



# OpenShift aaS on prem: Hosted Control Planes and Bare-Metal-Deployments

Daniel Danner  
Lead Architect IBM Fusion (D/A/CH)  
IBM



Proud sponsor of

Red Hat  
**Summit**


**Connect**

Darmstadt | 19. November 2024



# Agenda

- Bare Metal OpenShift
- Hosted Control Plane
- Storage and Backup for hosted Clusters
- Red Hat Advanced Cluster Manager



# The OpenShift Software Stack

## Bare Metal is Better

The Software Stack Architecture has a tremendous impact on

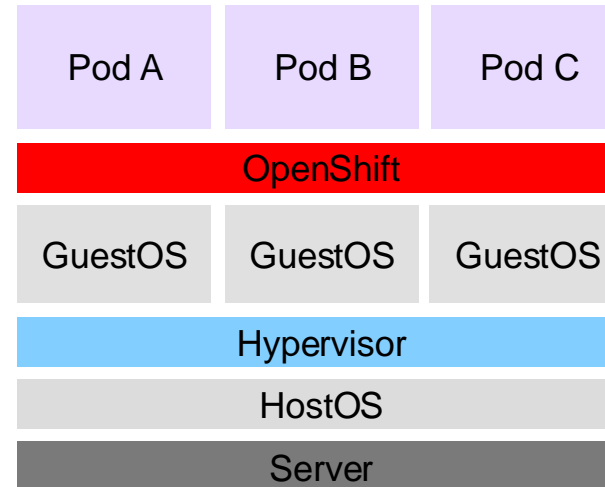
**Day2 Operations!**

### Bare Metal Advantages:

- Slimmer Software Stack
- Higher Performance
  - No IO penalty (VirtIO or SCSI)
  - NUMA memory mapping
- Lower Cost (TCO)
- Improved Security
- Reduced effort and complexity eases Day2 operations

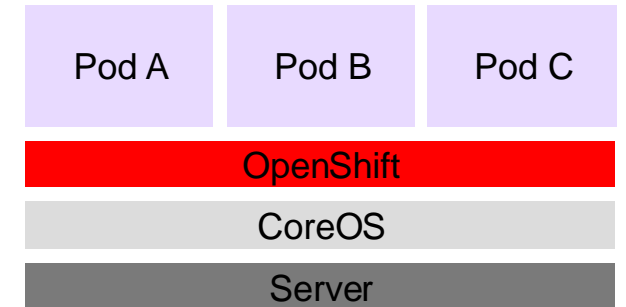


Hypervisor based  
Good



Virtualized Infrastructure  
with many layers

**Bare metal  
Better**



Bare metal  
with less layers

(1) VMs will add a performance penalty  
(2) OpenShift subscriptions licensing vs Socket based subscriptions

# The OpenShift Software Stack

## Bare Metal is Better

The Software Stack Architecture has a tremendous impact on

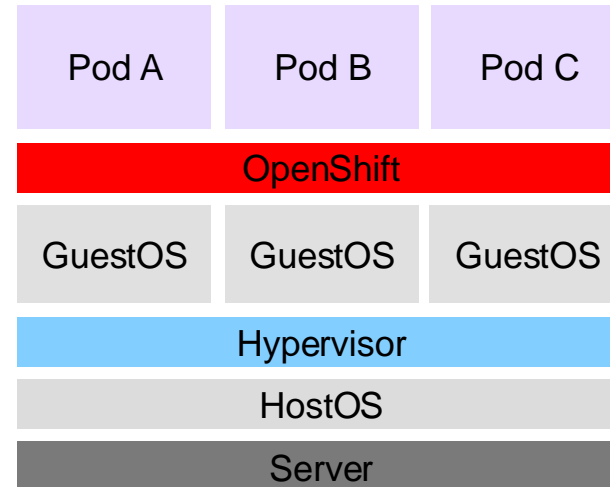
**Day2 Operations!**

### Bare Metal Challenges:

- Agility / Cluster lifecycle
- Networking setup redundancy
- Hardware monitoring (PDUs, Fans)
- Hardware management  
Bios / Managementprocessor (IMM)




Hypervisor based  
Good



Virtualized Infrastructure  
with many layers

**Bare metal  
Better**



Bare metal  
with less layers

(1) VMs will add a performance penalty  
(2) OpenShift subscriptions licensing vs Socket based subscriptions

# IBM Fusion HCI

“Bare-metal OpenShift in a Box”

- **Scalable**
- **Resilient**
- **Highly Available**

Infrastructure automation & management included!



# Fusion HCI

## 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 per server
- GPU servers
- Rack-mounted KVM console (9155-TF5)

### Bare metal disadvantages (without HCP):


- Each cluster requires 3 bare metal control nodes
- 1 Cluster per rack. Not very agile / flexible

Worker  
Control



# Agenda

- Bare Metal OpenShift
- Hosted Control Plane
- Storage and Backup for hosted Clusters
- ArgoCD and Red Hat Advanced Cluster Manager

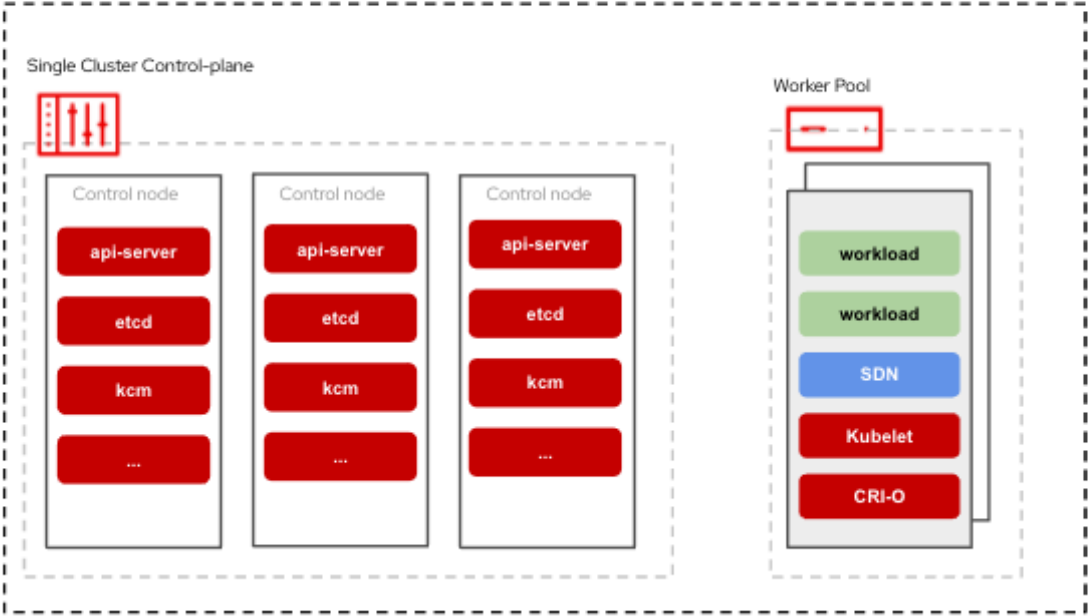


# Hosted Control Plane - Overview

## Standalone OpenShift

Control-Plane (CP) + Workers

Standalone OpenShift **Cluster** (dedicated CP nodes)



HCP removes the need for 3 dedicated control nodes for each OpenShift Cluster!

The control planes of the “Client” clusters are running on the Management Cluster.



# Hosted Control Plane - Overview

Hosted control plane clusters offer several advantages:

- Saves cost by removing the need to host dedicated control plane nodes
- Introduces separation between the control plane and the workloads, which improves **isolation** and reduces configuration errors that can require changes
- Decreases the cluster creation time by removing the requirement for control plane node bootstrapping
- Supports turn-key deployments or fully customized OpenShift Container Platform provisioning

## Sizing:

- 78 pods
- Three 8 GiB PVs for etcd
- Minimum vCPU: approximately 5.5 cores
- Minimum memory: approximately 19 GiB

# Hosted Control Plane - Overview

Because etcd writes data to disk and persists proposals on disk, its performance depends on disk performance.

Although etcd is not particularly I/O intensive, it requires a low latency block device for optimal performance and stability.

```
> oc get lvmclusters.lvm.topolvm.io -o yaml
apiVersion: lvm.topolvm.io/v1alpha1
kind: LVMCluster

spec:
  storage:
    deviceClasses:
      - deviceSelector:
          paths:
            - /dev/disk/by-path/pci-0000:61:00.0-nvme-1
          fstype: xfs
          name: hcp-etcd
          nodeSelector:
            nodeSelectorTerms:
              - matchExpressions:
                  - key: isf.ibm.com/lvm-node
                    operator: In
                    values:
                      - ""
```

```
status:
  deviceClassStatuses:
    - name: hcp-etcd
      nodeStatus:
        - deviceDiscoveryPolicy: Preconfigured
          devices:
            - /dev/nvme0n1
              name: hcp-etcd
              node: compute-1-ru5.fusion-hci.ibm1ab.de
              status: Ready
        - deviceDiscoveryPolicy: Preconfigured
          devices:
            - /dev/nvme0n1
              name: hcp-etcd
              node: compute-1-ru6.fusion-hci.ibm1ab.de
              status: Ready
        - deviceDiscoveryPolicy: Preconfigured
          devices:
            - /dev/nvme0n1
              name: hcp-etcd
              node: compute-1-ru7.fusion-hci.ibm1ab.de
              status: Ready
  ready: true
  state: Ready
```

# Hosted control plane on bare metal OpenShift

## Prerequisites

You can use multicluster engine operator with or without Red Hat Advanced Cluster Management. Red Hat Advanced Cluster Management also installs multicluster engine operator automatically and offers further multicluster capabilities.

Operator



```
apiVersion: metallb.io/v1beta1
kind: IPAddressPool
metadata:
  name: metallb
  namespace: metallb-system
spec:
  addresses: - 9.9.0.51-9.9.0.70
```



# Hosted control plane – deployment via hcp CLI

## [Cluster lifecycle with multicluster engine operator](#)

### [1.7.9.3.1. Creating a hosted cluster](#)

```
export CLUSTER_NAME="ddhcp-rhsummit"  
export PULL_SECRET="./pull-secret-dd.json"  
export MEM="48Gi"  
export CPU="8"  
export WORKER_COUNT="1"  
export ETCD_STORAGE="lvms-hcp-etcd"  
export SSHKEY="./id_rsa.pub"  
export OCPRELEASE="4.16.15-rc1+1"
```

```
- > ./create_ddhcp-rhsummit.sh
```

```
- 2024-10-11T16:40:47+02:00 INFO Applied Kube resource {"kind": "Namespace", "namespace": "", "name": "clusters"}  
- 2024-10-11T16:40:47+02:00 INFO Applied Kube resource {"kind": "Secret", "namespace": "clusters", "name": "ddhcp-rhsummit-pull-secret"}  
- 2024-10-11T16:40:47+02:00 INFO Applied Kube resource {"kind": "Secret", "namespace": "clusters", "name": "ddhcp-rhsummit-etcd-encryption-key"}  
- 2024-10-11T16:40:47+02:00 INFO Applied Kube resource {"kind": "Secret", "namespace": "clusters", "name": "ddhcp-rhsummit-ssh-key"}  
- 2024-10-11T16:40:47+02:00 INFO Applied Kube resource {"kind": "", "namespace": "clusters", "name": "ddhcp-rhsummit"}  
- 2024-10-11T16:40:47+02:00 INFO Applied Kube resource {"kind": "NodePool", "namespace": "clusters", "name": "ddhcp-rhsummit"}
```

```
dev/ocp-release.p{OCPRELEASE}
```

# Hosted control plane – deployment via hcp CLI

[Cluster lifecycle with multicluster engine operator](#)

[1.7.9.3.1. Creating a hosted cluster](#)

```
- > oc get hostedclusters -n clusters
```

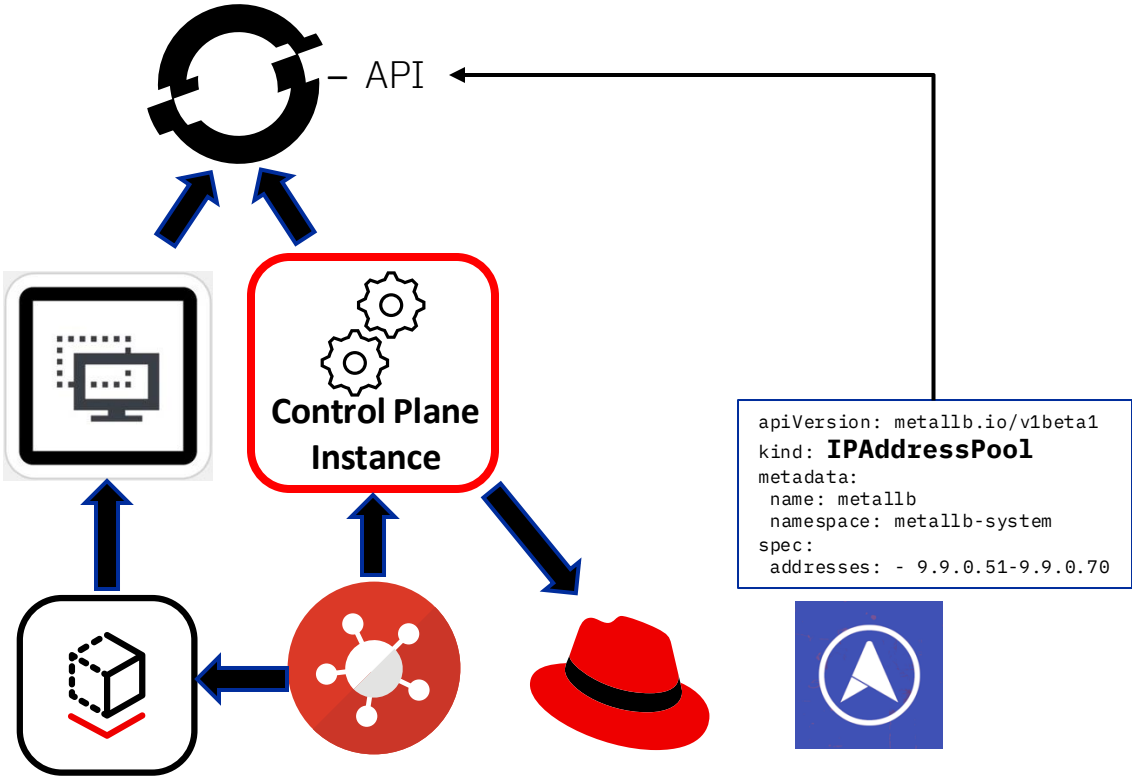
NAME	VERSION	KUBECONFIG	PROGRESS	AVAILABLE	PROGRESSING	MESSAGE
ddhcp-rhsummit control plane is available	4.16.15	ddhcp-rhsummit-admin-kubeconfig	Completed	True	False	The hosted
ddhcp3 control plane is available	4.15.33	ddhcp3-admin-kubeconfig	Completed	True	False	The hosted
ddhcp4 control plane is available	4.16.10	ddhcp4-admin-kubeconfig	Completed	True	False	The hosted

# Hosted control plane on bare metal OpenShift

## Building a hosted Cluster


Operator

```
#!/bin/bash
export CLUSTER_NAME="ddhcp-rhsummit"
export PULL_SECRET="./pull-secret-dd.json"
export MEM="48Gi"
export CPU="8"
export WORKER_COUNT="1"
export ETCD_STORAGE="lvms-hcp-etcd"
export SSHKEY="./id_rsa.pub"
export OCPRELEASE="4.16.15-multi"
export DISKSIZE="64"
#
# Valid releases: 4.14.38-multi , 4.15.35-multi ,
# 4.16.15-multi , 4.17.0-multi #
#
hcp create cluster kubevirt \
--name $CLUSTER_NAME \
--node-pool-replicas $WORKER_COUNT \
--pull-secret $PULL_SECRET \
--memory $MEM \
--cores $CPU \
--etcd-storage-class=${ETCD_STORAGE} \
--ssh-key ${SSHKEY} \
--release-image quay.io/openshift-release-dev/ocp-
release:${OCPRELEASE} \
--root-volume-size ${DISKSIZE}
```



# Agenda

- Bare Metal OpenShift
- Hosted Control Plane
- **Storage and Backup for hosted Clusters**
- Red Hat Advanced Cluster Manager



# Kubevirtcsi

=== Virtual Cluster ===

```
> oc get sc
```

osNAME	PROVISIONER	RECLAIMPOLICY	VOLUMEBINDINGMODE	ALLOWVOLUMEEXPANSION	AGE
kubevirt-csi-infra-default (default)	csi.kubevirt.io	Delete	Immediate	false	29m

```
> oc get pvc
```

NAME	STATUS	VOLUME	CAPACITY	ACCESS MODES	STORAGECLASS	VOLUMEATTRIBUTESCLASS
mariadb	Bound	pvc-27b40e9c-738d-41c9-800b-89b006705020	1Gi	RWO	kubevirt-csi-infra-default	<unset>

```
> oc debug node/ddhcp-rhsummit-b942de3d-wbnpt
```

```
[root@ddhcp-rhsummit-b942de3d-wbnpt ~]# mount |grep 27b40e9c-738d-41c9-800b-89b006705020  
/dev/sda on /var/lib/kubelet/pods/c7df7732-c8e4-43c5-a0dc-f9b853d44a3d/volumes/kubernetes.io~csi/pvc-27b40e9c-738d-41c9-800b-89b006705020/mount type ext4 (rw,relatime,seclabel)
```

```
[root@ddhcp-rhsummit-b942de3d-wbnpt ~]# lsblk
```

NAME	MAJ:MIN	RM	SIZE	RO	TYPE	MOUNTPOINTS
sda	8:0	0	1G	0	disk	/var/lib/kubelet/pods/c7df7732-c8e4-43c5-a0dc-f9b853d44a3d/volumes/kubernetes.io~csi/pvc-27b40e9c-738d-41c9-800b-89b006705020/mount

=== Hub Cluster ===

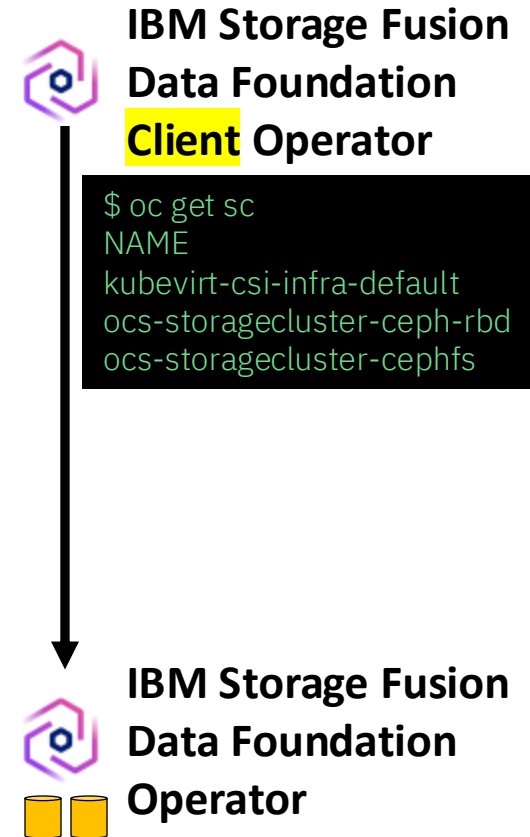
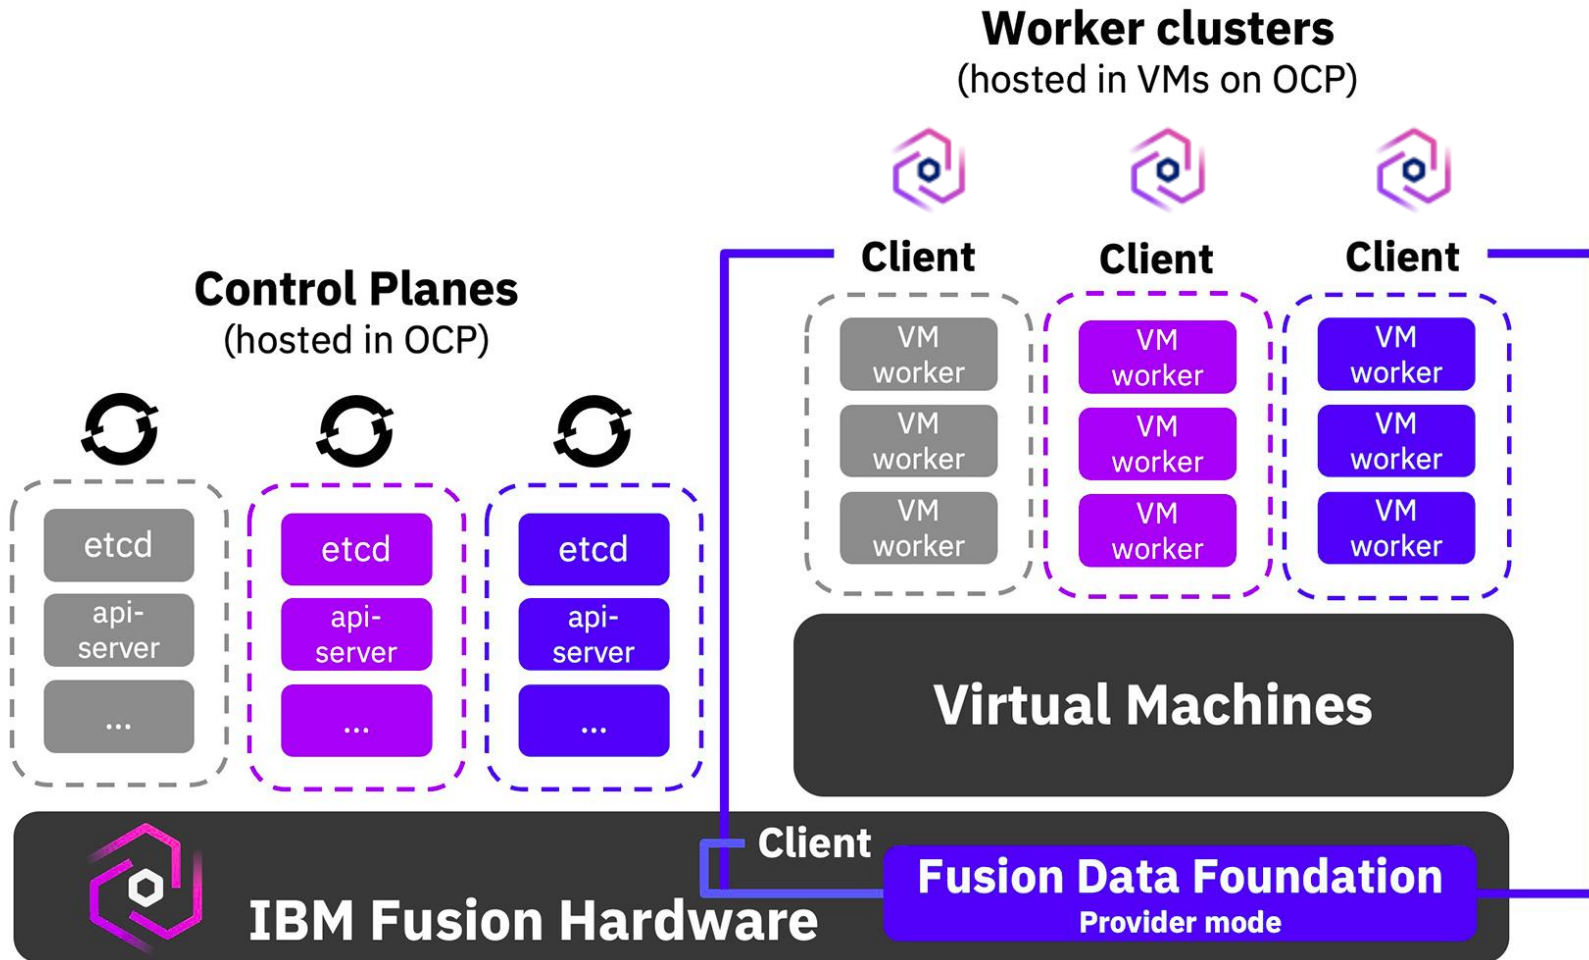
```
> oc get pvc -A |grep pvc-27b40e9c-738d-41c9-800b-89b006705020
```

clusters-ddhcp-rhsummit-60e5ebffd94f	1Gi	RWX	ocs-storagecluster-ceph-rbd	Bound	pvc-e67f4a00-17fa-4718-b988-60e5ebffd94f
--------------------------------------	-----	-----	-----------------------------	-------	--





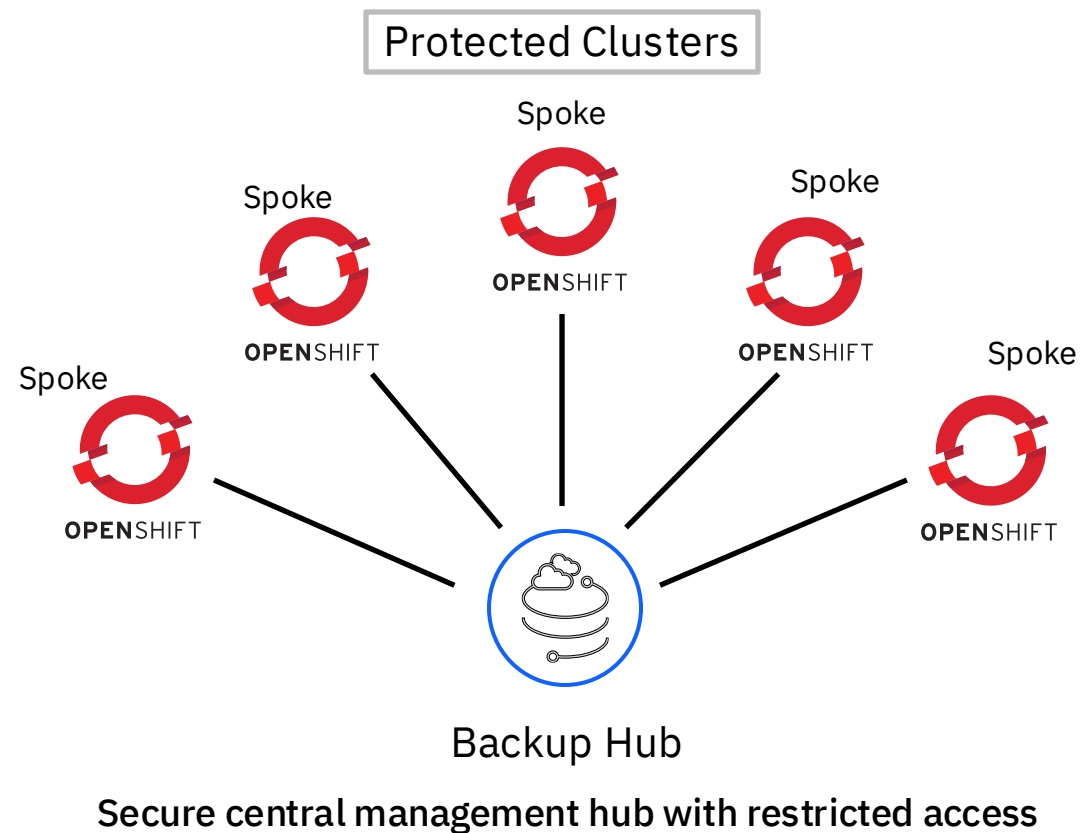
IBM Storage Fusion HCI System with **Fusion Data Foundation** (data provider mode) provides block/file/object storage on all Hosted Control Plane based clusters





Centralize management and monitoring of backups across multiple clusters

- **Simplify Management and Compliance**  
*Apply and enforce backup policies for all an organizations Kubernetes clusters from a central, secure management control plane.*
- **Security and Data Protection**  
*Separate application workloads from the backup service. If a workload cluster is compromised, backups are still safe.*
- **Efficient use of resources**  
*Deploy lightweight agent into clusters to be protected*

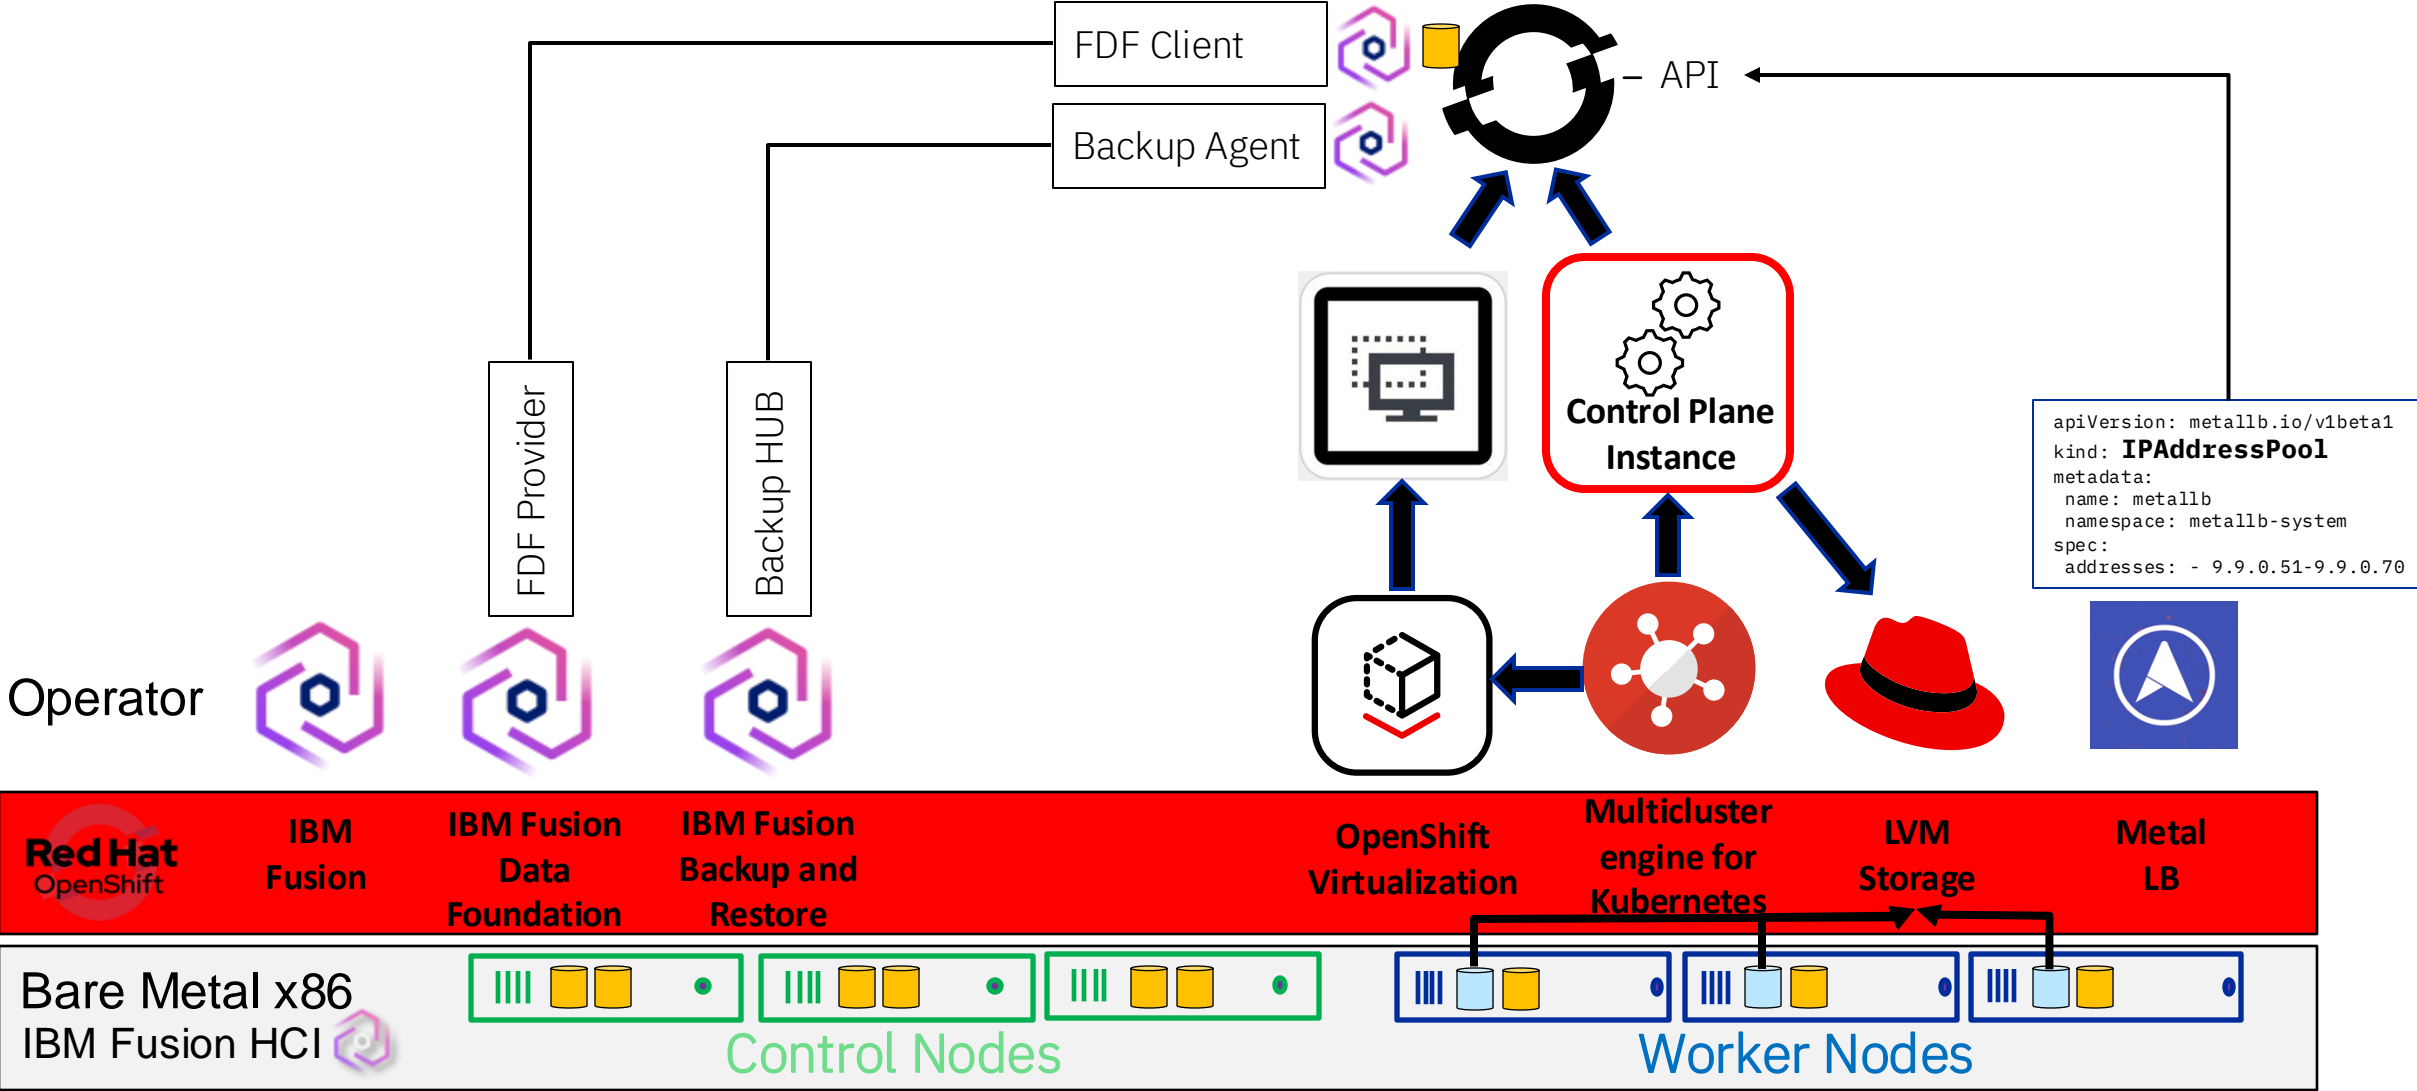


# IBM Fusion and Red Hat Data Services for OpenShift

Capability	Fusion	ODF Advanced	ODF Essentials
File, Block, Object Backed PV's (RWO, RWX, OBC)	✓	✓	✓
Available on all major platforms (x86, IBM Power, IBM Z & Linux ONE)	✓	✓	✓
Multi-Cloud Object Gateway (MCG)	✓	✓	✓
Compression & Deduplication for MCG	✓	✓	✓
Local snapshots & clones	✓	✓	✓
Basic cluster-wide encryption	✓	✓	✓
Cross-availability zone HA (3 zones)	✓	✓	✓
Internal mode deployment	✓	✓	✓
Advanced encryption with KMS support	✓	✓	
Stretch Cluster DR (OCP and ODF Stretched over 2 Zones + Arbitrator)	✓	✓	
Metro DR (sync replication)	✓	✓	
Regional DR (async replication)	✓	✓	
External mode deployment (shared storage)	✓	✓	
Mixed-workload support	✓	✓	
Global Data Platform (remote mount IBM Storage Scale gpfs filesystem)	✓		
<b>Fusion Data Foundation Provider mode:</b> Shared storage for other OpenShift clusters (consumer)	✓		
<b>Backup &amp; Restore</b> for Container and OpenShift Virtualization VMs	✓		
Application consistent Backup & Restore via recipes	✓		
Multi-Cluster Backup & Restore with Hub & Spoke design	✓		
Red Hat Advanced Cluster Management add-ons for Fusion Data Services (multi-cluster support)	✓		
Fusion Data Cataloging	✓		

# Hosted control plane on bare metal OpenShift

## Connecting Storage and Backup



# Fusion Data Foundation – Storage Clients

Red Hat OpenShift local-cluster Daniel-Danner

Administrator Home Operators Workloads Virtualization Migration Networking Storage

- Data Foundation
- Storage Clients**
- Backup & Restore
- Object Storage
- PersistentVolumes

### Storage clients

Generate client onboarding token Rotate signing keys

Filter Name Search by name...

Name ↑	Cluster ID	Openshift version ↕	Data Foundation version ↕	Last heartbeat ↕	
ocs-storagecluster	86a532ff-c01c-492e-9e2b-085076f3dc73	4.16.15	4.16.1	0 minutes ago	
ocs-storagecluster	956d0f45-0cd5-4c29-968d-160857cae76b	4.15.33	4.16.1	0 minutes ago	
ocs-storagecluster	7819d374-7ff5-4cea-800f-c2065adbd418	4.16.10	4.16.1	0 minutes ago	
ocs-storagecluster	9292eb66-36b7-4220-83ab-03d0e9e697df	4.16.5	4.16.0	0 minutes ago	
ocs-storagecluster	fd49eada-e900-4ab6-9991-176f44f6d20b	4.16.15	4.16.1	0 minutes ago	

# Fusion Backup & Restore - Topology

IBM Storage Fusion fusion-hci

Service upgrade available  
Version 4.16.1 is available for Data Foundation service. [View Services for details](#)

## Topology

Displays all clusters involved in Backup & Restore.


Connection status: Filter... Service status: Filter... Search [Connect cluster](#)

Name	Type	Connection status	Service status	Version	Backed up apps	Success rate	
apps.fusion-hci.ibmlab.de	Hub	-	Healthy	2.8.1	5/59	85.71 %	⋮
apps.ddhcp3.apps.fusion-hci.ibmlab.de	Spoke	Connected	Healthy	2.8.1	0/5	-	⋮
apps.ddhcp4.apps.fusion-hci.ibmlab.de	Spoke	Connected	Healthy	2.8.1	1/6	100.0 %	⋮
apps.fusion-sds.ibmlab.de	Spoke	Connected			0/0	-	⋮

Items per page: 25 1-4 of 4 items 1 1 of 1 page

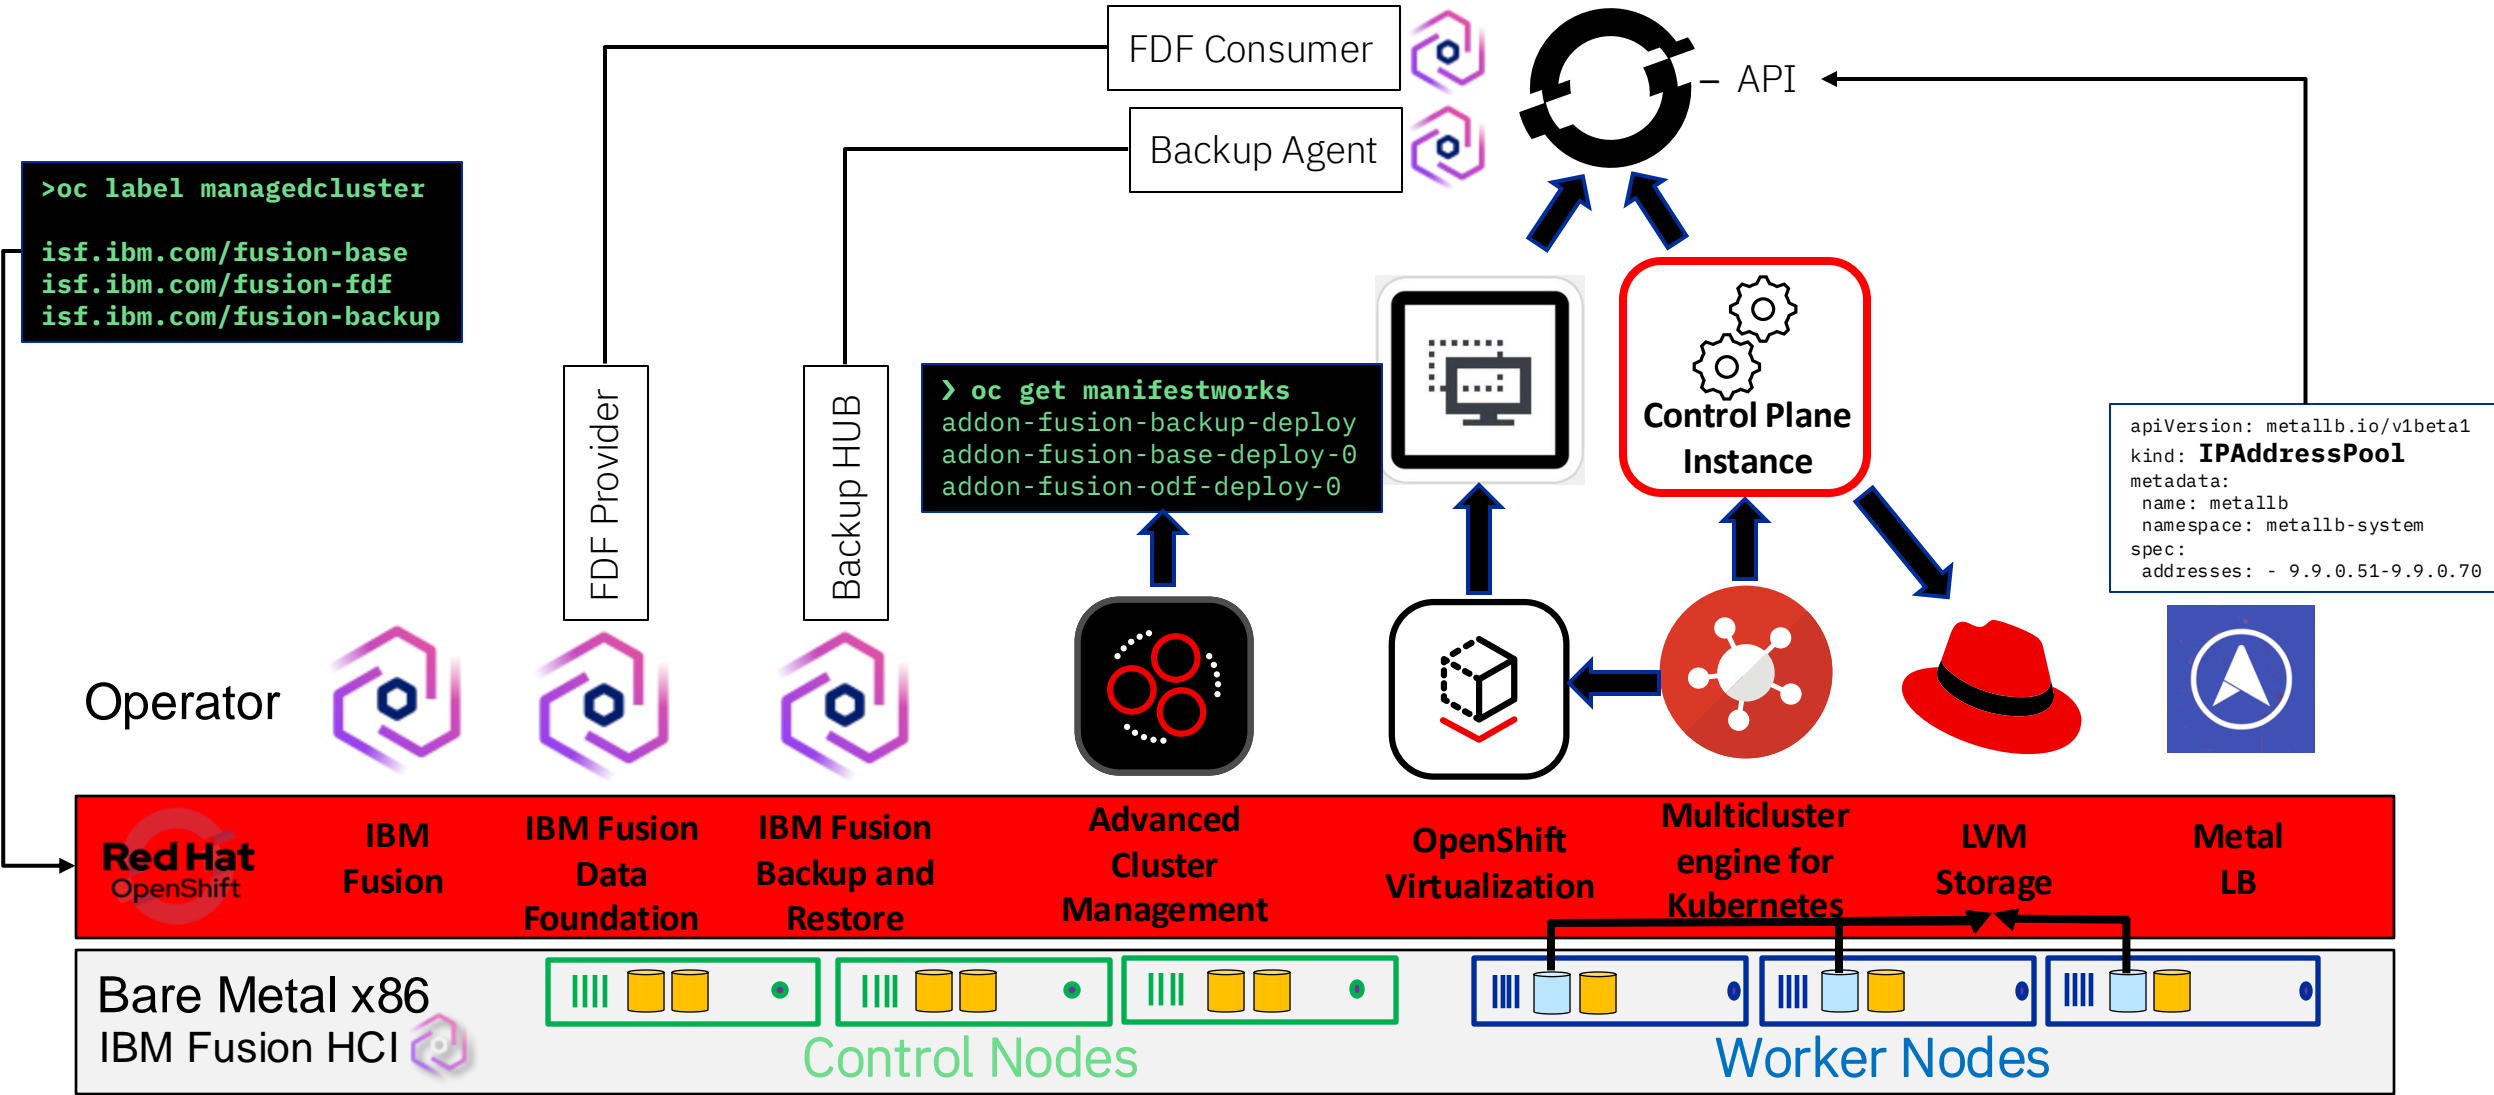
# Agenda

- Bare Metal OpenShift
- Hosted Control Plane
- Storage and Backup for hosted Clusters
- Red Hat Advanced Cluster Manager



# Hosted control plane on bare metal OpenShift

## Automation with ACM





# RHACM- Add-ons

Red Hat OpenShift | All Clusters | Daniel-Danner

Home | Infrastructure | Clusters | Automation | Host inventory | Applications | Governance | Credentials

Clusters > ddhcp-rhsummit

## ddhcp-rhsummit

Download kubeconfig | Actions

Overview | Nodes | Add-ons

Search

1 - 11 of 11

Name ↑	Status	Message
application-manager	Available	application-manager add-on is available.
cert-policy-controller	Available	cert-policy-controller add-on is available.
cluster-proxy	Available	cluster-proxy add-on is available.
config-policy-controller	Available	config-policy-controller add-on is available.
fusion-backup	Available	fusion-backup add-on is available.
fusion-base	Available	fusion-base add-on is available.
fusion-odf	Available	fusion-odf add-on is available.
governance-policy-framework	Available	governance-policy-framework add-on is available.
managed-serviceaccount	Available	managed-serviceaccount add-on is available.
search-collector	Available	search-collector add-on is available.
work-manager	Available	work-manager add-on is available.

1 - 11 of 11 items | 1 of 1 page

# RHACM- Add-ons Topology

The screenshot displays the Red Hat OpenShift console interface. On the left is a dark sidebar with navigation options: Home, Infrastructure, Applications, Governance, and Credentials. The main content area shows the 'All Clusters' dropdown and the breadcrumb 'Applications > ibm-fusion-hcp'. The page title is 'ibm-fusion-hcp', and the 'Topology' tab is selected. The topology diagram illustrates the following structure:

- OCP Application:** 'ibm-fusion-hcp' (Status: Success)
- Cluster:** (Status: Success)
- Child Resources:**
  - FusionServiceInstance:** 'odfmanager' (Status: Success)
  - Deployment:** 'fusion..r-controller' (Status: Success)
    - Replicaset:** 'fusion..r-controller' (Status: Success)
    - Pod:** 'fusion..r-controller' (Status: Success)
  - Odfcluster:** 'odfcluster' (Status: Success)

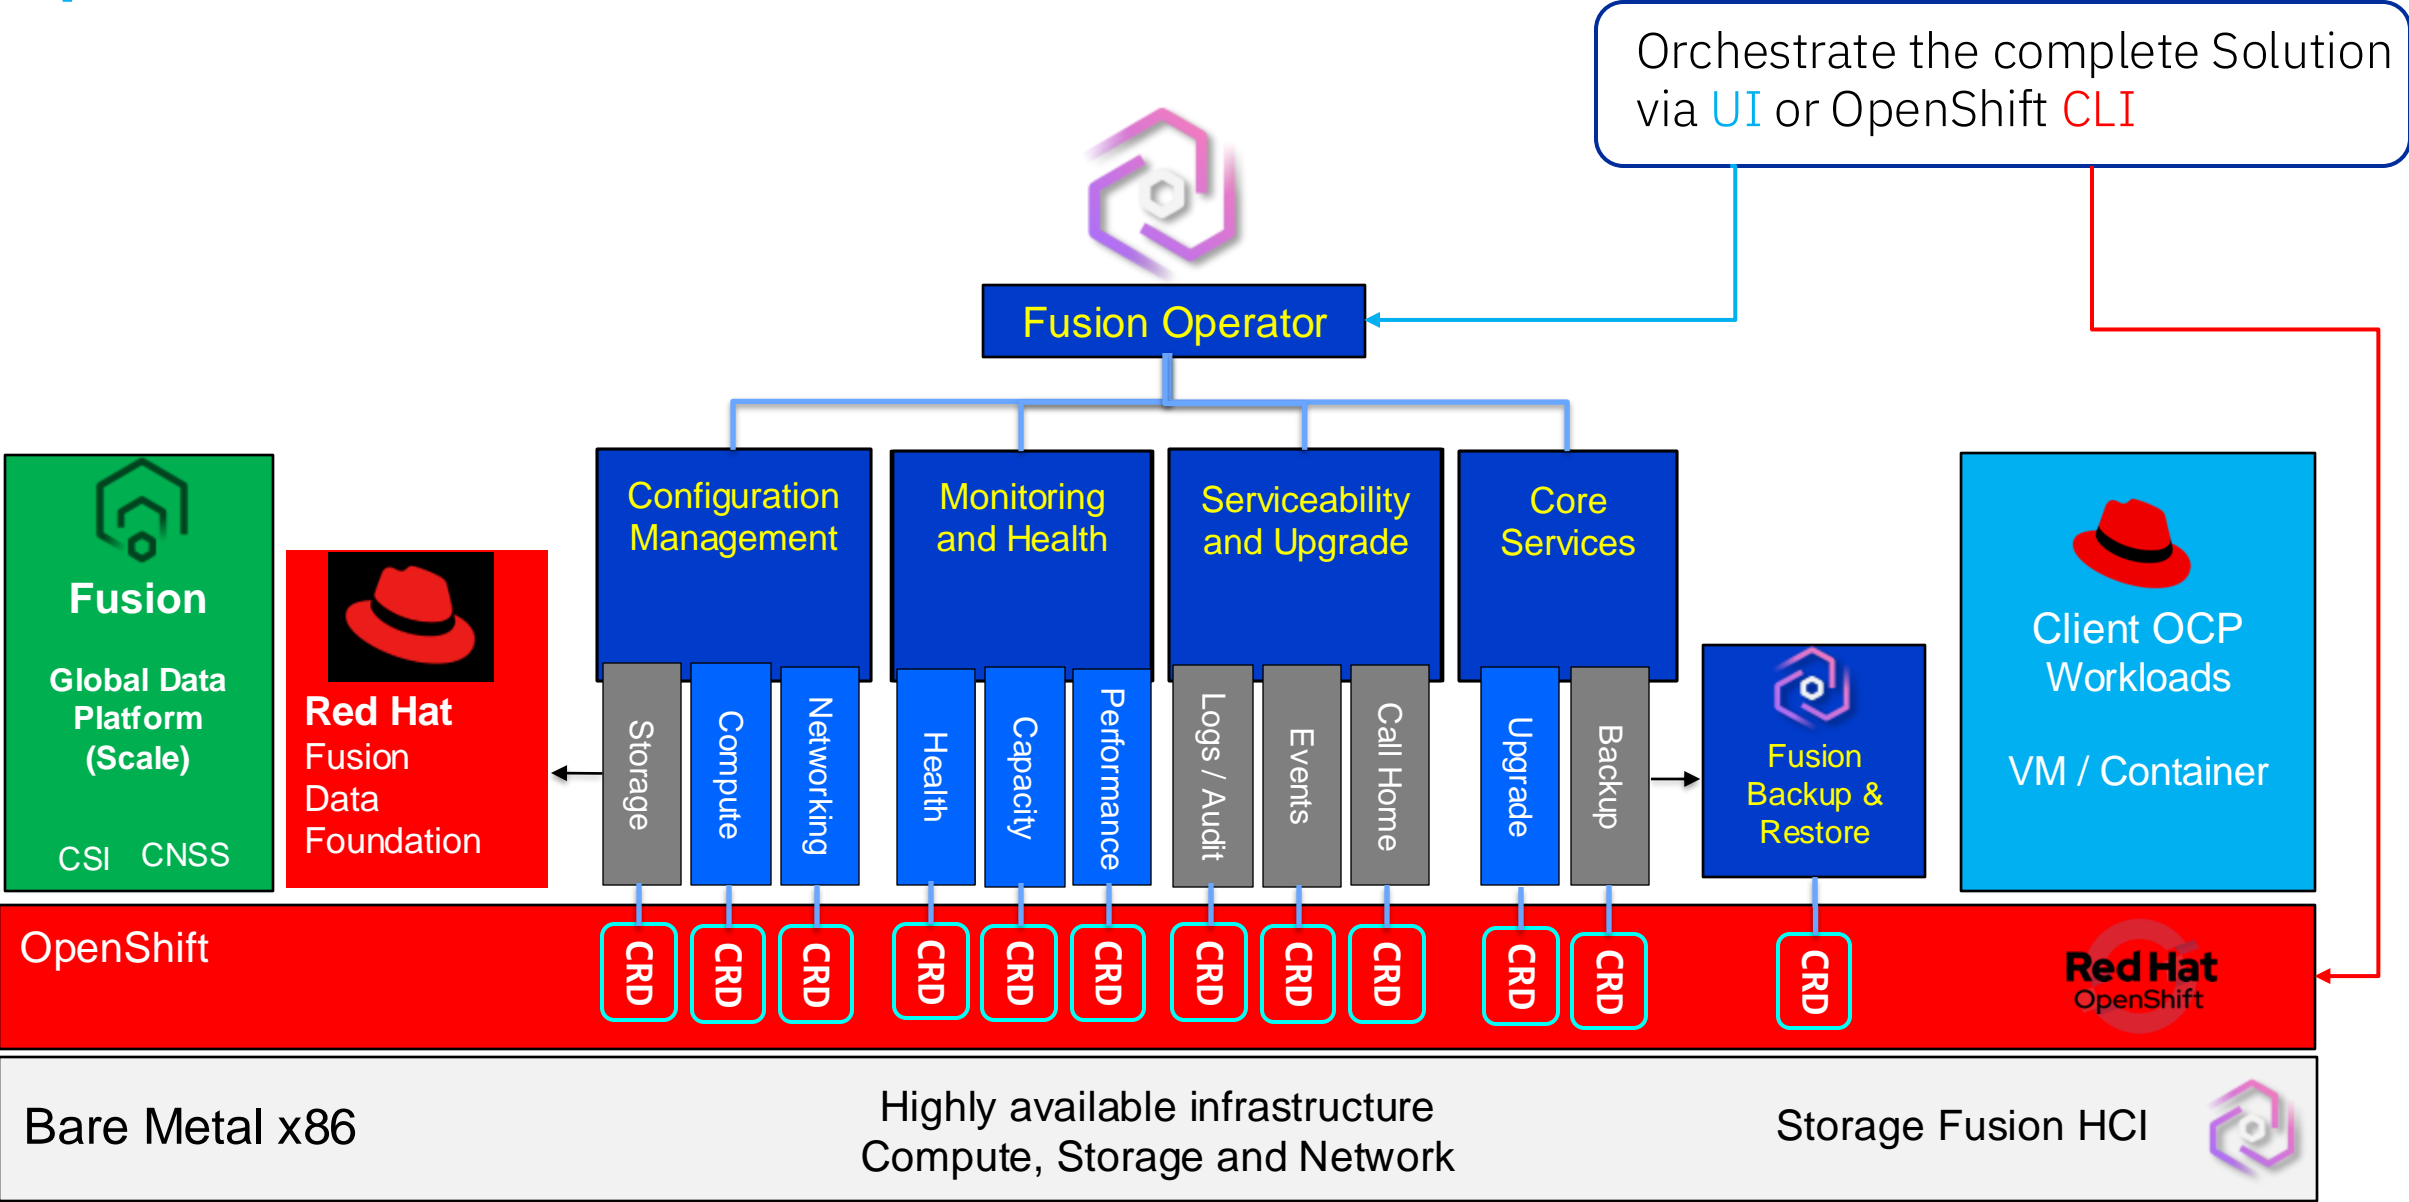
# RHACM- manifestwork

```
$ oc get manifestwork addon-fusion-base-deploy-0 -n rhcs-1 -o yaml
```

```
apiVersion: work.open-cluster-management.io/v1
kind: ManifestWork
...
spec:
  ...
  workload:
    manifests:
    - apiVersion: apps/v1
      kind: Deployment
      ...
      spec:
        app: fusion-agent
        spec:
          containers:
          - args:
            - fusion-base
            - --hub-kubeconfig=/var/run/hub/kubeconfig
            - --cluster-name=rhcs-1
            - --addon-namespace=open-cluster-management-agent-addon
            - --addon-name=fusion-base
            - --fusion-install-namespace=ibm-spectrum-fusion-ns
            - --catalog-source-name=ibm-operator-catalog
            - --catalog-source-image=icr.io/cpopen/isf-operator-software-catalog:2.8.1
            - --catalog-source-display-name=Customer Fusion Operator Catalog
          image: cp.icr.io/cp/isf/isf-hcp-manager@sha256:47dfa9f200
```

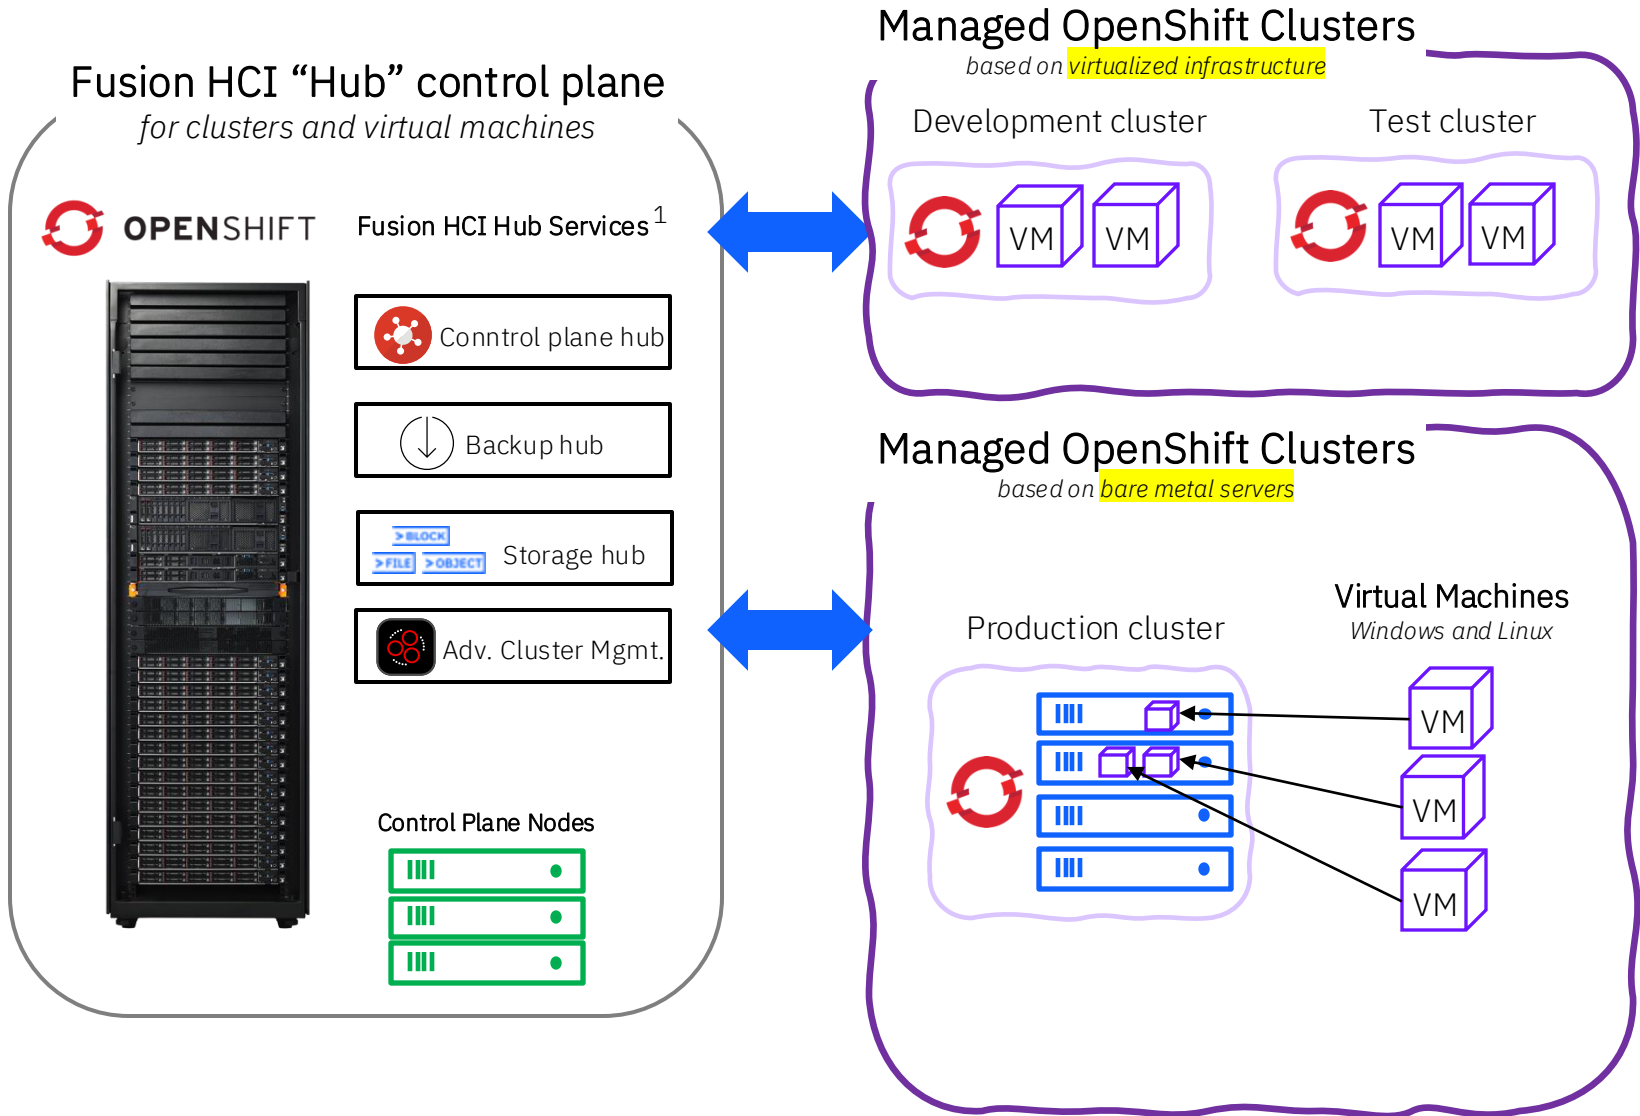

# IBM Fusion

## Complete Container Native Software Architecture



# Fusion HCI as a Control plane / Storage / Backup Hub

for multiple clusters and virtual machines



## Key Concepts

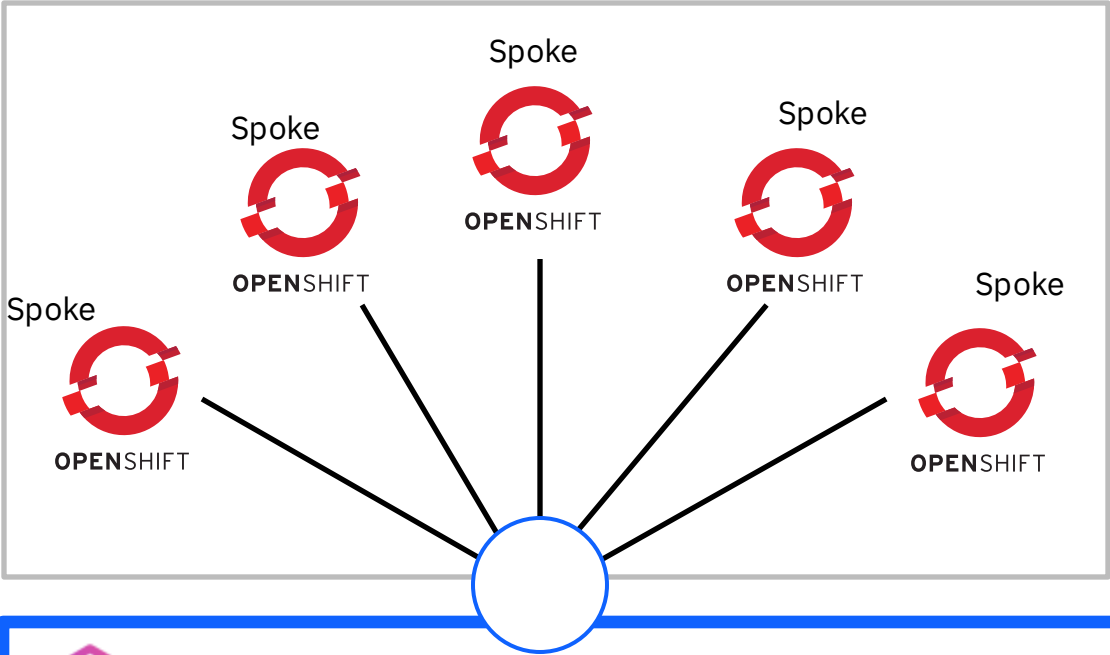
- Use Red Hat Advanced Cluster Manager to deploy hosted clusters
- Hosted clusters use FDF Client to consume storage from the Fusion HCI storage hub
- Administrators assign backup policies for hosted clusters
- Strong isolation / separation of concerns between hosted control planes and worker nodes
- Strong storage tenant isolation between hosted clusters (Data Foundation)
- Can also run Windows and Linux VMs with Kubevirt KVM and manage them from the Fusion HCI "Hub"
- Roadmap
  - Hosted clusters on bare-metal servers
  - DR of hosted clusters

<sup>1</sup> OpenShift and Advanced Cluster Manager are purchased separately

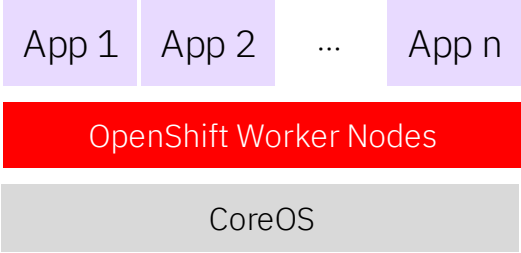
# OpenShift aaS on premise:

Based on Bare Metal OpenShift with Hosted Control Planes

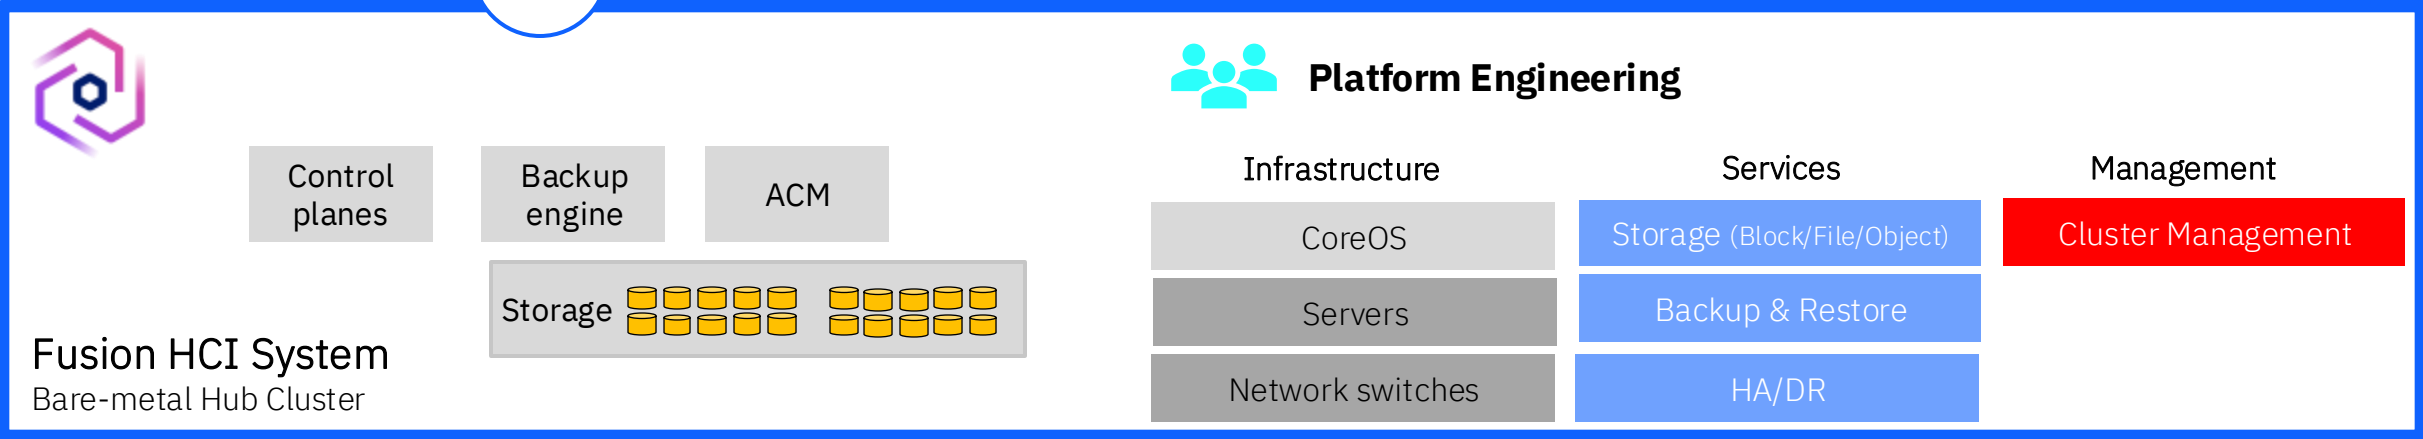
Managed spoke clusters provisioned on Fusion HCI System



- Application Developers
- Site Reliability Engineers
- Application DevOps Team



Interactions through APIs and Service Portals



Session: 12:05 - 12:35



[red.ht/rhsc24-de-s2](https://red.ht/rhsc24-de-s2)

# Jetzt Session bewerten!

Einfach QR-Code  
scannen, Session wählen  
und bewerten. **Vielen  
Dank!**



Thank you

