



DDN A³I[®] SOLUTIONS WITH NVIDIA DGX BASEPOD[™]

Fully-integrated and optimized infrastructure solutions for accelerated at-scale AI, Analytics and HPC featuring NVIDIA DGX H100 Systems

1. DDN A ³ I Enablement for NVIDIA DGX BasePOD	3
2. Solution Components	8
2.1. DDN AI400X2 Appliance	8
2.2. DDN Insight Software	9
2.3. NVIDIA DGX H100 System	10
2.4. NVIDIA Networking	11
3. DDN A ³ I Reference Architectures for DGX BasePOD	12
3.1. DGX BasePOD Network Recommendations	12
3.2. DGX BasePOD With 4 DGX H100 Systems	16
3.3. DGX BasePOD with 8 DGX H100 Systems	17
3.4. DGX BasePOD with 16 DGX H100 Systems	18
4. DDN A ³ I Solution Validation	19
4.1. DGX BasePOD System IO Performance	21
4.2. DGX BasePOD GPU IO Performance	23
5. Contact DDN for Additional Information	24

Executive Summary

DDN A³I Solutions are proven at-scale to deliver optimal data performance for Artificial Intelligence (AI), Data Analytics and High-Performance Computing (HPC) applications running on GPUs in NVIDIA DGX H100™ systems. This document describes fully validated reference architectures for scalable NVIDIA DGX BasePOD™ configurations. The solutions integrate DDN AI400X2 appliances and DDN Insight software with DGX H100 systems and NVIDIA network switches.

1. DDN A³I END-TO-END ENABLEMENT FOR NVIDIA DGX BASEPOD

DDN A³I solutions (Accelerated, Any-Scale AI) are architected to achieve the most from at-scale AI, Data Analytics and HPC applications running on DGX systems and DGX BasePOD. They provide predictable performance, capacity, and capability through a tight integration between DDN and NVIDIA systems. Every layer of hardware and software engaged in delivering and storing data is optimized for fast, responsive, and reliable access.

DDN A³I solutions are designed, developed, and optimized in close collaboration with NVIDIA. The deep integration of DDN AI appliances with DGX systems ensures a reliable experience. DDN A³I solutions are configured for flexible deployment in a wide range of environments and scale seamlessly in capacity and capability to match evolving workload needs. DDN A³I solutions are deployed globally and at all scale, from a single DGX system all the way to most of the largest NVIDIA DGX A100 and H100 SuperPOD™ clusters in operation today.

DDN brings the same advanced technologies used to power the world's largest supercomputers in a fully integrated package for DGX systems that's easy to deploy and manage. DDN A³I solutions are proven to maximize benefits for at-scale AI, Analytics and HPC workloads on DGX systems.

This section describes the advanced features of DDN A³I Solutions for DGX BasePOD.

1.1 DDN A³I SHARED PARALLEL ARCHITECTURE

The DDN A³I shared parallel architecture and client protocol ensures high levels of performance, scalability, security, and reliability for DGX systems. Multiple parallel data paths extend from the drives all the way to containerized applications running on the GPUs in the DGX system. With DDN's true end-to-end parallelism, data is delivered with high-throughput, low-latency, and massive concurrency in transactions. This ensures applications achieve the most from DGX systems with all GPU cycles put to productive use. Optimized parallel data-delivery directly translates to increased application performance and faster completion times. The DDN A³I shared parallel architecture also contains redundancy and automatic failover capability to ensure high reliability, resiliency, and data availability in case a network connection or server becomes unavailable.

1.2 DDN A³I STREAMLINED DEEP LEARNING WORKFLOWS

DDN A³I solutions enable and accelerate end-to-end data pipelines for deep learning (DL) workflows of all scale running on DGX systems. The DDN shared parallel architecture enables concurrent and continuous execution of all phases of DL workflows across multiple DGX systems. This eliminates the management overhead and risks of moving data between storage locations. At the application level, data is accessed through a standard highly interoperable file interface, for a familiar and intuitive user experience.

Significant acceleration can be achieved by executing an application across multiple DGX systems in a DGX SuperPOD simultaneously and engaging parallel training efforts of candidate neural networks variants. These advanced optimizations maximize the potential of DL frameworks. DDN works closely with NVIDIA and its customers to develop solutions and technologies that allow widely-used DL frameworks to run reliably on DGX systems.

1.3 DDN A³I MULTIRAIL NETWORKING

DDN A³I solutions integrate a wide range of networking technologies and topologies to ensure streamlined deployment and optimal performance for AI infrastructure. The latest generation NVIDIA Quantum InfiniBand (IB) and Spectrum Ethernet technology provide both high-bandwidth and low-latency data transfers between applications, compute servers and storage appliances.

DDN A³I Multirail enables grouping of multiple network interfaces on a DGX system to achieve faster aggregate data transfer capabilities. The feature balances traffic dynamically across all the interfaces, and actively monitors link health for rapid failure detection and automatic recovery. DDN A³I Multirail makes designing, deploying, and managing high-performance networks very simple, and is proven to deliver complete connectivity for at-scale infrastructure for DGX BasePOD deployments.

1.4 DDN A³I ADVANCED OPTIMIZATIONS FOR DGX H100 SYSTEM ARCHITECTURE

The DDN A³I client's NUMA-aware capabilities enable strong optimization for DGX systems. It automatically pins threads to ensure I/O activity across the DGX system is optimally localized, reducing latencies and increasing the utilization efficiency of the whole environment. Further enhancements reduce overhead when reclaiming memory pages from page cache to accelerate buffered operations to storage. The DDN A³I client software for DGX H100 systems has been validated at-scale with the largest DGX SuperPOD with DGX A100 and H100 systems deployments currently in operation.

1.5 DDN A³I HOT NODES

DDN Hot Nodes is a powerful software enhancement that enables the use of the NVME devices in a DGX system as a local cache for read-only operations. This method significantly improves the performance of applications if a data set is accessed multiple times during a particular workflow.

This is typical with DL training, where the same input data set or portions of the same input data set are accessed repeatedly over multiple training iterations. Traditionally, the application on the DGX system reads the input data set from shared storage directly, thereby continuously consuming shared storage resources. With Hot Nodes, as the input data is read during the first training iteration, the DDN software automatically writes a copy of the data on the local NVME devices. During subsequent reads, data is delivered to the application from the local cache rather than the shared storage. This entire process is managed by the DDN client software running on the DGX system. Data access is seamless and the cache is fully transparent to users and applications. The use of the local cache eliminates network traffic and reduces the load on the shared storage system. This allows other critical DL training operations like checkpointing to complete faster by engaging the full capabilities of the shared storage system.

DDN Hot Nodes includes extensive data management tools and performance monitoring facilities. These tools enable user-driven local cache management, and make integration simple with task schedulers. For example, training input data can be loaded to the local cache on a DGX system as a pre-flight task before the AI training application is engaged. As well, the metrics expose insightful information about cache utilization and performance, enabling system administrators to further optimize their data loading and maximize application and infrastructure efficiency gains.

1.6 DDN A³I MULTITENANCY

Through its built-in digital security framework, DDN A³I software makes it very easy to operate a DGX BasePOD as a secure multitenant environment. DDN A³I multitenancy makes it simple to share DGX systems across a large pool of users and still maintain secure data segregation. Multi-tenancy provides quick, seamless, dynamic DGX system resource provisioning for users. It eliminates resource silos, complex software release management, and unnecessary data movement between data storage locations. DDN A³I brings a very powerful multitenancy capability to DGX systems and makes it very simple for customers to deliver a secure, shared innovation space, for at-scale data-intensive applications.

1.7 DDN A³I CONTAINER CLIENT

Containers encapsulate applications and their dependencies to provide simple, reliable, and consistent execution. DDN enables a direct high-performance connection between the application containers on the DGX H100 system and the DDN parallel filesystem. This brings significant application performance benefits by enabling low latency, high-throughput parallel data access directly from a container. Additionally, the limitations of sharing a single host-level connection to storage between multiple containers disappear. The DDN in-container filesystem mounting capability is added at runtime through a universal wrapper that does not require any modification to the application or container.

Containerized versions of popular DL frameworks specially optimized for DGX systems are available from NVIDIA. They provide a solid foundation that enables data scientists to rapidly develop and deploy applications on DGX systems. In some cases, open-source versions of the containers are available, further enabling access and integration for developers. The DDN A³I container client provides high-performance parallelized data access directly from containerized applications on DGX system. This provides containerized DL frameworks with the most efficient dataset access possible, eliminating all latencies introduced by other layers of the computing stack.

1.8 DDN A³I S3 DATA SERVICES

DDN S3 Data Services provide hybrid file and object data access to the shared namespace. The multi-protocol access to the unified namespace provides tremendous workflow flexibility and simple end-to-end integration. Data can be captured directly to storage through the S3 interface and accessed immediately by containerized applications on a DGX system through a file interface. The shared namespace can also be presented through an S3 interface, for easy collaboration with multisite and multicloud deployments. The DDN S3 Data Services architecture delivers robust performance, scalability, security, and reliability features.

1.9 DDN A³I CSI DRIVER

The DDN CSI driver enables optimized data access for workloads managed by Kubernetes within DGX BasePOD. The driver provides direct control facilities for the container orchestrator for fully automated storage management. Several types of volumes and data access modes are supported. The driver also enables Kubernetes workloads to dynamically provision and deprovision storage resources based on the workload's requirements. The deep integration between DDN software and the orchestrator stack ensures most efficient management and utilization of storage resources. DDN engages in continuous software validation with CSI standards to ensure ongoing compatibility and integration of new capabilities as the standard evolves.

2. DDN A³I SOLUTIONS WITH NVIDIA DGX H100 SYSTEMS

The DDN A³I scalable architecture integrates DGX H100 systems with DDN AI shared parallel file storage appliances and delivers fully-optimized end-to-end AI, Analytics and HPC workflow acceleration on NVIDIA GPUs. DDN A³I solutions greatly simplify the deployment of DGX BasePOD configurations using DGX H100 systems in, while also delivering performance and efficiency for maximum GPU saturation, and high levels of scalability.

This section describes the components integrated in DDN A³I Solutions for DGX BasePOD.

2.1 DDN AI400X2 APPLIANCE

The AI400X2 appliance is a fully integrated and optimized shared data platform with predictable capacity, capability, and performance. Every AI400X2 appliance delivers over 90 GB/s and 3M IOPS directly to DGX H100 systems in DGX BasePOD. Shared performance scales linearly as additional AI400X2 appliances are integrated to DGX BasePOD. The all-NVMe configuration provides optimal performance for a wide variety of workload and data types and ensures that DGX BasePOD operators achieve the most from at-scale GPU applications, while maintaining a single, shared, centralized data platform.

The AI400X2 appliance integrates the DDN A³I shared parallel architecture and includes a wide range of capabilities described in section 1, including automated data management, digital security, and data protection, as well as extensive monitoring. The AI400X2 appliances enables DGX BasePOD operators to go beyond basic infrastructure and implement complete data governance pipelines at-scale.

The AI400X2 appliance integrates with DGX BasePOD over InfiniBand, Ethernet and RoCE. It is available in 60, 120, 250 and 500 TB all-NVMe capacity configurations. Optional hybrid configurations with integrated HDDs are also available for deployments requiring high-density deep capacity storage. Contact DDN Sales for more information.



Figure 1. DDN AI400X2 all-NVMe storage appliance.

2.2 DDN INSIGHT SOFTWARE

DDN Insight is a centralized management and monitoring software suite for AI400X2 appliances. It provides extensive performance and health monitoring of all DDN storage systems connected to DGX BasePOD from a single web-based user interface. DDN Insight greatly simplifies IT operations and enables automated and proactive storage platform management guided by analytics and intelligent software.

Performance monitoring is an important aspect of operating a DGX BasePOD efficiently. Provided the several variables that affect data I/O performance, the identification of bottlenecks and degradation is crucial while production workloads are engaged. DDN Insight provides deep real-time analysis across the entire DGX BasePOD cluster, tracking I/O transactions from applications running on DGX nodes all the way through individual drives in the AI400X2 appliances. The embedded analytics engine makes it simple for DGX BasePOD operators to visualize I/O performance across their entire infrastructure through intuitive user interfaces. These include extensive logging, trending, and comparison tools, for analyzing I/O performance of specific applications and users over time. As well, the open backend database makes it simple to extend the benefits of DDN Insight and integrate other AI infrastructure components within the engine, or export data to third party monitoring systems.

DDN Insight is available a software-installable package on customer-supplied management servers, and as a turnkey server appliance from DDN.

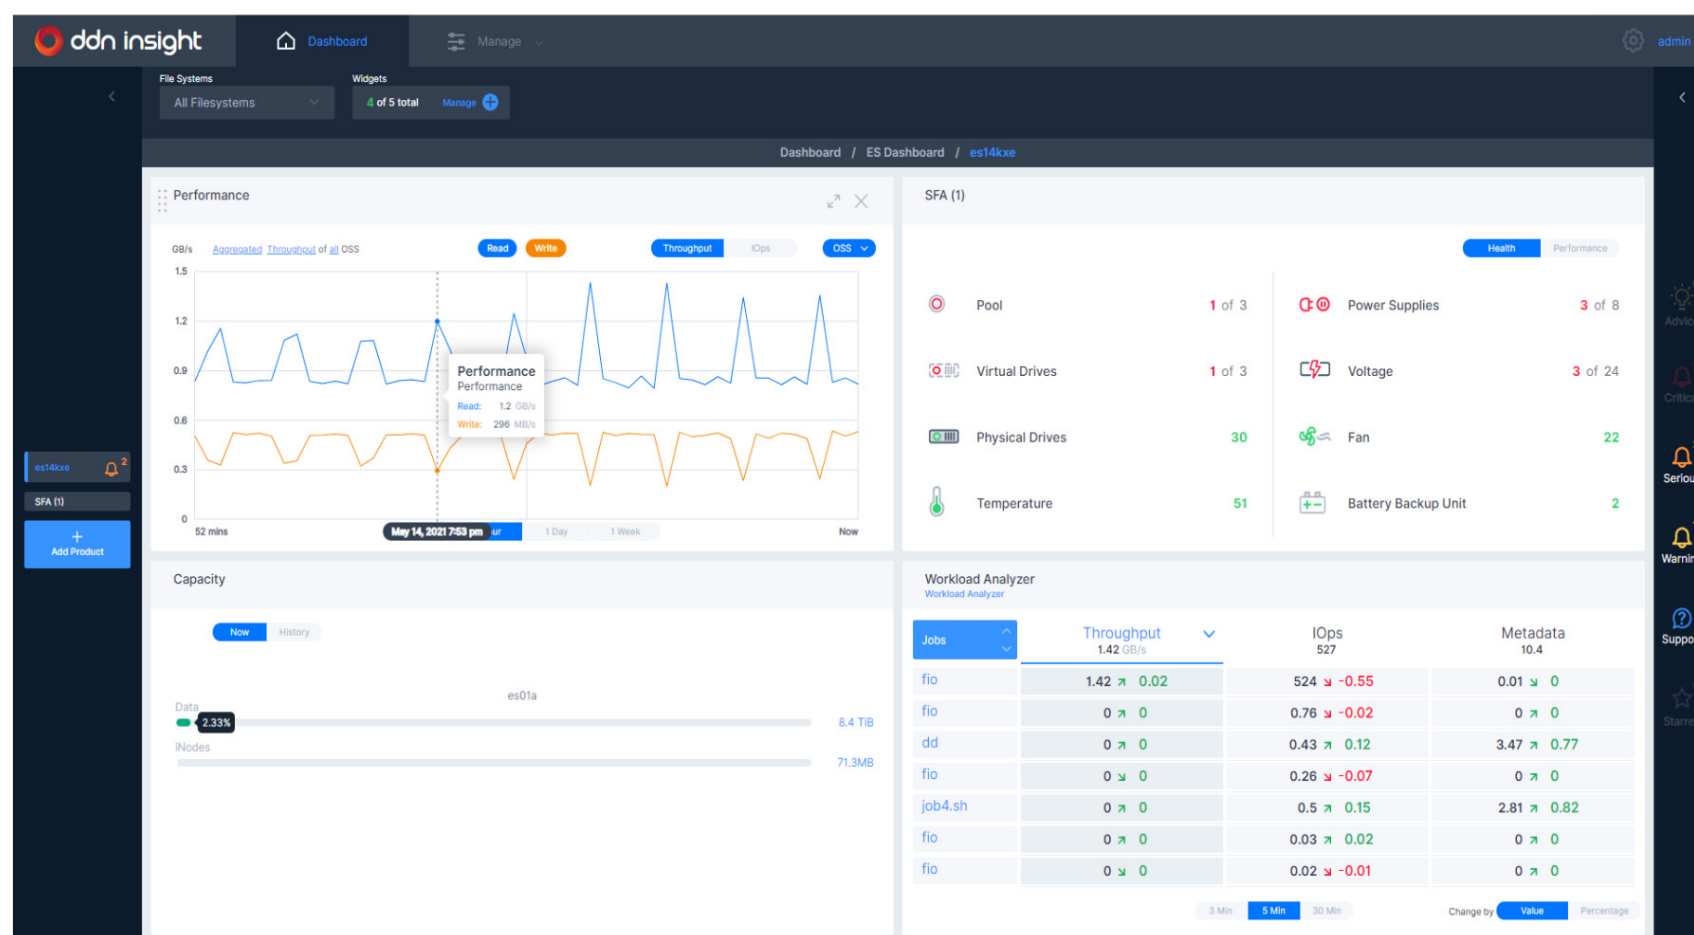


Figure 2. DDN Insight Workload Analyzer Tool.

2.3 NVIDIA DGX H100 SYSTEM

NVIDIA DGX H100 system is an AI powerhouse that enables enterprises to expand the frontiers of business innovation and optimization. The DGX H100 system, which is the fourth-generation NVIDIA DGX system, delivers AI excellence in an eight GPU configuration. The NVIDIA Hopper GPU architecture provides latest technologies such as the transformer engines and fourth-generation NVLink technology that brings months of computational effort down to days and hours, on some of the largest AI/ML workloads.

Some of the key highlights of the DGX H100 system over the DGX A100 system include:

- Up to 9X more performance with 32 petaFLOPS at FP8 precision.
- Server class x86-64 CPU with PCIe 5.0 support and DDR5 memory.
- 2X faster networking and storage @ 400 Gbps InfiniBand/Ethernet with NVIDIA ConnectX®-7 smart network interface cards (SmartNICs).
- 1.5X higher bandwidth per GPU @ 900 GBps with fourth generation of NVIDIA NVLink.
- 640 GB of aggregated HBM3 memory with 24 TB/s of aggregate memory bandwidth, 1.5X higher than DGX A100 system.



Figure 3. NVIDIA DGX H100 system.

2.4 NVIDIA NETWORKING SWITCHES

NVIDIA network switches provide optimal interconnect for DGX BasePOD. DDN recommends the NVIDIA Quantum-2 InfiniBand networking platform for both data-intensive compute and storage networks.

The NVIDIA Quantum-2 QM9700 InfiniBand Switch is recommended for DGX BasePOD compute and storage connectivity. It provides 64 ports of 400 Gb/s over 32 OSFP ports in a 1RU form factor. DDN recommends the QM9700 switch to deploy DGX BasePOD. Validated cabling configurations are detailed in section 3.1.3.



Figure 4. NVIDIA Quantum-2 QM9700 InfiniBand Switch.

The NVIDIA Spectrum-3 SN4600 Open Ethernet Switch is recommended for DGX BasePOD in-band management connectivity. It provides 64 QSFP56 200Gb ports in a 2 RU form factor.



Figure 5. NVIDIA Spectrum-3 SN4600 Open Ethernet Switch.

The NVIDIA Spectrum SN2201 Open Ethernet Switch is recommended for DGX BasePOD out-of-band management connectivity. It provides 52 ports with 48 RJ45 1Gb ports and 4 QSFP28 100Gb uplink ports in a 1 RU form factor



Figure 5. NVIDIA Spectrum-3 SN4600 Open Ethernet Switch.

3. DDN A³I REFERENCE ARCHITECTURES FOR DGX BASEPOD

configurations. DDN A3I solutions are fully validated with NVIDIA and already deployed with several DGX BasePOD and DGX SuperPOD customers worldwide.

The DDN AI400X2 appliance is a turnkey appliance for at-scale DGX deployments. DDN recommends the AI400X2 appliance as the optimal data platform for DGX BasePOD designs with the DGX H100 system. The AI400X2 appliances delivers optimal GPU performance for every workload and data type in a dense, power efficient 2RU chassis. The AI400X2 appliance simplifies the design, deployment, and management of a DGX BasePOD and provides predictable performance, capacity, and scaling. The appliance is designed for seamless integration with DGX systems and enables customers to move rapidly from test to production. As well, DDN provides complete expert design, deployment, and support services globally. The DDN field engineering organization has already deployed dozens of solutions for customers based on the A³I reference architectures.

As general guidance, DDN recommends the shared storage be sized to ensure at least 1 GB/s per second of read and write throughput for every NVIDIA H100 Tensor Core GPU in a DGX BasePOD (Table 1). This ensures sufficient performance for modern applications, including distributed training of large language models. These configurations can be adjusted and scaled easily to match more demanding workload requirements. Contact DDN Sales to review and discuss your recommendations for your specific application.

DGX BasePOD configuration	4 DGX H100	8 DGX H100	16 DGX H100
Recommended DDN storage	1 AI400X2	1 AI400X2	2 AI400X2
Shared read throughput	90 GB/s	90 GB/s	180 GB/s
Shared write throughput	65 GB/s	65 GB/s	130 GB/s
Per GPU read throughput	2.8 GB/s	1.4 GB/s	1.4 GB/s
Per GPU write throughput	2 GB/s	1 GB/s	1 GB/s

Table 1. Recommended baseline storage sizing for DGX BasePOD.

3.1.1 DDN A³I REFERENCE ARCHITECTURES FOR DGX BASEPOD

The DGX BasePOD reference design includes four networks:

Storage network. Provides connectivity between the AI400X2 appliances, the compute nodes and management nodes. Connects eight ports from each AI400X2 appliance. Connects two ports from each DGX H100 system, one each from two dual-port NVIDIA Mellanox ConnectX[®]-7 HCAs. Connects two ports from HCAs in the management nodes. The storage network can be InfiniBand IB or Ethernet. DDN recommends NVIDIA Quantum NDR 400Gbps InfiniBandIB for optimal performance and efficiency.

Recommended network connections for each DGX H100 system are shown on Figure 7 and recommended network connections for each AI400X2 appliance on Figure 8.

Compute network. Provides inter-node connectivity. Connects the four OSFP ports from each DGX H100 system. DDN recommends NVIDIA Quantum-2 400 Gb/s Infiniband for the compute network.

In-Band Management Network. Provides cluster provisioning, management and task scheduling. Connects two ports from each DGX H100 system, one each from two dual-port ConnectX[®]-7 HCAs. Connects two ports from HCAs in the management nodes. DDN recommends 200GbE for the in-band management network.

Out-of-Band Management Network. Provides management and monitoring connectivity for all DGX BasePOD components. DDN recommends 1GbE for the out-of-band management network.

3.1.1 DGX H100 SYSTEMS NETWORK CONNECTIVITY

For DGX BasePOD, DDN recommends ports 1 to 4 on the DGX H100 systems be connected to the compute network. Ports 5 and 7 should be connected to the in-band management network, which also serves as the storage network for Ethernet storage configurations. For InfiniBand storage configurations, ports 6 and 8 should be connected to the InfiniBand storage network. As well, the management BMC (“B”) port should be connected to the out-of-band management network.




Figure 7. Recommended DGX H100 system network port connections.

3.1.2 AI400X2 APPLIANCE NETWORK CONNECTIVITY

For DGX BasePOD, DDN recommends ports 1 to 8 on the AI400X2 appliance be connected to the storage network. As well, the management (“M”) and BMC (“B”) ports for both controllers should be connected to the out-of-band management network. Note that each AI400X2 appliance requires one inter-controller network port connection (“I”) using short ethernet cable supplied.




Figure 8. Recommended AI400X2 appliance network port connections.

3.1.3 AI400X2 APPLIANCE CABLING WITH NDR 400GBPS INFINIBAND SWITCHES

The AI400X2 appliance connects to the DGX BasePOD storage network with 8 200 Gbps InfiniBand interfaces. Particular attention must be given to the cabling selection to ensure compatibility between different InfiniBand connectivity and data rates. DDN and NVIDIA have validated the following cables to connect AI400X2 appliances with QM9700 switches. The use of splitter cables ensures most efficient use of switch ports.

Direct Attached Passive Copper Cables

MCP7Y60-Hxxx 980-9I46K-00Hxxx	NVIDIA DAC 1:2 splitter, InfiniBand 400Gb/s to 2x 200Gb/s, OSFP to 2x QSFP56 xxx indicates length in meters: 001, 01A (1.5m), 002
----------------------------------	--

Requires 4 cables per AI400X2 appliance.

Active Optical Cables

MFA7U10-Hxxx 980-9I117-00Hxxx	NVIDIA AOC 1:2 splitter, InfiniBand 400Gb/s to 2x 200Gb/s, OSFP to 2x QSFP56 xxx indicates length in meters: 003, 005, 010, 015, 020, 030
----------------------------------	--

Requires 4 cables per AI400X2 appliance.

3.2 DGX BASEPOD WITH 4 DGX H100 SYSTEMS

Figure 9 illustrates the DDN A3I architecture for DGX BasePOD deployed with a dedicated InfiniBand storage network. In this configuration, four DGX H100 systems are connected with an AI400X2 appliance through a pair of InfiniBand network switches. Every DGX H100 system and management node connects to each of the storage network switches via one NDR 400Gb/s InfiniBand link. Every AI400X2 appliance connects to each of the storage network switches via four InfiniBand links using the appropriate cable type. The DDN Insight server connects to the AI400X2 appliances over the 1GbE out-of-band management network. It does not require a connection to the InfiniBand storage network.

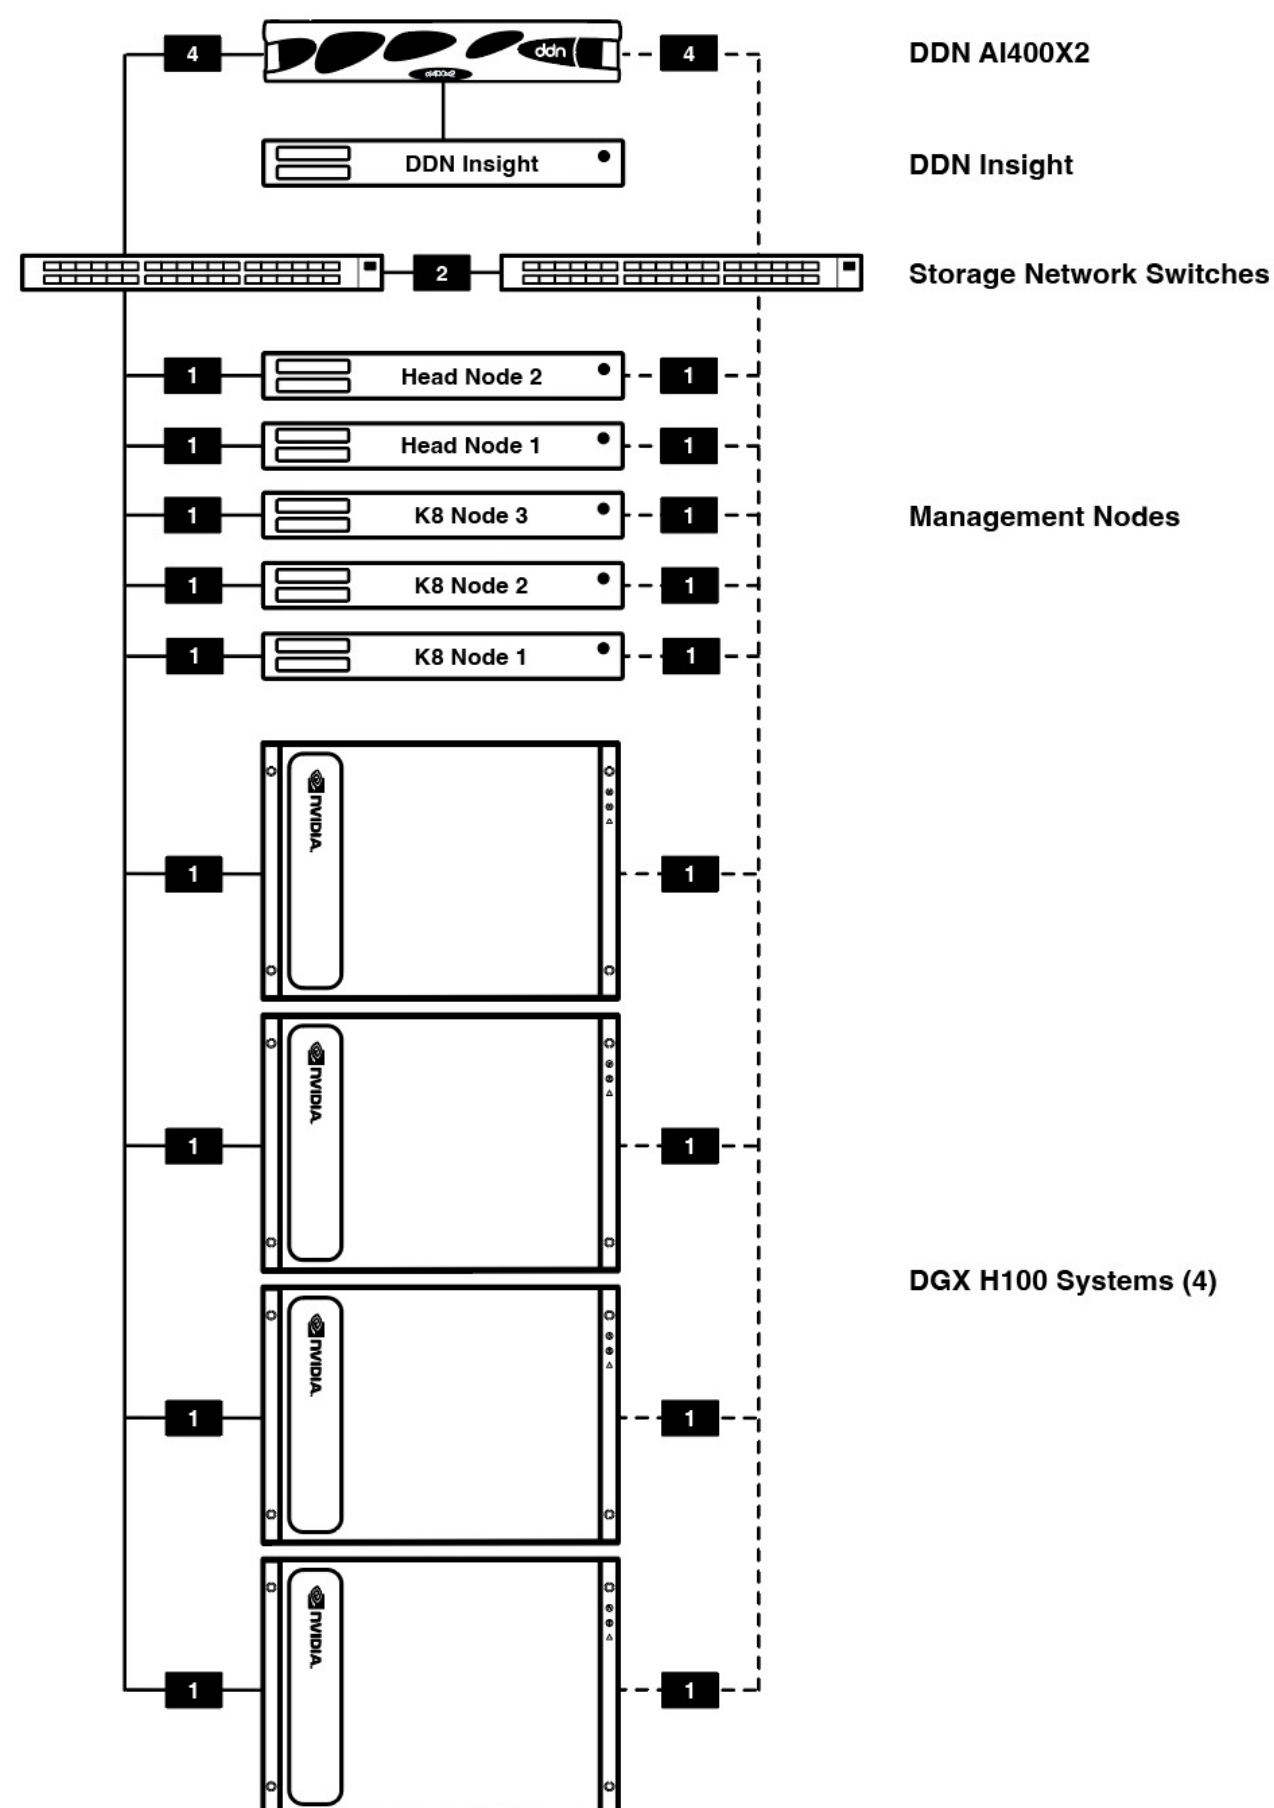


Figure 9. DDN A3I DGX BasePOD reference architecture with 4 DGX H100 systems (only storage and compute networks shown).

3.3 DGX BASEPOD WITH 8 DGX H100 SYSTEMS

Figure 10 illustrates the DDN A³I architecture for DGX BasePOD deployed with a dedicated InfiniBand storage network. In this configuration, eight DGX H100 systems are connected with an AI400X2 appliance through a pair of InfiniBand network switches. Every DGX H100 system and management node connects to each of the storage network switches via one NDR 400Gb/s InfiniBand link. The AI400X2 appliance connects to each of the storage network switches via four InfiniBand links using the appropriate cable type. The DDN Insight server connects to the AI400X2 appliances over the 1GbE out-of-band management network. It does not require a connection to the InfiniBand storage network.




Figure 10. DDN A³I DGX BasePOD reference architecture with 8 DGX H100 systems (only storage and compute networks shown).

3.4 DGX BASEPOD WITH 16 DGX H100 SYSTEMS

Figure 11 illustrates the DDN A³I architecture for DGX BasePOD deployed with a dedicated InfiniBand storage network. In this configuration, sixteen DGX H100 systems are connected with two AI400X2 appliances through a pair of InfiniBand network switches. Every DGX H100 system and management node connects to each of the storage network switches via one NDR 400Gb/s InfiniBand link. Every AI400X2 appliance connects to each of the storage network switches via four InfiniBand links using the appropriate cable type. The DDN Insight server connects to the AI400X2 appliances over the 1GbE out-of-band management network. It does not require a connection to the InfiniBand storage network.

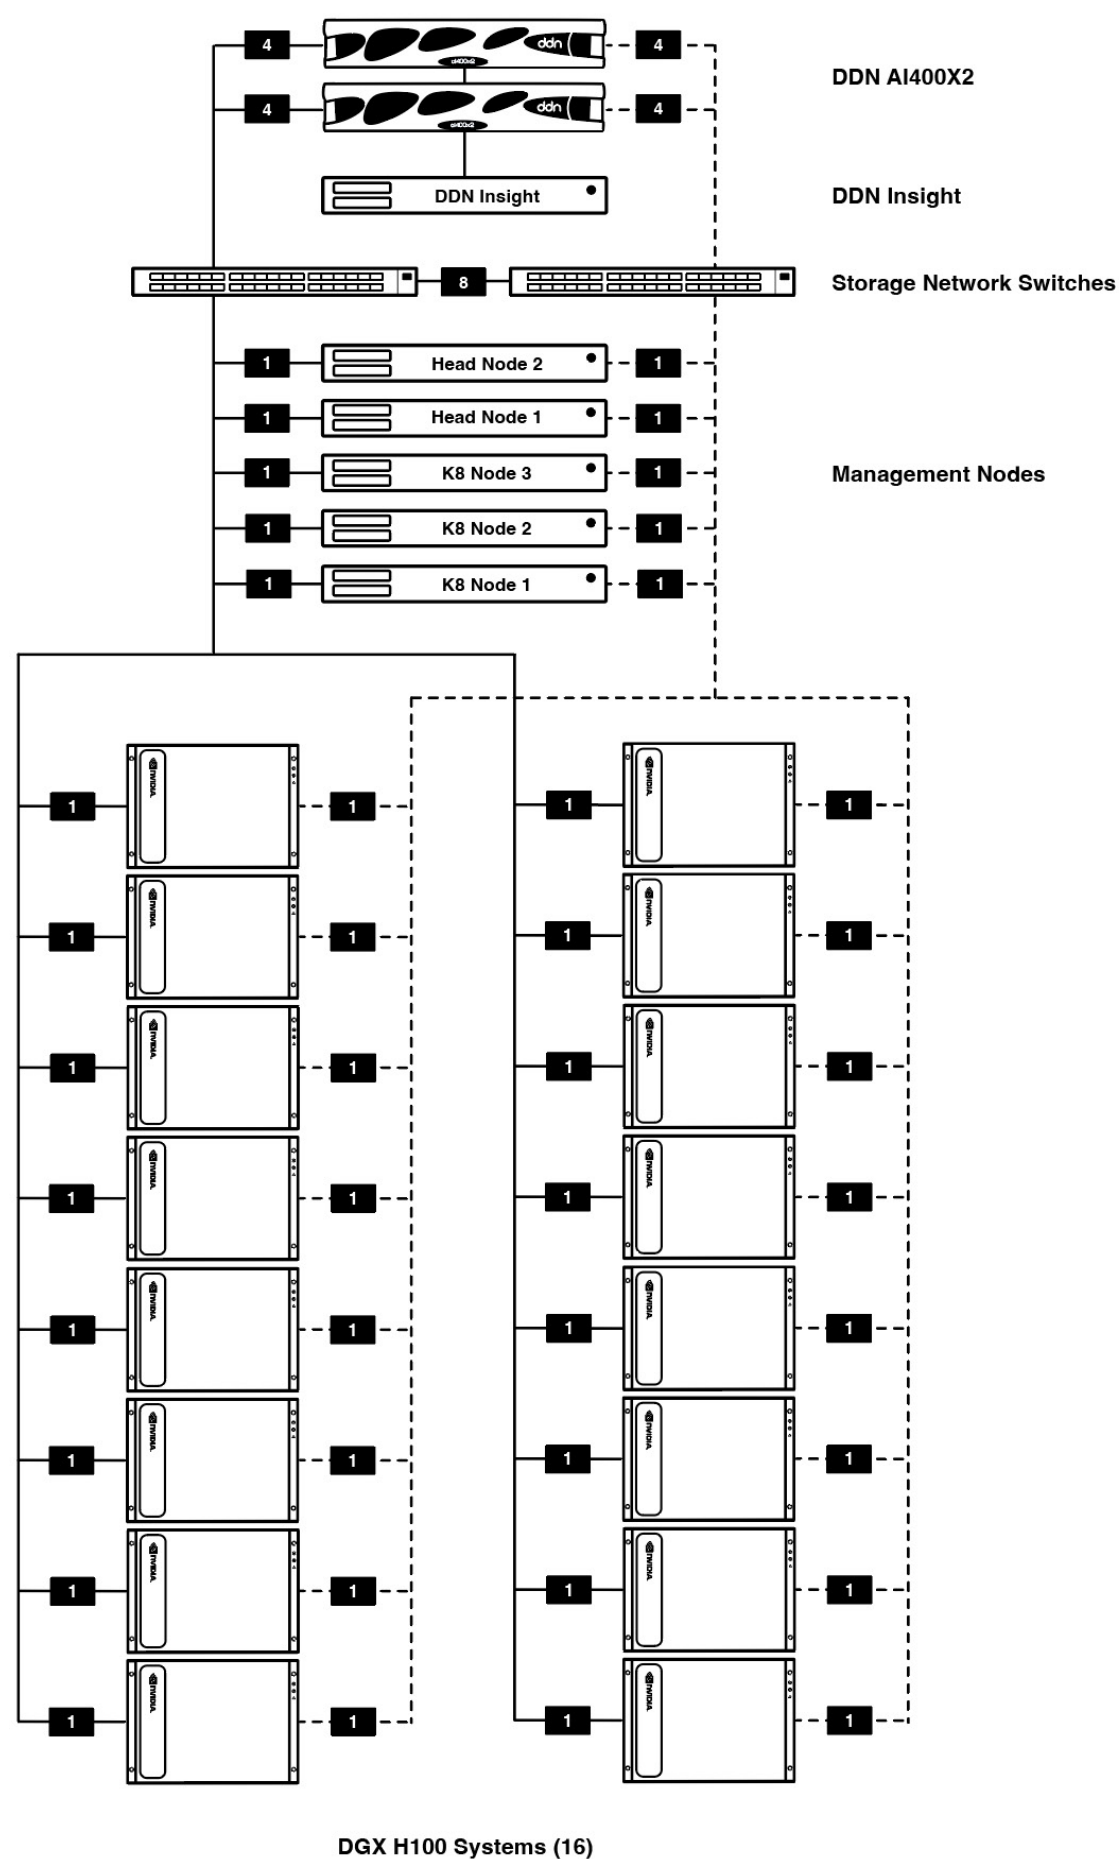


Figure 11. DDN A³I DGX BasePOD reference architecture with 16 DGX H100 systems (only storage and compute networks shown).

4. DDN A3I SOLUTIONS WITH NVIDIA DGX BASEPOD VALIDATION

DDN conducts extensive engineering integration, optimization, and validation efforts in close collaboration with NVIDIA and customers to ensure the best possible end-user experience using the reference designs in this document. The joint validation confirms functional integration, and optimal performance out-of-the-box for DGX BasePOD configurations.

Performance testing on the DDN A³I architecture is conducted with industry standard synthetic throughput and IOPS applications, as well as widely used DL frameworks and data types. The results demonstrate that with the DDN A³I shared parallel architecture, GPU-accelerated applications can engage the full capabilities of the data infrastructure and the DGX H100 systems. Performance is distributed evenly across all the DGX H100 systems in DGX BasePOD, and scales linearly as more DGX H100 systems are engaged.

This section details some of the results from recent at-scale testing integrating AI400X2 appliances with DGX H100 systems.

The tests described in this section were executed on sixteen DGX H100 systems each equipped with eight H100 GPUs running DGX OS Server Software 6.0.11 and two AI400X2 appliances running DDN EXAScaler 6.2.0.

For the storage network, all sixteen DGX H100 systems are connected to NVIDIA Quantum-2 QM9700 InfiniBand switches with two NDR 400Gb/s InfiniBand links, one per dual-ported adapter (see recommendation in section 3.1). The AI400X2 appliances are connected to the same network with eight HDR 200Gb/s InfiniBand links each.

This test environment allows us to demonstrate performance with a scalable DGX H100 BasePOD configuration using DGX H100 systems.

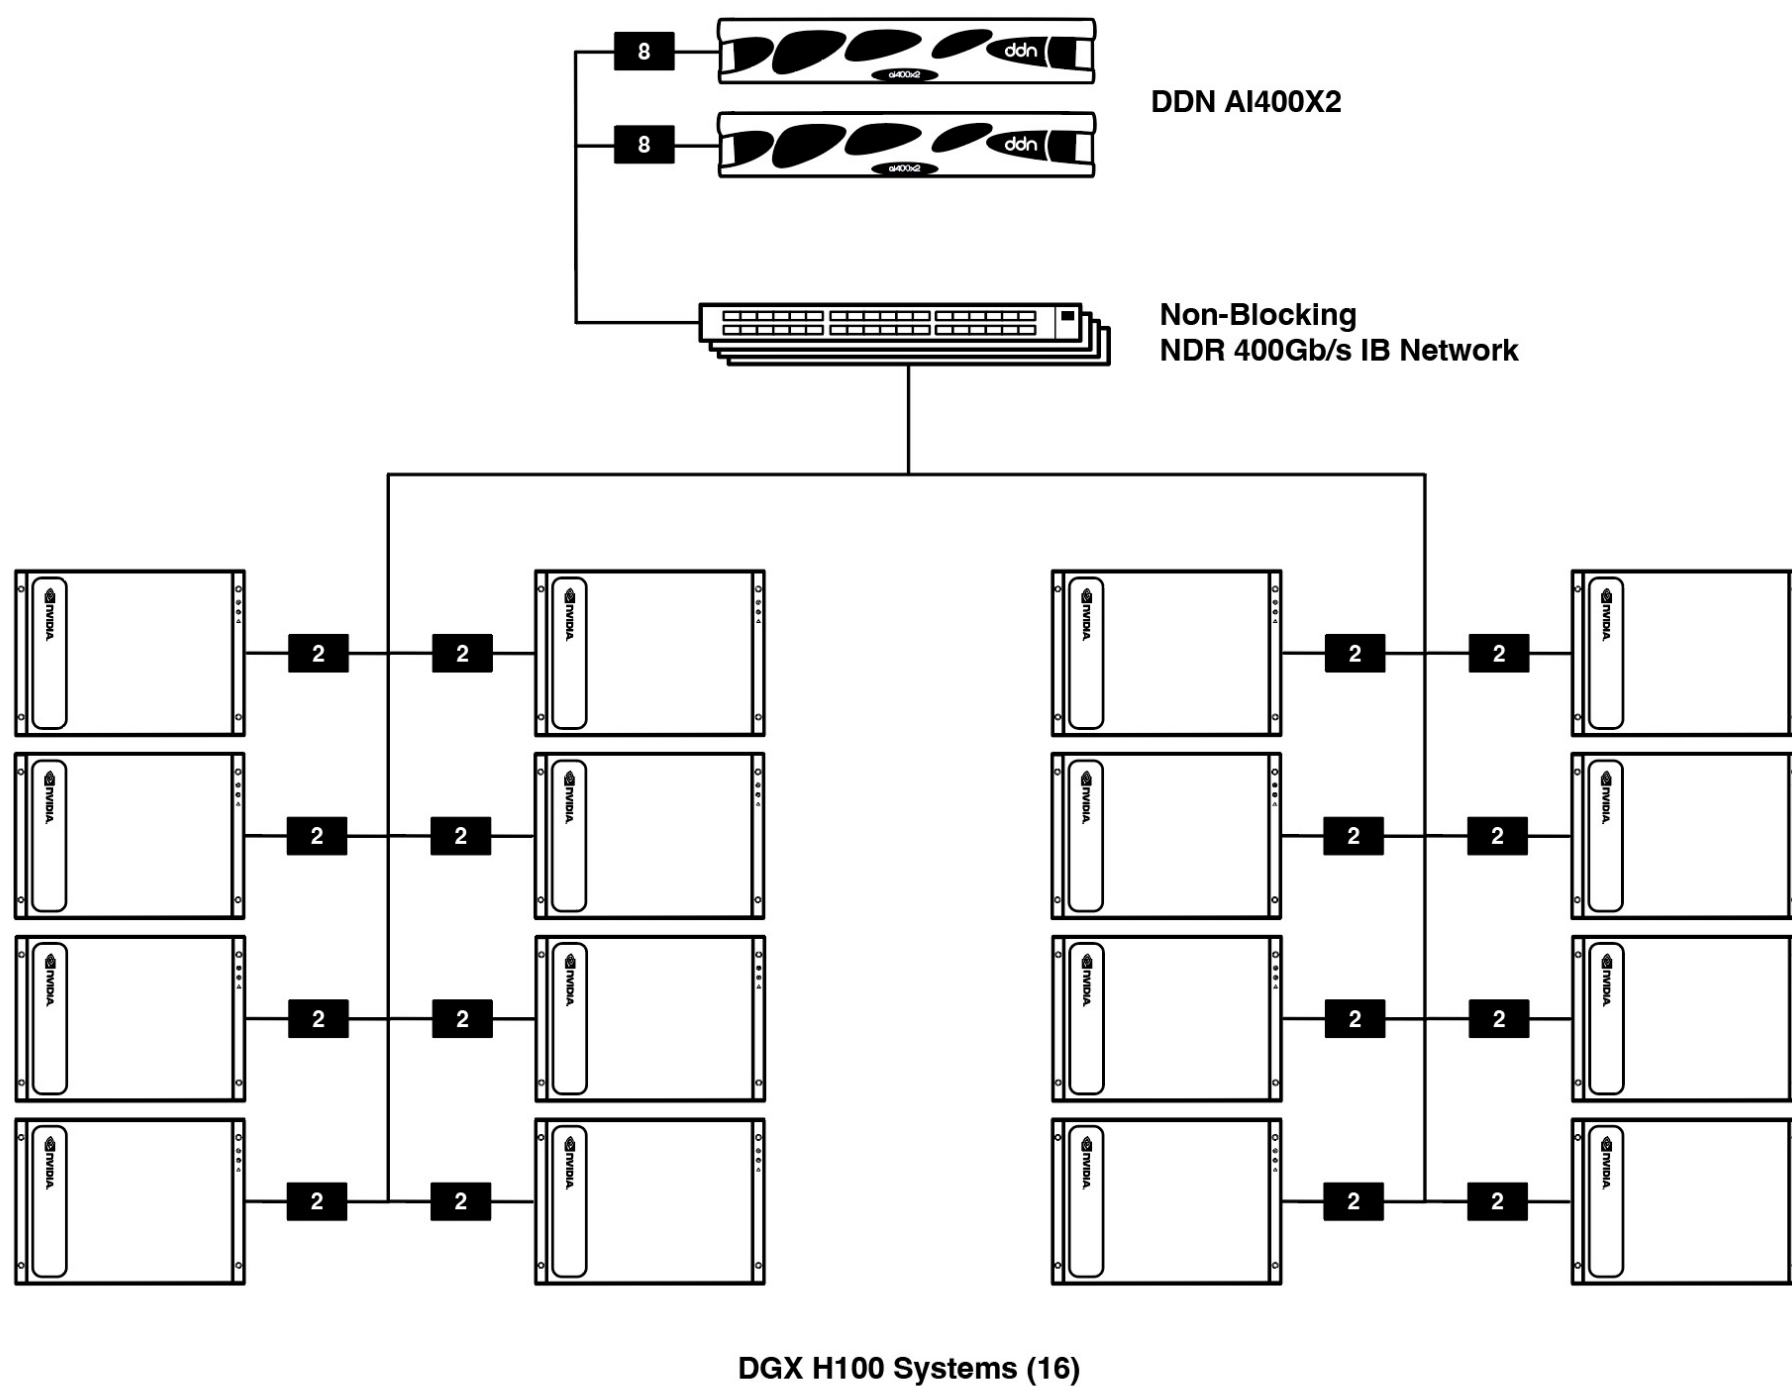


Figure 12. Test environment with sixteen DGX H100 systems and two AI400X2 (compute, in-band and out-of-band management networks not shown).

4.1 DGX BASEPOD SYSTEM PERFORMANCE VALIDATION

This first test demonstrates the peak performance of the DDN BasePOD reference architecture using the fio open-source synthetic benchmark tool. The tool is set to simulate a general-purpose workload without any performance-enhancing optimizations. Separate tests were run to measure both 100% read and 100% write workload scenarios

Below are the specific FIO configuration parameters used for these tests:

- blocksize = 1024k
- direct = 1
- iodepth= 128
- ioengine = posixaio
- bw-threads = 255

The AI400X2 appliance provides predictable, scalable performance. This test demonstrates the architecture's ability to deliver full throughput performance to a small number of clients and distribute the full performance of the DDN solution evenly as many DGX H100 systems are engaged.

In Figure 13, test results demonstrate that the DDN solution can deliver over 99 GB/s of read throughput and over 90 GB/s of write throughput to a single DGX H100 system. In this test, data is accessed through a single posix mount provided by the DDN shared parallel filesystem client installed on the DGX H100 system.





Figure 13. FIO throughput with DGX BasePOD configuration using a single DGX H100 system

In Figure 14, test results demonstrate that the DDN solution can deliver over 99 GB/s of read throughput and over 90 GB/s of write throughput to a single DGX H100 system. The DDN software evenly distributes the full read and write performance of the AI400X2 appliances with all sixteen DGX H100 systems engaged simultaneously. The DDN solution utilizes both network links on every DGX H100 system, ensuring optimal performance for a very wide range of data access patterns and data types.

FIO THROUGHPUT PERFORMANCE - PEAK PER DGX SYSTEM DGX H100 SUPERPOD CONFIGURATION WITH FOUR AI400X2



In Figure 14, test results demonstrate that the DDN solution can deliver over 99 GB/s of read throughput and over 90 GB/s of write throughput to a single DGX H100 system. The DDN software evenly distributes the full read and write performance of the AI400X2 appliances with all sixteen DGX H100 systems engaged simultaneously. The DDN solution utilizes both network links on every DGX H100 system, ensuring optimal performance for a very wide range of data access patterns and data types.

4.2 DGX BASEPOD GPU PERFORMANCE VALIDATION

This second test demonstrates IO performance scaling with 8 distinct instances of the synthetic benchmark tool running on every DGX system in the BasePOD. This validates that the reference architecture meets the recommended performance goal of at least 1 GB/s read and write throughput per DGX BasePOD using H100 GPUs.

Below are the specific FIO configuration parameters used for these tests:

- blocksize = 1024k
- direct = 1
- iodepth= 128
- ioengine = posixaio
- bw-threads = 255

In Figure 15, test results demonstrate that the DDN solution can fully distribute storage performance and delivers 1.4 GB/s of read and 1 GB/s write throughput to every H100 GPU in the DGX BasePOD used for testing. The results also demonstrate that the DDN solution delivers true linear performance scaling as all GPUs in the DGX BasePOD are engaged.

FIO THROUGHPUT PERFORMANCE - PEAK PER DGX SYSTEM DGX H100 SUPERPOD CONFIGURATION WITH FOUR AI400X2

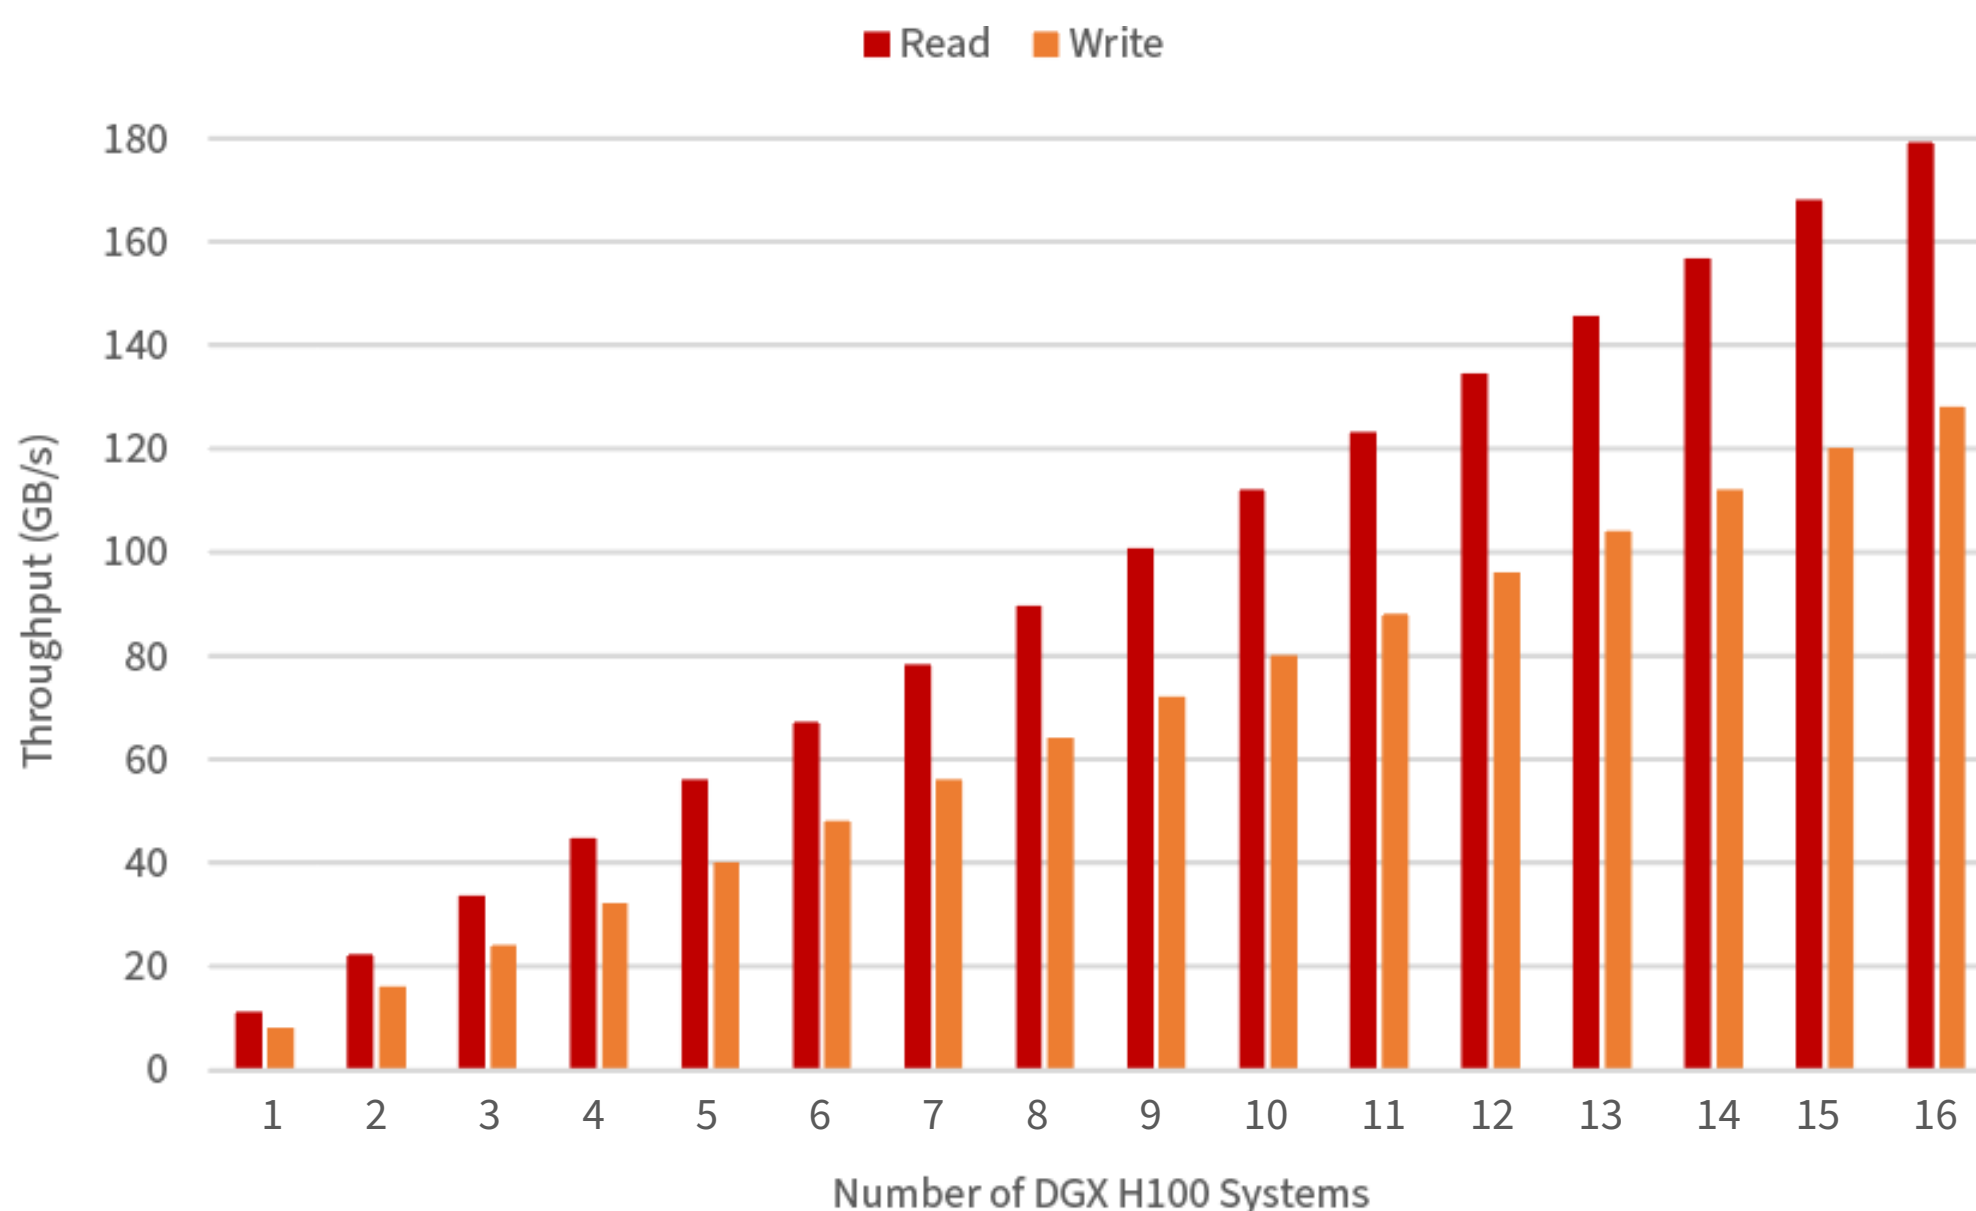


Figure 15. FIO throughput scaling with DGX BasePOD configuration, with eight distinct applications per DGX H100 system.

5. CONTACT DDN TO UNLEASH THE POWER OF YOUR DGX BASEPOD

DDN has long been a partner of choice for organizations pursuing at-scale data-driven projects. Beyond technology platforms with proven capability, DDN provides significant technical expertise through its global research and development and field technical organizations.

A worldwide team with hundreds of engineers and technical experts can be called upon to optimize every phase of a customer project: initial inception, solution architecture, systems deployment, customer support and future scaling needs.

Strong customer focus coupled with technical excellence and deep field experience ensures that DDN delivers the best possible solution to any challenge. Taking a consultative approach, DDN experts will perform an in-depth evaluation of requirements and provide application-level optimization of data workflows for a project. They will then design and propose an optimized, highly reliable and easy to use solution that best enables and accelerates the customer effort.

Drawing from the company's rich history in successfully deploying large scale projects, DDN experts will create a structured program to define and execute a testing protocol that reflects the customer environment and meet and exceed project objectives. DDN has equipped its laboratories with leading GPU compute platforms to provide unique benchmarking and testing capabilities for AI and DL applications.

Contact DDN today and engage our team of experts to unleash the power of your AI projects.