Solution Brief

CNode-X Accelerated Computing Platform

Stop Moving Data to AI. Bring AI to your Data.

The agentic era invalidates legacy infrastructure. The world is embarking on one of the largest technology refresh events in history as enterprises realize they need to uplevel their data infrastructure to feed agentic systems. Models need unrestricted access to all enterprise data instantly, yet 90% remains locked away in systems never designed for AI. Legacy architectures force enterprises to copy data to separate AI infrastructure, creating data silos, versioning conflicts, and compliance risks that prevent AI from reaching production.

VAST CNode-X introduces the industry's first end-to-end fully accelerated AI data stack. CNode-X combines NVIDIA RTX PRO 6000 Blackwell GPUs with VAST EBox data platform, making data systems an integral part of the compute framework. This embodies the NVIDIA AI Data Platform philosophy: bring AI capability to where your data already lives. It democratizes infrastructure capabilities that were previously the domain of large hyperscalers.

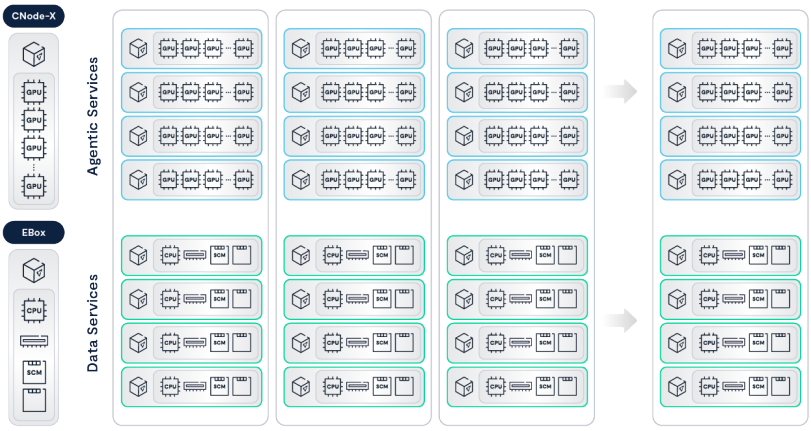

Although the form factor is hardware, this is a software partnership at its core. CNode-X expands GPU acceleration beyond training and inference into vector search, SQL analytics, and containerized AI model orchestration. Storage, database, event broker, functions, and AI models all run natively on the VAST AI Operating System making the idea of a thinking machine a reality.

Bring AI to Your Data
GPU compute runs where your data lives, eliminating data movement penalties.

Accelerate Everything
Industry's first fully accelerated AI data stack for all workloads.

Managed at Scale
Managed Kubernetes orchestrates GPU workloads with the data.

Real-Time, Multimodal Intelligence

By accelerating both compute and data services as one coherent system, CNode-X delivers real-time, multimodal intelligence across text, images, video, and unstructured enterprise content. For the first time in IT history, enterprises have tools to make sense of massive unstructured data reservoirs. GPU-accelerated workloads run locally with direct access to VAST EBox clusters that infinitely scale from petabytes to exabytes as your AI unlocks intelligence from your universally accessible data.

VAST integrates NVIDIA's acceleration libraries deep inside core platform services. NVIDIA cuVS brings GPU acceleration to vector search and clustering for real-time retrieval. Sirius provides GPU-native SQL acceleration. NVIDIA Nemotron models and NIM microservices enable scalable pipelines for video intelligence, document processing, and research. NVIDIA Context Memory Storage accelerates inference at scale for long-context, multi-agent workloads.

Architecture Built to Scale

This real-time processing becomes possible through VAST's breakthrough architecture. The platform runs as stateless containers orchestrated by VAST AI OS with managed Kubernetes. Every CNode-X sees every all data via direct NVMe-oF access without coordination overhead through VAST's Disaggregated Shared Everything architecture (DASE). Any node can access any data through a global namespace without sharding, east-west traffic, or resource coordination.

This eliminates the operational complexity of stitching together separate storage systems, vector databases, analytics engines, and AI infrastructure stacks. RAG pipelines that previously required separate GPU clusters, vector databases, and storage tiers now run as a unified workflow with all components accessing the same data corpus simultaneously.

Agents Built on Top of Data

With this foundation in place, enterprises can build agents directly on top of their data. CNode-X enables agent builders to deploy AI agents that access the complete enterprise data corpus in real time. Access policies flow automatically from raw data through vectors to inference with identity-aware security native to the platform. The same governance framework protecting files extends seamlessly to embeddings, database queries, and model inference without separate security layers.
The platform creates a recursive loop of learning where systems accumulate knowledge, receive new data, make correlations, and get smarter as those correlations accumulate. This continuous learning is driven by new data entering the system, enabling AI that evolves with your business.

From Experimentation to Production

These capabilities accelerate the path from experimentation to durable production systems. CNode-X provides a faster route to production for RAG pipelines, agentic systems, and continuous AI workloads. Development teams prototype on the same infrastructure that runs production systems, eliminating architectural rewrites and data migration projects that trap AI initiatives in pilot purgatory.

As workloads scale from terabytes to petabytes and from single models to multi-agent systems, the platform scales without infrastructure redesign. Organizations deploy dynamic agentic pipelines that process diverse data types and deliver actionable intelligence at scale.

This isn't an upgrade path. It's the foundation for graduating to a fundamentally different style of computing.

About VAST Data

VAST Data is the leading data platform software company bringing businesses into the AI era. By accelerating time-to-insight for workload-intensive applications, the VAST Data Platform delivers scalable performance, simplifies data management, and enhances productivity. Launched in 2019, VAST has become the fastest-growing data infrastructure startup in history.

Looking to simplify your data center and unlock insights from all of your data?
Contact us at hello@vastdata.com

©2026 VAST Data, Inc. All rights reserved. All trademarks belong to their respective owners.

02/2026