NVIDIA and IBM Global Data Fabric with ESS and Spectrum Scale

Highlights

- AI solutions are easier with IBM Global Data Fabric
- NVIDIA solutions run faster with ESS3200 80GB/s nodes
- Solutions are easy to deploy and can start with just one node
- Scales to 1000s of nodes
- Lowers resource requirements over 50% with GPU direct storage
- Flexibility with access via S3, NFS, SMB, GDS, POSIX, HDFS, CSI
- Global data protection for applications and containers
- Proven TB /s high performance throughput to avoid bottlenecks
- Container Native interface for high performance AI workloads

IBM Spectrum Scale and ESS and NVIDIA

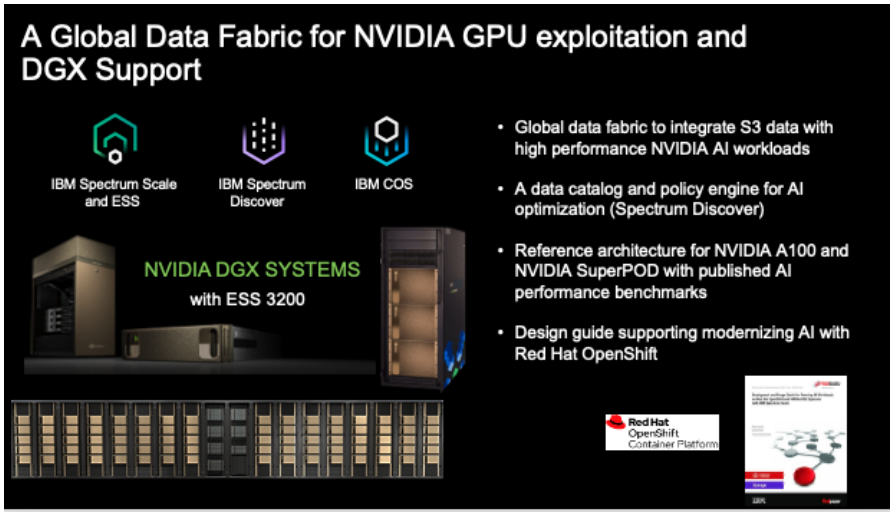

IBM Storage and NVIDIA

IBM Storage for Data and AI provides a high-performance global data platform that solves modern data challenges with a single source of truth and a global data fabric. IBM innovation and industry-leading software-defined storage management and data services provides workload agility, faster insights for AI and business results. Our storage connects multiple applications that each may contain different interfaces with multiple data sources in a single managed solution. Finally, you modernize without compromise as efficiency, protection and security are all designed to meet global enterprise requirements and current

investments with your infrastructure.

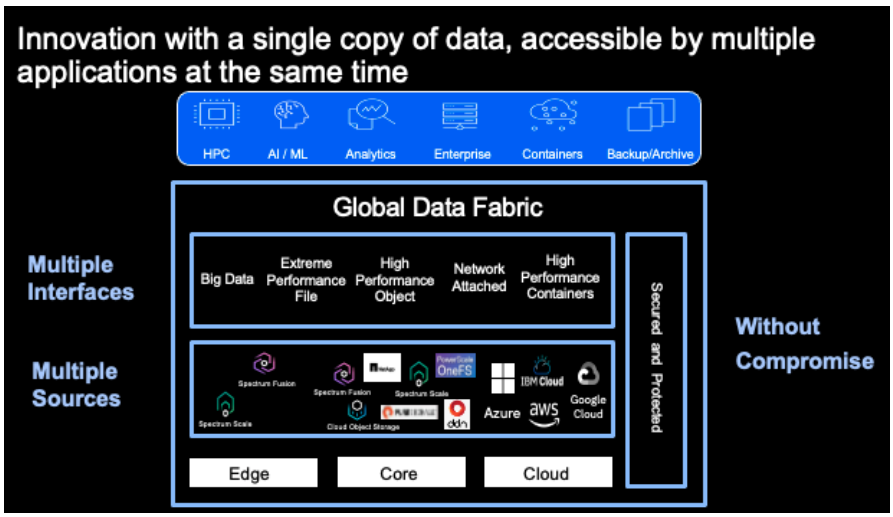

IBM Global Data Fabric

We provide the innovation with a single copy of data that is accessible by any application at the same time. We provide multiple interfaces from multiple sources that is all secured and protected from edge to core to cloud.

Our single source of truth provides value to the organization beyond the basic storage capabilities. We call our innovation the the Global Data Fabric.

The value we bring to NVIDIA solutions is we:

Eliminates duplicate data, simplify workflows, and lower cost as global access means no need to copy data from external sources

Improves application agility as global data access enables applications to run anywhere

Speeds AI results and accuracy because applications can be connected to more data with a single copy whether on the edge, core, or cloud,

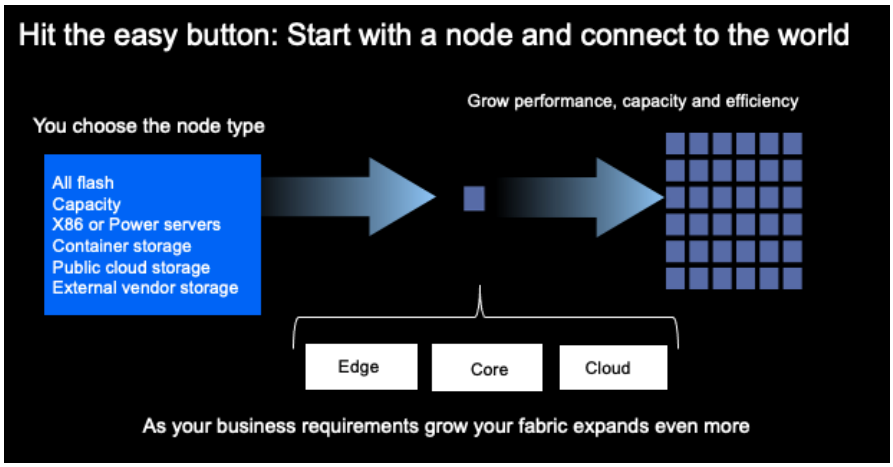

Hit the easy button

IBM Storage enables IBM's customers to start small with affordable solutions for early experimentation, and then expand the performance and capacity to support production AI, analytics, and commercial applications at virtually unlimited scale.

The smallest, entry level model of the IBM Elastic Storage System (ESS) 3200 is available with 48 terabytes (TB) and can scale out to suit any customer's needs, up to 8 yottabytes (YB) of capacity. The ESS 3200 can be deployed as part of an ESS cluster with existing ESS servers, including other ESS flash, hard disk drives (HDD) and hybrid models, and the ESS 3200 can be inserted into a Spectrum Scale environment that does not yet have ESS as part of the environment.

The ESS 3200 features all-NVMe solid-state drives (SSDs) for blisteringly fast performance, which scales linearly: 40 gigabytes per second (GB/sec) in a 2U form factor. Each additional ESS 3200 aggregates the performance and can add another 80 GB/sec of overall throughput to the Spectrum Scale system.

The ESS 3200 will deliver 80 GB/sec read sequential throughput under ideal conditions. Each ESS 3200 can deliver this performance, and by leveraging Spectrum Scale, the ESS 3200 aggregates parallel streaming performance of as many ESS 3200s as you have the network bandwidth for.

IBM can support up to 18 or 19 ESS 3200s in a single IBM 7965 rack. With up to 912 TB of usable NVMe extreme performance storage in a 2U24 form factor, a single rack ESS 3200 deployment with 20 ESS 3200s can deliver nearly 18.2 petabytes (PB) of usable storage and 1600 GB/sec sustained sequential workload in one NVMe-based extreme performance file storage system.

IBM Storage and Red Hat OpenShift and NVIDIA

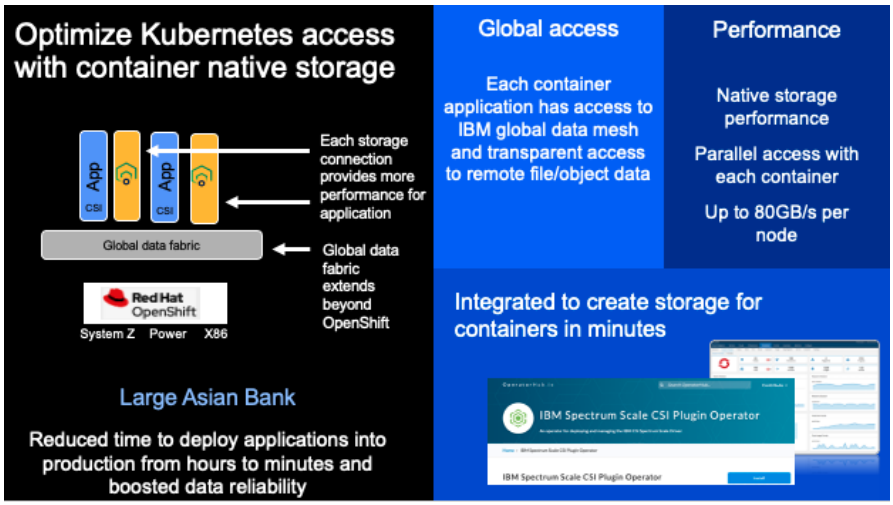

Kuberneties and IBM Storage

With IBM Spectrum Scale container native storage, DevOps or storage admins can easily configure storage for containers. Application performance and storage resources can be optimized as a Spectrum Scale container native storage can be deployed for many containers, or a container can have a dedicated storage interface to control the parallel access to the data. This provides higher performance with less server resources and leverages the high parallel throughput of IBM Spectrum Scale and Elastic Storage System nodes nodes. There are no capacity limitations on the local system and performance is not affected by the application server but is optimized by external storage resources. IBM Spectrum Scale provides the benefits of container native storage without the limits of using only local disk and server resources. IBM Spectrum Scale increases business agility for applications and helps eliminate data slios with a global data mesh that enables applications to access both file and object resources on the edge, in the core or even in the cloud.

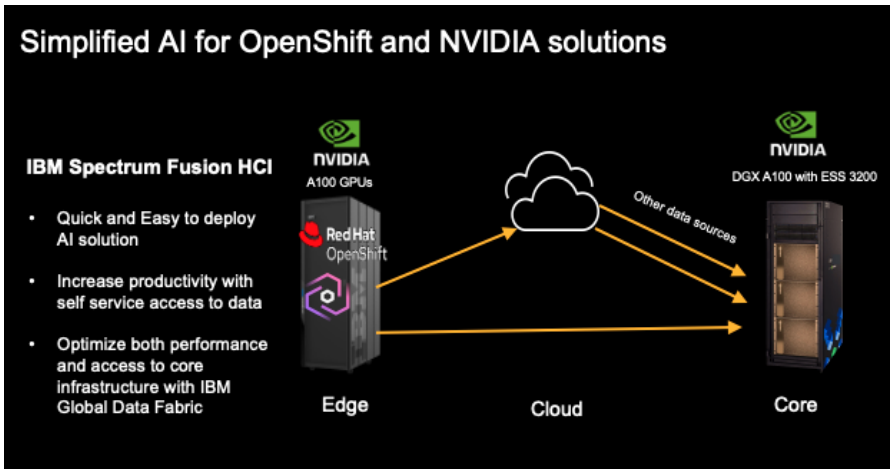

Simplified AI with OpenShift and container workloads

IBM Spectrum Fusion can also be used to optimize AI workflows with modern container applications. Data can be processed on the IBM Spectrum Fusion HCI and then data can be transparently leveraged in the Cloud (S3 data) or on another IBM Spectrum Scale system leveraging the high-performance IBM ESS NVMe flash system. If a customer is using NVIDIA for AI IBM Spectrum Scale can preprocess or process at the edge with NVIDIA A100 GPUs and `NVIDIA Red Hat Operator and tools and then transparently leverage that data in the cloud or directly on another larger more compute intensive NVIDIA DGX A100 system with the IBM ESS high performance storage system.

IBM Storage and NVIDIA

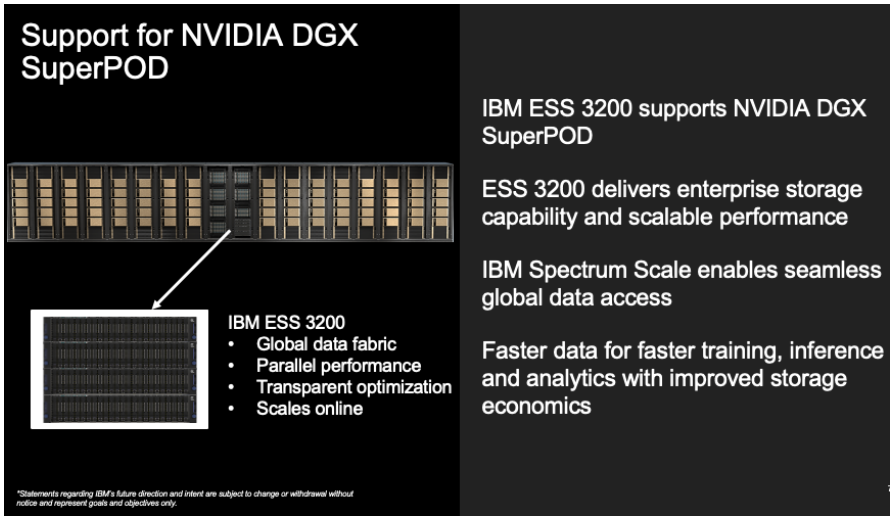

IBM ESS and SuperPOD

SuperPOD is NVIDIA's next generation cloud-native, multi-tenant AI supercomputer. IBM is announcing that we plan to support SuperPOD with ESS 3200 in 3Q21.

ESS 3200 is highly scalable: adding systems linearly adds throughput, which is a great match to the SuperPOD architecture,

Based on IBM Spectrum Scale, ESS is part of IBM's hybrid cloud strategy and so provides seamless access to all an organization's data around the world with enterprise storage services. This enables businesses to leverage the power of SuperPOD across all of their enterprise data.

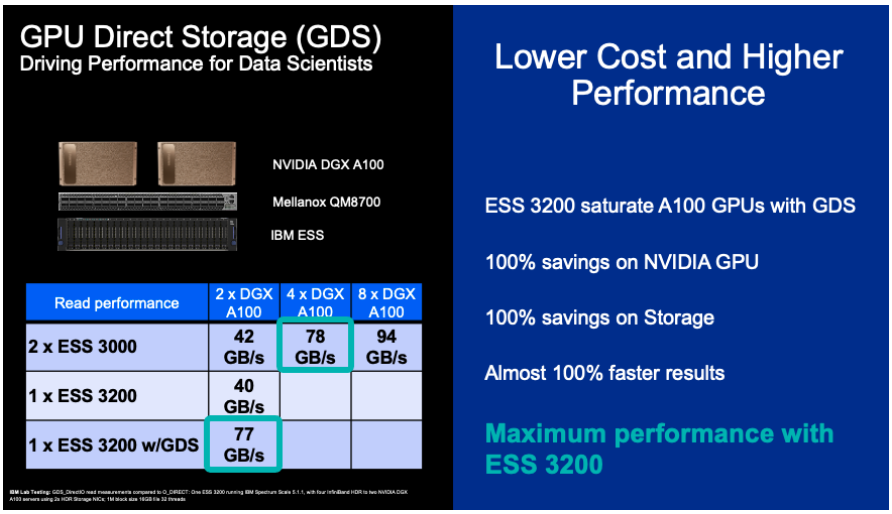

IBM and GPU direct storage

GPUDirect Storage is a technology from NVIDIA and Mellanox that enables a more direct path between GPUs and storage (RDMA). The result is higher bandwidth and lower overheads, which enables the (costly) GPU to do more AI, ML and analytics work.

In IBM measurements, we see that only one ESS 3200 with a pair of DGX A100 GPUs using GDS delivered the same data throughput for AI/ML as two ESS 3000s and four DGX A100s without GDS.That means ESS 3200 delivers enough throughput to saturate A100 GPUs using GDS to do the maximum amount of work, which means faster AI/ML results at lower cost.

To help customers with their deployments, IBM is updating our reference architecture for ESS with 2, 4, and 8-node configurations.

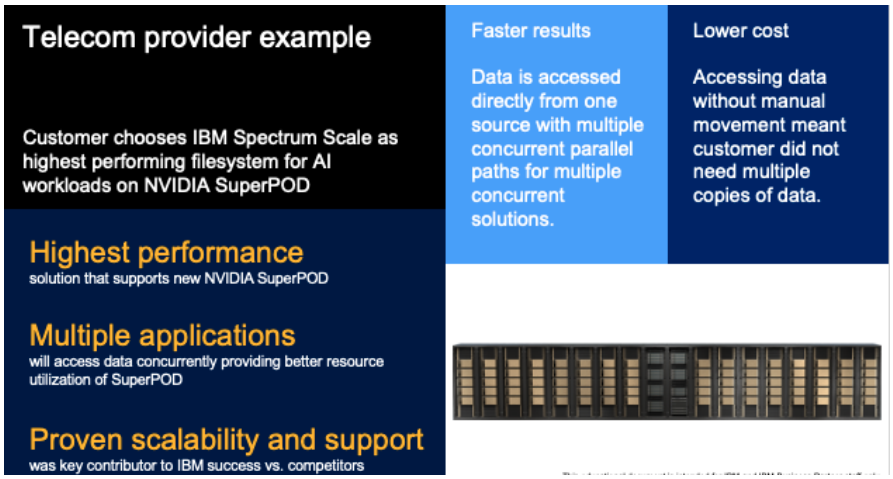

A customer story

A recent customer chose IBM Spectrum Scale as highest performing filesystem for AI workloads on NVIDIA SuperPOD. The value of the IBM solution was simple:

•Highest performance solution that supports new NVIDIA SuperPOD

•Multiple applications will access data concurrently providing better resource utilization of SuperPOD

•Proven scalability and support was key contributor to IBM success vs. competitors

Why IBM?

The value of Spectrum Scale and the Elastic Storage System with NVIDIA is simple: High-performance parallel data access with a single global data fabric connecting edge to core to public cloud in a single cluster with simple scalable building blocks optimized for maximum throughput, low latency and cost optimization. This makes Spectrum Scale second to none for AI and Big Data Analytics, HPC, and most any High-Performance Workload.

Next steps

→ Learn more about IBM Data and AI solutions

→ Reference architecture with IBM Storage and NVIDIA

→ Blog with IBM Storage and NVIDIA

For more information

For more information about our Global Data Fabric products

IBM Spectrum Scale

IBM Elastic Storage System

IBM Cloud Object Storage

IBM Spectrum Discover

IBM Spectrum Fusion

This document contains information pertaining to the following IBM products which are trademarks and/or registered trademarks of IBM Corporation:
IBM Spectrum Scale

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