



NVIDIA DGX SuperPOD: Next Generation Scalable Infrastructure for AI Leadership Reference Architecture Featuring NVIDIA DGX H200

Release latest

NVIDIA Corporation

Apr 14, 2025

Reference Architecture

1	Abstract	3
2	Key Components of the DGX SuperPOD	5
2.1	NVIDIA DGX H200 System	5
2.2	NVIDIA InfiniBand Technology	6
2.3	Runtime and System Management	6
2.4	Components	6
2.5	Design Requirements	7
2.5.1	System Design	8
2.5.2	Compute Fabric	8
2.5.3	Storage Fabric	8
2.5.4	In-Band Management Network	8
2.5.5	Out-of-Band Management Network	9
2.5.6	Storage Requirements	9
2.5.6.1	High-Performance Storage	9
2.5.6.2	User Storage	9
3	DGX SuperPOD Architecture	11
4	Network Fabrics	15
4.1	Compute Fabric	16
4.2	Storage Fabric	16
4.3	In-Band Management Network	17
4.4	Out-of-Band Management Network	19
5	Storage Architecture	21
6	DGX SuperPOD Software	23
6.1	NVIDIA Base Command	23
6.2	NVIDIA NGC	23
6.3	NVIDIA AI Enterprise	24
7	Summary	27
8	Notices	29
8.1	Notice	29
8.2	Trademarks	30
8.3	Copyright	30

Document Number	RA-11336-001 V2
-----------------	-----------------

Publication Date	2025-04-14
------------------	------------

Tip

The NVIDIA DGX SuperPOD: Next Generation Scalable Infrastructure for AI Leadership Reference Architecture Featuring NVIDIA DGX H200 is also available as a PDF.

Chapter 1. Abstract

NVIDIA DGX SuperPOD™ with NVIDIA DGX™ H200 systems is the next generation of data center architecture for artificial intelligence (AI). Designed to provide the levels of computing performance required to solve advanced computational challenges in AI, high performance computing (HPC), and hybrid applications where the two are combined to improve prediction performance and time-to-solution. DGX SuperPOD is based upon the infrastructure built at NVIDIA for internal research purposes and is designed to solve the most challenging computational problems of today. Systems based on the DGX SuperPOD architecture have been deployed at customer data centers and cloud-service providers around the world.

To achieve the most scalability, DGX SuperPOD is powered by several key NVIDIA technologies, including:

- ▶ NVIDIA DGX H200 system—to provide the most powerful computational building block for AI and HPC.
- ▶ NVIDIA NDR (400 Gbps) InfiniBand—bringing the highest performance, lowest latency, and most scalable network interconnect.
- ▶ NVIDIA NVLink® technology—networking technologies that connect GPUs at the NVLink layer to provide unprecedented performance for most demanding communication patterns.



The DGX SuperPOD architecture integrates NVIDIA software solutions including NVIDIA Base Command™, NVIDIA AI Enterprise, CUDA, and NVIDIA Magnum IO™. These technologies help keep the system running at the highest levels of availability, performance, and with NVIDIA Enterprise Support (NVEX), keeps all components and applications running smoothly.

This reference architecture (RA) discusses the components that define the scalable and modular architecture of DGX SuperPOD. The system is built on the concept of scalable units (SU), each containing 32 DGX H200 systems, which provides for rapid deployment of systems of multiple sizes. This RA includes details regarding the SU design and specifics of InfiniBand, NVLink network, Ethernet fabric topologies, storage system specifications, recommended rack layouts, and wiring guides.

Chapter 2. Key Components of the DGX SuperPOD

The DGX SuperPOD architecture has been designed to maximize performance for state-of-the-art model training, scale to exaflops of performance, provide the highest performance to storage and support all customers in the enterprise, higher education, research, and the public sector. It is a digital twin of the main NVIDIA research and development system, meaning the company's software, applications, and support structure are first tested and vetted on the same architecture. Using SUs, system deployment times are reduced from months to weeks. Leveraging the DGX SuperPOD design reduces time-to-solution and time-to-market of next generation models and applications.

DGX SuperPOD is the integration of key NVIDIA components, as well as storage solutions from partners certified to work in a DGX SuperPOD environment.

2.1. NVIDIA DGX H200 System

The NVIDIA DGX H200 system (Figure 2.1) is an AI powerhouse that enables enterprises to expand the frontiers of business innovation and optimization. The DGX H200 system delivers breakthrough AI performance with the most powerful chips ever built, in an eight GPU configuration. The NVIDIA Blackwell GPU architecture provides the latest technologies that brings months of computational effort down to days and hours, on some of the largest AI/ML workloads.




Figure 2.1: DGX H200 system

The primary highlight of the DGX H200 system over the DGX H100 system include:

- 1,128 GBs of aggregated HBM3 memory

2.2. NVIDIA InfiniBand Technology

InfiniBand is a high-performance, low latency, RDMA capable networking technology, proven over 20 years in the harshest compute environments to provide the best inter-node network performance. **InfiniBand** continues to evolve and lead data center network performance.

The latest generation InfiniBand, NDR, has a peak speed of 400 Gbps per direction with an extremely low port-to-port latency. It is backwards compatible with the previous generations of InfiniBand specifications. InfiniBand is more than just peak bandwidth and **low latency**. InfiniBand provides additional features to optimize performance including adaptive routing (AR), collective communication with SHARP™, dynamic network healing with SHIELD™, and supports several network topologies including fat-tree, Dragonfly, and multi-dimensional Torus to build the largest fabrics and compute systems possible.

2.3. Runtime and System Management

The DGX SuperPOD RA represents the best practices for building high-performance data centers. There is flexibility in how these systems can be presented to customers and users. NVIDIA Base Command Manager software is used to manage all DGX SuperPOD deployments.

DGX SuperPOD can be deployed on-premises, meaning the customer owns and manages the hardware as a traditional system. This can be within a customer's data center or co-located at a commercial data center, but the customer owns the hardware.

2.4. Components

The hardware components of DGX SuperPOD are described in [Table 2.1](#). The software components are shown in [Table 2.2](#).

Table 2.1: DGX SuperPOD / 4 SU hardware components

Component	Technology	Description
Compute nodes	NVIDIA DGX H200 system with eight H200 GPUs	The world's premier purpose-built AI systems featuring NVIDIA H200 Tensor Core GPUs, fourth-generation NVIDIA NVLink, and third-generation NVIDIA NVSwitch™ technologies.
Compute fabric	NVIDIA Quantum QM9700 NDR 400 Gbps InfiniBand	Rail-optimized, non-blocking, full fat-tree network with eight NDR400 connections per system
Storage fabric	NVIDIA Quantum QM9700 NDR 400 Gbps InfiniBand	The fabric is optimized to match peak performance of the configured storage array
Compute/storage fabric management	NVIDIA Unified Fabric Manager Appliance, Enterprise Edition	NVIDIA UFM combines enhanced, real-time network telemetry with AI powered cyber intelligence and analytics to manage scale-out InfiniBand data centers
In-band management network	NVIDIA SN4600C switch	64 port 100 Gbps Ethernet switch providing high port density with high performance
Out-of-band (OOB) management network	NVIDIA SN2201 switch	48 port 1 Gbps Ethernet switch leveraging copper ports to minimize complexity

Table 2.2: DGX SuperPOD software components

Component	Description
NVIDIA Base Command Manager	Comprehensive AI infrastructure management for AI clusters. It automates provisioning and administration and supports cluster sizes into the thousands of nodes.
NVIDIA AI Enterprise	Best-in-class development tools and frameworks for the AI practitioner and reliable management and orchestration for IT professionals
Magnum IO	Enables increased performance for AI and HPC
NVIDIA NGC	The NGC catalog provides a collection of GPU-optimized containers for AI and HPC
Slurm	A classic workload manager used to manage complex workloads in a multi-node, batch-style, compute environment

2.5. Design Requirements

DGX SuperPOD is designed to minimize system bottlenecks throughout the tightly coupled configuration to provide the best performance and application scalability. Each subsystem has been thoughtfully designed to meet this goal. In addition, the overall design remains flexible so that data center requirements can be tailored to better integrate into existing data centers.

2.5.1. System Design

DGX SuperPOD is optimized for a customers' particular workload of multi-node AI and HPC applications:

- ▶ A modular architecture based on SUs of 32 DGX H200 systems each.
- ▶ A fully tested system scales to four SUs, but larger deployments can be built based on customer requirements.
- ▶ Rack design can support two DGX H200 systems per rack, so that the rack layout can be modified to accommodate different data center requirements.
- ▶ Storage partner equipment that has been certified to work in DGX SuperPOD environments.
- ▶ Full system support—including compute, storage, network, and software—is provided by NVIDIA Enterprise Support (NVEX).

2.5.2. Compute Fabric

- ▶ The compute fabric is rail-optimized to the top layer of the fabric.
- ▶ The compute fabric is a balanced, full-fat tree.
- ▶ Managed NDR switches are used throughout the design to provide better management of the fabric.
- ▶ The fabric is designed to support the latest SHaRPv3 features.

2.5.3. Storage Fabric

The storage fabric provides high bandwidth to shared storage. It also has the following characteristics:

- ▶ It is independent of the compute fabric to maximize performance of both storage and application performance.
- ▶ Provides single-node bandwidth of at least 40 GBps to each DGX H200 system.
- ▶ Storage is provided over InfiniBand and leverages RDMA to provide maximum performance and minimize CPU overhead.
- ▶ It is flexible and can scale to meet specific capacity and bandwidth requirements.
- ▶ User-accessible management nodes provide access to shared storage.

2.5.4. In-Band Management Network

- ▶ The in-band management network fabric is Ethernet-based and is used for node provisioning, data movement, Internet access, and other services that must be accessible by the users.
- ▶ The in-band management network connections for compute and management servers operate at 100 Gbps and are bonded for resiliency.

2.5.5. Out-of-Band Management Network

The OOB management network connects all the base management controller (BMC) ports, as well as other devices that should be physically isolated from system users.

2.5.6. Storage Requirements

The DGX SuperPOD compute architecture must be paired with a high-performance, balanced, storage system to maximize overall system performance. DGX SuperPOD is designed to use two separate storage systems, high-performance storage (HPS) and user storage, optimized for key operations of throughput, parallel I/O, as well as higher IOPS and metadata workloads.

2.5.6.1 High-Performance Storage

HPS must provide:

- ▶ High-performance, resilient, POSIX-style file system optimized for multi-threaded read and write operations across multiple nodes.
- ▶ Native InfiniBand support.
- ▶ Local system RAM for transparent caching of data.
- ▶ Leverage local disk transparently for read and write caching.

2.5.6.2 User Storage

User storage must:

- ▶ Be designed for high metadata performance, IOPS, and key enterprise features such as checkpointing. This is different than the HPS, which is optimized for parallel I/O and large capacity.
- ▶ Communicate over Ethernet to provide a secondary path to storage so, that in the event of a failure of the storage fabric or HPS, nodes can still be accessed and managed by administrators in parallel.

Chapter 3. DGX SuperPOD Architecture

The DGX SuperPOD architecture is a combination of DGX systems, InfiniBand and Ethernet networking, management nodes, and storage. [Figure 3.1](#) shows the rack layout of a single SU. In this example, power consumption per rack exceeds 25 kW. The rack layout can be adjusted to meet local data center requirements, such as maximum power per rack and rack layout between DGX systems and supporting equipment to meet local needs for power and cooling distribution.

[Figure 3.2](#) shows an example management rack configuration with networking switches, management servers, storage arrays, and UFM appliances. Sizes and quantities will vary depending upon models used.

This reference architecture is focused on 4 SU units with 128 DGX nodes. DGX SuperPOD can scale to much larger configurations up to and beyond 64 SU with 2000+ DGX H200 nodes. See [Table 3.1](#) for more information.

Table 3.1: Larger SuperPOD component counts

SU Count	Node Count	GPU Count	InfiniBand Switch Counts			Cable Counts		
			Leaf	Spine	Core	Node-Leaf	Leaf-Spine	Spine-Core
4	128	1024	32	16	–	1024	1024	–
8	256	2048	64	32	–	2048	2048	–
16	512	4096	128	128	64	4096	4096	4096
32	1024	8192	256	256	128	8192	8192	8192
64	2048	16384	512	512	256	16384	16384	16384

Contact NVIDIA for more information regarding DGX SuperPOD solutions beyond four scalable units.



Figure 3.1: Complete single SU rack layout




Figure 3.2: Management rack configuration

Chapter 4. Network Fabrics

Building systems by SU provides the most efficient designs. However, if a different node count is required due to budgetary constraints, data center constraints, or other needs, the fabric should be designed to support the full SU, including leaf switches and leaf-spine cables, and leave the portion of the fabric unused where these nodes would be located. This will ensure optimal traffic routing and ensure that performance is consistent across all portions of the fabric.

DGX SuperPOD configurations utilize four network fabrics:

- ▶ Compute Fabric
- ▶ Storage Fabric
- ▶ In-Band Management Network
- ▶ Out-of-Band Management Network Each network is detailed in this section.

Figure 4.1 shows the ports on the back of the DGX H200 CPU tray and the connectivity provided. The compute fabric ports in the middle use a two-port transceiver to access all eight GPUs. Each pair of in-band management and storage ports provide parallel pathways into the DGX H200 system for increased performance. The OOB port is used for BMC access. (The LAN port next to the BMC port is not used in DGX SuperPOD configurations.)




Figure 4.1: DGX H200 network ports

4.1. Compute Fabric

Figure 4.2 shows the compute fabric layout for the full 127-node DGX SuperPOD. Each group of 32 nodes is rail-aligned. Traffic per rail of the DGX H200 systems is always one hop away from the other 31 nodes in a SU. Traffic between nodes, or between rails, traverses the spine layer.

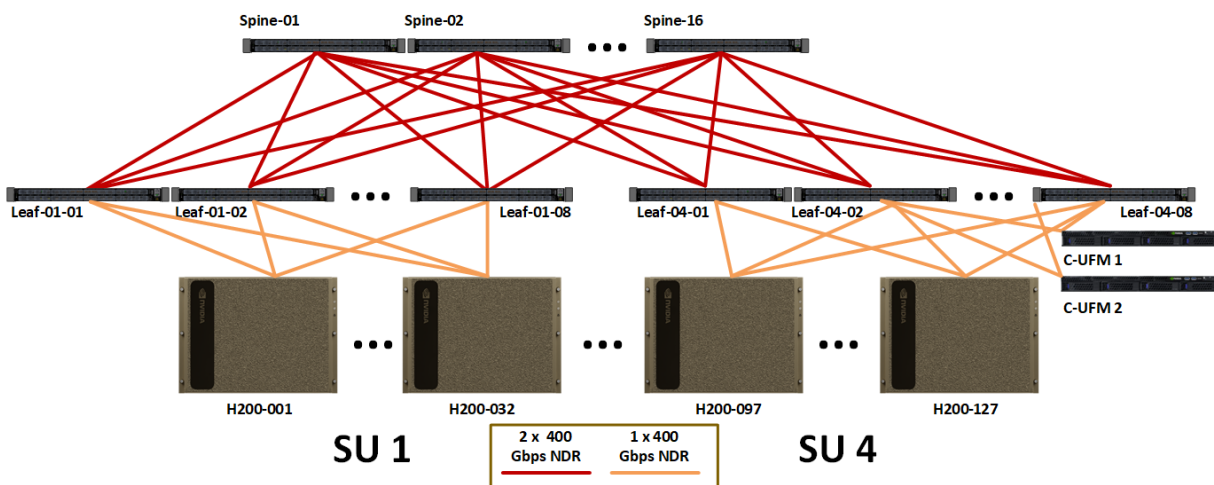


Figure 4.2: Compute fabric for full 127-node DGX SuperPOD

Table 4.1 shows the number of cables and switches required for the compute fabric for different SU sizes.

Table 4.1: Compute fabric component count

SU Count	Node Count	GPU Count	InfiniBand Switch Count		Cable Counts	
			Leaf	Spine	Compute UFM and Spine-Leaf	
1	31 ¹	248	8	4	252	256
2	63	504	16	8	508	512
3	95	760	24	16	764	768
4	127	1016	32	16	1020	1024

¹. This is a 32 node per SU design, however a DGX system must be removed to accommodate for UFM connectivity.

4.2. Storage Fabric

The storage fabric employs an InfiniBand network fabric that is essential to maximum bandwidth (Figure 4.3). This is because the I/O per-node for the DGX SuperPOD must exceed 40 GBps. High bandwidth requirements with advanced fabric management features, such as congestion control and AR, provide significant benefits for the storage fabric.

The storage fabric uses MQM9700-NS2F switches (Figure 4.4). The storage devices are connected at a

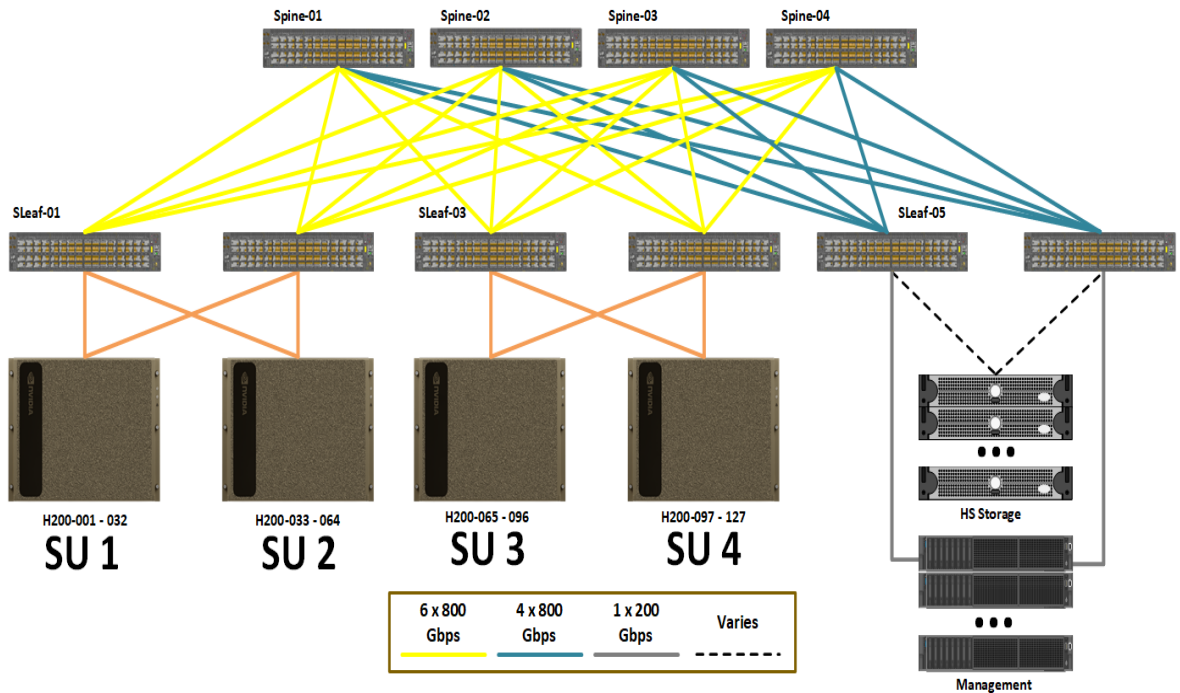


Figure 4.3: Storage fabric logical design

1:1 port to uplink ratio. The DGX H200 system connections are slightly oversubscribed with a ratio near 4:3 with adjustments as needed to enable more storage flexibility regarding cost and performance.



Figure 4.4: Figure 7. MQM9700-NS2F Switch

4.3. In-Band Management Network

The in-band management network provides several key functions:

- Connects all the services that manage the cluster.
- Enables access to the home filesystem and storage pool.
- Provides connectivity for the in-cluster services such as Base Command Manager, Slurm and to other services outside of the cluster such as the NGC registry, code repositories, and data sources.

Figure 4.5 shows the logical layout of the in-band Ethernet network. The in-band network connects the compute nodes and management nodes. In addition, the OOB network is connected to the in-band network to provide high-speed interfaces from the management nodes to support parallel operations to devices connected to the OOB storage fabric, such as storage.

The in-band management network uses **SN4600** switches (numref:sn4600c).

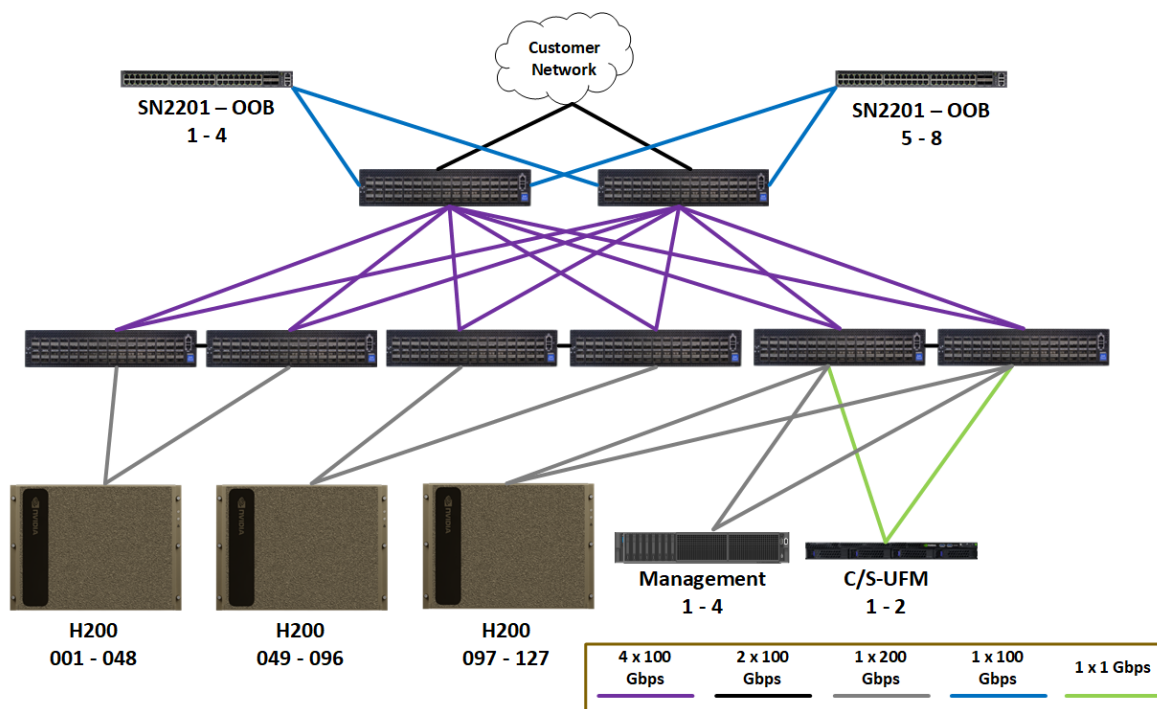


Figure 4.5: In-band Ethernet network



Figure 4.6: SN4600C switch

4.4. Out-of-Band Management Network

Figure 4.7 shows the OOB Ethernet fabric. It connects the management ports of all devices including DGX and management servers, storage, networking gear, rack PDUs, and all other devices. These are separated onto their own fabric because there is no use-case where users need access to these ports and are secured using logical network separation.

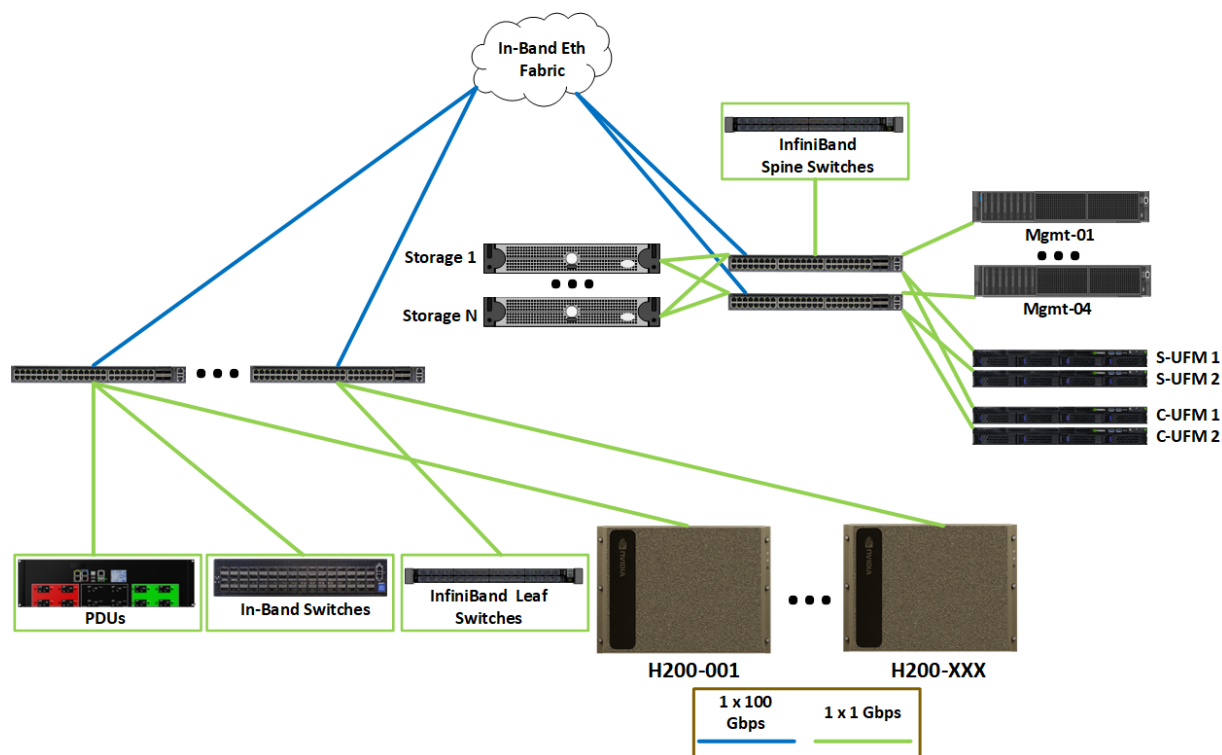


Figure 4.7: Logical OOB management network layout

The OOB management network uses SN2201 switches (Figure 4.8).



Figure 4.8: SN2201 switch

Chapter 5. Storage Architecture

Data, lots of data, is the key to development of accurate deep learning (DL) models. Data volume continues to grow exponentially, and data used to train individual models continues to grow as well. Data format, not just volume can play a key factor in the rate at which data is accessed so storage system performance must scale commensurately.

The key I/O operation in DL training is re-read. It is not just that data is read, but it must be reused again and again due to the iterative nature of DL training. Pure read performance still is important as some model types can train in a fraction of an epoch (ex: some recommender models) and inference of existing can be highly I/O intensive, much more so than training. Write performance can also be important. As DL models grow and time-to-train, writing checkpoints is necessary for fault tolerance. The size of checkpoint files can be terabytes in size and while not written frequently are typically written synchronously that blocks forward progress of DL models.

Ideally, data is cached during the first read of the dataset, so data does not have to be retrieved across the network. Shared filesystems typically use RAM as the first layer of cache. Reading files from cache can be an order of magnitude faster than from remote storage. In addition, the DGX H200 system provides local NVMe storage that can also be used for caching or staging data.

DGX SuperPOD is designed to support all workloads, but the storage performance required to maximize training performance can vary depending on the type of model and dataset. The guidelines in [Table 5.1](#) and [Table 5.2](#) are provided to help determine the I/O levels required for different types of models.

Table 5.1: Storage performance requirements

Perfor- mance Level	Work Description	Dataset Size
Good	Natural Language Processing (NLP)	Datasets generally fit within local cache
Better	Training Compressed Images, Compressed Audio and Text Data, such as LLM Training	Many to most datasets can fit within the local system's cache
Best	Training with large Video and Image files (such as AV replay), offline inference, ETL, generative networks such as stable diffusion, 3D images such as Medical U-Net, genomics workload and protein prediction such as AlphaFold	Datasets are too large to fit into cache, massive first epoch I/O requirements, workflows that only read the dataset once

Table 5.2: Guidelines for storage performance

Performance Characteristic ¹	Good (GBps)	Better (GBps)	Best (GBps)
Single node read	4	8	40
Single node write	2	4	20
Single SU aggregate system read	15	40	125
Single SU aggregate system write	7	20	62
4 SU aggregate system read	60	160	500
4 SU aggregate system write	30	80	250

Even for the best category in Table 5.2, it is desirable that the single node read performance is closer to the maximum network performance of 80 GBps.

Note

As datasets get larger, they may no longer fit in cache on the local system. Pairing large datasets that do not fit in cache with very fast GPUs can create a situation where it is difficult to achieve maximum training performance. NVIDIA GPUDirect Storage® (GDS) provides a way to read data from the local NVMe directly into GPU memory providing higher sustained I/O performance with lower latency. Using the storage fabric on DGX SuperPOD, a GDS-enabled application should be able to read data at over 27 GBps directly into the GPUs.

High-speed storage provides a shared view of an organization's data to all nodes. It must be optimized for small, random I/O patterns, and provide high peak node performance and high aggregate filesystem performance to meet the variety of workloads an organization may encounter. High-speed storage should support both efficient multi-threaded reads and writes from a single system, but most DL workloads will be read-dominant.

Use cases in automotive and other computer vision-related tasks, where high-resolution images are used for training (and in some cases are uncompressed) involve datasets that easily exceed 30 TB in size. In these cases, 4 GBps per GPU for read performance is needed.

While NLP and LLM cases often do not require as much read performance for training, peak performance for reads and writes are needed for creating and reading checkpoint files. This is a synchronous operation and training stops during this phase. If you are looking for best end-to-end training performance, do not ignore I/O operations for checkpoints. Consider at least ½ of the read performance as recommended write performance for LLM and large model use cases.

The preceding metrics assume a variety of workloads, datasets, and need for training locally and directly from the high-speed storage system. It is best to characterize workloads and organizational needs before finalizing performance and capacity requirements.

Chapter 6. DGX SuperPOD Software

DGX SuperPOD is an integrated hardware and software solution. The included software (Figure 6.1) is optimized for AI from top to bottom, from the accelerated frameworks and workflow management through to system management and low-level operating system (OS) optimizations, every part of the stack is designed to maximize the performance and value of DGX SuperPOD.

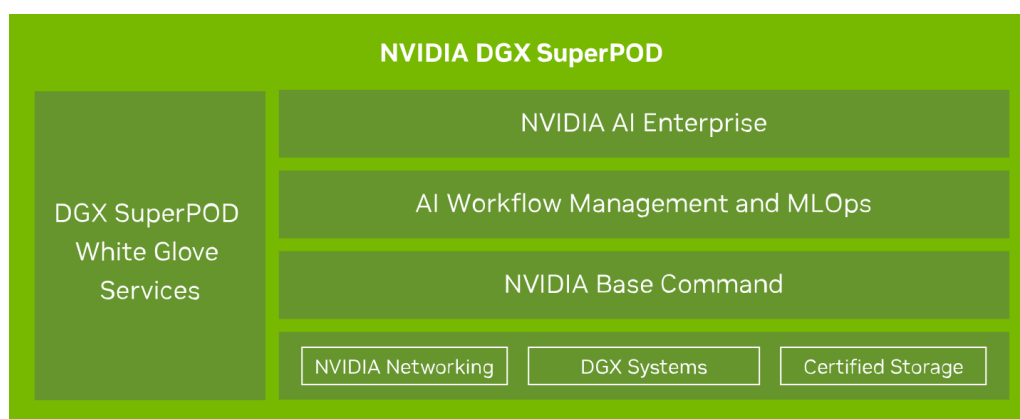


Figure 6.1: DGX SuperPOD high-level architecture

6.1. NVIDIA Base Command

NVIDIA Base Command powers every DGX SuperPOD, enabling organizations to leverage the best of NVIDIA software innovation. Enterprises can unleash the full potential of their investment with a proven platform that includes enterprise-grade orchestration and cluster management, libraries that accelerate compute, storage and network infrastructure, and an OS optimized for AI workloads.

6.2. NVIDIA NGC

NGC (Figure 6.2) provides software to meet the needs of data scientists, developers, and researchers with various levels of AI expertise.

Software hosted on NGC undergoes scans against an aggregated set of common vulnerabilities and exposures (CVEs), crypto, and private keys.

Software from the NGC catalog is tested and ensured to scale to multiple GPUs and in some cases, to scale to multi-node, ensuring users maximize the use of their DGX SuperPOD.

NGC CATALOG - GPU-OPTIMIZED SOFTWARE

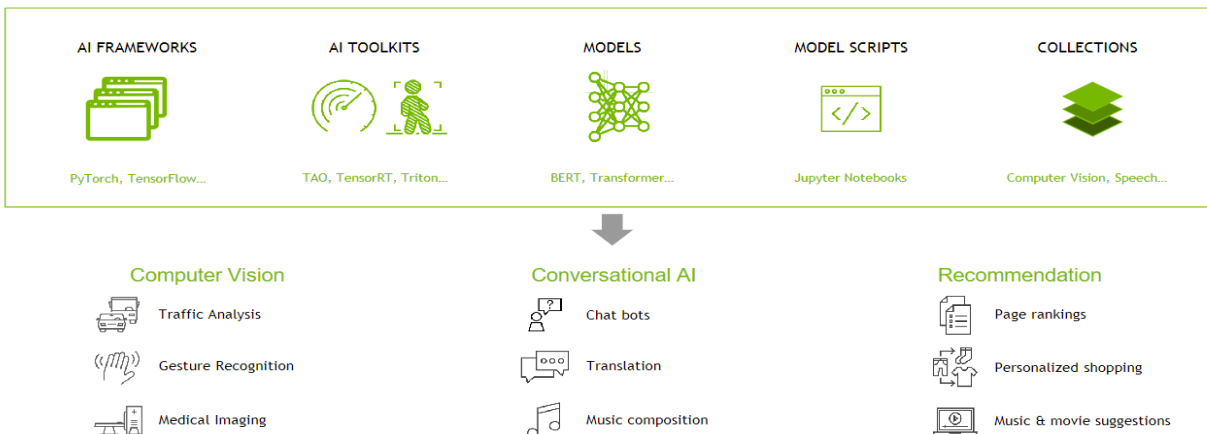


Figure 6.2: NGC catalog overview

6.3. NVIDIA AI Enterprise

NVIDIA AI Enterprise is the end-to-end software platform that brings generative AI into reach for every enterprise, providing the fastest and most efficient runtime for generative AI foundation models developed with the NVIDIA DGX platform. With production-grade security, stability, and manageability, it streamlines the development of generative AI solutions. NVIDIA AI Enterprise is included with DGX SuperPOD for enterprise developers to access pretrained models, optimized frameworks, microservices, accelerated libraries, and enterprise support.

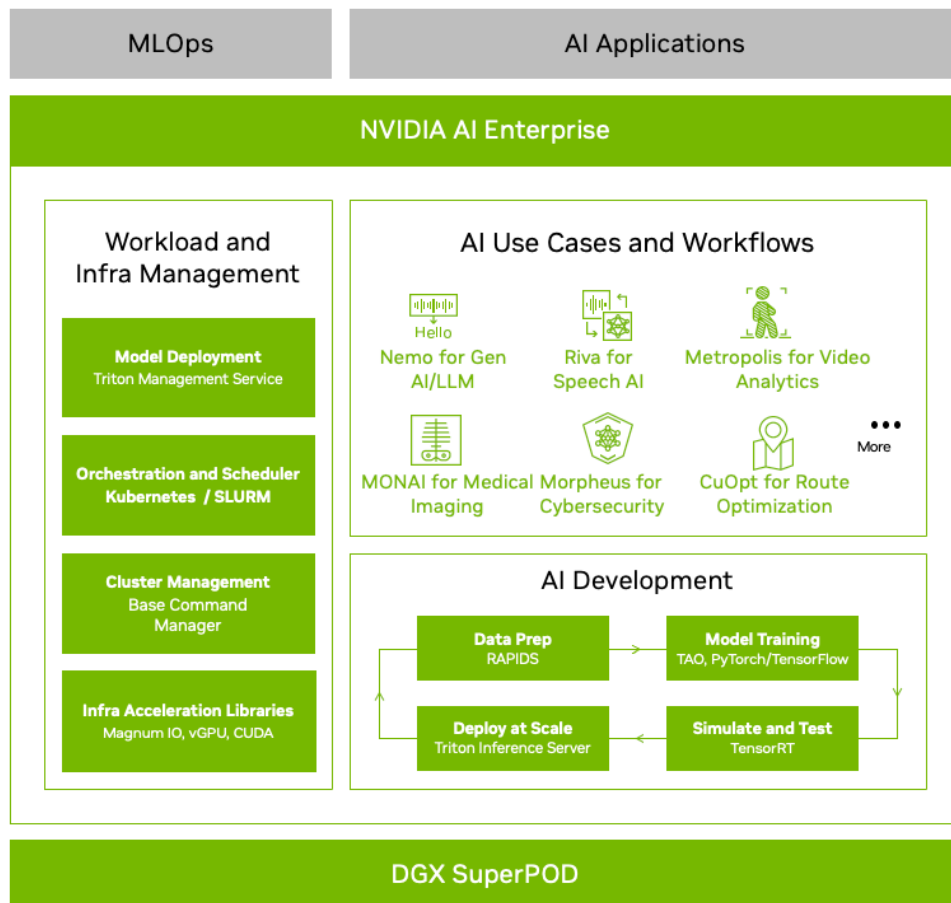


Figure 6.3: NVIDIA AI Enterprise Software Stack

Chapter 7. Summary

DGX SuperPOD with NVIDIA DGX H200 systems is the next generation of data center scale architecture to meet the demanding and growing needs of AI training. This RA document for DGX SuperPOD represents the architecture used by NVIDIA for our own AI model and HPC research and development. DGX SuperPOD continues to build upon its high-performance roots to enable training of the largest NLP models, support the expansive needs of training models for automotive applications, and scaling-up recommender models for greater accuracy and faster turn-around-time.

DGX SuperPOD represents a complete system of not just hardware but all the necessary software to accelerate time-to-deployment, streamline system management, proactively identify system issues. The combination of all these components keeps systems running reliably, with maximum performance, and enables users to push the bounds of state-of-the-art. The platform is designed to both support the workloads of today and grow to support tomorrow's applications.

Chapter 8. Notices

8.1. Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation (“NVIDIA”) makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer (“Terms of Sale”). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer’s sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application to avoid a default of the application or the product. Weaknesses in customer’s product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and

regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NONINFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

8.2. Trademarks

NVIDIA, the NVIDIA logo, NVIDIA DGX , and NVIDIA DGX SuperPOD are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

8.3. Copyright

© 2024-2025, NVIDIA Corporation. All rights reserved.

Copyright

©2024-2025, NVIDIA Corporation