IBM Storage and SDI

AI Ecosystem and Solutions with IBM Spectrum Scale

March 20, 2019

Piyush Chaudhary
piyushc@us.ibm.com

Spectrum Scale Development

Agenda

- IBM Spectrum Storage for AI with NVIDIA DGX
- IBM Spectrum Storage for AI with Power 9 and Watson Machine Learning Accelerator
- Client use case: DATIS – Private Cloud for Analytics/AI
- Business partner & IBM solution: CANCOM KFM

IBM Storage and SDI

IBM Spectrum Storage for AI with NVIDIA DGX

Converged Solution for Data Science Productivity

Introducing IBM Spectrum Storage for AI With NVIDIA® DGX

A Scalable, software-defined infrastructure powered by IBM Spectrum Scale and NVIDIA DGX-1 systems. IBM Spectrum Storage for AI with NVIDIA DGX is a powerful engine for your data pipeline.

The workhorse of an AI data infrastructure on which companies can build their shared data service.

High-Performance to feed the GPUs

- NVMe throughput of 120GB/s in a rack
- Over 40GB/s sustained random read per 2U

Composable to grow as needed

- Up to 9 DGX-1 servers (72 GPUs) in a rack
- Storage scale-out from a single 300TB node to 8 Exabytes and a Yottabyte of files

Extensible for the AI Data Pipeline

- Support for any tiered storage, including Cloud and Tape

NVIDIA DGX – What is it?

NVIDIA DGX-1:
THE ESSENTIAL TOOL OF AI

Highest Performance, Fully Integrated System

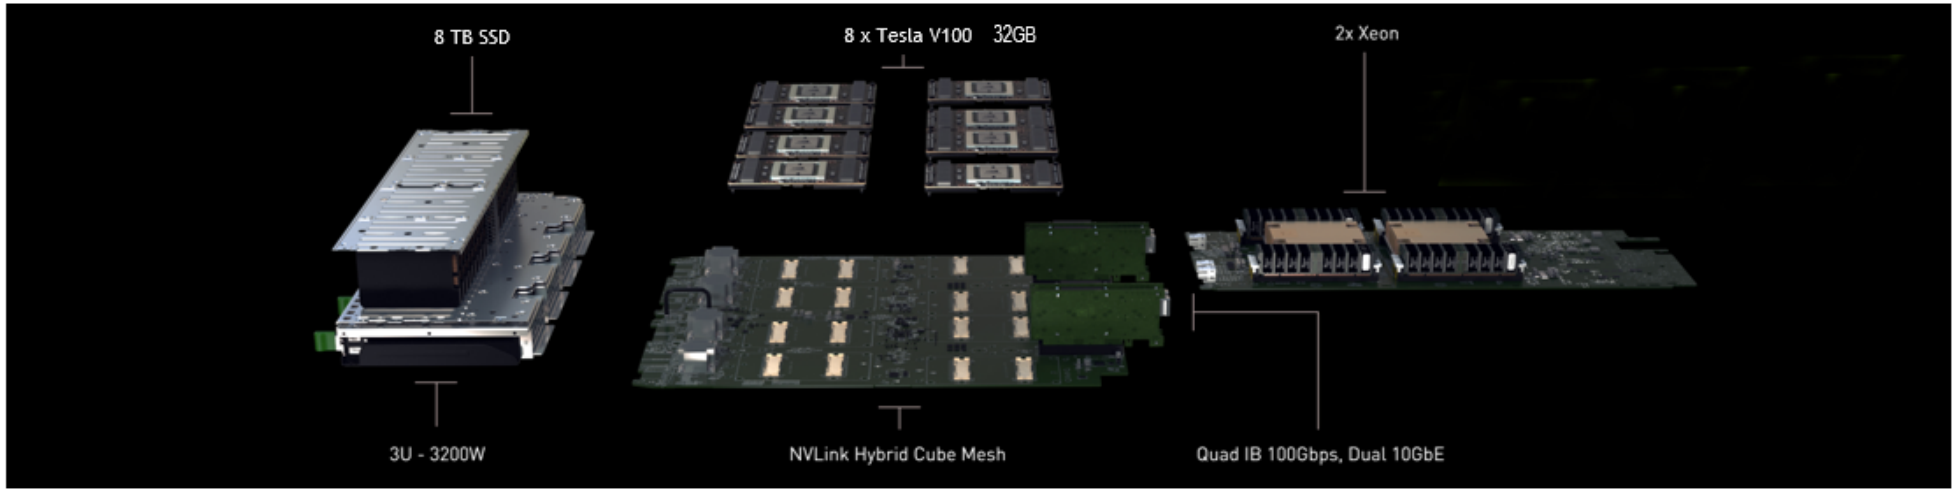

1 PFLOPS | 8x Tesla V100 32GB | 300 GB/s NVLink Hybrid Cube Mesh

2x Xeon | 8 TB RAID 0 | Quad 100Gbps, Dual 10GbE | 3U — 3500W

Why IBM Spectrum Storage for AI with NVIDIA DGX?

Challenges	The IBM Difference
Growth of data	Unrivalled scalability for the billions of files and objects in a single name space
Cost of data	SDS approach allows COTS hardware to be used for up to 10X cost savings and automated movement across tiers
Operational cost of managing data	SDS automates basic tasks and empowers the system administrator to manage data cost effectively at the Petabyte and Exabyte level
Performance	Parallel performance. Nothing is faster
Security	Built-in encryption and data protection
Simplicity	Integrates with Open Architectures like OpenStack, POSIX, NFS, etc
Ease of deployment	Predefined bundles. Simple GUI
Confidence	Over 4,000 clients and more than 100,000 systems

IBM Spectrum Storage for AI with NVIDIA DGX
System Specifications

Hardware:

- **DGX-1 or DGX-2 Servers** – purpose-built solutions for ML/DL integrating eight or sixteen NVIDIA Tesla V100 Tensor Core GPUs

- **IBM Spectrum Scale NVMe appliance** – 2U building block appliance built on FS9100/FAB3 chassis with IBM Spectrum Scale SW (Target GA mid 2019)

OR

- **IBM All Flash ESS** – Any IBM ESS GSxS models

- **Mellanox IBM Networking**

Software:

- **The NVIDIA DGX software stack** - optimized for maximized GPU-accelerated training performance, including the new RAPIDS framework to accelerate data science workflow

- **IBM Spectrum Scale v5** - the leading software-defined file storage, architected specifically for AI workloads with enhanced small file, metadata and random IO performance.

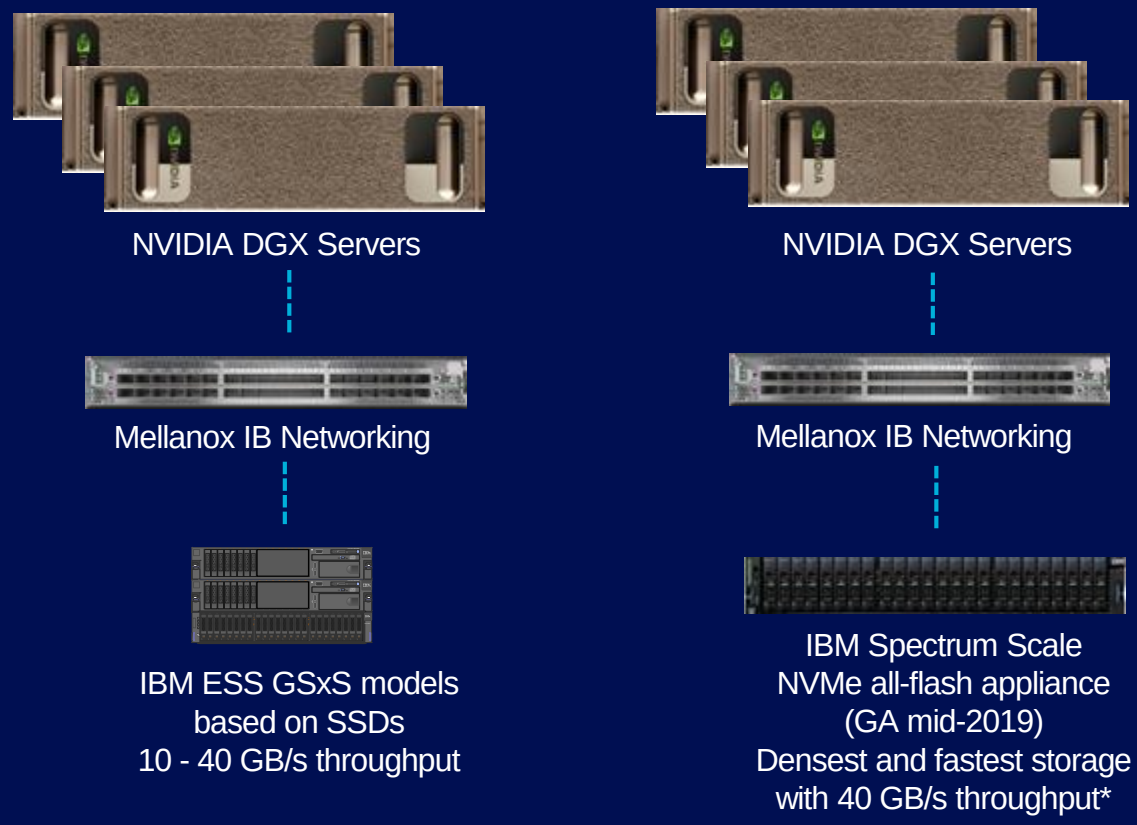

Note:
IBM will support any rightly sized and supported IBM Spectrum Scale configuration in this solution, while benchmark testing is done on a selected configuration mentioned in the reference architecture. Contact your IBM Storage and NVIDIA Deep Learning partner or IBM technical sales team if you need assistance in sizing the solution.

IBM Spectrum Storage for AI with NVIDIA DGX Architectures

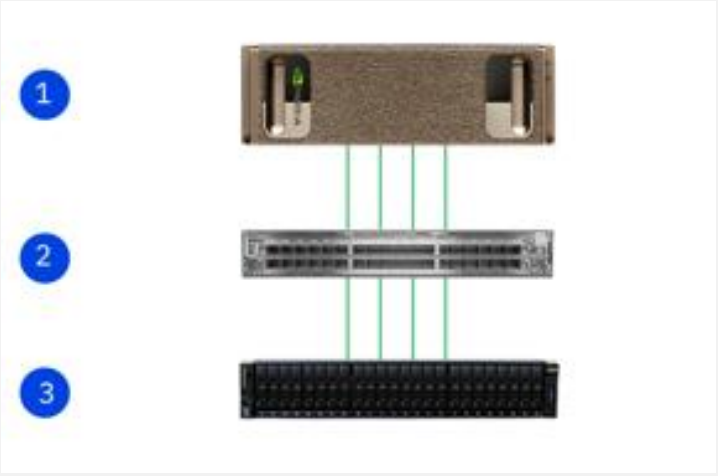

One DGX-1 server to one Spectrum Scale NVME all-flash appliance configuration with one Mellanox switch

Reference Architecture

1 NVIDIA DGX-1 server

2 Mellanox EDR IB network switch

3 IBM Spectrum Scale NVMe all-flash appliance

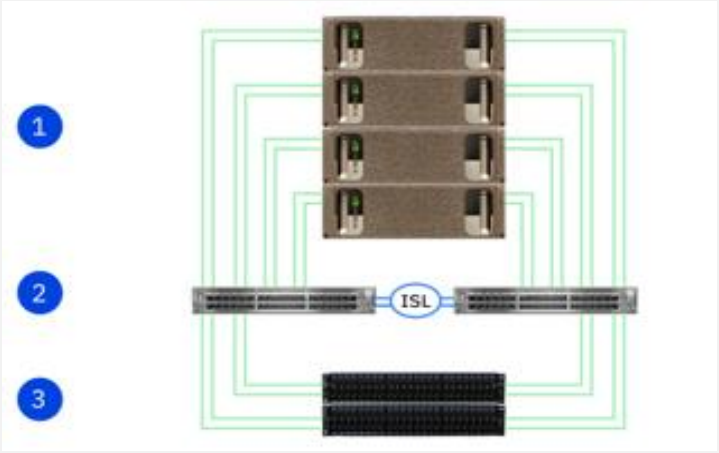

Four DGX-1 servers in a IBM two Spectrum Scale NVMe all-flash appliance unit configuration with two Mellanox switches

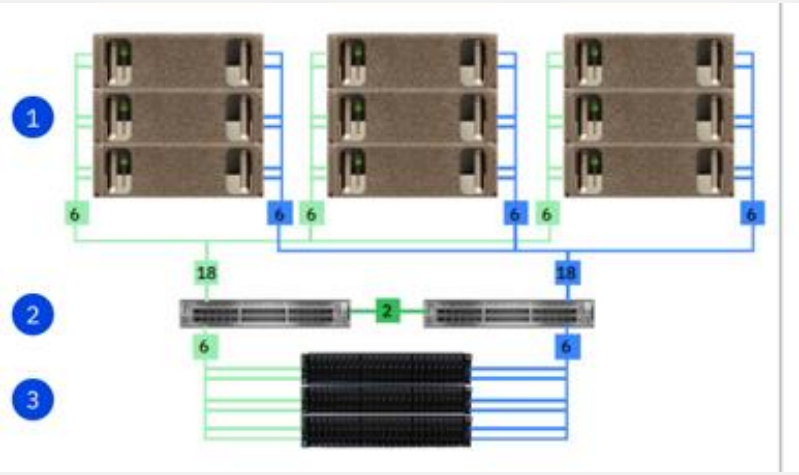

Nine DGX-1 servers in a IBM three Spectrum Scale NVMe all-flash appliance unit configuration with two Mellanox switches

IBM Spectrum Storage for AI with NVIDIA DGX Throughput Scaling

Near Linear Scaling by adding 40GB/s per 2U appliance

No need for downtime or reconfiguration

Best in class throughput potential

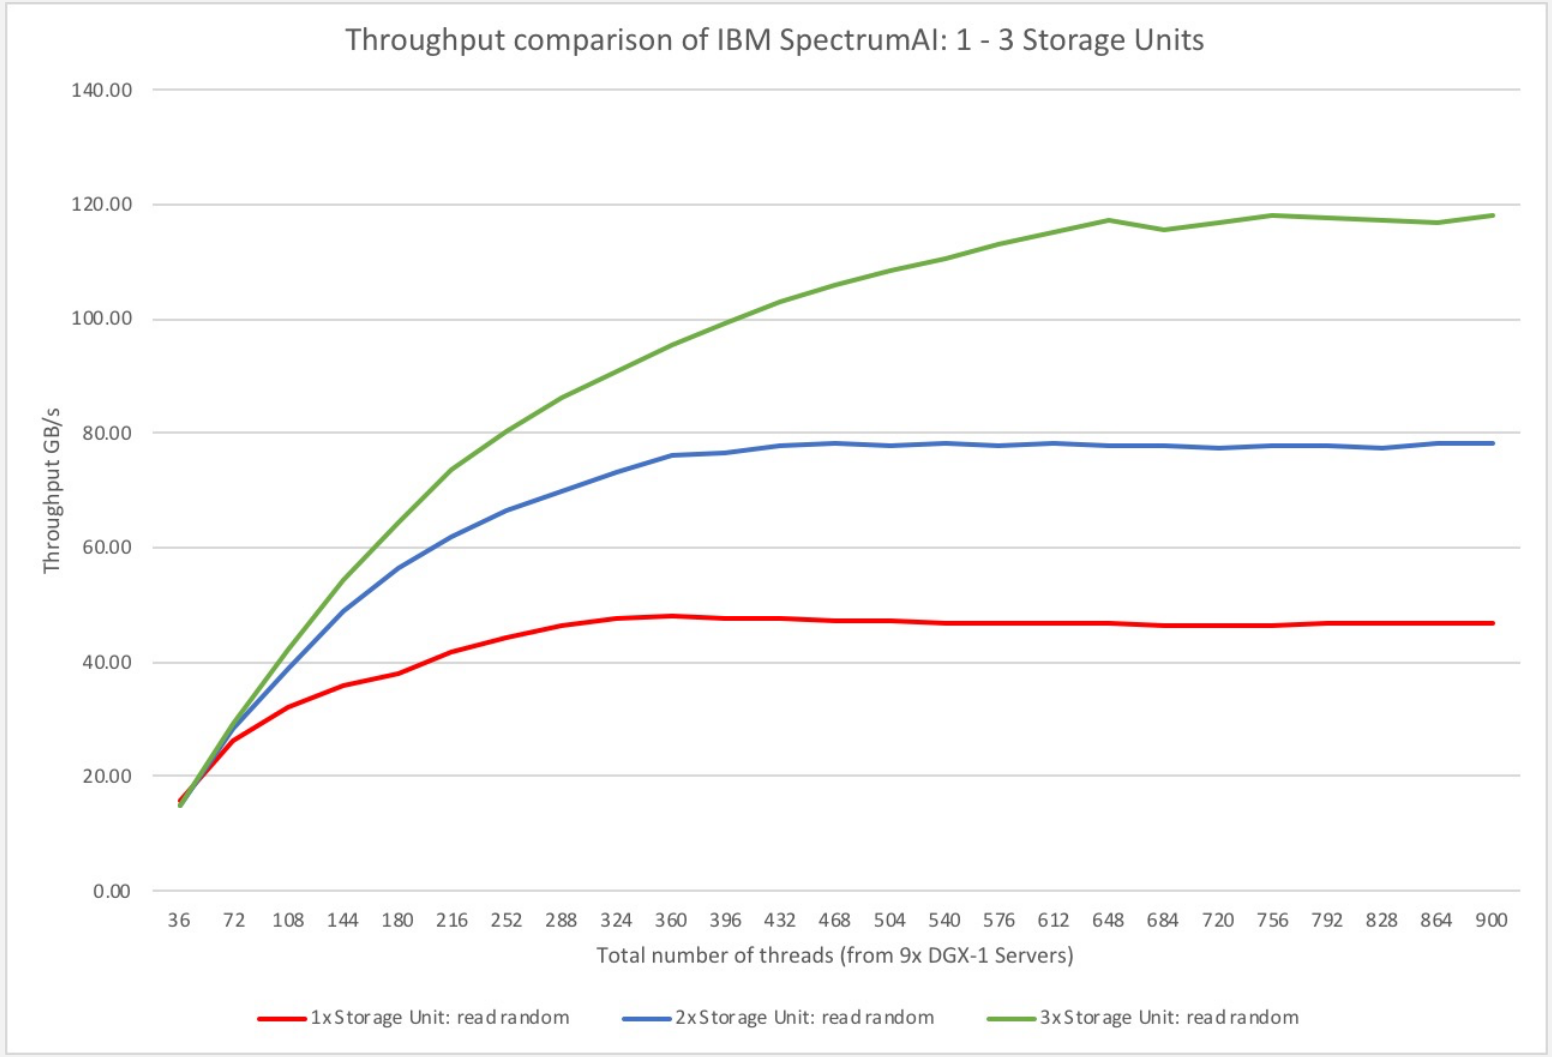

IBM Spectrum Scale v5 on NVMe Flash

Delivering the random read throughput needed for AI.

Sustained random read performance about 40 GB/s in 2U NVMe array

Sequential Reads take advantage of prefetch for highest peak performance

As workload expands, all data access patterns become random

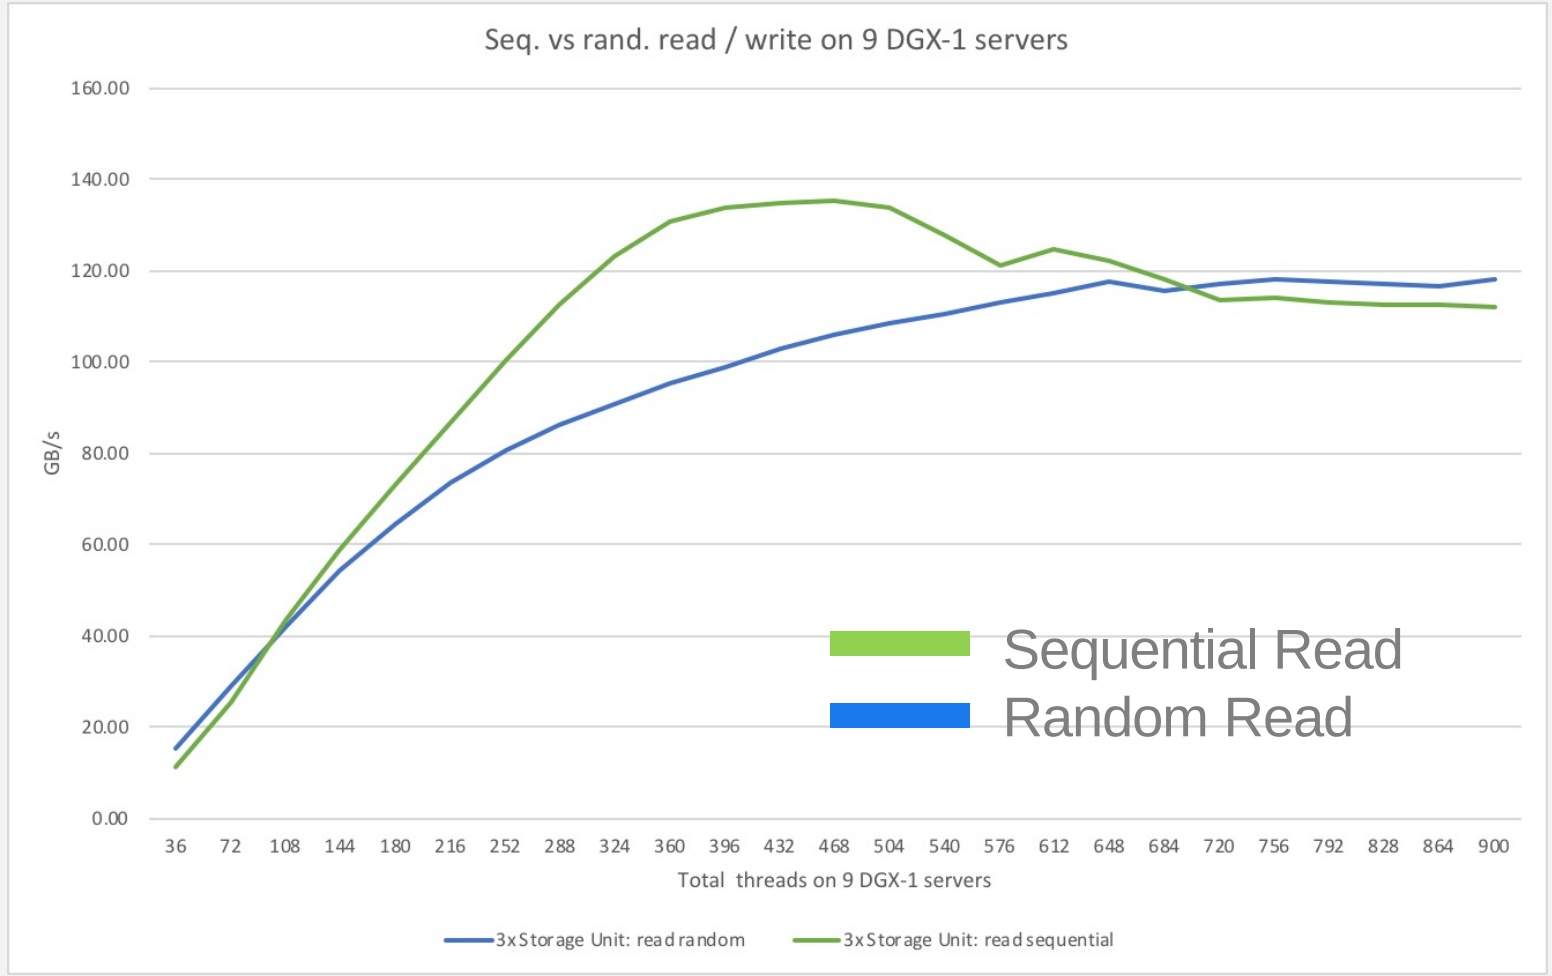

IBM Storage and SDI

IBM Spectrum Storage for AI with Power 9 and Watson Machine Learning Accelerator

IBM One AI

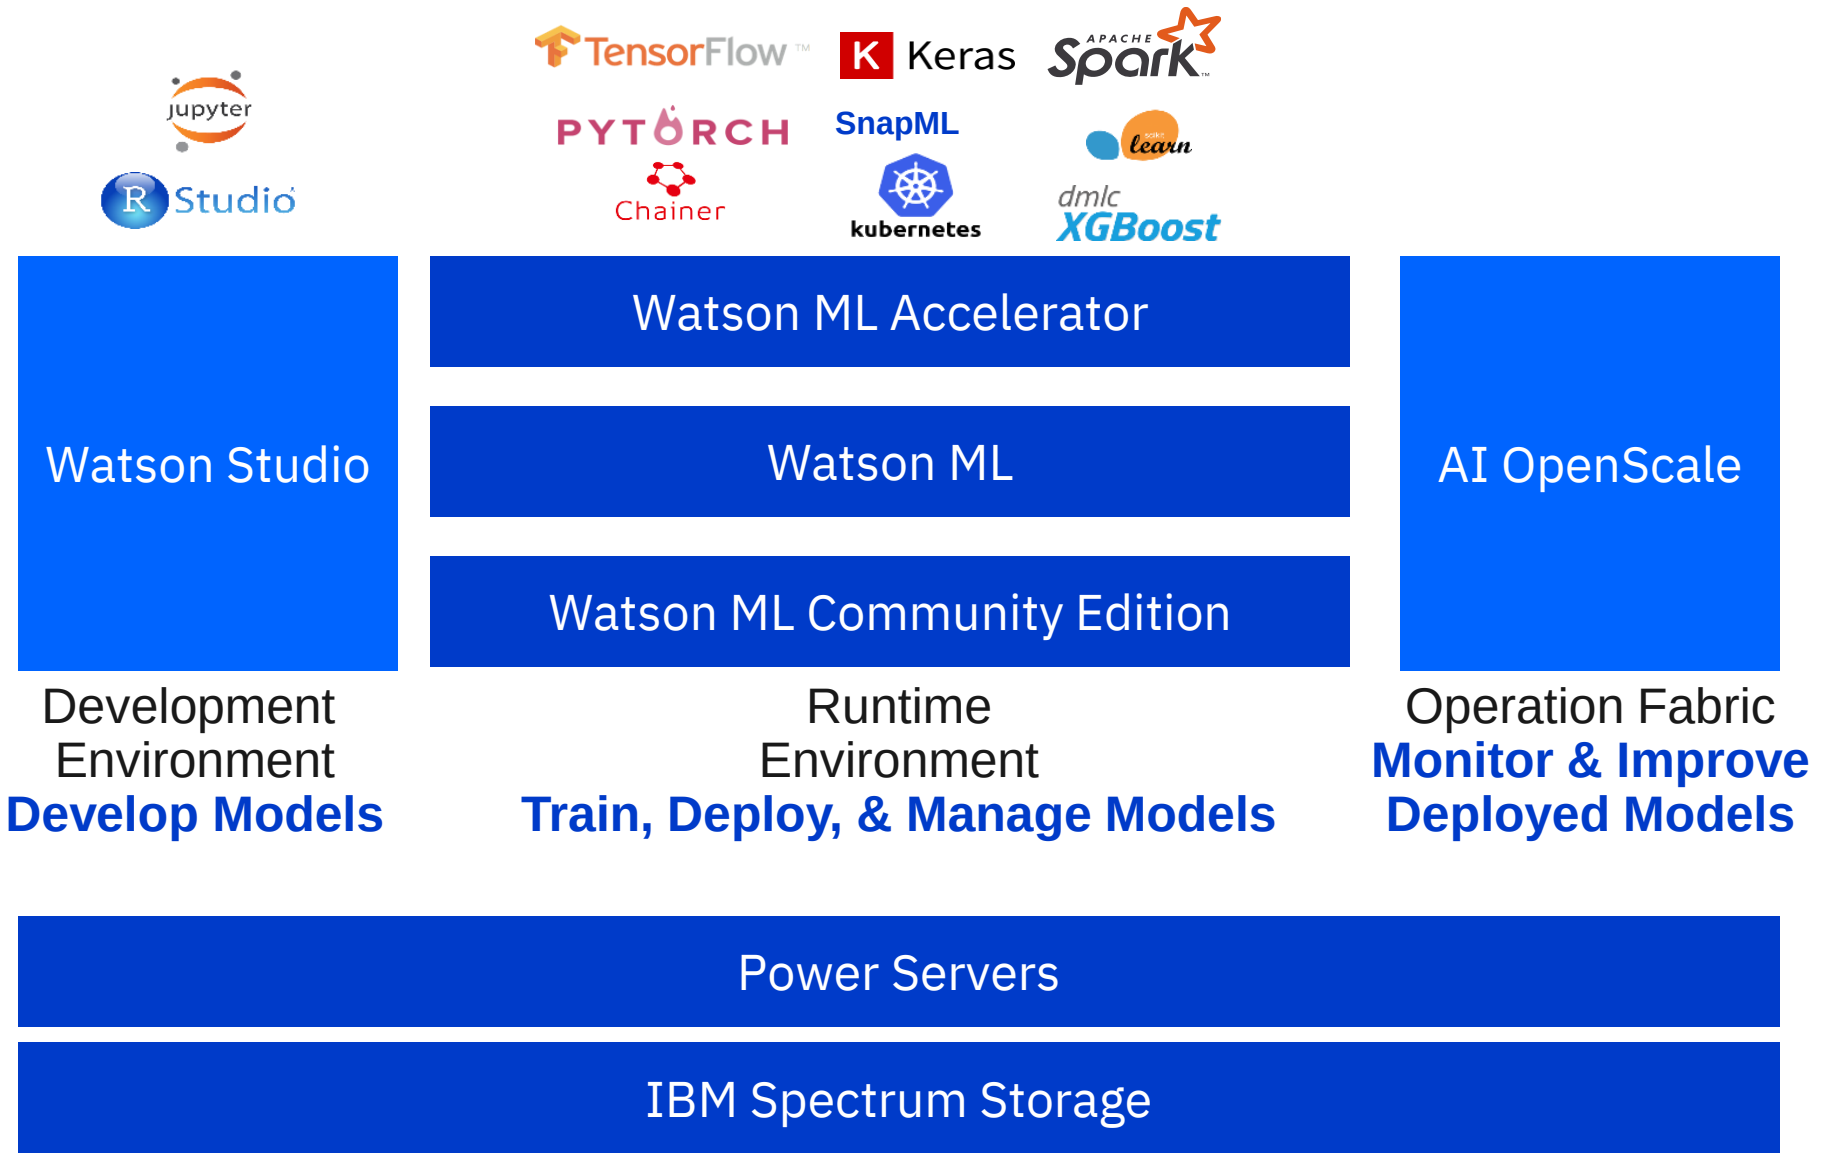
jupyter
R Studio
TensorFlow
Keras
Spark
PYTORCH
SnapML
learn
Chainer
kubernetes
dmlc XGBoost
Watson Studio
Watson ML Accelerator
Watson ML
Watson ML Community Edition
AI OpenScale
Development Environment
Develop Models
Runtime Environment
Train, Deploy, & Manage Models
Operation Fabric
Monitor & Improve Deployed Models
Power Servers
IBM Spectrum Storage

Summary & Proofpoints

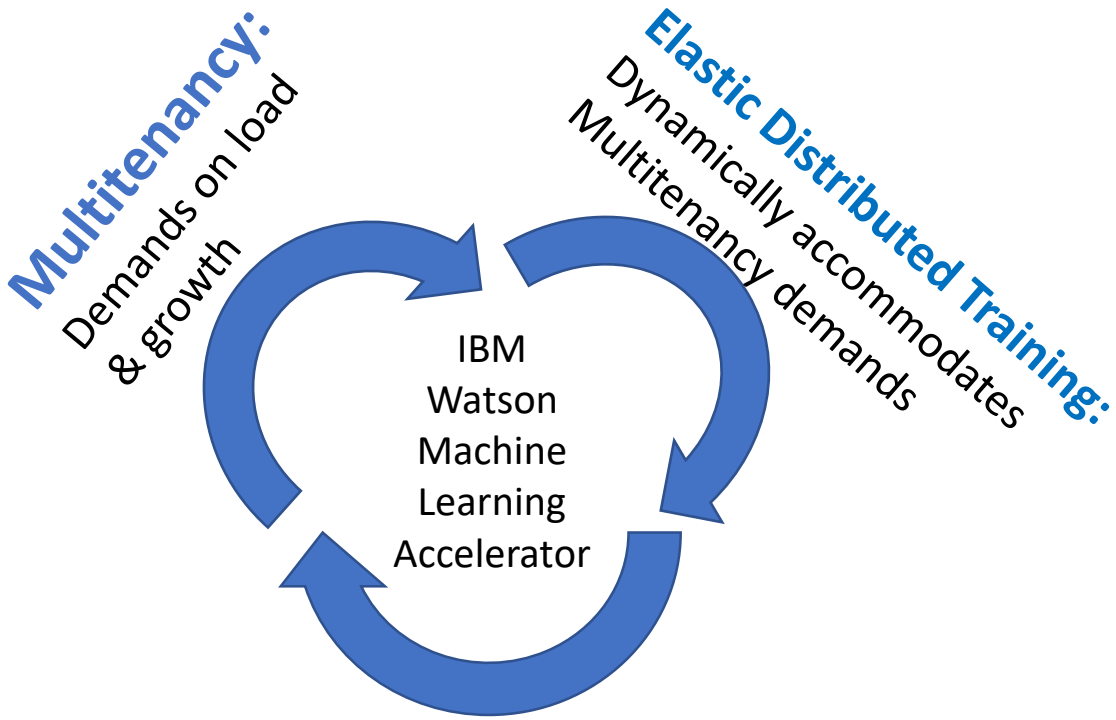

Scalability:
Power architecture
Satisfies growth requirements

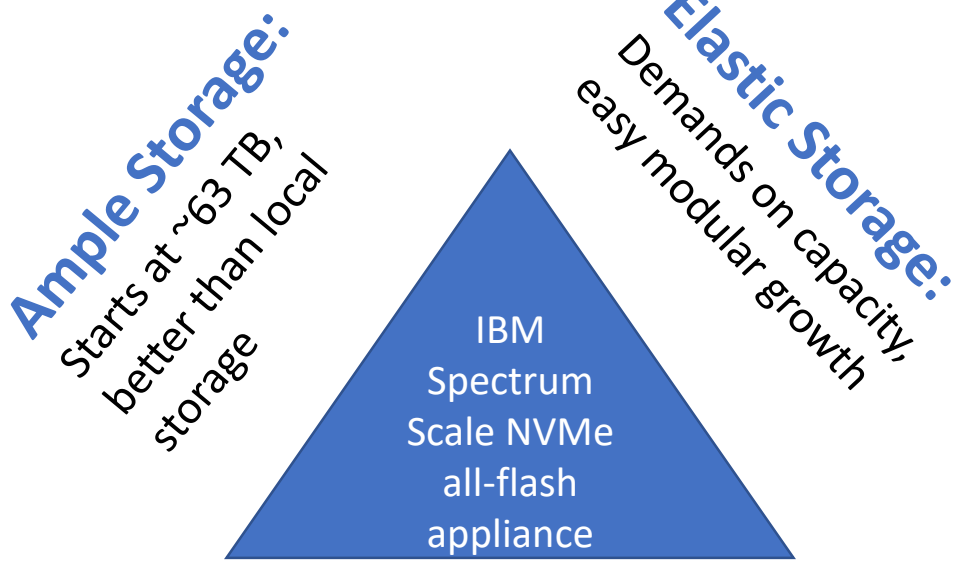

Fast Storage:
Storage keeps up with the AI workload demands, as fast as local storage

Elastic Solution End-to-End:
Elastic server nodes, Elastic compute AI, Elastic storage

Power9 with IBM Watson Machine Learning Accelerator (WMLA)

- Up to 8 x Power9 AC922 hosts, in this environment
- GTX model, water cooled
- 512 GM RAM, per Power9
- 6 GPU per Power9 host, up to 48 GPU in this environment
- Single dual-ported IB adapter per Power9 host

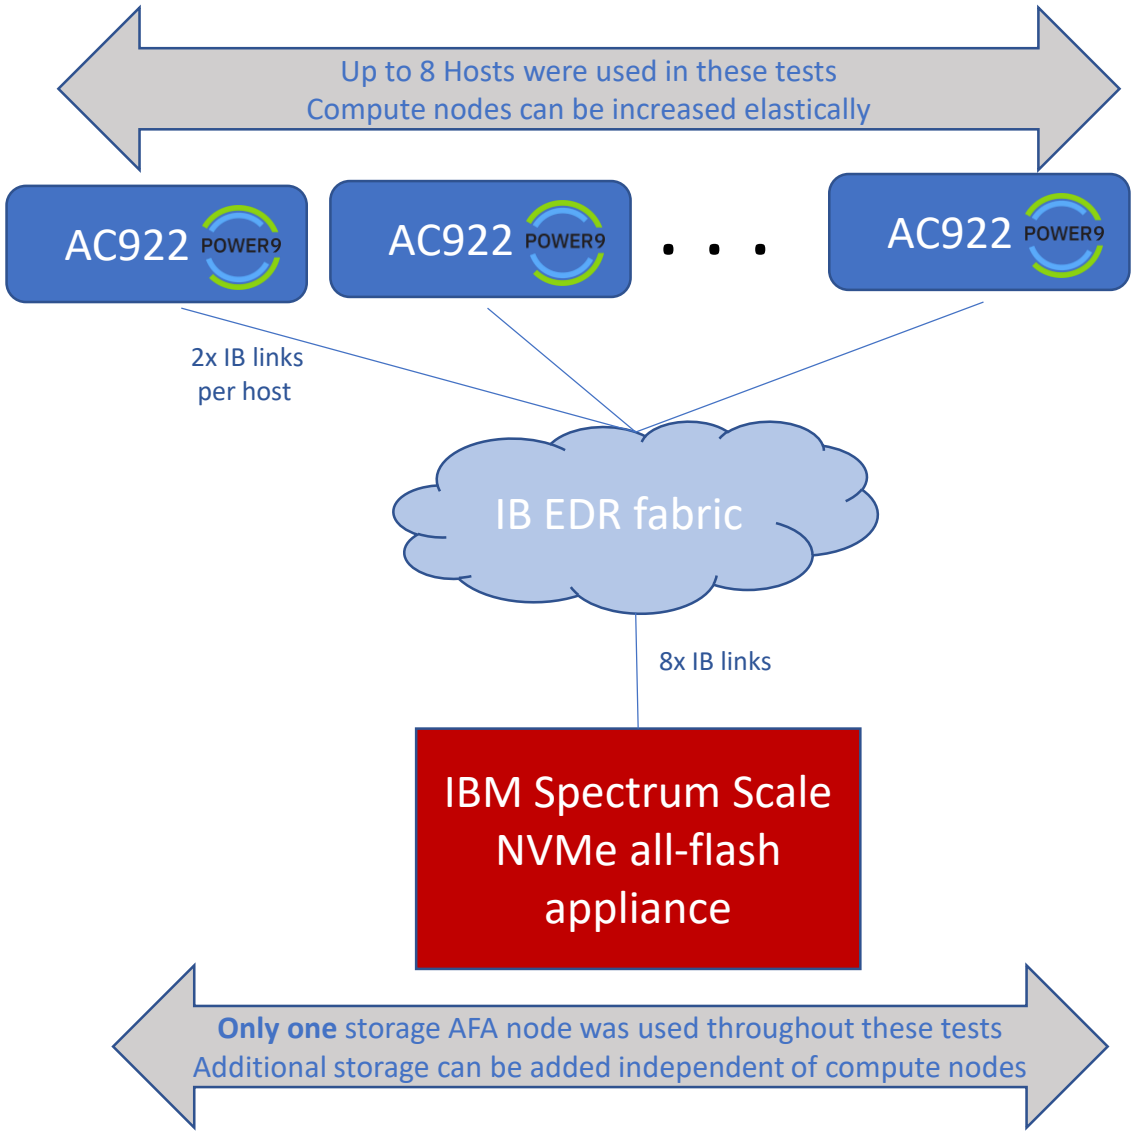

IBM Storage: Spectrum Scale NVMe all-flash appliance

- *NVMe-based Storage provides more than ample performance for these AI benchmarks which saturate the GPUs.*
  - Single AFA Storage uses 2U rack space and provides ~63 TB of user capacity
  - Max Read from storage: Over 35 GB/s, assuming enough network adapters
  - Storage can be increased in a linear fashion to meet capacity and/or performance requirements.

IBM Spectrum Scale – Parallel Architecture for Performance Scaling

OAK RIDGE National Laboratory

Summit System

- 4608 nodes, each with:
  - 2 IBM Power9 processors
  - 6 Nvidia Tesla V100 GPUs
  - 608 GB of fast memory
  - 1.6 TB of NVMe memory
- 200 petaflops peak performance for modeling and simulation
- 3.3 ExaOps peak performance for data analytics and AI

2.5 TB/sec
Throughput to storage architecture

250 PB
HDD storage capacity

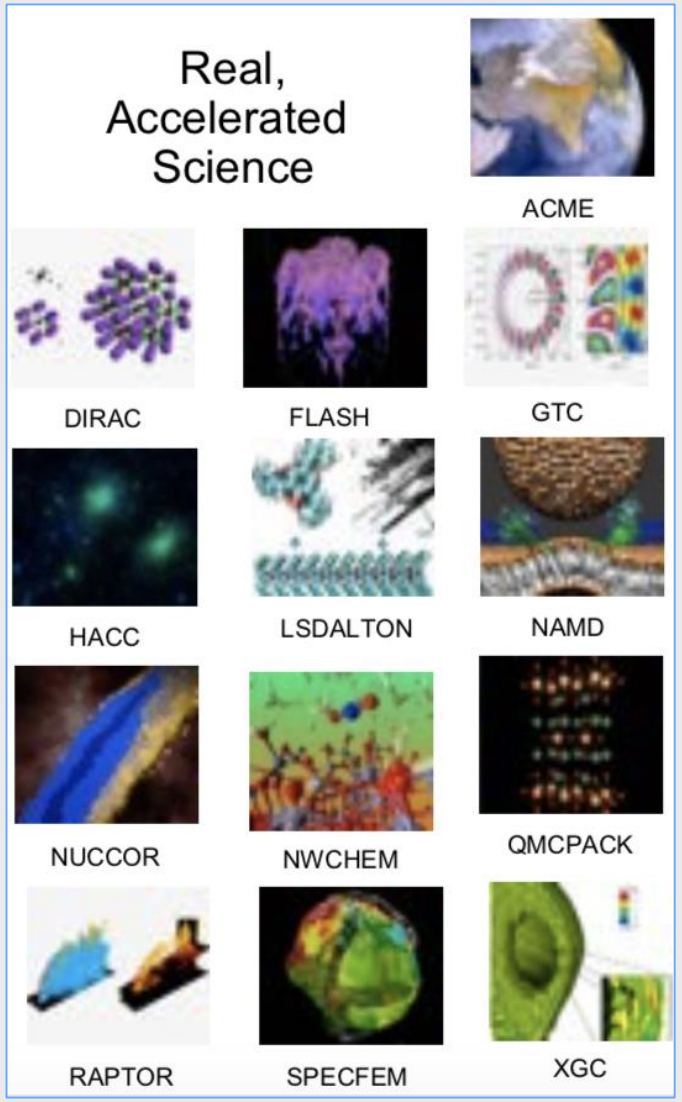

IBM Storage and SDI

Thank you

ibm.com/storage