IBM Spectrum Storage for AI with NVIDIA DGX Systems

Proven Infrastructure Solution for AI workloads

Contents

Introduction

Artificial intelligence (AI) and deep learning (DL) are the engines that are rapidly powering innovation across industries from healthcare to autonomous vehicles and agriculture. By 2020, IBM projects that the world's volume of digital data will exceed 44 zettabytes [1]. Organizations that recognize the value of their data for decisions and actions are turning to DL systems that can rapidly ingest, accurately interpret, and quickly provide key data insights from the volumes of new data generated now and in the future.

Enterprises are increasing investment in AI research and innovation, with related patents growing more than 30% and academic papers by 13% during the last decade [2]. Arguably, only one kind of enterprise will survive and thrive in the future— the data-driven enterprise.

Highly performant and scalable DL systems must excel at data ingestion, data training, and verification, and deliver inferences and classifications while handling the growing demands of DL in the organization. The DL algorithm accuracy may improve by increasing the neural network size, and data quantity and quality used for the model training step. However, this accuracy improvement comes with a significant increase in computational complexity and increased demand on DL compute, storage, and network resources.

NVIDIA has led the AI computing revolution, leveraging the power of the modern GPU with its massive processor core count and parallel architecture, uniquely suited to the massively parallelized operations that are core to DL, and which exceed the limitations of traditional CPU based architectures. IBM delivers industry-leading, high-performance, low-latency, and cost effective all-flash storage and software defined clustering with proven scalability that enables massively parallel processing of different data types such as image, audio, video, text, or time series data.

Together NVIDIA and IBM provide an integrated, individually scalable compute and storage solution with end-to-end parallel throughput from flash to GPU for accelerated DL training and inference. This paper is intended for enterprise leaders, solution architects, and other readers interested in learning how the IBM Spectrum Storage for AI with NVIDIA® DGX™ systems simplifies and accelerates AI. The scalable infrastructure solution integrates the NVIDIA DGX-1™ systems and NVIDIA DGX-2™ systems with IBM® Spectrum Scale™ file storage software which powers the IBM Elastic Storage Server (ESS) family of storage systems that includes the new IBM Elastic Storage System (ESS 3000).

The paper also demonstrates how this solution provides linear storage performance while scaling from one to nine DGX-1 systems or one to three DGX-2 systems with both synthetic workloads and ImageNet data.

[1] https://www.ibm.com/blogs/systems/ibm-and-nvidia-further-collaboration-to-advance-open-source-gpu-acceleration/

[2] https://web.luxresearchinc.com/hubfs/18%20for%202018/Lux%20Research%2018%20for%202018.pdf

Components of IBM Spectrum Storage for AI with NVIDIA DGX Systems - *Converged Solution*

The NVIDIA DGX Systems Family

The NVIDIA DGX systems family of supercomputers is a purpose-built fully integrated hardware and software solution for AI applications. The NVIDIA DGX POD™ implements data center design best practices for DL, which have been incorporated in the reference architecture. Both the DGX-1 system and the DGX-2 system are described in this paper.

The DGX-1 system and DGX-2 system are powered by the DGX system software which includes the operating system (DGX system OS) and NVIDIA GPU Cloud (NGC) Deep Learning Stack with fully-optimized versions of today's DL containers engineered by NVIDIA for maximized GPU-accelerated performance. This software stack facilitates rapid development and deployment on a single DGX system (or multiple DGX systems), and multi-GPU and multi-system scale-up of applications on the DGX system platform saving time and developer effort. The software components offered with the DGX POD include cluster management, libraries, frameworks, and workload scheduling which can be leveraged for management of the architecture described in this paper. Visit NVIDIA GPU Cloud (NGC) to learn more and get started.

NVIDIA DGX-1 System

Figure 1*: NVIDIA DGX-1 System*

Each DGX-1 system (Figure 1) integrates eight NVIDIA Tesla™ V100 Tensor Core GPUs configured in a hybrid cube mesh topology that uses NVIDIA NVLink™ technology. This fabric provides high bandwidth low-latency GPU-to-GPU communication that enables scalable multi-GPU training while eliminating the PCIe-based interconnect bottleneck found in traditional architectures that result in non-linearity of performance as more GPUs are used for training. The DGX-1 system includes four Mellanox VPI cards enabling EDR InfiniBand or 100 GbE network ports for multi-node clustering with high speed RDMA capability for exceptional storage-to-DGX system data rates.

NVIDIA DGX-2 System

Each DGX-2 system (Figure 2) integrates sixteen Tesla V100 Tensor Core GPUs configured in a NVIDIA NVSwitch™ fabric x 16 GPU interconnect fabric that uses NVLink technology. This high bandwidth low-latency GPU-to-GPU fabric provides 2 petaFLOPS of DL computing capacity by eliminating bottlenecks and intermediary GPU hops, and enables scalable multi-GPU training while eliminating the PCIe-based interconnect bottleneck found in traditional architectures.

The DGX-2 system includes eight Mellanox VPI cards enabling EDR InfiniBand or 100 GbE network ports for multi-node clustering with high speed RDMA capability for exceptional storage-to-DGX system data rates.

Figure 2: *NVIDIA DGX-2 System*

Mellanox InfiniBand Network

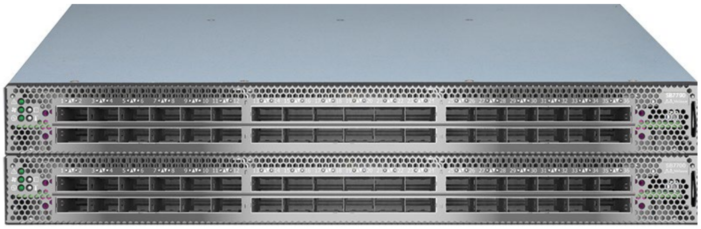

Figure 3: *Mellanox SB7700 Series*

Current challenges facing data-driven organizations are the need to ingest and analyze an exponentially increasing deluge of data in ever shorter amounts of time as well as supporting more complex models and simulations where, often, actionable insight and responses need to be made in real time. To meet the demands of these current and emerging data challenges, it is necessary to scale out the datacenter in an efficient way.

For this reference architecture, the IBM Spectrum Scale on NVMe storage is attached to the DGX-1 or DGX-2 systems by a Mellanox EDR InfiniBand network to provide the most efficient scalability of the GPU workloads and datasets beyond a single DGX system while providing the inter-node communications between DGX systems.

By delivering the fastest data speeds, lowest latency, and intelligent offload capabilities, InfiniBand has maintained its leadership as the best choice to connect the world's top HPC and AI supercomputers. Mellanox delivers a complete, end-to-end interconnect solution for high performance GPU clusters used for AI workloads, leveraging advanced technologies such as RDMA, GPUDirect® technology to accelerate GPU-to-GPU communications, SHARP™ technology to accelerate machine learning algorithms by providing state-of-the-art in-network computing capabilities.

IBM Spectrum Scale Powered Storage Systems

IBM and NVIDIA recently conducted two benchmarks that tested the Non-Volatile Memory Express (NVMe) flash storage component of IBM Spectrum Storage for AI with DGX-1 systems and DGX-2 systems. Each benchmark included leading IBM software defined file storage, IBM Spectrum Scale, which powers the IBM Elastic Storage Server (ESS) family of storage systems and the NVMe all-flash ESS 3000. ESS is a family of pre-integrated storage hardware offerings. The IBM Spectrum Storage for AI with NVIDIA DGX Systems converged solution can be composed using any of the ESS all-flash models or the ESS 3000 (Figure 4).

IBM Elastic Storage System 3000

Figure 4: *IBM Elastic Storage System 3000*

IBM Elastic Storage System 3000 combines the performance of flash and end-to-end NVMe with the rich features of IBM Spectrum Scale, along with several high-speed attachment options such as 100 Gb/s IB— all in a powerful 2U storage system.

ESS 3000 was also tested with other industry standard NVMe flash drive options such as with Samsung NVMe drives with the DGX-2 system tests. With each of these drive options IBM Spectrum Scale on NVMe is designed to be the market leader in all-flash performance, and scalability with bandwidth of around 40 GB/s per ESS 3000 unit, around 120 GB/s sustained reads for a cluster of three ESS 3000 units, and 100 microseconds latency.

Providing data-driven multicloud storage capacity, ESS 3000 is deeply integrated with the software defined capabilities of IBM Spectrum Scale storage™ allowing you to plug it into your AI data pipeline easily.

ESS 3000 supports a variety of connectivity options and other industry standard NVMe flash drive options. It is well designed to address the full range of AI workloads and business use cases.

IBM Spectrum Scale and IBM Spectrum Scale RAID

IBM Spectrum Scale is the industry leader in high performance parallel file system software. It powers the world's #1 and #2 fastest supercomputers, named Summit and Sierra respectively, as well as other supercomputers. And the world's #1 and #2 fastest supercomputers, named Summit and Sierra respectively, among other supercomputers are powered by IBM Spectrum Scale. "Summit will have the capacity of 30B files and 30B directories and will be able to create files at a rate of over 2.6 million I/O file operations per second. That is opening every book in the US Library of Congress in 10 seconds." (https://www.ibm.com/blogs/systems/fastest-storage-fastest-system-summit)

A key ability that Spectrum Scale provides is a single namespace (or data plane) so that each data source can add data to the repository using NFS, SMB, Object, or a POSIX interface. This single data plane allows the data prep tools to access the data in place – no copying required. AI training can access the data in place also, as can the inference applications, all with no copying, all through industry standard interfaces.

Another key strength is 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. IBM provides full AI pipeline support. For more information see the Enterprise Data Pipeline section of this paper.

IBM Spectrum Scale RAID is a software implementation of storage erasure code technologies within IBM Spectrum Scale that provides sophisticated data placement and error-correction algorithms to deliver high levels of storage performance, availability, and reliability. Shared file systems are created from the Network Shared Disks (NSD) defined with IBM Spectrum Scale RAID. This file system can be accessed concurrently by all compute nodes in the configuration to efficiently meet the capacity and performance requirements of modern scale out applications like AI.

Along with the most proven scalable capabilities to handle enterprise workloads, IBM Spectrum Scale RAID as part of IBM Spectrum Storage for AI delivers enhanced data protection and support to run GPU processor-intensive AI workloads at the highest performance levels required for multi-system GPU clusters such as the DGX POD.

Scaling with GPUs

IBM Spectrum Storage for AI with DGX Systems provides scalability for capacity and performance. Deployments can start with a single ESS 3000 and a single DGX-1 or DGX-2 system or multiples of each as workloads demand, all while providing a cost-effective solution.

Figure 5 depicts possible ESS 3000 configurations in DGX POD racks for DGX-1 systems and DGX-2 systems. Each configuration can be optimized for performance and is easily scalable for expansion of existing production workflows and additional DL applications. ESS 3000 has minimal power requirements at 2U and 6U configurations allowing maximum flexibility when designing rack space and addressing power requirements.

3:1 Configuration **6:2 Configuration** **9:3 Configuration** **3:3 Configuration**

Figure 5: *Example scale-out configurations*

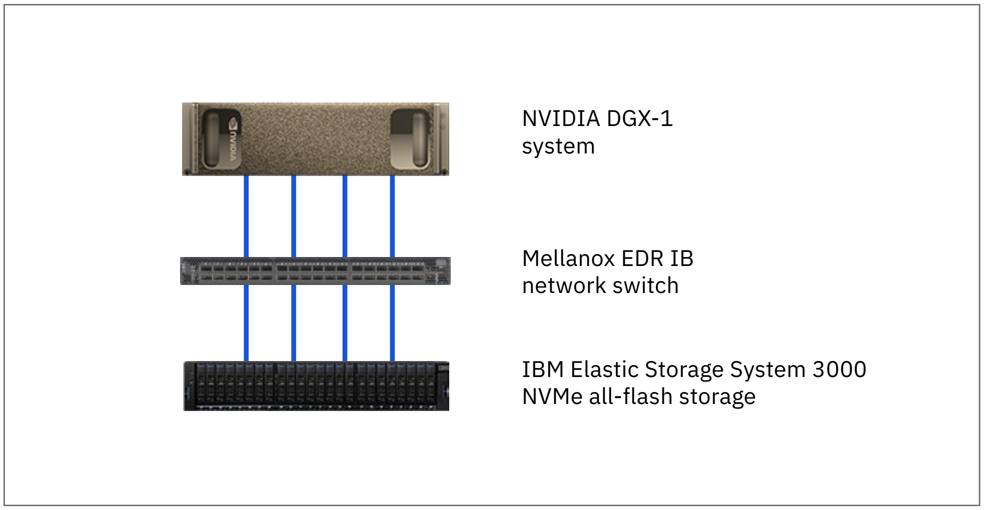

Figure 6: *1 DGX-1 system + 1 ESS 3000*

Figure 6 illustrates a one DGX-1 system to one ESS 3000 configuration.

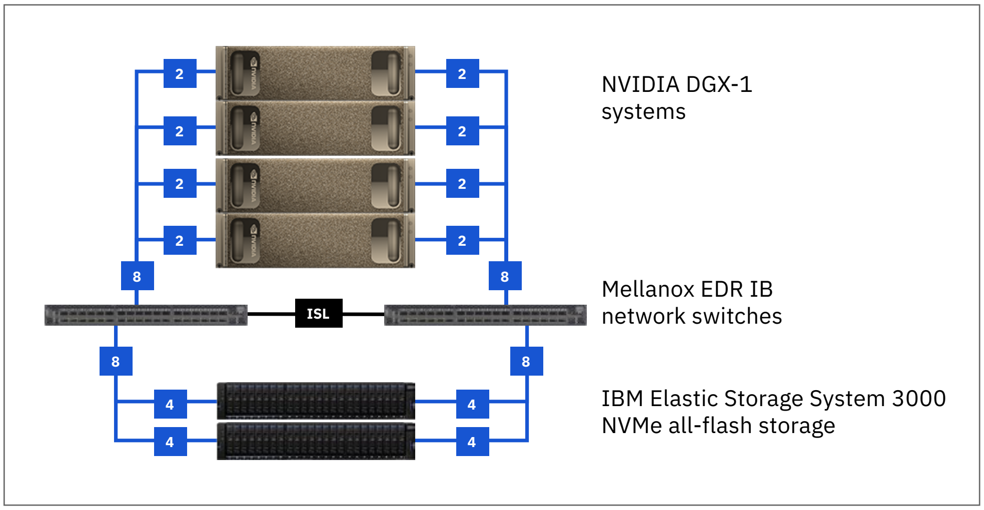

Figure 7: *IBM Spectrum Storage for AI with NVIDIA in a 4 DGX-1 System + 2 ESS 3000 unit configuration*

Figure 7 illustrates four DGX-1 systems in a two ESS 3000 unit configuration example to show another potential architecture option with either two switches as shown, or one switch if desired. Other network topologies such as fat tree should also be considered when sizing new or scaling up existing 10 infrastructure performance beyond what a flat network topology provides.

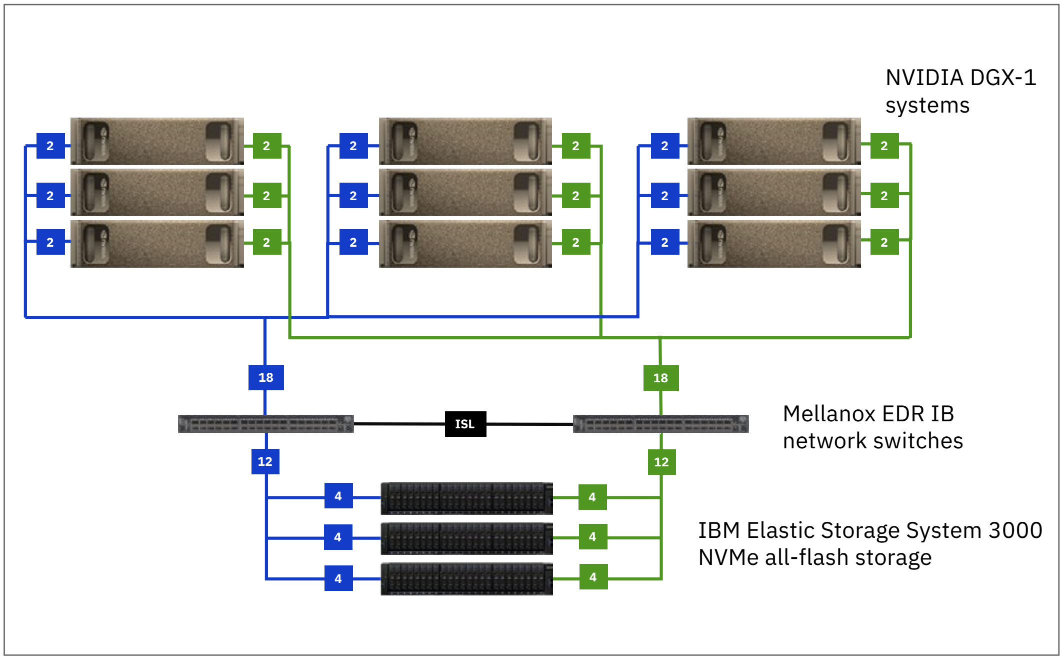

Figure 8: *9 DGX-1 systems + 3 ESS 3000 units*

Figure 8 depicts a full DGX POD configuration with three ESS 3000 units. This is the full reference architecture used in the IBM Elastic Storage System 3000 with DGX-1 systems benchmark, however only two ESS 3000 units were necessary to obtain the DGX-1 system model images/sec results reported in this paper.

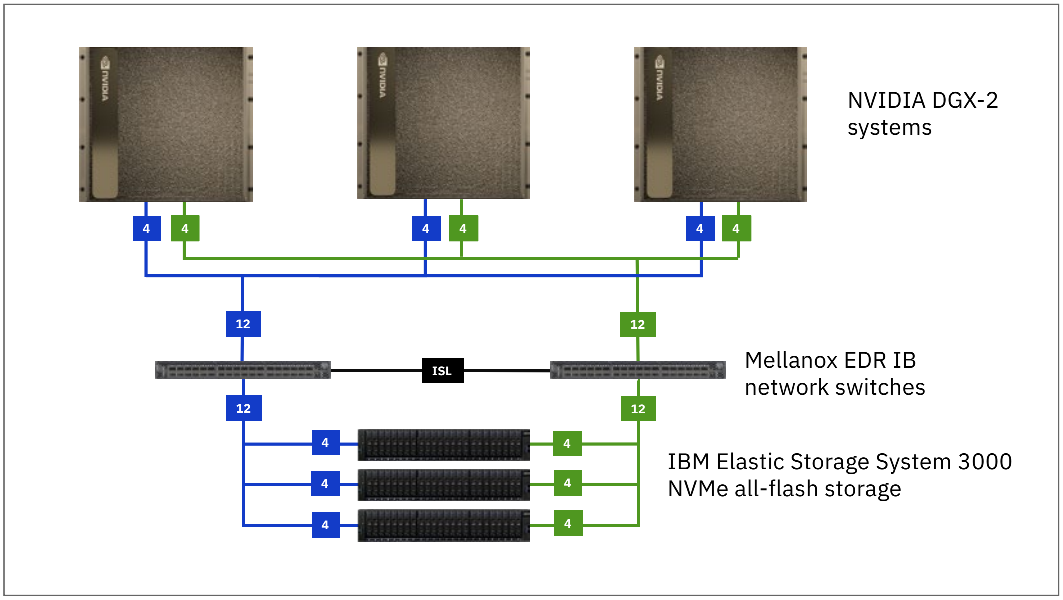

Figure 9: *3 DGX-2 systems + 3 ESS 3000 units*

Figure 9 depicts three DGX-2 systems with three ESS 3000 units. This is the reference architecture used in the IBM Elastic Storage System 3000 with NVIDIA DGX-2 systems benchmark, however only two ESS 3000 units were necessary to obtain the DGX-2 model images/sec results reported in this paper.

IBM Spectrum Storage for AI with NVIDIA DGX Systems- *Reference Architecture with DGX-1*

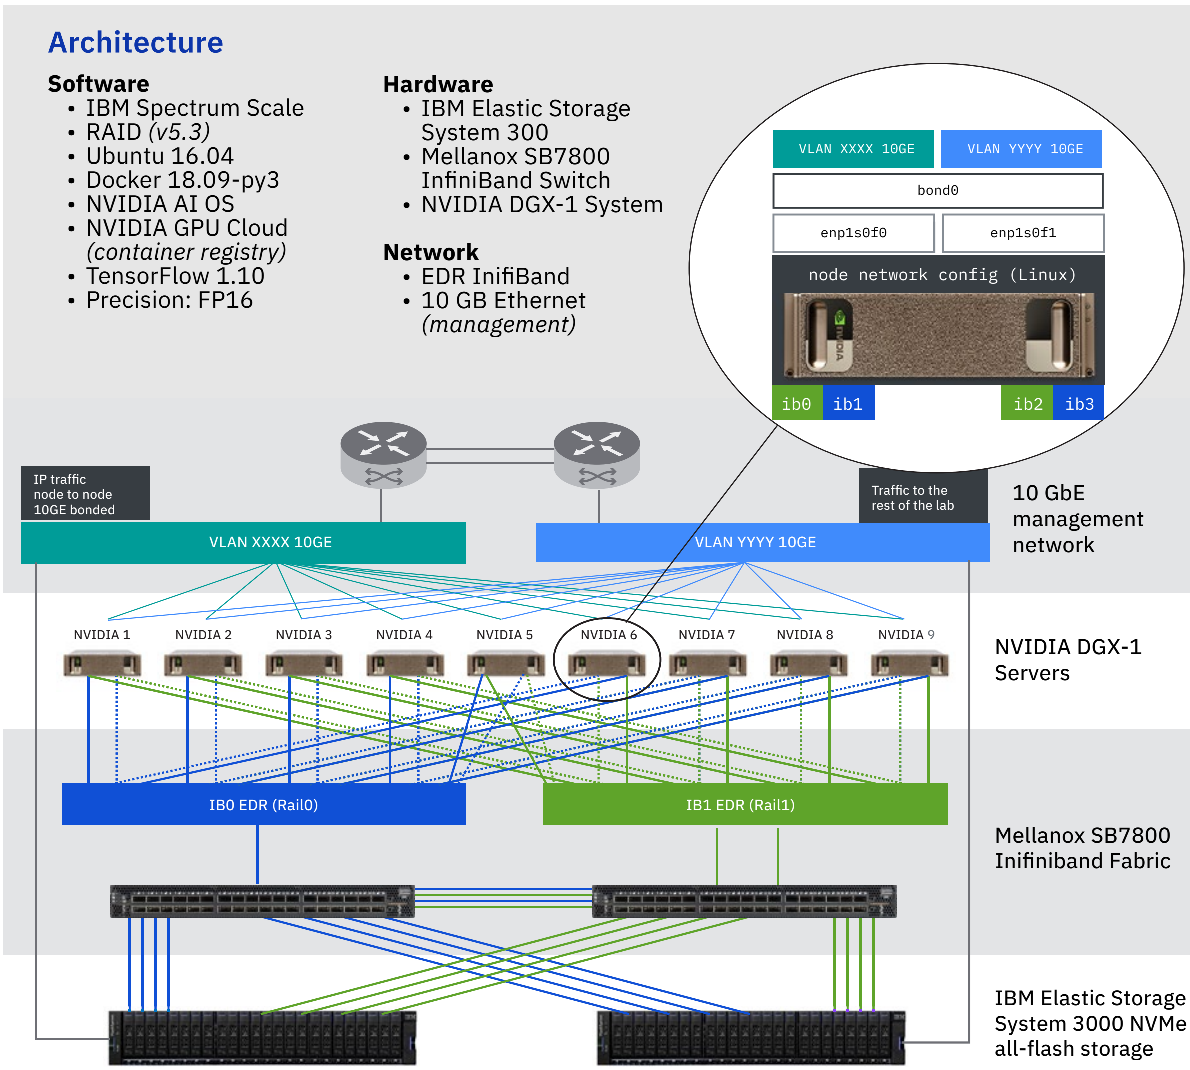

Figure 10*: DGX-1 System Solution Reference Architecture diagram*

Figure 10 illustrates the end-to-end DL reference architecture of IBM Spectrum Storage for AI with NVIDIA DGX Systems solution environment as deployed for DGX-1 system testing. In this cutting-edge DGX-1 system test environment, IBM Spectrum Scale RAID v5 is installed on the ESS 3000 base Linux OS. The IBM Spectrum Scale RAID software is generally available as part of the Spectrum Scale software stack for IBM Elastic Storage Server (ESS) deployments. As configured, each ESS 3000 unit provides a pair of fully redundant NSD systems within the Spectrum Scale cluster. The ESS 3000 is connected over EDR InfiniBand with eight links to two Mellanox SB7800 EDR switches. The DGX-1 systems also connect with four links to the two InfiniBand switches. In addition to the high speed EDR InfiniBand fabric, the ESS 3000 and DGX-1 systems are connected by a 10GbE management network to Ethernet switches.

IBM Spectrum Storage for AI with NVIDIA DGX Systems- *Reference Architecture with DGX-2*

Architecture

Software
- IBM Spectrum Scale RAID *(v5.3)*
- Ubuntu 18.04
- Docker 19.04-py3
- NVIDIA AI OS
- NVIDIA GPU Cloud *(container registry)*
- TensorFlow 1.13.1
- Precision: FP16

Hardware
- IBM Elastic Storage System 3000
- Mellanox SB7800 InfiniBand Switch
- NVIDIA DGX-2 System

Network
- EDR InifiBand
- 10 GB Ethernet *(management)*

Figure 11 illustrates the end-to-end DL reference architecture of IBM Spectrum Storage for AI with NVIDIA DGX systems solution environment as deployed for DGX-2 system testing. In this DGX-2 system test environment, similar to what was tested separately with DGX-1 systems, IBM Spectrum Scale RAID v5 is installed on the ESS 3000 base Linux OS. The IBM Spectrum Scale RAID software is generally available as part of the Spectrum Scale software stack for IBM Elastic Storage System (ESS) deployments.

As configured, each ESS 3000 provides a pair of fully redundant NSD systems within the Spectrum Scale cluster. Each ESS 3000 unit is connected over EDR InfiniBand with eight links to two Mellanox SB7800 fabric interconnect switches. The DGX-2 systems also connect with eight links to the EDR InfiniBand switch. In addition to the high speed EDR InfiniBand fabric, the ESS 3000 and DGX-2 systems are connected by a 10GbE management network to Ethernet switches.

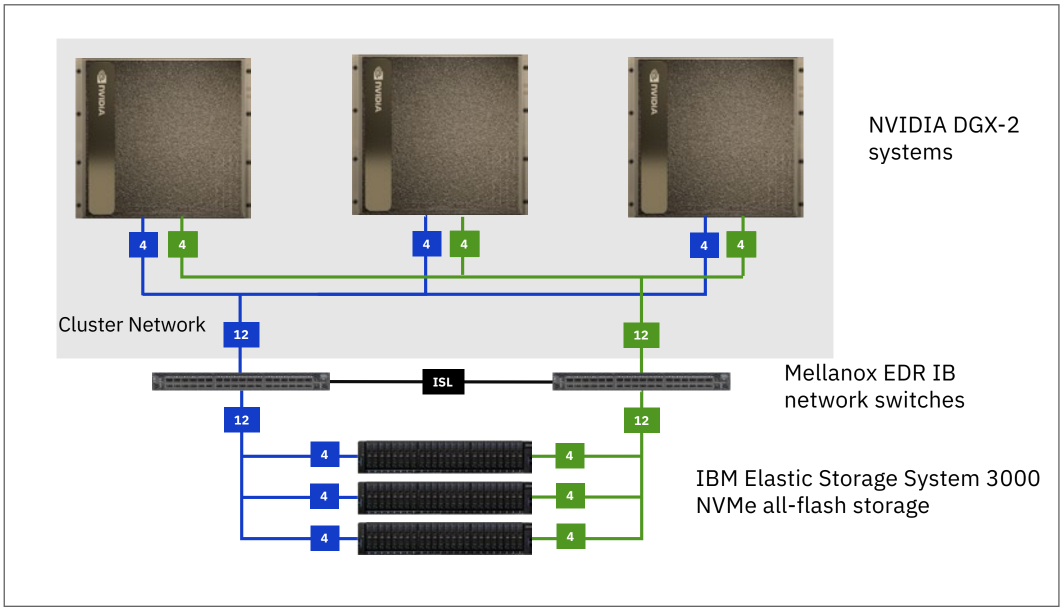

Figure 11*: DGX-2 system solution reference architecture*

Storage and Network Performance Testing

Both the DGX-1 system and the DGX-2 system benchmarks are primarily focused on providing system performance capabilities to help DL development teams plan their infrastructure effectively. For example, the DGX-1 system benchmark tests focus on DL model training performance as well as DL model inference performance and the incremental and total throughput capabilities when scaling up from one DGX-1 system to a full DGX POD with IBM Elastic Storage System 3000.

Likewise, the DGX-2 system benchmark tests focus on DL model training performance as well as DL model inference performance and the incremental and total throughput capabilities when scaling up from one DGX-2 system to three DGX-2 systems with ESS 3000. For both benchmarks the performance was tested with synthetic throughput test applications such as IOR and fio and with the DL framework TensorFlow using several models such as ResNet-50, ResNet-152, Inception-v3, and other networks with the ImageNet Large Scale Visual Recognition Challenge 2012 (ILSVRC2012) dataset.

The benchmarks utilize NVIDIA NGC containers. Each AI container has the NGC Software Stack, a pre-integrated set of GPU-accelerated software. The stack includes the chosen application or framework, NVIDIA CUDA Toolkit, NVIDIA DL libraries, and a Linux OS — all tested and tuned to work together immediately with no additional setup.

For developers, the NVIDIA Deep Learning SDK offers powerful tools and libraries for the development of DL frameworks such as Caffe2, Cognitive toolkit, MXNet, PyTorch, TensorFlow, and others. These frameworks rely on GPU-accelerated libraries such as cuDNN and NCCL to deliver high-performance multi-GPU accelerated training. Developers, researchers, and data scientists can get easy access to NVIDIA optimized DL framework containers, performance tuned and tested for NVIDIA GPUs. This eliminates the need to manage packages and dependencies or build DL frameworks from source.

TensorFlow benchmarks were performed using scripts made available by TensorFlow on GitHub. See https://github.com/tensorflow/benchmarks/tree/master/scripts/tf_cnn_benchmarks

System Throughput Results with DGX-1 Systems

The total system throughput performance for one to nine DGX-1 systems with one to three ESS 3000 units shows the performance scales to keep all DGX-1 system GPUs saturated for each benchmark easily. The results demonstrate that the ESS 3000 solution maximizes the potential throughput of the data infrastructure, scaling linearly from around 40 GB/s read performance for one ESS 3000 unit and also with multiple ESS 3000 units.

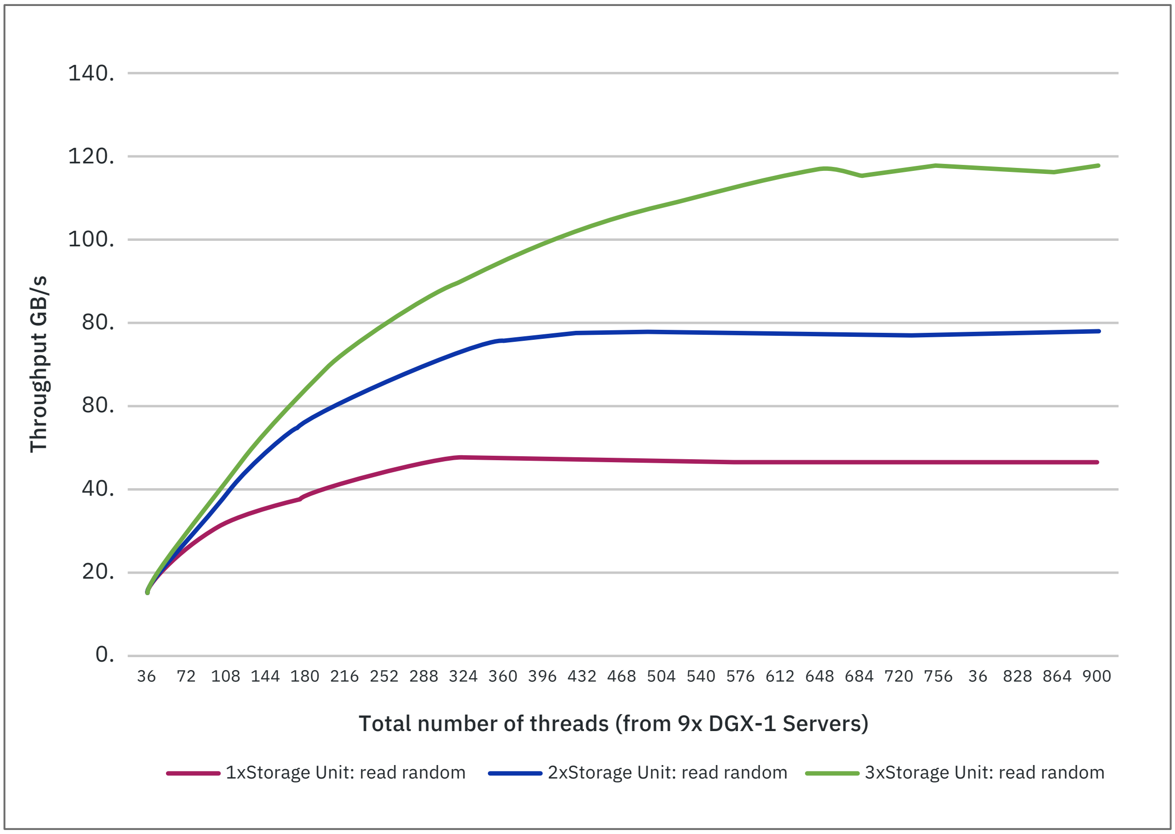

***Figure 12**: Scalable throughput using fio benchmark and DGX-1 systems*

Figure 13 shows that the ESS 3000 can efficiently feed the DGX systems to achieve full GPU utilization at nearly 100% average for the DL workloads while the overall bandwidth demands for a single DGX-1 system are not taxing the Spectrum Scale NVMe appliances. Extrapolating storage throughput capabilities and requirements from both charts shows that a single ESS 3000 as tested can handle the model training workloads presented in this paper.

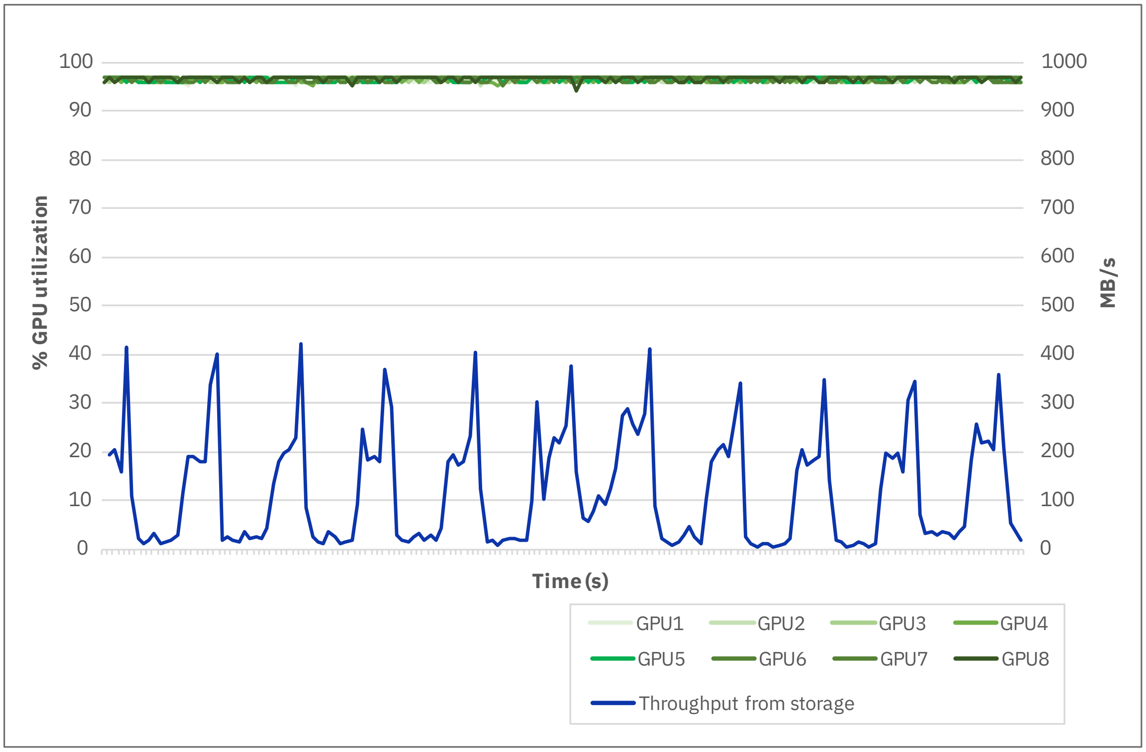

***Figure 13**: Single DGX-1 system GPU utilization vs. IO bandwidth (ResNet-50)*

To demonstrate the flexibility of the ESS 3000 storage solution, additional throughput tests were run for sequential versus random IO access patterns (Figure 14). Sequential read performance versus random read performance shows some prefetch advantage which fades when the number of job threads increases. ESS 3000 shows robust throughput capabilities regardless of the IO type.

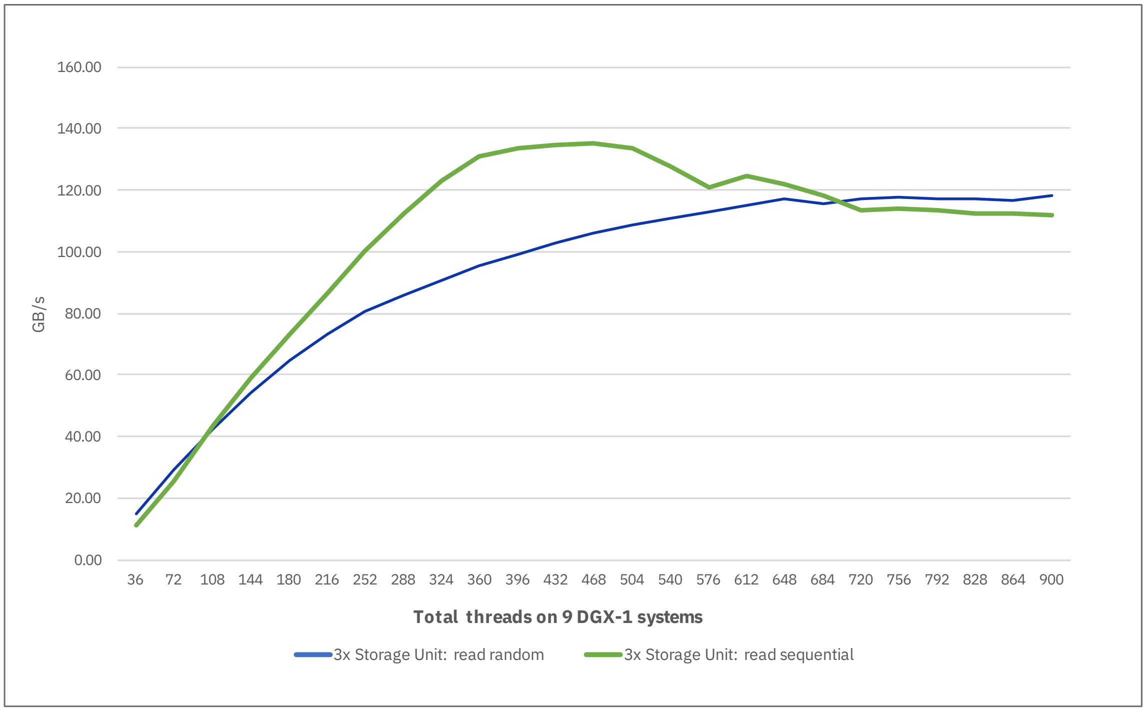

***Figure 14**: Sequential vs. Random Read throughput with fio benchmark and DGX-1 systems*

Overall, the ESS 3000 with DGX-1 systems throughput results show that ESS 3000 performs well with one to nine DGX-1 systems making full use of the GPUs.

This performance capability provides development teams the ability to add compute resources when needed knowing the storage system performance can accommodate their increasing workload demands and, if needed, by seamlessly adding additional ESS 3000 units into the AI storage cluster.

System Throughput Results with DGX-2 Systems

In Figure 16, three DGX-2 systems are needed to drive the single storage unit to maximum throughput.

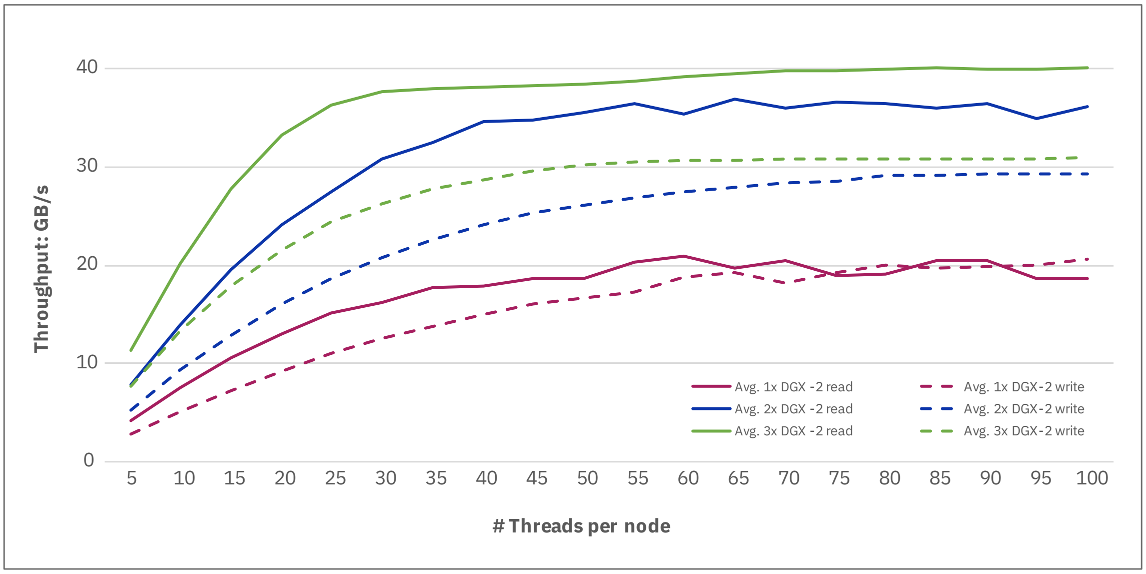

Figure 16*: Throughput benchmark for single ESS 3000 with DGX-2 systems*

The chart highlights the combined throughput capabilities of three DGX-2 systems with one, two, and three ESS 3000 units.

As previously tested with DGX-1 systems, two ESS 3000 units can provide up to 80 GB/s of random read filesystem bandwidth, and three ESS 3000 units can provide up to 120 GB/s throughput. However, two ESS 3000 storage units are all that are needed to allow three DGX-2 systems to scale nearly linearly from one, two, and three systems to reach their maximum combined 60 GB/s system bandwidth. Further testing for linear scaling from one, two, and three ESS 3000 NVMe storage units, similar to what was tested with DGX-1 systems, was not possible because of the limited number of DGX-2 systems available for testing.

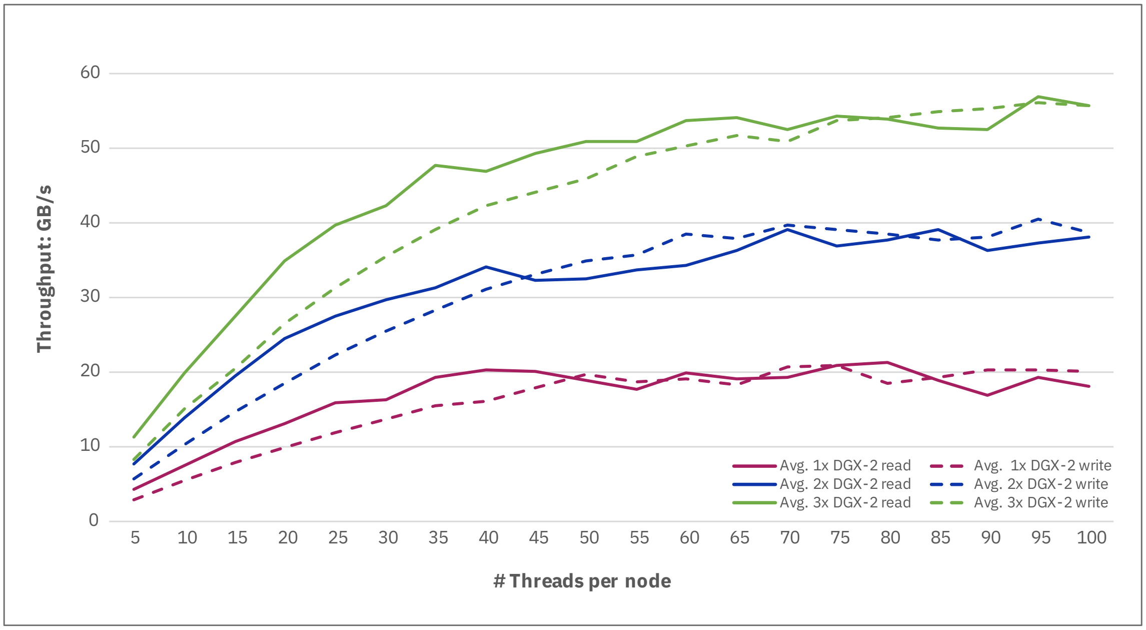

***Figure 17**: Throughput benchmark for three ESS 3000 units with three DGX-2 systems*

Next we looked at model training and inference testing that utilize the DGX-2 systems GPUs. We observed that, similar to the DGX-1 system test results, the NVMe flash storage can efficiently feed the DGX-2 systems to achieve full GPU utilization at nearly 100% average for the DL workloads while the overall bandwidth demands for a single DGX-2 system are not taxing the ESS 3000 units.

Once again, the throughput results for the ESS 3000 with DGX-2 systems show that ESS 3000 performs well with one to three DGX-2 systems, making full use of the GPUs. This performance capability provides development teams the ability to add compute resources when needed knowing the storage system performance can accommodate their increasing workload demands and, if needed, by seamlessly adding additional ESS 3000 units into the AI storage cluster.

Training Results – Single DGX-1 System

Figure 18 shows the images per second training throughput with AlexNet, ResNet-50, ResNet-152, Inception-v3, LeNet, Inception-v4, and GoogLeNet models using different numbers of GPUs on a single DGX-1 system with a separate container per GPU and comparing training runs between the IBM Spectrum Scale filesystem and local RAM disk only runs. As shown, ESS 3000 effectively feeds the DGX-1 system GPUs in the Spectrum Scale cluster, keeping the DGX-1 systems fully saturated with data for maximum training capabilities for all models.

Model training results for the single DGX-1 system also show that overall there is minimal to no penalty when comparing Spectrum Scale file system performance to internal system RAM disk performance for DL models with GPUs.

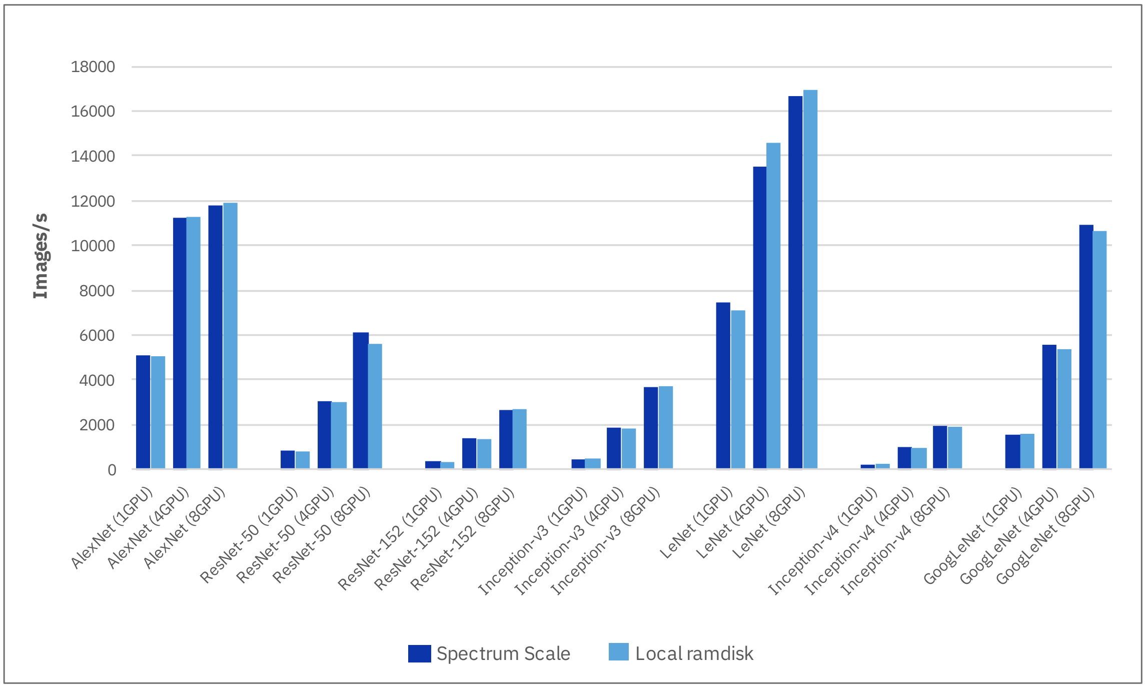

Figure 18: *Model-to-GPU performance, and Spectrum Scale filesystem vs. RAM disk performance with a single DGX-1 system*

As shown, some models scale up with linearity as the number of GPUs increase while others present a consistent non-linear scale up pattern whether using ESS 3000 storage or local RAM disk. This indicates that the scalability in these cases is not constrained by storage IO whether local or shared storage, but rather by a pattern of the DL model scalability within the compute infrastructure itself.

Training Results – Multiple DGX-1 Systems

For multiple DGX-1 systems with separate containers, ESS 3000 demonstrates linear scale-up to full saturation of all DGX-1 system GPUs simultaneously running from one to nine DGX-1 systems for a total of 72 GPUs providing the following performance results of aggregate images/sec using ImageNet datasets for the models shown in Figure 19.

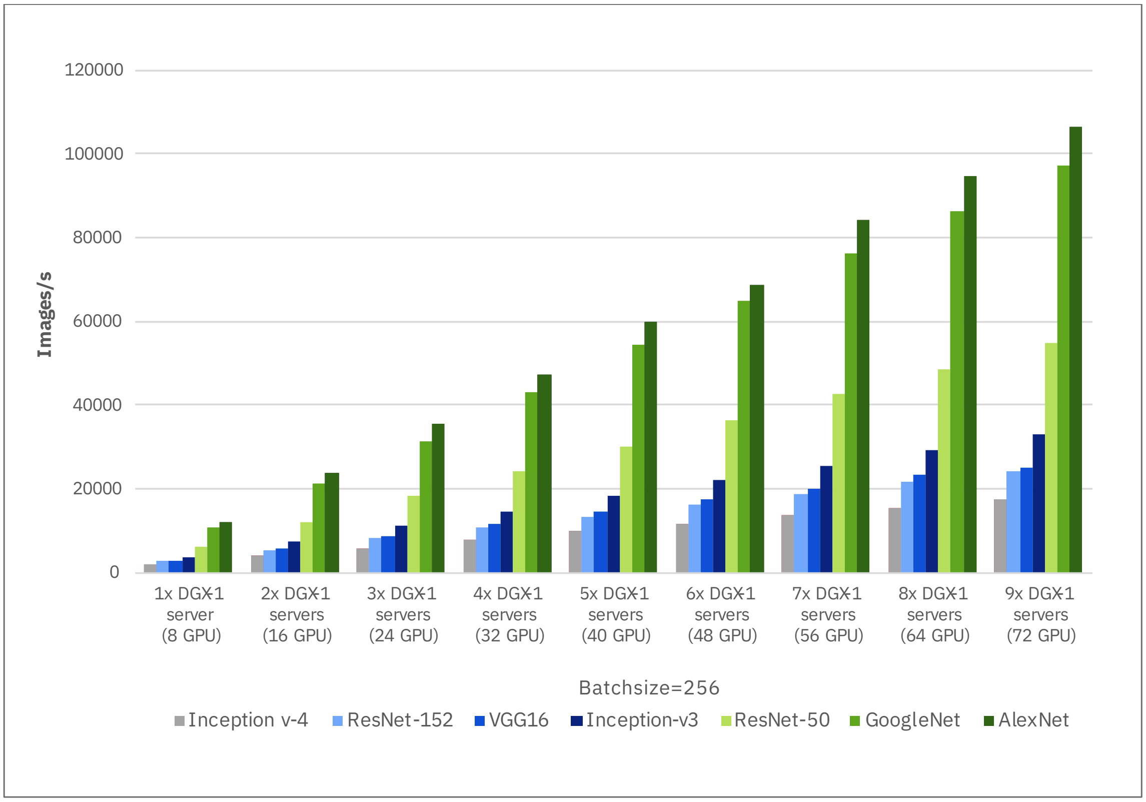

***Figure 19**: Training rates with TensorFlow models and multiple DGX-1 systems*

The image processing rates with multiple DGX-1 systems (Figure 19) demonstrate the scalability of training application performance with ESS 3000 with Inception-v3, ResNet-50, GoogLeNet, and AlexNet models using eight GPUs on each DGX-1 system. The ESS 3000 solution shows linear scale-up when adding additional DGX-1 systems beginning with one DGX-1 system with eight GPUs and ramping up to nine DGX-1 systems with consistent full saturation of GPUs for all 72 GPUs tested.

Inference Results – Multiple DGX-1 Systems

For multiple DGX-1 systems with separate containers, ESS 3000 demonstrates linear scale-up to full saturation of all the DGX-1 system GPUs simultaneously running, from one to nine DGX-1 systems, for a total of 72 GPUs providing the following performance results of aggregate images/sec inference using ImageNet datasets for the models shown in Figure 20.

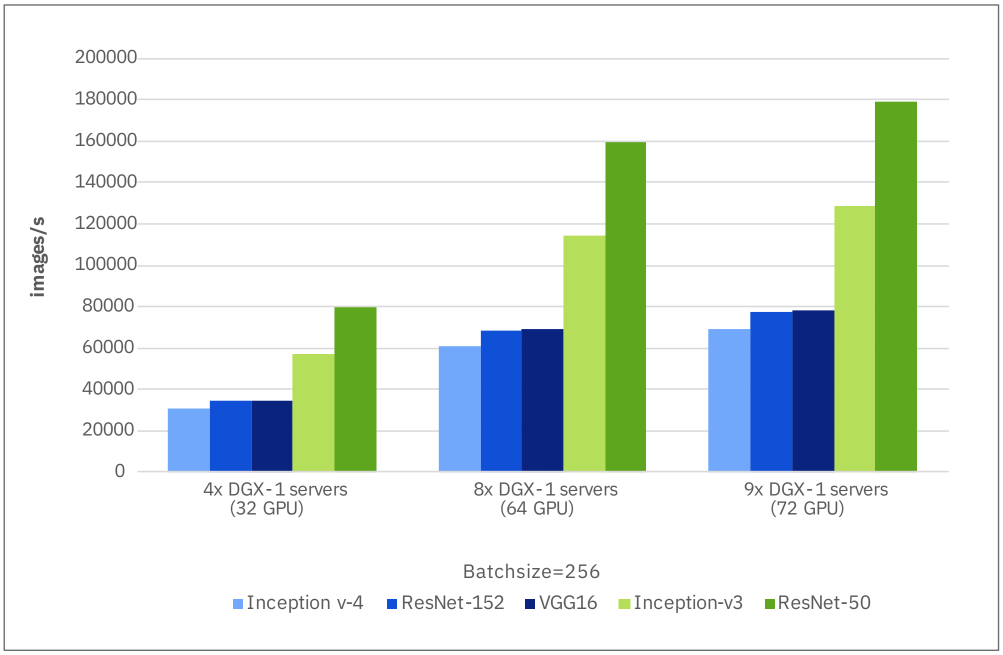

Figure 20*: Inference rates for multiple DGX-1 systems with TensorFlow models*

As tested, inference image processing rates are between 1.5x to almost 4x the training rates of the corresponding TensorFlow models. The DGX-1 system with ESS 3000 solution provides data scientists the ability to run in mixed training and inference mode on a single DGX-1 system as needed, dedicating one or two GPUs to inference and the remaining GPUs in the DGX-1 system to training jobs.

Training Results – Multiple DGX-2 Systems

For multiple DGX-2 systems with separate containers per GPU, ESS 3000 demonstrates linear scale-up to full saturation of all 16 DGX-2 system GPUs simultaneously running from one to three DGX-2 systems for a total of 48 GPUs. This provides the following performance results of aggregate images per second training throughput with ResNet-50, ResNet-152, Inception-v3, Inception-v4, and VGG-16 models and also compares training runs between synthetic (artificial test data is generated inside GPUs, no CPUs or data transfer involved), IBM Spectrum Scale filesystem, and local RAM disk only runs.

As shown in Figure 21, ESS 3000 effectively feeds the DGX-2 system GPUs in the Spectrum Scale cluster, keeping the DGX-2 systems fully saturated with data for maximum training processing capabilities for all models.

Model training results for the DGX-2 systems also show that overall there is minimal to no penalty when comparing Spectrum Scale file system performance to internal system RAM disk performance for DL models with GPUs.

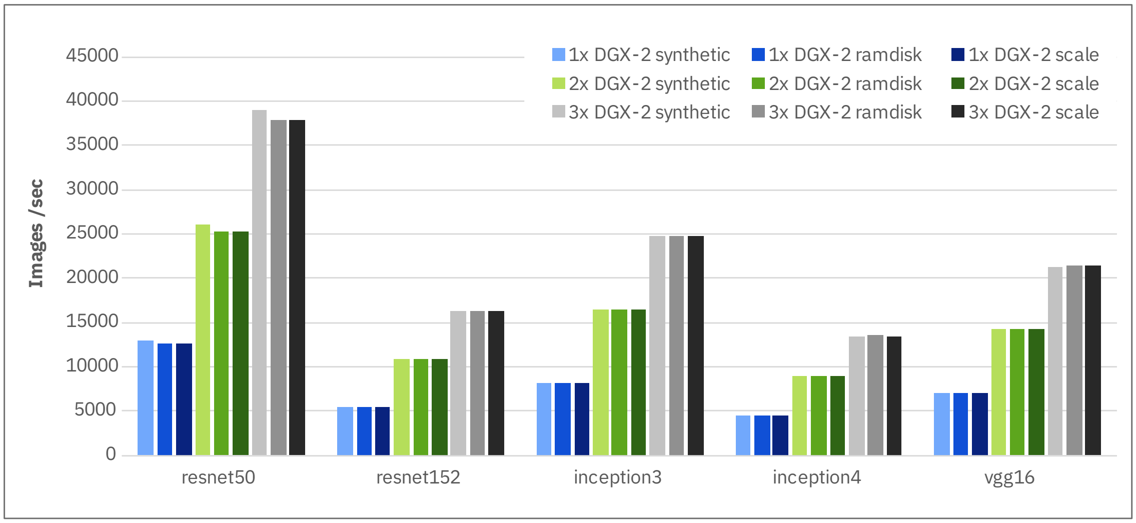

***Figure 21**: Model-to-GPU performance, and Spectrum Scale filesystem vs. synthetic vs. RAM disk performance with DGX-2 systems*

As shown previously with DGX-1 systems, with DGX-2 systems some models scale up with near linearity as the number of systems/GPUs increase while others present a consistent non-linear scale-up pattern whether using ESS 3000 storage, synthetic (artificial test data is generated inside GPUs, no CPUs or data transfer involved), or local RAM disk. This indicates that the scalability in these cases is not constrained by storage IO whether local or shared storage, but rather by a pattern of the DL model scalability within the compute infrastructure itself.

Overall, the ESS 3000 solution shows near linear scale-up when adding additional DGX-2 systems beginning with one DGX-2 system with sixteen GPUs and ramping up to three DGX-2 systems with consistent full saturation of GPUs for all 48 GPUs tested.

Inference Results – Multiple DGX-2 Systems

For multiple DGX-2 systems with separate containers, ESS 3000 demonstrates near linear scale-up to full saturation of all the DGX-2 system GPUs simultaneously running, from one to three DGX-2 systems for a total of 48 GPUs providing the following performance results of aggregate images/sec inference using ImageNet datasets for the models shown in Figure 22.

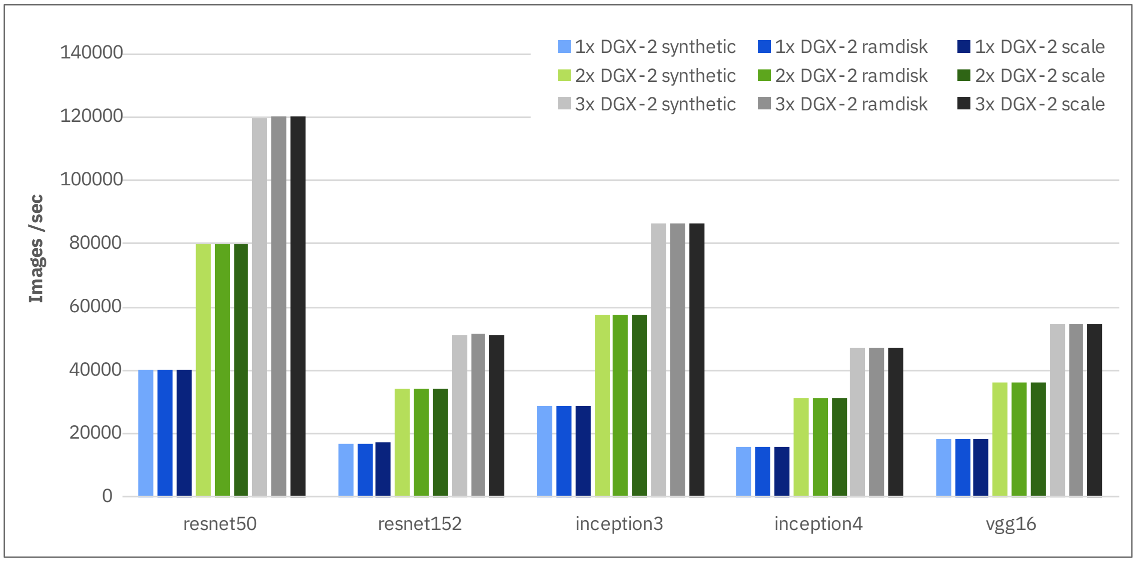

***Figure 22**: Inference rates for multiple DGX-2 systems with TensorFlow models and synthetic RAM disk or Spectrum Scale*

Similar to DGX-1 systems, inference image processing rates for DGX-2 systems are between 1.5x to almost 4x the training rates of the corresponding TensorFlow models regardless of synthetic, RAM disk, or Spectrum Scale filesystem storage types. And once again similar to the DGX-1 system, test results show that the DGX-2 system with ESS 3000 solution provides data scientists the ability to run in mixed training and inference mode on a single DGX-2 system as needed, dedicating the desired number of GPUs to inference and the remaining GPUs in the DGX-2 system to training jobs.

Enterprise Data Pipeline

The productivity of the data science team depends on the ready availability of the latest development frameworks, ample compute power and data accessibility. While performance is important, it is not the only consideration. Data preparation and ingest can consume most of the AI development timeline. For each project, data must be extracted from other sources and properly organized so that it can be used for model training. Once a model is developed, the data must be retained for traceability. The value of data grows with diverse use across multiple users, systems, and models. Data scientist productivity depends on the efficacy of the overall data pipeline as well as the performance of the infrastructure used for running AI workloads. Moreover, the underlying storage and network technologies play a crucial role in both these workflow aspects.

As organizations move from prototyping AI to deploying it in a production environment, the first challenge is to embed AI into the existing enterprise data pipeline or build a data pipeline that can leverage existing data repositories. A typical such enterprise data pipeline with storage requirements for each stage in the pipeline is represented in Figure 23.

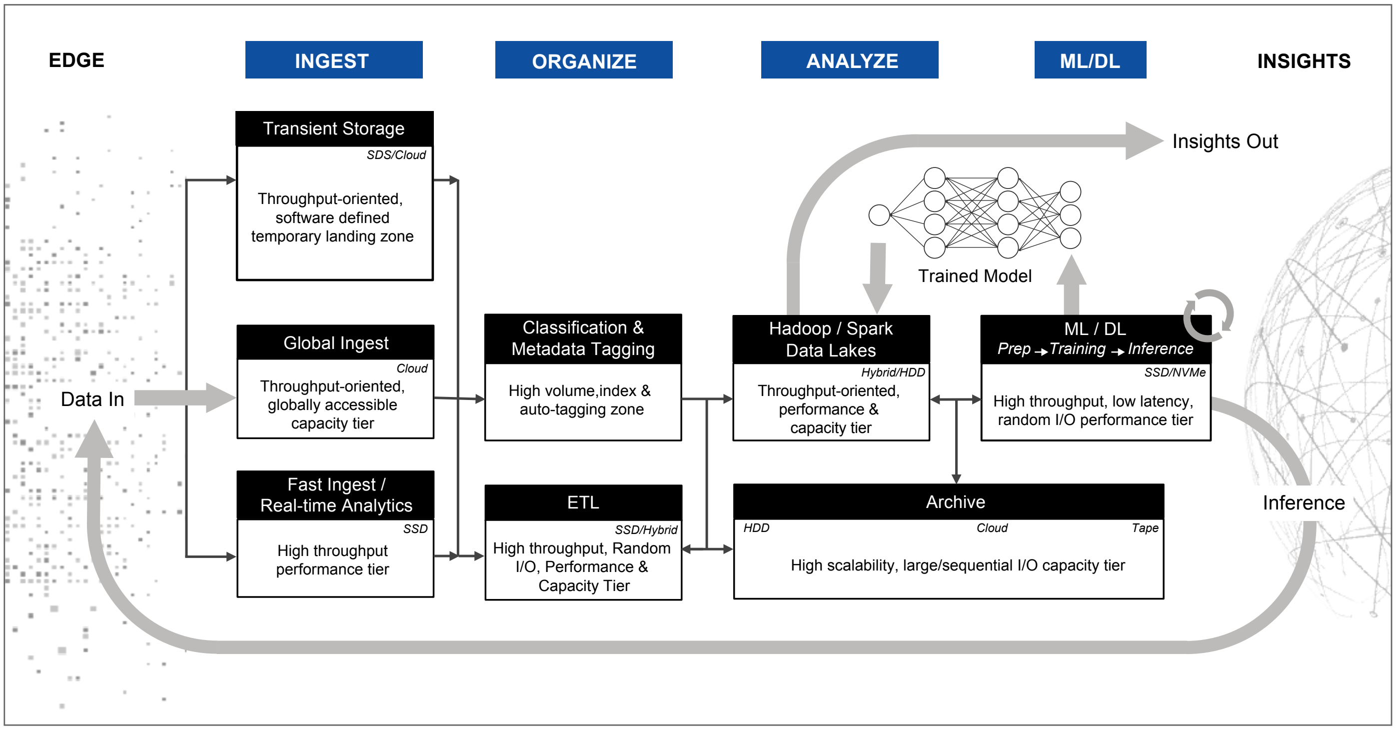

Figure 23: *An enterprise data pipeline with storage requirements*

IBM delivers a comprehensive portfolio of software defined storage products (Figure 24) that enable customers to build their enterprise data pipelines with the right performance and cost characteristics for each stage. This includes IBM Cloud Object Storage for global ingest and geographically dispersed repositories, a variety of different Elastic Storage Server models, and Elastic Storage System 3000 powered by IBM Spectrum Scale for high performance file storage and scalable common data lakes. Additionally, IBM Spectrum Archive enables direct file access to data stored on tape for inexpensive archives.

A new addition to this portfolio is IBM Spectrum Discover. Spectrum Discover is metadata management software that provides data insight for exabyte scale unstructured storage both on premises and in the cloud. Spectrum Discover easily connects to IBM Cloud Object Storage, IBM Spectrum Scale, and to third party storage through NFS and S3 interfaces to rapidly ingest, consolidate, and index metadata for billions of files and objects. IBM Spectrum Discover fulfills a leading role in the data classification phase of the overall data pipeline, but also provides capabilities that support data governance requirements and enable storage optimization along the pipeline.

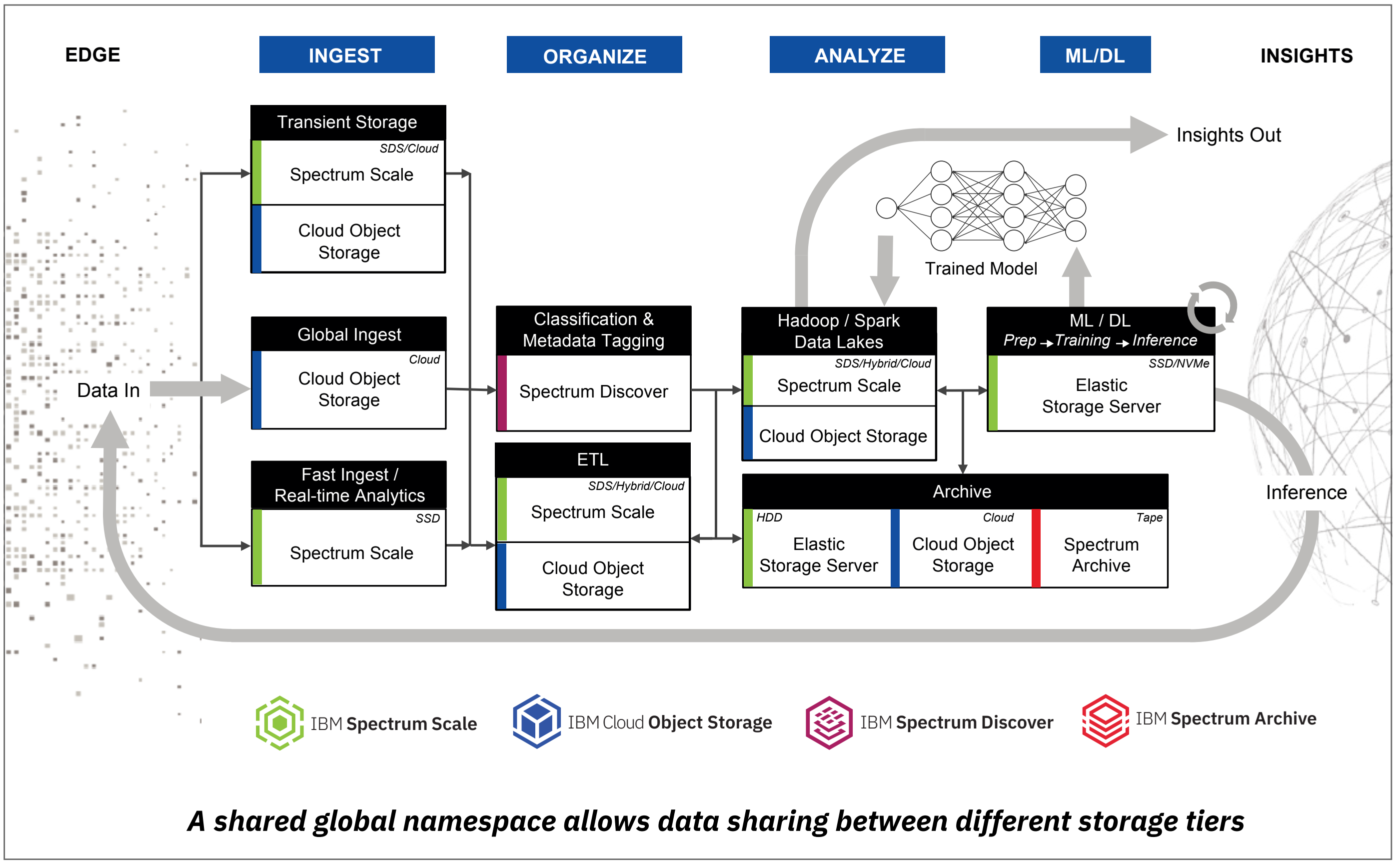

***Figure 24**: IBM solutions for the AI data pipeline*

Customers who take this end-to-end data pipeline view when choosing storage technologies can benefit from improved data governance. Data science teams benefit from rapid time-to-insight with a minimum number of data copies. Infrastructure teams benefits from simplified management, scalability and improved TCO.

If we zoom into this pipeline, represented by Figure 25 below, to focus more on the ML/DL stage, this phase is an iterative process itself. Once a trained neural network model is developed, it needs to be tested and retrained continuously to keep it current and to improve its accuracy. As projects grow beyond the first test systems, appropriate storage and networking infrastructures are needed so these ML/DL systems can sustain the growth and eventually deliver the required insights to make business decisions.

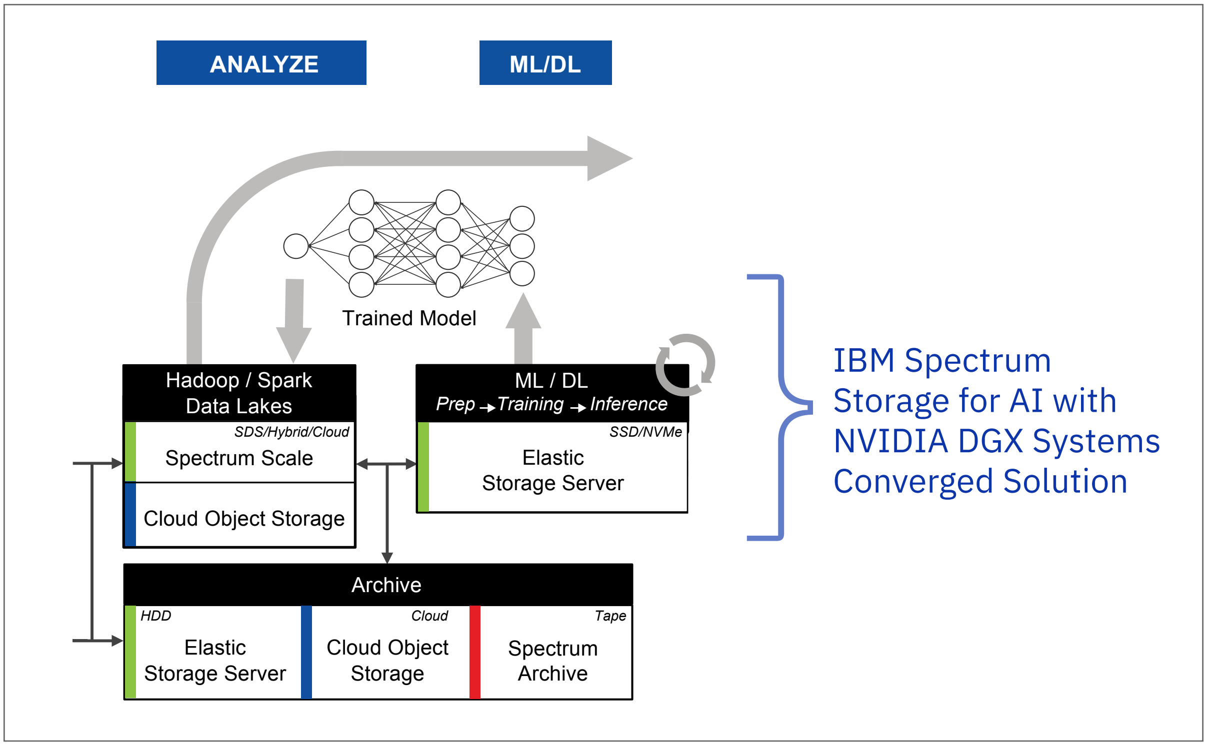

Figure 25: *IBM Spectrum Storage for AI with NVIDIA DGX System(s) Converged Solution for ML/DL stage*

The design goal for the IBM Spectrum Storage for AI with NVIDIA DGX Systems converged solution is to reduce the complexity and time required to plan and execute growth for these ML/DL sub-stages (data prep, model training, and inference).

Conclusion

The IBM Spectrum Storage for AI with ESS 3000 architecture over a Mellanox EDR InfiniBand fabric provides leading-edge performance with DGX-1 and DGX-2 systems for DL workload training and inference and servicing high performance parallel processing with high bandwidth and low latency for full utilization of GPUs when running on multiple DGX-1 or DGX-2 systems as our tests have demonstrated.

IBM Spectrum Storage for AI and the NVIDIA DGX POD systems integrated with the NGC Software Stack in combination with the IBM Spectrum family of software defined storage solutions 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.

Additional Resources

Introduction Guide to the IBM Elastic Storage Server
http://www.redbooks.ibm.com/redpapers/pdfs/redp5253.pdf

NVIDIA DGX SuperPOD Delivers World Record Supercomputing to Any Enterprise
https://devblogs.nvidia.com/dgx-superpod-world-record-supercomputing-enterprise/

NVIDIA DGX-2 POD: From Concept to World-Record Setting Supercomputer in Three Weeks
https://www.nvidia.com/en-us/data-center/resources/nvidia-dgx-superpod-reference-architecture

Acknowledgements

We gratefully acknowledge the contributions to the IBM Spectrum Scale Reference Architecture by our partners from NVIDIA, Jacci Cenci, Satinder Nijjar, Darrin Johnson, Robert Sohigian, and Tony Paikeday. We would also like to thank the team at Arrow Electronics.

We could not have completed this study without the support and guidance of our IBM team members Dave McDonnell, Brian Porter, Constantine Arnold and Udayasuryan A Kodoly. Our sincere appreciation and thanks to all these individuals, who provided insight and expertise that greatly assisted in the research for this paper.

IBM Systems
3039 Cornwallis Road
RTP, NC 27709