

H3C Unveils Full-Stack AI Infrastructure at NAVIGATE 2026 to Address Global AI Cost & Efficiency Challenges

Hangzhou, China – May 22, 2026 — At the NAVIGATE 2026 International Summit, global digital infrastructure provider H3C has officially launched its next-generation full-stack AI infrastructure solutions for the global market. Headlined by the flagship H3C UniPoD S80000 Series SuperPod, the launch brings disruptive innovations across six core technological domains, targeting critical pain points of soaring computational costs and cluster efficiency bottlenecks in the Agentic AI era.

To tackle these industry-wide hurdles, H3C has upgraded its "AI in ALL" technical strategy, adopting "Computing × Connectivity" as its core roadmap. Moving beyond fragmented hardware aggregation, the integrated full-stack system optimizes overall fabric efficiency and significantly reduces Total Cost of Ownership (TCO), enabling global enterprises to gain unprecedented price-performance in the emerging Token Economics landscape.

Built exclusively for trillion-parameter AI models, the UniPoD S80000 Series SuperPod delivers four core strengths. It features ultra-high density: each node integrates 1 CPU and 4 AI accelerators with full liquid cooling, paired with an 800V HVDC power supply to support deployments exceeding 350kW per rack, maximizing density while minimizing PUE. Its ultimate interconnection leverages 102.4Tbps switching silicon for nanosecond-level latency, boosting cluster bandwidth by 4x for 256-card setups and over 10x for 1024-card configurations to eliminate fabric congestion. Software optimization enables time-sharing multiplexing for "inference by day, training by night," lifting per-card inference efficiency by 3x; FlashCheckPoint ensures data persistence within 10 minutes, securing an effective training time ratio above 96%. The open architecture scales seamlessly from 32 up to 16,384 cards, natively supporting multi-vendor heterogeneous compute (CPU, GPU, NPU, DPU) and mainstream frameworks.

H3C also rolled out cross-domain upgrades under six pillars. In compute and storage, the Intel-co-developed UniServer G7 Server Series and X20000 Series AI-Native Storage deliver 200GB/s node bandwidth, 90% acceleration in TTFT, 35% faster data preparation and 30% less GPU idle time. For network and security, 102.4T single-chip AI switches, the world's first enterprise-grade Wi-Fi 8 AP and SecPath M9000 AI Firewall create an "AI-against-AI" security fabric. The CloudOS-powered AI cloud platform unifies general computing and MaaS capabilities to break down compute silos.

Positioned as a foundational architect of the AI era, H3C will advance "Computing × Connectivity" innovation, empowering global enterprises to turn data into business value amid the generative intelligence boom.