NVIDIA DGX BasePOD: The Infrastructure Foundation for Enterprise AI

Reference Architecture

Featuring NVIDIA DGX H100 and A100 Systems

Document History

RA-11126-001

Version	Date	Authors	Description of Change
01	2022-09-23	Craig Tierney, Premal Savla, Robert Sohigian, Shawn Kaiser, and Tony Paikeday	Initial release
02	2022-09-30	Shawn Kaiser	Update to 40 node design and leaf change for 30 nodes
03	2022-10-05	Shawn Kaiser	Figure changes
04	2023-01-23	Craig Tierney, Premal Savla, Robert Sohigian, Shawn Kaiser, Scott Ellis	Added NDR200 configuration to NVIDIA DGX™ A100 NVIDIA A100 system
05	2023-01-31	Craig Tierney, Chris Kawalek, Robert Sohigian, and Shawn Kaiser	Miscellaneous updates
06	2023-02-15	Craig Tierney, Robert Sohigian, and Shawn Kaiser	Added NVIDIA DGX H100 system designs
07	2023-03-01	Craig Tierney, Robert Sohigian, and Shawn Kaiser	Reference architecture improvements

Abstract

The number of use cases for AI within an enterprise, including examples such as language modeling, cybersecurity, autonomous systems, and healthcare, continues to expand quickly. Not only have the number of use cases grown, but model complexity and data sources also are growing. The system required to process, train, and serve these next generation models must also grow. Training models commonly use dozens of GPUs for evaluating and optimizing different model configurations and parameters. Training data must be readily accessible to all the GPUs for these new kinds of workloads. In addition, organizations have many AI researchers that must train numerous models simultaneously. Enterprises need the flexibility for multiple developers and researchers to share these resources as they refine and bring their AI stack to production.

NVIDIA DGX BasePOD™ provides the underlying infrastructure and software to accelerate deployment and execution of these AI workloads. By building upon the success of NVIDIA DGX systems, DGX BasePOD is a prescriptive AI infrastructure for enterprises, eliminating the design challenges, lengthy deployment cycle, and management complexity traditionally associated with scaling AI infrastructure. Powered by NVIDIA Base Command™, DGX BasePOD provides the essential foundation for AI development optimized for enterprise.

This reference architecture discusses the key components of DGX BasePOD and provides a prescriptive design for DGX BasePOD solutions.

Contents

1.   DGX BasePOD Overview

DGX BasePOD is an integrated solution consisting of NVIDIA hardware and software components, MLOps solutions, and third-party storage. Leveraging best practices of scale-out system design with NVIDIA products and validated partner solutions, customers can implement an efficient and manageable platform for AI development. The designs in this DGX BasePOD reference architecture (RA) support developer needs, simplify IT manageability, and infrastructure scaling from two nodes to dozens with certified storage platforms from an industry-leading ecosystem. Optional MLOps solutions can be integrated with DGX BasePOD to enable a full stack solution to shorten AI model development cycles and speed the ROI of AI initiatives.

Figure 1 highlights the various components of NVIDIA DGX BasePOD. Each of these layers is an integration point that users typically would have to build and tune before an application could be deployed. The designs in the RA simplify system deployment and optimization using a validated prescriptive architecture.

Figure 1. Layers of integration for DGX BasePOD
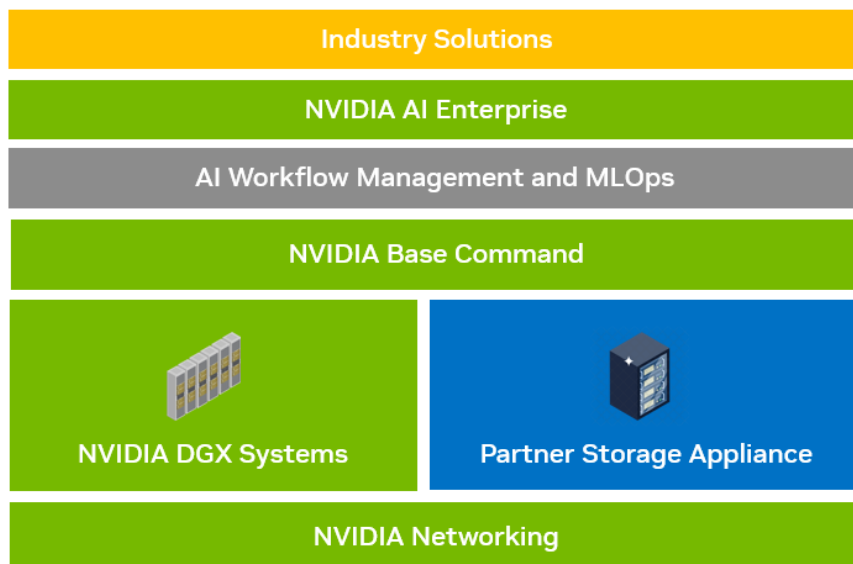
1.1   NVIDIA Networking

InfiniBand and Ethernet technologies enable networking functionality in DGX BasePOD. Proper networking is critical to ensuring that DGX BasePOD does not have any bottlenecks or suffer performance degradation for AI workloads. For more information on the products and technologies that enable this, refer to NVIDIA Networking.

1.2   Partner Storage Appliance

DGX BasePOD is built on a proven storage technology ecosystem. As NVIDIA validated storage partners introduce new storage technologies into the marketplace, they will qualify these new offerings with DGX BasePOD to ensure design compatibility and expected performance for known workloads. Every storage partner has performed rigorous testing to ensure that applications receive the highest performance and throughput when deployed with DGX BasePOD.

1.3   NVIDIA Software

1.3.1   NVIDIA Base Command

NVIDIA Base Command (Figure 2) powers every DGX BasePOD, 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 operating system (OS) optimized for AI workloads.

**Figure 2. NVIDIA Base Command features and capabilities with DGX BasePOD**
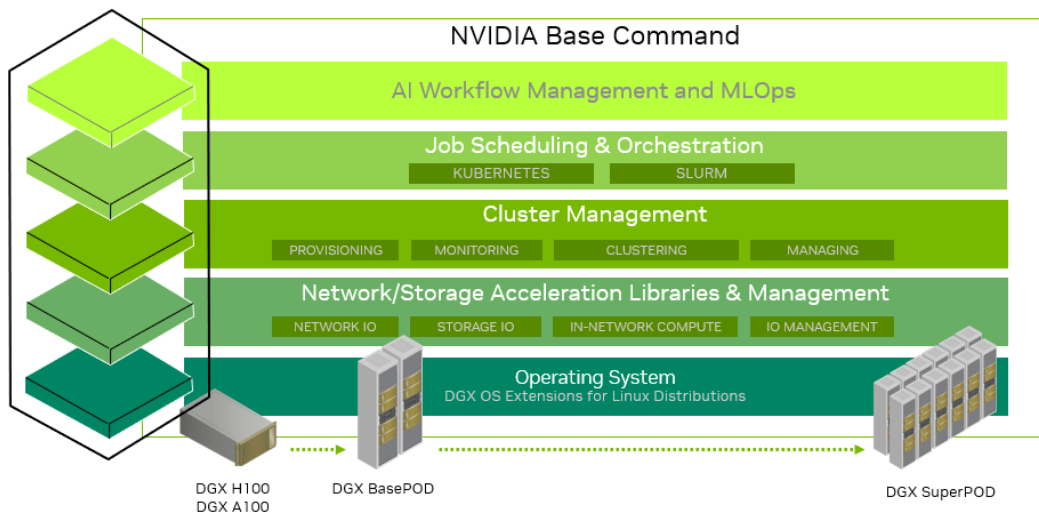
DGX BasePOD hardware is further optimized with acceleration libraries that know how to maximize the performance of AI workload across a GPU, the DGX system and an entire DGX cluster, speeding data access, movement, and management from system I/O to storage to network fabric.

Base Command provides integrated cluster management from installation and provisioning to ongoing monitoring of systems—from one to hundreds of DGX systems. Base Command also supports multiple methods for workflow management. Either Slurm or Kubernetes can be used to allow for optimal scheduling and management of system resources within a multi-user environment.

1.3.2   NVIDIA NGC

NVIDIA NGC™ (Figure 3) provides software to meet the needs of data scientists, developers, and researchers with various levels of AI expertise.

Figure 3. NGC catalog overview
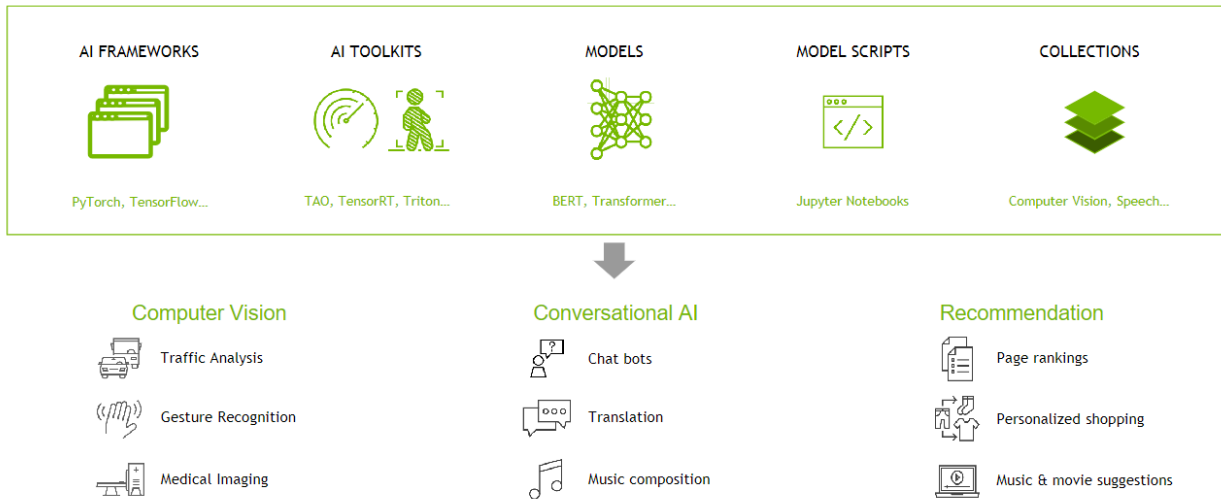
Software hosted on NGC undergoes scans against an aggregated set of common vulnerabilities and exposures (CVEs), crypto, and private keys. It is tested and designed to scale to multiple GPUs and in many cases, to multi-node, ensuring users maximize their investment in DGX systems.

1.3.3   NVIDIA AI Enterprise

NVIDIA AI Enterprise is a suite of AI and data analytics software optimized for the development and deployment of AI. NVIDIA AI Enterprise includes proven, open-sourced containers and frameworks such as NVIDIA RAPIDS™, NVIDIA TAO Toolkit, NVIDIA TensorRT™ and NVIDIA Triton Inference Server, which are certified and supported to run on DGX systems. NVIDIA AI Enterprise is included with DGX systems and is used in combination with NVIDIA Base Command and NVIDIA NGC.

2. Core Components

The compute, HCA, and switch resources form the foundation of the DGX BasePOD. The specific components used in the DGX BasePOD Reference Architectures are described in this section.

2.1 NVIDIA DGX Systems

NVIDIA DGX BasePOD configurations use DGX A100 and DGX H100 systems. The systems are described in the following sections.

2.1.1   NVIDIA DGX A100 System

The NVIDIA DGX A100 system (Figure 4) offers unprecedented compute density, performance, and flexibility in the world's first 5 petaFLOPS AI system.

Figure 4. DGX A100 system

Key specifications of the DGX A100 system are:

> Eight NVIDIA A100 GPUs.

> 40 GB or 80 GB GPU memory options.

> Six NVIDIA NVSwitch™ chips.

> Dual AMD EPYC™ 7742 CPUs, 128 total cores, 2.25 GHz (base), 3.4 GHz (max boost).

> Up to 2 TB of system memory.

> Eight NVIDIA ConnectX-6 or ConnectX-7 InfiniBand HCAs.

> Two 1.92 TB M.2 NVMe drives for DGX OS, eight 3.84 TB U.2 NVMe drives for storage/cache.

> 6.5 kW max power.

More information about DGX A100 systems is available here.

The rear ports of the DGX A100 CPU tray are shown in Figure 5. Four of the single-port ConnectX-7 HCAs are used for the InfiniBand compute fabrics. Each pair of dual-port ConnectX-7 HCAs provide parallel pathways to the storage and management fabrics. The out-of-band port is used for BMC access.

Figure 5. DGX A100 port mapping for various functions

2.1.2   NVIDIA DGX H100 System

The DGX H100 system (Figure 6) is the latest iteration of DGT Systems and the AI powerhouse that is accelerated by the groundbreaking performance of the NVIDIA H100 GPU.

Figure 6. DGX H100 system

Key specifications of the DGX H100 system are:

> Eight NVIDIA H100 GPUs.

> 80 GB GPU memory.

> Four NVIDIA NVSwitch™ chips.

> Dual 56-core fourth Gen Intel® Xeon® capable processors with PCIe 5.0 support.

> 2 TB of DDR5 system memory.

> Four OSFP ports serving eight single-port NVIDIA ConnectX-7 VPI, three dual-port NVIDIA ConnectX-7 VPI.

> Two 1.92 TB M.2 NVMe drives for DGX OS, eight 3.84 TB U.2 NVMe drives for storage/cache.

> 11.3 kW max power.

The rear ports of the DGX H100 CPU tray are shown in Figure 7. Four of the OSFP ports serve eight ConnectX-7 HCAs for the InfiniBand compute fabric. Each pair of dual-port ConnectX-7 HCAs provide parallel pathways to the storage and management fabrics. The out-of-band port is used for BMC access.

Figure 7. DGX H100 port mapping for various functions

2.2   NVIDIA Networking Adapters

NVIDIA DGX H100 systems are equipped with NVIDIA® ConnectX®-7 HCAs. NVIDIA DGX A100 systems are available with ConnectX-7or ConnectX-6 HCAs. Both are described in this section.

2.2.1   NVIDIA ConnectX-7 HCAs

The ConnectX-7 HCA (Figure 8) is the latest ConnectX HCA line. It can provide 25/50/100/200/400G of throughput. NVIDIA DGX systems use both the single and dual-port ConnectX-7 HCAs to provide flexibility in DGX BasePOD deployments with NDR. Specifications are available here.

Figure 8. NVIDIA ConnectX-7 HCA

2.2.2   NVIDIA ConnectX-6 HCAs

ConnectX-6 HCAs (Figure 9) can provide 10/25/40/50/100/200G of throughput. NVIDIA DGX systems use both the single and dual-port ConnectX-6 HCAs to provide flexibility in DGX BasePOD deployments with HDR. Specifications are available here.

Figure 9. NVIDIA ConnectX-6 HCA

2.3 NVIDIA Networking Switches

DGX BasePOD configurations can be equipped with four types of NVIDIA networking switches. The switches are described in this section, with how the switches are being deployed in the Reference Architectures section.

2.3.1 NVIDIA QM9700 Switch

NVIDIA QM9700 switches (Figure 10) with NDR InfiniBand connectivity power the compute fabric in NDR BasePOD configurations. ConnectX-7 single-port adapters are used for the InfiniBand compute fabric. Each NVIDIA DGX system has dual connections to each QM9700 switch, providing multiple high-bandwidth, low-latency paths between the systems.

Figure 10. NVIDIA QM9700 switch

2.3.2 NVIDIA QM8700 Switch

NVIDIA QM8700 switches (Figure 11) with HDR InfiniBand connectivity power the compute fabric in HDR BasePOD configurations. ConnectX-6 single-port adapters are used for the InfiniBand compute fabric. Each NVIDIA DGX system has dual connections to each QM8700 switch providing multiple high-bandwidth, low-latency paths between the systems.

Figure 11. NVIDIA QM8700 switch

2.3.3   NVIDIA SN4600 Switch

NVIDIA SN4600 switches (Figure 12) offer 128 total ports (64 per switch) to provide redundant connectivity for in-band management of the DGX BasePOD. The NVIDIA SN4600 switch can provide for speeds between 1 GbE and 200 GbE.

For storage appliances connected over Ethernet, the NVIDIA SN4600 switches are also used. The ports on the NVIDIA DGX dual-port HCAs are used for both in-band management and storage connectivity.

Figure 12. NVIDIA SN4600 switch

2.3.4   NVIDIA SN2201 Switch

NVIDIA SN2201 switches (Figure 13) offer 48 ports to provide connectivity for out-of-band management. Out-of-band management provides consolidated management connectivity for all components in BasePOD.

Figure 13. NVIDIA SN2201 switch

3. Reference Architectures

DGX BasePOD is a flexible solution that offers multiple prescriptive architectures. These architectures are adaptable to support the evolving demands of AI workloads.

3.1 DGX A100 System with HDR

The components of the DGX BasePOD are described in Table 1.

Table 1. DGX BasePOD components—DGX A100 HDR

Component	Technology
Compute nodes (2-40)	NVIDIA DGX A100 system with eight 40 or 80 GB A100 GPUs, with HDR InfiniBand
Compute fabric	NVIDIA Quantum QM8700 HDR 200 Gbps InfiniBand switch
Management and storage fabric	NVIDIA SN4600 switch
Out-of-band management fabric	NVIDIA SN2201 switch
Control plane	Minimum requirements (each server): > 2 × AMD EPYC 7352 CPU > 512 GB memory > 1 × 6.4 TB NVMe for storage > 2 × 480 GB M.2 RAID for OS > 4 × 200 Gbps network > 2 × 100 GbE network

3.1.1   System Architecture—DGX A100 HDR

Figure 14 depicts the architecture for the DGX BasePOD for up to ten nodes with HDR InfiniBand. DGX BasePOD with DGX A100 systems uses four compute connections from each node running at HDR200. The complete architecture has three networks, an InfiniBand-based compute network, an Ethernet fabric for system management and storage, and an out-of-band Ethernet network.

Figure 14. DGX BasePOD with up to ten systems—DGX A100 HDR
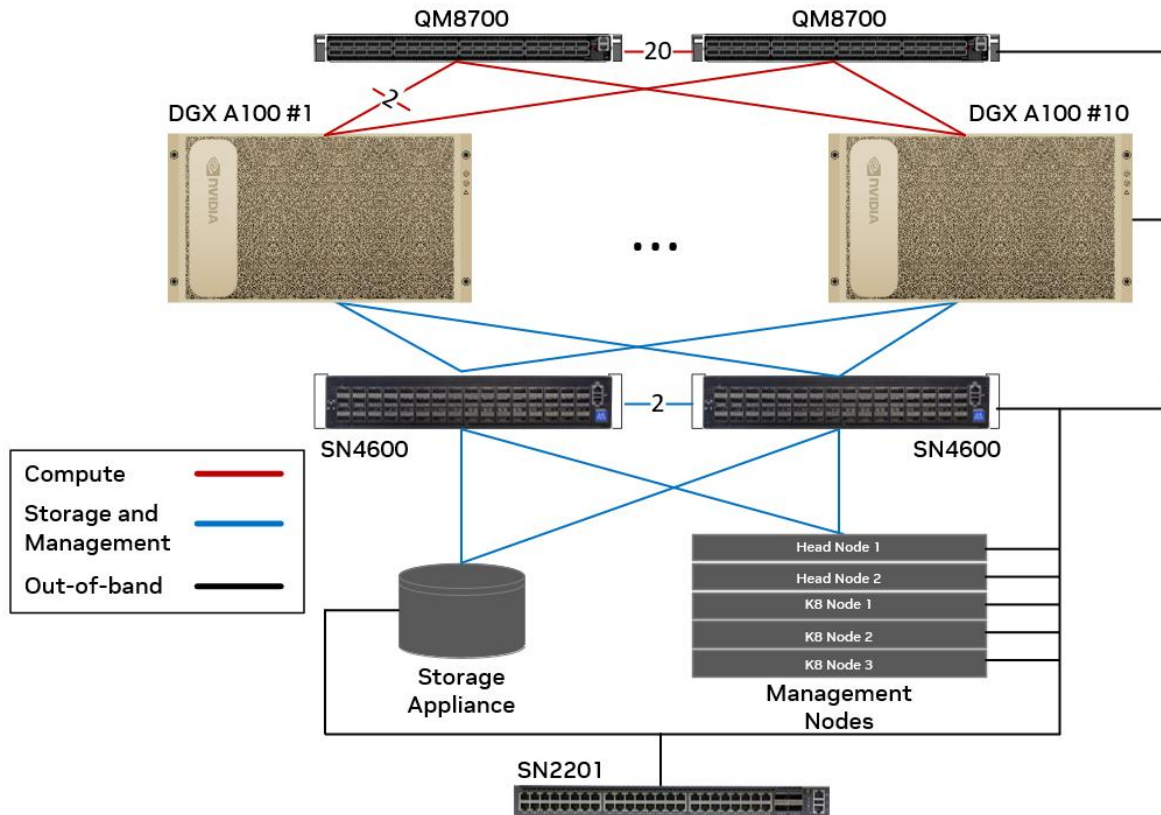
Included in the reference architecture are five CPU-only servers for system management. Two of these systems are used as the head nodes for Base Command Manager. The three additional systems provide the platform to house specific services for the deployment. This could be login nodes for a Slurm-based deployment, or Kubernetes master nodes supporting an MLOps-based partner solution. See Table 1 for control plane server requirements.

Figure 15 represents the DGX BasePOD architecture for larger configurations, up to 40 nodes, with HDR InfiniBand. To support the larger node count, the compute InfiniBand fabric requires scaling to two levels.

Figure 15. DGX BasePOD configurations with up to 40 systems—DGX A100 HDR
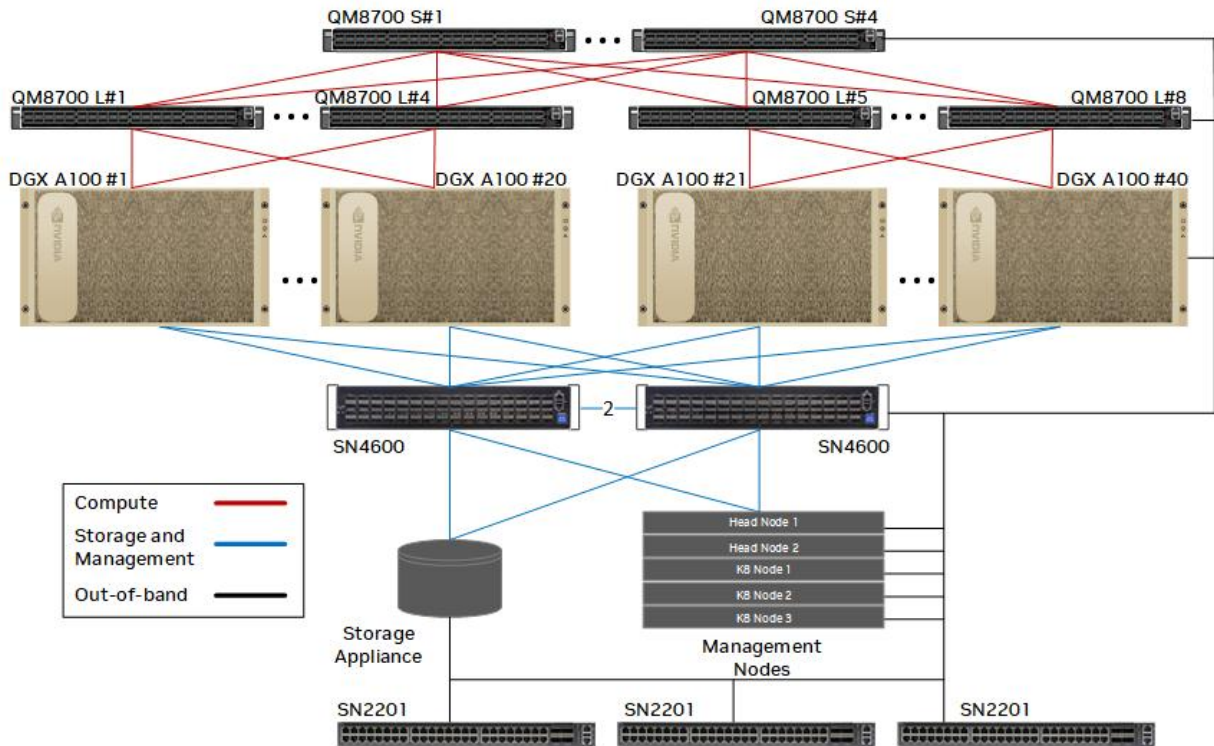
3.1.2   Switches and Cables—DGX A100 HDR

Table 2Table shows the number of cables and switches required for various deployments of DGX BasePOD with DGX A100 systems and HDR networking. These designs are built with active optical cables or direct attached copper. Alternatively, DGX BasePOD may be deployed with transceivers and fiber cables.

Table 2. Switches and cables—DGX A100 HDR

Component	Part Number	DGX A100 Systems			
		10	20	30	40
QM8700 InfiniBand switches, compute fabric	920-9B110-00FH-0MD	2	6	12	12
AOC from DGX A100 to compute fabric leaf switch	MFS1S00-HxxV	40	80	120	160
AOC from compute fabric leaf to spine InfiniBand switches	MFS1S00-HxxV	0	80	120	160
HDR InfiniBand DAC from leaf to leaf of compute fabric	MCA1J00-HxxxE	20	0	0	0
SN2201 switches for OOB management	MSN2201-CB2FC	1	2	3	3
SN4600 switches for in-band management and storage	920-9N302-00FA-0C0	2	2	2	2
1 GbE Cat 6 cables for OOB management system to switch	No specific requirement	33	57	83	107
200 GbE AOC for DGX A100 systems for in-band and storage	MFS1S00-Cxxx	20	40	60	80
200 GbE DAC for in-band fabric ISL	MCP1650-VxxxE26	2	2	2	2
100 GbE cables OOB to in-band switches	MFA1A00-Cxxx	2	4	6	6
BCM management servers	Varies	5	5	5	5
100 GbE AOC for management servers to in-band management switches	MFA1A00-Cxxx	10	10	10	10

3.2   DGX A100 System with NDR200

DGX BasePOD is a flexible solution that offers multiple prescriptive architectures. These architectures are adaptable to support the evolving demands of AI workloads.

The components of the DGX BasePOD are described in Table 3.

Table 3. DGX BasePOD components—DGX A100 NDR200

Component	Technology
Compute nodes (2-32)	NVIDIA DGX A100 system with eight 40 or 80 GB A100 GPUs and NDR200 InfiniBand networking
Compute fabric	NVIDIA Quantum QM9700 NDR400 Gbps InfiniBand switch
Management and storage fabric	NVIDIA SN4600 switches
Out-of-band management fabric	NVIDIA SN2201 switches
Control plane	Minimum requirements (each server): > 2 × AMD EPYC 7352 CPU > 512 GB memory > 1 × 6.4 TB NVMe for storage > 2 × 480 GB M.2 RAID for OS > 4 × 200 Gbps network > 2 × 100 GbE network

3.2.1　System Architecture—DGX A100 NDR200

Figure 16 depicts the architecture for the DGX BasePOD for up to thirty-two nodes with NDR InfiniBand. DGX BasePOD with DGX A100 uses four compute connections from each node running at NDR200. The complete architecture has three networks, an InfiniBand-based compute network, an Ethernet fabric for system management and storage, and an out-of-band Ethernet network.

Figure 16. DGX BasePOD with up to 32 systems—DGX A100 NDR200
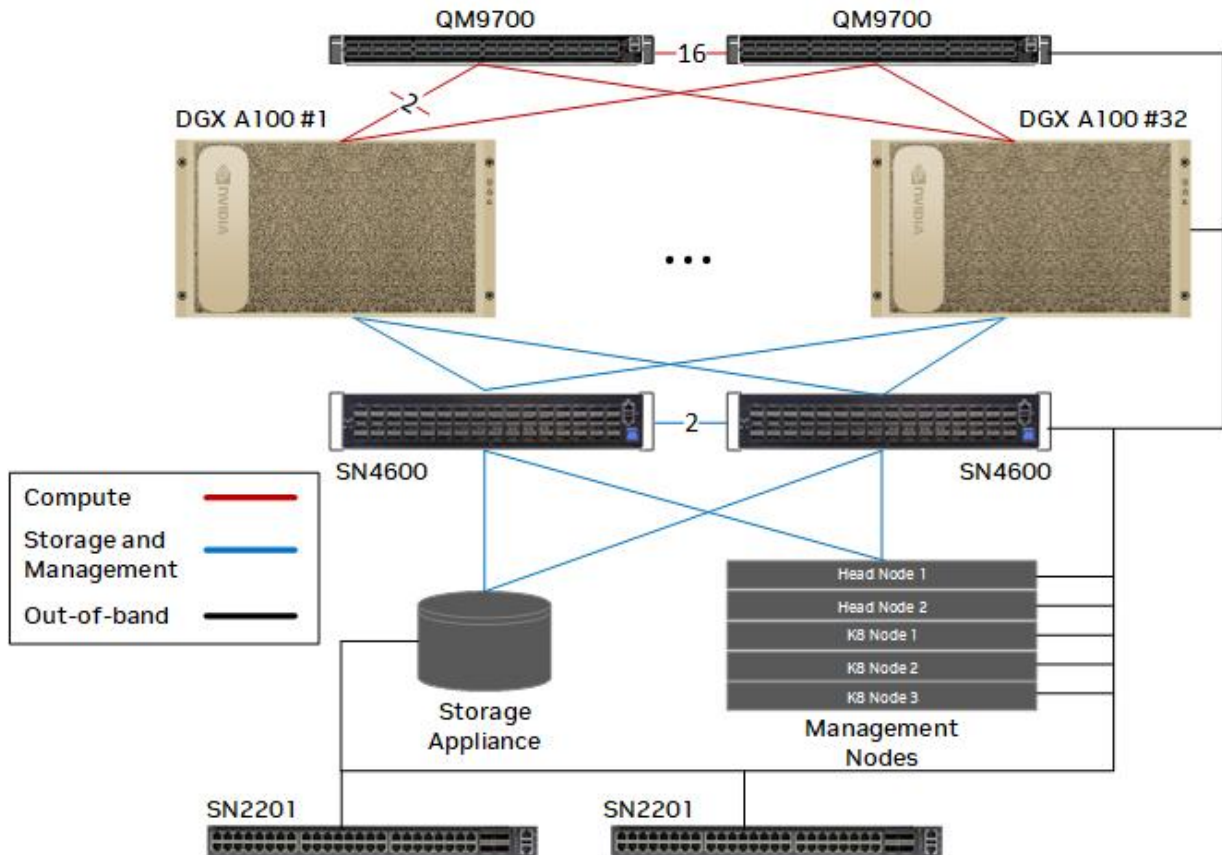
Included in the reference architecture are five dual-socket x86 servers for system management. Two nodes are used as the head nodes for Base Command Manager. The three additional nodes provide the platform to house specific services for the deployment. This could be login nodes for a Slurm-based deployment, or Kubernetes for MLOps-based partner solutions. Any OEM server that meets the minimum requirements for each node described Table 3 can be used.

3.2.2 Switches and Cables—DGX A100 NDR200

Table 4 shows the number of cables and switches required for various deployments of DGX BasePOD. These designs are built with active optical cables or direct attached copper. Alternatively, DGX BasePOD may be deployed with transceivers and fiber cables.

Table 4. Switches and cables—DGX A100 NDR200

Component	Part Number	DGX A100 Systems		
		10	20	32
QM9700 InfiniBand switches	QM9700	2	6	12
NDR200 MPO split-fiber InfiniBand cable from DGX A100 to leaf switch	MFP7E40-N0xx	20	40	64
Single Port OSFP transceiver for DGX A100	MMA4Z00-NS400	40	80	128
Dual Port OSFP transceiver for switch	MMA4Z00-NS	20	40	64
DAC links from leaf-to-leaf NDR InfiniBand DAC from leaf-to-leaf	MCP4Y10-Nxxx	10	20	32
SN2201 switches	MSN2201-CB2FC	1	1	2
SN4600 switches	920-9N302-00FA-0C0	2	2	2
1 GbE Cat 6 cables	No specific requirement	31	34	51
200 GbE AOC for DGX A100 systems from DGX to Ethernet switch	MFS1S00-Cxxx	20	40	60
200 GbE DAC for ISL	MCP1650-VxxxE26	2	2	2
100 GbE cables OOB to in-band	MFA1A00-Cxxx	2	2	4
BCM management servers	Varies	5	5	5
100 GbE AOC for management servers	MFA1A00-Cxxx	10	10	10

3.3   DGX H100 System with NDR200

DGX BasePOD is a flexible solution that offers multiple prescriptive architectures. These architectures are adaptable to support the evolving demands of AI workloads.

The components of the DGX BasePOD are described in Table 5.

Table 5. DGX BasePOD components—DGX H100 NDR200

Component	Technology
Compute nodes (2-16)	NVIDIA DGX H100 system with eight 80 GB H100 GPUs and NDR200 InfiniBand networking
Compute fabric	NVIDIA Quantum QM9700 NDR400 Gbps InfiniBand switch
Management and storage fabric	NVIDIA SN4600 switches
Out-of-band management fabric	NVIDIA SN2201 switches
Control plane	Minimum requirements (each server): > 2 × Intel Xeon 5318Y > 512 GB memory > 1 × 6.4 TB NVMe for storage > 2 × 480 GB M.2 RAID for OS > 4 × 200 Gbps network > 2 × 10 GbE network

3.3.1 System Architecture—DGX H100 NDR200

Figure 17 depicts the architecture for the DGX BasePOD for up to 16 nodes with NDR InfiniBand. BasePOD with DGX H100 systems uses eight compute connections from each node running at NDR200. The complete architecture has three networks, an InfiniBand-based compute network, an Ethernet fabric for system management and storage, and an out-of-band Ethernet network.

Figure 17. DGX BasePOD with up to 16 systems—DGX H100 NDR200
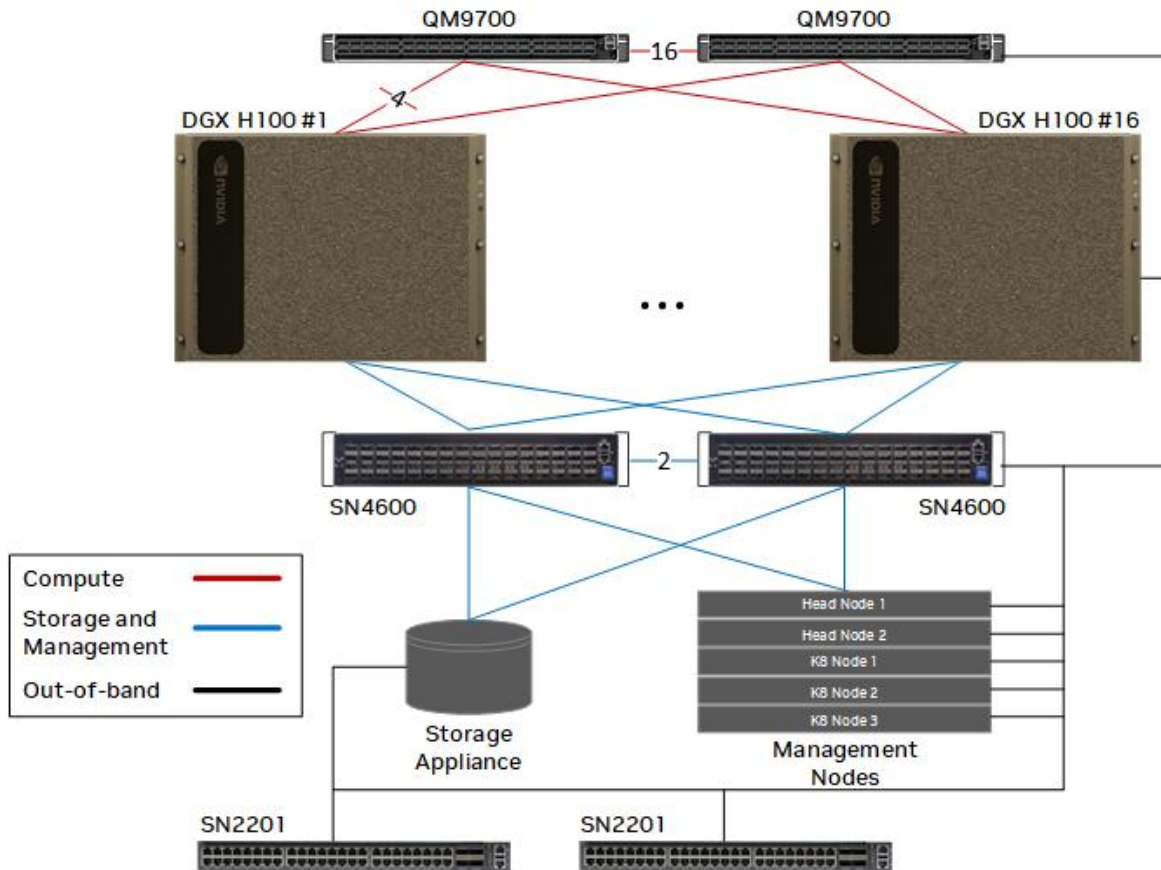
Included in the reference architecture are five dual-socket x86 servers for system management. Two nodes are used as the head nodes for Base Command Manager. The three additional nodes provide the platform to house specific services for the deployment. This could be login nodes for a Slurm-based deployment, or Kubernetes for MLOps-based partner solutions. Any OEM server that meets the minimum requirements for each node described in Table 5 can be used.

3.3.2 Switches and Cables—DGX H100 NDR200

Table 6 shows the number of cables and switches required for various deployments of DGX BasePOD. These designs are built with active optical cables or direct attached copper. Alternatively, DGX BasePOD may be deployed with transceivers and fiber cables.

Table 6. Switches and cables—DGX H100 NDR200

Components	Part Number	DGX H100 Systems		
		4	8	16
QM9700 InfiniBand switches	QM9700	2	2	2
NDR200 MPO InfiniBand cable from DGX A100 systems to leaf switch	MFP7E40-N0xx	16	32	64
Dual Port twin-OSFP transceiver for DGX H100 system	MMA4Z00-NS-FLT	16	32	64
Dual Port OSFP transceiver for switch	MMA4Z00-NS	8	16	32
NDR InfiniBand DAC from leaf to leaf	MCP4Y10-Nxxx	4	8	16
SN2201 switches	MSN2201-CB2FC	1	2	2
SN4600 switches	920-9N302-00FA-0C0	2	2	2
1 GbE Cat 6 cables	No specific requirement	29	45	77
200 GbE AOC for DGX H100 systems	MFS1S00-Cxxx	8	16	32
200 GbE DAC for ISL	MCP1650-VxxxE26	2	2	2
100 GbE cables OOB to in-band	MFA1A00-Cxxx	2	4	4
BCM management servers	Varies	5	5	5
100 GbE AOC for management servers	MFA1A00-Cxxx	10	10	10

4. Summary

Every enterprise wants to leverage AI to improve their products, services, and processes. But many struggle with how to operationalize AI at scale. Production-ready AI infrastructure requires blending a complex set of leading-edge hardware and software into a complete solution. This takes time and expertise to design, is difficult to deploy, and expensive to support across a multilayered technology stack from a variety of vendors.

With DGX BasePOD, NVIDIA has done the work for you–solving the complexity of designing AI infrastructure, systemizing it to power AI development and deployment, and simplifying its management. NVIDIA DGX BasePOD incorporates tested and proven design principles into an integrated AI infrastructure solution that incorporates best-of-breed NVIDIA DGX systems, NVIDIA software, NVIDIA networking, and an ecosystem of high-performance storage to enable AI innovation for the modern enterprise.

NVIDIA Corporation | 2788 San Tomas Expressway, Santa Clara, CA 95051
http://www.nvidia.com