IBM Spectrum Scale Strategy

Ted Hoover Program Director, Spectrum Scale



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 File & Object Data



IBM Spectrum Scale – Accomplishments over last 12 months



Access Services

Modernizing and Containerizing protocols

• High Performance S3

Machine Learning / AI / GPU acceleration

• Maximize GPU performance for Enterprise AI and Analytic environments

Containerization

• Spectrum Fusion SDS/HCI

Caching Services

Spectrum Scale AFM

• Policy-based tiering to object storage: AWS, Azure, Google

Performance and Scalability

- ESS 3500 NVME performance, HDD Hybrid/ Capacity
- Throughput and IOPS improvements

Management Services

Visibility, control and automation

- Ease of use and Automation
- Ansible playbooks
- Proactive monitoring

Reliability Availability & Serviceability (RAS)

 Call Home: protocols and network

Security Services

Security

- Multifactor authentication
- Additional QRadar integration: Access Denied events

Resiliency

- Storage Cyber Resiliency Assessment Tool
- Cyber Incident Response Storage Assessment

Ingest or access data with high performance S3 interface

IBM Spectrum Scale Data Access Services (DAS) - - High Performance Object Protocol

- fast AI results for S3 cloud native applications
- scalable solution for ingesting high performance S3 object data from remote locations
- scale performance and capacity
 as needed
- container native deployment for easy OpenShift integration
- applications can now optimize with the interface they need to access all the data they require (example: ingest S3 and access via file)*

GB/s to TB/s performance for S3 object data





Container Native Storage Access

https://www.spectrumscaleug.org/event/ssugdigital-persistent-storage-for-containers-with-spectrum-scale/



| Scalability: | Containerized compute cluster can scale with the OpenShift cluster |
|-------------------|---|
| Speed: | R/W benchmarks of Spectrum Scale CSI have shown same performance as non-containerized Spectrum Scale |
| Container 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. |
| Automation: | 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 |

Data Caching Services with active file management (AFM)

2 Data Caching Services

Spectrum Scale Active File Management - Transparent data caching , enabling tiering and sharing of data across clusters

Spectrum Scale AFM – Use Cases

- Investment protection Break down storage silos, easily leverage multi-vendor and multi-cloud resources
- Increase application agility Accessing data from edge to core to cloud
- Quickly scale your data From resources you choose with performance you require
- Faster access to remote data transparently caching remote data locally when needed



Data Virtualization

- Integrate legacy file and object data stores into a single file system to breakdown legacy data silos
- Create a **High-Performance Tier** for analytics



Data Collaboration

- Geo-distributed collaboration on data transparently shared between data centers, the cloud and edge sites
- Consistent cache provides a single source of truth with no stale data copies



Data Resilience

- Provides a Disaster Recovery solution for business continuity
- Air gap solution for DR
- Create an Active-Passive site relationship with failover and automatic data reconciliation on failback



Hybrid cloud / Bursting

- Dynamically increase computation resources in the cloud or at another site
- Burst site sees all data at home site and fetches data transparently on demand

IBM Storage / © 2022 IBM Corporation

Spectrum Scale DevOps: Strategy

Reusable infrastructure

Extend to provide administrative commands, ready for further reuse



Spectrum Scale on the Cloud

Access Data from Multiple Interfaces Access Data from Many Sources Deliver on the Value of Spectrum Scale

Hybrid Cloud Use Cases

- Backup / Archive
- Tiering
- Bursting
- Data Sharing

Deployment Models

- Lift and shift
- Container Native
- Managed Service
- Hybrid

Workload Enablement

• Analytics, AI, Containers

Ecosystem Integration

Cloud Toolkit



GPU Direct Storage

June 28, 2021: Initial release from NVIDIA



NVIDIA Magnum IO

- Family of I/O Optimizations for GPU accelerated data centers.
- GPU Direct RDMA: Access peer node's memory without copying to host memory
- GPU Direct Storage: Transfer data to/from GPU directly from storage without involving CPU and CPU memory

CUDA Toolkit

- GDS will be in the CUDA toolkit
- Development environment for GPU accelerated applications
- Libraries, compilers, debuggers, optimizers, and tools
- Leading GPU compute platform since 2006

GDS for Applications

- Invoked using the CUDA Toolkit (cuFile) API
- APIs must be explicitly called by the applications
- Storage must be GDS enabled. If not, GDS call falls back to regular data movement.

Why it matters

- AI, HPC, Analytics are data hungry and require a very high data throughput.
- GPUs are starved by slow I/O (and NFS is particularly slow)

• GDS is support on InfiniBand

| Benchmark (10 client Nodes) ESS3200 – 2 Building Blocks | Starting Baseline Prefetch Enabled (default) with 5.1.2 | SC21 Submission * Prefetch Disabled with 5.1.2 | ISC22 Submission** Prefetch Enabled (default) + hints with 5.1.3 |
|---|---|--|--|
| ior-easy-write | 103.6 | 106.4 | 109.47 |
| mdtest-easy-write | 187.9 | 195.6 | 174.86 |
| ior-hard-write | 3.2 | 4.3 | 32.7 |
| mdtest-hard-write | 19.3 | 22.3 | 22.12 |
| find | 2469.3 | 1185.2 | 2113.28 |
| ior-easy-read | 149.6 | 88.1 | 148.93 |
| mdtest-easy-stat | 267.2 | 272.2 | 335.48 |
| ior-hard-read | 1.9 | 29.3 | 28.77 |
| mdtest-hard-stat | 264.7 | 266.9 | 340.53 |
| mdtest-easy-delete | 114.2 | 113.4 | 174.19 |
| mdtest-hard-read | 251.3 | 205.4 | 407.59 |
| mdtest-hard-delete | 22.3 | 20.5 | 29.73 |
| BW Score | 17.5 | 33.0 Rank | 62.58 Rank 30 |
| IOPS Score | 158.9 | 143.5 28 SC21 | 193.58 In ISC22 |
| Total Score | 52.8 | 68.8 | 110.07 |

IOR Bandwidth – GiB/s mdtest/find - kIOPS

Improvements Result from:

Configuration and Tuning

Code Changes to Improve

Performance (e.g. hints)

Newly Added Hints Called

from Benchmark:

IOR hard read – FGRS hint

IOR hard write – FGWS hint

* SC21 list: https://io500.org/list/sc21/ten

** ISC22 list: https://io500.org/list/isc22/ten

IBM Elastic Storage System 3500

The simplest and fastest way to deploy a global data platform for AI and Hybrid Cloud workloads

Manage next generation and traditional workloads with simultaneous high-performance file and object data access services to the same data

Optimize local and remote access and simplify DR with global hybrid cloud data services

Speed access to critical data with Intelligent and automated data management services

Protect against cyber threats with Cyber-secure data services for unstructured data including end to end encryption and identification to recovery

Lower RTO times with proven data protection and data resiliency services

IBM Breaks Storage Performance Barriers for AI and Hybrid Cloud Workloads and Accelerates Recovery Times for Cyber Threats



up to 500+YB per cluster up to 30M IOPS per rack up to 91GB/s per node up to 1.8TB+/S per rack

Data Security Services - Active Protection for Cyber Resiliency

4 Data Security Services



- Cyber Resiliency Assessment Tool, Probes 100s of different controls and best practices
- Cyber Incident Response Storage Assessment (CIRSA)

RECOVER 🗸

- Instant access with Spectrum Scale AFM
- Spectrum Scale and Spectrum Protect recover multi-petabyte filesystems in hours
- CyberVault 4Q22
- QRadar Incident Forensics

IBM Global Data Platform

PROTECT



- Multifactor Auth, RBAC, Privileged Access Monitoring
 (IBM Security Verify)
- Safeguarded Copy, Logical air gap
- Scan snapshots for signs of ransomware (CyberVault)
- Log all Admin & user actions



- QRadar and Splunk SIEM integration
- File Audit Logging, Watch Folders
- Analyze backup data for signs of ransomware (Spectrum Protect)
- Reporting: QRadar User behavior analytics



- Automated action upon threat detection (QRadar)
 - Snapshot, Block Session , Etc..
- Alerts automatically prioritized based severity of the threat and criticality of the assets involved

Data Management Services to address distributed storage challenges and optimize time to result

- Virtually connect data end points and simplifying access pattern over any data to
 - Abstract data access across global data platform
 - reduce data copies
- Maintain only a single copy of the data with one global namespace
- Prefetch and Tier data to the right Storage tier to meet user requirements
- Provides and activates global automatic policy enforcement for increased data protection
- Utilizes augmentation of metadata to enable dynamic, intelligent and automated data orchestration
- Provides automatic enrichment to contextualize data with semantics and knowledge







A new software defined IBM Spectrum Fusion

Integrated OpenShift data services platform



IBM Spectrum Fusion HCI NEW

OpenShift data services platform software



Coming attractions

OpenShift data services platform software on Public Cloud



Spectrum Fusion HCI

Turnkey Red Hat OCP private cloud

- Fast to deploy, simple to scale and manage
- Optimized for containers

Kubernetes-native data services

- CSI and CNI
- High performance parallel file system

Integrated backup/restore

- Backup persistent data to remote vSnap & S3
- Policy driven backups



Key Solution Features

1. Bare metal OpenShift

- Eliminates cost, performance, and management overhead of unneeded hypervisor
- 2. Commodity x86 storage rich 1U servers
 - Populated with high performance NVMe flash drives
- 3. NVIDIA A100 GPUs to accelerate AI/ML
- 4. Global data platform services
 - Eliminate duplicate data and ad-hoc data management
 - Transparently access data anywhere

5. Single point of contact for solution support





https://www.ibm.com/storage/artificial-intelligence