IBM Storage Scale on the GPU Cloud

Re-thinking commodity server deployments

Scale User Group; SC23

Douglas O'Flaherty, Global Ecosystem Sales
douglasof@us.ibm.com

Sanjay Sudam, WW Storage Ecosystem Sales, IBM Technology Sales Infrastructure
sanjay.sudam@in.ibm.com

IBM

Disclaimer

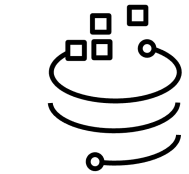

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

IBM reserves the right to change product specifications and offerings at any time without notice. This publication could include technical inaccuracies or typographical errors. References herein to IBM products and services do not imply that IBM intends to make them available in all countries.

IBM's Global Data Platform for AI with NVIDIA®

Engineered and optimized for data science productivity

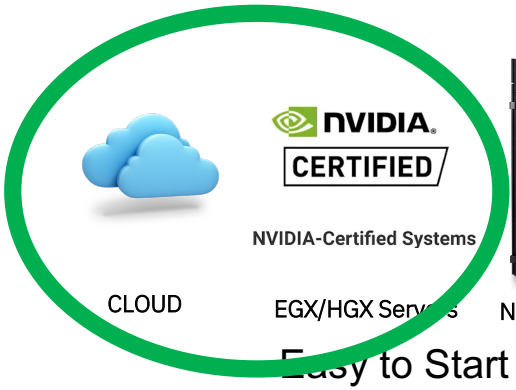

NVIDIA DGX BasePOD

NVIDIA DGX SuperPOD

Easy to Start

Easy to Expand

Parallel Performance Access Services	Multi-site/Multi-vendor Caching Services	Increased Efficiency Management Services	Cyber Resilient Security Services
Speeds AI Results	*Connects AI Data*	*Optimizes AI Data*	Protects AI Data
Up to 1.8TB/s and 30M IOPS per rack*	Breaks down silos with a Global Data Platform	Policy based data placement and reduction	IBM Safeguarded Copy and Cyber Vault

* IBM ESS 3500

Cloud Storage Architecture for Hybrid Platform

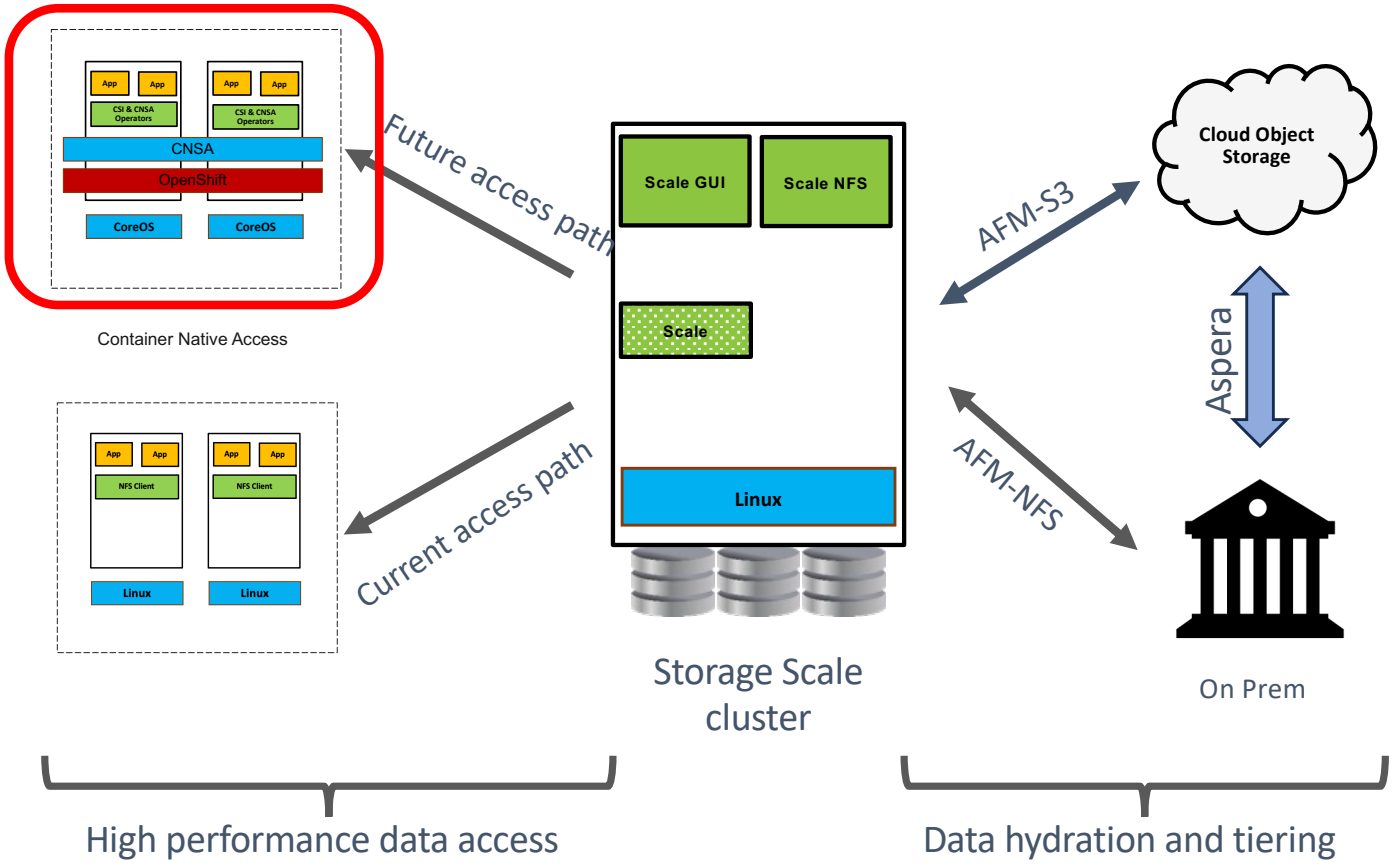

Current Solution

- LSF scheduler
- VSI/Bare metal compute cluster
- Storage Scale based high performance storage
- Data hydration with
  - Storage Scale AFM
  - LSF Data Manager
  - Aspera

Converged Solution

- Container native hybrid cloud
- Co-exists with current solution
- Leverages Container Native Storage Scale access (CNSA)
- CNSA feature to be supported in official Storage Scale Cloud solution

Ephemeral Storage with IBM Storage Scale

High-performance data is shared across multiple systems of GPU server nodes.

This replaces networked data storage or managing multiple copies on each server node. The solution makes multi-GPU computation, training, and inference easier to manage and faster to scale.

IBM Storage Scale on GPU-as-a-service is the fastest shared file system for 3-6 GPU nodes. The solutions does not require additional hardware.

Target Use Cases: AI, ML, DL

- Easy self-deployment
- Low cost for end-users
- Can be delivered by service providers

Ref Arch based on Oracle Cloud Infrastructure (OCI)

https://blogs.oracle.com/cloud-infrastructure/post/accelerate-ai-ml-workloads-oci-nvidia-ibm

OCI A100 GPU Cluster + NVIDIA GDS + IBM Spectrum Scale

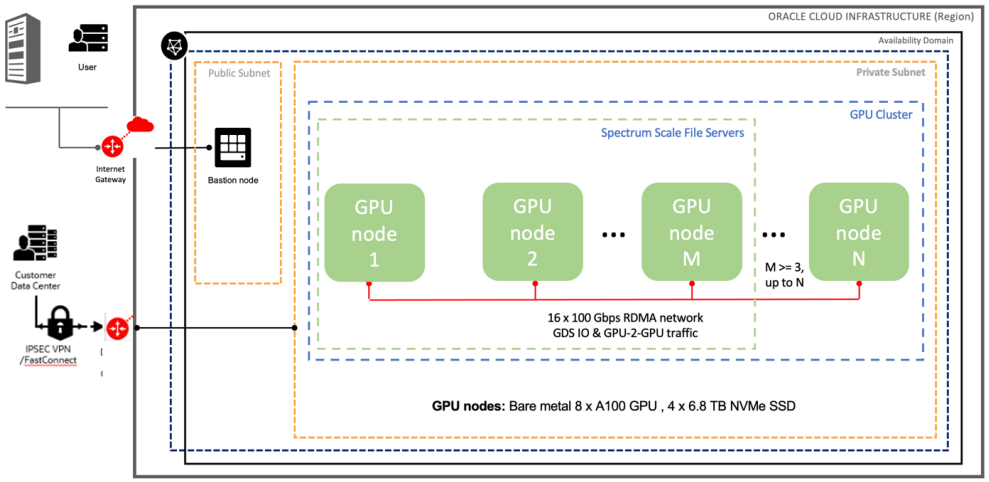

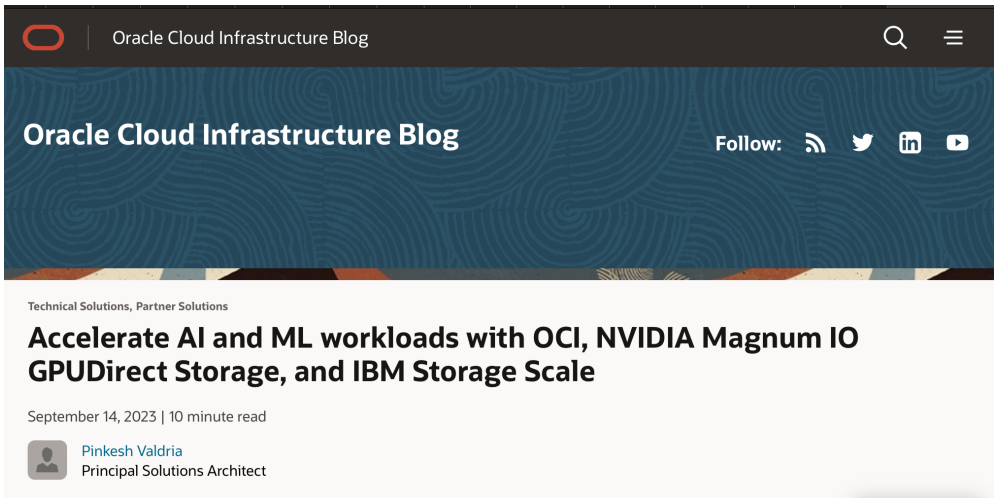

Reference Architecture based on OCI in detail

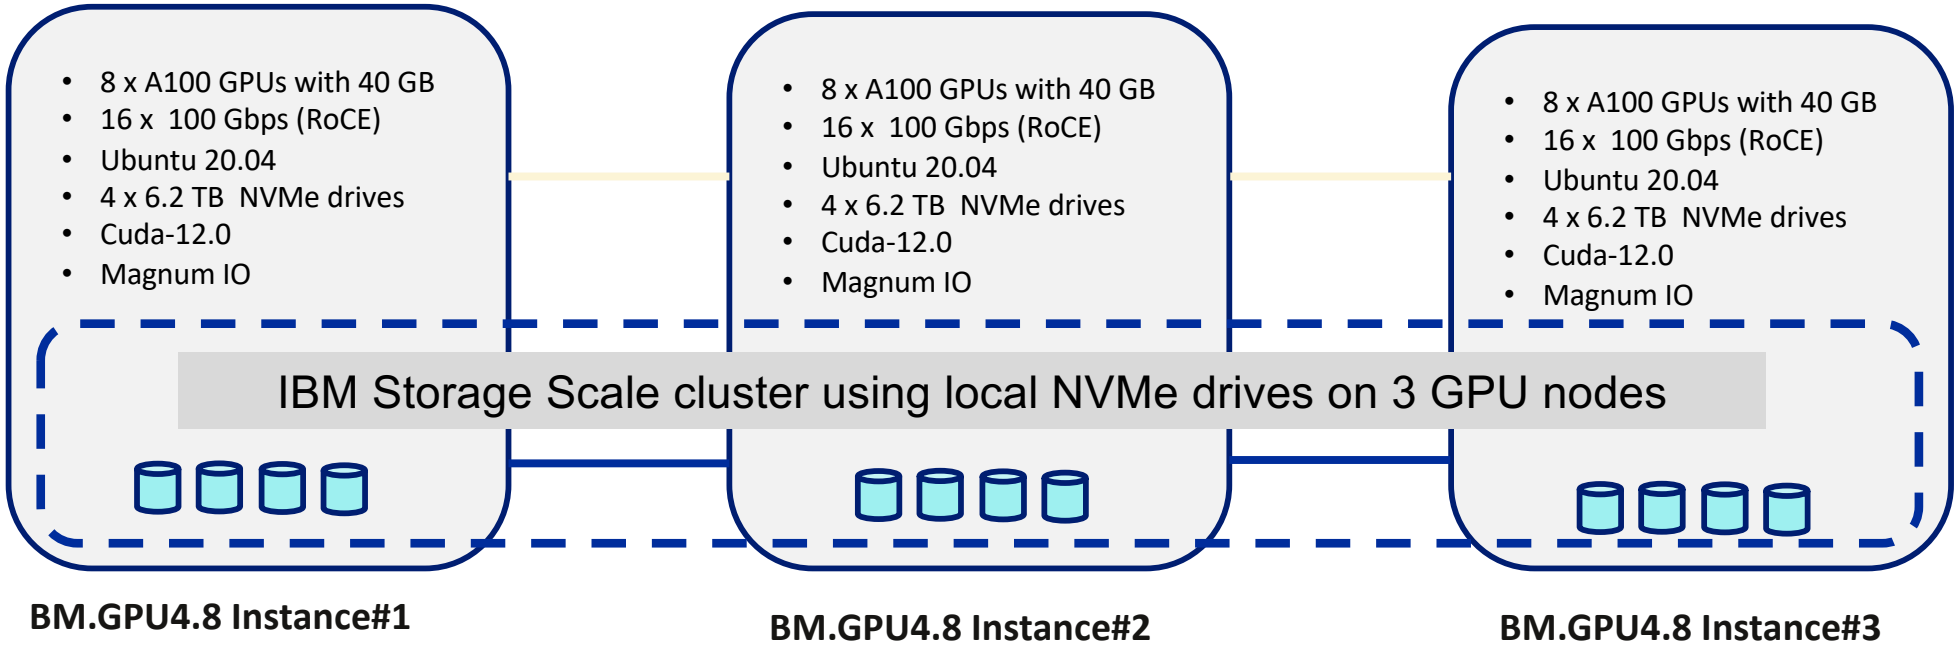

IBM Storage Scale configuration

- GPFS file system created across 12 NVMe drives configured as 12 NSD drives
- No gpfs level replication. High performance parallel scratch shared storage for on demand AI training workloads.
- 75 TB usable NVMe drive capacity
- **70 GB/s gds read performance from 3 node Storage Cluster.**
- **30 GB/s gds write performance from 3 node Storage Cluster**

GDSIO PoC – gds vs non-gds IO comparison

GiB/sec	x=0 GDS		x=1 CPU Only		x=2 CPU_GPU		Remarks
	BW	Latencies	BW	Latencies	BW	Latencies	
Write	29.87 GiB/s	2091.230 usecs	24.58 GiB/s	2991.685 usecs	24.65 GiB/s	2534.993 usecs	**20% more BW; 25% reduction in latencies;** NVMe drives are bottleneck during GDS.
Read	44.85 GiB/s	1393.354 usecs	20.89 GiB/s	2545.627 usecs	19.16 GiB/s	3269.147 usecs	**2X BW; 50% reduction in Latencies;** NVMe drives are bottleneck during GDS

GPFS File system block size = 4M, gdsio IO transfer size = 1M

gdsio xfer_type :

0 : Storage ➔ GPU Direct I/O (GDS)
1 : Storage ➔ CPU
2 : Storage ➔ CPU ➔ GPU
3 : Storage ➔ CPU ➔ GPU_ASYNC
4 : Storage ➔ PAGE_CACHE ➔ CPU ➔ GPU
5 : Storage ➔ GPU_ASYNC
6 : Storage ➔ GPU_BATCH

DeepCam Inference Benchmark

DeepCam Inference benchmark is used to show the performance of an application used to identify hurricanes and atmospheric rivers in climate simulation data. We saw 1.5x higher bandwidth (GB/s) and throughput (samples/sec) with GDS enabled for single node (8 GPUs) and three node (24 GPUs) inference run.

Table-1: DeepCam Inference benchmark results for single node and 3 node test

Number of Nodes	IO type	GDS Enabled	GDS Disabled	GDS Gain
1 Node (8 GPUs)	Max Bandwidth GB/s	35.59	23.94	1.48x
	Max Throughput Samples/second	674.88	459.76	1.47x
3 Nodes (24 GPUs)	Max Bandwidth GB/s	81.53	57.31	1.42x
	Max Throughput Samples/second	1546.12	1086.64	1.42x

User Experience

1. Data appears as if on local drive simplifying application environment.
2. Shared file system enables file locking, common read/write, and data management for job management.
3. Faster than networked storage and supports GPU Direct Storage.
4. Delivers full the performance of the fast, local NVMe drives already included in the GPU nodes.
5. Solution is software only. No additional systems or networking to deploy or manage.
6. Self-service provisioning by end-user.
7. No specific expertise required.
8. Option to read/write directly from Object Storage for permeance.
9. Data destruction when deprovisioned. Data on systems is unavailable/unreadable once systems are shutdown, or filesystem shut down.

Technical Requirements

- Minimum 3 nodes for scratch high performance tier offering.
  - Converged infrastructure for both GPU Compute and high performance storage with local NVMe drives
  - Scale as high performance scratch tier.
  - Object Storage as persistent capacity data lake.
  - 4 NVMe drives per node;
  - 2 x100 GbE Ethernet. RoCE for accelerated performance with GDS
  - Ubuntu/RHEL