

# IBM Spectrum Scale Strategy

## SC22 Dallas TX, Nov 13<sup>th</sup> 2022


Ted Hoover  
Product Manager, Storage for Data and AI, Spectrum Scale

Wayne Sawdon  
CTO Spectrum Scale & ESS

# IBM's Global Data Platform for File & Object Data



## 1 Data Access Services



## 2 Data Caching Services

### Global Data Platform

(powered by Spectrum Scale)

Local Cache

Local Cache

Local Cache

Local Cache



Investment protection

Object Storage

File Storage

NextGen workloads



File & Object Storage

(NetApp, PowerScale, etc)



IBM COS



Spectrum Scale



Spectrum Fusion

## 3 Data Management Services

## 4 Data Security Services

Identify



Protect



Detect




Respond



Recover



# IBM Spectrum Scale – Accomplishments over last 12 months



## Access Services

### Modernizing and Containerizing protocols

- High Performance S3
- HPO Multi Protocol Data Access (POSIX, S3, CSI)

### Machine Learning / AI / GPU acceleration

- Maximize GPU performance for Enterprise AI and Analytic environments
- GDS Write Acceleration

### Containerization

- Enhanced Fusion Support

## Caching Services

### Spectrum Scale AFM

- Policy-based tiering to object storage: AWS, Azure, Google
- AFM to S3 Performance
- AFM Recovery Improvements

### Performance and Scalability

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

## Management Services

### Visibility, control and automation

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

### Reliability Availability & Serviceability (RAS)

- Call Home: protocols and network
- Enhanced stretch cluster monitoring
- Detection of hung/unresponsive nodes

## Security Services

### Security

- Multifactor authentication
- Additional QRadar integration: Access Denied events
- Encryption -Using Hashicorp/ KMIP
- SED support(ECE)

### Resiliency

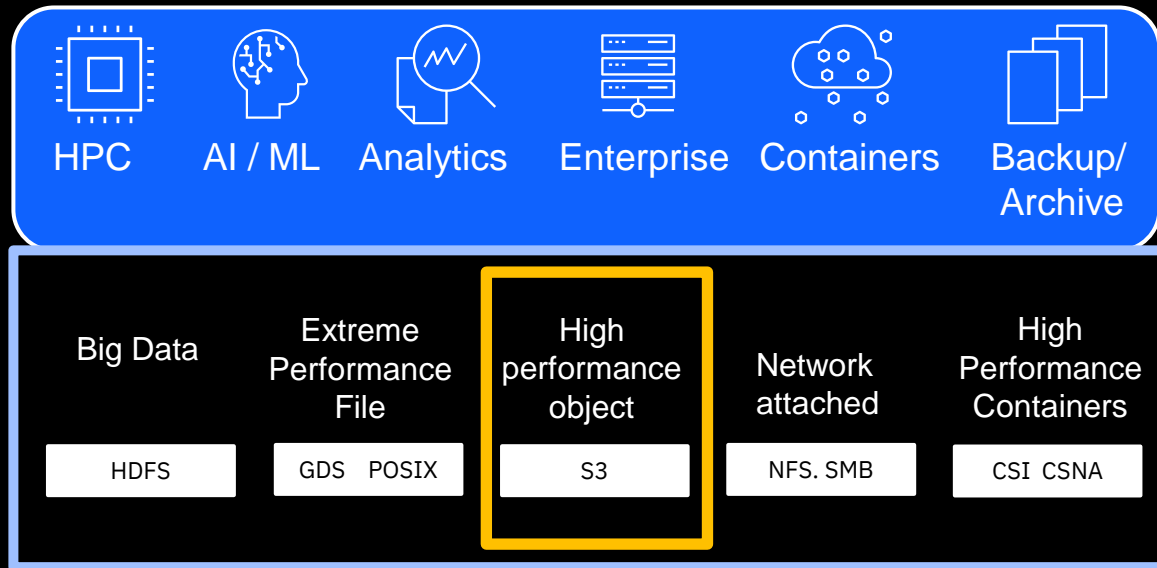
- Storage Cyber Resiliency Assessment Tool
- CyberVault: Blueprints and automation to provide monitoring, validation and recovery of safeguarded copies<sub>3</sub>

# 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



Global Data Platform

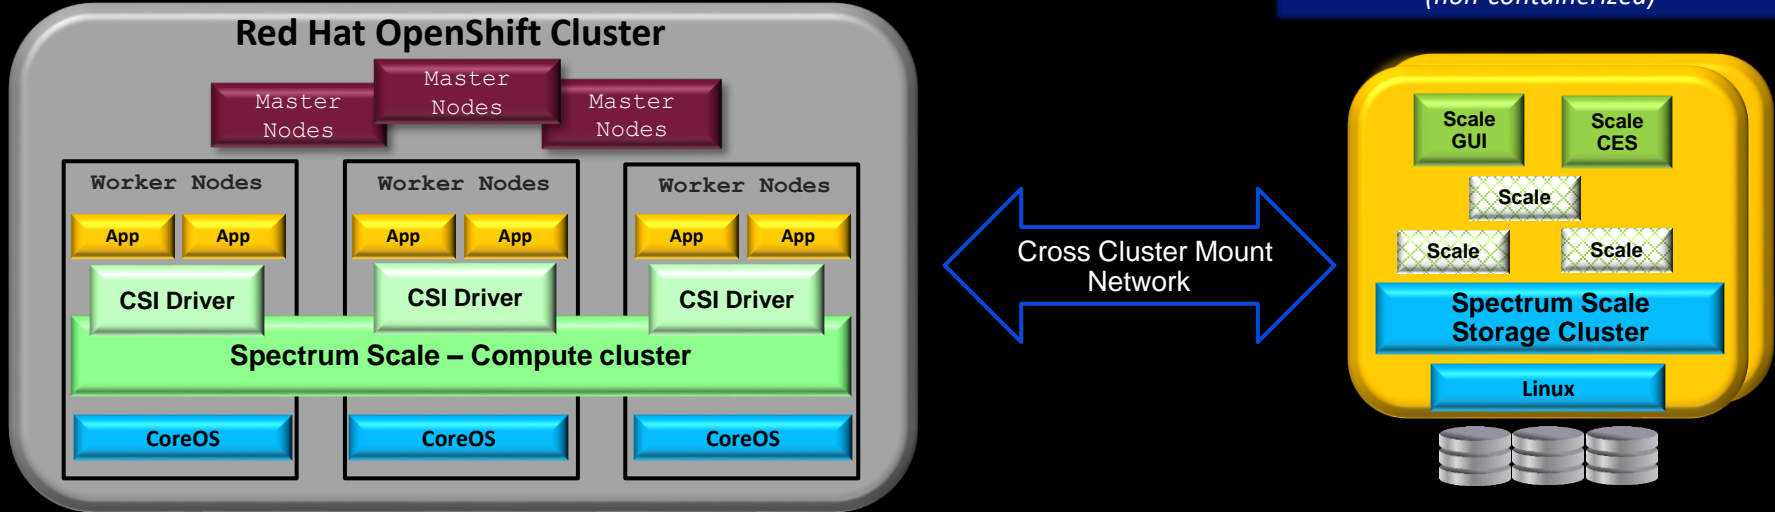
# Container Native Storage Access

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



Spectrum Scale in a container + CSI

Existing Spectrum Scale storage cluster  
(non-containerized)



**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*

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

- **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

## Spectrum Scale AFM – Use Cases



### 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



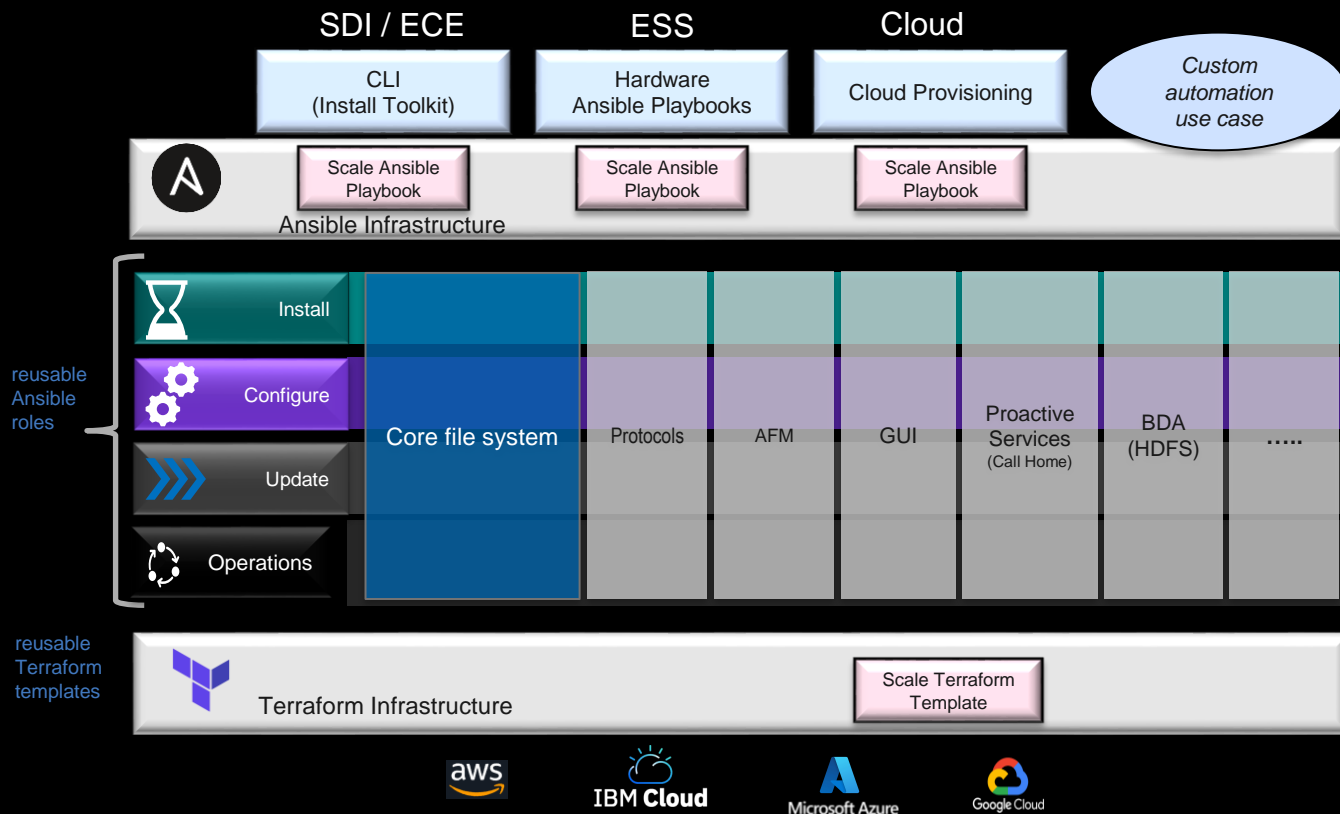
### Hybrid cloud / Bursting

- 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

# 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

Spectrum Scale CloudKit

Data Sources  
and Locations



File and  
Object



Kubernetes



Edge



Core Data  
Center



Public  
Cloud



Tape/Cloud

Data and AI  
Outcomes

POSIX /  
GDS



High  
Performance  
AI

Cloud  
Native S3



Backup /  
Archive

NFS /  
SMB



Enterprise Apps

HDFS

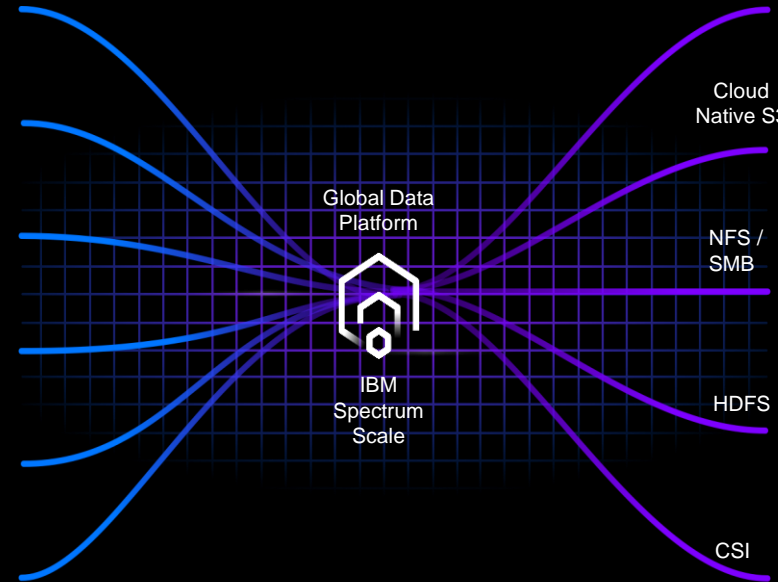


High Performance  
Analytics

CSI

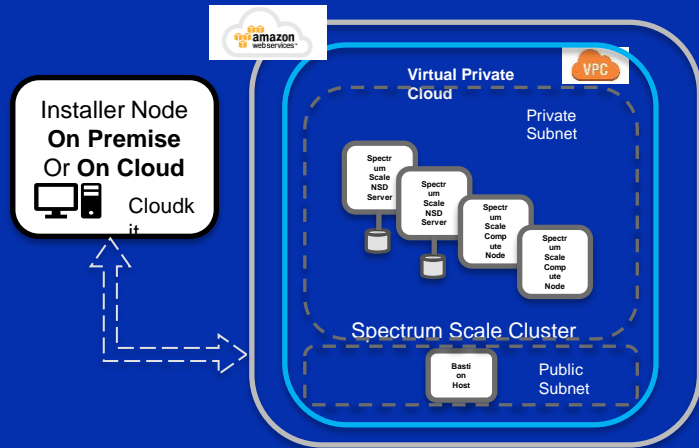


Containers  
Hybrid Cloud





# Spectrum Scale CloudKit



## Advantages

- Cloud agnostic (initial support for AWS)
- Easy to use, guided interface
- Rapid deployment and configuration of Spectrum Scale on the Cloud (in minutes)
- Flexible deployment (supports multiple deployment models)
- Ability to provision multiple clusters from a single CloudKit instance

## What is Spectrum Scale Cloudkit?

- Command Line Interface tool to create Spectrum Scale clusters on the cloud
- Provides end to end automation to create and bring up a Spectrum Scale cluster on public clouds
  - Automates infrastructure provisioning on the cloud
  - Automates the deployment of Spectrum Scale on the cloud
  - Applies Spectrum Scale best practises for deploying on the cloud

## CloudKit can Provision

- Single cluster with Spectrum Scale (Storage) NSD servers and Compute nodes
- Separate Spectrum Scale Storage Cluster (with all NSD servers) and Spectrum Scale Compute Cluster
- Separate Spectrum Scale Storage Cluster (with all NSD servers)

# IBM Spectrum Scale Cloud Integration

Cloud Offering on various Public Cloud Platforms

**File Storage As a Service**  
(Partner Managed)



**Marketplace Offering BYOL**  
(Bring Your Own License)



**Self Managed**  
(Bring Your Own License)

**Spectrum Scale  
CloudKit**  
(Tech Preview 4Q22)



## Multi-Cloud Offering

**Feature/Function support for Cloud/Hybrid use cases**

- **High Performant Filesystem**
- **Data Access Services - Multi-protocol**
- **Data Caching Services - AFM Support Speeds data access to collaborators and resources around the world**
- **Data Management Services – Policy driven data lifecycle management**
- **Security Services**

## IDENTIFY



- 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

## 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

## DETECT



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

## RESPOND





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

## IBM Global Data Platform



# 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
- **GDS** is support on **InfiniBand**

## 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)

<b>IO500 Benchmark (10 client Nodes) ESS3200 – 2 Building Blocks</b>	<b>Starting Baseline Prefetch Enabled (default) with 5.1.2</b>	<b>SC21 Submission * Prefetch Disabled with 5.1.2</b>	<b>ISC22 Submission** Prefetch Enabled (default) + hints with 5.1.3</b>
ior-easy-write	103.6	106.4	109.47
mdtest-easy-write	187.9	195.6	174.86
<b>ior-hard-write</b>	<b>3.2</b>	<b>4.3</b>	<b>32.7</b>
<b>mdtest-hard-write</b>	<b>19.3</b>	<b>22.3</b>	<b>22.12</b>
find	2469.3	1185.2	2113.28
<b>ior-easy-read</b>	<b>149.6</b>	<b>88.1</b>	148.93
mdtest-easy-stat	267.2	272.2	<b>335.48</b>
<b>ior-hard-read</b>	<b>1.9</b>	<b>29.3</b>	<b>28.77</b>
mdtest-hard-stat	264.7	266.9	<b>340.53</b>
mdtest-easy-delete	114.2	113.4	174.19
mdtest-hard-read	251.3	205.4	<b>407.59</b>
<b>mdtest-hard-delete</b>	<b>22.3</b>	<b>20.5</b>	<b>29.73</b>
<b>BW Score</b>	<b>17.5</b>	<b>33.0</b>	<b>62.58</b>
<b>IOPS Score</b>	<b>158.9</b>	<b>143.5</b>	<b>193.58</b>
<b>Total Score</b>	<b>52.8</b>	<b>68.8</b>	<b>110.07</b>

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

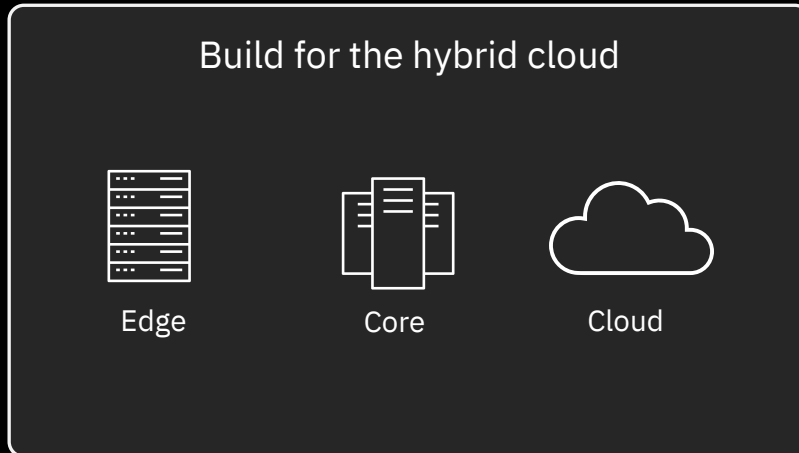
up to **30M IOPS** per rack

up to **91GB/s** per node

up to **1.8TB+/s** per rack

# Data 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



**IBM Spectrum Fusion**

*Coming attractions*

OpenShift data services platform software on Public Cloud



**IBM Spectrum Fusion**



# 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

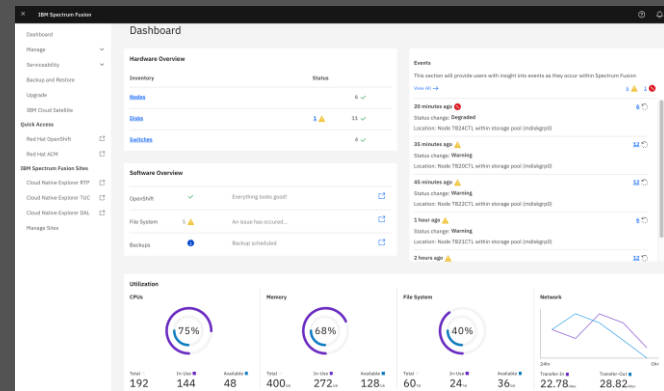
- Populated with high performance NVMe flash drives

### 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



# Spectrum Discover



- Scans and catalogs Metadata from heterogenous storage systems
- Enables a holistic view of data across the entire enterprise

Data sources:  
Files, Objects other data



Spectrum Scale



Spectrum Protect



IBM Cloud Object Storage



Spectrum Archive



*New capability included on 2.0.5*

*NFSv4 as data source on Spectrum Discover Operator-based version*



## New data source



- Ability to scan NFSv4 volumes by establishing new data connection for structured data
- NFSv4 Access Control List (ACL) and access control entities types reading and index data access.
- Ability to index, tag and structure access control, minimum of read-only access necessary.

Data Insight



Search



Reporting



Dashboard

- Simple to deploy (OpenShift Operator OR VMware virtual appliance)
- Metadata curation
- Custom metadata tagging
- Automatic indexing
- Policy-Engine
- Application SDK

Use Cases

## Large-Scale Analytics

- Data discovery
- Dataset identification
- Data pipeline progression/orchestration

## Data Governance

- Data inspection
- Data classification (tagging)
- Data clean-up
- IBM WKC Integration

## Data Optimization

- Archive / tiering / moving
- Duplicate data removal
- Trivial data removal



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