

IBM Storage Reference Architecture with NVIDIA DGX A100 Systems

Version 2 Release 1

Douglas O'Flaherty
Sanjay Sudam

Content

Content.....	ii
About this document	3
Scope	3
Prerequisites.....	3
Introduction.....	4
The Artificial Intelligence Ladder	5
Solution Components.....	6
Solution Validation	11
Solution Sizing Guidance	18
Conclusion	20
Get more Information	20
Authors	21
Acknowledgements.....	21
Notices.....	21



About this document

This document is intended to facilitate the deployment and configuration of the NVIDIA DGX™ A100 systems using IBM Spectrum Scale, IBM Elastic Storage System (ESS) 3000 and IBM Elastic Storage System (ESS) 3200. It is the latest in the series of IBM and NVIDIA DGX Reference Architectures for each generation of NVIDIA DGX systems.

The information in this document is distributed on an “as is” basis without any warranty that is either expressed or implied. Support assistance for the use of this material is limited to situations where IBM Spectrum Scale® and/or IBM ESS 3000 are supported and entitled and where the issues are specific to a blueprint implementation.

Scope

This blueprint guide provides a solutions architecture and related solution configuration parameters with the following essential components:

- IBM Spectrum Scale
- IBM Elastic Storage System 3000 and 3200
- NVIDIA DGX A100 Systems
- NVIDIA InfiniBand (IB) Switches

This blueprint does not:

- Replace any official manuals and documents issued by IBM
- Replace any official manuals and documents issued by NVIDIA

The NVIDIA DGX POD™ with IBM storage reference architecture does not describe data and metadata integration of data ingest, pre-processing, archiving, tiering, and legacy data stores. Though these use-cases are supported by IBM Spectrum Scale, describing them is not in the scope of this document.

Prerequisites

This technical paper assumes the following prerequisites:

- Basic knowledge of IBM Elastic Storage Server (ESS)3000 series
- Basic knowledge of IBM Spectrum Scale
- Basic knowledge of DGX A100 systems
- Basic knowledge of NVIDIA switches
- Basic knowledge of InfiniBand Networking Technologies

Introduction

The adoption of GPU-accelerated computing infrastructure has become imperative to support deep learning (DL) and other computationally intensive workloads. GPUs enable working with large data sets and real-time streaming that is not possible on other platforms. However, accelerated computing needs to be matched with storage and networking that is also capable of large, diverse, and real-time data management.

The IBM Storage with NVIDIA DGX A100 systems reference architecture provides a prescriptive solution for IT administrators to deploy a validated solution with shared, extensible storage that is designed for DL, inference, data exploration, and other computationally and I/O-intensive work.

For the IT administrator, this reference architecture provides the specifications for networking, storage, and infrastructure that have been proven to enable improved scalability, performance, and cost-effective manageability.

For the data scientist, team productivity, data reuse, and logical data locality can be enhanced through shared storage that can integrate with and enhance the DL workflow.

Designed as a scalable unit that can be scaled to meet growing business and research needs, the IBM Storage with NVIDIA DGX A100 systems supports organizations looking to satisfy the needs for multiple workloads, from providing “as-a-Service” access for small interactive jobs, to supporting cluster-wide jobs that make the full use of multi-GPU and multi-node resources. IBM Storage provides a storage framework with scalability, extensibility, and the enterprise attributes required, including data protection, data tiering, and hybrid cloud integration.

The Artificial Intelligence Ladder

Data is the fuel for artificial intelligence (AI). The best AI is built on a foundation of data that is collected and organized as carefully as it is analyzed and then infused into the business. Organizations are challenged with gaining insights from their data for many reasons. Data silos make it difficult to get a holistic view of all your information, limiting the value of AI. Current infrastructure that was not built for AI is not flexible enough to respond to new demands without adding complexity.

Every successful AI project goes through a multi-step process (Figure 1) that starts with having the right data and progresses to using AI broadly.

Collect

Data is the fuel that powers AI, but it can become trapped or stored in a way that makes it difficult or cost prohibitive to maintain or expand. IBM Storage makes data simple and accessible for a hybrid cloud infrastructure.



Figure 1: The AI Ladder

Organize

AI can only be as good as the data it relies on. Businesses must fully understand what data they have so they can leverage it for AI and other organizational needs, including compliance, data optimization, data cataloging and data governance.

Analyze

Organizations must plan for opportunities beyond the development of AI; they need to build AI infrastructures with confidence in scalability, high performance, and extensibility to the data lake and the storage catalog.

Infuse

Business challenges can become an opportunity to explore, understand, predict, and bring an AI infrastructure to your entire organization. IBM Storage is empowering customers to use data and AI storage to leverage that infrastructure in more ways that bring value to the organization.

Solution Components

IBM Storage for Data and AI

IBM Storage for Data and AI covers the complete use case for Big Data and AI applications across hybrid cloud. The integrated approach provides a management framework and unified view of enterprise and research data across multiple data sources and applications. With the ability to support the entire AI ladder, IBM reduces costs and enables data scientists to rapidly identify and use relevant data. Highly flexible and extremely scalable, IBM Storage for Data and AI may be deployed on and integrate with multiple platforms and clouds.

There are three independent, yet integrated offerings in the software-defined portfolio.

- **IBM Spectrum Scale** provides scalable file storage, multi-protocol access, and data peering with IBM COS.
- **IBM Cloud Object Storage** (IBM COS) is an efficient and secure data lake for object applications or backing multiple services.
- **IBM Spectrum Discover** provides the data catalogue and custom data tagging across platforms.

IT and end-users can use a programmatic interface for data discovery, storage efficiency, and workflow optimization. IBM Storage for Data and AI provides multiple deployment options with IBM solutions, software-defined storage, or on public cloud providers such as IBM Cloud, AWS, Azure and Google.

IBM Spectrum Scale, IBM COS, and IBM Spectrum Discover provide the data scientist and the IT administrator performance, optimized data tiering, and data tracking that can reduce costs and improve productivity. Together, they represent the Storage for AI and Big Data part of the IBM storage portfolio that covers primary block storage with all-flash arrays, modern data protection, and software-defined storage. (Figure 2).

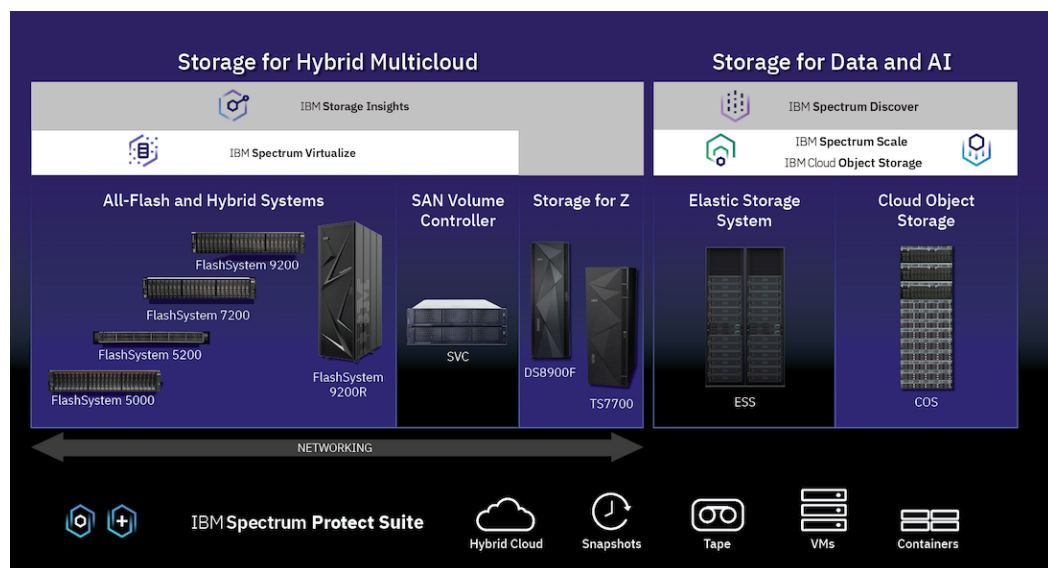


Figure 2: IBM Storage Portfolio

IBM Spectrum Scale

IBM Spectrum Scale is an industry leader in high-performance file systems. The underlying parallel filesystem provides scalable throughput and low latency data access, as well as superior metadata performance. Unlike NFS systems that can easily bottleneck, the distributed architecture of a parallel filesystem provides reliable performance for multi-user sequential and random read or write. IBM Spectrum Scale has been proven on the largest AI clusters in the world, including the US National Labs supercomputers Summit and Sierra, as well as the Circe supercomputer built by NVIDIA.

IBM Spectrum Scale v5 has been optimized for modern AI workloads. Remarkably fast directory and file metadata management is required for the many AI/ML workloads that distribute data across multiple directories or many files.

The latest release of IBM Spectrum Scale v5.1.1 provides Tech Preview support for NVIDIA GPUDirect Storage (GDS). GDS provides allows CUDA applications to perform remote memory reads directly from storage to the GPU and bypassing the system memory. By eliminating the bounce through the CPU memory, GDS reduces latency, increases effective bandwidth, and offloads the CPU for greater system efficiency.

With IBM Spectrum Scale, a DGX POD deployment may share data with existing infrastructure, such as an HPC cluster, Hadoop/Spark, or a structured data source. IBM Spectrum Scale creates a single namespace (or data plane) across systems. For users, it is a single repository that can access NFS, SMB, Object, or high-performance native POSIX filesystem. This single data plane allows the data administrator, analyst, or data scientist to access all the data in place. The entire data pipeline, from ingest to inference can be completed without having to make additional copies or move data between systems. Multiple clusters can be integrated into a single namespace to provide rapid local access to logically or geographically distributed data.

Although not in the scope of this Reference Architecture, an DGX POD may also be backed up by other storage. IBM Spectrum Scale enables data to be tiered automatically and transparently to and from more cost-effective storage, including hard disk drive (HDD), tape, and cloud. The software provides automated, policy driven, or workload directed data tiering to optimize performance and costs. When deployed with IBM Cloud Object Storage, or AWS S3, IBM Spectrum Scale has high-performance object tiering for transparent bi-directional data sharing. Data may be written or read as objects and still be available to be run on a DGX POD with IBM Spectrum Scale.

IBM Spectrum Scale provides Container Native Access and Operators to support Kubernetes driven DevOps and Data Ops practices. In addition, IBM Spectrum Scale provides enterprise features such as call-home proactive support, encryption, and audit file logging that works with Qradar and Splunk SEIM platforms.

IBM Elastic Storage System (ESS) 3200

IBM recently updated the IBM Elastic Storage System all-flash solution with faster

processing and updated networking. The IBM Elastic Storage System 3200 (Figure 3) combines the performance of NVMe storage technologies with the reliability and the rich features of IBM Spectrum Scale, along with several high-speed attachment options such as 100 Gb/s Ethernet and InfiniBand (IB) — all in a powerful 2U storage system.



Figure 3: IBM Elastic Storage System 3200

IBM Spectrum Scale on NVMe is designed to be the market leader in all-flash performance, and scalability with a bandwidth of around 80 GB/s per NVMe all-flash appliance and 100 microseconds latency. This doubles the throughput of the IBM ESS 3000. Providing data-driven multicloud storage capacity, the NVMe all-flash appliance is deeply integrated with the software defined capabilities of IBM Storage for Data and AI to seamlessly plug it into an analytics, scalable cluster, or AI workload.

Available with multiple drive options and advanced erasure coding, the IBM ESS 3200 provides options to optimize costs for different installation sizes. As with all IBM Spectrum Scale solutions, capacity and performance can be scaled. Combining ESS 3200 systems provides nearly linear performance scalability. IBM ESS 3200 solutions may also be used as an all-flash NVMe performance tier on slower storage, including tape or object storage. Table 1 shows some key specifications of the IBM ESS 3000 and IBM ESS 3200 systems.

	IBM ESS 3000	IBM ESS 3200
System Features	<ul style="list-style-type: none"> • Dual 2-socket Storage Controllers, Active/Active • 384 GB or 768 GB memory per controller • De-Clustered RAID supporting erasure coding schemas: 3-way replication, 4-way replication, 4+2P, 4+3P, 8+2P, 8+3P 	<ul style="list-style-type: none"> • Dual 2-socket Storage Controllers, Active/Active • 384 GB or 768 GB memory per controller • De-Clustered RAID supporting erasure coding schemas: 3-way replication, 4-way replication, 4+2P, 4+3P, 8+2P, 8+3P
Performance	<ul style="list-style-type: none"> • Sequential read performance up to 42G B/s • Sequential write performance up to 32 GB/s 	<ul style="list-style-type: none"> • Sequential read performance up to 80G B/s • Sequential write performance up to 58 GB/s
Networking	<ul style="list-style-type: none"> • EDR IB, up to 12 ports • 100G Ethernet, up to 12 ports 	<ul style="list-style-type: none"> • HDR IB, up to 12 ports • 100G Ethernet, up to 12 ports
Drive Support	12 or 24 NVMe SSDs (1.92 TB, 3.84 TB, 7.68 TB or 15.36 TB)	12 or 24 NVMe SSDs (1.92 TB, 3.84 TB, 7.68 TB or 15.36 TB)

Table 1: IBM Elastic Storage System 3000 Specifications

For detailed ESS3200 specifications, please refer the link <https://www.ibm.com/products/elastic-storage-system>

NVIDIA DGX A100 System

The DGX A100 system (Figure 4) is the universal system for all AI workloads, offering unprecedented compute density, performance, and flexibility in the world's first 5 peta FLOPS AI system. The DGX A100 system features the world's most advanced accelerator, the NVIDIA A100 Tensor Core GPU, enabling enterprises to consolidate training, inference, and analytics into a unified, easy-to-deploy AI infrastructure.




Figure 4: DGX A100 system

NVIDIA Mellanox Quantum QM8700 InfiniBand Switch

The InfiniBand converged fabric used for both compute and storage utilizes NVIDIA® Mellanox Quantum™ HDR 200 Gb/s InfiniBand Smart Switches (Figure 5).



Figure 5: NVIDIA Mellanox QM8700 HDR 200 Gb/s InfiniBand Switch

Each switch includes 40 QSFP56 ports, which are used for both communication to each DGX A100 system, as well as between switches in the compute fabric. All connections are HDR, maximizing the bandwidth between network elements. No InfiniBand partitions, or other segmentation is used, with the QM8700 switch providing the subnet manager for the compute fabric. Connection to the out-of-band management ports on the switch to the out-of-band management fabric may be done if needed, but is not critical to the operation of the DGX POD

NVIDIA NGC

The NVIDIA NGC™ (Figure 6) provides a range of options that meet the needs of data scientists, developers, and researchers with various levels of AI expertise. These users can quickly deploy AI frameworks with containers, get a head start with pre-trained models or model training scripts, and use domain specific workflows and Helm charts for the fastest AI implementations, giving them faster time-to-solution.

Spanning AI, data science, and HPC, the container registry on NGC features an extensive range of GPU-accelerated software for NVIDIA GPUs. The NGC hosts containers for the top AI and data science software. Containers are tuned, tested, and optimized by NVIDIA. Other containers for additional HPC applications and data analytics are fully tested and made available by NVIDIA as well. NGC containers provide powerful and easy-to-deploy software proven to deliver the fastest results, allowing users to build solutions from a tested framework, with complete control.

NGC offers step-by-step instructions and scripts for creating DL models, with sample performance and accuracy metrics to compare your results. These scripts provide expert guidance on building DL models for image classification, language translation, text-to-speech and more. Data scientists can quickly build performance-optimized models by easily adjusting hyperparameters. In addition, NGC offers pre-trained models for a variety of common AI tasks that are optimized for NVIDIA Tensor Core GPUs and can be easily re-trained by updating just a few layers, saving valuable time.

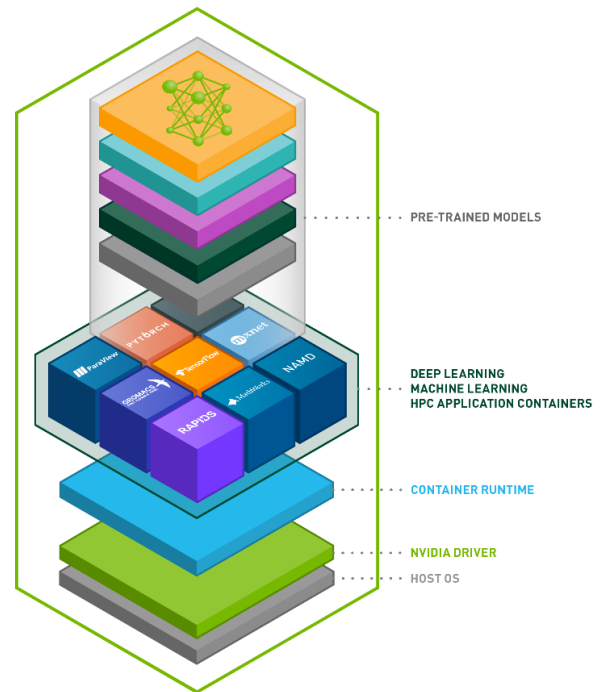


Figure 6: NGC Components

Solution Validation

This section describes the lab architecture, configuration, and validation of the NVIDIA DGX A100 systems with the IBM Spectrum Scale on IBM ESS3000 with QM8700 switches (Figure 7).




Figure 7: Converged Architecture for NVIDIA DGX A100 systems with IBM Storage

The converged infrastructure was created using QM8700 switches for both Compute fabric and Storage fabric to simplify network administration and deployment.

The DGX POD has Compute and Storage fabric networks:

- **Compute fabric.** Connects the eight 200 Gb NVIDIA Mellanox ConnectX®-6 HCAs from each DGX A100 system through separate network planes for inter-node communication. These adapters are configured in InfiniBand mode.
- **Storage fabric.** Connects two 200 Gb/s NVIDIA Mellanox ConnectX-6 HCAs from each DGX A100 system for storage communication purpose. This design document recommends InfiniBand for more throughput and network simplification.

NVIDIA RAP LAB Testing

This section describes the various performances tests run on the IBM ESS 3000 platform in the NVIDIA RAP Lab.

Follow on testing of the IBM ESS3200 performance took place in the IBM DGX Lab. IBM ESS 3200 results are covered in the next section.

NVIDIA Collective Communications Library (NCCL) scalability test

Various tests are performed to validate the scalability of GPU to GPU communication across the multiple DGX A100 systems in the POD. Multi-node job scheduling requires high performance and low latency RDMA communication between the nodes in the DGX POD. Test results show the scalable bandwidth between the NVIDIA DGX A100 systems required for multi-node DL workloads in Figure 8.




Figure 8: NCCL bandwidth performance for DGX A100 systems

Fio Bandwidth Test:

Flexible IO tester (FIO) is an open-source benchmark tool used for testing the storage bandwidth. FIO bandwidth tests were performed to measure the storage bandwidth from each DGX A100 system. Tests were simulated to run with both 100% reads and 100% writes with the below parameters

```
direct=1
ioengine=posixaio
iodepth=32
blocksize=1024k
size=4194304k
numjobs=120
```

The total system throughput performance for one to eight DGX A100 systems with the two IBM ESS 3000 units shows the performance scaling for reads (Figure 9) and writes (Figure 10) as systems are added. Two IBM ESS 3000s provide sufficient bandwidth to meet the throughput needs of eight DGX A 100 systems running the tested workloads.

As expected, the scalability of read performance is roughly linear through 5 NVIDIA DGX A100 servers. The performance tapers off as the systems approach the maximum throughput of the 2 ESS 3000 servers and the storage network connections. Should more read or write performance be needed, a third IBM ESS 3000 can be easily added to the network to deliver additional throughput.

Figure 9 shows the total system read throughput scales to deliver of 94 GB/s with two ESS 3000 units and eight DGX A100 systems.




Figure 9: 8 node DGX A100 read performance with 2 x IBM ESS 3000 systems

Figure 10 shows the total system write throughput scales to 62 GB/s with two ESS units and eight DGX A100 systems.

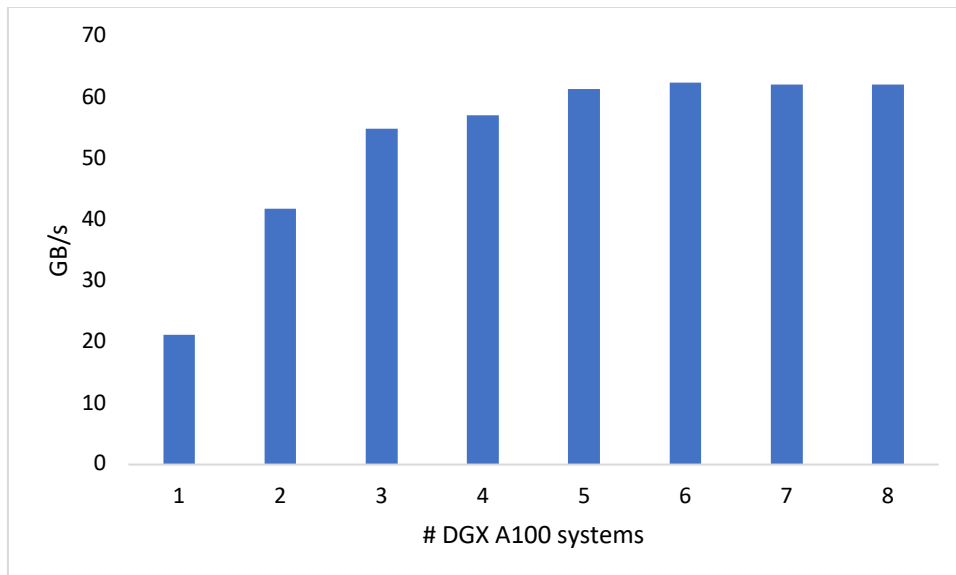


Figure 10: 8 node DGX A100 write performance with 2 x IBM ESS 3000 systems

IBM Spectrum Scale solutions scale almost linearly as storage systems are added. Results of the testing in Figure 11 demonstrates the read performance with the addition

of a second IBM ESS 3000. Each IBM ESS 3000 delivers around 48 GB/Sec and scales almost linearly with the addition of an IBM ESS 3000 system to the same IBM Spectrum Scale storage instance.

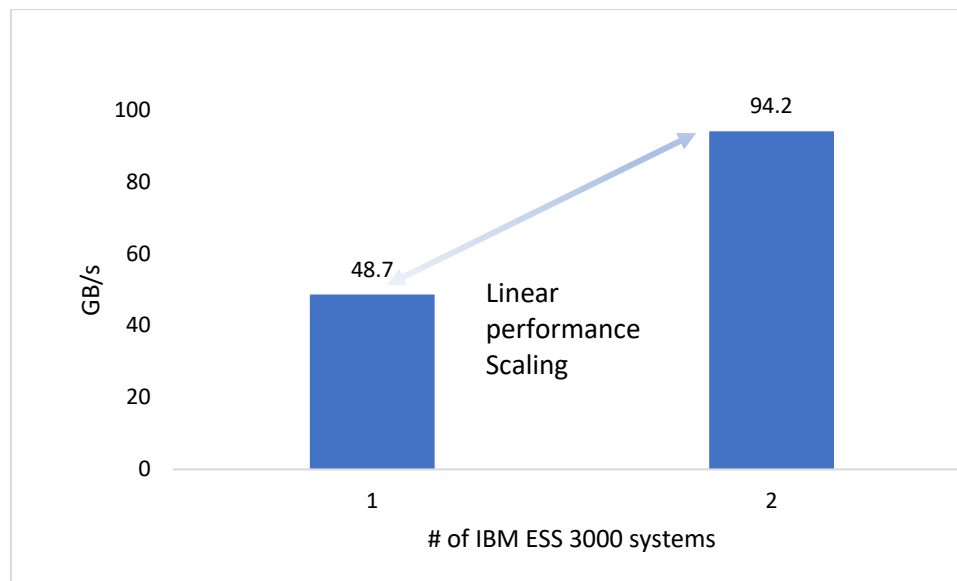


Figure 11: IBM ESS 3000 performance Scaling linearly

The results demonstrate that the ESS 3000 solution maximizes the potential throughput of the data infrastructure, scaling linearly from around 48 GB/s read performance for one ESS 3000 unit and also with two ESS 3000 units in only 4 units of rack space. Due to the composable scalability, extremely large, high-performance storage can be delivered using IBM ESS solutions with IBM Spectrum Scale.

MLPerf test

MLPerf is the industry standard set of benchmark implementations of neural networks and it measures how fast a system can perform training and inferencing of the DL models. Configurations were tested with one, two, four and eight DGX A100 systems and two IBM ESS 3000 systems to measure the operation of DL workloads. This test used the MXNet implementation of ResNet-50 along with the ImageNet dataset.

The training results were measured at Epoch 0 and compared to the overall run average time for measuring the storage system read bandwidth during the validation. Epoch 0 is the most I/O intensive portion of the MLPerf benchmark run and time to insight is dependent on the Storage system performance.

IBM storage solutions demonstrate the linear scaling of the DGX POD with the additional DGX A100 systems added to the infrastructure. The results of the MLPerf performance in Figure 12 demonstrate the Multi-GPU scaling of the solution with the IBM Storage systems without data bottlenecks. It delivers linear scaling of time to insights as more DGX A100 systems are added to the DGX POD and reduces the time to insights to under 7 minutes with eight DGX A100 systems running with the two IBM ESS 3000 systems.

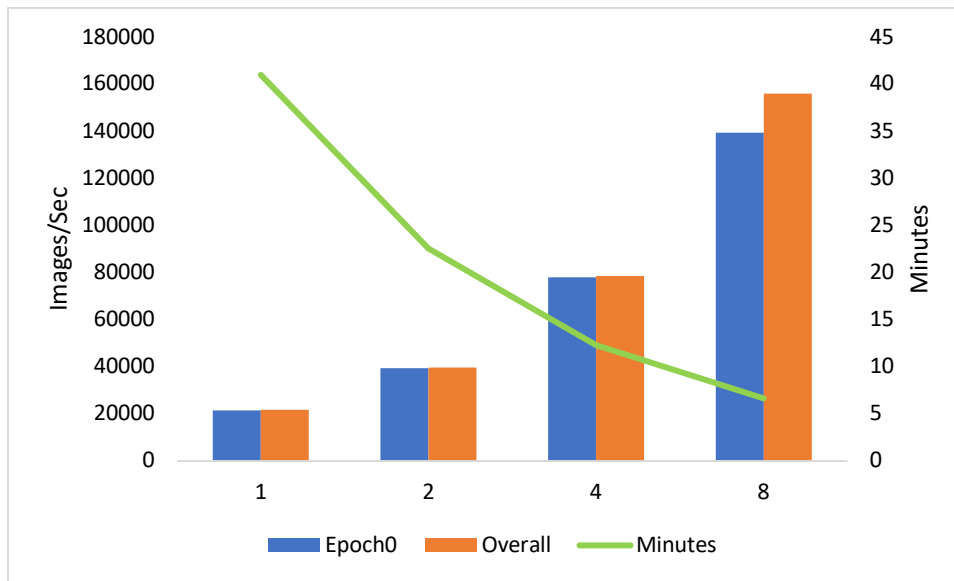


Figure 12: MLPerf ResNet-50 Performance

ESS 3200 IBM LAB Testing

This section describes the performances tests run on the IBM ESS3200 platform in the IBM Lab. The IBM Lab used the NVIDIA RAP Lab configuration as a guide for a two node NVIDIA DGX A100 cluster.

The expected throughput requirement of two NVIDIA DGX servers should be approximately 40GB/s in normal operation based upon the NVIDIA RAP Lab testing above. Therefore, a single ESS3200 capable of 80GB/s was used for the testing.

Lab Architecture

A two node NVIDIA DGX A100 cluster was tested with a single IBM ESS3200 and a single QM8700 InfiniBand switch. The network was HDR InfiniBand.

The DGX POD has Compute and Storage fabric networks:

- **Compute fabric.** Connects the eight 200 Gb NVIDIA Mellanox ConnectX®-6 HCAs from each DGX A100 system through separate network planes for inter-node communication. These adapters are configured in InfiniBand mode.
- **Storage fabric.** Connects two 200 Gb/s NVIDIA Mellanox ConnectX-6 HCAs from each DGX A100 system for storage communication purpose. This design document recommends InfiniBand for more throughput and network simplification.

The IBM ESS3200 was connected to the storage fabric with four 200 Gb/s NVIDIA Mellanox ConnectX-6 HCAs.

Bandwidth Test:

IBM used the NVIDIA provided GDSIO¹ to enable comparison testing of standard data throughput against GDS enabled results. Similar to FIO, GDSIO allows for the specification of target files, file sizes, IO sizes, number of IO threads, and other related IO parameters. Unlike FIO, it enables GPUDirect Storage (GDS), including validating the data that is directly written to the GPU memory.

Bandwidth tests were performed to measure the storage bandwidth from each DGX A100 system without GDS. Tests were simulated to run with both 100% reads with the below parameters

blocksize=1024k

size=16GB

threads=32 (two per GPU)

The throughput for two DGX A100 systems was 39GB/s².

This is almost exactly the FIO results demonstrated with the two ESS3000s in the NVIDIA RAP Lab. It confirms the decision to size these tests with a single ESS3200, which is capable of delivering the throughput needed by the pair of NVIDIA DGX A100 servers.

Magnum IO GPUDirect Storage Testing:

Testing with GDS enabled resulted in significant improvements in throughput, as well as reduced latency and lower CPU utilization. All tests were performed across the InfiniBand storage fabric, which is 2x HDR (200Gb/s) network interfaces on each NVIDIA DGX server.

For reference, the theoretical maximum, no-overhead throughput is 25GB/s per network interface, or 50GB/s per server. Of course, there is overhead in the InfiniBand protocols and data processing, so this is an unachievable number when measuring the actual data throughput. A GA version of NVIDIA GDS was installed on the DGX systems and enabled with the IBM Spectrum Scale 5.1.1. Tech Preview GDS support.

Using the same parameters as the non-accelerated throughput benchmarks, the results at 1.96x more throughput, or 77GB/s³.

Additional testing using GDSIO at different number of threads and block sizes demonstrated similar results.

NVIDIA GPUDirect Storage (GDS) is part of the MagnumIO data enhancement program from NVIDIA. GDS is enabled through the CUDA framework for programming and executing accelerated code on NVIDIA GPUs.

¹ GDSIO documentation : <https://docs.nvidia.com/gpudirect-storage/configuration-guide/index.html#benchmarking-gds>

² O_Direct read measurements performed by IBM Labs: One ESS 3200 running IBM Spectrum Scale 5.1.1, with four InfiniBand HDR to two NVIDIA DGX A100 servers using the two storage fabric HDR InfiniBand network connections.

³ Also reported as 71GiB/s for the GDS accelerated throughput, compared to 36GiB/s for the standard data path.

IBM Spectrum Scale 5.1.1 provides GPUDirect Storage as Tech Preview today, so our clients can test their own code, understand the impact to the applications, and plan for the storage demands of the future.

Sizing Considerations

In the two-node configuration, a single ESS 3200 has demonstrated the ability to deliver more than enough performance to support two NVIDIA DGX A100s requirements for normal operations. Furthermore, the single ESS 3200 has demonstrated the ability to support GDS accelerated requirements across two NVIDIA DGX servers.

The 2x improved performance of the ESS 3200 over the ESS 3000⁴ provides the additional bandwidth to support additional NVIDIA DGX servers in conventional applications and/or to deliver the performance needs of GDS accelerated applications. Should additional performance be required, the ESS solutions have demonstrated the near linear scalability of IBM Spectrum Scale. Larger systems can be built with confidence.

	2 x DGX A100 Read perf	4 x DGX A100 Read perf	8 x DGX A100 Read perf
2x ESS 3000	42 GB/s	78 GB/s	94 GB/s
1x ESS 3200	40 GB/s		
1x ESS 3200 w. GDS¹	77GB/s		

¹ GPUDirect Storage Tech Preview

Figure 13: Read performance of ESS 3000, ESS 3200, and ESS 3200 with GPUDirect Storage⁵

The ability for Magnum IO GPUDirect Storage to increase data throughput by roughly 2x over the normal data path in a DGX A100 is being met by the advances in IBM ESS Storage.

⁴ <https://newsroom.ibm.com/2021-04-27-IBM-Launches-Advanced-Storage-Solutions-Designed-to-Simplify-Data-Accessibility-Availability-Across-Hybrid-Clouds>

⁵ RAP Lab performance results with FIO and IBM Lab results with GDSIO using configurations presented in this paper

Solution Sizing Guidance

This reference architecture intends to provide the various configuration guidelines to the customers and partners planning to implement a DL infrastructure with the NVIDIA DGX A100 systems and IBM Storage systems. We focused on three rack configurations sized for different starting points (Figure 14). These are building blocks with which data and IT teams can growth from two to eight nodes. Though the deployment architecture is not in scope for this paper, the single name space of IBM Spectrum Scale also readily adapts to multi-rack scaling and larger clusters.



Figure 14: Two, four, and eight DGX A100 system configurations

2:1 configuration

Two DGX A100 systems with one IBM ESS 3000 systems

Figure 15 depicts the two DGX A100 system configuration with the one IBM ESS 3000 systems using QM8700 IB switches. This configuration delivers up to 40 GB/s read performance from a single ESS 3000 system. Each DGX A100 system delivers 20 GB/s read performance using 200 Gb/S HDR connection configured for the storage usage.




Figure 15: 2:1 configuration – Two DGX A100 systems with one IBM ESS 3000 system

4:1 configuration

Four DGX A100 systems with one IBM ESS 3000 systems

Figure 16 depicts the 4 DGX A100 system configuration with the one IBM ESS 3000 systems using QM8700 IB switches. This configuration delivers up to 48 GB/s read performance from a single ESS 3000 system, which is around 12 GB/s per DGX A100 system.




Figure 16: 4:1 configuration – Four DGX A100 systems with one IBM ESS 3000 system

8:2 configuration

8 DGX A100 systems with two IBM ESS 3000 systems

Figure 17 depicts the eight DGX A100 system configuration with two IBM ESS 3000 systems using QM8700 IB switches. This configuration delivers up to 94 GB/s read performance from a single GPFS file system configured using two IBM ESS systems. This configuration delivers more than 10 GB/s read performance per DGX A100 system.

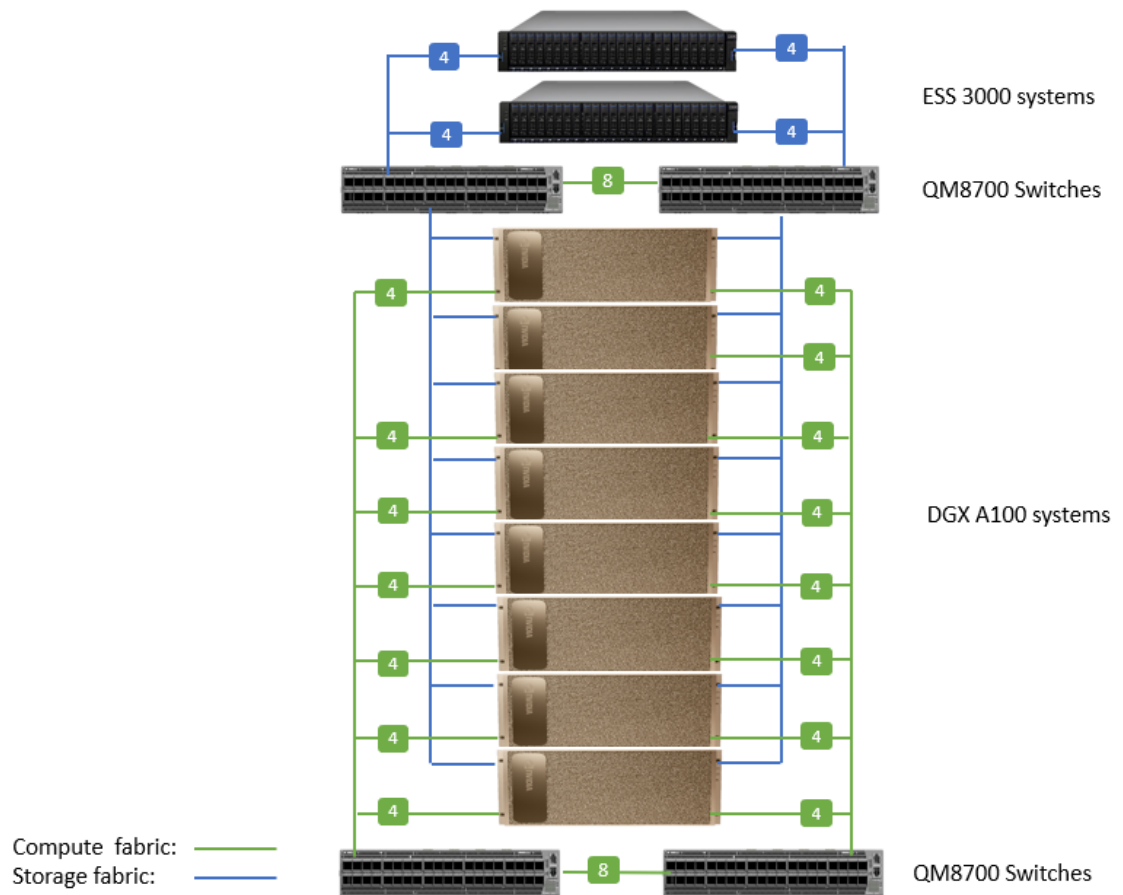


Figure 17: 8:2 configuration – Eight DGX A100 systems with two IBM ESS 3000 systems

Conclusion

The IBM Storage architecture over an NVIDIA IB fabric provides converged infrastructure for both compute and Storage fabric. IBM Storage provides a leading-edge performance for DL workload with high bandwidth and low latency for full utilization of GPUs when running on multiple DGX A100 systems as our tests have demonstrated.

IBM Storage for Data and AI and the NVIDIA DGX POD systems integrated with the NGC Software Stack provides the workload consolidation, data preparation and management, and process automation that organizations seek to streamline end-to-end AI data pipeline development and ease integration into existing infrastructure

Get more Information

To learn more about the IBM, Mellanox, and SAS products and capabilities, contact your IBM representative or IBM Business Partner, or visit the following websites:

- IBM Spectrum Scale:
<https://www.ibm.com/products/scale-out-file-and-object-storage>
 - IBM Elastic Storage System
<https://www.ibm.com/us-en/marketplace/elastic-storage-system-3000>
 - NVIDIA DGX A100 Systems
<https://www.nvidia.com/en-in/data-center/dgx-a100/>
-

Authors

Douglas O’Flaherty is the **Global Ecosystem Leader** for IBM Storage. His team works with clients and partners to develop solutions that deliver business value and technical excellence. He presents frequently at conferences and events on how to adopt leading-edge technology in enterprise IT.

Sanjay Sudam is the Senior Solution Architect for the Data and AI solutions using IBM Storage systems. He is responsible for creating the reference architectures and solution blueprints with the IBM storage portfolio and external partner solutions. Sanjay has created an end-to-end reference architecture for AI, analytics, cloud data, Data Protection, Digital Video Surveillance and Media Entertainment solutions with the IBM portfolio.

Acknowledgements

We gratefully acknowledge the contributions to the IBM Storage Reference Architecture by our partners from NVIDIA, Jacci Cenci, Richard Salazar, Darrin Johnson, Robert Sohigian, Nachiket Karmarkar, Fred Herard, Brian Dowdy, Akshay Patel and Matthew Sherad.

We could not have completed this without the support and guidance of our IBM team members Bill Martinson, Vanessa Vera Nava, Pallavi Galgali, Matt Geiser, Piyush Chaudhary, Kumaran Rajaram, Jay Vaddi and John Lewars.

Our sincere appreciation and thanks to all these individuals, who provided insight and expertise that greatly assisted for this paper.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements, or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing, or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at

<http://www.ibm.com/legal/copytrade.shtml>

NVIDIA, the NVIDIA logo, DGX, DGX A100 are registered trademarks or trademarks of NVIDIA, Inc., or its subsidiaries in the United States and/or other jurisdictions.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Other product, company or service names may be trademarks or service marks of others.

References in the publication to IBM products or services do not imply that IBM intends to make them available in all countries in the IBM operates.