

# IBM Spectrum Fusion Technical Update - 2H22

*OpenShift Enterprise Data Services*



**IBM**  
Spectrum Fusion





Data services  
for mission-  
critical Cloud  
Paks and  
containers

Share data  
across the  
edge, core  
and cloud

High  
performance  
for AI/ML  
workloads

# IBM Spectrum Fusion

A **data services platform** for  
Red Hat OpenShift, everywhere



**Access any data, anywhere**  
with a secure global data fabric that is  
**simple to use | consistent everywhere | strategic**

# IBM Spectrum Fusion

Unlocks the Full Potential of Red Hat OpenShift



RED HAT<sup>®</sup>  
OPENS SHIFT



IBM  
Spectrum Fusion



Application portability



Security, resiliency, and backup



Integration with existing storage



Enterprise hardened



Unified File | Block | Object



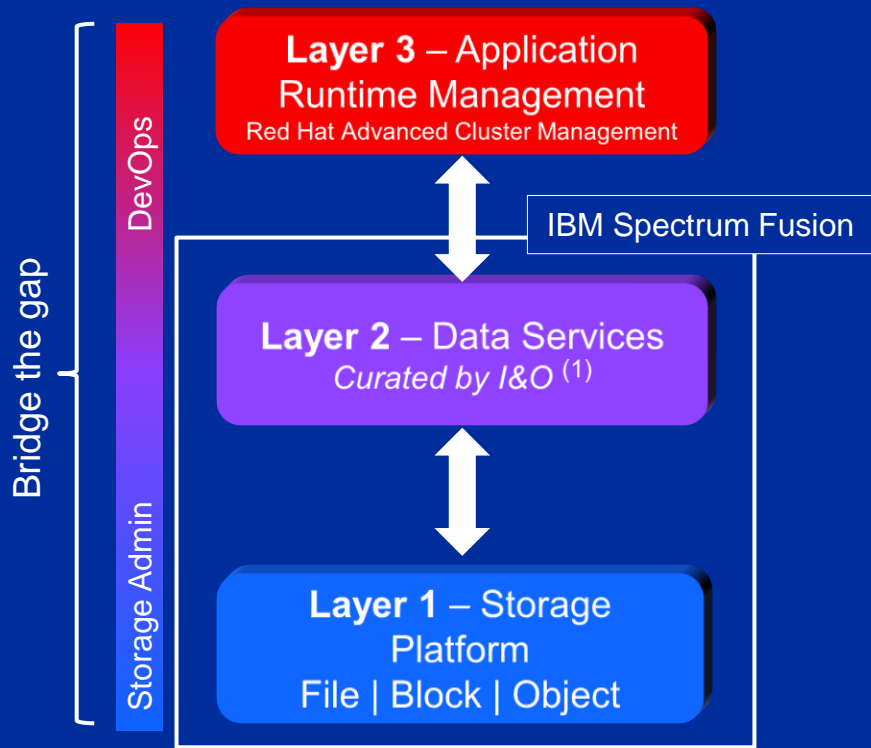
Data discovery and sharing



# IBM Spectrum Fusion

Container data services that are *simple to use, consistent everywhere, and strategic*

- **Protect application data**  
Configure application backup policies; recover applications to any point in time
- **Ensure application availability**  
Configure cross-zone data replication; manage availability with policies to RPO / RTO objectives
- **Access any data, anywhere**  
Manage access to data with policies; connect applications to any data source, anywhere



# IBM Spectrum Fusion with two deployment options

*OpenShift appliance or standalone software*



## IBM Spectrum Fusion HCI

*OpenShift Appliance + Spectrum Fusion data services platform*



Bare-metal Red Hat OpenShift cluster-in a box

## IBM Spectrum Fusion

*Stand-alone software*



vmware®



Google Cloud



IBM Cloud

X86 Bare Metal



# Spectrum Fusion – OpenShift Container Native Data Services

## Fusion Data Services

HA

Disaster Recovery

Backup & Restore

Discovery & Catalog

Encryption

Migration



Client OCP Workloads



OCP Virtualization

## Kubernetes Services

## OpenShift Container Platform

## Fusion Storage Services

HCI (Appliance)

Spectrum Scale (ECE)

HCI Management  
Networking  
Compute  
Storage (Scale ECE)

SDS (Software Defined Storage)

OpenShift Data Foundation (internal mode)

Spectrum Scale (remote-mount)



# IBM Spectrum Fusion HCI

A better way to run mission-critical applications on bare-metal OpenShift

## IBM Cloud Paks

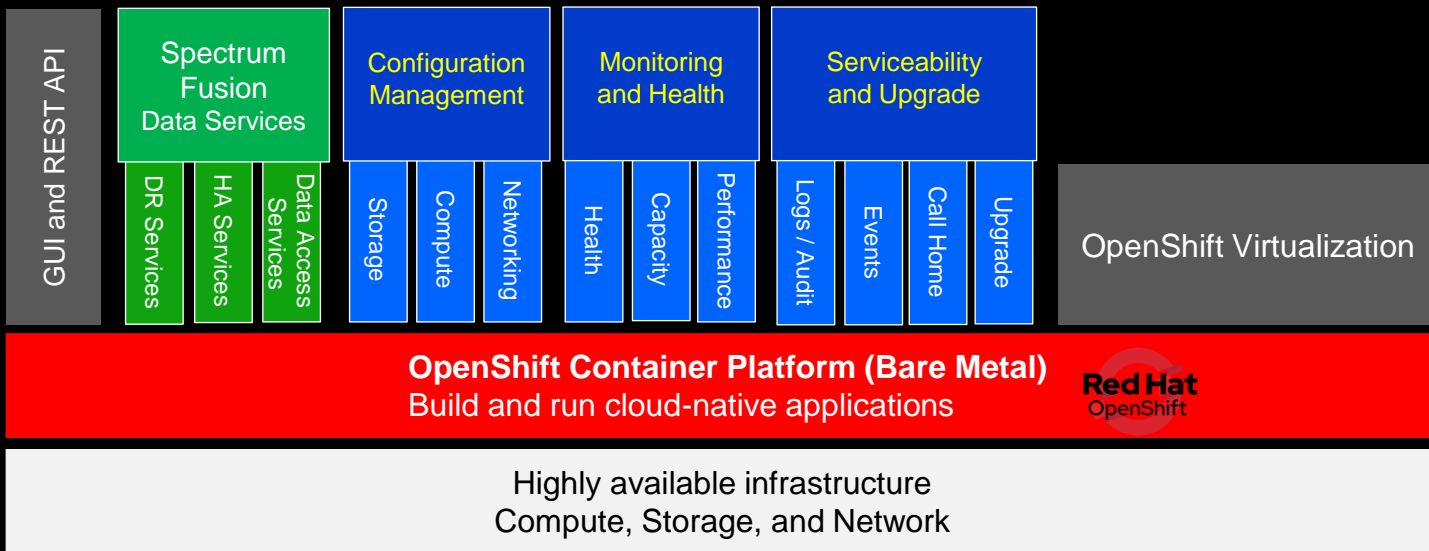
- Data, Security, Integration, Network Automation, Business Automation, Watson AI Ops

## Client/Partner

- Any OpenShift application
- NVIDIA CUDA
- Edge applications
- ...

## Open Source

- AI/ML: TensorFlow, PyTorch, NumPy, run.ai, etc.
- PostgreSQL, MongoDB
- ...






**IBM**  
Spectrum Fusion

**BETTER**

**TOGETHER**



**IBM**  
Cloud Paks

## CORE CAPABILITIES

Container-native data services  
platform for **Red Hat OpenShift** &  
**IBM Cloud Paks**



Tested and validated across ALL  
Cloud Pak Services\*

\*CP4BA – 3Q22



Built from proven **IBM Spectrum**  
**Storage** and **Red Hat** components



**Directly provision** storage volumes through  
the Spectrum Fusion CSI interface



**Backup and restore application state** to  
object stores for data resiliency



Application resiliency via **sync. and async.**  
**replication** across zones and regions



**Global Data Platform** – any application, any  
data source, any protocol, anywhere



**Simple, efficient, centralized management**  
of all OpenShift data services

“Spectrum Fusion plus Cloud Paks is **the** fastest way to simplify your storage.”

— Mike Gilfix, CPO, Cloud Paks



# USE CASES

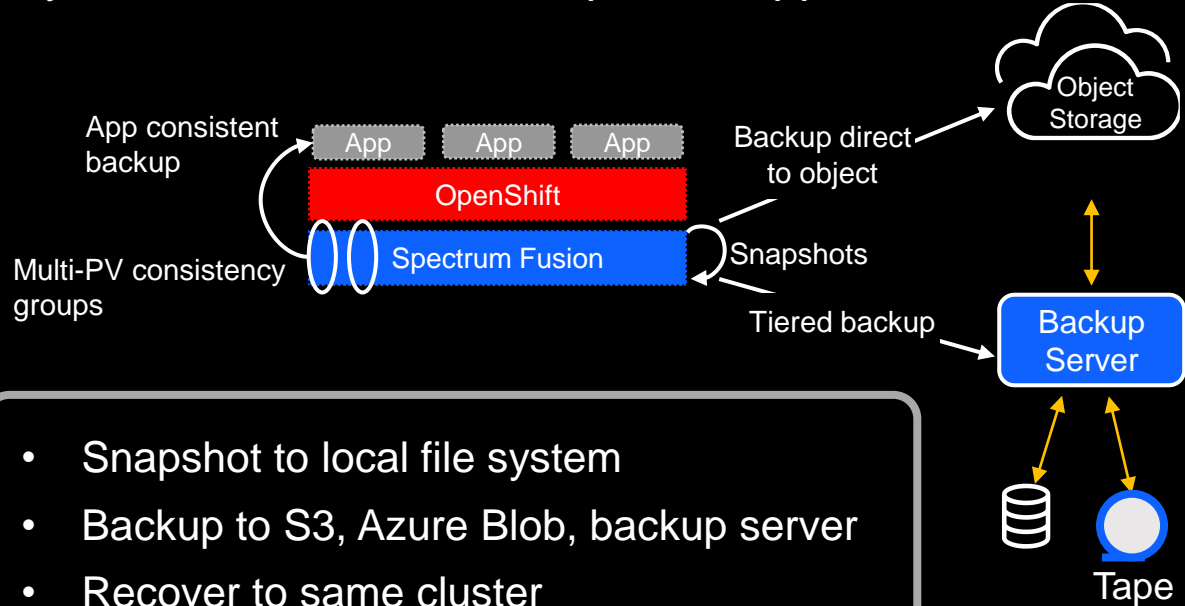
# Protect application data

Enable I&O teams to efficiently deliver data services to protect application data

Backup applications.

Recover to any point in time.

- Snapshot to local file system
- Backup to S3, Azure Blob, backup server
- Recover to same cluster
- Crash consistent
- Application consistent (quiesce hooks)
- Policy driven (schedule, target)

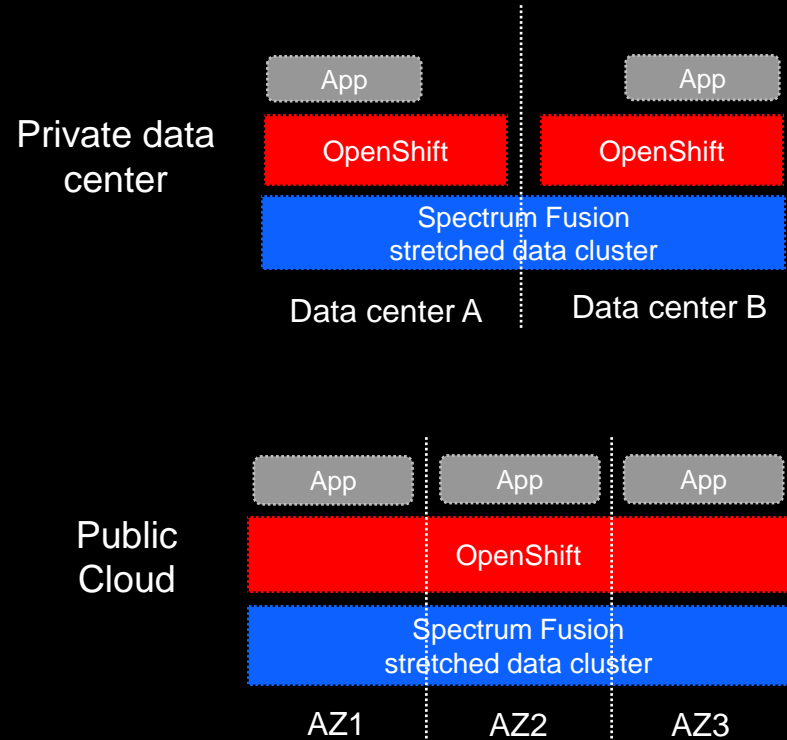


# Ensure application availability

Enable I&O teams to implement easy to use HA / DR services

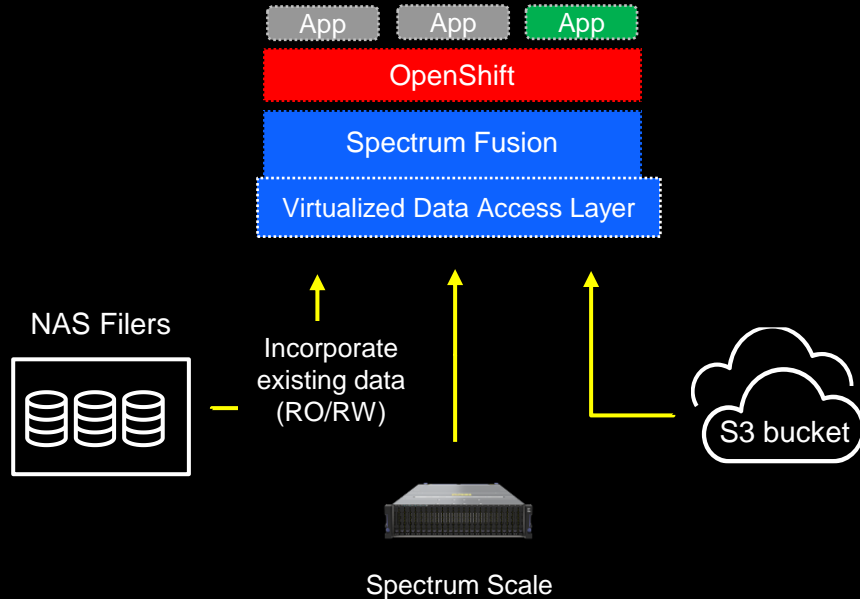
## Metro DR, Multi-Availability Zone

- Synchronous data replication
- Distance restriction – regions must be connected by a high-bandwidth, low latency link
- “Tie-breaker” application manages fail-over



# Enable access to any data, anywhere from edge to core to cloud

Enable I&O teams to increase application agility while protecting existing investments



- **Virtualize data access**  
Integrate multi-vendor and multi-cloud file and object data stores into a single file system
- **Access without data copies**  
Efficiently and securely share data across edge, core, and cloud by applying policies to the Spectrum Fusion virtualized data access layer
- **Cache remote data for performance**  
In a two-way consistent performance tier to avoid stale data, and improve data governance and security