



NetApp Verified Architecture

NetApp AFF A90 Storage Systems with NVIDIA DGX SuperPOD

Design guide

David Arnette, NetApp
February 2025 | NVA-1175-DESIGN

In collaboration with



Abstract

The NVIDIA DGX SuperPOD™ with NetApp AFF® A90 storage systems combines the world-class computing performance of NVIDIA DGX systems with NetApp cloud-connected storage systems to enable data-driven workflows for machine learning (ML), artificial intelligence (AI) and high-performance technical computing (HPC). This document describes the high-level architecture of the DGX SuperPOD solution using NetApp AFF A90 storage systems with an Ethernet storage fabric.

TABLE OF CONTENTS

Executive summary	3
Program Summary	3
Solution overview	3
Target audience	4
Solution technology.....	4
Use case summary	6
Solution Architecture	7
Solution Hardware Requirements	9
Software requirements.....	9
Solution verification	9
Conclusion	10
Where to find additional information	10

LIST OF TABLES

Table 1) Hardware requirements	9
Table 2) Software requirements.....	9
Table 3) SuperPOD workload examples.....	10
Table 4) DGX SuperPOD performance guidelines.	10

LIST OF FIGURES

Figure 1) NVIDIA DGX SuperPOD with NetApp AFF A90 storage systems.....	3
Figure 2) NVIDIA Base Command Software.....	6
Figure 3) Storage fabric topology.....	7
Figure 4) pNFS, session trunking, FlexGroups and GDD.....	8

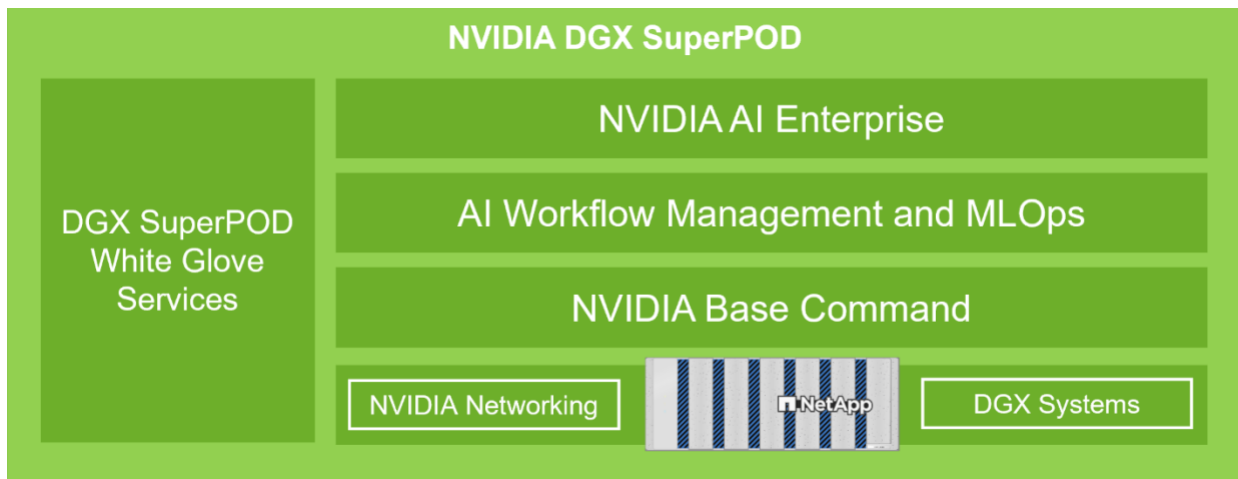
Executive summary

With the proven computing performance of NVIDIA DGX SuperPOD combined with NetApp's industry-leading data security, data governance and multi-tenancy capabilities, customers can deploy the most efficient and agile infrastructure for next-generation workloads. This document describes the high-level architecture and key features that help customers deliver faster time-to-market and return on investment for AI/ML initiatives.

Program Summary

NVIDIA DGX SuperPOD offers a turnkey AI data center solution for organizations, seamlessly delivering world-class computing, software tools, expertise, and continuous innovation. DGX SuperPOD delivers everything customers need to deploy AI/ML and HPC workloads with minimal setup time and maximum productivity. Figure 1 shows the high-level components of DGX SuperPOD.

Figure 1) NVIDIA DGX SuperPOD with NetApp AFF A90 storage systems.



DGX SuperPOD provides the following benefits:

- Proven performance for AI/ML and HPC workloads
- Integrated hardware and software stack from infrastructure management and monitoring to pre-built deep learning models and tools.
- Dedicated services from installation and infrastructure management to scaling workloads and streamlining production AI

Solution overview

As organizations embrace artificial intelligence (AI) and machine learning (ML) initiatives, the demand for robust, scalable, and efficient infrastructure solutions has never been greater. At the heart of these initiatives lies the challenge of managing and training increasingly complex AI models while ensuring data security, accessibility, and resource optimization. The evolution of agentic AI and sophisticated model training requirements has created unprecedented demands on computational and storage infrastructure. Organizations must now handle massive datasets, support multiple concurrent training workloads, and maintain high-performance computing environments while ensuring data protection and regulatory compliance. Traditional infrastructure solutions often struggle to meet these demands, leading to operational inefficiencies and delayed time-to-value for AI projects. This solution offers the following key benefits:

- **Scalability.** The NVIDIA DGX SuperPOD with NetApp AFF A90 storage systems delivers unparalleled scalability through its modular architecture and flexible expansion capabilities. Organizations can seamlessly scale their AI infrastructure by adding DGX compute nodes and AFF A90 storage systems without disrupting existing workloads or requiring complex reconfigurations.
- **Data Management and access.** The NVIDIA DGX SuperPOD with NetApp AFF A90 storage systems is based on NetApp ONTAP that excels in data management through its comprehensive suite of enterprise-grade features. Using ONTAP's snapshot and FlexClone capabilities, teams can instantly create space-efficient copies of datasets and vector databases for parallel development and testing. FlexCache and Snapmirror replication technologies enable streamlined, space-efficient and automated data pipelines from data sources across the enterprise, and multi-protocol access to data using NAS and object protocols enables new workflows optimized for ingest and data engineering tasks.
- **Security.** NetApp AFF A90 storage systems deliver enterprise-grade security through multiple layers of protection. At the infrastructure level, the solution implements robust access control mechanisms, including role-based access control (RBAC), multi-factor authentication, and detailed audit logging capabilities. The platform's comprehensive encryption framework protects data both at rest and in transit, utilizing industry-standard protocols and algorithms to safeguard intellectual property and maintain compliance with regulatory requirements. Integrated security monitoring tools provide real-time visibility into potential security threats, while automated response mechanisms help mitigate risks before they can impact operations.

Target audience

This solution is intended for organizations with HPC and AI/ML workloads that require deeper integration into broad data estates and traditional IT infrastructure tools and processes.

The target audience for the solution includes the following groups:

- IT and line of business decision makers planning for the most efficient infrastructure to deliver on AI/ML initiatives with the fastest time to market and ROI.
- Data scientists and data engineers who are interested in maximizing efficiency for critical data-focused portions of the AI/ML workflow.
- IT architects and engineers who need to deliver a reliable and secure infrastructure that enables automated data workflows and compliance with existing data and process governance standards.

Solution technology

NVIDIA DGX SuperPOD includes the servers, networking and storage necessary to deliver proven performance for demanding AI workloads. NVIDIA DGX™ H200 and NVIDIA DGX B200 systems provide world-class computing power, and NVIDIA Quantum and Spectrum™ InfiniBand network switches provide ultra-low latency and industry leading network performance. With the addition of the industry-leading data management and performance capabilities of NetApp ONTAP storage, customers can deliver on AI/ML initiatives faster and with less data migration and administrative overhead. The following sections describe the storage components of the DGX SuperPOD with AFF A90 storage systems.

NetApp AFF A90 storage systems with NetApp ONTAP

The NetApp AFF A90 powered by NetApp ONTAP data management software provides built-in data protection, anti-ransomware capabilities, and the high performance, scalability and resiliency required to support the most critical business workloads. It eliminates disruptions to mission-critical operations, minimizes performance tuning, and safeguards your data from ransomware attacks. NetApp AFF A90 systems deliver-

- **Performance.** The AFF A90 easily manages next-generation workloads like deep learning, AI, and high-speed analytics as well as traditional enterprise databases like Oracle, SAP HANA, Microsoft SQL Server, and virtualized applications. With NFS over RDMA, pNFS and session trunking,

customers can achieve the high level of network performance required for next-generation applications using existing data center networking infrastructure and industry-standard protocols with no proprietary software. Granular Data Distribution enables individual files to be distributed across every node in the storage cluster, and when combined with pNFS delivers high-performance parallel access to datasets contained in a single large file.

- **Intelligence.** Accelerate digital transformation with an AI-ready ecosystem built on data-driven intelligence, future-proof infrastructure, and deep integrations with NVIDIA and the MLOps ecosystem. Using ONTAP's snapshot and FlexClone capabilities, teams can instantly create space-efficient copies of datasets for parallel development and testing. FlexCache and Snapmirror replication technologies enable streamlined, space-efficient and automated data pipelines from data sources across the enterprise. And multi-protocol access to data using NAS and object protocols enables new workflows optimized for ingest and data engineering tasks. Data and training checkpoints can be tiered to lower-cost storage to avoid filling primary storage. Customers can seamlessly manage, protect, and mobilize data, at the lowest cost, across hybrid cloud with a single storage OS and the industry's richest data services suite.
- **Security.** The NVIDIA DGX SuperPOD with NetApp ONTAP Storage delivers enterprise-grade security through multiple layers of protection. At the infrastructure level, the solution implements robust access control mechanisms, including role-based access control (RBAC), multi-factor authentication, and detailed audit logging capabilities. The platform's comprehensive encryption framework protects data both at rest and in transit, utilizing industry-standard protocols and algorithms to safeguard intellectual property and maintain compliance with regulatory requirements. Integrated security monitoring tools provide real-time visibility into potential security threats, while automated response mechanisms help mitigate risks before they can impact operations. NetApp ONTAP is the only hardened enterprise storage that's validated to store top-secret data.
- **Multi-tenancy.** NetApp ONTAP delivers the widest array of features to enable secure multi-tenant usage of storage resources. Storage Virtual Machines provide tenant-based administrative delegation with RBAC controls. Comprehensive QoS controls guarantee performance for critical workloads while enabling maximum utilization, and security features such as tenant-managed keys for volume-level encryption guarantee data security on shared storage media.
- **Reliability.** NetApp eliminates disruptions to mission-critical operations through advanced reliability, availability, serviceability, and manageability (RASM) capabilities, delivering the highest uptime available. For more information see the [ONTAP RASS whitepaper](#). In addition, system health can be optimized with AI-based predictive analytics delivered by Active IQ and Data Infrastructure Insights.

NVIDIA DGX B200 Systems

NVIDIA DGX™ B200 is a unified AI platform for develop-to-deploy pipelines for businesses of any size at any stage in their AI journey. Equipped with eight NVIDIA Blackwell GPUs interconnected with fifth-generation [NVIDIA NVLink™](#), DGX B200 delivers leading-edge performance, offering 3X the training performance and 15X the inference performance of previous generations. Leveraging the [NVIDIA Blackwell architecture](#), DGX B200 can handle diverse workloads—including large language models, recommender systems, and chatbots—making it ideal for businesses looking to accelerate their AI transformation.

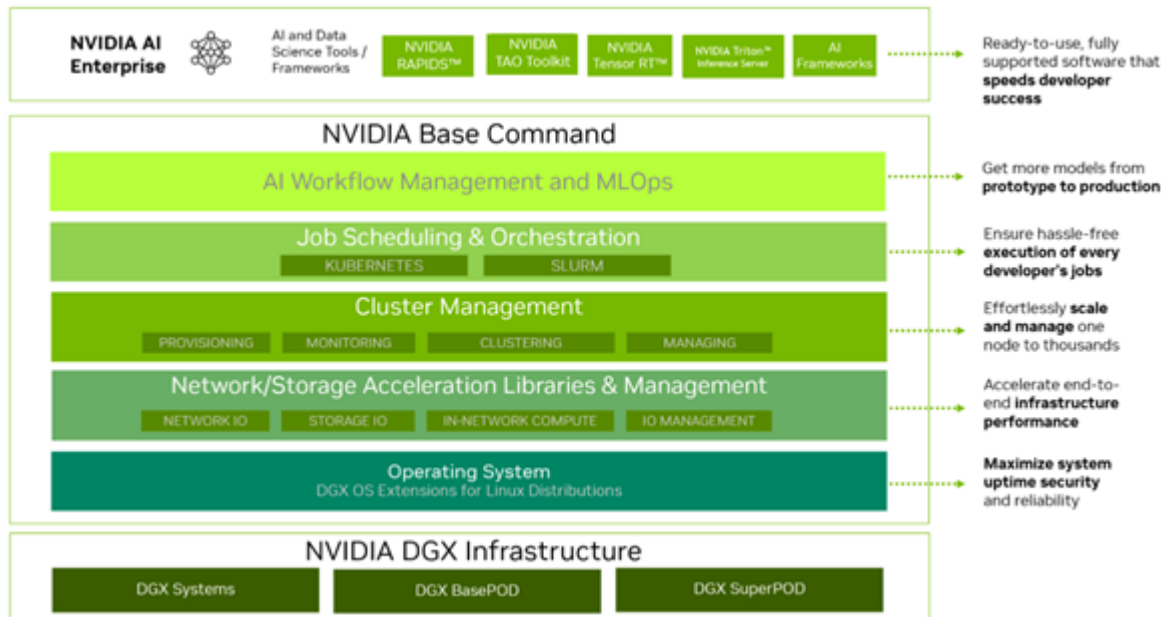
NVIDIA Spectrum SN5600 Ethernet switches

The SN5600 smart-leaf, spine, and super-spine switch offers 64 ports of 800GbE in a dense 2U form factor. The SN5600 enables both standard leaf/spine designs with top-of-rack (ToR) switches as well as end-of-row (EoR) topologies. The SN5600 offers diverse connectivity in combinations of 1 to 800GbE and boasts an industry-leading total throughput of 51.2Tb/s.

NVIDIA Base Command software

NVIDIA Base Command™ powers the NVIDIA DGX platform, enabling organizations to leverage the best of NVIDIA AI innovation. With it, every organization can tap the full potential of their DGX infrastructure with a proven platform that includes AI workflow management, enterprise-grade cluster management, libraries that accelerate compute, storage, and network infrastructure, and system software optimized for running AI workloads. Figure 2 shows the NVIDIA Base Command software stack.

Figure 2) NVIDIA Base Command Software.



NVIDIA Base Command Manager

NVIDIA Base Command Manager offers fast deployment and end-to-end management for heterogeneous AI and high-performance computing (HPC) clusters at the edge, in the data center, and in multi- and hybrid-cloud environments. It automates provisioning and administration of clusters ranging in size from a couple of nodes to hundreds of thousands, supports NVIDIA GPU-accelerated and other systems, and enables orchestration with Kubernetes. Integrating NetApp AFF A90 storage systems with DGX SuperPOD requires minimal configuration of Base Command Manager to system tuning and mount parameters for optimal performance, but no additional software is required to deliver highly available multi-path access between DGX systems and the AFF A90 storage system.

Use case summary

NVIDIA DGX SuperPOD is designed to meet the performance requirements of the most demanding workloads at the largest scale.

This solution applies to the following use cases:

- Machine learning at massive scale using traditional analytics tools.
- Artificial intelligence model training for Large Language Models, computer vision/image classification, fraud detection and countless other use cases.
- High performance computing such as seismic analysis, computational fluid dynamics and large-scale visualization.

Solution Architecture

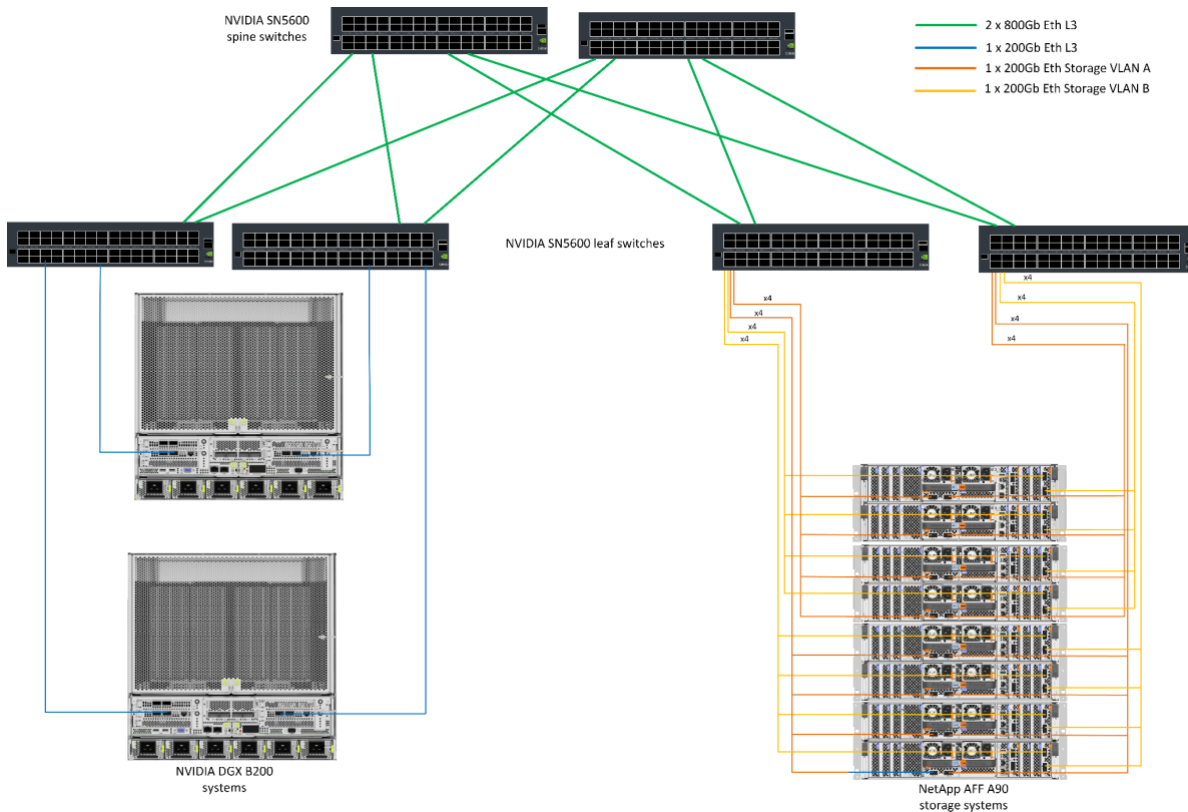
DGX SuperPOD is based on the concept of a Scalable Unit (SU) that includes 32 DGX B200 systems and all of the other components necessary to deliver the required connectivity and eliminate any performance bottlenecks in the infrastructure. Customers can start with a single or multiple SUs and add additional SUs as needed to meet their requirements. This document describes the storage configuration for a single SU, and Table 1 shows the components needed for larger configurations.

The DGX SuperPOD reference architecture includes multiple networks, and the AFF A90 storage system is connected to several of them. For more information about DGX SuperPOD networking please refer to the [NVIDIA DGX SuperPOD Reference Architecture](#).

For this solution the high-performance storage fabric is an Ethernet network based on the NVIDIA Spectrum SN5600 switch with 64 800Gb ports in a Spine/Leaf configuration. The In-band network provides user access for other functions such as home directories and general file shares and is also based on SN5600 switches, and the out of band (OOB) network is for device-level system administrator access using SN2201 switches.

The storage fabric is a leaf-spine architecture where the DGX systems connect to one pair of leaf switches and the storage system connects to another pair of leaf switches. Multiple 800Gb ports are used to connect each leaf switch to a pair of spine switches, creating multiple high-bandwidth paths through the network for aggregate performance and redundancy. For connectivity to the AFF A90 storage system, each 800Gb port is broken into four 200Gb ports using the appropriate copper or optical breakout cables. To support clients mounting the storage system with NFS over RDMA the storage fabric is configured for RDMA over Converged Ethernet (RoCE), which guarantees lossless packet delivery in the network. Figure 3 shows the storage network topology of this solution.

Figure 3) Storage fabric topology.

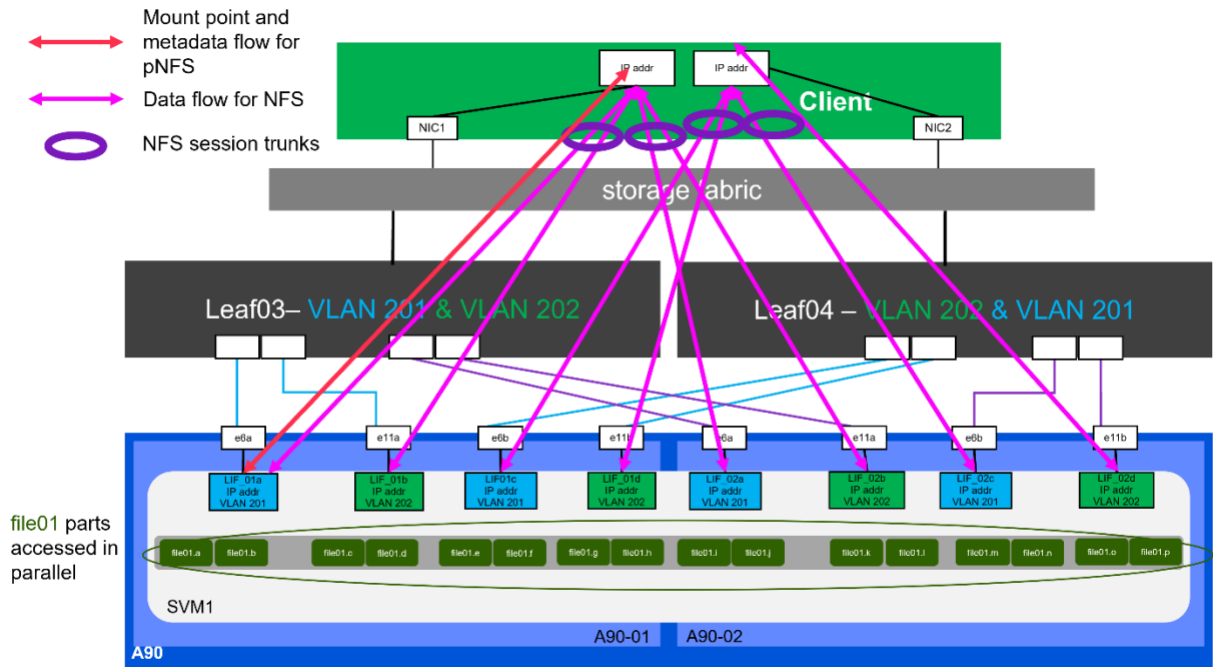


The NetApp AFF A90 storage system is a 4RU chassis containing 2 controllers that operate as high availability partners (HA Pair) for each other, with up to 48 2.5 inch form-factor solid state disks (SSD). Each controller is connected to both the SN5600 storage leaf switches using four 200Gb Ethernet connections, and there are 2 logical IP interfaces on each physical port. The storage cluster supports NFS v4.1 with Parallel NFS (pNFS) that enables clients to establish connections directly to every controller in the cluster. Additionally, session trunking combines the performance from multiple physical interfaces into a single session, enabling even single-threaded workloads to access more network bandwidth than is possible with traditional ethernet bonding. Combining all of these features with RDMA enables the AFF A90 storage system to deliver low latency and high throughput that scales linearly for workloads leveraging NVIDIA GPUDirect Storage™.

For connectivity to the in-band network the AFF A90 controllers have additional 200Gb Ethernet interfaces configured in a LACP interface group providing general NFS v3 and v4 services as well as S3 access to shared filesystems if desired. All of the controllers and storage cluster switches are connected to the OOB network for remote administrative access.

To enable high performance and scalability, the storage controllers form a storage cluster that enables the entire performance and capacity of the cluster nodes to be combined into a single namespace called a FlexGroup with data distributed across the disks of every node in the cluster. With the new Granular Data Distribution feature released in ONTAP 9.16.1, individual files are separated and distributed across the FlexGroup to enable the highest levels of performance for single-file workloads. Figure 4 below shows how pNFS and NFS session trunking work together with FlexGroups and GDD to enable parallel access to large files leveraging every network interface and disk in the storage system.

Figure 4) pNFS, session trunking, FlexGroups and GDD.



5

© 2024 NetApp, Inc. All rights reserved. — NETAPP CONFIDENTIAL —

This solution leverages multiple Storage Virtual Machines (SVM) to host volumes for both high-performance storage access as well as user home directories and other cluster artifacts on a management SVM. Each SVM is configured with network interfaces and FlexGroup volumes and QoS policy is implemented to ensure performance for the Data SVM. For more information on FlexGroups, Storage Virtual Machines and ONTAP QoS capabilities please refer to the [ONTAP documentation](#).

Solution Hardware Requirements

Table 1 lists the storage hardware components that are required to implement one, two, four or eight scalable units. For detailed hardware requirements for servers and networking please see the [NVIDIA DGX SuperPOD Reference Architecture](#).

Table 1) Hardware requirements.

SU Size	AFF A90 systems	Storage Cluster Interconnect switches	Usable capacity (typical with 3.8TB SSD)	Max useable capacity (with 15.3TB NVMe SSD)	RU (typical)	Power (typical)
1	4	2	555TB	13.75PB	18	7,300 watts
2	8	2	1PB	27.5PB	34	14,600 watts
4	16	2	2PB	55PB	66	29,200 watts
8	32	4	4PB	110PB	102	58,400 watts

NOTE: NetApp recommends a minimum of 24 drives per AFF A90 HA pair for maximum performance. Additional internal drives, larger capacity drives and external expansion drive shelves enable much higher aggregate capacity with no impact to system performance.

Software requirements

Table 2 lists the software components and versions that are required to integrate the AFF A90 storage system with DGX SuperPOD. DGX SuperPOD also involves other software components that are not listed here. Please refer to the [DGX SuperPOD release notes](#) for complete details.

Table 2) Software requirements.

Software	Version
NetApp ONTAP	9.16.1
NVIDIA BaseCommand Manager	10.24.11
NVIDIA DGX OS	6.3.1
NVIDIA OFED driver	MLNX_OFED_LINUX-23.10.3.2.0 LTS
NVIDIA Cumulus OS	5.10

Solution verification

This storage solution was validated in multiple stages by NetApp and NVIDIA to ensure that performance and scalability meet the requirements for NVIDIA DGX SuperPOD. The configuration was validated using a combination of synthetic workloads and real-world ML/DL workloads to verify both maximum performance and application interoperability. Table 3 below provides examples of typical workloads and their data requirements that are commonly seen in DGX SuperPOD deployments.

Table 3) SuperPOD workload examples.

Level	Work Description	Data Set Size
Standard	Multiple concurrent LLM or fine-tuning training jobs and periodic checkpoints, where the compute requirements dominate the data I/O requirements significantly.	Most datasets can fit within the local compute systems' memory cache during training. The datasets are single modality, and models have millions of parameters.
Enhanced	Multiple concurrent multimodal training jobs and periodic checkpoints, where the data I/O performance is an important factor for end-to-end training time.	Datasets are too large to fit into local compute systems' memory cache requiring more I/O during training, not enough to obviate the need for frequent I/O. The datasets have multiple modalities and models have billions (or higher) of parameters.

Table 4 shows performance guidelines for the example workloads above. These values represent the storage throughput that can be generated by these workloads under ideal conditions.

Table 4) DGX SuperPOD performance guidelines.

Performance Characteristic	Standard (GBps)	Enhanced (GBps)
Single SU aggregate system read	40	125
Single SU aggregate system write	20	62
4 SU aggregate system read	160	500
4 SU aggregate system write	80	250

Conclusion

The NVIDIA DGX SuperPOD with NetApp **AFF A90 storage systems** represents a significant advancement in AI infrastructure solutions. By addressing key challenges around security, data management, resource utilization, and scalability, it enables organizations to accelerate their AI initiatives while maintaining operational efficiency, data protection, and collaboration. The solution's integrated approach eliminates common bottlenecks in AI development pipelines, enabling data scientists and engineers to focus on innovation rather than infrastructure management.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- [NVA-1175 NVIDIA DGX SuperPOD with NetApp AFF A90 Storage Systems Deployment Guide](#)
- [NVIDIA DGX B200 SuperPOD Reference Architecture](#)
- [NVIDIA DGX H200 SuperPOD Reference Architecture](#)
- [NVIDIA BaseCommand Software](#)
- [NVIDIA Spectrum SN5600 Ethernet switches](#)
- [NVIDIA DGX SuperPOD release notes](#)
- [NetApp Documentation](#)
- [NetApp AI Solutions Documentation](#)
- [NetApp ONTAP software](#)
- [NetApp Install and Maintain AFF Storage Systems](#)
- [NFS over RDMA](#)
- [What is pNFS](#) (older doc with great pNFS info)

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright information

Copyright © 2025 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.


The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data—Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

Händlerinformation INNEO Solutions GmbH inneo@inneo.com inneo.com	
Deutschland: IT-Campus 1 73479 Ellwangen +49 (0) 7961 890-0	Schweiz: Ruchstuckstrasse 21 8306 Brüttsellen +41 (0) 44 805 1010