-TB) provides massive cache tier
- 100+ GB/s aggregate throughput serves high concurrent user request loads
- Redundant power and RAID protection ensure content availability[11]

Vantageo 22T1: AMD EPYC 9005/9004 Platform

Platform Overview

The Vantageo 22T1 is a 2U dual-socket AMD EPYC SP5 platform supporting both current-generation EPYC 9005 (Zen 5, 3nm process) and prior-generation EPYC 9004 (Zen 4, 5nm process) processors. The platform targets AI-driven data center workloads, GPU-accelerated computing, cloud infrastructure, and applications requiring maximum core density and memory bandwidth[19][20].

Technical Specifications

Processor Architecture

- Sockets: 2× AMD SP5 (LGA 6096)
- Processor Support: AMD EPYC 9005 series (Zen 5) and EPYC 9004 series (Zen 4)
- Core Options: Up to 192 cores per socket (384 cores dual-socket), 384 threads per socket
- TDP Support: Up to 500W per socket (varies by processor model)

- Manufacturing Process: 3nm (EPYC 9005) or 5nm (EPYC 9004)[20][21]

Memory Subsystem

- Capacity: 24 DDR5 DIMM slots total (12 per CPU, 1DPC configuration)
- Maximum Memory: Up to 6,144 GB (6 TB) using 256 GB RDIMMs or LRDIMMs
- Memory Technology: DDR5 RDIMM, LRDIMM support
- Memory Speed: Up to DDR5-6000 (EPYC 9005), DDR5-4800 (EPYC 9004)
- Architecture: 12-channel DDR5 per processor for maximum bandwidth[19][21]

Storage Configuration

- Front Bays: 12× 2.5" hot-swap drive bays supporting NVMe, SAS, or SATA
- Storage Flexibility: Mix of NVMe, SAS, and SATA drives in different tiers
- PCIe Gen5 NVMe Support: High-speed storage interfaces for Gen5 NVMe SSDs
- M.2 Support: PCIe Gen4 M.2 slots for boot drives or caching[19]

Expansion and I/O

- PCIe Slots: 4× PCIe Gen4 ×16 slots (full-height, full-length)
 - Slot 4: PCIe Gen4 ×16 (×16 link) from CPU0
 - Additional slots distributed across CPUs for GPU or storage controller expansion
- GPU Support: Up to 4 dual-slot GPUs for AI acceleration workloads
- Networking:
 - 2× 10GbE BASE-T LAN ports (onboard)
 - 1× 1GbE management LAN port (dedicated)
- Rear I/O: 2× USB 3.2 Gen1 (Type-A), 1× VGA, 1× COM, 2× RJ45 LAN, 1× management LAN[19]

Server Management

- BMC: ASPEED AST2600 Baseboard Management Controller
- Standards: Industry-leading BMC for remote monitoring, troubleshooting, and firmware updates
- Security: TPM 2.0 header with SPI interface (optional TPM 2.0 module)[19]

Power and Physical

- Form Factor: 2U rackmount (19" standard)
- Power Supply: Redundant high-efficiency power supplies (specifications vary by configuration)
- Cooling: Optimized airflow design for up to 500W TDP processors and multi-GPU configurations[19]

AMD EPYC 9005/9004 Processor Advantages

EPYC 9005 Series (Zen 5, 3nm)

AMD's 5th generation EPYC processors deliver substantial improvements over previous generations:

1. **AI and Machine Learning Performance**

- Full 512-bit data paths in core for AI workloads
- Instruction pipeline optimizations and additional integer/floating-point units
- Up to 1.9× AI inference throughput vs. previous generation (dual 9965 CPUs)[22]
- Leadership AI inference performance for language models up to 13B parameters running on CPU-only infrastructure[23]

2. **GPU Acceleration Host CPU Role**

- Optimized high-frequency variants (e.g., EPYC 9575F) for GPU-accelerated servers
- Up to 13% faster time-to-first-token and 6.6% higher overall inference throughput vs. Intel Xeon 6960P in 8-GPU configurations[23]
- Maximizes ROI of GPU investments by eliminating CPU bottlenecks in data feeding[23]

3. **Core and Thread Density**

- Up to 192 cores, 384 threads per socket (384 cores, 768 threads dual-socket)
- Up to 5 GHz boost frequencies on select SKUs
- Up to 384 MB L3 cache for large working sets[21][24]

4. **Memory and I/O**

- 12 DDR5-6000 memory channels per socket, up to 6 TB total capacity
- Up to 160 PCIe 5.0 lanes per socket for maximum I/O bandwidth
- CXL support for memory expansion and accelerator connectivity[24]

EPYC 9004 Series (Zen 4, 5nm)

For customers seeking cost-optimized performance:

- Up to 96 cores, 192 threads per socket (Zen 4 architecture)
- DDR5-4800 memory support with 12-channel architecture
- PCIe Gen5 support for storage and accelerators
- Proven track record in enterprise deployments since 2022[19][21]

Performance Positioning vs. Intel

AMD positions EPYC 9005 as delivering:

- Higher core counts (192 vs. 128 maximum for Intel Xeon 6 P-cores)
- Greater memory capacity (up to 6 TB vs. 4 TB typical Intel configurations)
- More PCIe lanes (up to 160 vs. 136 single-socket Intel)
- Competitive or superior AI inference performance depending on workload[23][24]

Graid SupremeRAID Integration

The 22T1 platform supports Graid SupremeRAID SR-1000 for NVMe RAID acceleration identical to the 2260 series. For 12-bay front storage configurations, SupremeRAID provides:

- **Aggregate Throughput:** Up to 100+ GB/s with 12 Gen5 NVMe drives in RAID5/10
- **IOPS:** Millions of random IOPS without CPU overhead
- **Data Protection:** Enterprise RAID 5/6/10 with near-JBOD performance
- **Scalability:** Add more NVMe via PCIe expansion cards while maintaining SupremeRAID management[7][10]

Target Workloads and Use Cases

AI Training and Inference

- Extreme core density (up to 384 cores dual-socket) for parallel AI model training
- High memory bandwidth feeds large model weights and training batches
- GPU slots support NVIDIA A100/H100 or AMD Instinct accelerators
- SupremeRAID eliminates storage bottlenecks when feeding multi-GPU training pipelines[19][23]

Database Processing

- 12-channel memory architecture and large L3 cache optimize in-memory databases
- High core counts enable massive parallel query execution
- NVMe storage supports high-transaction OLTP workloads[24]

Virtualization and Cloud Computing

- High core density maximizes VM consolidation ratios
- Large memory capacity supports memory-intensive virtualized workloads
- PCIe Gen5 and NVMe provide fast storage for virtual machine images[19][21]

Scientific Computing and HPC

- Leadership core count for parallel simulations and modeling
- High memory bandwidth and capacity support large datasets
- PCIe Gen5 connectivity for InfiniBand/Ethernet HPC fabrics[24]

Comparative Analysis: Vantageo vs. HPE/Dell

Storage Performance Leadership

Platform	Max Throughput	Max IOPS	Bottleneck
Vantageo 2260 + Graid SR-1000	100-220 GB/s	16M+	CPU/PCIe fabric
Vantageo 22T1 + Graid SR-1000	100-220 GB/s	16M+	CPU/PCIe fabric
US Vendor 1 + SmartRAID	~46 GB/s	3.5M	RAID controller
US Vendor 2 + PERC11	Tens of GB/s	10M	RAID controller

Table 2: Storage subsystem performance comparison

Key Advantage: Vantageo platforms with SupremeRAID scale NVMe performance linearly without controller bottlenecks, delivering 2-4× higher aggregate throughput in dense NVMe configurations (16-24 drives). This is critical for AI training data pipelines, large-scale databases, and content delivery workloads[3][4][8][10].

Total Cost of Ownership

Initial Acquisition Cost

- Vantageo servers typically price competitively with Tier-2 OEMs and below Tier-1 (HPE/Dell) premium pricing
- Graid SupremeRAID SR-1000 adds incremental cost vs. basic RAID controllers, offset by eliminating need for expensive high-end RAID cards

Performance per Dollar

- Higher NVMe throughput enables fewer servers for same aggregate performance, reducing:
 - Data center space and power costs
 - Networking infrastructure requirements
 - Software licensing costs (per-server or per-core licensing models)

Operational Efficiency

- SupremeRAID's minimal CPU overhead (<2%) frees compute resources for revenue-generating workloads vs. software RAID consuming 20-40% CPU[9]
- Lower power consumption per GB/s throughput improves data center PUE

Support and Service

Vantageo provides comprehensive support across India:

- 200+ support service locations with nationwide coverage
- 283 highly skilled support engineers across India
- 24/7/365 helpdesk with multiple response tiers:
 - 4-hour response, 6-hour call-to-resolution CTR (critical)
 - 4-hour response, 24-hour call-to-resolution CTR (standard)
 - Next Business Day call-to-resolution CTR (basic)
- 11 spares depots in metro cities with next-flight-out parts delivery
- Consulting and deployment services available[11][25]

Implementation Considerations

Sizing Guidelines

For 2260 Series (Intel Xeon 6)

Choose 2260 when workloads require:

- Best single-threaded performance (P-core variants)
- Intel-specific software certifications or optimizations

- Maximum memory bandwidth with MRDIMM support (8000 MT/s)
- Highest per-core AI inference performance

For 22T1 (AMD EPYC 9005/9004)

Choose 22T1 when workloads require:

- Maximum core density (up to 384 cores dual-socket)
- Highest aggregate memory capacity (up to 6 TB)
- Most PCIe lanes for multi-GPU or storage expansion
- Cost-optimized performance with EPYC 9004 processors

Graid SupremeRAID Deployment

Optimal Use Cases

- 8+ NVMe drives per server (cost-justification threshold)
- AI/ML workloads with high-throughput training data requirements
- Databases with intense random I/O patterns
- Content delivery requiring 100+ GB/s aggregate throughput

Alternative Considerations

- Configurations with <6 NVMe drives may use Intel VROC or software RAID
- Workloads not sensitive to storage throughput can use basic RAID controllers
- Budget-constrained deployments can deploy SupremeRAID in a subset of nodes

NVMe Drive Selection

Performance Tier (AI/HPC)

- Samsung PM1743 (Gen5, up to 13 GB/s sequential read)
- Micron 7500 series (Gen5, high endurance)
- Kioxia CM7 series (Gen5, data center optimized)

Balanced Tier (Cloud/Virtualization)

- WD Ultrastar DC SN840 (Gen4, proven reliability)
- Intel D5-P5520 (Gen4, consistent performance)
- Micron 7450 series (Gen4, cost-effective)

Capacity Tier (Storage Arrays)

- Samsung PM9A3 (Gen4, high capacity points)
- WD Red SN700 (Gen3, cost-optimized)

Conclusion

The Vantage 2260 (Intel Xeon 6) and 22T1 (AMD EPYC 9005/9004) platforms represent next-generation enterprise server architecture, addressing the performance limitations of traditional designs in NVMe-dense configurations. By integrating Graid SupremeRAID SR-1000 GPU-accelerated RAID, these platforms deliver:

1. **2-4× higher storage throughput** vs. comparable to top US servers vendors equipped with Smart Array or PERC RAID controllers
2. **Linear NVMe performance scaling** from 8 to 24+ drives without controller bottlenecks
3. **Minimal CPU overhead** (<2%) while maintaining enterprise RAID data protection
4. **Competitive total cost of ownership** through performance consolidation and operational efficiency

For AI training pipelines requiring 100-200+ GB/s protected storage throughput, high-performance databases, content delivery networks, and compute-intensive HPC workloads, Vantageo platforms with SupremeRAID offer measurable advantages over incumbent solutions.

Organizations evaluating enterprise server refresh cycles should benchmark actual workload performance on Vantageo platforms against current US Vendors deployments to quantify these benefits in their specific environments.

References

- [1] DiGiCOR Australia. (2024). RC-GRT-SR-BUN-1000-FD12 product page. <https://www.digicor.com.au/rc-grt-sr-bun-1000-fd12>
- [2] Graid Technology. (2021). SupremeRAID SR-1000 Spec Sheet v1.2.1. <https://download.graidtech.com/document/spec-sheet/sr-1000/>
- [3] HPE. (2024). HPE SmartRAID Storage Controller specifications. <https://edshop.edsystem.eu/fileshare/indexes/specification/>
- [4] Dell Technologies. (2023). NVMe Performance Increases for Next-Generation PowerEdge Servers with PERC11 Controller. Dell Technologies Info Hub. <https://infohub.delltechnologies.com/>
- [5] Reddit r/sysadmin. (2018). Discussion: HP Smart Array and NVMe SSD RAID performance limitations. <https://www.reddit.com/r/sysadmin/comments/92gaph/>
- [6] Boston India. (2023). SupremeRAID by Graid Technology – Next Generation RAID architecture overview. <https://www.bostonindia.in/blog/2023/06/07/supremeraid.aspx>
- [7] Graid Technology. (2023). GraidTech SupremeRAID SR-1000 Spec Sheet v1.2.9. <https://download.graidtech.com/document/spec-sheet/sr-1000/>
- [8] Storage Review. (2021). GRAID SupremeRAID SR-1000 Review - performance testing. <https://www.storagereview.com/review/graid-supremeraid-sr-1000-review>
- [9] Graid Technology. (2025). Technical Comparison: SupremeRAID vs. Software RAID. <https://graidtech.com/post/technical-comparison-supremeraid-tm-vs-software-raid>
- [10] Storage Review. (2024). Performance and Resiliency: Graid SupremeRAID for AI and HPC Workloads. <https://www.storagereview.com/review/performance-and-resiliency-graid-supremeraid-for-ai-and-hpc-workloads>
- [11] Vantageo. (2024). Vantageo 2260-RM Server Data Sheet. <https://vantageo.com/pdfs/2260-rm.pdf>

- [12] Intel. (2026). Intel Xeon 6 Architecture - Performance and Efficiency Cores. <https://www.intel.com/content/www/us/en/support/articles/000098612/>
- [13] Unicom Engineering. (2025). Intel Xeon 6 Powers Next-Generation High-Performance Computing. <https://www.unicomengineering.com/blog/intel-xeon-6-processors-the-ideal-solution-for-high-performance-computing-hpc/>
- [14] Vantageo. (2024). Vantageo 2260-RT Server Data Sheet. https://www.vantageo.in/files/Vantageo_2260-RT.pdf
- [15] EnGenius. (2026). What are Intel Xeon 6 Processors. <https://www.engeniustech.com/what-are-intel-xeon-6-processors/>
- [16] Kingston. (2024). Intel Xeon 6 Guide: Key Features, Benefits & Use Cases. <https://www.kingston.com/en/blog/servers-and-data-centers/intel-xeon-6-guide-features-benefits-use-cases>
- [17] Intel. (2026). Intel Xeon Processors - Server, Data Center, and AI Processors. <https://www.intel.com/content/www/us/en/products/details/processors/xeon.html>
- [18] Intel. (2024). Why choose Intel Xeon 6 processors as host CPUs? Product Brief. <https://www.intel.com/content/dam/www/central-libraries/us/en/documents/2024-12/xeon-6-processors-with-p-cores-ig.pdf>
- [19] Vantageo. (2024). Vantageo 22G1-RG Server Data Sheet (AMD EPYC 9004 reference architecture). <https://vantageo.com/pdfs/22g1-rg.pdf>
- [20] Red Hat Ecosystem Catalog. (2026). Vantageo 22T1 certification page. <https://catalog.redhat.com/en/hardware/system/detail/307397>
- [21] ASUS. (2024). AMD EPYC 9005 Processors overview. <https://www.asus.com/event/amd-epyc-9005-servers/>
- [22] AMD. (2024). AMD EPYC 9005 Series Processor Datasheet. <https://www.amd.com/content/dam/amd/en/documents/epyc-business-docs/datasheets/amd-epyc-9005-series-processor-datasheet.pdf>
- [23] AMD. (2025). 5th Generation AMD EPYC Server CPUs product page. <https://www.amd.com/en/products/processors/server/epyc/9005-series.html>
- [24] AMD. (2025). AMD EPYC Processor Overview. <https://www.amd.com/en/partner/articles/epyc-processor-overview.html>
- [25] Vantageo. (2024). Vantageo India Enterprise Servers - Support Services. <https://www.vantageo.in/enterprise-servers>