IBM

ADVANCED TECHNOLOGY GROUP (ATG)

Accelerate with ATG: IBM Fusion and RedHat OpenShift Virtualization

Take your Data Center to the Next Level

Lloyd Dean

Principal Storage Technical Specialist
IBM Storage for Red Hat OpenShift SME
IBM Advanced Technology Group (ATG)
IBM Technology Sales, Americas
Email: LloydDean@IBM.com

IBM

Accelerate with ATG Technical Webinar Series

Advanced Technology Group experts cover a variety of technical topics.

Audience: Clients who have or are considering acquiring IBM Storage solutions. Business Partners and IBMers are also welcome.

To automatically receive announcements of upcoming Accelerate with IBM Storage webinars, Clients, Business Partners and IBMers are welcome to send an email request to accelerate-join@hursley.ibm.com.

2024 Upcoming Webinars – Register Here!

Boost your IBM Storage Ceph Knowledge: New Ceph 7.1 Features – August 27th, 2024

Important Links to bookmark:

ATG Accelerate Site: https://ibm.biz/BdSUFN

ATG MediaCenter Channel: https://ibm.biz/BdfEgQ

Offerings

Client Technical Workshops

- **IBM Fusion & Ceph: A Deep Dive into Next Gen Storage: August 21st – 22nd in Atlanta, GA**
- **IBM FlashSystem Deep Dive & Advanced Functions: September 18th – 19th in Paramus, NJ**
- IBM DS8900F Advanced Functions
- IBM Cyber Resiliency with IBM Storage Defender
- IBM Storage Scale System & Storage Scale Workshop

TechZone Test Drive / Demo's

- IBM Storage Scale and Storage Scale System GUI
- IBM Storage Virtualize Test Drive
- IBM DS8900F Storage Management Test Drive
- Managing Copy Services on the DS8000 Using IBM Copy Services Manager Test Drive
- IBM DS8900F Safeguarded Copy (SGC) Test Drive
- IBM Cloud Object Storage Test Drive - (Appliance based)
- IBM Cloud Object Storage Test Drive - (VMware based)
- IBM Storage Protect Live Test Drive
- IBM Storage Ceph Test Drive - (VMware based)

Please reach out to your IBM Representative or Business Partner for more information.

***IMPORTANT* The ATG team serves clients and Business Partners in the Americas, concentrating on North America.**

Registration Open!

Storage @ IBM TechXchange Conference 2024

October 21-24, 2024
Mandalay Bay | Las Vegas
#IBMTechXchange

Key Learnings

- Practical how-to advice
- Patterns and best practices
- Success stories, IBM PoV, proven techniques

Featured Products

IBM Storage Defender
IBM Storage Fusion
IBM Storage Scale + IBM Storage Ceph
IBM Tape + IBM SAN
IBM Storage FlashSystem + IBM Storage DS8000

Collaborate. Learn. Play.

Community

IBM Champions
User Groups
Tech Peers
Business Partners

Sandbox

Network
Learn
Collaborate
Play

Roadmaps

Go deep with people in the know and set the stage for where IBM is going in the future

Breakout Sessions

Trends and Directions
User Groups
Product Deep Dives
Meet the Expert
Professional Development
Show the Code
Birds of a Feather
Academic/Research

Accelerate your Career

Labs (Instructor-Led, Self-paced)
IBM Certification Testing
Earn up to 25 hours in CPE credits

https://www.ibm.com/community/ibm-techxchange-conference/

Game On! →

Accelerate with ATG Survey

Please take a moment to share your feedback with our team!

You can access this 6-question survey via Menti.com with code 1708 6924 or

Direct link https://www.menti.com/alwhyze7z1gz

Or

QR Code

ADVANCED TECHNOLOGY GROUP (ATG)

Accelerate with ATG:

Co Speakers & Accelerate Fusion SMEs covering Q&A

Shu Mookerjee - Senior Storage Technical Specialist
Andrew Rice – Storage Technical Specialist

Meet the Speakers

Lloyd Dean is an IBM Principal Storage Technical Specialist in IBM Storage Solutions. Lloyd has held numerous senior technical roles at IBM during his 25 plus years at IBM. Lloyd most recently is leading efforts in the Advanced Technology Group as the IBM Storage for Red Hat OpenShift focal and as a Hybrid Cloud storage solution SME covering IBM Block, File and Object storage solutions and their use cases supporting IBM Cloud Paks. Lloyd is also the ATG Fusion focal and SME.

Andrew Rice is an Infrastructure/Storage Engineer with over 17 years of experience implementing cloud infrastructure design, storage solutions, and virtualization. Andrew's expertise extends across IBM's storage portfolio primarily in IBM Storage Scale, Fusion, Storage Protect, IBM Flashsystems, and encompasses technical proficiencies in VMware and Red Hat OpenShift.

Meet the Speakers

Shu Mookerjee is a Level 2 Certified Technical Specialist with over twenty years at IBM, working in a variety of roles including sales, management and technology. For the last decade, he has focused exclusively on storage and has been the co-author of four (4) Redbooks. Currently, Shu is part of the Advanced Technology Group where he provides education, technical guidance, Proofs of Concept and Proofs of Technology to IBMers, business partners and clients.

ADVANCED TECHNOLOGY GROUP (ATG)

IBM Fusion and RedHat OpenShift Virtualization

Take your Data Center to the Next Level

Storage Concepts

Storage Usage in VMWare vs OpenShift

VMWare Datastore Storage

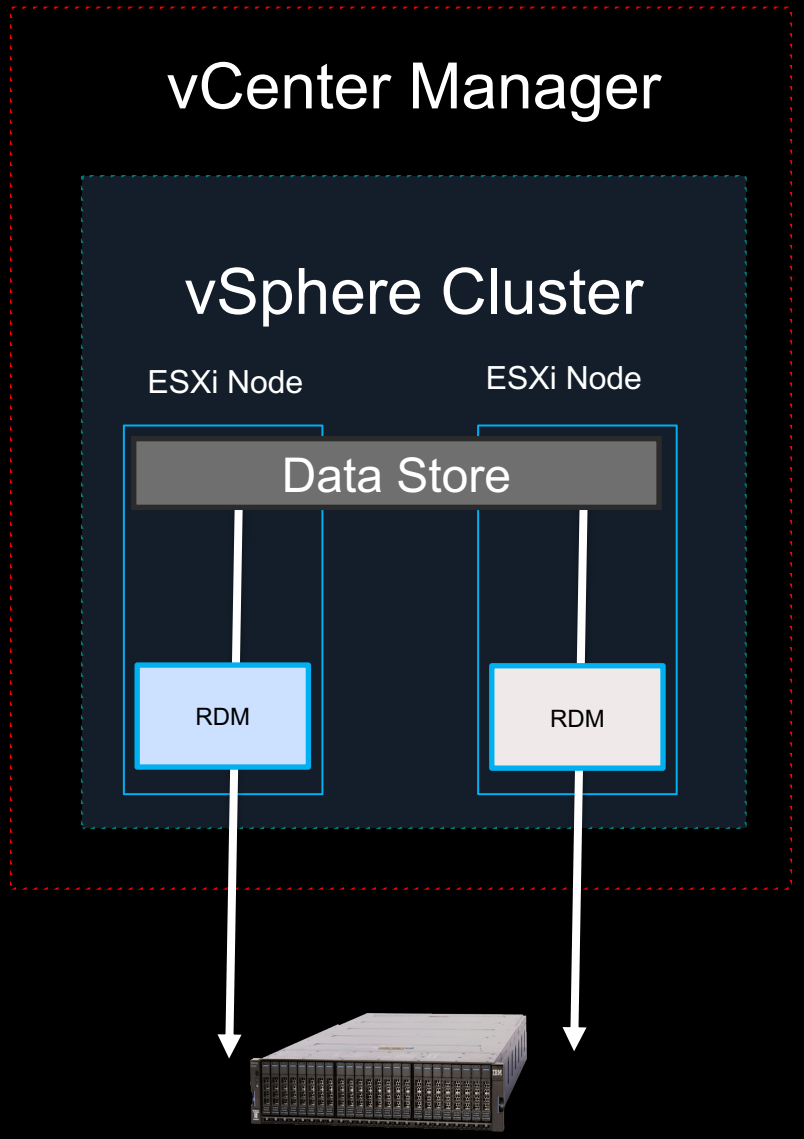

Datastore storage is a VMWare abstraction construct that allows you to deploy and use one or more attached block storage volume(s) in an ESXi or vSphere cluster configuration.

The attached block storage can be provided over Fibre Channel, iSCSI, NVMe/TCP, or IP with NFS.

Container-ready Storage

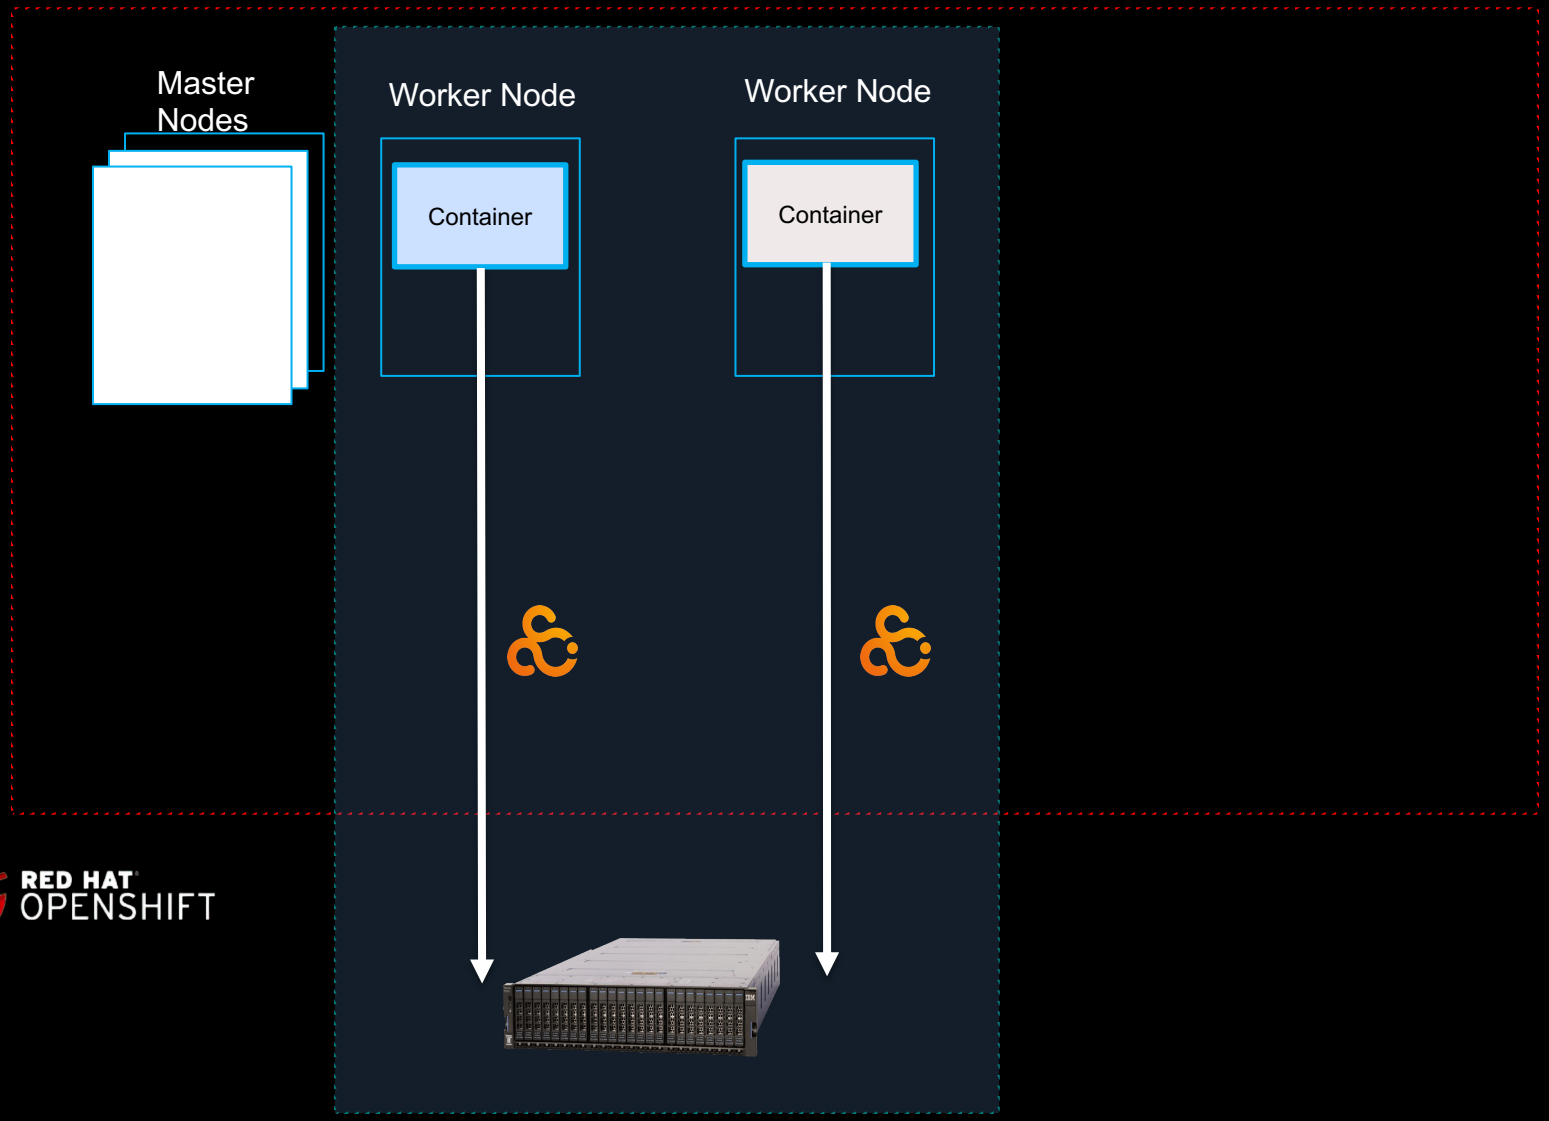

Container-ready Storage is the next-step to scale up/out and leverage external storage.

Through the IBM CSI (Container Storage Interface) drivers for Block, FlashSystems or DS8000 can become the default choice of container-ready storage!

This block storage can be provided over FC, iSCSI, NVMe/TCP

Container Native Storage

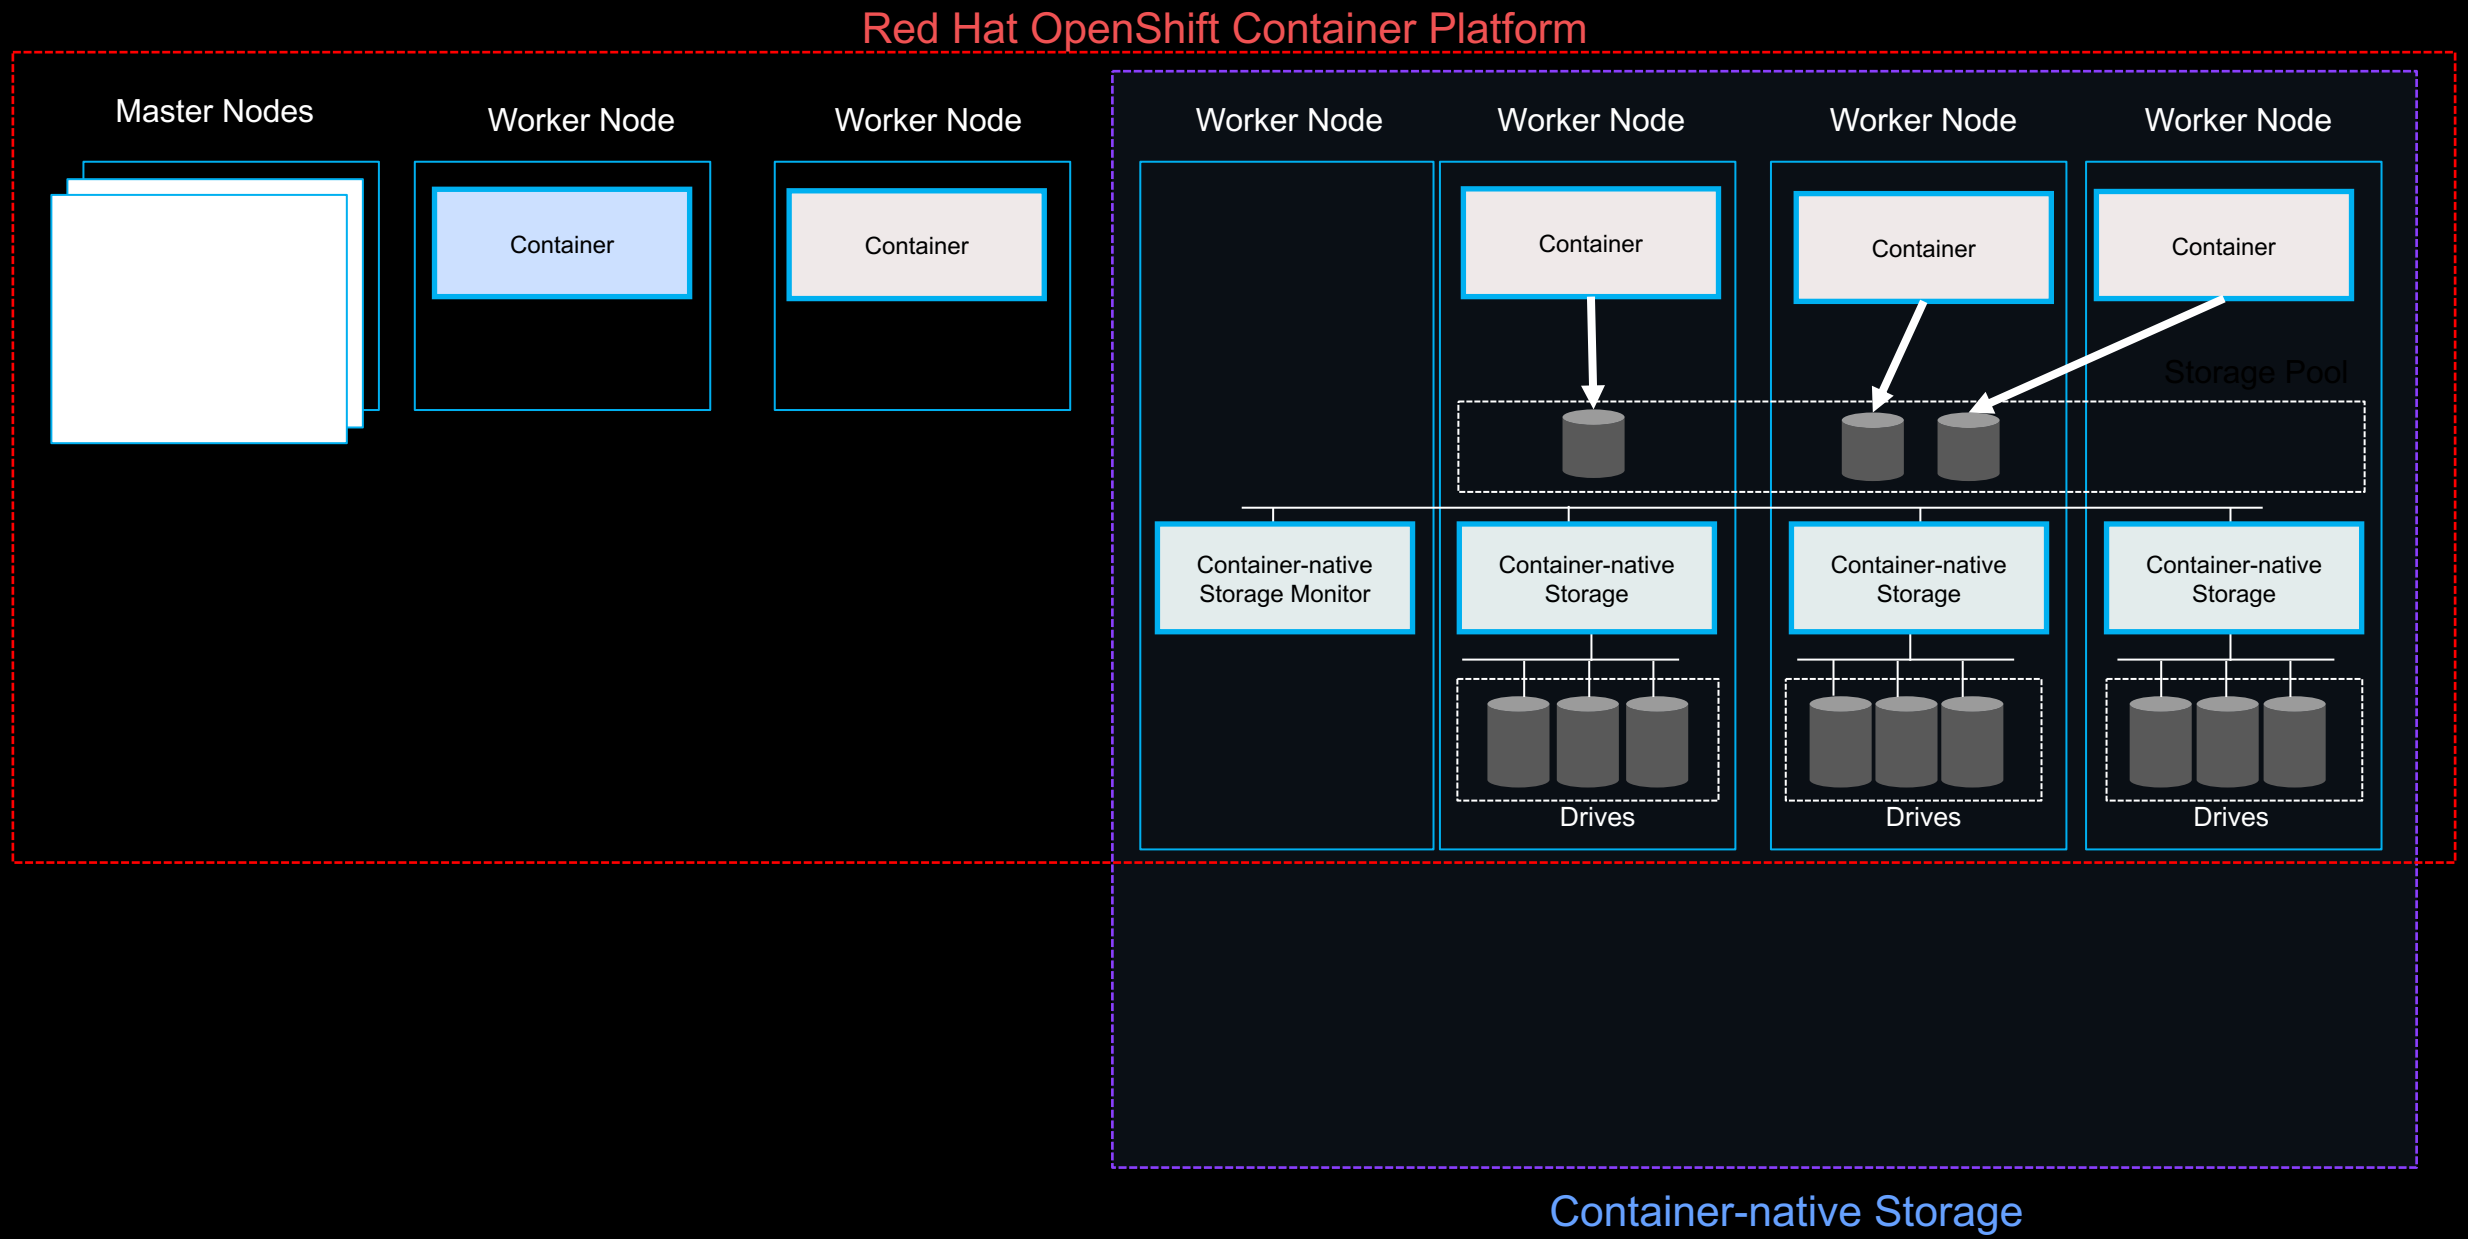

Container-ready and Container-native Storage

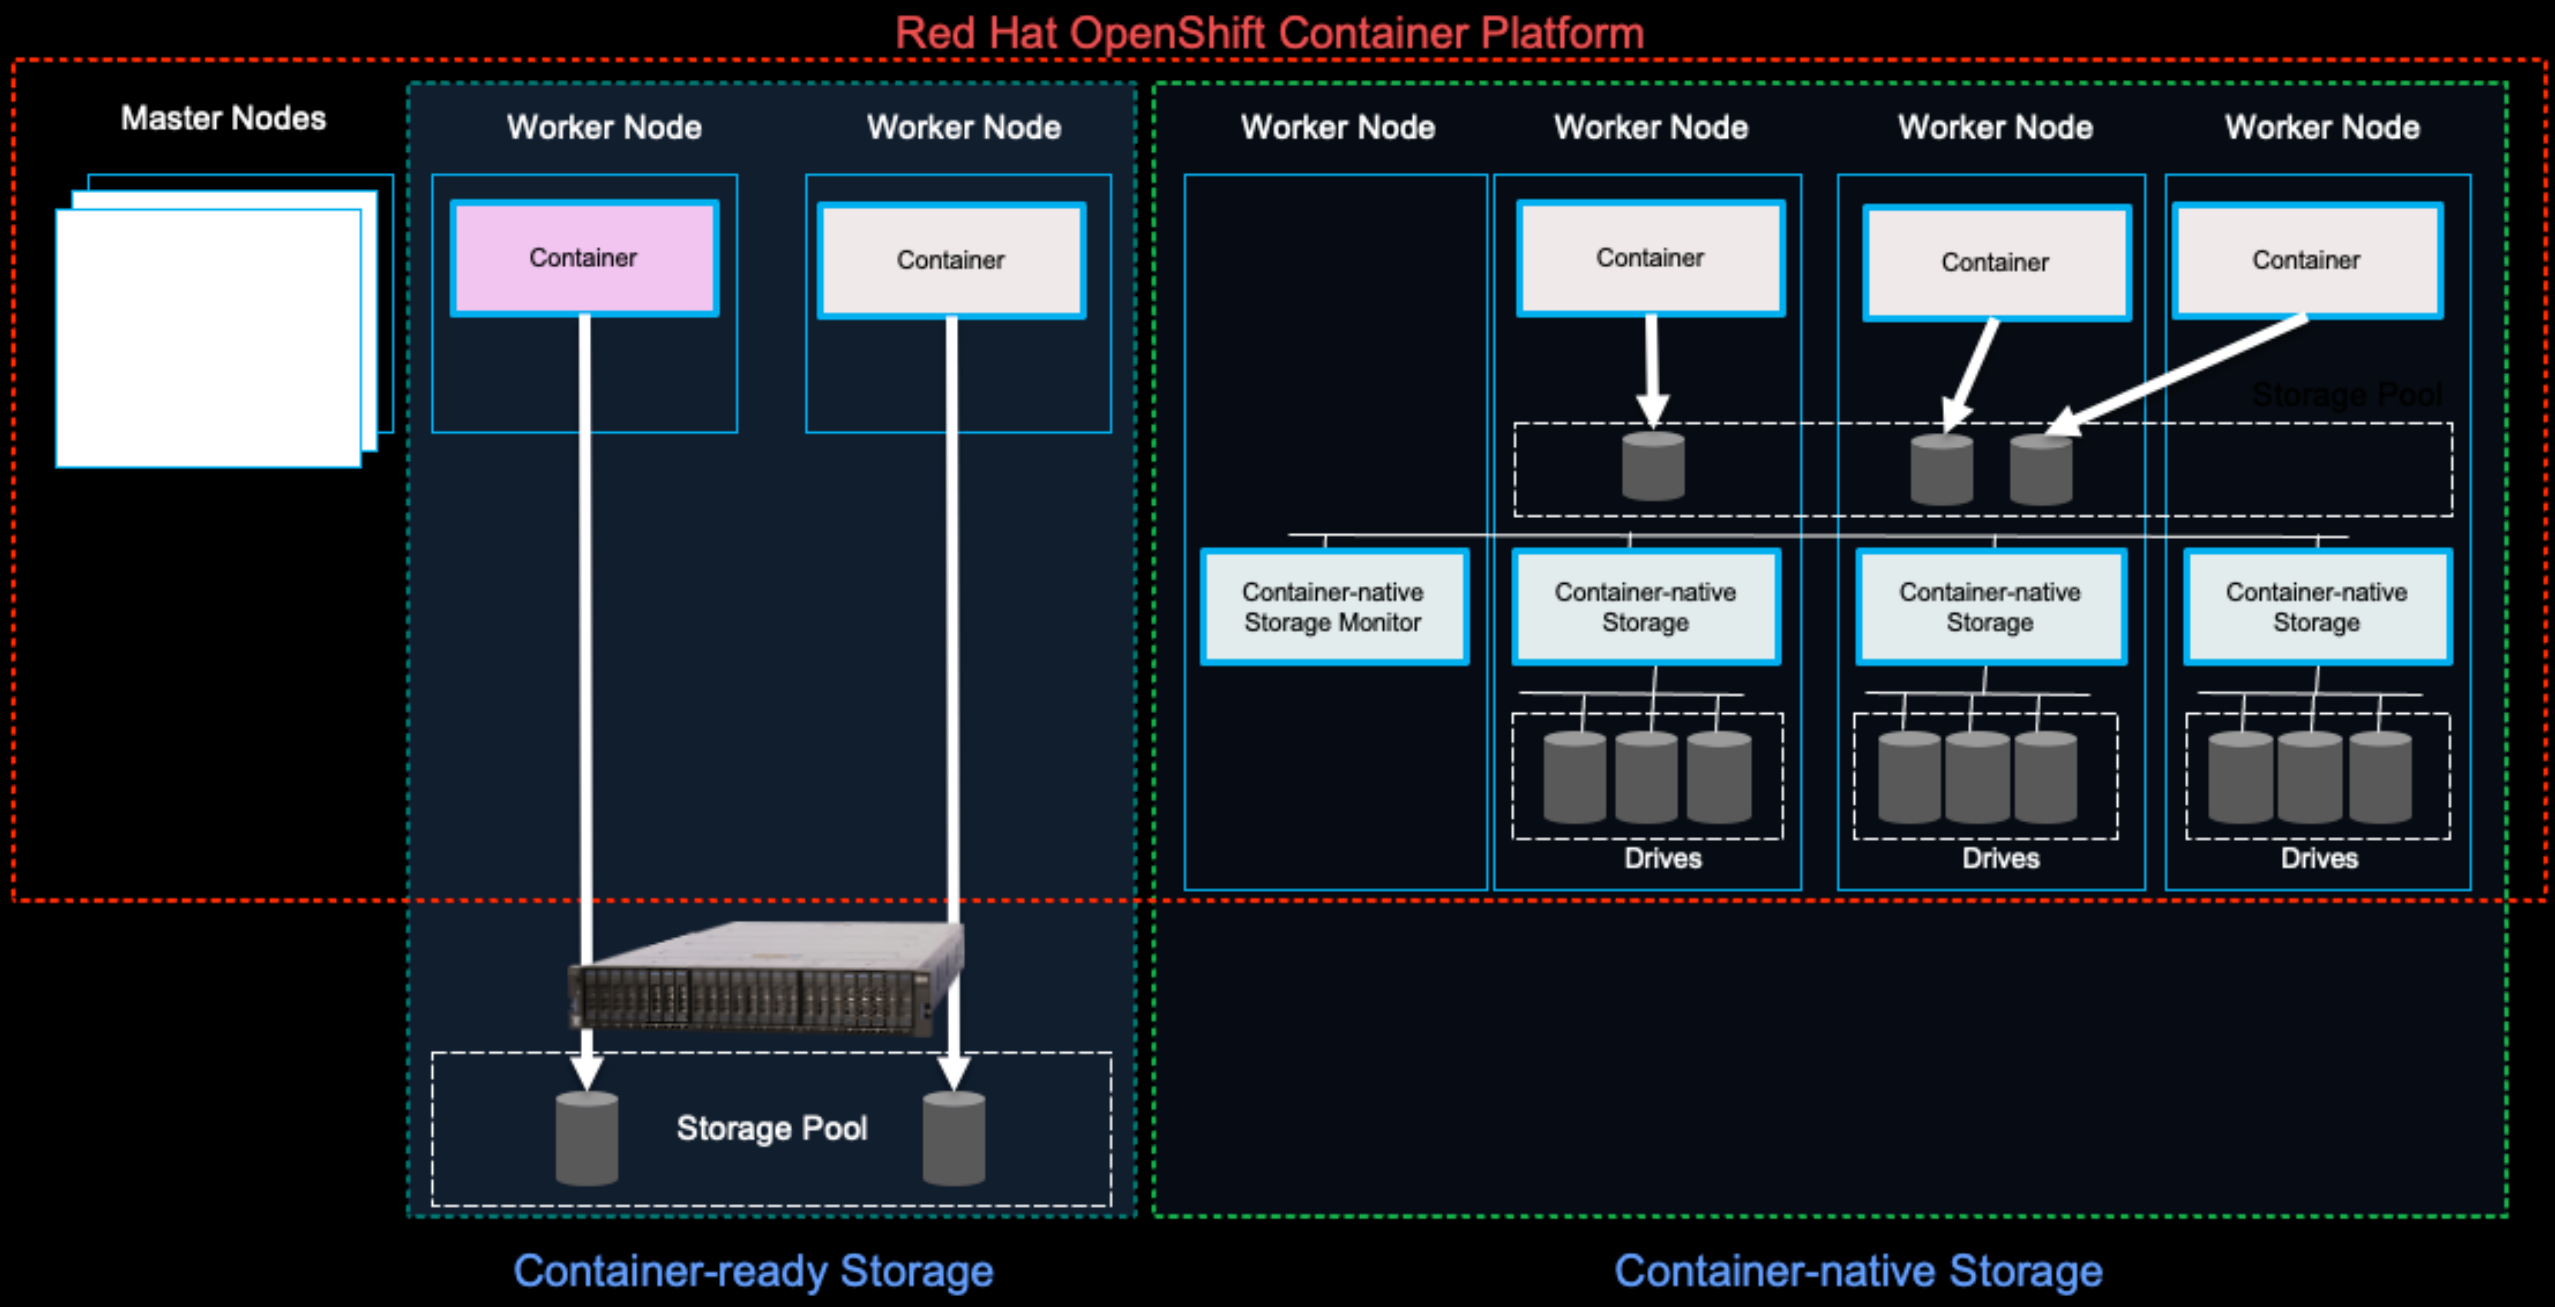

Goal: Enable OpenShift environments to UNIFY operations for containers and VMs

Consistent management experience

Red Hat OpenShift Virtualization

- Create and manage Linux and Windows VMs
- Import and clone existing VMs
- Manage network and storage attached to virtual machines
- Migrate VMs between nodes
- Import VMs from VMWare using Red Hat MTV toolkit

OPENSHIFT

IBM Fusion

- Shared persistence
- Simplify workload co-location
- Backup and restore for VMs
- Disaster recovery for storage
- Bare metal performance on HCI
- One IBM support approach

IBM Fusion

Why OpenShift Virtualization for VMs on IBM Fusion vs VMWare

Terminology Comparison

Feature	OpenShift Virtualization	vSphere
Where VM disks are stored	PVC / PV	Datastore
Policy based storage	StorageClass	Storage Policy Based Management (SPBM)
Non-disruptive VM migration	Live migration	vMotion
Non-disruptive VM storage migration	N/A	Storage vMotion
Active resource balancing	Pod eviction policy, descheduler	Dynamic Resource Scheduling (DRS)
Physical network configuration	NMstate Operator, Multus	vSwitch / DvSwitch
Overlay network configuration	OVN-Kubernetes, CNI partners, Multus	NSX-T
Host / VM metrics	OpenShift Metrics and Monitoring	vCenter, vRealize Operations

Red Hat

The Fusion 5

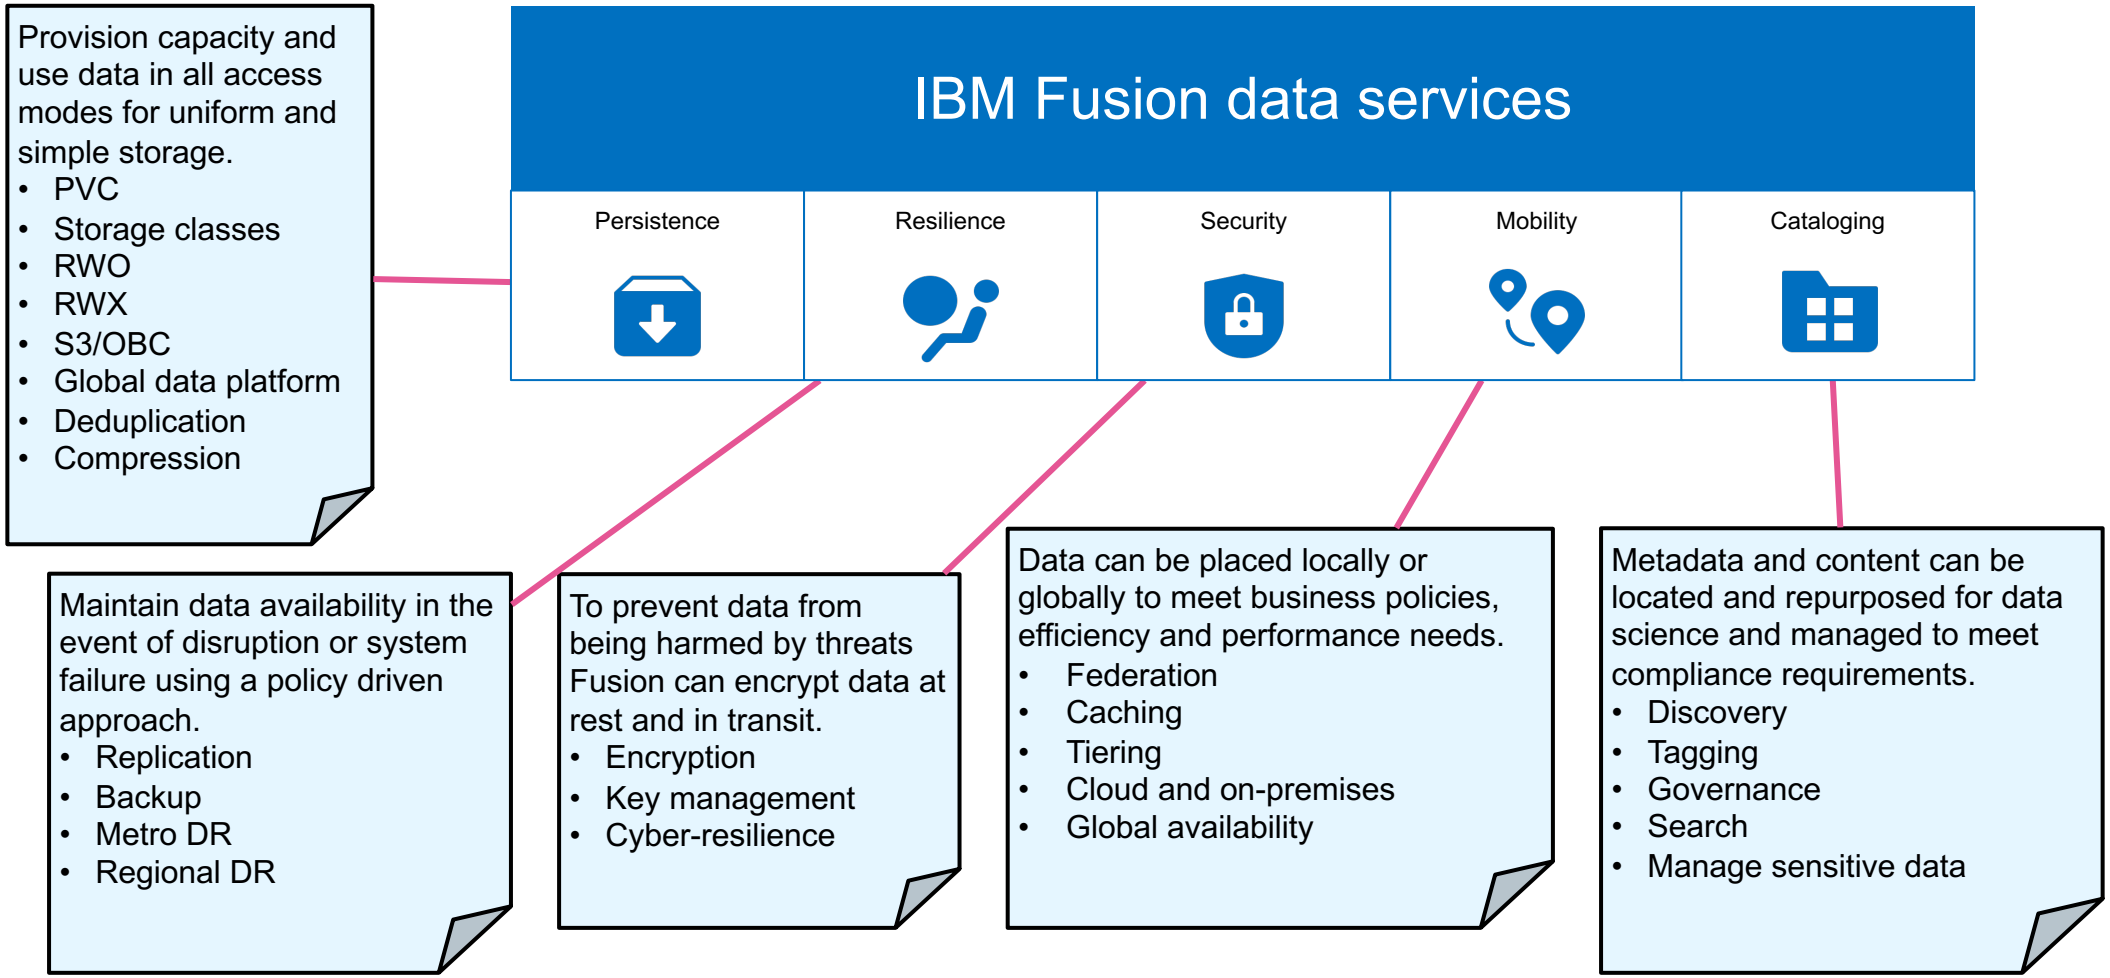

Containers are not virtual machines

Containers are not virtual machines

- Containers are process isolation
- Kernel namespaces provide isolation and cgroups provide resource controls
- No hypervisor needed for containers
- Contain only binaries, libraries, and tools which are needed by the application
- Ephemeral

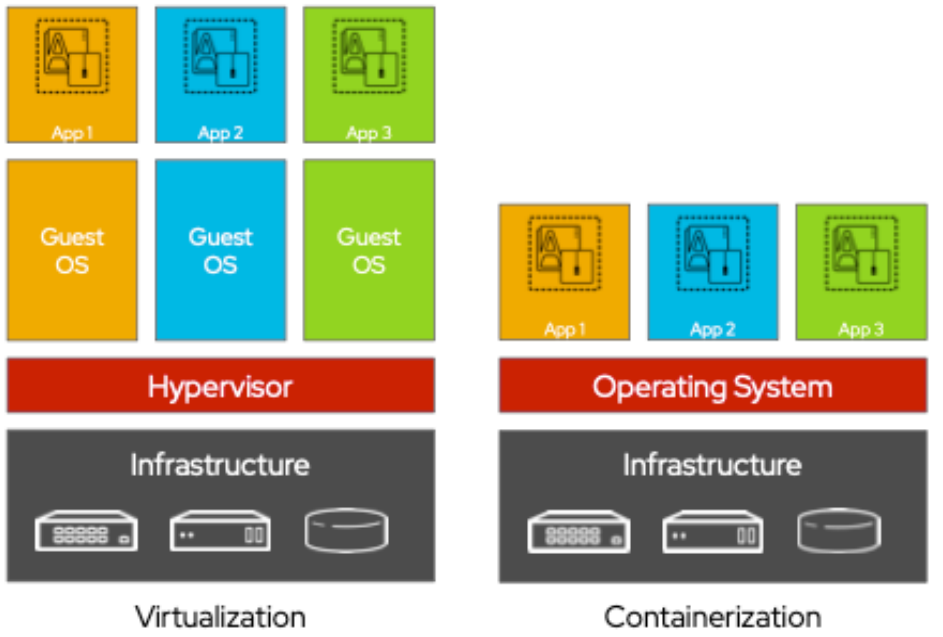

Red Hat

Virtual machines can be put into containers

Virtual machines can be put into containers

- A KVM virtual machine is a process
- Containers encapsulate processes
- Both have the same underlying resource needs:
  - Compute
  - Network
  - (sometimes) Storage

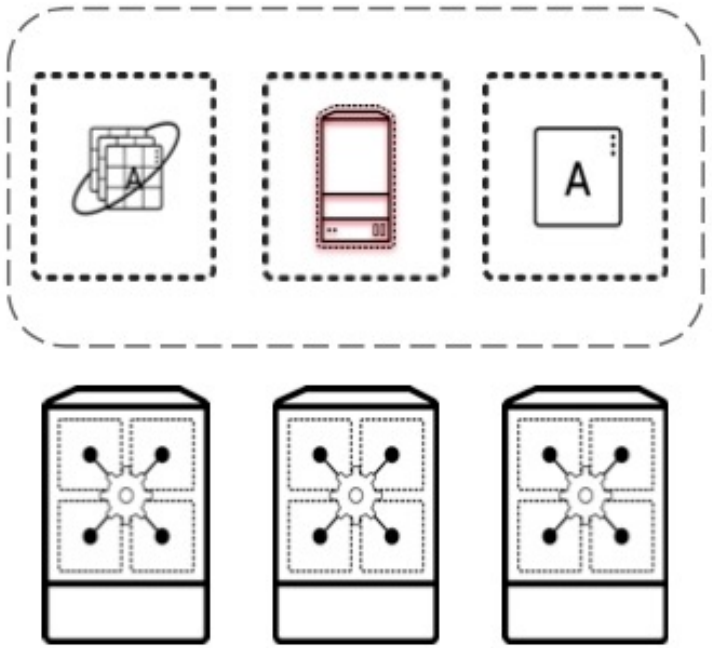

Red Hat

OpenShift Virtualization

OpenShift Virtualization

- Virtual machines
  - Running in containers, managed as Pods
  - Using the KVM hypervisor
- Scheduled, deployed, and managed by Kubernetes
- Integrated with container orchestrator resources and services
  - Traditional Pod-like SDN connectivity and/or connectivity to external VLAN and other networks via multus
  - Persistent storage paradigm (PVC, PV, StorageClass)

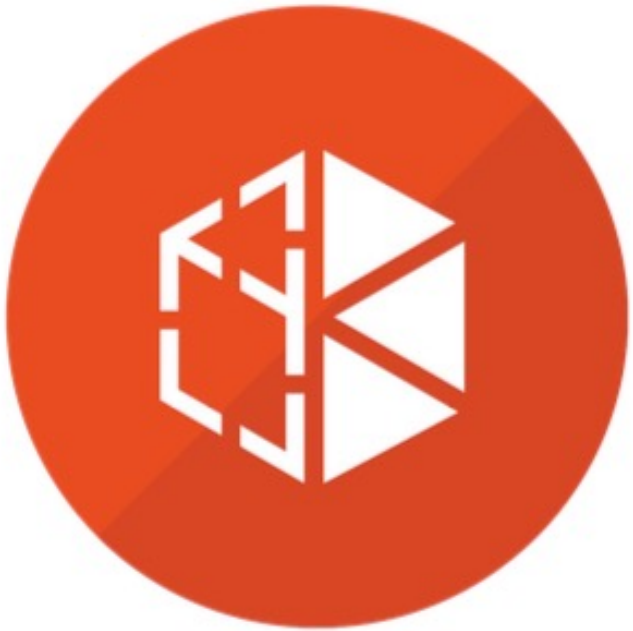

Red Hat

OpenShift Virtualization uses KVM

OpenShift Virtualization uses KVM

- OpenShift Virtualization uses KVM, the Linux kernel hypervisor
- KVM is a core component of the Red Hat Enterprise Linux kernel
  - KVM has 10+ years of production use: Red Hat Virtualization, Red Hat OpenStack Platform, and RHEL all leverage KVM, QEMU, and libvirt
- QEMU uses KVM to execute virtual machines
- `libvirt` provides a management abstraction layer
- Currently supported on x86 bare metal
- For other platforms contact Product Management for roadmap

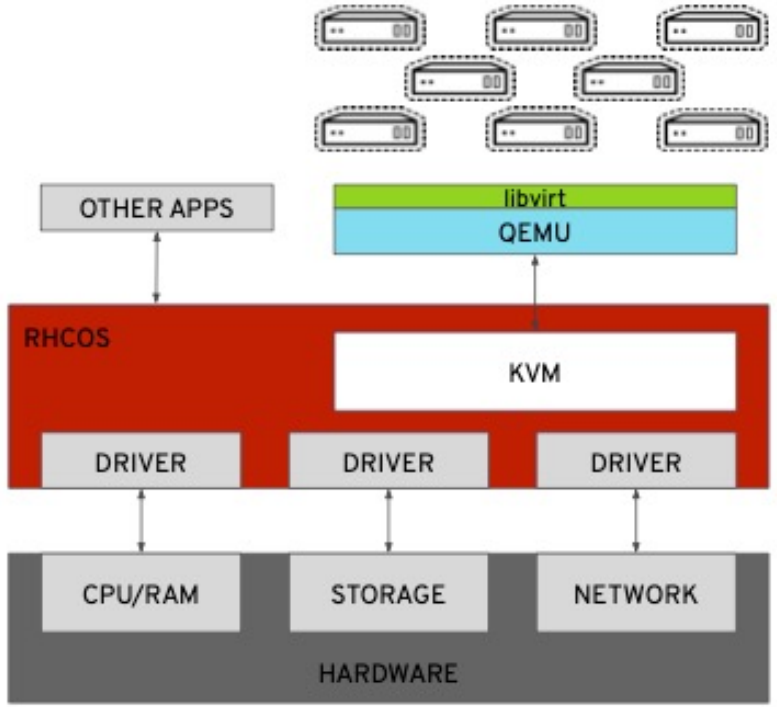

Red Hat

Virtual Machine Management

Virtual Machine Management

- Create, modify, and destroy virtual machines, and their resources, using the OpenShift web interface or CLI
- Use the `virtctl` command to simplify virtual machine interaction from the CLI

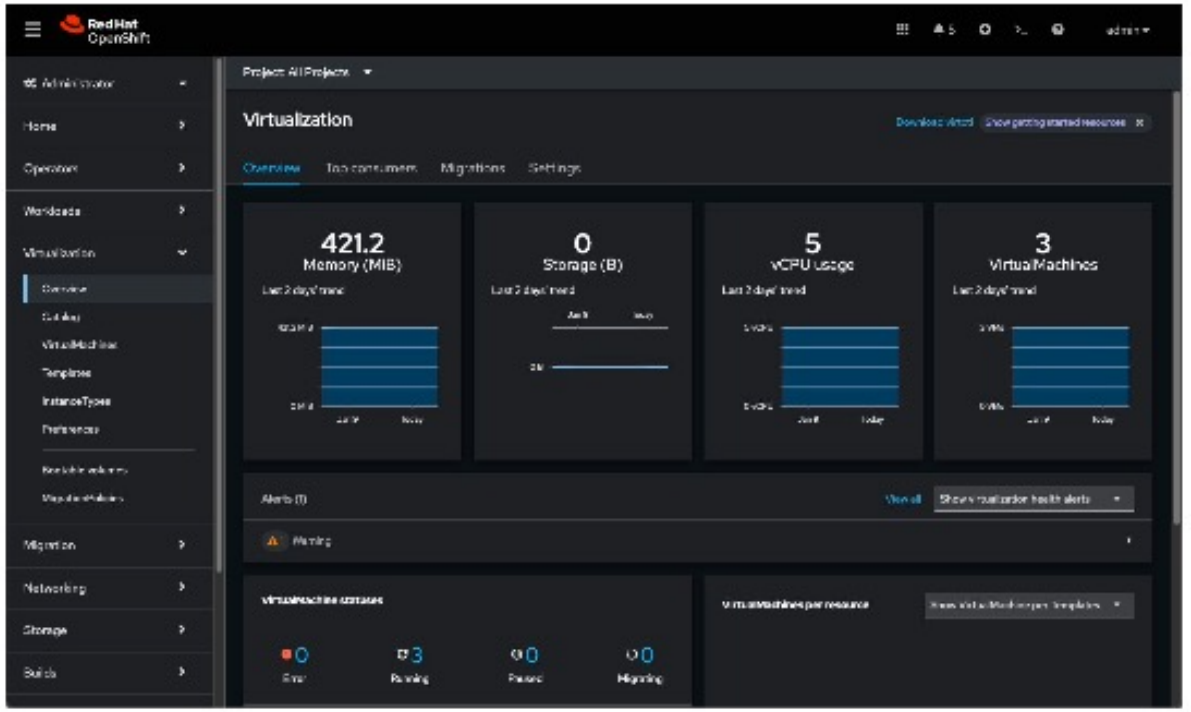

Red Hat

IBM Fusion, two offerings

Customer Apps

IBM Cloud Paks

- Data
- Integration
- Security
- Cloud Satellite
- Business automation
- Watson AIOps
- Network automation

Databases

- Cassandra
- MongoDB
- RabbitMQ
- Elastic search
- PostgreSQL
- Spark

Off the shelf

- Pega
- Mulesoft
- TIBCO

Custom apps

- Home grown

Offering

Fusion software

Data services for stateful OpenShift applications

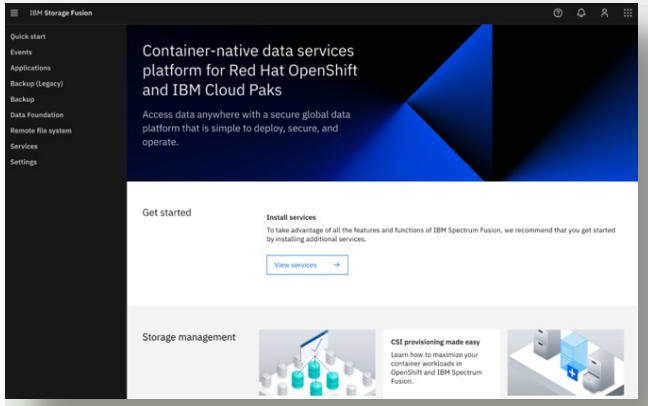

Consistent experiences

- APIs
- Fusion console
- Data protection
- Disaster Recovery
- Fusion Data Foundation (FDF)

Fusion HCI System

Integrated Application Platform for OpenShift

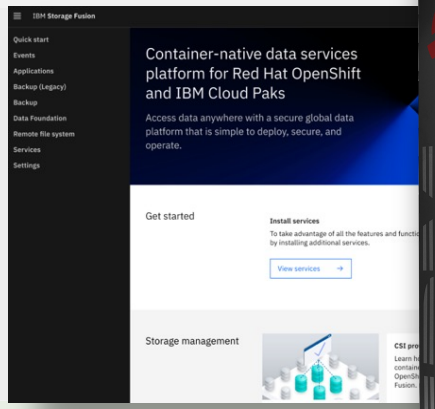

Built on

OPENSHIFT

Deployments

- AWS EBS
- Azure block
- Persistent disk
- IBM block
- SAN, vSAN
- DS8K
- FlashSystems
- Bare metal
- Server-attached drives

Hyper-Converged Infrastructure for OpenShift

Switches, servers, storage

Elements of IBM Storage Fusion HCI System

GPU nodes
3x NVIDIA A100 80GB (current)
8x NVIDIA L40S (future)

AFM nodes
Used for storage acceleration of object access

High speed switches
100 GbE
Dedicated storage network

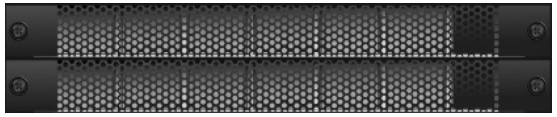

Management switches
1 GbE
Used for appliance management/monitoring

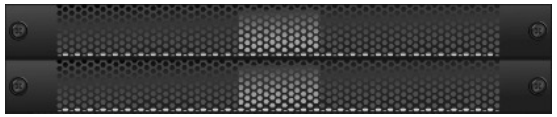

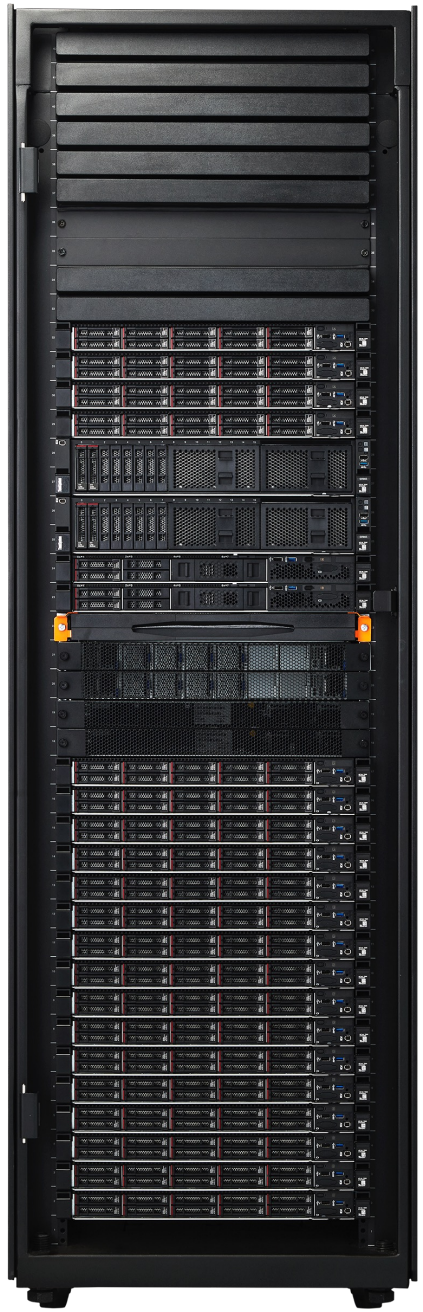

Why Fusion HCI for OpenShift Virtualization

Compute/storage nodes
32 or 64 core
256/512/1024/2048 GB memory
Provide compute for workloads
Storage for persistent volumes
Tier 1 cache for storage acceleration

OpenShift Virtualization requires **Bare-Metal hosts**

New Gen 2 Nodes

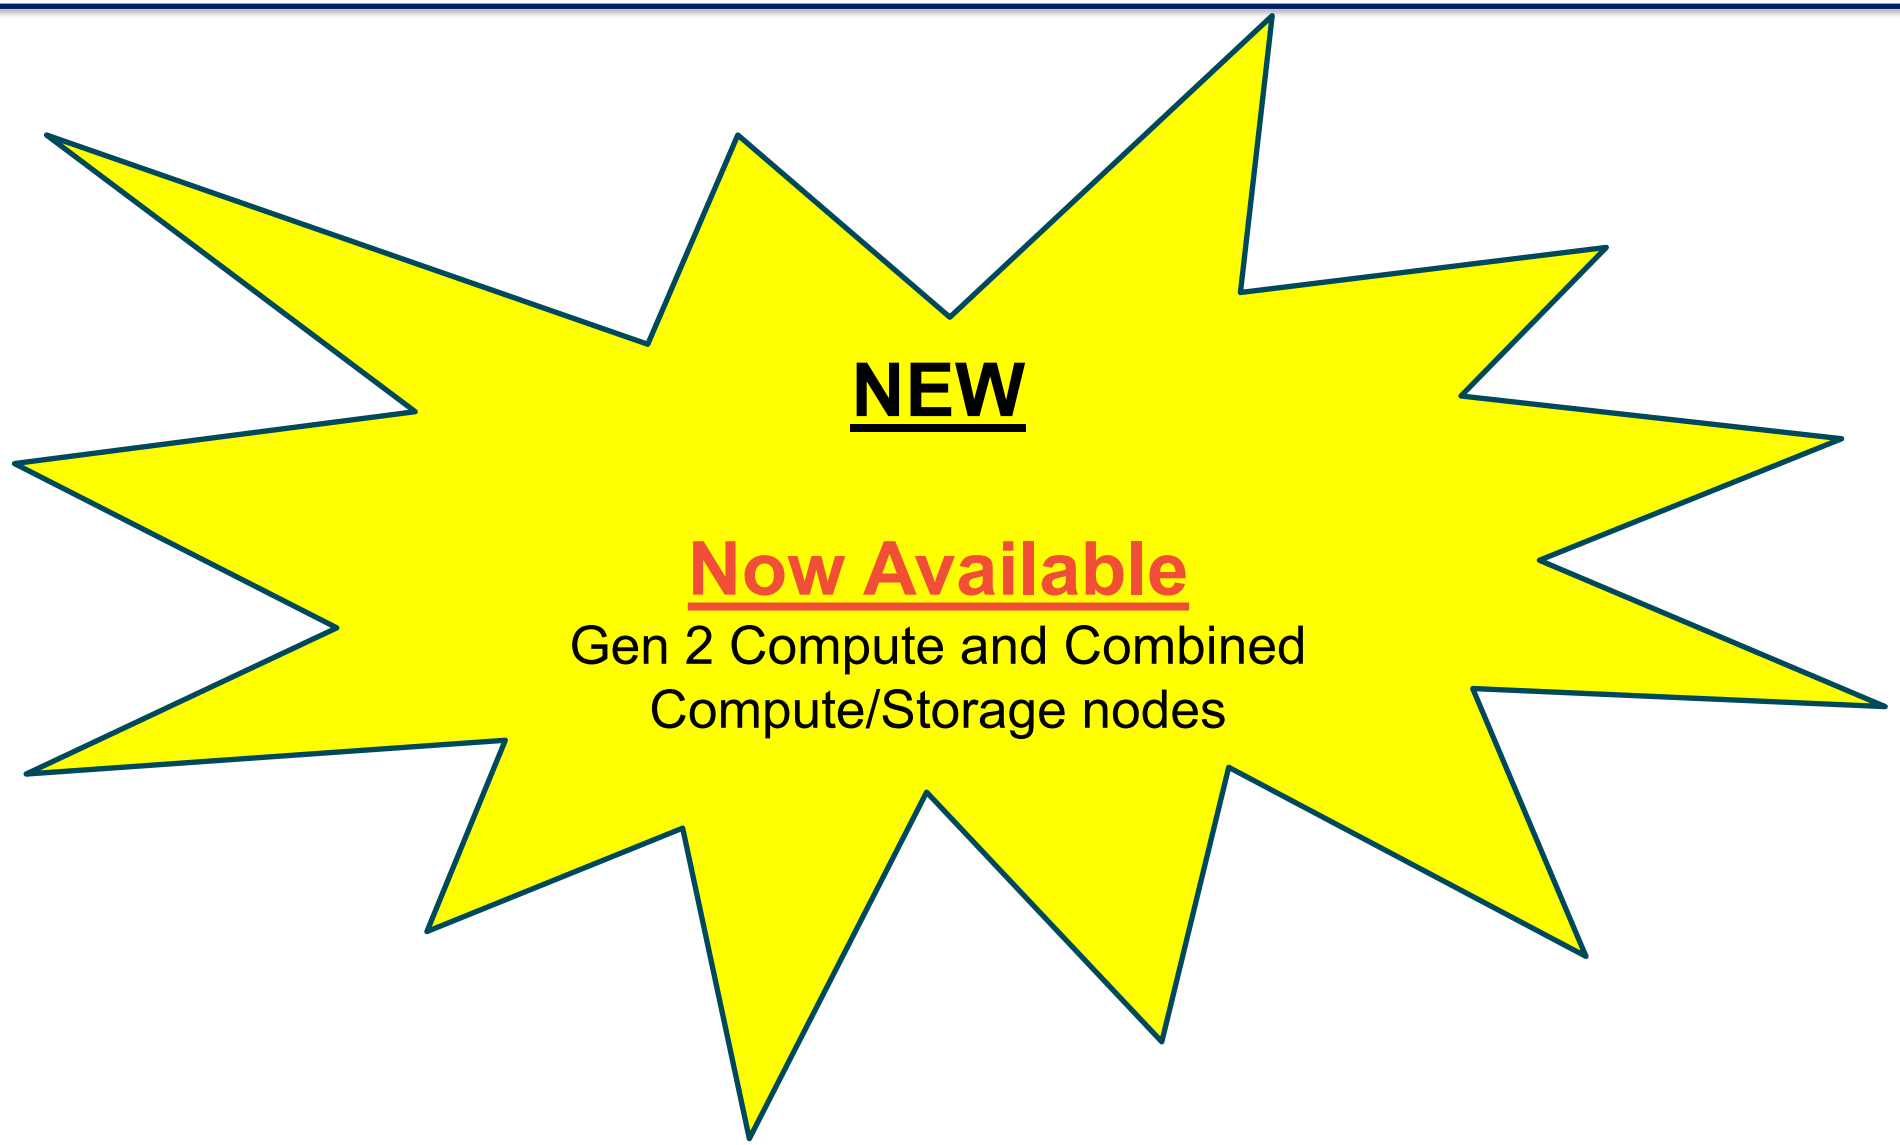

Two new servers with two drive size options

9155-C10 – Fusion HCI 32 core dual-socket server

- Lenovo SR630 V3 PCIe 5 planer
- 2x Intel Gold 6426Y 16C 2.5GHz 185W CPU (“Sapphire Rapids”)
- Hyper-Thread SMT=2
- 256 GB DDR5 memory, upgradeable to 512 GB
- 2x 960GB M.2 OS drives (RAID 1)
- Drive options: 3.84 or 7.68 TB - configure from 0 to 10 drives (1)

The 9155-C10 replaces the following Generation 1 MTMs

- 9155-C00 – 32 core compute server
- 9155-C01 – 32 core compute/storage server
- 9155-F01 – AFM server

9155-C14 – Fusion HCI 64 core dual-socket server

- Lenovo SR630 V3 PCIe 5 planer
- 2x Intel Gold 6438N 32C 2.0GHz 205W CPU (“Sapphire Rapids”)
- Hyper-thread SMT=2
- 1024 GB DDR5 memory, upgradable to 2048 GB
- 2x 960GB M.2 OS drives (RAID 1)
- Drive options: 3.84 or 7.68 TB – configure from 0 to 10 drives (1)

The 9155-C14 replaces the following Generation 1 MTMs

- 9155-C04 – 64 core compute server
- 9155-C05 – 64 core compute/storage server

(1) All storage servers in a cluster must be configured with the same size and number of drives

- **Base configuration (grey)**
- 2x Ethernet 100 GbE high-speed switches
- 2x Ethernet 1 GbE management switches
- 6x Compute with storage servers with 2 NVMe drives per server

- **Options (blue)**
- 42U rack
- Additional servers up to a max of 16
  - Compute with storage servers or compute only servers
  - 32 core or 64 core with and without storage
- NVMe drives
  - Add up to 8 additional NVMe drives
- GPU servers
- Rack-mounted KVM console (9155-TF5)

- **Multi-rack cluster**
- Connect-up to 3 racks to create large OpenShift clusters
- Spine switches for multi-rack configurations (9155-S01)

For Hardware Assisted High Availability

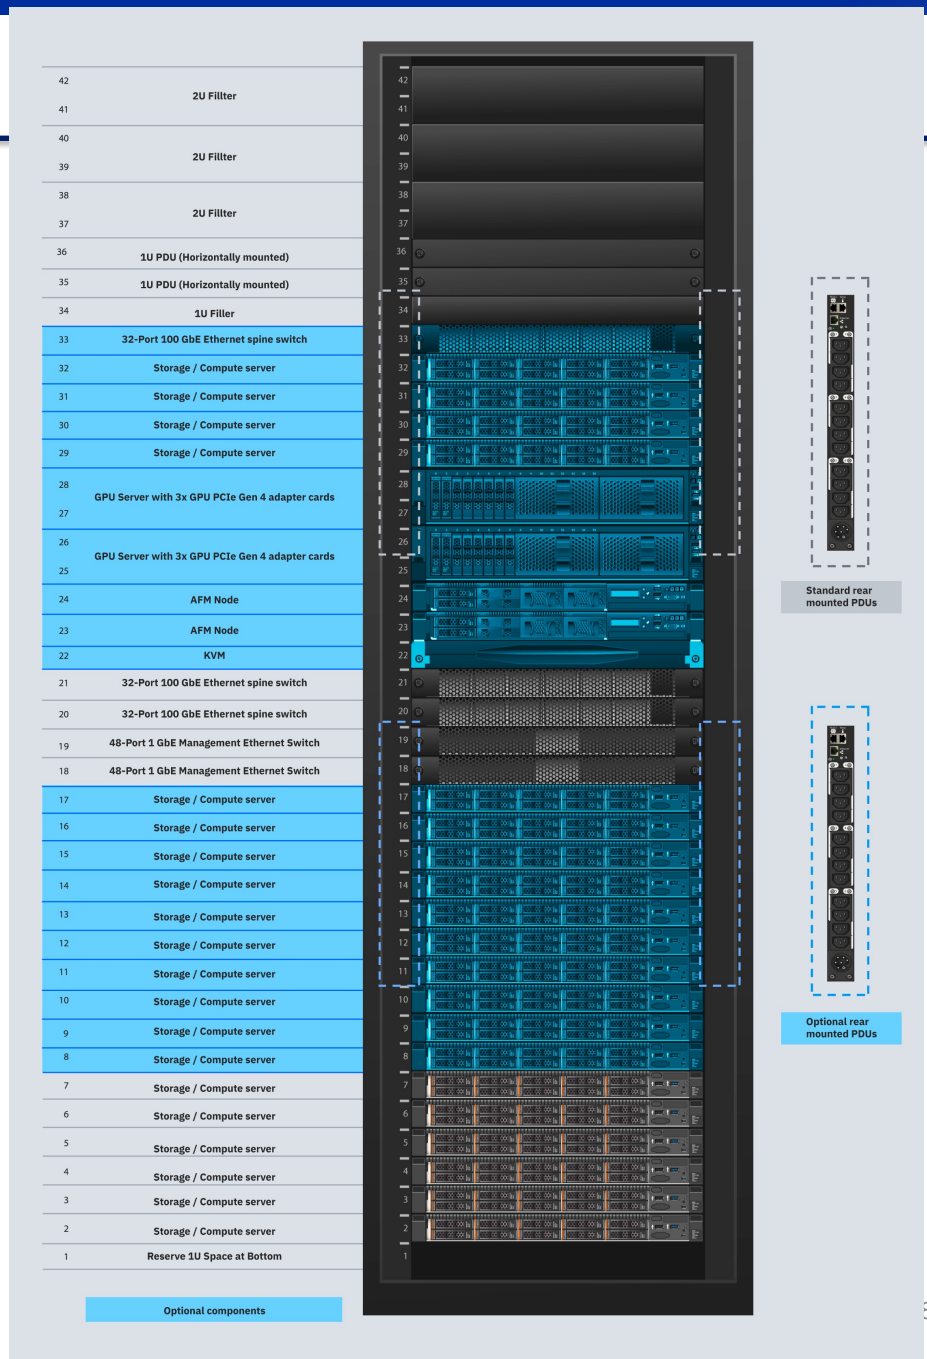

Flexible configuration options

6-node Fusion HCI 1 rack (min size)
- 55 TiB usable storage
- 96 cores (192 vCPU)

16-node Fusion HCI 1 rack (max size)
- Up to 716 TiB usable storage
- Up to 1024 cores (2048 vCPU)

48-node Fusion HCI 3 rack (max size)
- Up to 2148 TiB usable storage
- Up to 3072 cores (6144 vCPU)

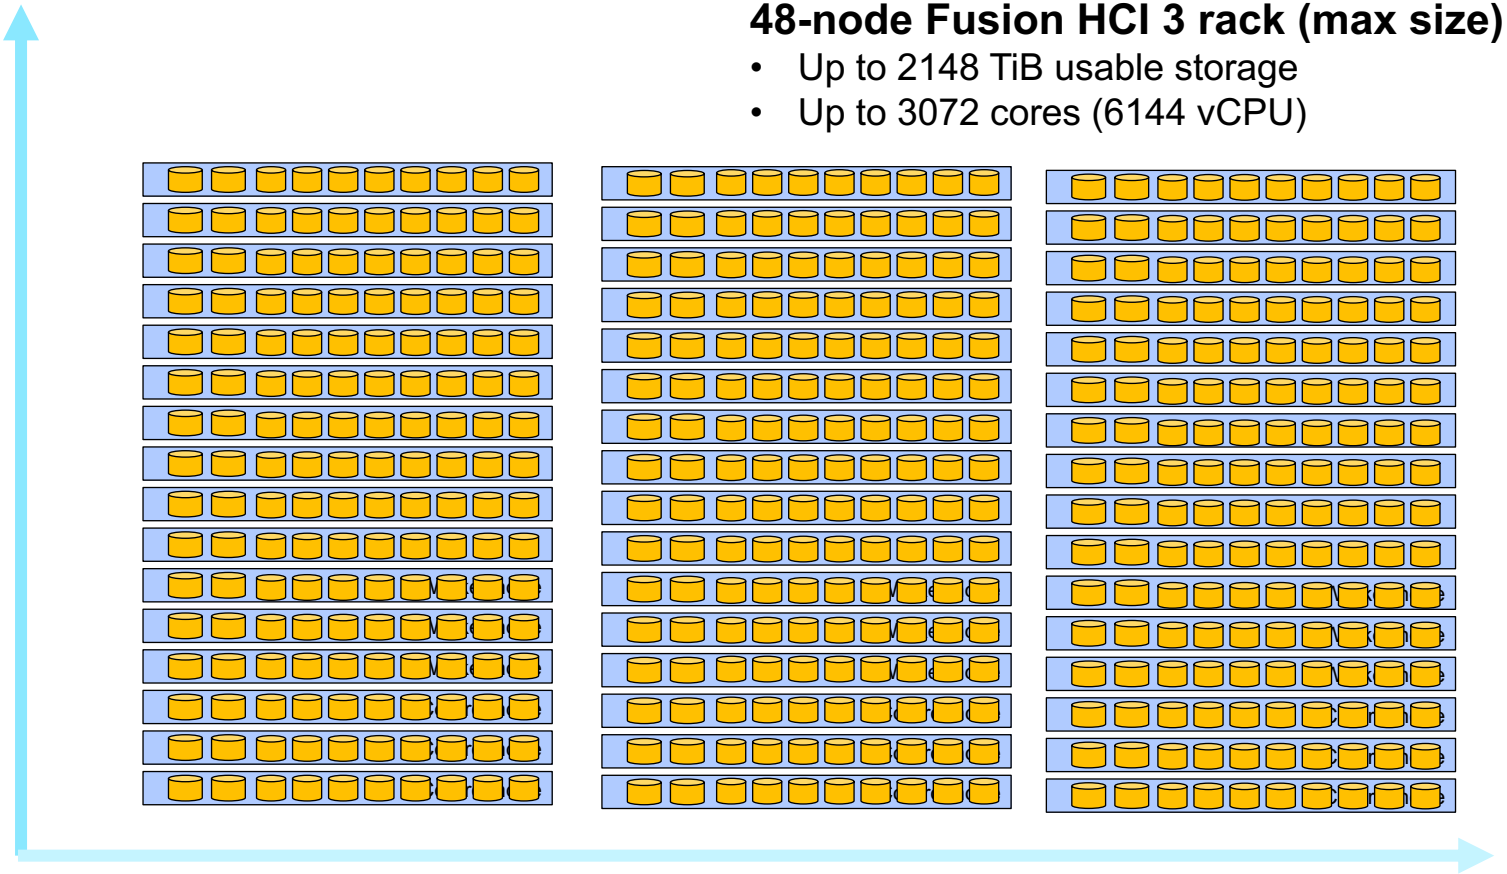

Persistence - Primary storage

Fusion can be deployed using one or both primary storage technologies. Fusion Data Foundation must be used when internal storage is requi d. It also has the capability to integrate with external IBM Storage Ceph.

IBM Storage Scale wi Storage Fusion software ed is available only in remote mount configuration.

More than one IBM Storage Fusion cluster can access the external IBM Storage Scale cluster which simplifies data management for those clients who already use it.

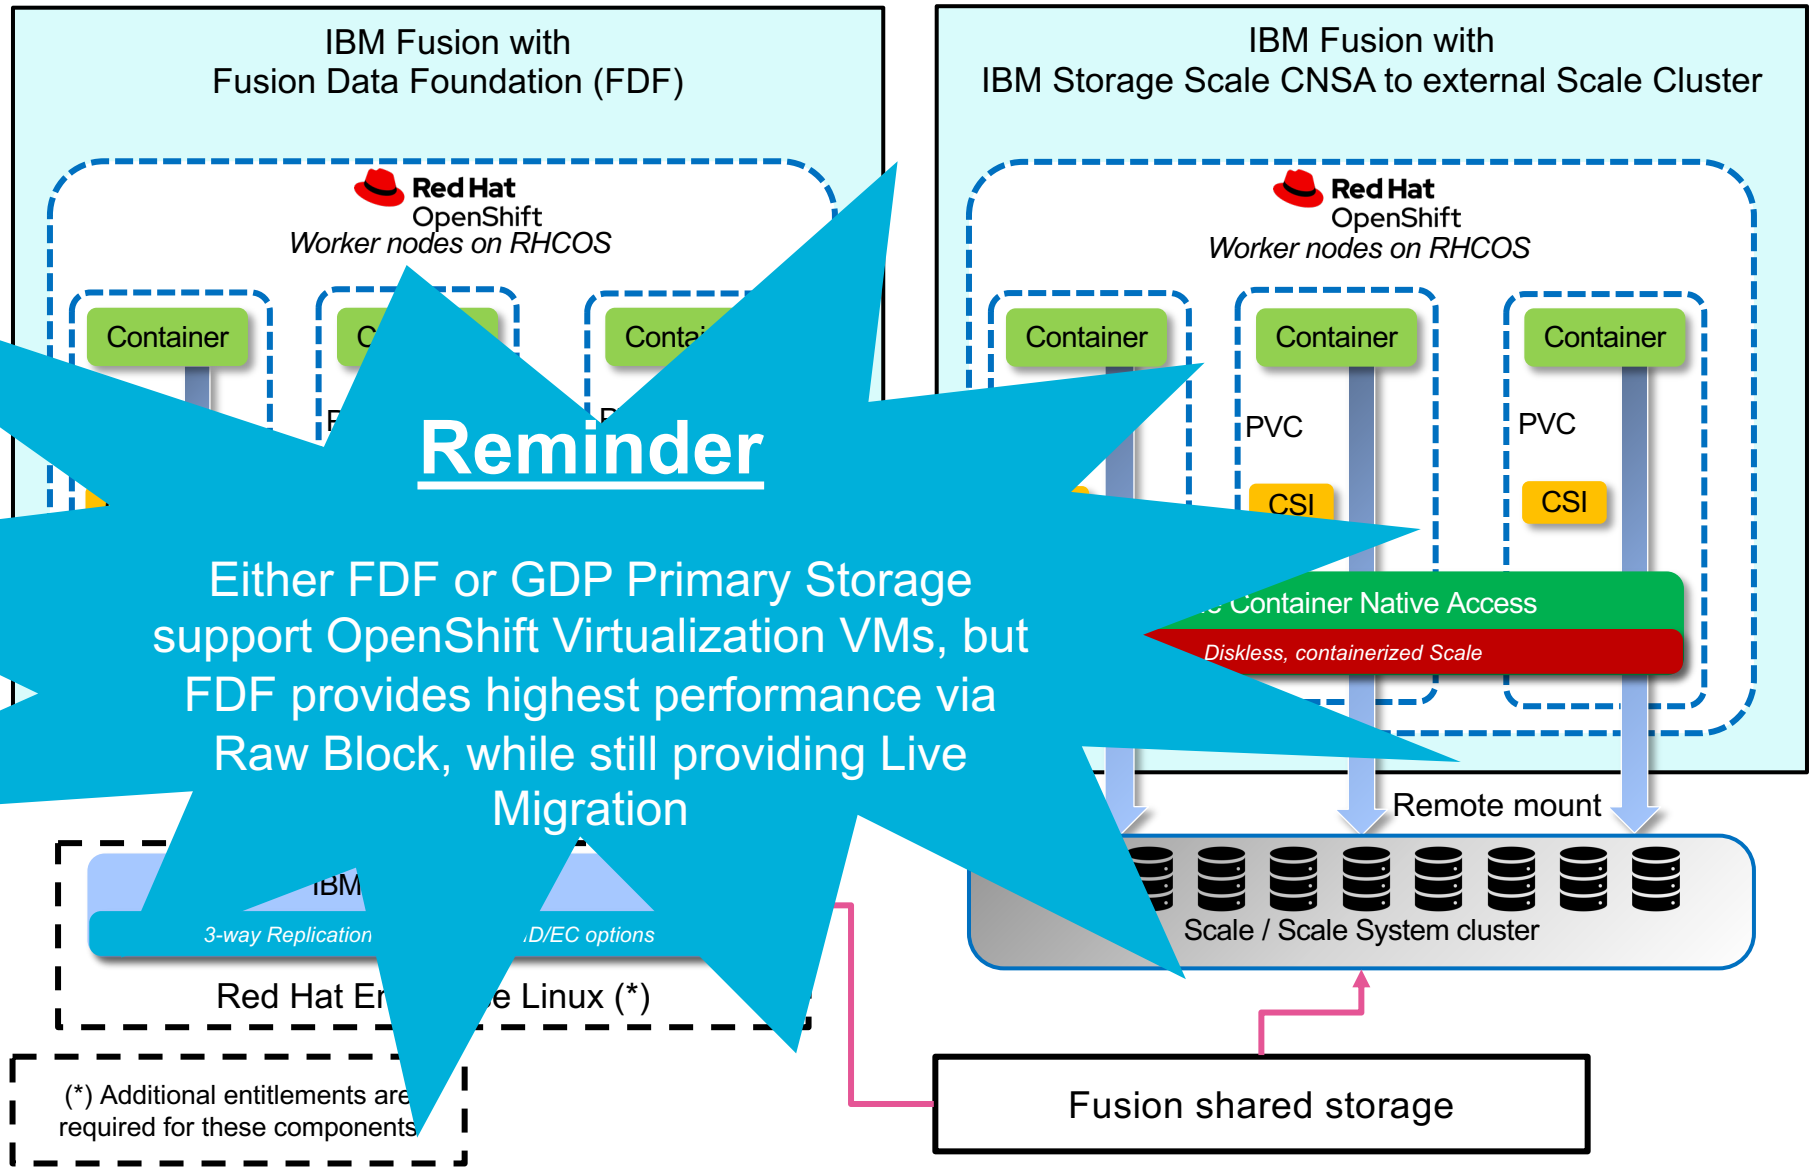

Resilience - Application aware backup and restore

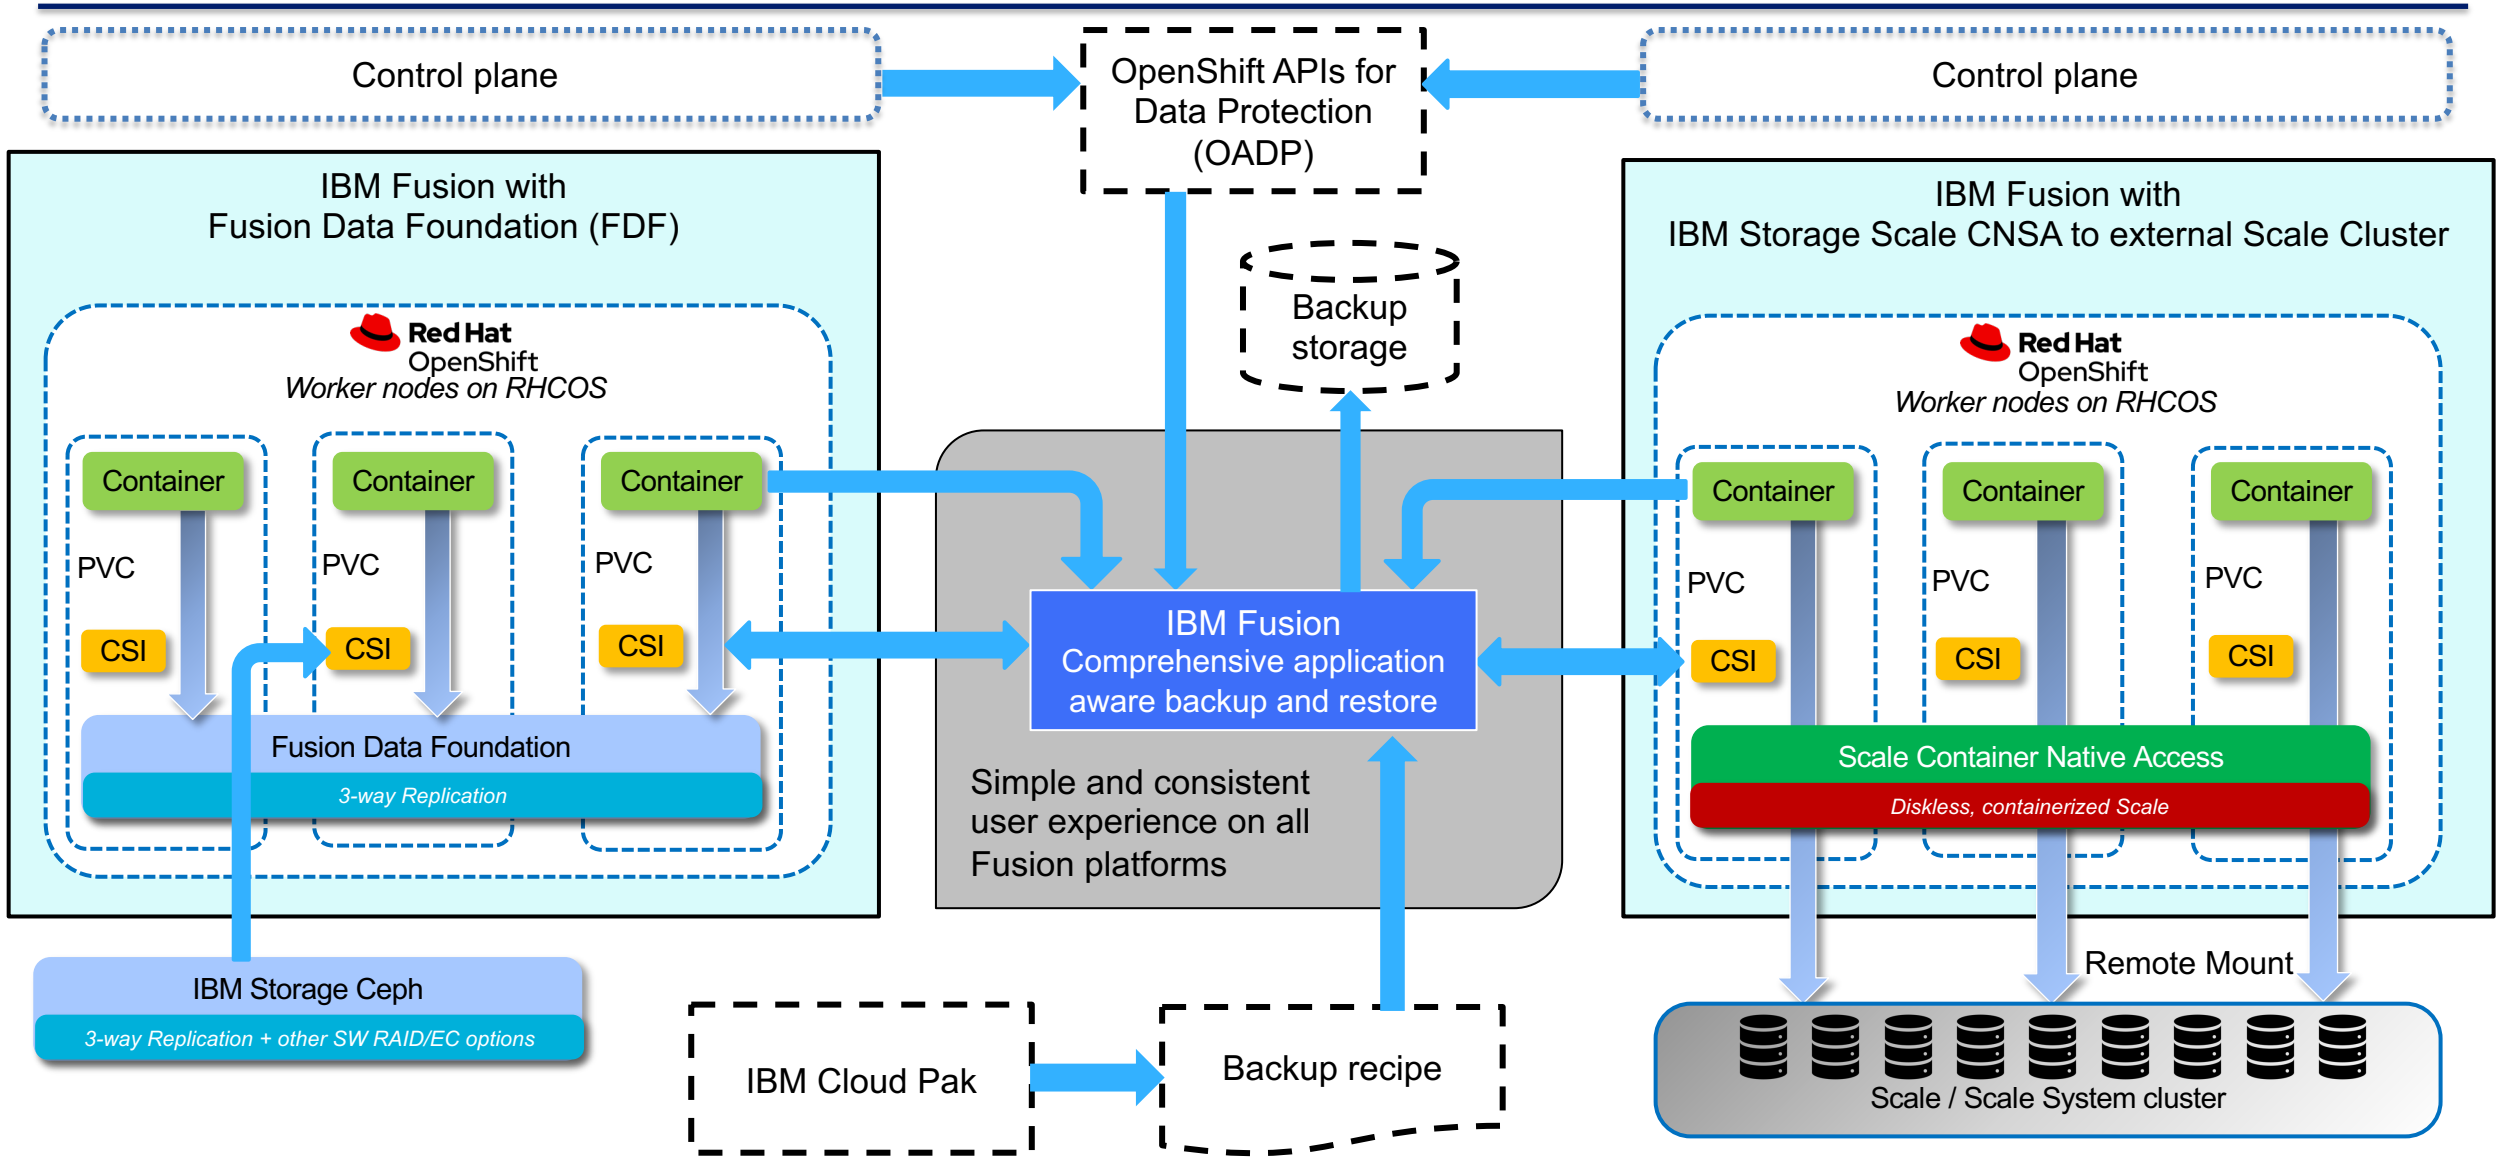

Resilience – Metro DR

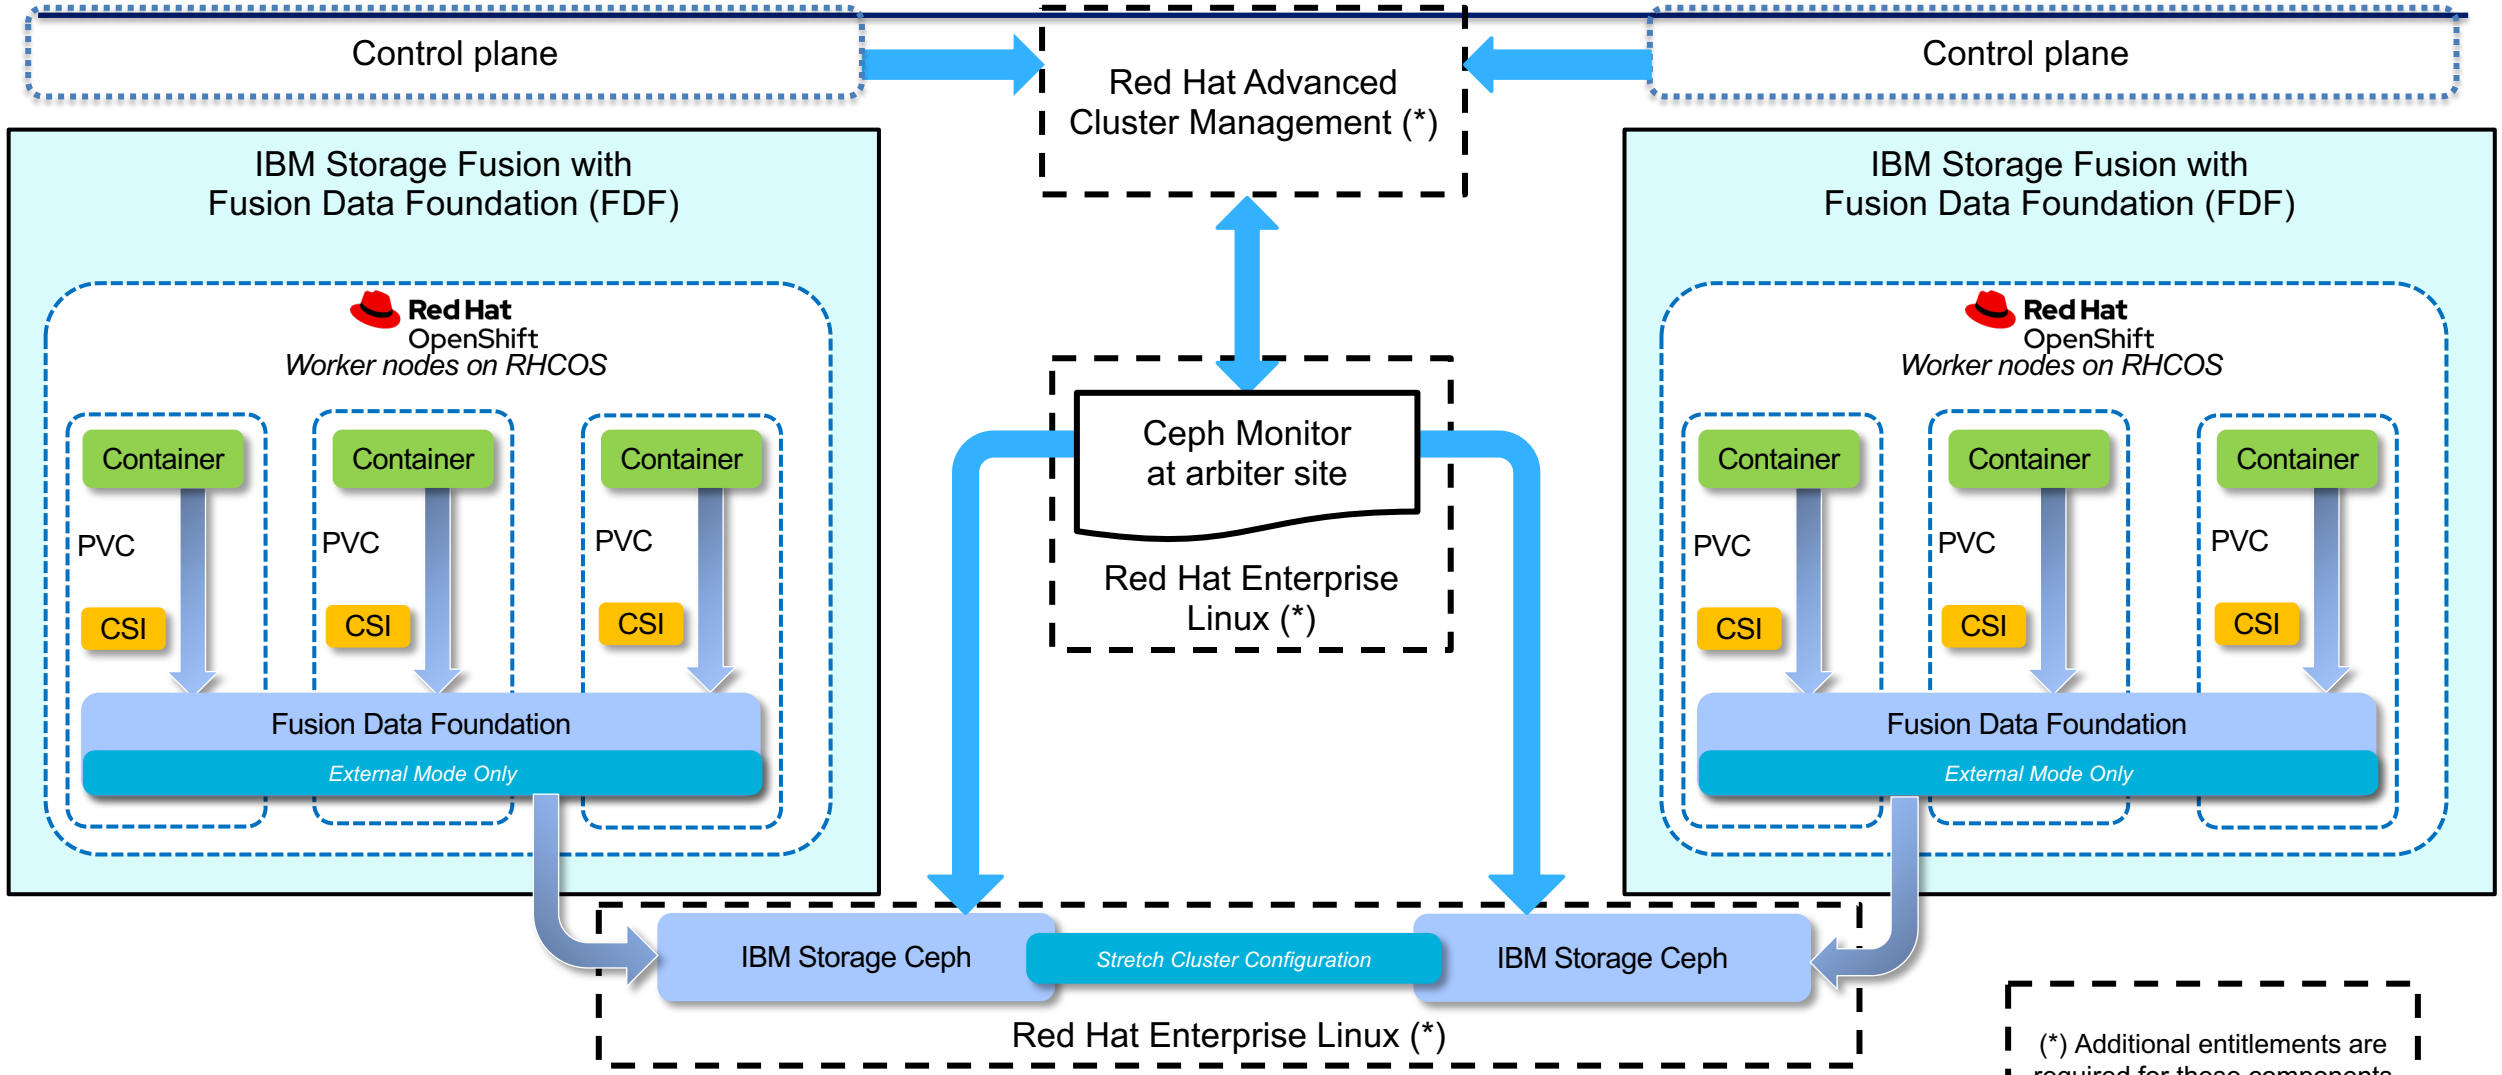

Use case – Recovery Point Objective (RPO) is close to zero

Resilience – Regional DR

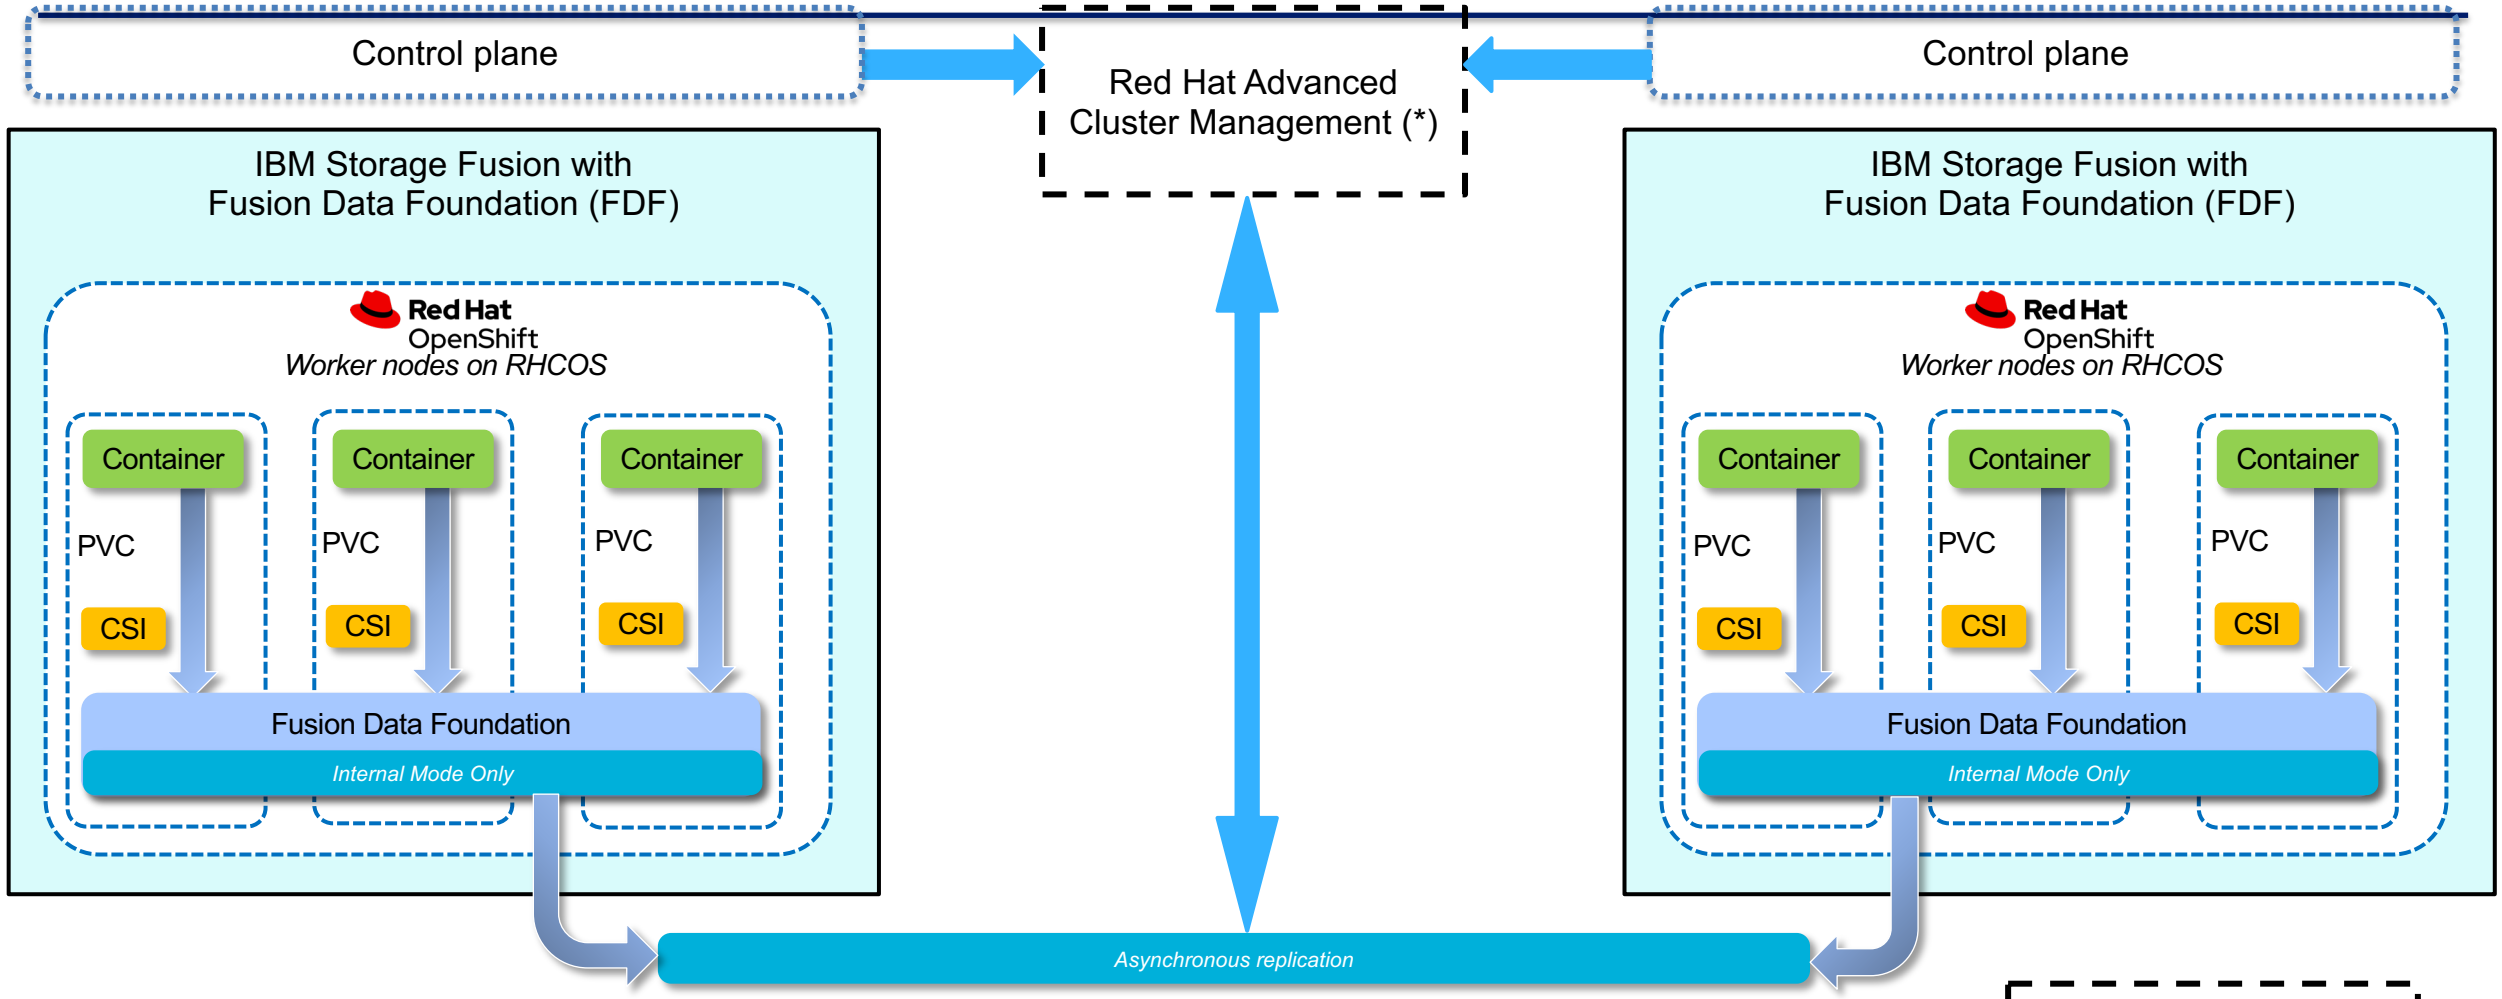

Use case – Recovery Point Objective (RPO) is minutes or hours

Next Steps – Additional Assistance from Red Hat

Red Hat Offerings to assist

- Virtualization Migration Workshop
- VirtualizationTraining andCertificationBundle
- OpenShiftVirtualizationProof of Value
- OpenShiftVirtualizationMigrationFactory

For any or all of these Red Hat offerings reach out to your Red Hat Account Manager

Demo Time

Demo Time

Accelerate with ATG Survey

Please take a moment to share your feedback with our team!

You can access this 6-question survey via Menti.com with code 1708 6924 or

Direct link https://www.menti.com/alwhyze7z1gz

Or

QR Code