IBM Spectrum Scale Strategy Days 2021

Ted Hoover Program Director Spectrum Scale

Wayne Sawdon CTO Spectrum Scale & ESS



High performance parallel data access with enterprise data services connecting edge to core to public cloud in a single federated cluster



# **Strategic Directions**



## Workload Performance

- Analytics
- o Al
- o HPC

## Enterprise Ready Platform

- Security
- DevOps
- **Container Ready**

Hybrid Multi Cloud

- Hybrid Cloud Solutions
- Edge Enabled

## Multi Cloud Data Management

- Central Management
- Cross Geo Data Movement
- o 5G Networks

## High Performance Storage

Spectrum Scale has a 20 year history of providing storage for the world's fastest and largest supercomputers

As hardware and workloads evolve additional investment is required.

Performance improvements regardless of deployment model or workload:

 National Labs, Data Center, Cloud, Hybrid Clouds, Edge Computing, Bare Metal, Fully Containerized

## Workloads

- Analytics
- Al
- HPC



## **Summit System**

- 4608 nodes,
- 3.3 ExaOps
- 2.5 TB/sec
- 250 PB storage

- RDMA
- GPU Direct
- IOPS, Direct I/O
- MCoT
- HDFS
- Object
- AFM WAN



# ESS 5000 Highlights

Bigger, Faster, Denser Capacity Storage



Third-generation, capacity-oriented Elastic Storage System: New throughput-optimized Power9 Data Servers:

- 9 PCle slots, up to 3 Network Adapters per Data Server
- Substantial improvement in performance, 55 GB/s Throughput
- Can support up to 9 enclosures in single building block.
- Logtip on NVDIMM, improvement in small write performance

#### 2 HDD models for maximum flexibility:

- SCx : same 4U106 HDD enclosure as Gen.2 GLxC models
- SLx : new 5U92 enclosure, up to 25% denser than Gen.2 GLxS

Containerized S/W stack to simplify installation and upgrade





# Spectrum Scale DevOps: Strategy



## **Reusable infrastructure**

#### Extend to provide administrative commands, ready for further reuse



# Evolving Storage Market

## Traditional Storage

<ul> <li>Deliver underlying infrastructure needs to support enterprise</li> </ul>	Container-Ready Storag	e
requirements. Centralized administration for	<ul> <li>Leverage existing investments in traditional storage to support</li> </ul>	Container-Native Storage
organization. <b>Examples:</b> DS8900, FlashSystem,	<ul> <li>Container deployments.</li> <li>Allows use of snapshots, clones,</li> </ul>	Storage deployed inside containers with enterprise level data management services to support mission critical applications deployed in containers.
IBM Spectrum Scale	and replication but doesn't take advantage of container framework and related benefits.	
	• Not optimized for Kubernetes so can be a bottleneck to achieving increased agility and elasticity.	<ul> <li>Direct attach and external storage support varying performance and capacity</li> </ul>
	<b>Examples:</b> DS8900, FlashSystem, IBM Spectrum Scale	needs.
		<ul> <li>Kubernetes control plane allows self service capabilities driving higher levels of efficiencies.</li> </ul>
		Spectrum Scale Container Native     Storage Access

0

# Container Native Storage Access https://www.spectrumscaleug.org/event/ssugdigital-persistent-storage-for-containers-with-spectrum-scale/



ontainer Native:	Classic Spectrum Scale has been separated into its fundamental components and built from the ground up with containerization of each
	component in mind. Spectrum Scale now 'lives' next to customer application containers.
utomation:	Spectrum Scale and CSI operators allow automated cluster and storage provisioning

**Flexibility:** Existing Spectrum Scale, ESS, ECE, clusters are used as storage via a remote mount, independent of OpenShift

**Open standards:** CSI provides an open standard for direct access to Spectrum Scale storage

C

A global file system built for the hybrid cloud and containers



## Hybrid cloud data Performance

Globally accessible pool of data across the enterprise lowers cost

Home > IBM Spectrum Scale CSI Plugin Operator

IBM Spectrum Scale CSI Plugin Operator

Performance is optimized with parallel access and enhanced throughput to storage resources

DevOps can create storage for containers in minutes

**IBM Spectrum Scale CSI Plugin Operator** 

# Hybrid Multi Cloud Spectrum Scale Use Cases

- Cloud Storage Global, consistent and secure data access or sharing using AFM
- Primary Storage Array Integration (synching, Tiering and backup) with Public Cloud using Spectrum Scale AFM S3
- Global namespace that can be deployed and span across on Premise and Public Cloud
- Hybrid Cloud Storage as Part of the Platform by tight integration with other IaaS and PaaS services such as IBM satellite, AWS outpost, Azure stack, Google Anthos.
- Primary storage deployed in colocation facilities with direct access to public cloud for faster data access/sharing for live workload running both on premise and Public Cloud





## Primary Enterprise Transformation Motions



March 22, 2021

IBM

## Example workload scenario

## IBM



- Edge server (or, directly on the intelligent camera) on the plant floor processes video data
  - Recognizing good and bad welds
- Alerting the reject/re-work processes when needed
- Reporting centrally on quality results

## Hybrid Multi Cloud: Edge Enabled



Build once ... deploy anywhere ... innovate anywhere ... move freely ... at the core ... at the edge



Data Management challenges at the Edge



# Distributed computing at scale Cloud computing

thousands of heterogenous servers & devices in thousands of remote locations with thousands of locations, servers & devices continuously added lots of homogenous servers in fewer data centers with minimal change to the server environment

## Massive scale, variability and rate of change requires autonomous management

# Hybrid Multi Cloud Data Fabric

The **hybrid data fabric** enables the right data to be made available across any cloud and on-premises, with integrated security, at the optimum cost while meeting non-functional requirements including performance, governance, risk and compliance. It is a loosely coupled collection of distributed services.



Discover Compliance Protection Migration Provenance Monitor

Address the complexities of data management across the clouds

## Edge Computing...

## IBM



© IBM Corporation, 2020

