White paper

NCP High-Performance Storage Reference Design for NVIDIA HGX Servers

Designed & Validated for NVIDIA HGX Servers with VAST Data EBoxes

Table of Contents

Introduction **4**

VAST AI OS with NVIDIA HGX Servers **5**

VAST Data: Hardware Components **6**

Key specifications of each VAST EBox include: 7

The Converged Network 7

VAST AI OS: Software Architecture **8**

VAST DataStore 10

VAST DataBase 10

VAST DataSpace 10

VAST Management Service (VMS) 11

Multi-Tenancy 12

NVIDIA Recommended High Performance Storage (HPS) Sizing for NCP **13**

VAST Data Design for NCP Deployments **15**

Recommended Storage Networking 15

VLAN Configuration Setup **15**

NCP Deployments with 1024 GPUs **16**

VAST Data Solution Performance Validation **22**

Additional HPS Features and Considerations for NCP 22

Summary **23**

Appendix A: Consolidate Storage Silos, Accelerate Insight **24**

Appendix B: The Imperative for Secure Multi-Tenant Storage **25**

Zero Trust: The VAST AI OS meets the Data Pillar Framework 25

VAST Data's Secure Multi-Tenant Data Platform 26

Advanced Tenancy Capabilities 27

Benefits of a Zero-Trust Multi-Tenant Data Platform 28

Cost Optimization and Performance Benefits 28

Compliance and Auditability 28

Appendix C: VLAN Configuration and Verification 29

Introduction

The NVIDIA® Cloud Partner Reference Design (NCP RD) for Artificial Intelligence (AI) training and inference, featuring NVIDIA HGX™ Servers, is the next generation of data center AI architecture for cloud service providers (CSPs). This document covers the high-performance storage (HPS) reference design for the EBox integrated model from VAST Data.

About VAST Data and NVIDIA Cloud Partners

The VAST AI Operating System (VAST AI OS) is an integrated set of storage and data services that enable service providers with a comprehensive catalog of data-centric offerings, including multi-protocol data ingest, accelerated data preprocessing for feature engineering, high-performance storage (HPS) for model training, edge data management for model serving, and end-to-end cataloging to meet audit and compliance requirements. Compared to traditional AI infrastructure with siloed storage solutions, the VAST AI OS enables NVIDIA Cloud Partners (NCPs) to develop and deploy these data services to market faster while effortlessly meeting availability service level agreements (SLAs), and fortifying security with a true Zero Trust framework that fully supports multi-tenant, fine-grained quality-of-service (QoS) controls, and tenant-specific data encryption.

While the VAST AI OS is robust in capabilities, this document—in collaboration with NVIDIA—was created to detail a validated reference design that specifically adheres to the NCP Reference Design (NCP RD) specifications for HPS. The remainder of this document provides an overview of the VAST Data solution (both hardware and software), along with exact configurations that have been validated to exceed the performance requirements necessary for NCPs building AI infrastructure at scale.

VAST AI OS with NVIDIA HGX Servers

VAST Data's support of the NCP RD builds on years of experience supporting large-scale HGX Servers deployed with some of the largest NCPs in the world. This fully integrated, turnkey design is validated at scale and is backed directly by the VAST Data global support services team.

The VAST AI OS is the foundation for some of the world's largest GPU cloud deployments.

The VAST AI OS uniquely marries the high performance and scalability demanded by AI workloads with the simplicity of enterprise manageability and resiliency needed by CSPs. The result is an ideal platform for service operators thanks to VAST Data's industry leadership in the HPS market in terms of production uptime as well as total cost of ownership (TCO). The VAST AI OS gives customers native support for industry standard protocols like NFS, SMB, and S3, enhancing the agility of CSPs to support demanding AI and HPC workloads seamlessly and efficiently.

The VAST AI OS, built on the VAST Disaggregated Shared Everything (DASE) architecture, is engineered to fully exploit the performance and operational characteristics of NVMe flash storage.

By pioneering new, low-overhead erasure codes and advanced data reduction technology, VAST reduces the costs of all-NVMe storage to a level comparable to the combined cost of tiered, multi-cluster solutions. This allows NCPs to realize the performance benefits of flash at archive economics. With the VAST AI OS, NCPs can leverage the simplicity of standard file and object protocols while streamlining AI pipelines and eliminating the complexity and security challenges of managing a multi-tenant environment.

Benefits for NVIDIA Cloud Partners:

Service provider-grade reliability

Observed production systems realizing 99.9999% availability

Secure multi-tenant

Per tenant data encryption, flexible network segmentation, robust audit capabilities

Fine-grained workload isolation
Granular QoS policies prevent multi-tenant I/O contention ("noisy neighbor")

Infrastructure cost savings
Consolidate storage silos into a single, affordable all-flash solution and permanently eliminate the overhead of data copies and sprawl

Improved operational efficiency

Online operations (such as upgrades) never require downtime, leading to better customer experience with fewer admins required

Multi-Protocol Support
Simultaneous access to data via NFS and Object Storage protocols along with interchangeability in reading and writing to meet growing demands

VAST AI Operating System: Hardware Components

VAST offers a single, cohesive, software-defined AI Operating System that unifies the entire stack, from an AI application runtime down to high-performance storage. Built on the revolutionary DASE architecture, VAST eliminates the bottlenecks inherent in older designs, providing a purpose-built environment for AI agents to live, learn, and act. This integrated software approach delivers significant simplicity, reduces costs, and enables end-to-end optimization for performance and efficiency.

Our DASE architecture enables true parallelism, eliminating bottlenecks and simplifying infrastructure. Deploy one OS from edge to cloud, manage all your AI data and workloads, and accelerate your journey to building and deploying transformative AI applications and intelligent agents.

The DASE architecture, however, is often misunderstood and equated with a specific hardware architecture (i.e., the physical disaggregation of CPU controllers from SSDs). In reality, the innovation behind DASE is a software architecture comprising individual microservices (packaged as software containers) that can run on separate, dedicated systems or co-installed on the same physical server.

VAST deployments are flexible and can be built on dedicated, diskless, CPU (x86) servers – aka CNodes – to serve as a clustered I/O controller and storage management system, while a separate JBOF chassis – aka DBox – houses large amounts of NVMe drives with the sole purpose of servicing low-level SSD access.

The VAST EBox – a VAST reference design using industry-standard (x86) servers – is an alternative way to build large, scale-out VAST clusters. With this architecture, each EBox simultaneously serves both roles (CNode and DBox) in a single, integrated appliance. Both hardware deployment models have full software feature parity and no impact on scalability or aggregate performance.

For environments where performance-density is more critical (HPS requirements for NCPs fit this criteria), VAST EBox clusters tend to lower datacenter rack space requirements, power consumption, and total solution cost. EBox servers are readily available from a wide range of tier-1 OEM manufacturers.



Key specifications each VAST EBox include:

- **CPU:** Single AMD Genoa 9454P UP (48 Cores / 96 Threads) CPU, 2.75GHz
- **I/O Connectivity:** Dual-port Connect-X 7 for frontend (client-facing) & backend (internal-facing)
- **Memory:** 384GB (12 × 32 GB) DDR5-4800
- **Drives:** 8 x U.2/U.3 Hyperscale QLC drives and 2 x SCM drives
- **Total Capacity:** Up to 246 TB maximum raw capacity per EBOX
- **Form Factor:** 1 Rack Units (RU)
- **Power Consumption:** 0. 65 kW Active

The Converged Network

In the converged data network design, both the front-end and back-end network traffic for the EBoxes are hosted on the same SN5600/5610 Leaf switches.

The front-end traffic from the logical CNode containers serves all protocol access (NFS/S3/SMB/Block/Kafka) to the GPU nodes and also enables the VAST Management System (VMS) to be reached on the data network via a dedicated mgmt-data Virtual IP (VIP). This dedicated mgmt-data VIP facilitates API access, e.g., Kubernetes CSI/COSI and Audit Catalog.

NVMe-oF is the main back-end connectivity between EBoxes and forms the backbone of the VAST DASE architecture. This technology allows VAST EBoxes to access storage devices in other EBoxes across a high-speed network, as if they were locally connected in the same physical server. By leveraging Remote Direct Memory Access (RDMA), NVMe-oF provides the essential high-speed data transfer capabilities required to meet the demands of data-intensive applications. NVMe-oF is integral to ensure that each EBox can efficiently interact with every SSD across all EBoxes in the deployment, regardless of physical location within the data center. This setup supports the disaggregated storage architecture of the VAST AI OS, where data processing and capacity resources are scaled and managed independently.

VAST AI Operating System: Software-Defined

The VAST AI OS is a unified platform for data management, AI orchestration, and intelligent automation. It combines storage, compute, and agentic runtime capabilities built on robust physical components of the VAST DASE architecture. Utilizing logical CNodes containers to encapsulate operational logic, each DNode container processes I/O requests efficiently, ensuring high performance and streamlined data handling. The platform's unique use of wide erasure code stripes that span multiple EBoxes and similarity-based data reduction across the cluster significantly reduces storage overhead and total cost of ownership, making it a cost-effective solution for large-scale data environments.

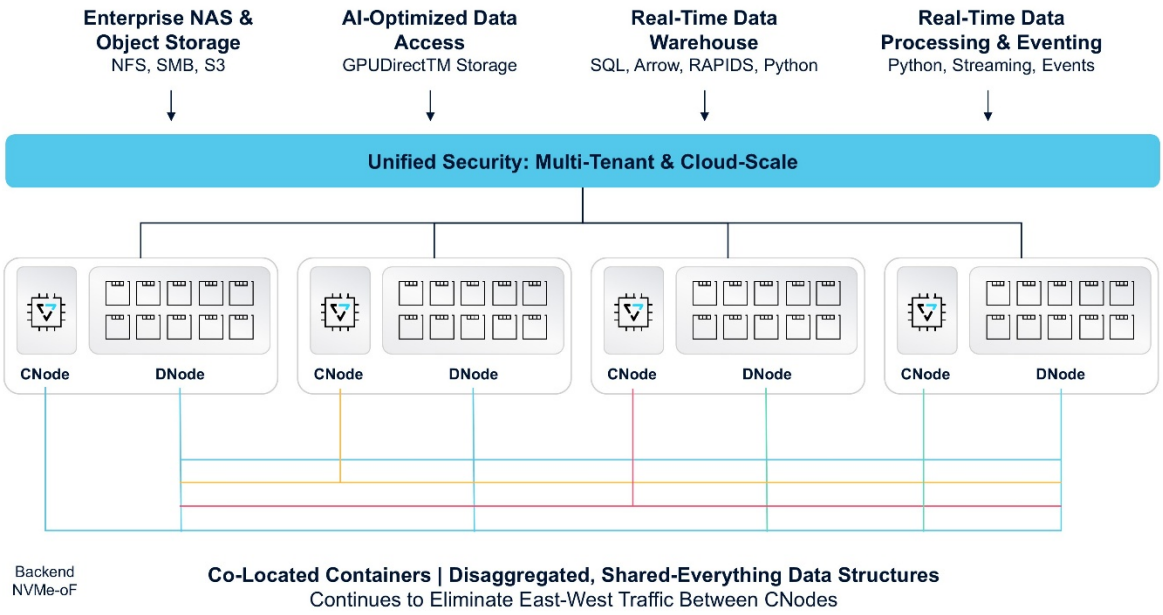

Figure 1. VAST's DASE architecture disaggregates CPUs and connects them to a globally accessible pool of SCM and ultra-dense NVMe flash SSD

Communication within the platform is facilitated through a semantic layer on high-speed SCM drives, allowing CNodes to operate without cross-talk and ensuring data integrity. This setup enhances performance and supports simultaneous access to data via multiple protocols, including NFS, SMB, S3, NFSoRDMA, and NVIDIA Magnum IO GPUDirect™ Storage (GDS). The platform's robust security framework adheres to zero-trust principles, incorporating per-tenant encryption and comprehensive access controls, including QoS, ABAC, RBAC, and audit logging.

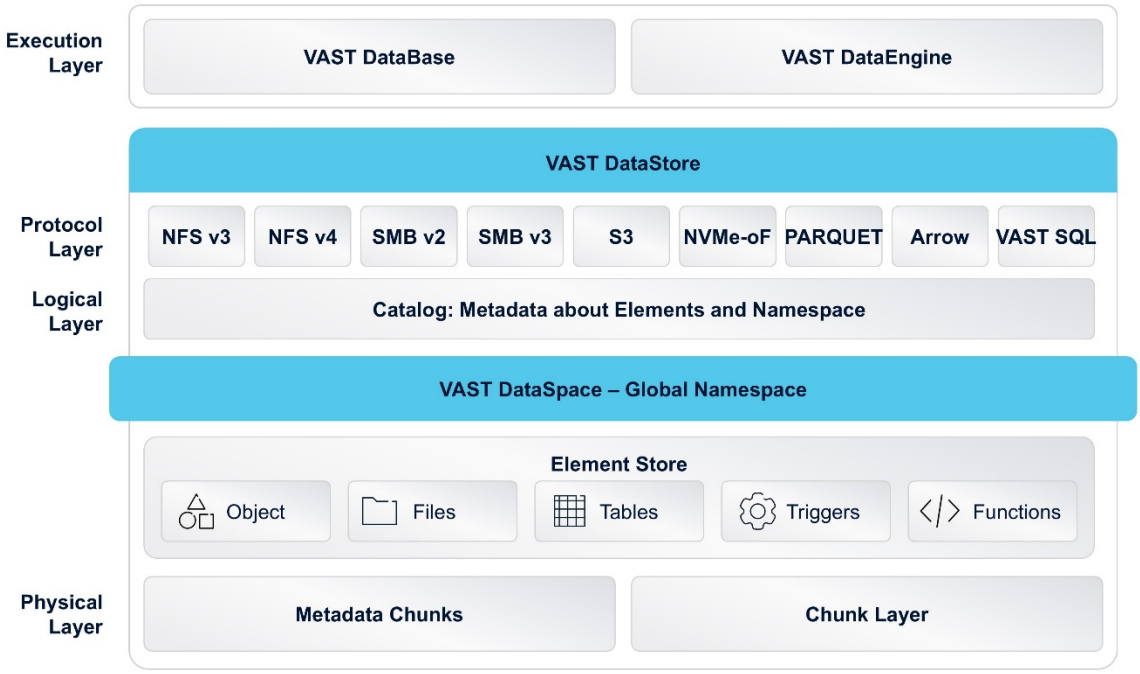

Figure 2. VAST AI OS is a single namespace solution that provides multi-protocol file and object storage (DataStore), high-performance data warehousing (DataBase), and multi-site global data access (DataSpace)

The VAST Data AI OS provides a secure, scalable, and economically efficient foundation for data storage. While there is a multitude of capabilities within the platform worth discussing, there are three critical components that will be discussed in further detail in this section: VAST DataStore, VAST DataBase, VAST DataSpace, and VAST DataEngine. These components collectively expand an NCP's offerings and platform capabilities for managing and analyzing unstructured and structured data sets across AI and data pipelines.



VAST DataStore

The VAST DataStore is a scalable, HPS solution with support for multiple file (NFS, SMB) and object (S3) interfaces. This multi-protocol support and interchangeability enable simultaneous access to data for reading and writing from either NFS or S3 storage protocols. For peak performance during intensive tasks like AI model training, the DataStore facilitates RDMA access over NFS (v3/v4) and employs NFS multipathing, enabling terabyte-per-second throughput from a single NFS mount point.

Notably, VAST collaborated with NVIDIA to integrate GDS through standard NFS-over-RDMA, marking a significant technological advancement. With its native implementation of an S3-compatible object store, the VAST DataStore can also simultaneously be leveraged for data collection and ingest as part of a broader data lake offering.

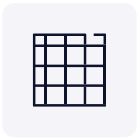

VAST DataBase

The VAST DataBase is an exabyte-scale data warehouse capable of delivering robust performance for both transactional and analytical workloads. It excels in environments with intensive data preprocessing needs by accelerating data access for frameworks such as Apache Spark, Trino, and NVIDIA Data Science (via RAPIDS-enabled applications). Featuring advanced predicate pushdowns (filtering), the VAST DataBase optimizes query processing and retrieval times, thus streamlining ETL/ELT stages and shortening AI pipelines end-to-end. It also powers the VAST Catalog, which automatically indexes metadata for every file and object within the DataStore, enhancing capabilities for audit, reporting, and data provenance essential for AI regulatory compliance.



VAST DataSpace

VAST DataSpace extends the platform's capabilities into a cohesive environment for global data management, offering a unified namespace across on-prem and multiple regions across NCP data centers. It ensures consistent and reliable data access with policies that optimize data placement and synchronization while maintaining high performance.

DataSpace supports multiple consistency modes for global read/write access and point-in-time data snapshots, facilitating complex data operations and recovery scenarios. The VAST DataSpace not only enhances data availability and disaster recovery but also simplifies management across diverse and geographically dispersed infrastructure, essential for large-scale, data-intensive applications.



VAST DataEngine

VAST DataEngine makes it possible to create and execute event-driven workflows using triggers, serverless functions, and containerized compute engines directly on VAST—powering real-time, governed pipelines without operational complexity. It includes a compute and execution layer which includes InsightEngine and AgentEngine and SyncEngine to see your entire data estate and provide scalable, high-performance data migration.

InsightEngine is a framework that accelerates the development of real-time AI data pipelines by eliminating integration steps and data movement while enabling automated, governed AI pipelines.

AgentEngine leverages DataEngine capabilities to create, manage, and orchestrate AI agents. It executes application runtime operations and manages container runtime, orchestration, scheduling, and observability.

SyncEngine is a data router that performs as a discovery and migration tool for unstructured data, built to connect your enterprise data from disparate sources to fuel AI pipelines. Included in the VAST AI OS, it combines two critical functions:

1. A universal data catalog to see your entire data estate
2. A scalable, high-performance migration engine to move data easily

VAST Management Service (VMS)

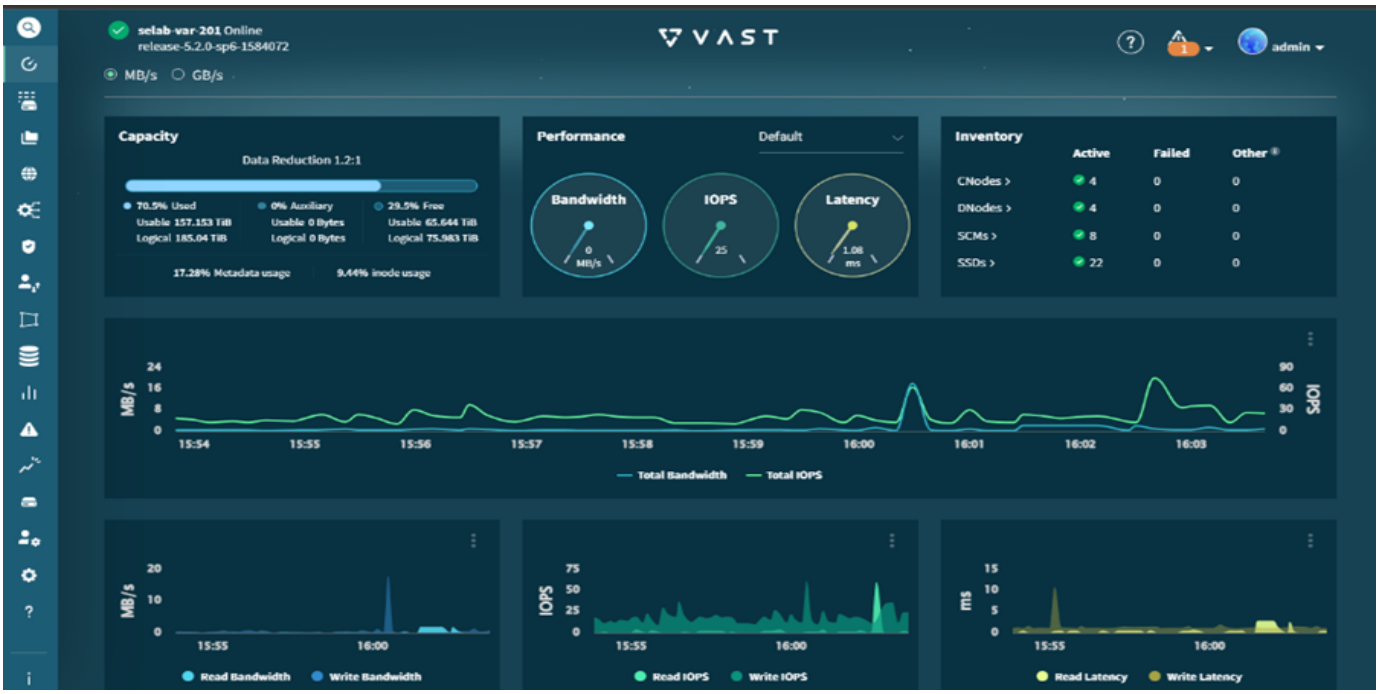

The VAST Management System (VMS) provides AI platform teams with deep, continuous observability into the storage and data layers that feed training, fine-tuning, and inference pipelines. Using VMS, MLOps teams can provision and monitor multi-protocol access, such as NFS and S3 to the same underlying data, allowing ML engineers to use S3 for data loading while simultaneously leveraging high-performance file access for other parts of your AI workflow components.

In real time, the User Activity view provides exactly which jobs, users, or clients are generating load against the data platform — not just the top noisy tenants, but any consumer. In practice, that lets teams pinpoint why a training job is underperforming or why a fine-tune run stalled (for example, a bulk feature extraction or an unexpected data export saturating I/O) and intervene before GPU hours are wasted. The VMS Data Catalog gives powerful, SQL-based query capabilities to instantly find, manage,

and version massive training datasets, dramatically accelerating the data preparation phase. It also enables AI teams to monitor in real-time IO operations for AI workloads that include checkpoint operations and model loading and unloading. For longer-term AI operations, VMS becomes a performance flight recorder. The VAST Analytics dashboard tracks detailed metrics — latency, IOPS, throughput, capacity, and even component-level health — and lets teams build custom dashboards to correlate storage behavior with ML events like checkpoint bursts, evaluation sweeps, or inference surges.

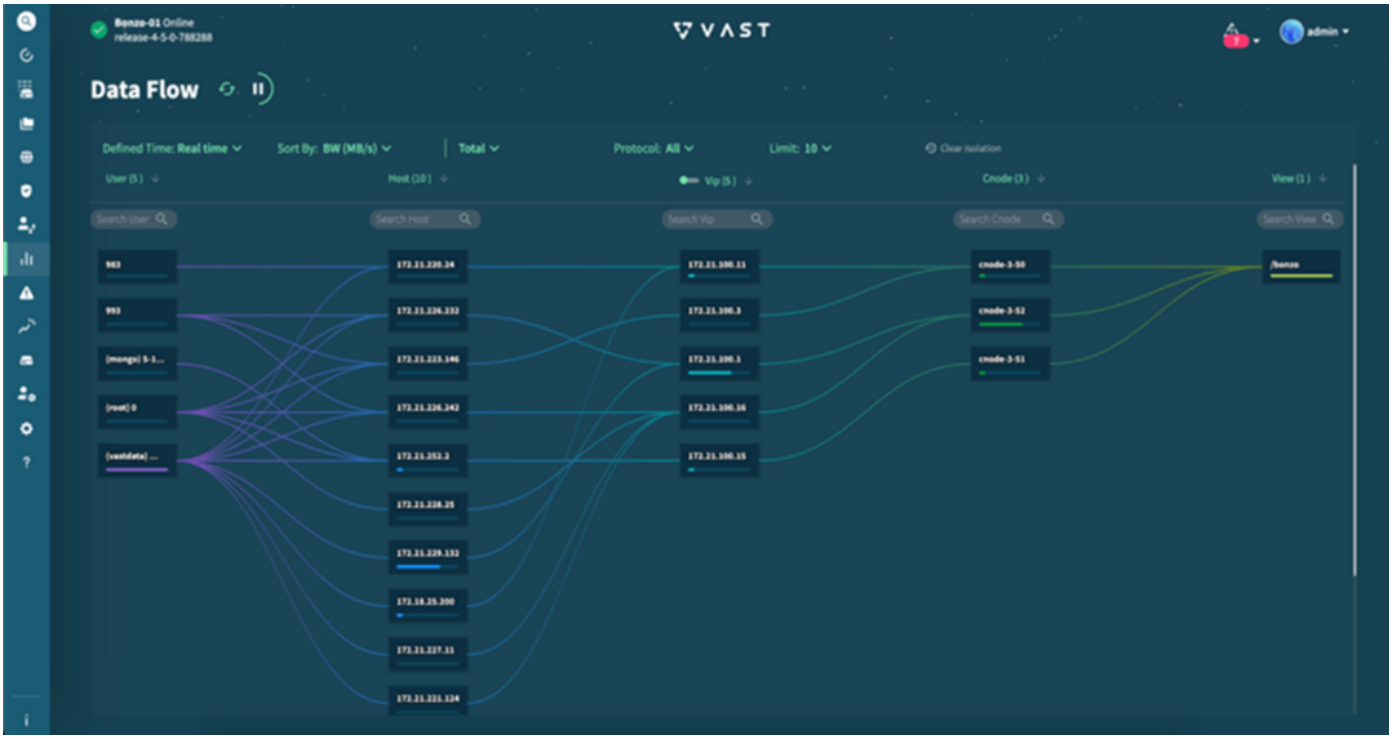

During model fine-tuning and RLHF (Reinforcement learning through Human Feedback) -style adaptation, VMS treats model artifacts — including incremental LoRA (Low Rank Adoption) weights and intermediate checkpoints — as first-class, governed assets. Checkpoints can be captured, snapshotted, and promoted without being scattered across individual servers or lost on someone's scratch volume, so AI teams get controlled lineage: which data produced which tuned model. That same policy-driven management (access control, encryption, retention) is critical when fine-tuning on regulated or customer data.

VAST Management System (VMS) gives AI and ML teams a single, unified data control plane that spans training, fine-tuning, and inference.

Multi-Tenancy

For NVIDIA Cloud Partners, VAST provides a secure, scalable, and multi-tenant architecture to integrate VAST data clusters. Tenants represent isolated administrative and data domains within the VAST cluster. Each tenant maintains a distinct directory hierarchy and can integrate with identity providers such as LDAP or Active Directory, or use local accounts. It also includes unique encryption keys, quotas, quality of service (QoS) settings, and the flexibility to define firewall network rules per tenant.

Multi-tenant isolation can be achieved using two separate approaches.

- Shared VIP pool + source-IP mapping (recommended at scale): One VIP pool per protocol serves all tenants, and access is gated by Tenant IP Range allowlists.
- Per-tenant VIP pools: Each tenant gets its own VIP pool(s), typically paired with VRF (Virtual Routing and Forwarding)/VLAN. This provides hard isolation and enables future support for overlapping client IP spaces.

Attribute	Shared VIP Pool + Source-IP Mapping	Per-Tenant VIP Pools
Isolation Type	Logical (IP allowlists)	Network (VRF/VLAN)
Client IP Requirements	Non-overlapping (or NAT)	Overlaps supported
Scalability	Very high, minimal config	Lower, higher admin overhead
Operational Complexity	Low	High
Security / Compliance	Moderate (ACL-based)	Strong (hard network isolation)
DNS Management	Centralized	Per-tenant entries
Ideal For	Large-scale shared platforms	Regulated / high-security tenants

More information is available in Appendix B

For AI workloads, clients should be configured with a multi-path NFS driver using NFS-over-RDMA and spreading traffic across both local ports and remote ports. Additionally, Priority Flow Control (PFC) using DSCP, as well as ECN, is implemented in the converged network fabric to enable lossless behavior and, if using VXLAN, propagating DSCP from the inner header to the outer header.

NVIDIA Recommended High Performance Storage (HPS) Sizing for NCP

To properly size the HPS system, it is essential to balance the performance requirements for read and write bandwidth according to the demands of training, fine-tuning, and inference operations being utilized.

For clusters larger than 1,024 GPUs, one additional EBox is added to the cluster to host the VAST AI OS Management Service (VMS). The VMS is responsible for all system and cluster management operations across a VAST Cluster. It is a highly available service that runs as a Linux container within a VAST Cluster. The capabilities of the VMS are explained in the VAST Management Service (VMS) section.

While the VMS is hosted on an EBox it does not participate in serving data to clients. In the event of a VMS hosting EBox failing, the VMS container is transitioned to another EBox within the cluster.

The standard storage performance guidance for seven common compute scale points, or scalable units (SUs), is shown in Table 1 below.

Description	Number of GPUs						
	1,024	2,304	4,096	8,192	16,384	29,952	41,472
Read throughput (GB/s)	160	360	640	1,280	2,560	4,680	6,480
Write throughput (GB/s)	80	180	320	640	1,280	2,340	3,240
Storage Configuration							
EBox	10	22	39	77	153	279	386
VAST Management Node (VMS)	0	1	1	1	1	1	1
Number of Namespaces	1	1	1	1	1	1	1
Number of 400G storage ports	10	23	40	78	154	280	387
Number of in-band management connections	0	0	0	0	0	0	0
Number of out-of-band management connections	10	23	40	78	154	280	387
Number of rack units	10	23	40	78	154	280	387
Power (KW)	7.7	17.71	28.6	49.4	98.15	182	251.55
Cooling (BTU/hr)	24,395.8	51,009	97587.25	167194.9	33,4901.7	62,1009.7	85,8324.2

Table 1. Guidance for standard HPS aggregate storage performance

VAST Data Design for NCP Deployments

Recommended Storage Networking

The VAST Data solution adheres to the NVIDIA Converged Network Design and connects to the NVIDIA Spectrum SN5600/5610 leaf switches. As part of the NCP RD, SN5600/5610 leaf switches are included in the bill of materials (BOM) articulated by NVIDIA.

When deploying the VAST AI OS, each VAST EBox has a NVIDIA ConnectX®-7 (CX-7) network adapter that connects to one 400 GbE ports on the SN5600/5610 leaf switches. Table 1 describes the number of 400 GbE ports required on the SN5600/5610 leaf switches to support the VAST HPS solution at a standard performance level.

There are multiple copper and optical breakout cable options from NVIDIA for connecting SN5600/5610 leaf switches to the EBox CX-7 network adapters. As a default, VAST recommends MMA4Z00-NS, OSFP finned-top twin-port 2x400G SR8 multimode transceivers in the SN5600/5610 connecting to a MMA1Z00-NS400 QSFP112 flat-top SR4 transceiver in the CX-7 using 2, split MPO-12/APC to MPO-12/APC cables. This results in one EBox consuming one port on the SN5600/5610. In addition, for distances less than five meters, NVIDIA offers passive and active copper cable options.

VLAN Configuration Setup

For each VAST deployment, the following VLANs must be configured on the HPS switching infrastructure prior to installation:

- Internal VAST (EBox-to-EBox) communication (default: VLAN 69)
- Discovery and expansion services (default: VLAN 10)
- VIP pool allocation for client access (default: VLAN 100)

Please see Appendix C for sample switch configuration and verification commands.

For AI workloads, clients should be configured with a multi-path NFS driver using NFS-over-RDMA and spreading traffic across both local ports and remote ports. Additionally, Priority Flow Control (PFC) using DSCP, as well as ECN, is implemented in the converged network fabric to enable lossless behavior and, if using VXLAN, propagating DSCP from the inner header to the outer header.

NCP Deployments with 1024 GPUs

Figure 3 illustrates the logical networking view of VAST Data's reference design for NCP deployments with 1,024 GPUs (1 SU), 10 EBoxes. EBoxes directly connect to the SN5600/5610 converged network switches using 400 GbE links. The 200 GbE interface from EBoxes is used for both user and data traffic operating in shared mode.

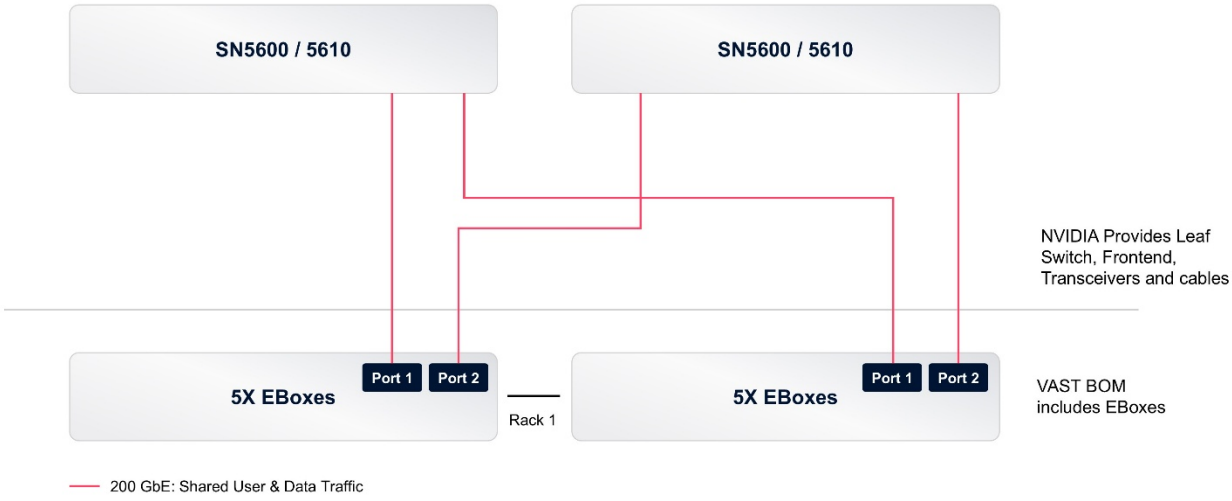

Figure 3. Logical view for VAST Networking Configuration for 1024 GPUs with 10 EBoxes

The table below provides the configuration for 1,024 GPUs

GPU Count	EBoxes	Number of 200GbE ports required	Number of 400GbE ports used
1,024	10	20	10

Table 2: VAST Sizing for 1,024 GPUs with Network ports utilization

As summarized in Table 2 for the 1,024 GPU sizing, VAST storage configuration utilizes 10 400GbE ports in the SN5600/5610 leaf switches, and these are split into 20x 200GbE connections.

VAST recommends connecting each EBox interface to different ports on the Nvidia SN5600/5610 HPS leaf switches to provide network redundancy, as shown below in Figure 4, physical view diagram. This implementation view translates the logical plan into a concrete deployment diagram with actual switches SN5600/SN5610 and EBox stacks. It also shows how an AI-optimized infrastructure scales to support large GPU clusters by combining NVIDIA's high-performance network fabric with VAST's storage fabric. At the top, the SN5600/5610 switches form the NVIDIA Ethernet Storage Fabric Network, the high-bandwidth path that connects to the VAST Data EBoxes Fabric. In this configuration, 10 EBoxes support 1,024 GPUs.

NCP Deployments with 1,024 GPUs:
10 EBoxes

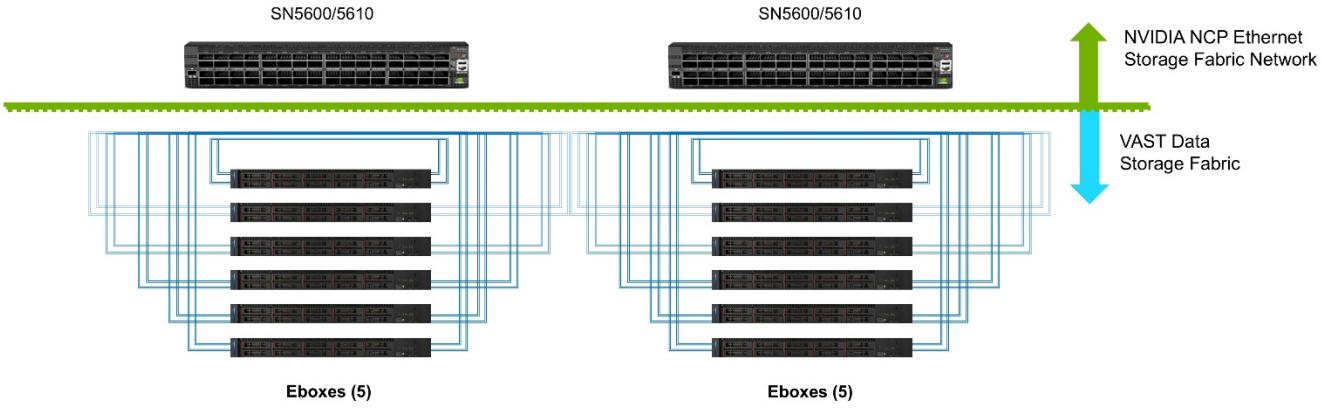

The storage network is logically split internally, with the client data access outbound in green and internal communications between the EBoxes remaining in the VAST Data storage fabric.

While VAST Data configurations can scale linearly to meet the performance needs of any workload, the HPS storage requirements are primarily driven by the size of the model being trained and the dataset associated with it. The next section discusses the scalability of these designs to meet 16K and 40K GPU configuration. For NCPs targeting LLM training, fine-tuning, and Inference workloads, the standard high-performance storage configuration as listed in sizing table 1 is ideal, as it exceeds the necessary performance targets for maximum GPU utilization.

NCP Deployments with 16,384 GPUs

Figure 5 below shows the logical view for VAST Data's reference design for NCP deployments with 16,384 GPUs using 154 EBoxes. VAST EBoxes directly connect to the converged SN5600/5610 network switches dedicated to high-performance storage traffic. The VAST configuration utilizes 154 400GbE ports from the storage network fabric switches. The 200 GbE interface from EBoxes is used for both user and data traffic operating in shared mode.

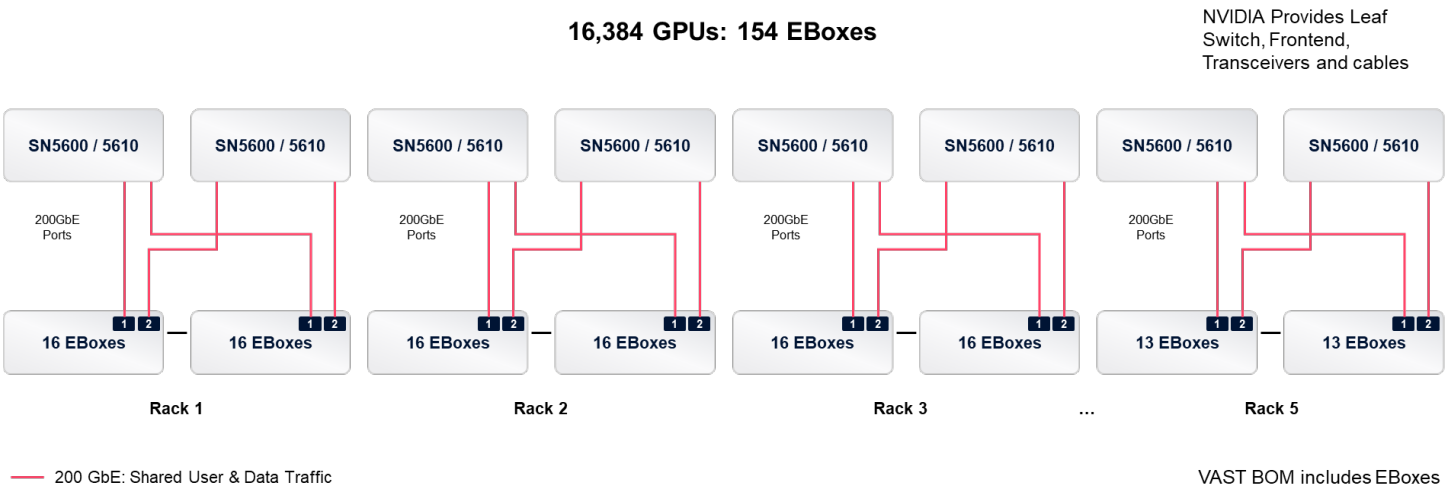

Figure 5. Logical view for VAST Networking Configuration for 16,384 GPUs with 154 EBoxes

NVIDIA supplies the SN5600/5610 leaf switches, front-end transceivers, and cables. The VAST BOM includes the EBoxes for a high-performance storage appliance.

The table below provides the configuration for 16,384 GPUs

GPU Count	EBoxes	Number of 200GbE ports required	Number of 400GbE ports used
16,384	154	308	154

Table 3: VAST Sizing for 16,384 GPUs with Network ports utilization

As summarized in Table 3 for 16,384 GPU sizing, VAST storage configuration utilizes 154 400GbE ports in the SN5600/5610 leaf switches, and these are split into 308x 200GbE connections.

VAST recommends connecting each EBox interface to different ports on the NVIDIA SN5600/5610 HPS leaf switches to provide network redundancy, as shown below in Figure 5, physical view. This implementation view translates the logical plan into a concrete deployment diagram with actual switches SN5600/SN561 and EBoxes. This diagram also shows how an AI-optimized infrastructure scales to support large GPU clusters of 16,384 GPUs by combining NVIDIA's high-performance network fabric

with VAST's storage fabric using 154 EBoxes. At the top, the SN5600/5610 switches form the NVIDIA Ethernet Storage Fabric Network, the high-bandwidth path that connects to the VAST Data EBoxes Fabric.

The following diagram outlines the networking design for the recommended VAST configuration:

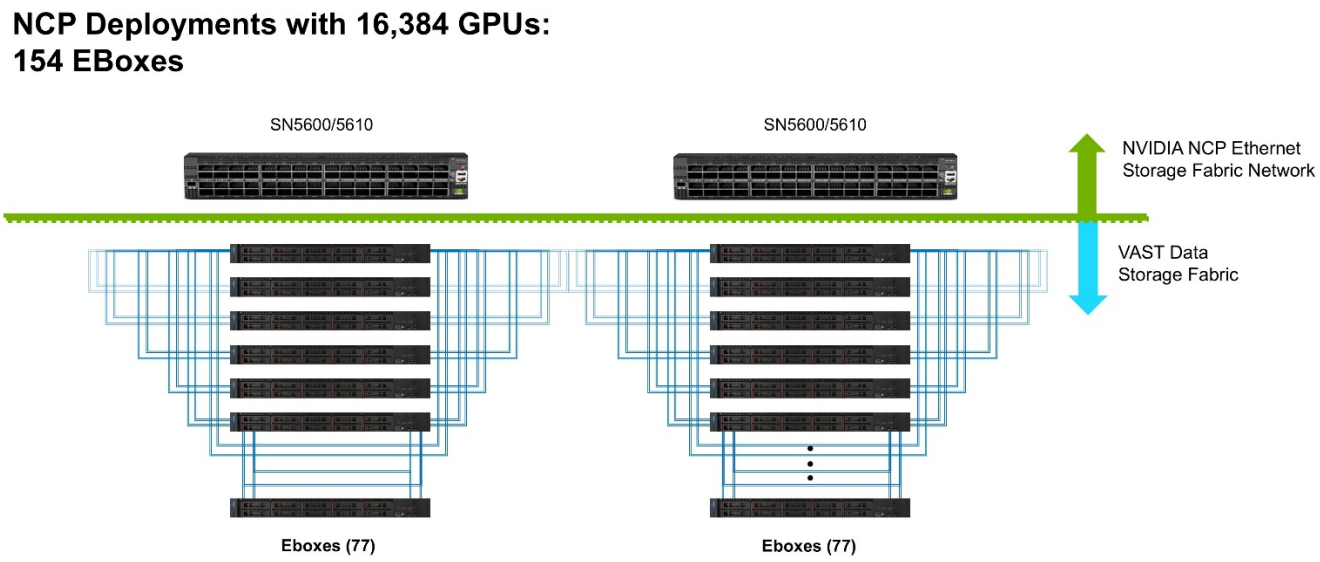

Figure 6. Physical or Implementation view for VAST Networking Configuration for 16,384 GPUs with 154 EBoxes

NCP Deployments with 41,472 GPUs

Figure 6 below shows a logical network design for NCP deployments with 41,472 GPUs using 387 EBoxes. VAST EBoxes directly connect to the SN5600/5610 storage network fabric, utilizing 387 400GbE ports. The 200 GbE interface from EBoxes is used for both user and data traffic operating in shared mode.

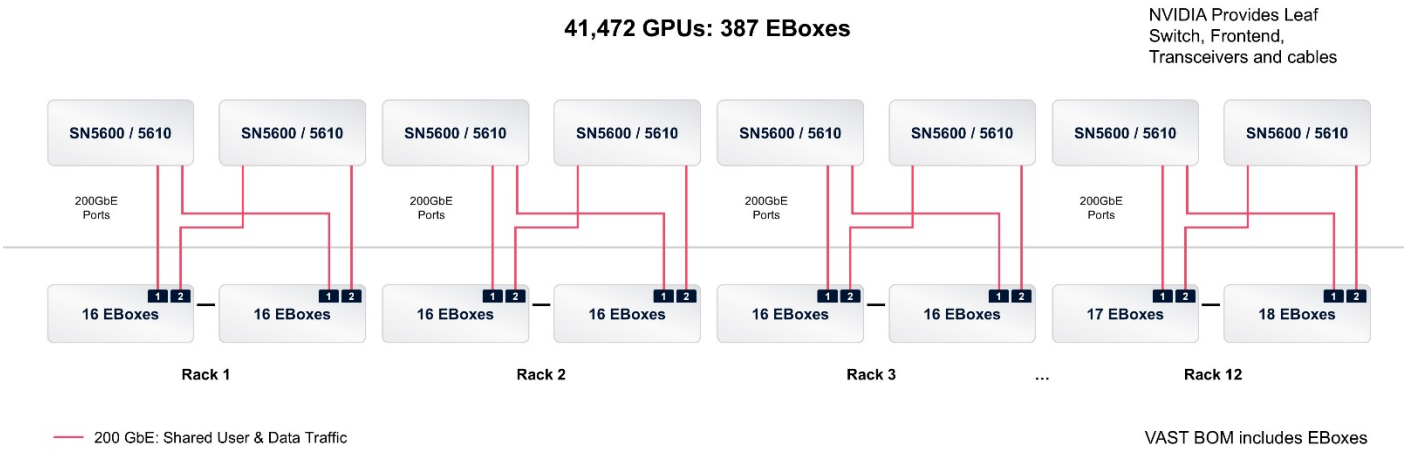

Figure 7. Logical View for VAST Networking Configuration for 41,472 GPUs with 387 EBoxes

The table below provides the configuration for 41,472 GPUs

GPU Count	EBoxes	Number of 200GbE ports required	Number of 400GbE ports used per Leaf switch
41,472	387	774	387

Table 4: VAST Sizing for 41,472 GPUs with Network ports utilization

As summarized in Table 4 for the 41,472 GPU sizing, VAST storage configuration utilizes 387 400GbE ports in the SN5600/5610 leaf switches, and these are split into 774x 200GbE connections.

VAST recommends connecting each EBox interface to different ports on the NVIDIA SN5600/5610 HPS leaf switches to provide network redundancy, as shown below in Figure 8, physical view. This implementation view translates the logical plan into a concrete deployment diagram with the actual switches, SN5600/SN561, and EBoxes. This diagram also shows how an AI-optimized infrastructure scales to support large GPU clusters of 41,472 GPUs by combining NVIDIA's high-performance network fabric with VAST's storage fabric using 387 EBoxes. At the top, the SN5600/5610 switches form the NVIDIA Ethernet Storage Fabric Network, the high-bandwidth path that connects to the VAST Data EBoxes Fabric.

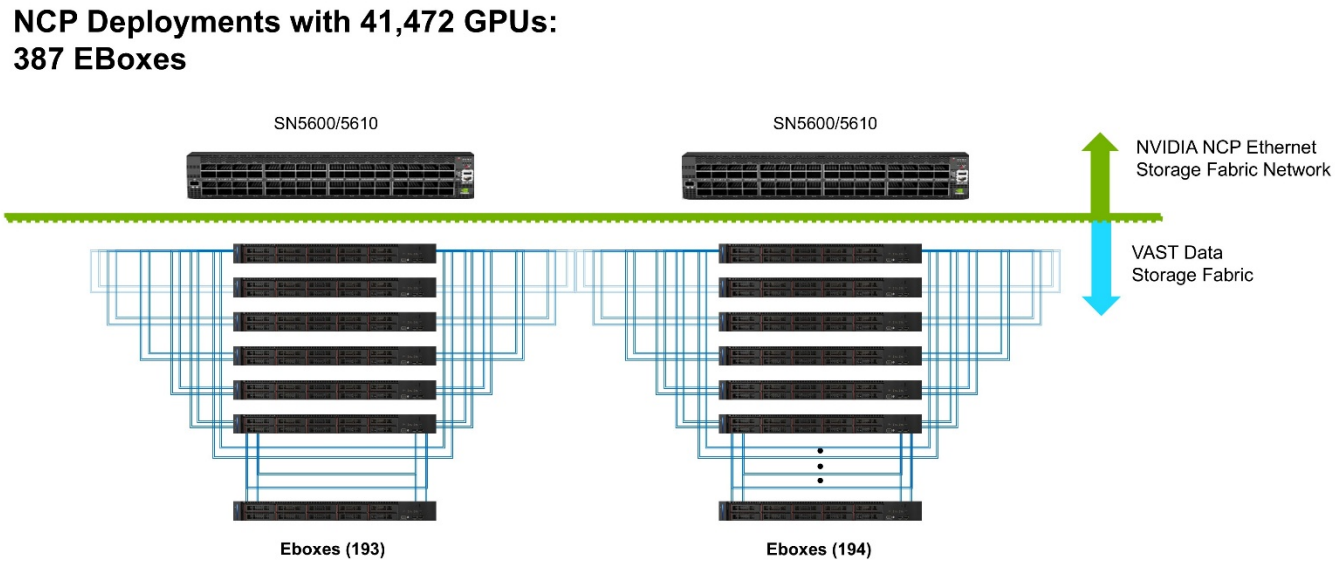

Figure 8. Physical Implementation view for VAST Networking Configuration for 41,472 GPUs with 387 EBoxes

VAST Data Solution Performance Validation

The solution outlined in this reference design document was validated by NVIDIA on a **VAST Data v5.2** cluster. A benchmarking test suite from NVIDIA was utilized to characterize overall cluster performance typical of long-running AI model training jobs across various compute scale points. All test results were submitted to and validated by NVIDIA.

> **For a more in-depth look on checkpointing, please see:**
> https://www.vastdata.com/blog/a-checkpoint-on-checkpoints-in-llms

In addition to the benchmark results, VAST has provided NVIDIA real-world production examples—hosted at NCPs— showing load profiles of various AI model training jobs running across six thousand NVIDIA GPUs in parallel, training one trillion parameter models backed by VAST Data for HPS.

Lastly, a variety of scaling, resilience, and QoS tests were performed. It was demonstrated that the system could linearly scale based on the number of GPUs (i.e., from 1K to 2K) without rebuilding or reinstalling, and that cluster performance was not negatively impacted while expanding. The VAST cluster could also tolerate loss of network interfaces, CNodes, and DBox with only a proportional impact, no greater. Finally, QoS levels could be enforced while read and write operations were ongoing. These results were also reviewed and validated by NVIDIA as exceeding expectations.

Additional HPS Features and Considerations for NCP

The NCP reference design is inherently multitenant with secure network isolation on the fabric. HPS designs must support the following:

- Scale throughput, capacity, and tenant support when required
- Highly resilient to multiple failures that may occur and be highly available to ensure data access continues when unexpected failures occur, with minimal performance impact during any failure.
- Secure and isolated data paths with quality-of-service settings for each tenant
- Storage management must also be secure to allow multiple tenant administrators to manage and access the assigned isolated tenant data space
- Software and Firmware upgrades operations are done without loss of data access with minimal performance impact

Some cloud partners may also require HPS support for the following features:

- Different data APIs, such as S3 and others
- Encryption for data security with key management

- Data Reduction services, such as compression or duplication
- Data management features for backup and migration of data in and out of the HPS system

The VAST AI OS meets and exceeds both mandatory and recommended features listed above in full compliance with all of the NCP design guides requirements.

Summary

The VAST AI OS pioneered an architecture that delivers not only performance and scale but also operational efficiency and ease of use. With flexible configurations to support over 16K GPUs in a single cluster, NCPs can confidently pair VAST Data solutions with any scale of large AI infrastructure deployments. Based on the NCP RD, featuring HGX Servers, this jointly validated VAST Data reference design meets and exceeds the intensive demands of AI, high-performance computing, and other data-intensive applications.

This document details a validated reference design from VAST Data that specifically adheres to the NCP RD specifications for HPS. The VAST AI OS also provides additional data services—such as multi-protocol data ingest, accelerated ELT for feature engineering—useful in helping CSPs accelerate data pipelines for their customers, all underpinned by a multi-tenant Zero Trust security model. This holistic approach not only streamlines AI pipelines by consolidating storage across various data stages but also significantly minimizes complexity and costs, thereby boosting overall system performance and security for CSPs.



For more information on Universal Storage and how it can help you solve your application problems, reach out to us at hello@vastdata.com.

Appendix A: Consolidate Storage Silos, Accelerate Insight

Because the VAST AI OS supports file, object, and tabular data formats, organizations can consolidate high-performance storage (training data), data management systems (feature stores, model artifacts), and data lake (data collection and preprocessing) storage requirements onto a single platform. By eliminating data transfers across separate storage systems, AI training, inference, and model serving processes are simplified, enhancing speed and reliability. This convergence of data formats and workflow stages streamlines the entire AI data lifecycle, allowing organizations to focus on innovation rather than managing complex storage architectures and cumbersome data copy management tasks.

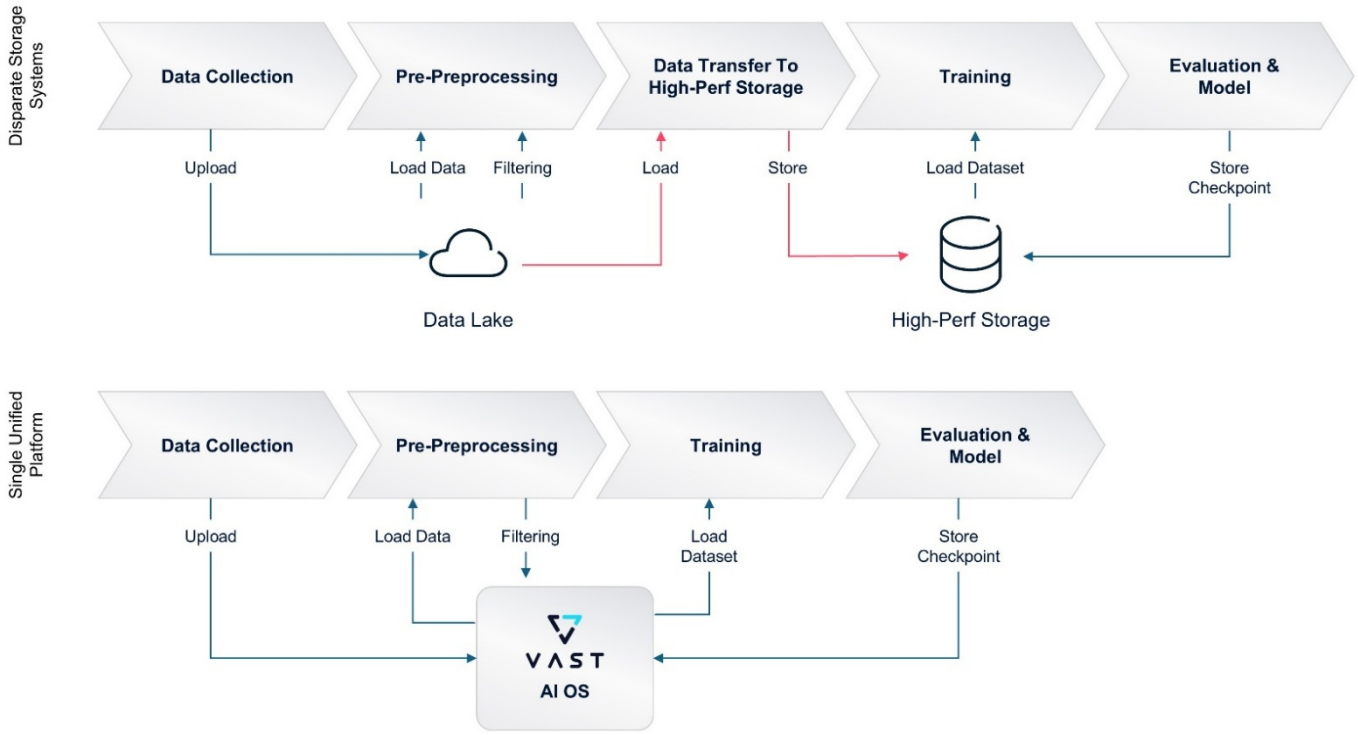

Figure 5. (Top) Typical workflow of an AI pipeline and corresponding storage locations requiring data transfer between storage systems. (Bottom) Simplified AI workflow with the VAST AI OS. Data is accessed in place and is available over both file and object protocols

By consolidating the various storage systems into a single, unified platform, VAST enables NVIDIA Cloud Partners to offer a highly differentiated solution that accelerates the entire AI pipeline. This empowers NCPs to create and deliver services to their customers faster and with lower operational overhead than traditional multi-cluster, tiered storage solutions. Additionally, NVIDIA Cloud Partners can improve the security posture of their service offerings by reducing the attack surface and eliminating the need to manage complex permissions across multiple systems.

Appendix B: The Imperative for Secure Multi-Tenant Storage

When deploying the VAST AI OS in an NCP RD environment, especially if constructing a technology stack around NVIDIA HBX servers, selecting the right data services layer is essential. This layer must not only provide the high performance needed to power NVIDIA GPUs but also incorporate multi-tenant storage capabilities to enhance the service's efficiency, security, and scalability. By implementing full-stack multi-tenant, including at the data services level, cloud providers can significantly reduce both capital expenditures (CAPEX) and operational expenditures (OPEX), leading to reduced administrative costs and increased profit margins.

Zero Trust: The VAST AI OS meets the Data Pillar Framework

Encryption-in-Flight SQL over HTTPS, NFS4 w/ KRB5P, NFS w/ TLS, S3 HTTPS	**Secure Multi-Tenancy** Cloud-Scale Logical Data Segmentation
Encryption-at-Rest FIPS 140-2 Validated, 140-3 Capable	**VLAN-Based Segmentation** Create and map VAST Views to specific VLANs/Pools
External Key Management Thales CypherTrust 2.9+, Fortanix DSM, Hashicorp Vault Enterprise	**Ransomware-Proof Snapshots** Indestructible Snapshots to avoid snap policy tampering
Access Audit User Access (NFS, SMB, S3), Admin Audit	**STIG** STIG profile MAC 1 - Mission Critical Classified
RBAC Active Directory, Openldap, FreeIPA	**Maintained OS** CIQ Rocky 8
Syslog Support Rsyslog w/ Elastic, Splunk & Qradar	**Index, Search, Catalog** VAST Catalog

VAST AI OS delivers a powerful foundation for secure multitenancy, enabling organizations to isolate tenant data while maintaining efficient, high-performance access. Unlike traditional architectures that rely on complex sharded systems, VAST decouples storage and compute, providing a shared data infrastructure without sacrificing the strict segregation required for tenant security. This model allows administrators to easily implement universal security controls across all tenants, without increasing operational complexity as the number of customers grows. Each tenant is fully isolated, ensuring that no unauthorized data access occurs across domains, critical for compliance with Zero Trust principles like those outlined in NIST 800-207.

VAST Data's Secure Multi-Tenant Data Platform

The VAST Data AI OS is the industry's only secure multi-tenant data platform built for large-scale AI clouds. It delivers a comprehensive solution for NVIDIA Cloud Partners seeking to implement a zero-trust, multi-tenant infrastructure. The platform integrates multi-category security (MCS) and a hybrid model of Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC), providing granular access controls and ensuring strict data isolation among tenants.

The table below highlights Zero Trust Data pillar implementation with VAST AI OS.

Data Pillar	VAST AI OS Features	Description
Authorization	ABAC	Attribute Based Authentication in-conjunction with Role Based
Categorize	Catalog	Index, search, and query data in realtime all of the information stored on the VAST DataStore
Log	Audit Logging	Console Access, Admin, and Protocol Access (S3, NFS, SMB)
Tag	Tagging	Tag any object/file with a unique tag for categorization or authorization
Respond	Event Forwarding	Integrate with Syslog, Splunk, ArcSight, Qradar, Elastic, Sumologic, etc
Encrypt	At Rest & In-Flight	FIPS 140-2 Validated & FIPS 140-3 Capable
Automate	Automation	K8S CSI, Manilla, Cinder Plugins for OpenStack, Rest API, Python SDK
Monitoring & Observability	Telemetry	System Capacity, Performance, Health and Data Flow

Advanced Tenancy Capabilities

The VAST AI OS offers a range of advanced tenancy capabilities that enable cloud providers to deliver secure, isolated, and dynamically provisioned AIaaS offerings:

- **Dynamic Provisioning:** Supports dynamic provisioning of storage resources, allowing cloud providers to allocate and deallocate storage capacity on demand based on tenant requirements
- **Network Segmentation**: Provides robust network segmentation capabilities, creating isolated virtual networks for each tenant to ensure that tenants' data traffic remains separate and secure
- **Micro-Segmentation:** Supports micro-segmentation, enabling fine-grained control over access to specific data sets and resources within a tenant's environment
- **Data Separation via View and Tenant-Level Encryption:** Offers data separation at both the view and tenant levels, with tenant-level encryption ensuring that each tenant's data is encrypted with unique keys, further enhancing data isolation and security
- **Automated Data Labeling:** Administrators can define labels/attributes at the "View" level for data classification and authorization
- **Granular Access with Hybrid RBAC-ABAC:** Combines RBAC and ABAC for finer-grained access control, supported with extensive audit capabilities and unique encryption keys for data isolation
- **Data Encryption and Key Management:** Uses external key management with KMIP for encrypting data at rest and in transit, providing secure key exchange and management
- **In-Flight Encryption:** Based on the TLS protocol 1.3, offering enhanced security for encrypted communications using FIPS-validated encryption ciphers
- **NVIDIA Magnum IO GPUDirect Storage Encryption:** Supports GDS with end-to-end encryption, enabling secure and high-performance data transfer between storage and GPUs
- **Object Storage Encryption:** Supports server-side encryption (SSE) and client-side encryption for enhanced data protection
- **Crypto Erase:** The ability to destroy encryption keys and purge data in a way that meets NIST 800.88 controls that ensures that a data set is purged

Benefits of a Zero-Trust Multi-Tenant AI OS

- **Enhanced Data Security and Privacy:** Ensures data isolation and protection, providing robust security mechanisms that keep each tenant's data private and secure, critical for complying with industry regulations such as GDPR and HIPAA.
- **Cost Efficiency and Resource Optimization:** Allows for the efficient use of storage resources by pooling them and dynamically allocating these resources among multiple tenants based on demand, reducing the need for dedicated storage systems and overprovisioning.
- **Scalability and Flexibility:** Offers scalable storage options that can easily expand to meet increasing demands, ensuring the AIaaS offering can support a wide range of applications, from small-scale experiments to large, enterprise-level AI deployments.
- **Simplified Management and Maintenance:** Managing a single, multi-tenant storage infrastructure is far simpler and more cost-effective than maintaining separate storage systems for each customer, reducing administrative overhead and the risk of human error.
- **Performance and Availability:** Designed for high-performance workloads, these solutions deliver the high throughput and low latency required for AI and ML tasks, with built-in redundancy and failover capabilities to ensure high availability and data protection.
- **Compliance and Auditability:** Provides the necessary tools and frameworks to ensure that data handling and processing meet regulatory standards, including features for data encryption, access controls, and audit trails, which are vital for demonstrating compliance with regulations and for conducting audits.

Cost Optimization and Performance Benefits

The VAST AI OS is designed to optimize costs and maximize performance for CSPs offering AI-as-a-Service (AIaaS). Key benefits include simplified AI infrastructure, HPC performance with enterprise simplicity, data insights & management, and flexible scalability.

Compliance and Auditability

The VAST AI OS provides the necessary tools and frameworks to ensure that data handling and processing meet regulatory standards, including features for data encryption, access controls, and audit trails, which are vital for demonstrating compliance with regulations and for conducting audits.

Appendix C: VLAN Configuration and Verification

For each VAST deployment, the VLANs 10, 69 & 100 are used for discovery/expansion, internal communication & client access, respectively.

Below is the step-by-step VLAN Configuration verification process for the SN5600/5610 HPS leaf switches. To ensure the VAST EBoxes can communicate with each other, the following commands can be used for verification:

```
nv set bridge domain br_default untagged 10
nv set bridge domain br_default vlan 10 vni 10
nv set bridge domain br_default vlan 69 vni 69
nv set bridge domain br_default vlan 100 vni 100
nv set interface swpX-XX bridge domain br_default
nv set evpn enable on
nv set interface vlan69 ip address XX.XX.XX.XX/XX
nv set interface vlan69 ip address XX.XX.XX.XX/XX
nv set interface vlan69 ip address XX.XX.XX.XX/XX
nv set interface vlan69 vlan 69
nv set interface vlan10 vlan 10
nv set interface vlan100 vlan 100
nv set interface vlan100 ip vrf STORAGE
nv set interface vlan100 ip address XX.XX.XX.XX/XX
nv set interface vlan100 ip vrr address XX.XX.XX.XX/XX
nv set nve vxlan arp-nd-suppress on
nv set nve vxlan decapsulation dscp action preserve
nv set nve vxlan enable on
nv set nve vxlan encapsulation dscp action copy
nv set nve vxlan source address XX.XX.XX.XX/XX
nv set qos roce enable on
nv set router bgp autonomous-system XXX
nv set router bgp enable on
nv set router bgp router-id XX.XX.XX.XX
nv set vrf STORAGE evpn enable on
nv set vrf STORAGE evpn vni 4001
nv set vrf STORAGE router bgp address-family ipv4-unicast enable on
nv set vrf STORAGE router bgp address-family ipv4-unicast redistribute connected enable on
nv set vrf STORAGE router bgp address-family ipv4-unicast route-export to-evpn enable on
nv set vrf STORAGE router bgp autonomous-system XX.XX.XX.XX
nv set vrf STORAGE router bgp enable on
nv set vrf STORAGE router bgp router-id XX.XX.XX.XX
nv set vrf default router bgp neighbor swpXX-XX peer-group LEAF-SPINE
nv set vrf default router bgp neighbor swpXX-XX type unnumbered
nv set vrf default router bgp peer-group LEAF-SPINE address-family ipv4-unicast enable on
nv set vrf default router bgp peer-group LEAF-SPINE address-family l2vpn-evpn enable on
nv set vrf default router bgp peer-group LEAF-SPINE bfd detect-multiplier 3
nv set vrf default router bgp peer-group LEAF-SPINE bfd enable on
nv set vrf default router bgp peer-group LEAF-SPINE bfd min-rx-interval 300
nv set vrf default router bgp peer-group LEAF-SPINE bfd min-tx-interval 300
nv set vrf default router bgp peer-group LEAF-SPINE description eBGP
nv set vrf default router bgp peer-group LEAF-SPINE graceful-shutdown on
nv set vrf default router bgp peer-group LEAF-SPINE remote-as external
```

To ensure the VAST EBoxes can communicate with each other, the following commands can be used for verification:

1. Connectivity validation (Leaf ↔ C-Node/D-Node): Ensure all cabled leaf-facing ports are in Forwarding state; port without equipment blocking.

```
cumulus@leaf-13-154:mgmt:~$ nv show bridge domain br_default brief
Bridge info:
    mac-address         : b0:cf:0e:46:df:fd
    type                : vlan-aware
    encap               : 802.1Q
    ageing              : 1800
    stp mode            : rstp
    vlan-vni-offset     : 0

Bridge Vlan Info :
untagged      tagged
------------- ------------------------------------------------------
10            100,69

Bridge Port Info:
Port           State
-------------- ---------------
swp1           forwarding
swp2           forwarding
swp3           blocking
<.....>
```

2. Underlay connectivity validation (Leaf ↔ Spine): Verify BGP neighbors are Established and PfxRcd/PfxSnt > 0. Any 0 values indicate a fabric/underlay fault.

```
leaf-13-154# show ip bgp ipv4 unicast summary
BGP router identifier 10.10.200.22, local AS number 65122 vrf-id 0
BGP table version 217
RIB entries 97, using 21 KiB of memory
Peers 7, using 140 KiB of memory
Peer groups 1, using 64 bytes of memory

Neighbor           V         AS   MsgRcvd   MsgSent   TblVer  InQ OutQ  Up/Down State/PfxRcd   PfxSnt Desc
peerlink.4094      4          0         0         0        0    0    0    never         Idle        0 N/A
spine-7-175(swp27) 4      65298 799817924 799861287        0    0    0 4d21h40m          288       49 N/A
spine-7-171(swp28) 4      65298 801264959 801314426        0    0    0 05w3d06h          288       49 N/A
```

3. Overlay connectivity Validation (Leaf ↔ Spine / EVPN): Ensure BGP L2VPN-EVPN is Established with PfxRcd/PfxSnt > 0 (Type-2/Type-5 routes). 0 indicates overlay failure.

```
leaf-13-154# show bgp l2vpn evpn summary
BGP router identifier 10.10.200.22, local AS number 65122 vrf-id 0
BGP table version 0
RIB entries 401, using 88 KiB of memory
Peers 7, using 140 KiB of memory
Peer groups 1, using 64 bytes of memory

Neighbor           V         AS   MsgRcvd   MsgSent   TblVer  InQ OutQ  Up/Down State/PfxRcd   PfxSnt Desc
peerlink.4094      4          0         0         0        0    0    0    never         Idle        0 N/A
spine-7-175(swp27) 4      65298 799916097 799959465        0    0    0 4d21h46m         7536     8199 N/A
spine-7-171(swp28) 4      65298 801363082 801412604        0    0    0 05w3d06h         7538     8199 N/A
```

4. MAC Address advertised check between Leaf (VTEP)<>Leaf (VTEP). We must see in the output [remote-vtep] IPs of other leafs, if no IPs exist, then the advertised MAC Address is not working.

```
cumulus@leaf-13-154:mgmt:~$ nv show evpn vni 69
                   operational   applied
-----------------  ------------  -------
route-advertise
  svi-ip           off
  default-gateway  off
[remote-vtep]      10.10.200.3
[remote-vtep]      10.10.200.4
[remote-vtep]      10.10.200.5
[remote-vtep]      10.10.200.6
<.......>
vlan               69
bridge-domain      br_default
tenant-vrf         default
vxlan-interface    vxlan48
mac-count          1346
host-count         1837
remote-vtep-count  37
local-vtep         10.10.200.22
```

Learn more about the VAST Data AI OS.