Storage Fusion
Deep Dive / Day2 Operations

Storage Scale German User Meeting 2023
Sindelfingen, Germany – March 22-23, 2023

Daniel Danner
Lead Architect for Storage Fusion (D/A/CH)
daniel.danner@ibm.com

Agenda

- Fusion Overview
- Fusion Metro DR
- Fusion Day 2 Operations

What about below the red line?

Many choices to run OpenShift in the cloud
as a service or create your own environment.

But how easy to run OpenShift in your data center?
On the edge?

Can you connect the them together?

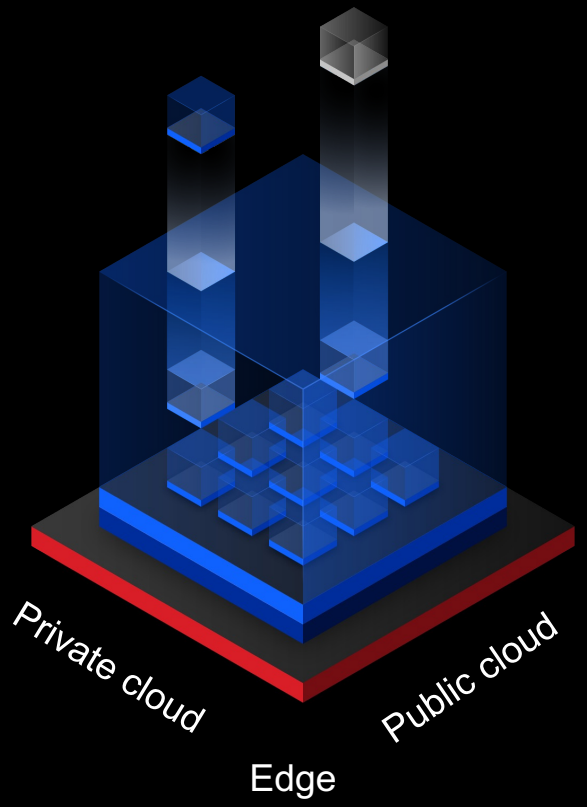

Introducing IBM Storage Fusion

We've designed Storage Fusion to address the
data management challenges
being faced as customers transitions to Containers and OpenShift

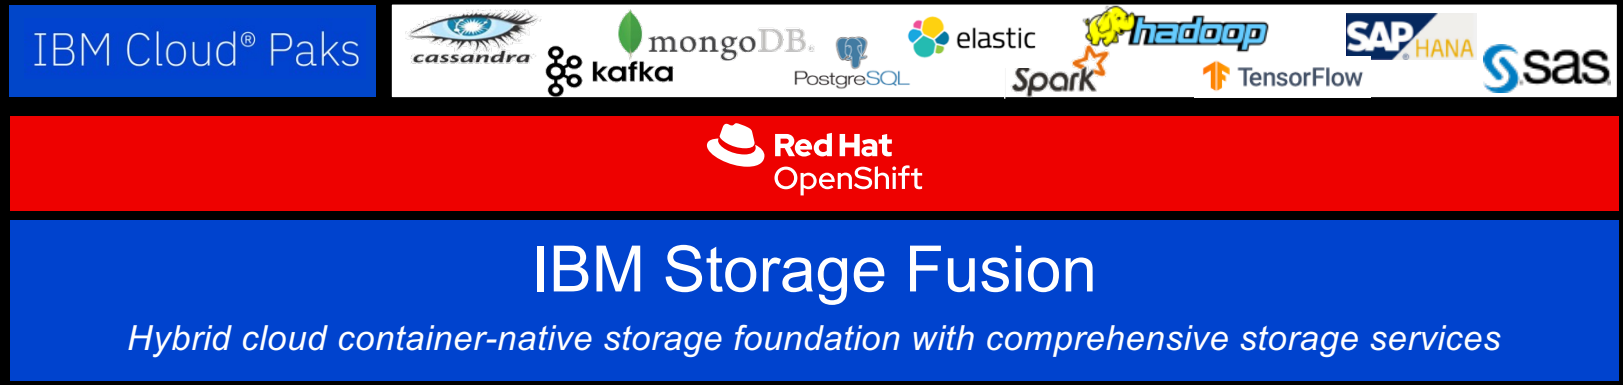

IBM Storage Fusion with two deployment options:
OpenShift appliance or *stand-alone software*

OpenShift Appliance + data services platform

IBM Storage Fusion HCI

OpenShift data services platform software

Red Hat OpenShift

vmware

Bare Metal x86

IBM Storage Fusion

OpenShift data services platform software on Public Cloud

Red Hat OpenShift

Microsoft Azure

aws

IBM Cloud

Google Cloud

IBM Storage Fusion

IBM Storage Fusion
Architecture

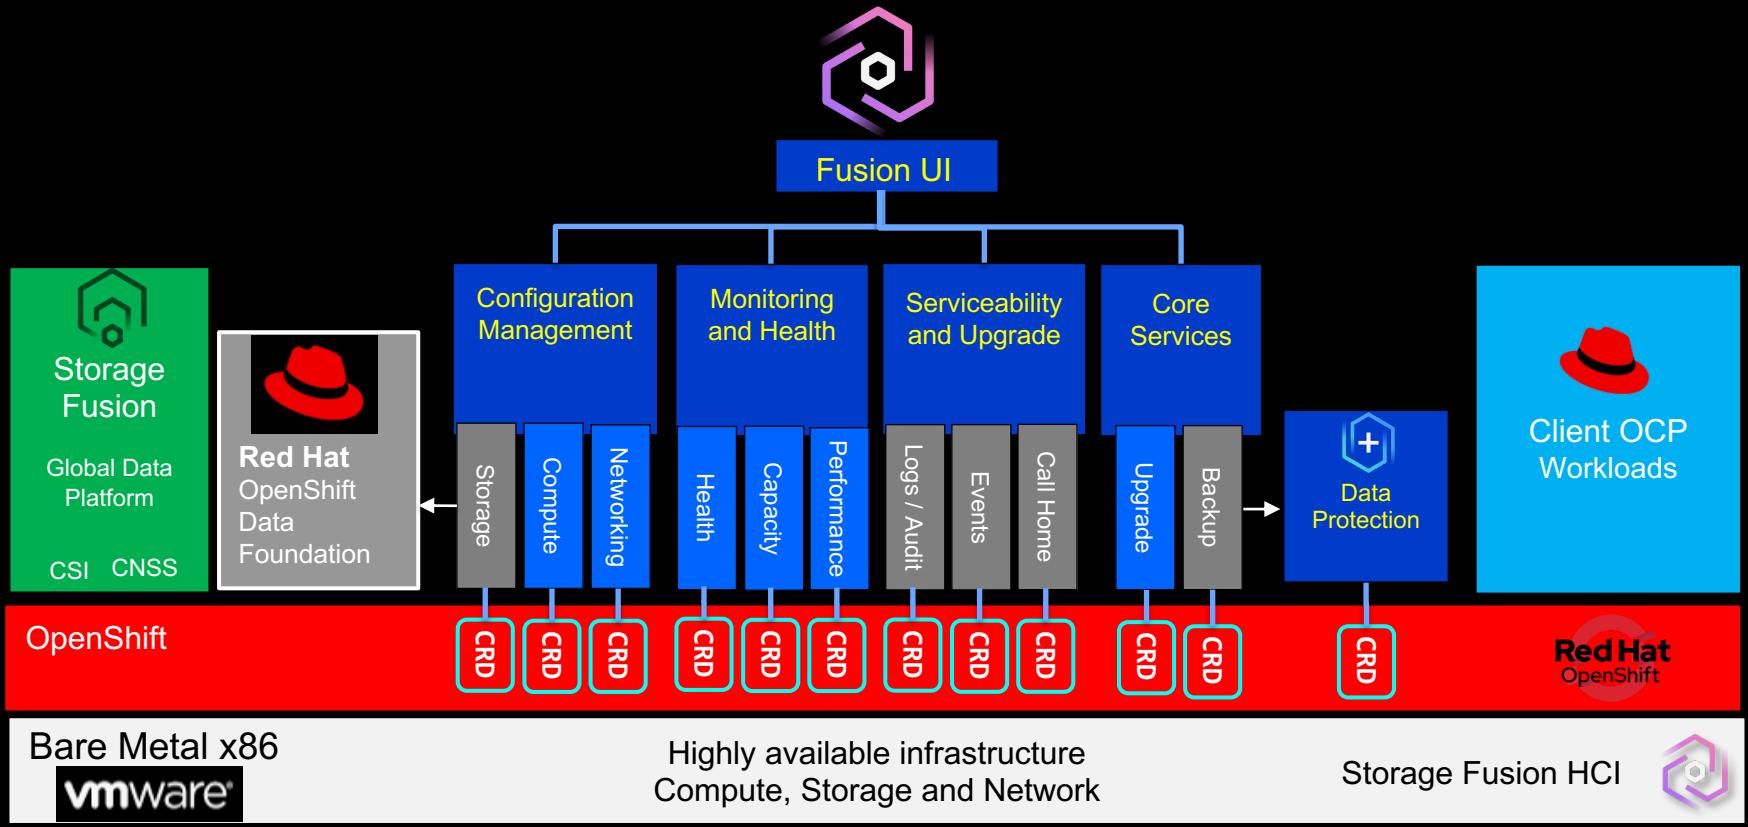

IBM Storage Fusion HCI System

An Engineered System for OpenShift

Integrated | Resilient | Built for OpenShift

“Bare-metal OpenShift in a Box”

What is it?

An engineered compute system
with OpenShift installed

- **Bare-metal x86 servers**
- **Ethernet switches**
  - OpenShift application network (25 GbE)
  - Dedicated storage network (100 GbE)
- **Hyper-converged storage architecture**
- **Engineered for performance and resiliency**
- **Supported by IBM (SPOC)**

“Bare-metal OpenShift in a Box”

Why?

Time to Value

You could spend six months designing an architecture and six months proving it, or

Eliminate the guess work and go with a proven solution.

Deploy a high resiliency OpenShift cluster on bare metal in less than a day

Accelerate getting OpenShift clusters into production

Why?

Reduce Risk

Eliminate risk of poor cluster design and missed performance objectives

Engineered for resiliency: ensure continuous operation of business-critical applications

Supported by IBM

Why?

Reduce Total Cost of Ownership

No need to design/prove/maintain

- *a hardware architecture*
- *a resilient network*
- *a resilient and high-performance storage layer*
- *an OpenShift deployment model*

Eliminate the virtualization tax

- *Performance, cost, and operational overhead*

Consolidate skills to a common platform

- *Build once, run anywhere*
- *Innovate anywhere, with anyone's technology*
- *Consistent everywhere*

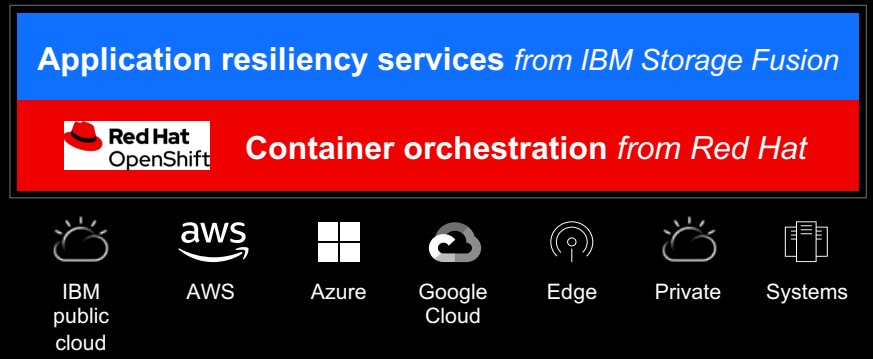

Fusion HCI System rack-up

Base configuration (smallest thing clients can order today)

- 42U rack, pre-loomed networking, 6 PDUs
- 2x Ethernet high-speed switches (100GbS)
- 2x Ethernet management switches
- 6x Storage/Compute servers with 2 NVMe drives/server
  - Server in RU7 is the provisioner node, connected to the KVM
  - Servers in RU2, RU3, and RU4 become the OpenShift control plane

Scale out options:

- 32c or 64c multithread nodes
- Storage rich and compute only nodes
- Memory options 8 GB to 32 GB RAM per physical core
- Up to 20 nodes per rack, in-field scalable
- GPU servers, each with 3x NVIDIA A100 GPUs
- Increased storage by adding pairs of drives to storage/compute servers
  - 7.68TB NVMe PCIe Gen4 drives/server to a max of 10 drives/server
- AFM (Active File Manager) NFS/S3 gateway nodes

42	3U Filler
41	
40	
39	3U Filler
38	
37	
36	1U PDU (horizontally mounted)
35	1U PDU (horizontally mounted)
34	2U Filler
33	
32	Storage/Compute Server
31	Storage/Compute Server
30	Storage/Compute Server
29	Storage/Compute Server
28	GPU Server with 3x GPU PCIe Gen 4 adapter cards
27	
26	GPU Server with 3x GPU PCIe Gen 4 adapter cards
25	
24	AFM Node
23	AFM Node
22	KVM Service Console
21	32-port 100 GbE Ethernet Switch
20	32-port 100 GbE Ethernet Switch
19	48-port 1 GbE Management Ethernet Switch
18	48-port 1 GbE Management Ethernet Switch
17	Storage/Compute Server
16	Storage/Compute Server
15	Storage/Compute Server
14	Storage/Compute Server
13	Storage/Compute Server
12	Storage/Compute Server
11	Storage/Compute Server
10	Storage/Compute Server
9	Storage/Compute Server
8	Storage/Compute Server
7	Storage/Compute Server
6	Storage/Compute Server
5	Storage/Compute Server
4	Storage/Compute Server
3	Storage/Compute Server
2	Storage/Compute Server
1	Reserve 1U Space at Bottom
	Required
	Optional

Fusion HCI System

Scale compute and storage

Scale compute
- 6 nodes to 20 nodes

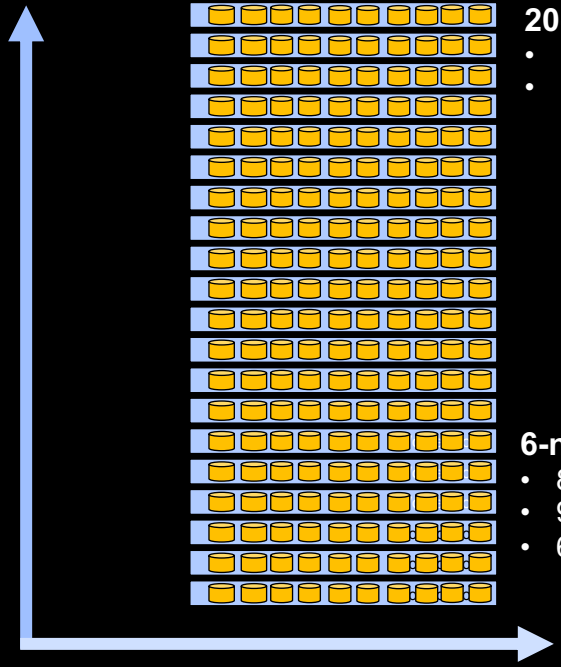

20-node Fusion HCI rack (max size)
- Up to 1,396 TB RAW, 935 TB usable
- Up to 1088 cores (2176 vCPU)

6-node Fusion HCI rack (min size)
- 83 TB raw, 55 TB usable
- 96 cores (192 vCPU) total
- 60 cores (120 vCPU) usable

Scale storage
- In 83 TB raw, 55 TB usable increments

Worker node options

Compute only node
- 32 core / 64 vCPU
- 64 core / 128 vCPU

Compute and storage node
- 32 core / 64 vCPU
- 64 core / 128 vCPU

GPU node options

3 or 6 NVIDIA 3x A100 GPUs

AFM Protocol Servers

1 or 2 AFM protocol servers

Fusion HCI System

Value of storage architecture

Scale-out parallel file system

- Enormous scalability with software-based declustered RAID protection
- Very high performance no additional RAID hardware
- Performance scales as you add nodes

Durable and robust storage

- ❖ Distributes data across nodes and drives for higher durability *without* the cost of replication
- End to end checksum identifies and corrects errors introduced by network or media
- Withstands multiple failures; Rapid recovery and rebuild

High storage efficiency

- 66% efficient with default 4+2p erasure coding
- 75% efficient with 8+3p optional erasure coding
- Lower cost than 33% efficient 3-way replication

Fusion HCI System

High performance for demanding workloads

	Provisioning modes		Persistence		File RWX	File RWO [3]	AFM S3 Caching	AFM NFS Caching	Scale out parallel file system [4]	Erasure coding
	Dynamic	Static	Zone [1]	Multi-zone [2]						
Fusion HCI	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
NFS	X	✓	X	X	✓	X	X	X	X	X
vSphere Volume	✓	✓	X	X	✓	✓	✓ (3rd party)	X	X	X

1) Fusion Metro DR synchronous data replication
2) Fusion Regional DR asynchronous data replication (roadmap 1H23)
3) Ideal for Cloud Paks and other applications that specify a need for block storage
4) High throughput, low latency performance for modern AI/ML and analytics workloads

IBM Storage
Fusion HCI System

Application High Availability

IBM Storage Fusion &
Red Hat Advanced Cluster Management

RH ACM

- OCP management
- Application management
- Policy management

Scenarios

- Monitor Storage Fusion systems
- Change configurations
- Upgrade Storage Fusion systems
- Deploy applications across all Storage Fusion systems
- Policy Management

Red Hat
Advanced Cluster Management

IBM Storage Fusion

IBM Storage Fusion

IBM Storage Fusion

IBM Storage Fusion

IBM Storage Fusion

IBM Storage Fusion &
Red Hat Advanced Cluster Management

Migrate workloads
from one location to another
without duplicating data

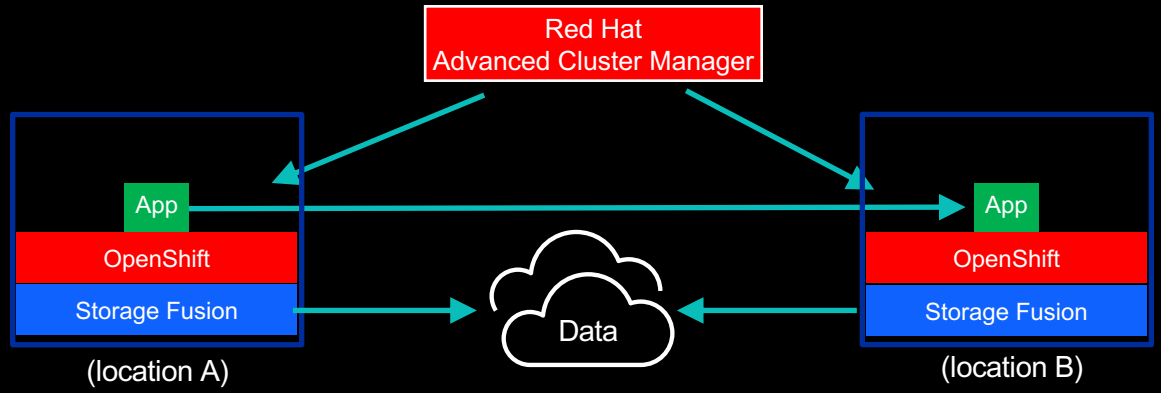

Please keep in mind!

RHACM is a **Deployment and Configuration Management Tool!**

It does not take care of mirroring or moving persistent data (PVCs) to another Openshift Cluster!

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 high-bandwidth, low latency link
- “tie-breaker” application manages fail-over

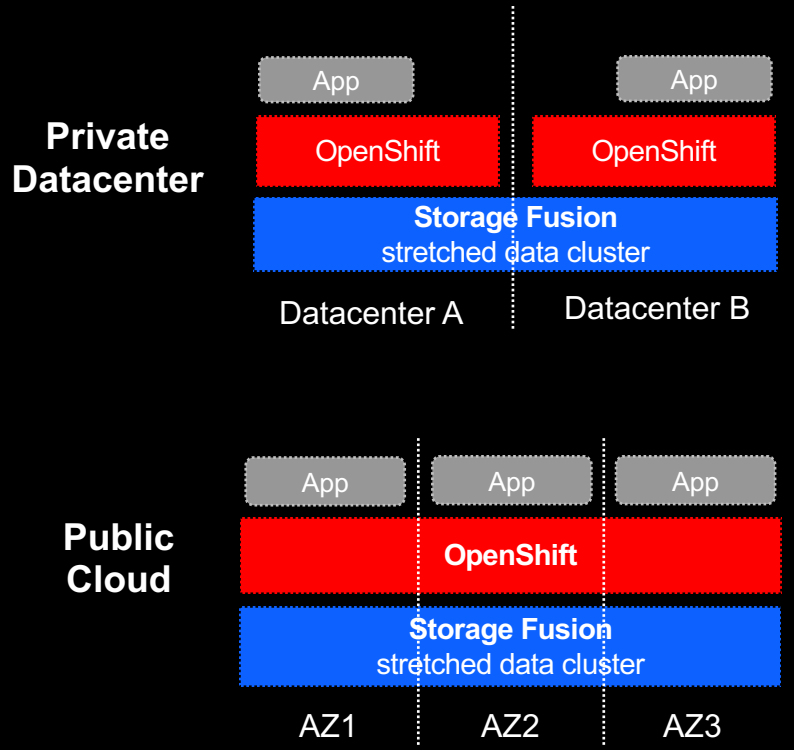

AZ – Availability Zone

Data Replication for Application HA

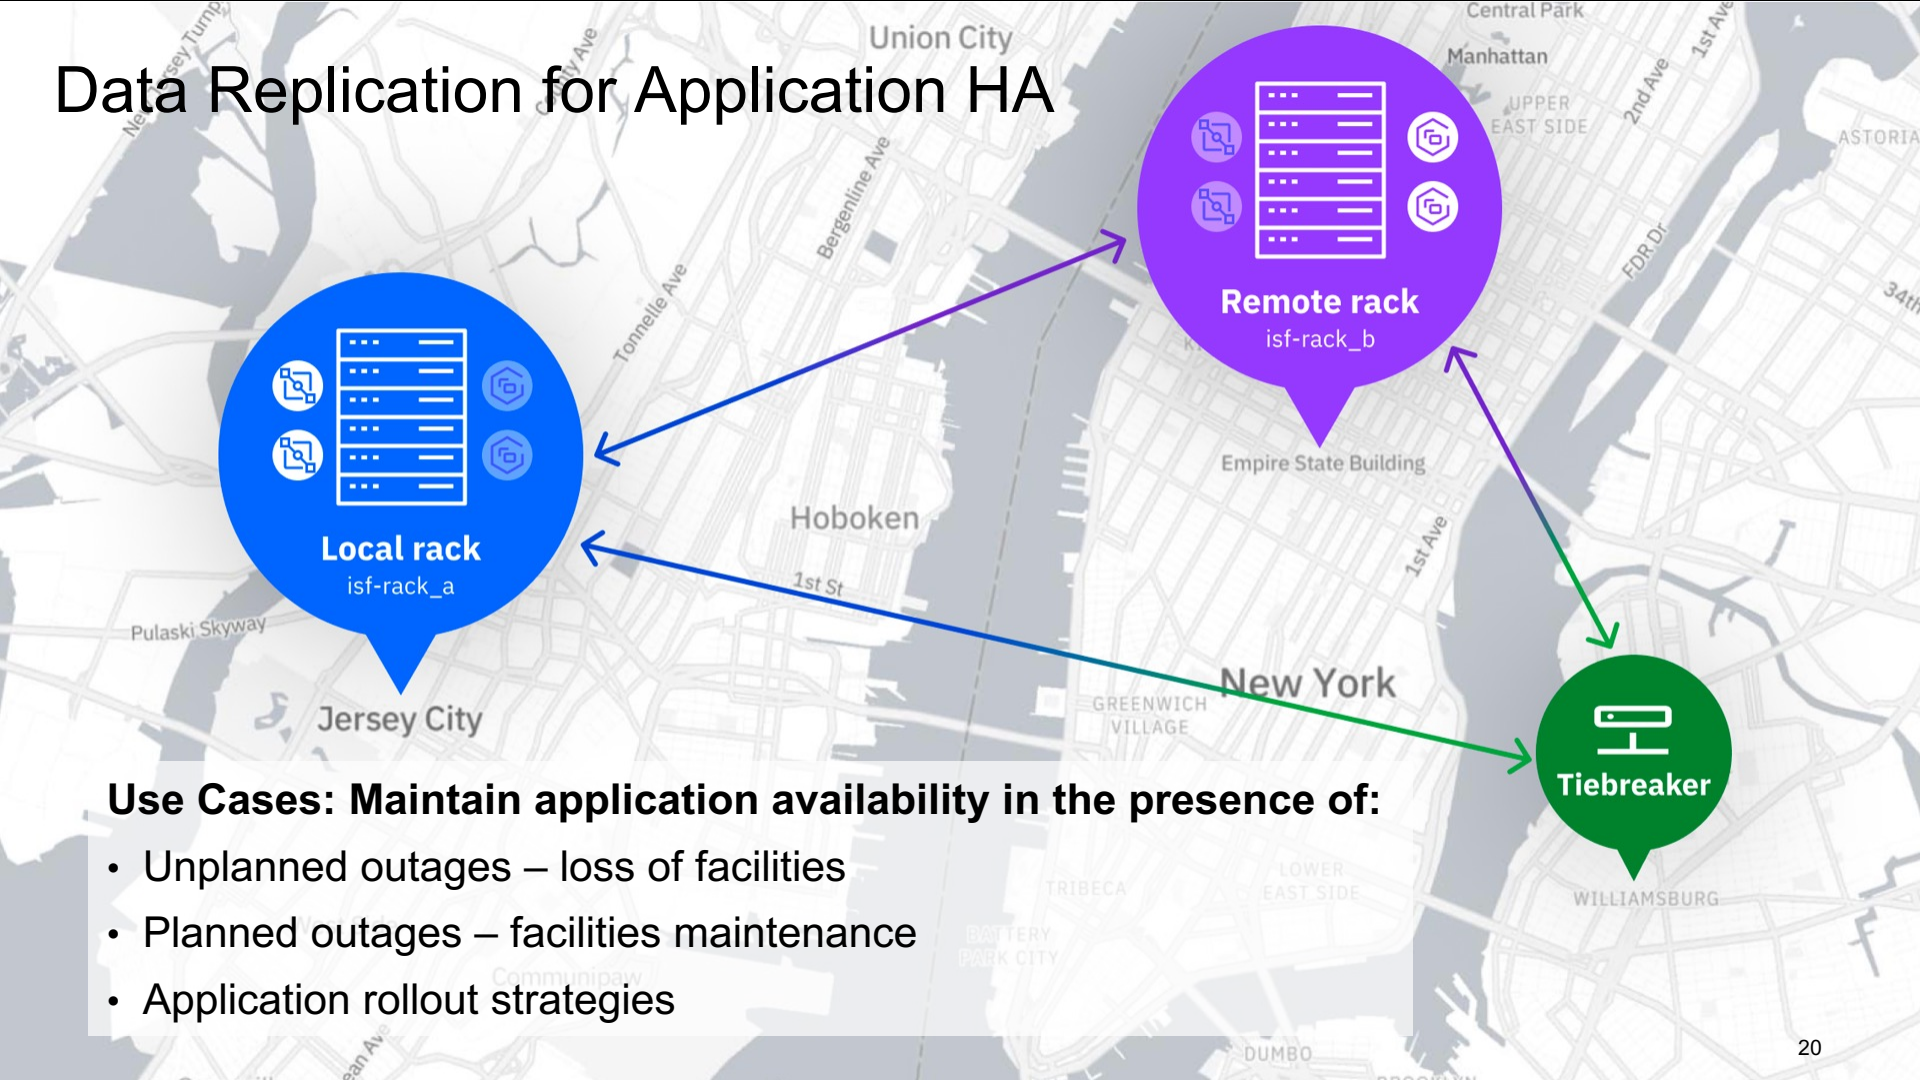

Use Cases: Maintain application availability in the presence of:

- Unplanned outages – loss of facilities
- Planned outages – facilities maintenance
- Application rollout strategies

Application High Availability

Metro DR replication

- Pair Fusion HCI racks in a Metropolitan area
- Replicate data synchronously between the two racks
- RPO = 0, RTO is variable

Regional DR replication *

- Pair Fusion HCI racks world-wide
- Replicate data asynchronously between the two racks
- RPO = variable

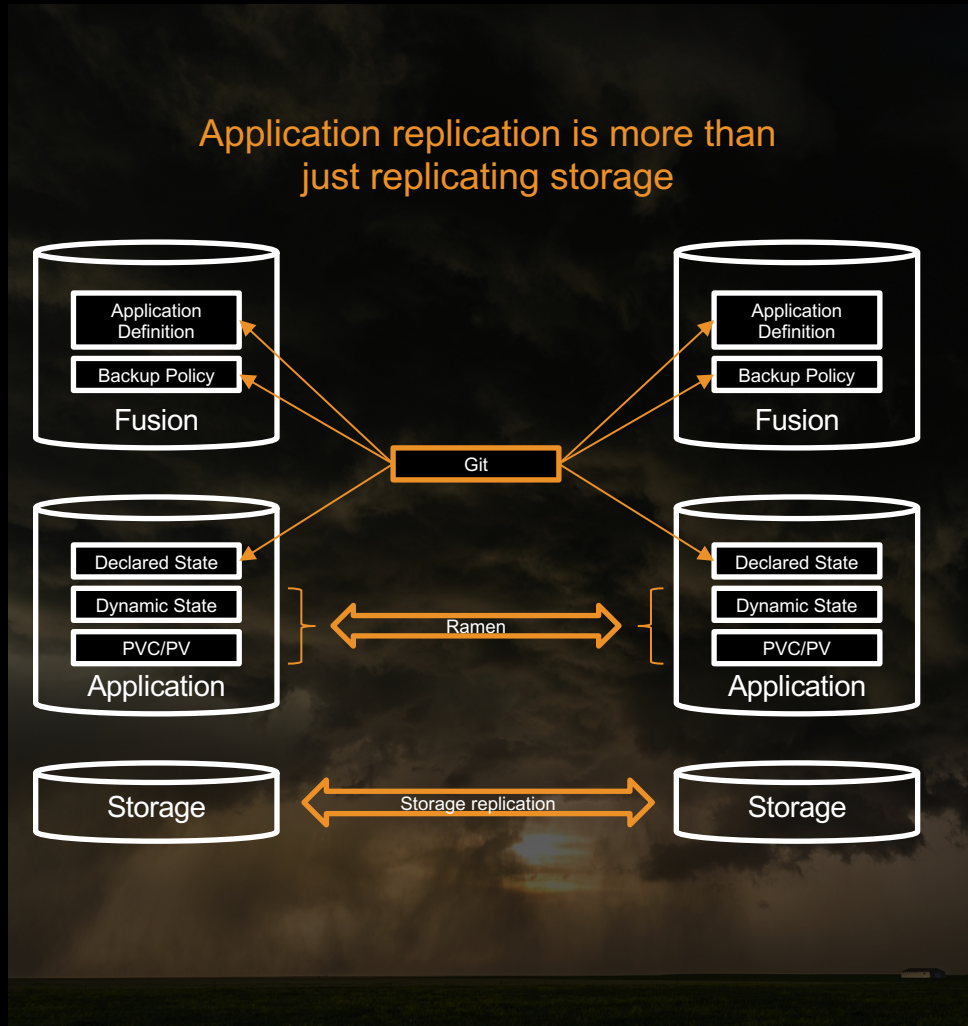

* Roadmap: mid 2023

Storage Fusion Metro DR in-action

- HCI A and HCI B see the exact same data
  RPO is zero
- Trigger relocations manually or automated
  GUI, CLI, K8s Custom Resource, RHACM (roadmap)
- Relocation is a ‘pull’ action
  RTO in seconds

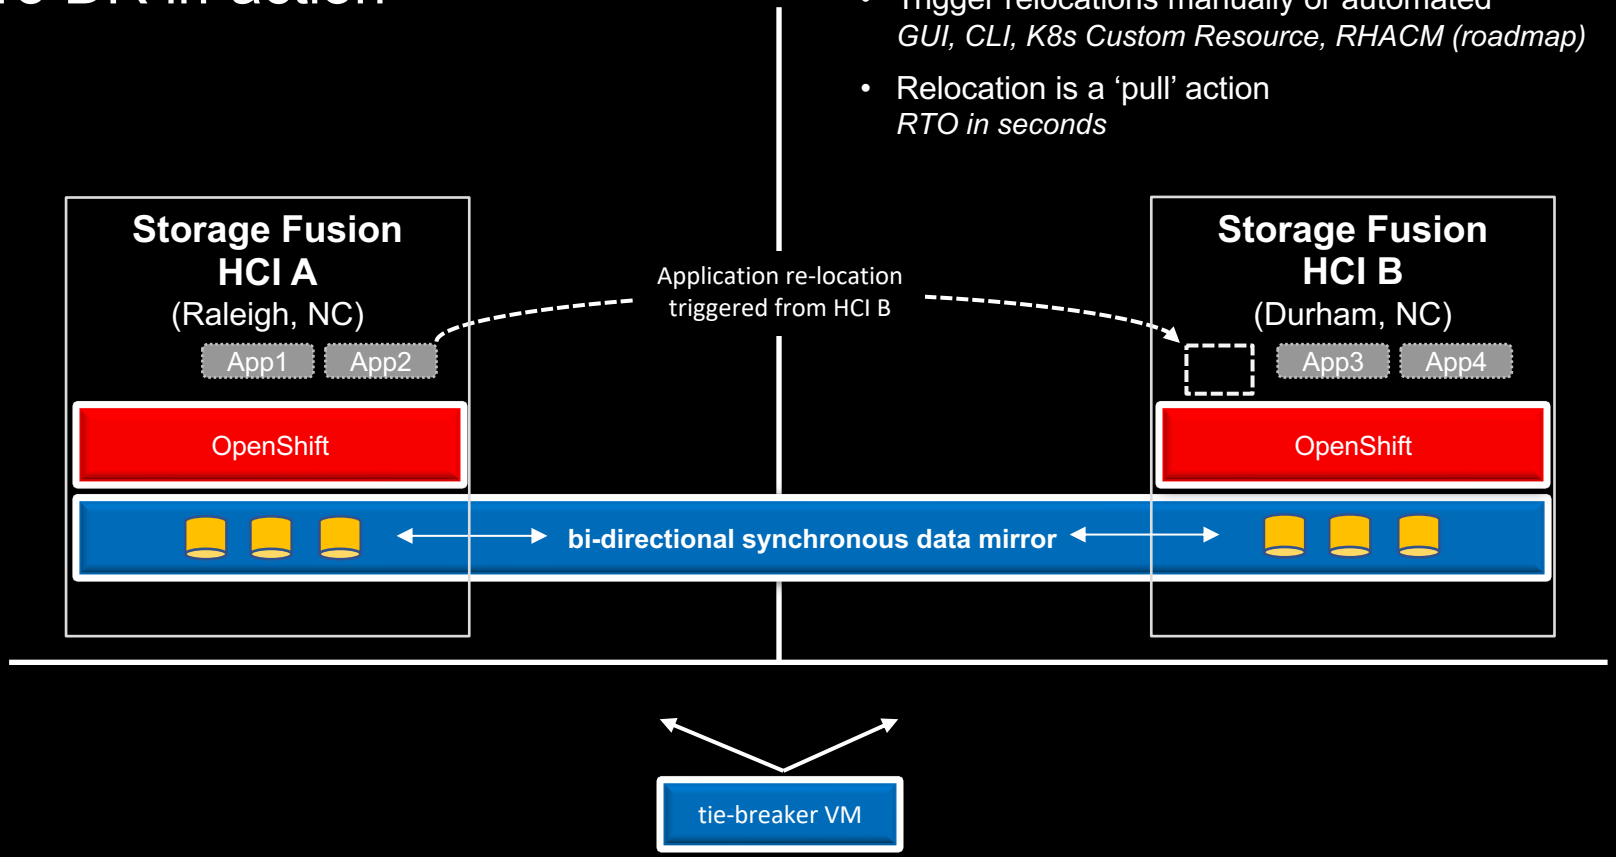

Storage Fusion HCI System

Enable Metro-DR per Application

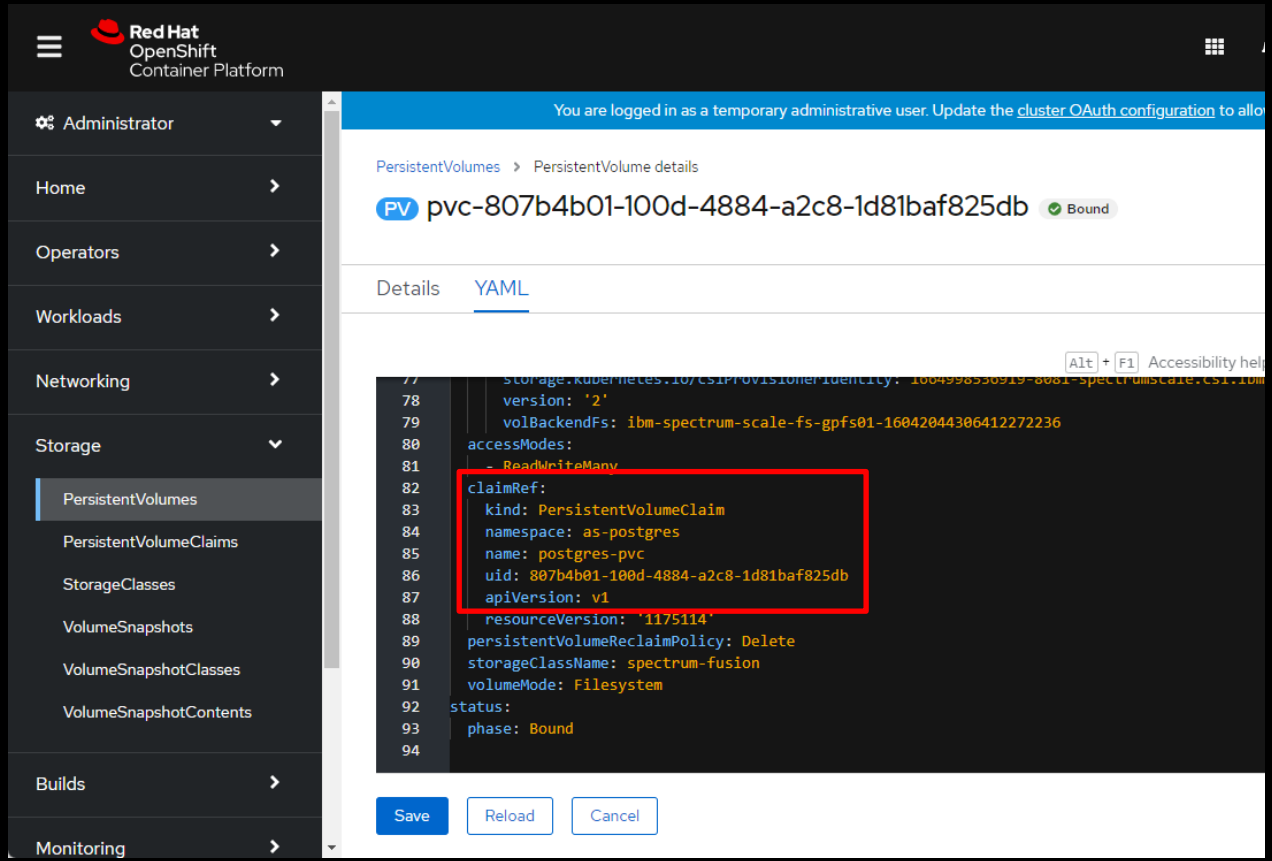

IBM Storage Fusion will:

A) Create a second, sync copy of all storage class content

B) leverage Ramen DR to replicate the PV metadata

When the Application is deployed on ClusterB (including the PVC), Openshift will verify if an existing PV has a claimRef to the PVC.
Only if not a new PV will be generated

-> ensure replication is active prior to deploying the application on Cluster B

Storage Fusion HCI System

Enable Metro-DR per Application

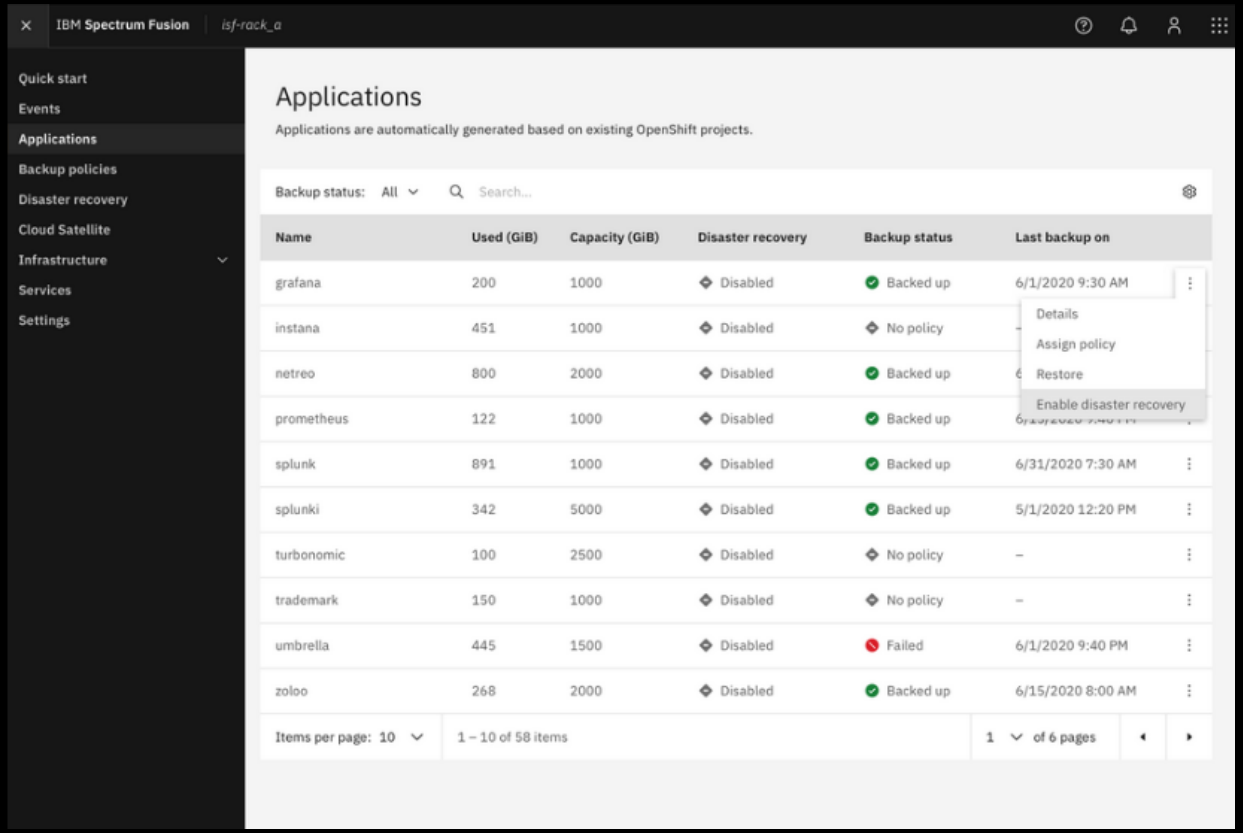

Storage Fusion HCI System

The Metro-DR configuration is part of the Storage Fusion setup

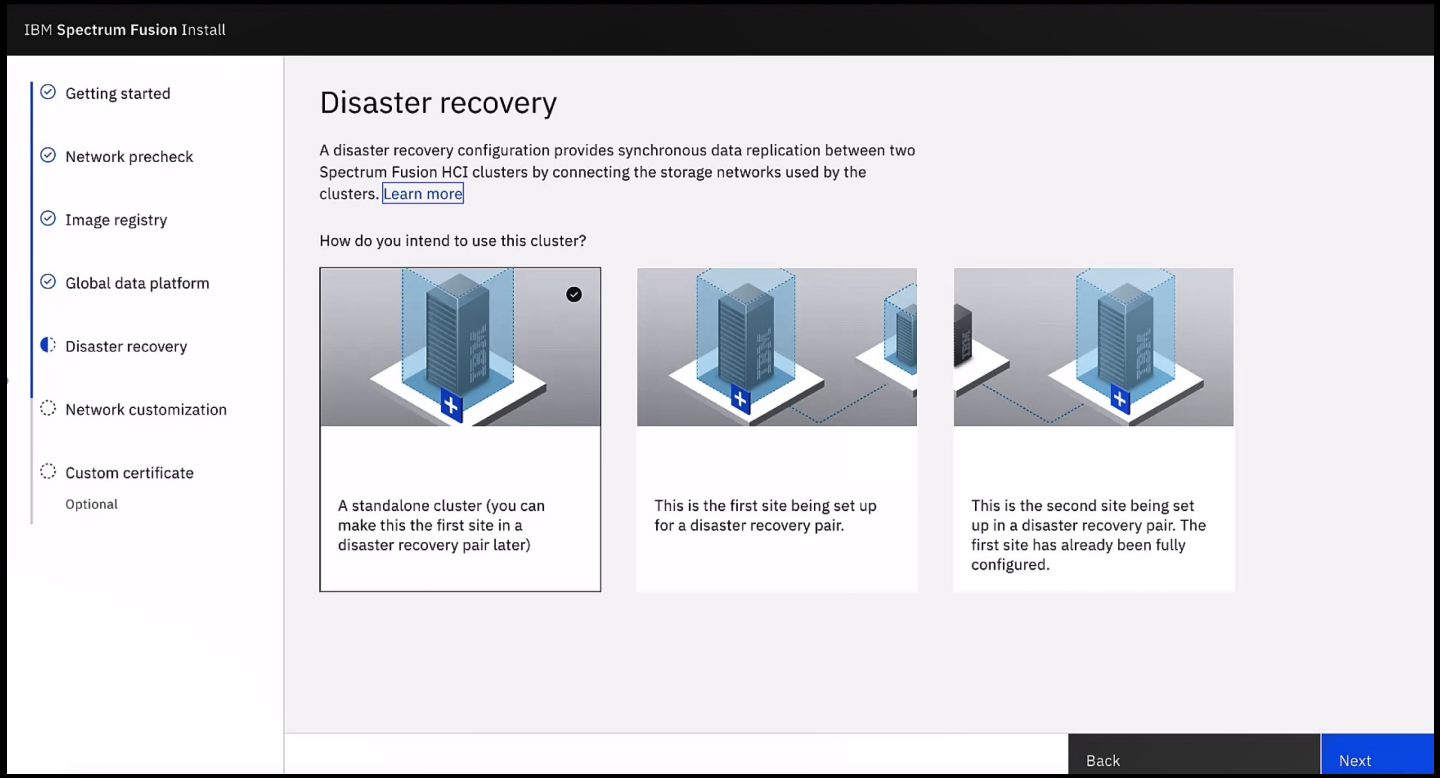

Storage Fusion HCI System

Automate install of the tie-breaker

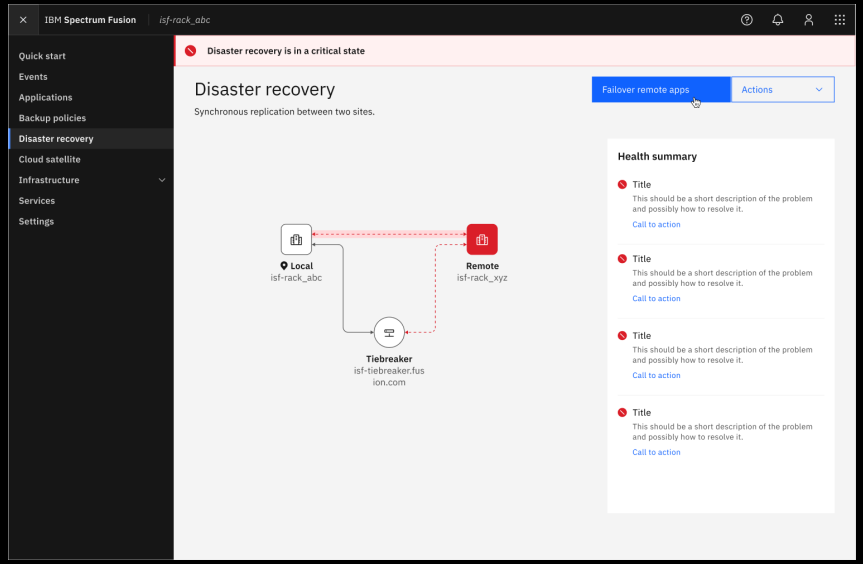

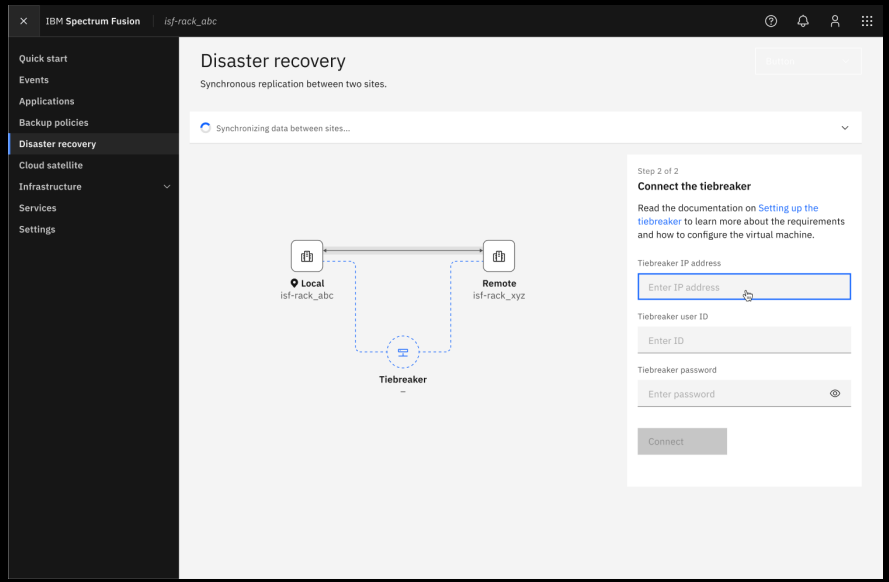

Simplified monitoring of the MetroDR configuration

Regional disaster recovery (ROADMAP 2023)

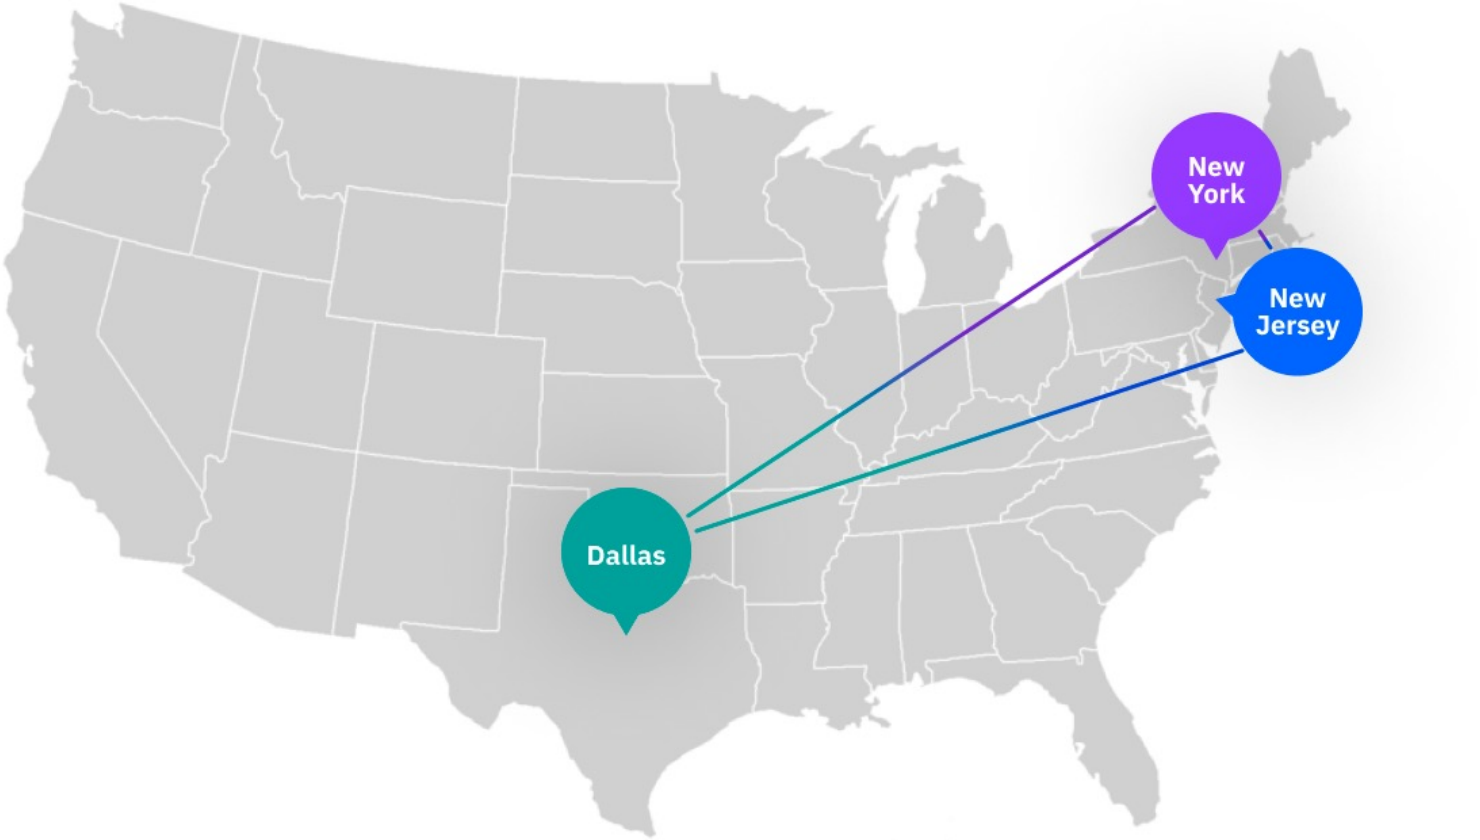

IBM Storage Fusion

Storage Fusion HCI System
Day 2 Operations

What are Day 2 operations?

Imagine you're moving into a house. 

Day 1 operations are moving into the house (installation / deployment).

Day 2 operations are the "housekeeping" stage of a systems life cycle.
The care and feeding of the system,
maintaining the overall stability and health in production.

- Monitoring
- Troubleshooting
- Updating
- Increasing capacity
- Backups and Restores

IBM Storage Fusion HCI

Day2 Operations

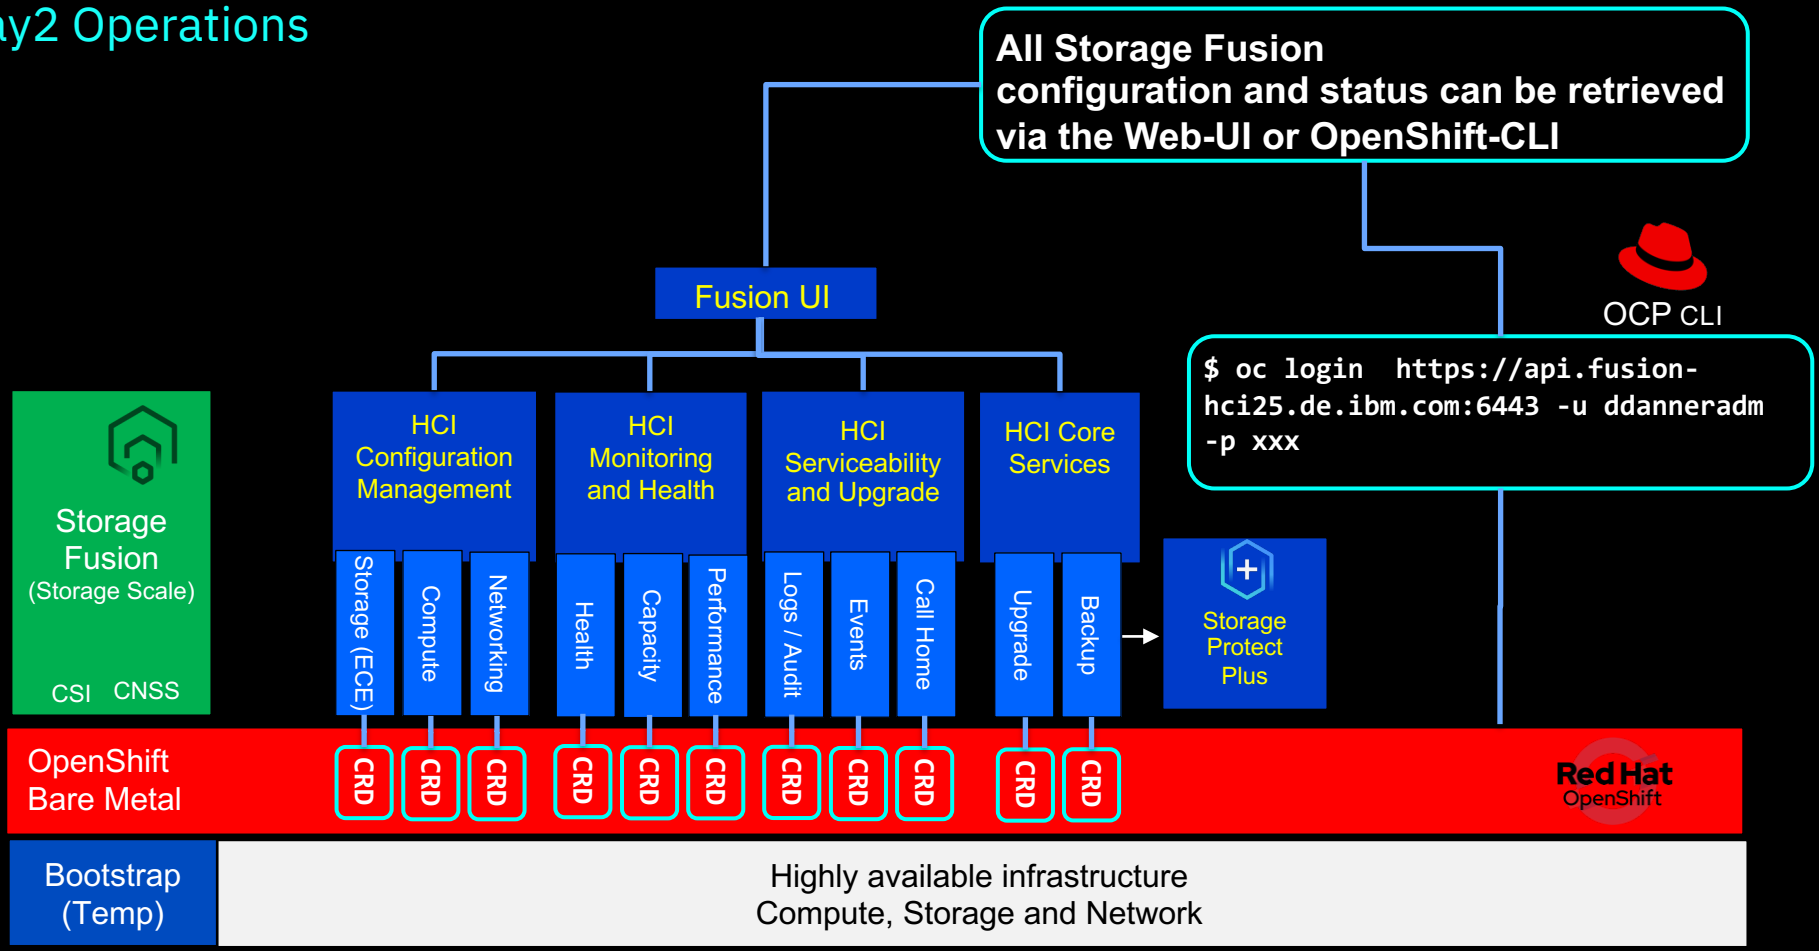

Fusion HCI System

Single Pane of Glas UI

Integrated lifecycle management

- Red Hat OpenShift Operators are used for managing the hardware of Fusion HCI System
- The inventory of all compute nodes and Ethernet switches is presented using an intuitive GUI
- The status of the hardware is shown along with menus of available actions
- Regular updates of the firmware of all components are provided
- Rolling updates

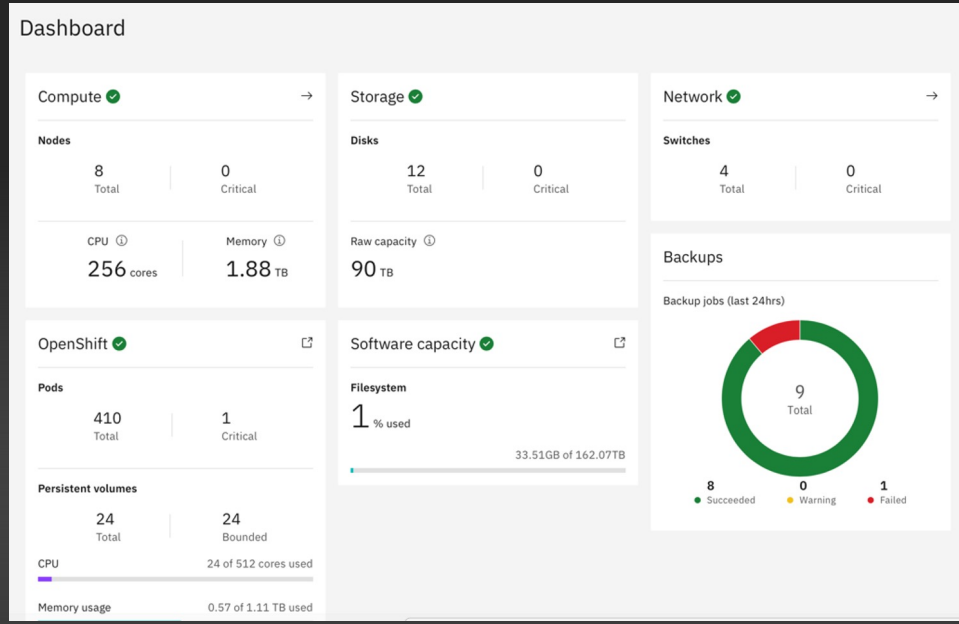

Fusion HCI System

Two ways to manage, and it's the same everywhere!

Management UI

- Guided experience
- Simplified workflows

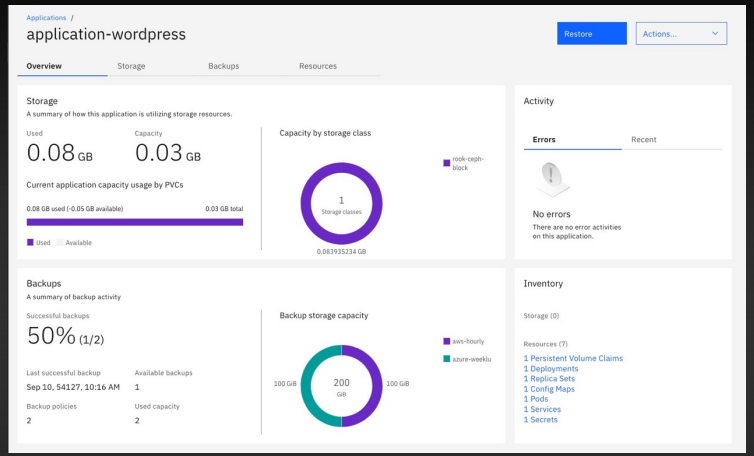

"As code" via custom resources

- Kubernetes native APIs
- Check configuration policies into git for easy deployment into multiple clusters (GitOps)

```yaml
apiVersion: data-protection.isf.ibm.com/v1alpha1
kind: PolicyAssignment
metadata:
  name: wordpress-aws-daily
  namespace: ibm-spectrum-fusion-ns
spec:
  application: wordpress
  backupPolicy: aws-daily
  runNow: true
```

Enable I&O administrators to define and deliver data services that can be consumed by application teams on demand

Storage Fusion

Provides a UI and As Code to easily ***define backup locations***

Adding a backup location

As code

```yaml
apiVersion: data-protection.isf.ibm.com/v1alpha1
kind: BackupStorageLocation
metadata:
  name: aws-bucket
  namespace: ibm-spectrum-fusion-ns
spec:
  type: aws
  credentialName: "aws-credentials"
  params:
    bucket: backup-bucket
    region: us-east-1
```

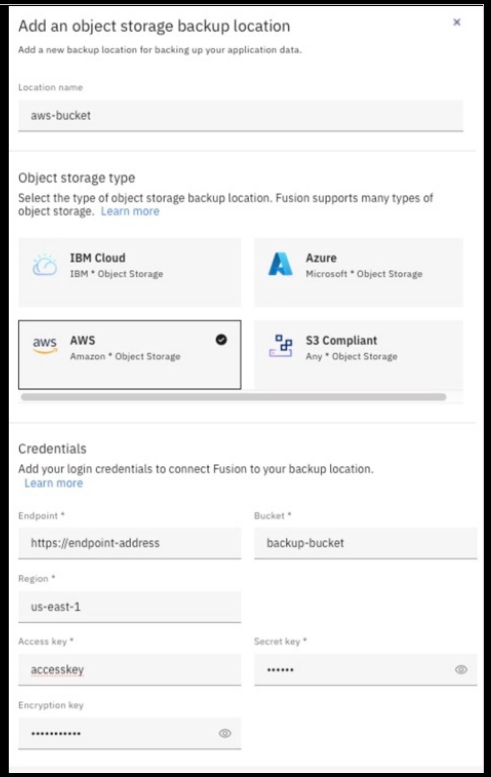

Storage Fusion

Provides a UI and As Code to easily ***assign backup policies***

Creating a policy

As code

```yaml
apiVersion: data-protection.isf.ibm.com/v1alpha1
kind: BackupPolicy
metadata:
  name: aws-daily
  namespace: ibm-spectrum-fusion-ns
spec:
  backupStorageLocation: aws-bucket
  provider: isf-ibmspp
  retention:
    number: 7
    unit: day
  schedule:
    cron: 0 12 * * *
    timezone: America/Los_Angeles
```

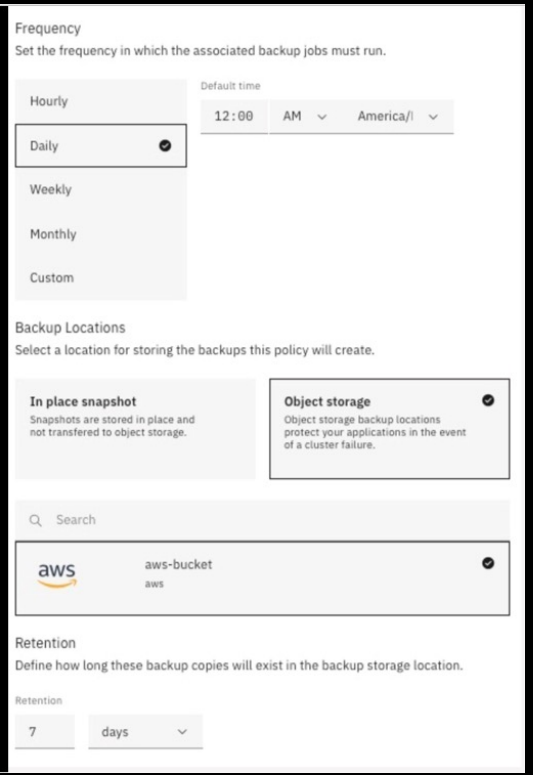

Storage Fusion

Provides a UI and As Code to easily ***restore an application***

33

Restoring an application

As code

```yaml
apiVersion: data-protection.isf.ibm.com/v1alpha1
kind: Restore
metadata:
  name: wordpress-restore-2022-03-01
  namespace: ibm-spectrum-fusion-ns
spec:
  backup: wordpress-backup-2022-03-01
```

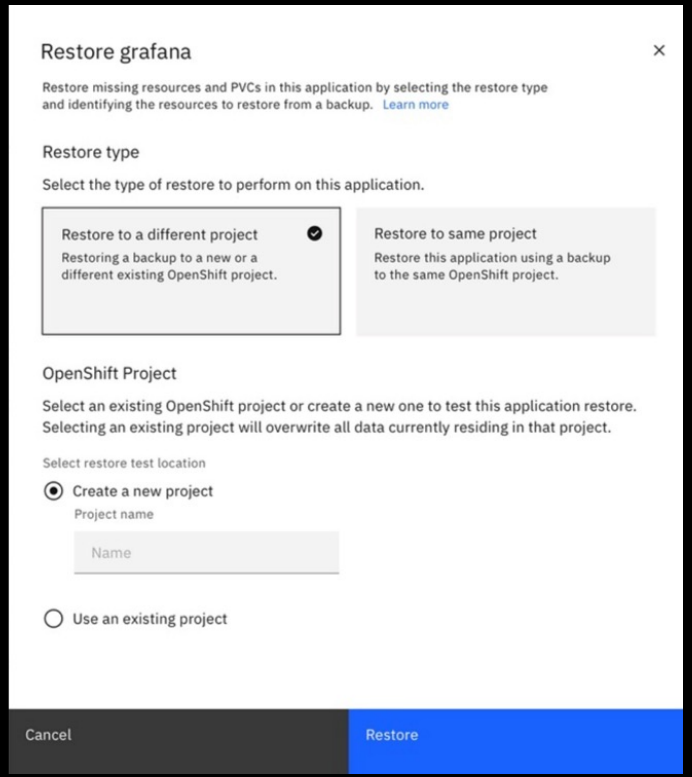

Backup Status

Validate the current Configuration and Status of your Backups

```
$ oc get backuppolicies.data-protection.isf.ibm.com
NAME              PROVIDER     SCHEDULE       RETENTION   RETENTIONUNIT
daily-snap-10d    isf-ibmspp   00 1  * * *    10          days
hourly20days      isf-ibmspp   35 *  * * *    20          days
weekly-full-10d   isf-ibmspp   00 0  * * 0    10          days

$ oc get backupstoragelocations.data-protection.isf.ibm.com
NAME                PROVIDER     PHASE
cos-ibmcloud-hs     isf-ibmspp   on-line
cos-onprem-zone4    isf-ibmspp   on-line
in-place-snapshot   isf-ibmspp   on-line

$ oc get fbackup
NAME                                                                PROVIDER     PHASE       STARTTIME
ENDTIME                EXPIRATION             OBJECTSNUM   OBJECTSPROCESSED

ddannerdemo-hourly20days-ddannerdemo-16585185701                    isf-ibmspp   Completed   2022-07-22T19:35:00Z
2022-07-22T19:38:07Z   2022-08-11T19:35:00Z   2

ddannerdemo-hourly20days-ddannerdemo-16585221701                    isf-ibmspp   Completed   2022-07-22T20:35:00Z
2022-07-22T20:38:07Z   2022-08-11T20:35:00Z   2
```

Storage Fusion OpenShift Integration

Display all Events from the UI on OpenShift

```
$ oc get events -n ibm-Storage-fusion-ns
```

```
LAST SEEN   TYPE      REASON            OBJECT                          MESSAGE

10d         Warning   ISFEventManager   deployment/eventmanager         BMYNW0055 - CumulusLinkDOWN, hardware
address: mgmt1-78ac99a:78AC95A, port: swp47 - 607
```

Network Configuration and Status

Access the current Network configuration and Status form the OCP CLI

```
$ oc get switches
NAME              AGE
hspeed1-78ac99a   35d
hspeed2-78ac99a   35d
mgmt1-78ac99a     35d
mgmt2-78ac99a     35d

$ oc get switches hspeed1-78ac99a -o yaml
...
spec:
  cutomerPorts:
  - "31"
  - "32”
...
 cutomerPortStatus:
  - inputPkt: "387925128"
    outputPkt: "126215691"
    portNumber: "31"
    portState: UP
    speed: 100G

  - inputPkt: "0"
    outputPkt: "0"
    portNumber: "32"
    portState: ADMDN
    speed: N/A
```

```
status:
  availableFirmwareVersion: 4.3.0
  currentFirmwareVersion: 4.3.0

  temperatureStatus:
    state: OK

  fanStatus:
    state: OK

$ oc get vlans rack-vlans -o yaml
 torVlanSpec:
  - Description: Internal Storage VLAN
    multipleLinks: "no"
    vlanId: "3201"
    vlanName: Internal Storage
    vlanType: Internal Storage VLAN
  - Description: Internal Management VLAN
    multipleLinks: "no"
    vlanId: "4091"
    vlanName: Internal Management
    vlanType: Internal Management VLAN
  - Description: Data VLAN on TOR switch
    multipleLinks: "no"
    vlanId: "40"
    vlanName: vlan40
```

Firmware Package Details

See the Firmwaredetails of each Component in great detail

```
$ oc get computefirmwares  firmware-compute-1 -o yaml

status:
  availableFirmwareVersion: isf-2.1.0
  currentFirmwareVersion: isf-2.1.0

  firmwareInfos:

  - description: XCC Firmware
    name: BMC (Primary)
    version: D8OT16J-3.01

 - description: UEFI Firmware/BIOS
    name: UEFI
    version: D8E116E-2.01

  - description: NVMe DiskDrive Firmware
    name: NVMeDiskDriveFirmware
    version: CQ36

  - description: DETA Power Supply Firmware
    name: DETA Firmware
    version: "4.51”
```

```
  - description: Mellanox Firmware Bundle
    name: ThinkSystem Mellanox ConnectX-6 HDR100/100GbE
QSFP56 2-port PCIe VPI Adapter (Mellanox Firmware Bundle)
    version: 20.28.2006

- description: Mellanox uEFI driver
    name: ThinkSystem Mellanox ConnectX-4 Lx 10/25GbE SFP28
2-port PCIe Ethernet Adapter (Mellanox uEFI driver)
    version: 14.21.17

  - description: Mellanox Base Firmware
    name: ThinkSystem Mellanox ConnectX-4 Lx 10/25GbE SFP28
2-port PCIe Ethernet Adapter (Mellanox Base Firmware)
    version: 14.28.2006
version: 14.21.17

...


  ocpNodeName: compute-1.fusion-hci.lab.escc.de.ibm.com
  systemProdName: 7D2XCTO1WW
  systemSerialNumber: xxxxxxxx
  updateRequired: false
```

Update and Software Status

Access the current Software Versions and see if a update is needed from the OCP CLI

```
$  oc get updatemanagers

NAME      AVAILABLE   UPGRADING   ERROR
version   True        False       False


$  oc get updatemanagers  -o yaml

  status:
    componentsState:
      fusionSwState:
        availableVersion: 2.2.1
        currentVersion: 2.2.1

      nodeState:
        nodes:
        - availableFirmwareVersion: isf-2.1.0
          currentFirmwareVersion: isf-2.1.0
          nodeName: firmware-compute-2
          systemSerialNumber: J101LNY2
          updateAvailable: "no"

      openshiftState:
        currentVersion: 4.8.19
        updateAvailable: "no"

      openshiftState:
        currentVersion: 4.8.19
        updateAvailable: "no”

      scaleState:
        availableUpgradeVersion: 5.1.3.1
        currentVersion: 5.1.3.1
        updateAvailable: "no"
        upgradeCompleted: false

      sppState:
        currentVersion: 10.1.10.2709
        updateAvailable: "no"
        upgradeCompleted: false

      switchState:
        switches:
        - availableFirmwareVersion: 4.3.0
          currentFirmwareVersion: 4.3.0
          switchName: hspeed1-78ac99a
          switchRole: primary
          switchSerialNumber: M1LM11E000M
          switchType: hspeed
          updateAvailable: "no"
```

Updating the OpenShift Software Stack

The following slides explore how different OpenShift Stacks are updated

1. Fusion HCI
2. OpenShift on a Hypervisor
3. OpenShfit on Bare Metal

As required by common security audits (e.g. PCI-DSS)
the full software stack need to be updated
including the Firmware of involved infrastructure components.

Updating Storage Fusion HCI

Storage Fusion HCI
Has full control of the underlaying Hardware Stack.
Therefore, we can update the following components:

- Software Defined Storage
- Backup Tool
- Server Firmware (NVMe drives, Network Cards, Bios …)
- Top of Rack Switches

HCI Advantages
- All updates can be managed via the OCP and Fusion UI
- IBM delivers the full Software Stack
- This Stack is fully tested and the interoperability is ensured
- IBM provides support for the full Software Stack

Storage Fusion HCI

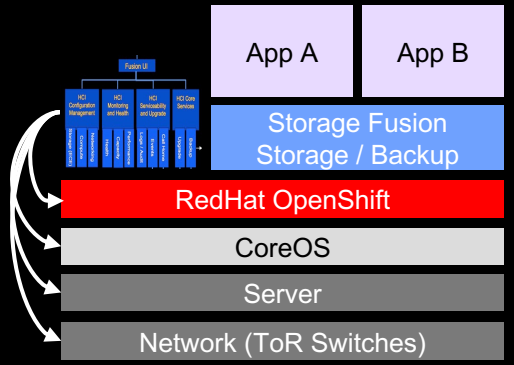

Bare metal OpenShift Stack

Updating OpenShift on a Hypervisor

OpenShift running in VMs

- This software stack has multiple layers with multiple UIs
- Multiple vendors are involved for support and licensing
- The software packages and update instructions have to be obtained from multiple sources
- Interoperability of the components have to be ensured and tested. (e.g. Backup Tool -> Storage CSI -> OCP version)
- In many organizations, multiple teams are involved. Timing a change management has to be coordinated to avoid outages

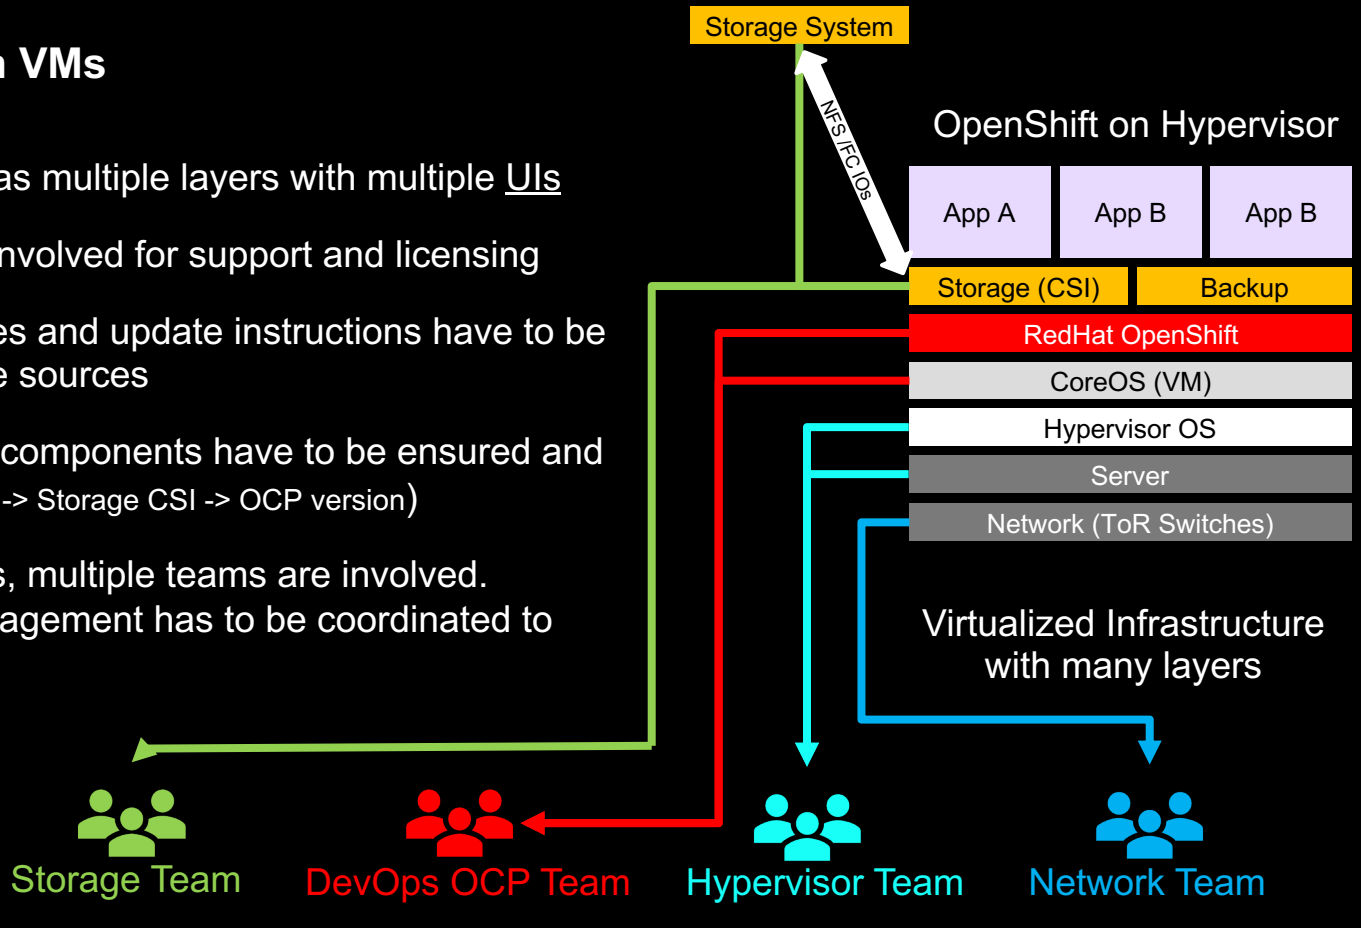

Updating Bare Metal OpenShift

Storage Fusion SDS on Bare Metal

Bare Metal OpenShift
Can be installed via RH IPI or assisted installer.
OpenShift has no control of the Hardware underneath.
Firmware updates have to be done outside of OpenShift.

Storage Fusion SDS for Bare Metal
Requires a user provisioned bare metal OCP Cluster.
The Storage Fusion Operator is installed on this cluster, providing storage and data protection services.
Storage Fusion SDS does only update its own components.

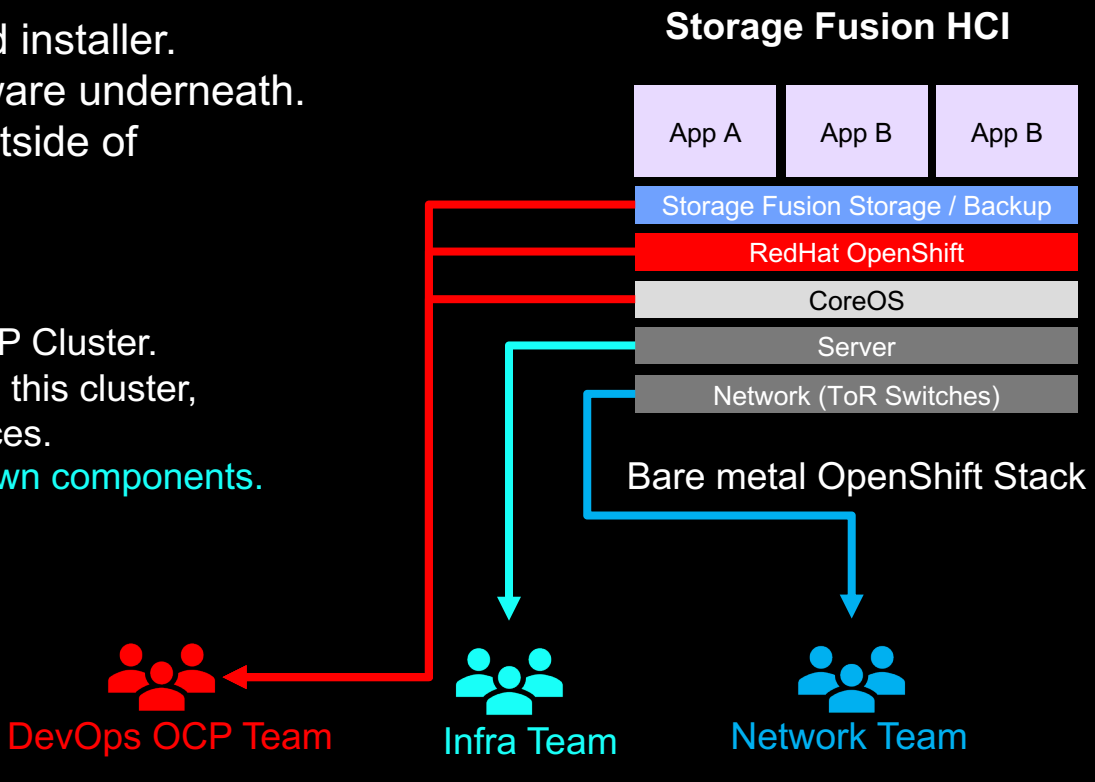

Storage Fusion HCI Upgrade Procedure

The Storage Fusion HCI upgrade procedure is divided into four steps:

1. Fusion management software
2. Fusion Services update (Storage and Backup Software)
3. Compute nodes and Switches firmware
4. OpenShift Container Platform update

Fusion management software update (1/4)

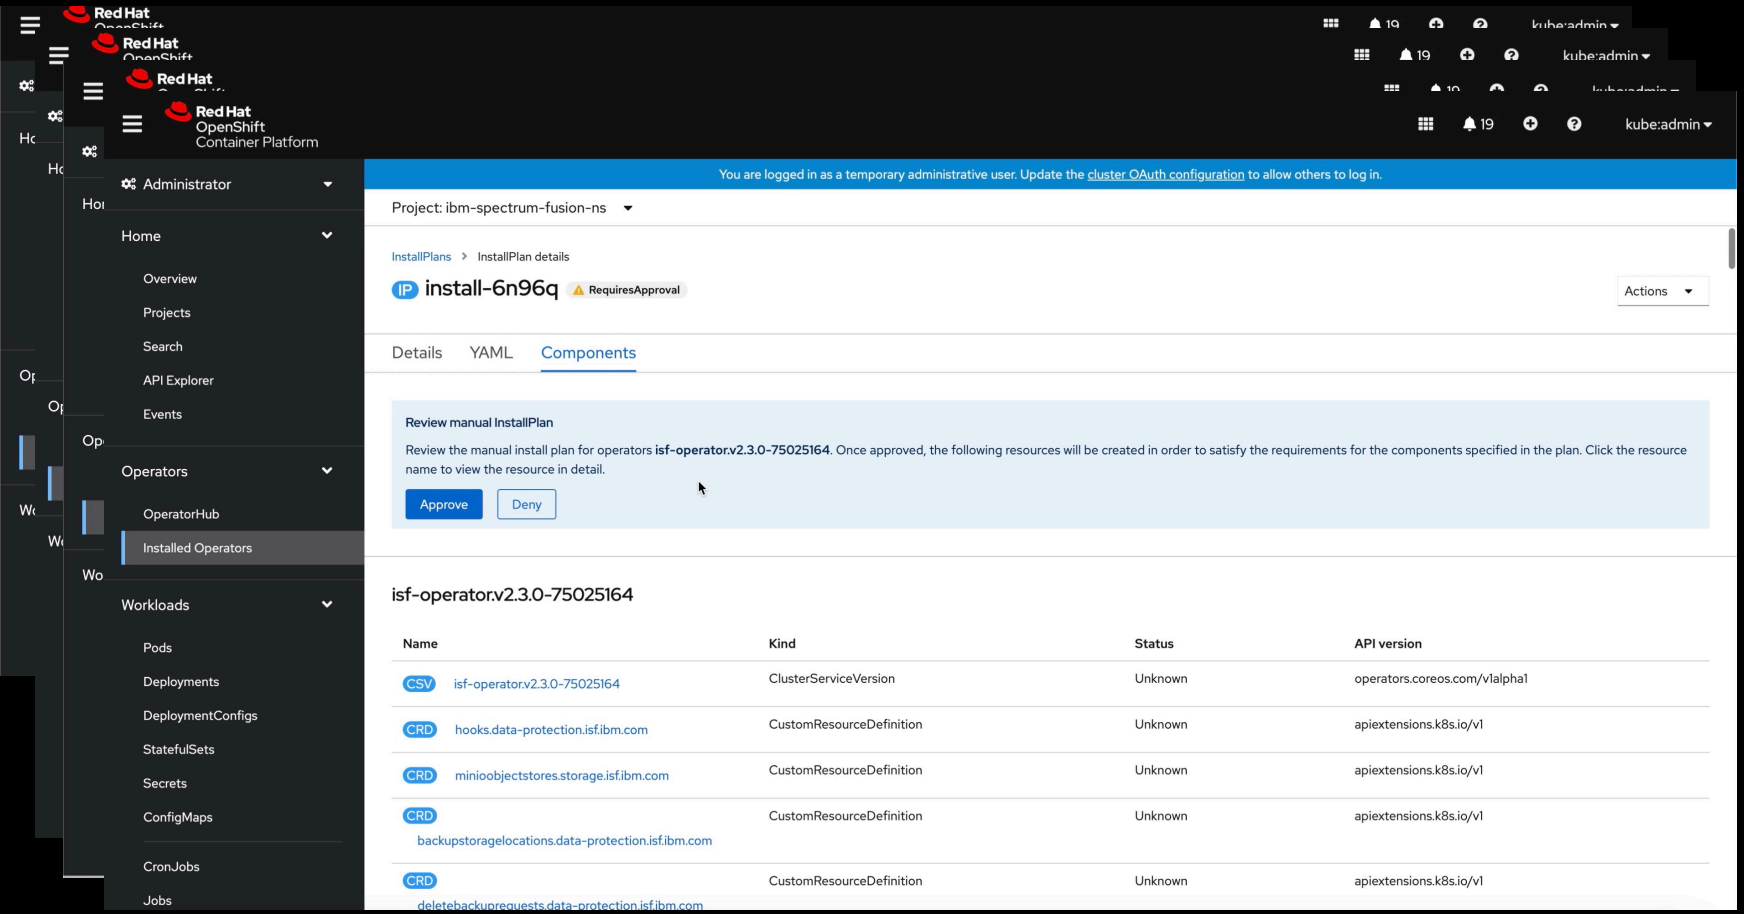

Fusion Services update (2/4)

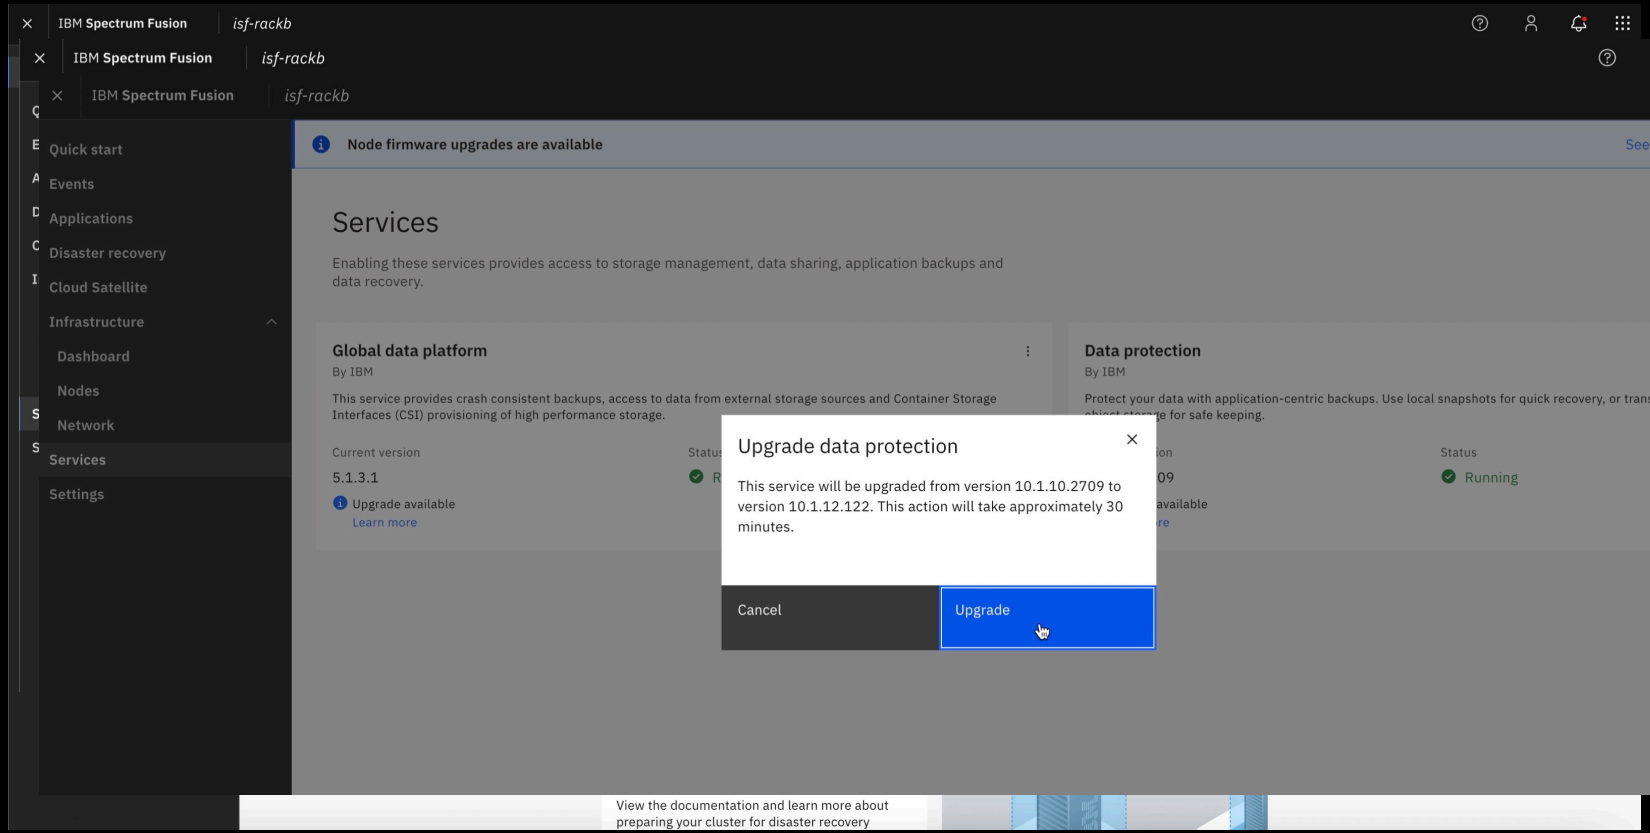

Compute and Switch Firmware upgrade (3/4)

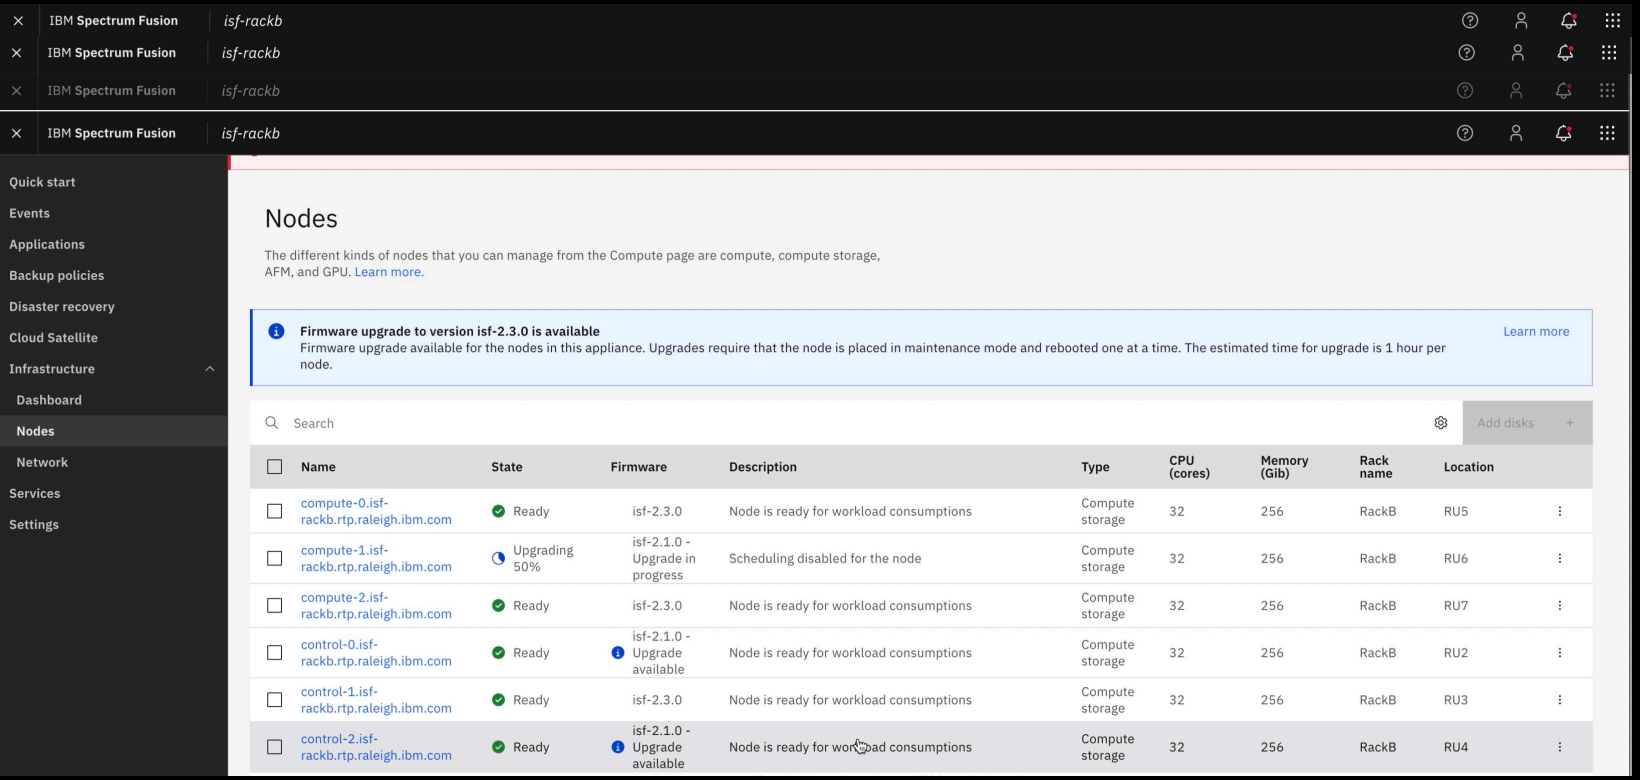

OpenShift upgrade (4/4)

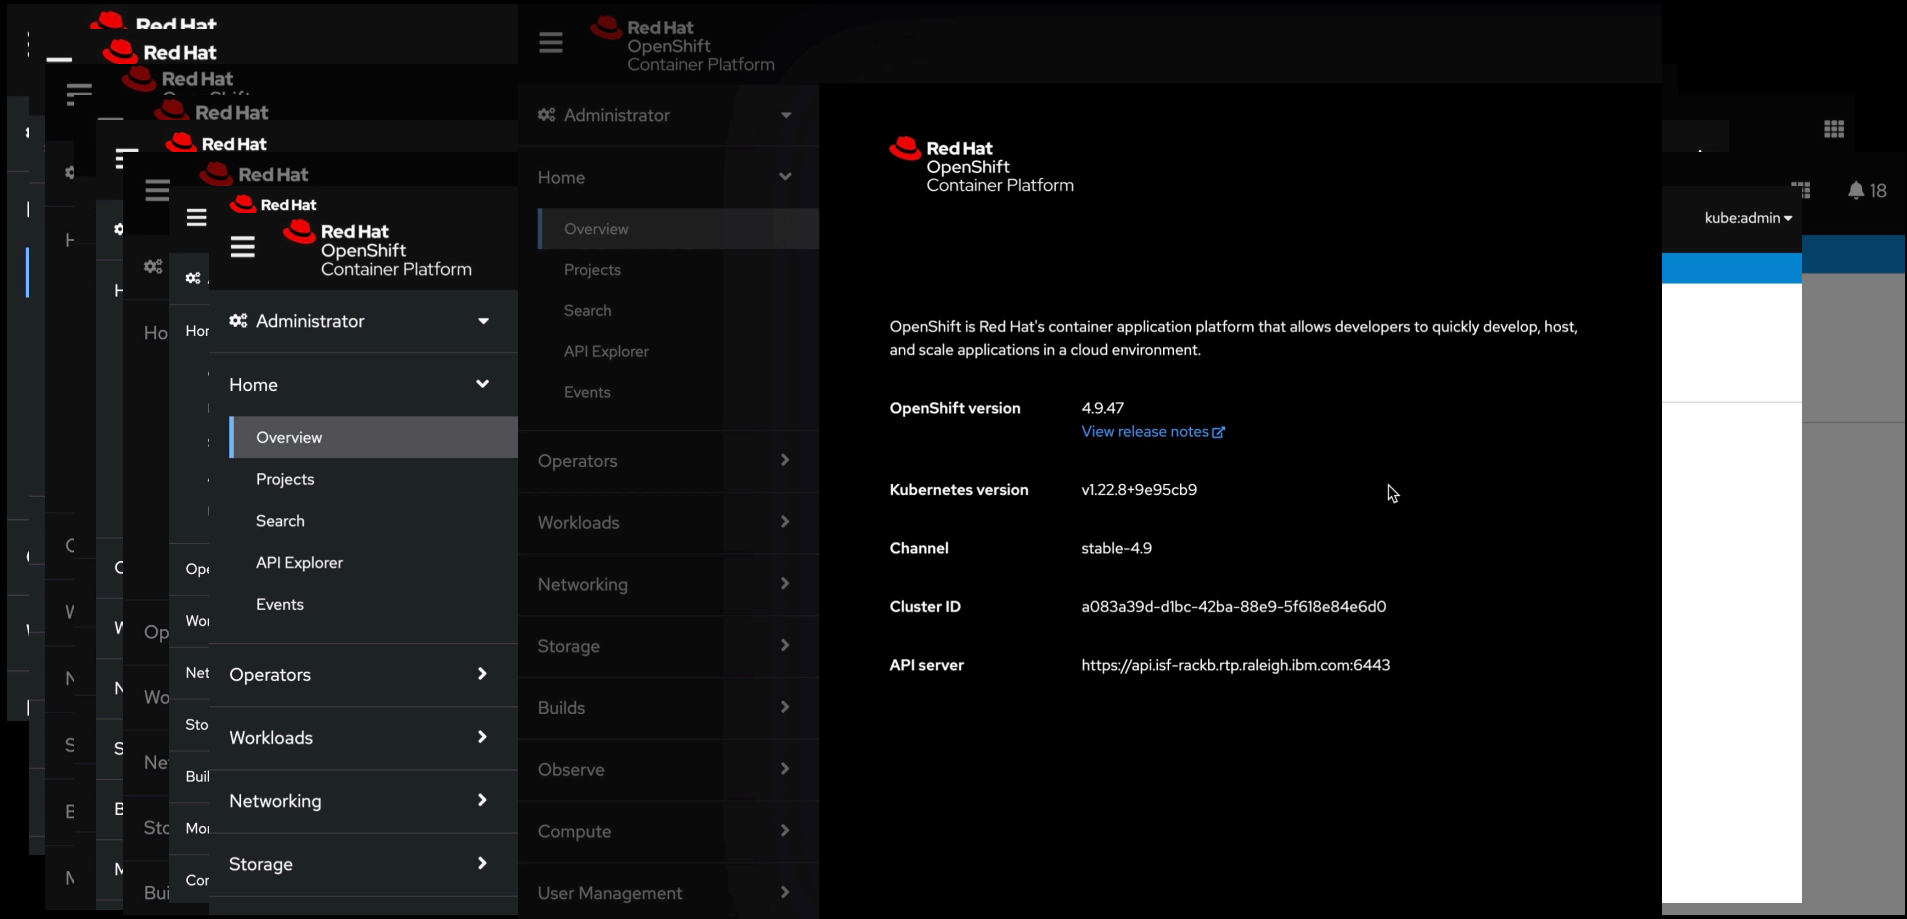

Fusion HCI System

Integrated HCI lifecycle management services simplifies day-2 administration

A Fusion HCI System Administrator can

- Add storage to the cluster without the assistance of a storage administrator
- Apply firmware updates to switches and servers without the assistance of network or data center engineers
- Apply software and firmware updates in a rolling update fashion, without application downtime
- Monitor status of hardware and systems software

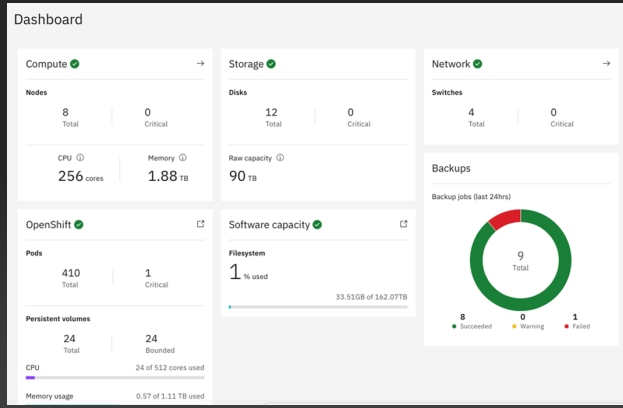

Redbooks


Draft Document for Review February 28, 2023 6:17 pm REDP-5688-00
Redbooks
ibm.com/redbooks
IBM Spectrum Fusion Product Guide
Jon Herd
Daniel Danner
Axel Koester
Cloud
Storage
IBM
Redpaper


Draft Document for Review March 10, 2023 9:09 pm REDP-5708-00
Redbooks
ibm.com/redbooks
IBM Storage Fusion HCI System
Metro Sync DR Use Case
Pallavi Singh
Prakash Chandraya
Seema Pandhre
Cloud
Storage
In partnership with
IBM Academy of Technology
IBM
Redpaper

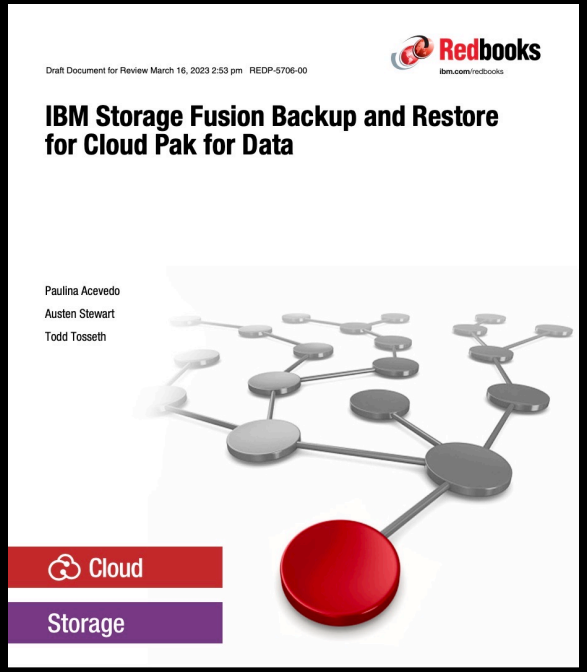
Draft Document for Review March 16, 2023 2:53 pm REDP-5706-00
Redbooks
ibm.com/redbooks
IBM Storage Fusion Backup and Restore
for Cloud Pak for Data
Paulina Acevedo
Austen Stewart
Todd Tosseth
Cloud
Storage
IBM
Redpaper


IBM

IBM Storage Fusion

*Get container workloads out of pilot and **into production faster***

CORE CAPABILITIES

Container Native Storage
Directly provision storage volumes through the Storage Fusion CSI interface

Backup and restore
applications and VMs to object stores for data resiliency

Application resiliency
via sync. or async. replication across availability zones and regions

Simple, centralized management
of all OpenShift data services.
Single pane of glass UI

Storage Fusion HCI

Turnkey Bare-metal OpenShift cluster-in-a-box

Bare Metal is better!
Save the Hypervisor license and performance tax.
Benefit from cheaper OpenShift subscriptions by licensing sockets instead of cores.

Fully automated installation
From power on to a production ready OpenShift in 5 hours.

Scalable
Starts small with 6 servers and scales up to 20.
Start with 50 TB usable capacity on grow online to 900TB.

One point of contact
Hardware and Software support from IBM.
For the full SW stack, including OpenShift, SDS Storage and Backup.

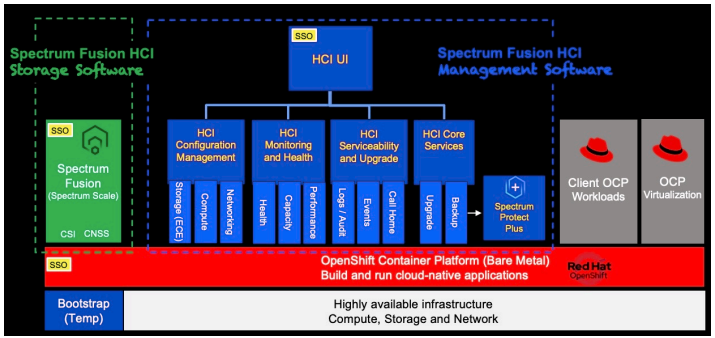

IBM Storage Fusion HCI

Höchste Leistung und Sicherheit für KI- und Containeranwendungen

Was ist Storage Fusion HCI?

Eine hyperkonvergente Hardware- und Softwarelösung für Red Hat OpenShift, die einfach und sicher eine hybride Containerplattform bereitstellt.

Welche Kernfunktionen bietet Storage Fusion HCI unternehmenskritischen KI & Containeranwendungen?

Container Native Storage
Direkte Bereitstellung von performantem Speicher über die Storage Fusion CSI-Schnittstelle

Einfache Sicherung und Wiederherstellung
von Containern und VMs für maximale Datenausfallsicherheit mit IBM Storage Protect Plus

Ausfallsicherheit der Anwendung
durch synchrone oder asynchrone Replikation

Einfaches, zentrales Management
der Appliance, der Datensicherung sowie der OpenShift-Datendienste über eine einzige Benutzeroberfläche

Weitere Informationen: https://www.ibm.com/de-de/products/Storage-fusion

Warum entscheiden sich Kund_innen für Storage Fusion HCI?

✓ **Hervorragender ROI**
Bare Metal OpenShift ist schneller und günstiger, da es ohne Hypervisorebene auskommt. **Ermöglicht 80% Einsparung bei der OpenShift-Subscription**

✓ **Einfaches Deployment von Red Hat OpenShift**
vom Einschalten bis zum **produktionsbereiten OpenShift** in wenigen Stunden

✓ **Herausragende Leistung**
Meistert anspruchvolle Workloads mit bis zu 20 Servern, hochperformantem Datenzugriff und optionalen NVIDIA A100 GPUs

✓ **Data Lake Management**
Schnelles **Einbinden von sowohl weit entfernten Daten als auch Cloudspeicher** durch die Integration von IBM Storage Scale

✓ **Einfaches Infrastruktur Management**
Flexibles Skalieren der Server & Speicherkapazität, sowie **Hardware- und Software-Support von IBM** über den gesamten Stack, einschließlich OpenShift, Storage und Backup