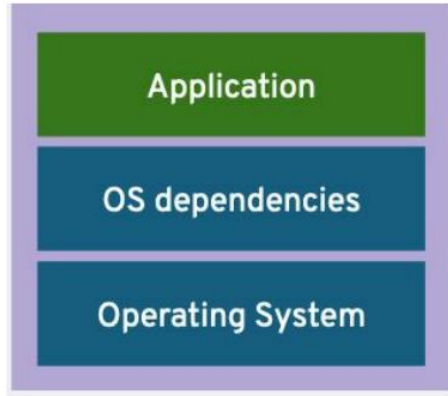



“OpenShift Container Storage: Stockage résilient pour Containers & Hybrid Cloud (OCS) ”

Laurent-Xavier Murgier, BDM SDS & Open HCI : lmurgier@redhat.com

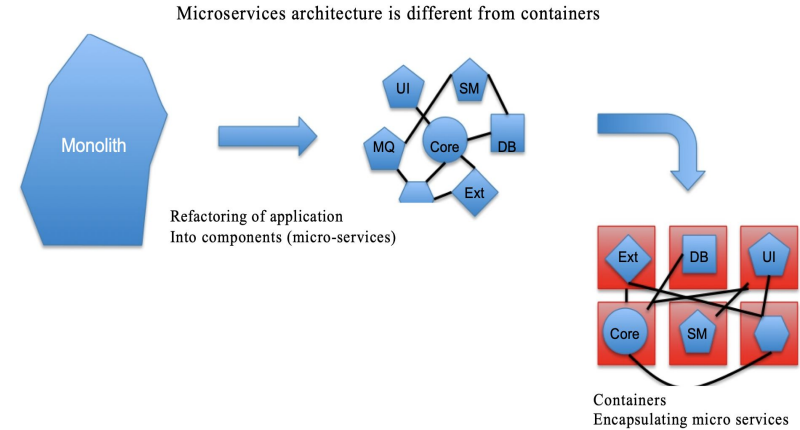
CONTAINER DEFINITION : CONTEXT



LEGACY APPLICATION



MODERN APPLICATION




MICRO-SERVICES & CONTAINERS

NEED FOR PERSISTENT STORAGE

CONTAINERS NEED STORAGE

Containers are not persistent by default. App data (process) is lost when containers die.



OCS : RESOURCES PROVISIONING &

OCS : part of OCP

Dashboards

- Overview
- Storage**

Details

Property	Value
Name	cluster-name
Provider	AWS Cloud
Region	region-name
Availability zone	zone-name
OpenShift version	v4.0
Kubernetes version	v1.12.4+670b342
Operating system	RHEL Server

Health [See all](#) **Compliance** [See all](#)

Cluster is healthy ✓ Cluster is compliant ✓

Utilization

Resource	Used	Available
CPU	70%	2.2 available of 6.3 GHz
Memory	25%	76.8 available of 1.05 Ti
Storage	57%	120 available of 279 Ti
Network	6%	9.4 available of 10 Gbps

Top consumers

Pod/VM	CPU time
pod-1	19.3%
vm-1	17.2%
pod-27	15.0%
pod-82	13.2%
pod-23	4.3%
vm-23	3.1%
pod-23	2.7%
vm-23	2.4%

Inventory

Resource	Count
3 Hosts	
24 Disks	
106 Pods	
102 VMs	

Utilization [6 hours](#)

Resource	6 hour average
CPU	42%
Memory	62.3 Gi
Network Transfer	1.4 Gbps external 3.5 Gbps internal
Disk IO R/W	604.6 Mi/s
Disk usage	4.2 Ti
Number of pods	23

Events

A few seconds ago - 1 time in the last 24 hours

pod-name-1
From kubelet ip-10-0-147-109.ec2.internal pulling image "openshift/hello-openshift"

2 minutes ago - 2 times in the last 24 hours

host-name-1
From kubelet ip-1-0-120-109.ec2.internal CPU utilization over 50%, migrated 2 pods to other hosts.

10 minutes ago - 1 time in the last 24 hours

host-name-1
From kubelet ip-1-0-120-109.ec2.internal Memory utilization over 80%. Migrating workloads to other hosts.

Why OCS is natively integrated with OCP

ELASTICITY, DATA PORTABILITY, PROTECTION AND LOWER COSTS

- CONTAINERS ARE SYSTEM PROCESSES, & CONTAINER STATES ARE VOLATILE BY DEFAULT
- CONTAINERS NEED PERSISTENT STORAGE
- LEGACY STORAGE DOESN'T FEET CORRECTLY WITH OCP NEEDS FOR SEVERAL REASONS
- PROVISIONING & HANDLING OF CONTAINER STORAGE BECOMES COMPLEX SOON WITH LEGACY STORAGE
- DIFFERENT STORAGE NEEDS IN OPENSIFT REQUIRE DIFFERENT SOLUTIONS – TOO COMPLEX FOR LEGACY

STATIC PROVISIONING

Storage Admin creates storage volumes upfront

Selects a predefined volume based upon claim, nearest available size No automated housekeeping - causing administrative burden

Error Prone due to increasing complexity and resulting administrative overhead

- **REDHAT OCS** - NOT JUST STORAGE FEATURE – IT'S **RESOURCE ORCHESTRATION & STORAGE MANAGER**

DYNAMIC PROVISIONING : ALL OPENSIFT RESOURCES

OpenShift user requests for storage by persistent volume claim (PVC)

Storage system does the needful in an automated way



STORAGE OPTIONS FOR CONTAINERS



STORAGE ARRAYS AND APPLIANCES CLOUDS

- Scalability and high availability fall short of customer needs
- Vendor lock in and high TCO
- Monolithic appliance model
- **Tech Refresh impact** on OCP users & Devops eg“Service Interruption”



SILOED OR POINT PLAY STORAGE SOLUTIONS

- No hybrid cloud support
- No unified control plane (K8s)
- Lack of automated or dynamic handling of storage needs
- No aligned with OCP releases cadence



PORTABLE STORAGE ACROSS ON-PREM / MULTIPLE PUBLIC *OpenShift Container Storage* *Red Hat*

- Runs in public clouds
- Single Vendor Support model
- Seamless user experience for developers and DevOps
- Full integration and all automated



OCS is the OPENSIFT “BACK END” :

- **REGISTRY STORE**

Where container base images reside.

By default not redundant, therefore possible point of failure. OCS resolves this.

- **PERSISTENT BLOCK STORAGE FILE STORAGE FOR CONTAINERS**

Specific storage type for specific workloads that require certain performance i.e. Database workloads, Logging where Elastic or equivalents are involved.

- **PERSISTENT FILE STORAGE FOR CONTAINERS**

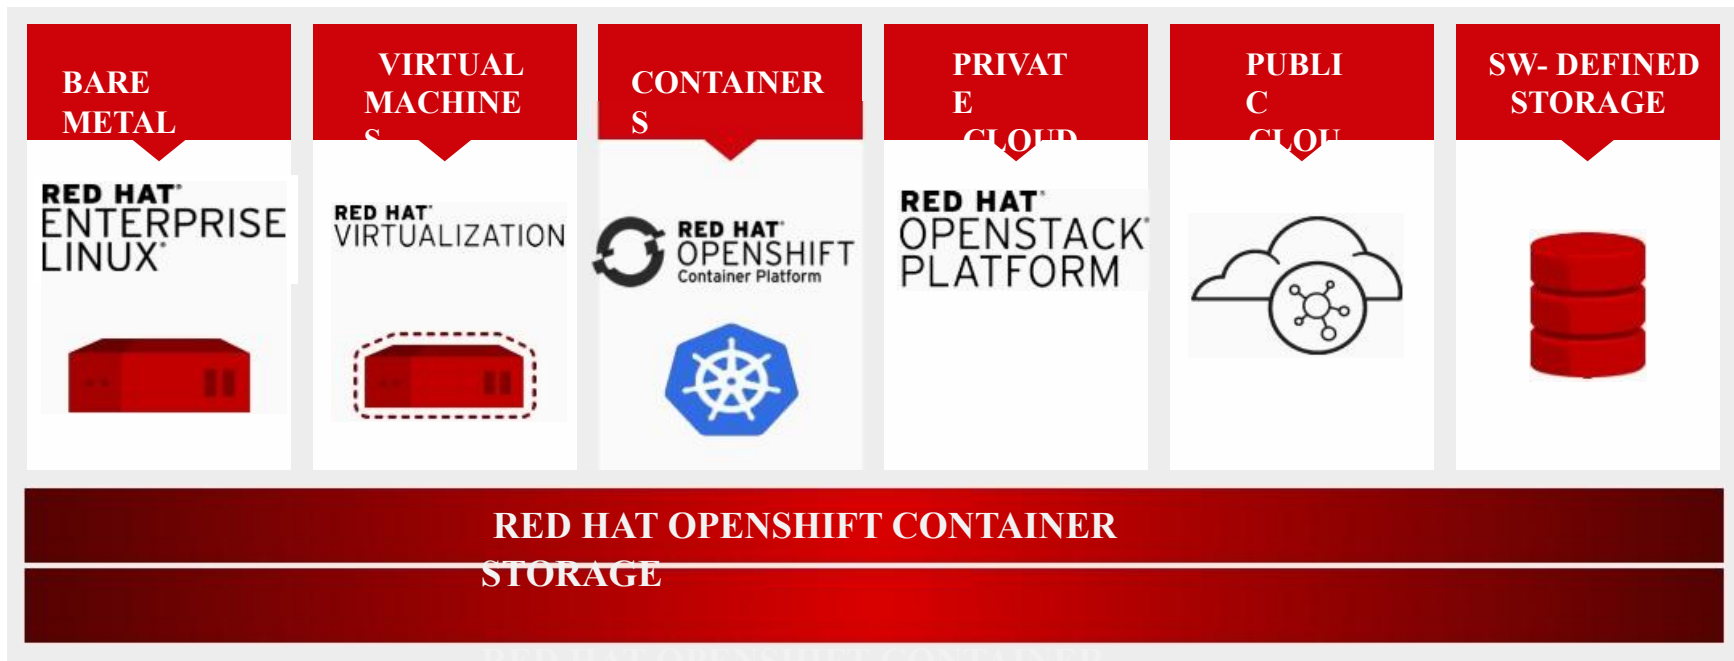
Container application state is held in this persistent file storage

- **PERSISTENT OBJECT STORAGE FOR CONTAINERS**

Container and Cloud application state running with S3 Protocol eg Big Data

Consistent Storage Experience Across Hybrid Cloud

APPLICATION PORTABILITY AND LOWER COSTS



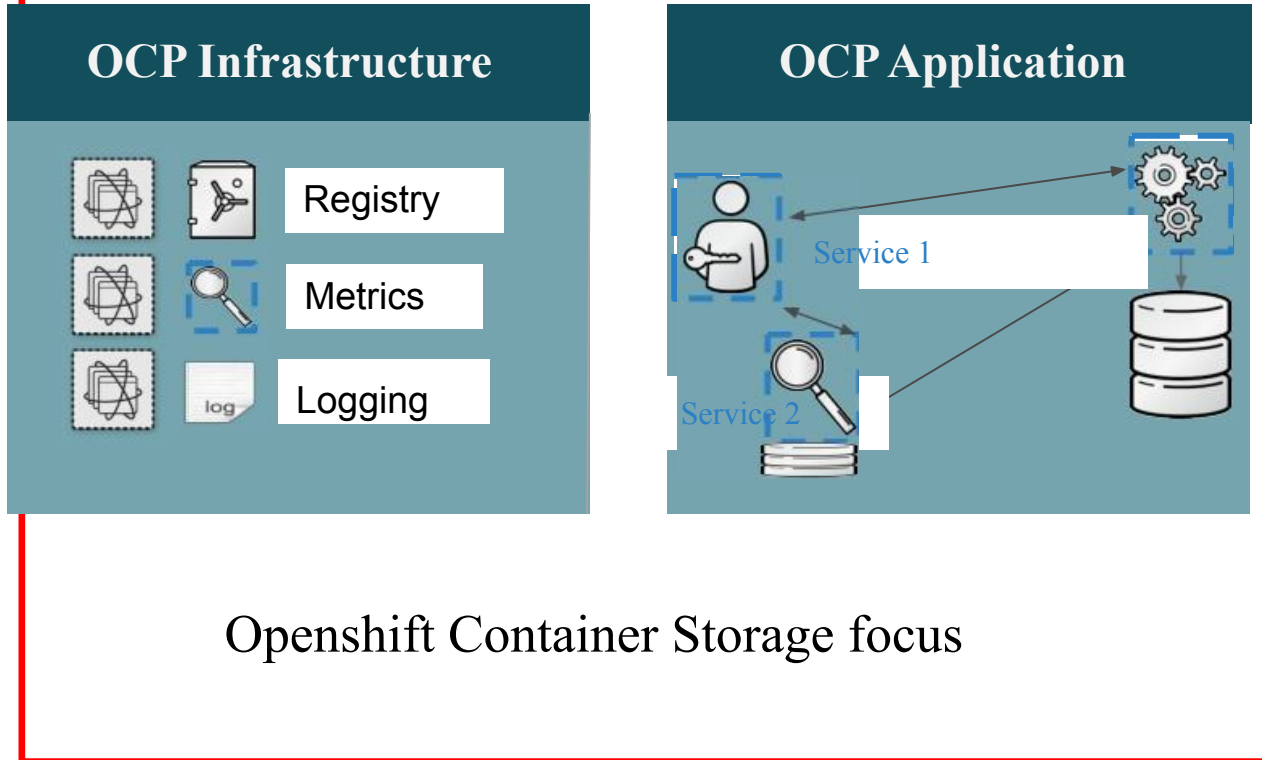
RED HAT OPENSIFT CONTAINER
PLATFORM

A low-angle, upward-looking photograph of several modern skyscrapers. The image is heavily overlaid with a semi-transparent red filter, which is darker in some areas and lighter in others, creating a dramatic, high-contrast effect. The perspective makes the buildings appear to converge towards the top center of the frame.

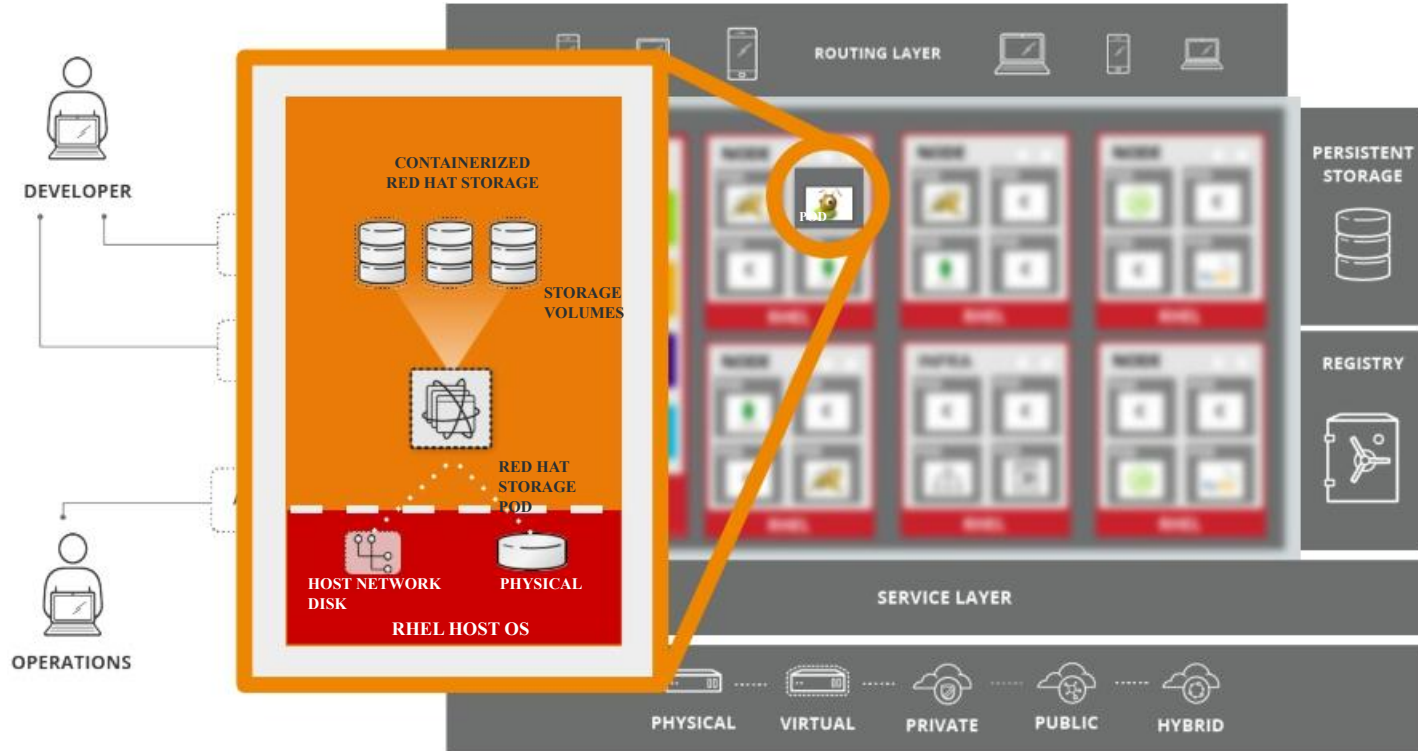
OCS Gen 3

OpenShift Container Storage
Based on GlusterFS technology

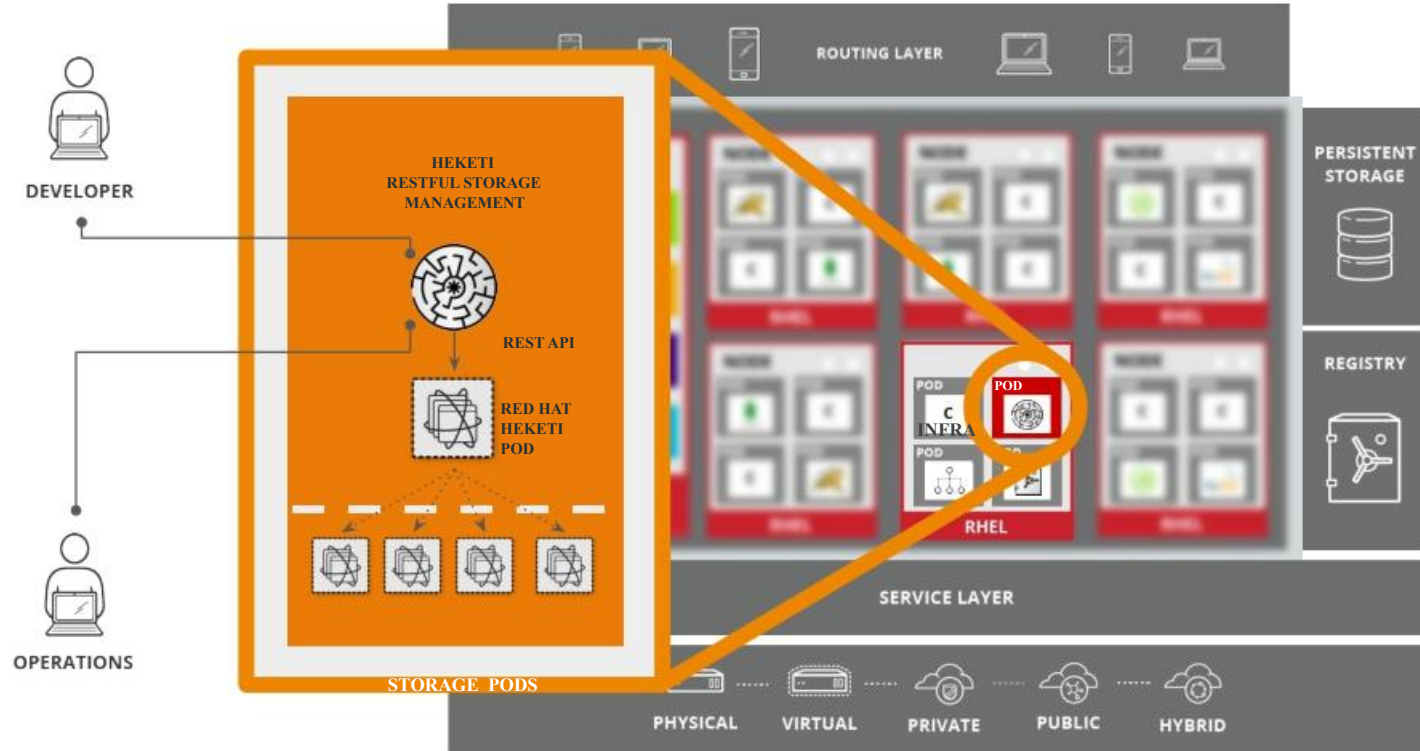
OPENSIFT PLATFORM STORAGE NEEDS



OCS Data Plane



OCS Control Plane



OCS product features

OCP Gen 3.x

Most part of our
customers are here !

Integration	Features	CNS 3.9	OCS 3.10	OCS 3.11
kubernetes	Support RWX, RWO, ROX	✓	✓	✓
	Dynamic provisioning	✓	✓	✓
	PVC resize (oc edit pvc)	✓	✓	✓
Openshift	Prometheus storage API metrics	□	✓	✓
	Deploy with OCP ansible playbook	□	✓	✓
	PVC resize (web-console)	□	□	✓
	Storage class volume options	□	✓	✓
	Infra support registry, metrics, logging	✓	✓	✓
Storage	Block storage with iSCSI support	✓	✓	✓
	File sharing with glusterfs-fuse	✓	✓	✓
	Object with S3/Swift (tech preview)	✓	✓	✓
	Snapshot and geo-replication	✓	✓	✓
	Arbiter volume (replica 2 + metadata)	□	✓	✓
Infrastructure	Public Azure, AWS, GCP	✓	✓	✓
	Private Openstack	✓	✓	✓
	Virtualization (VMW, RHV)	✓	✓	✓



TRUSTED ENTERPRISE KUBERNETES

- Trusted Host, Content, Platform
- Full Stack Automated Install
- Over the Air Updates & Day 2 Mgt

A CLOUD-LIKE EXPERIENCE, EVERYWHERE • Hybrid, Multi-Cluster Management

- Operator Framework
- Operator Hub & Certified ISVs

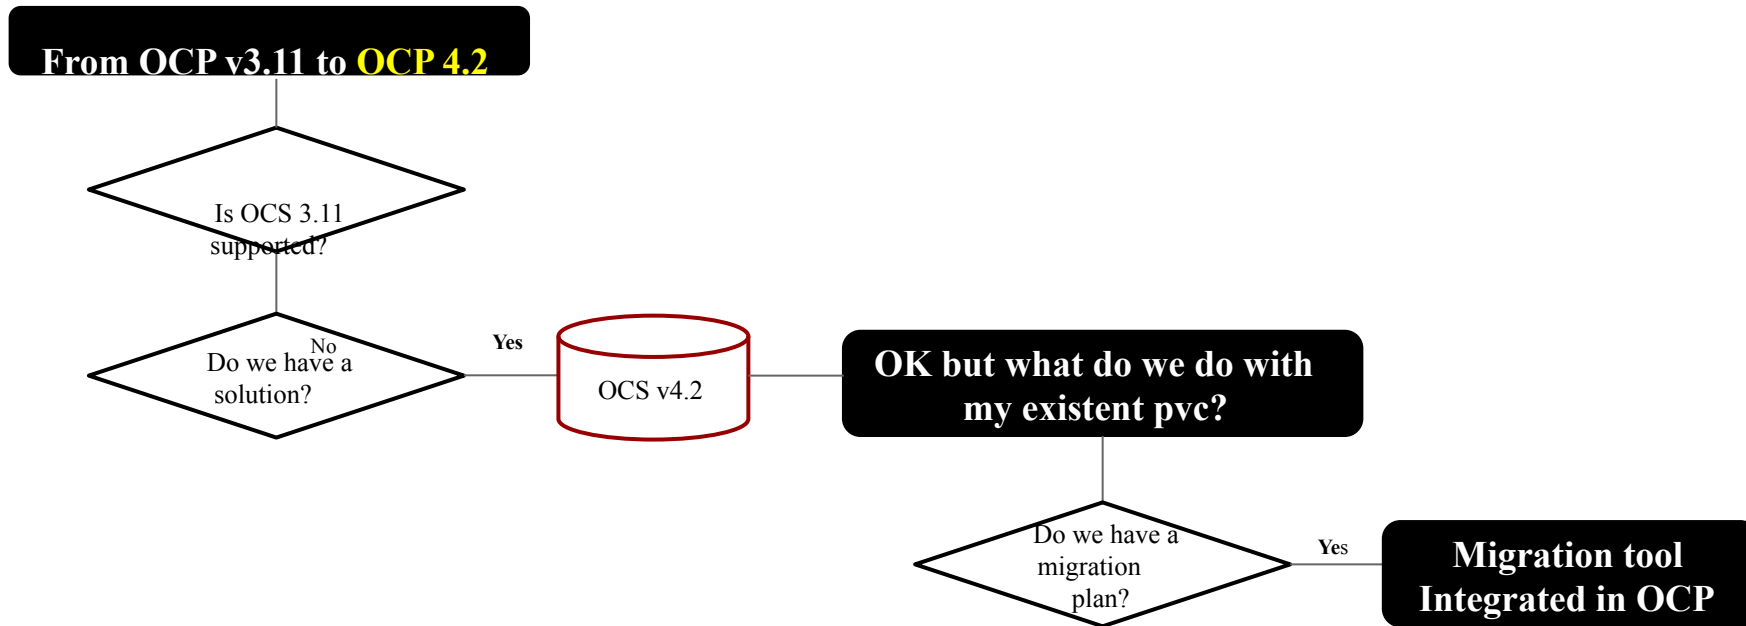
EMPOWERING DEVELOPERS TO INNOVATE

- OpenShift Service Mesh (Istio)
- OpenShift Serverless (Knative)
- CodeReady Workspaces (Che)



It's OCS 3.11 ready for OCP 4.2?

Use case for existent OCP 3.11 to OCP 4.2



OPENSIFT OPERATOR FRAMEWORK




Picture source: pexels.com

Migrating to OpenShift 4 = Migration Tool



Tooling and advice for moving from OpenShift 3.x to 4.x



OPENSIFT OPERATOR FRAMEWORK

OpenShift Gen v4.2 uses **OPERATORS** to manage EVERY ASPECT of the cluster.

This includes operators that manage **essential** Kubernetes project components like the api server, scheduler, and controller manager.

Additional operators for components like the cluster-autoscaler, cluster-monitoring, web console, dns, ingress, networking, node-tuning, and authentication are included to provide management of the entire platform.

App migration experience

Using open source tooling based on Velero

Velero is an upstream project previously known as Ark. Check out [this video](#) if you are curious and want to get a sneak peek at our capabilities.

What's moved during a migration

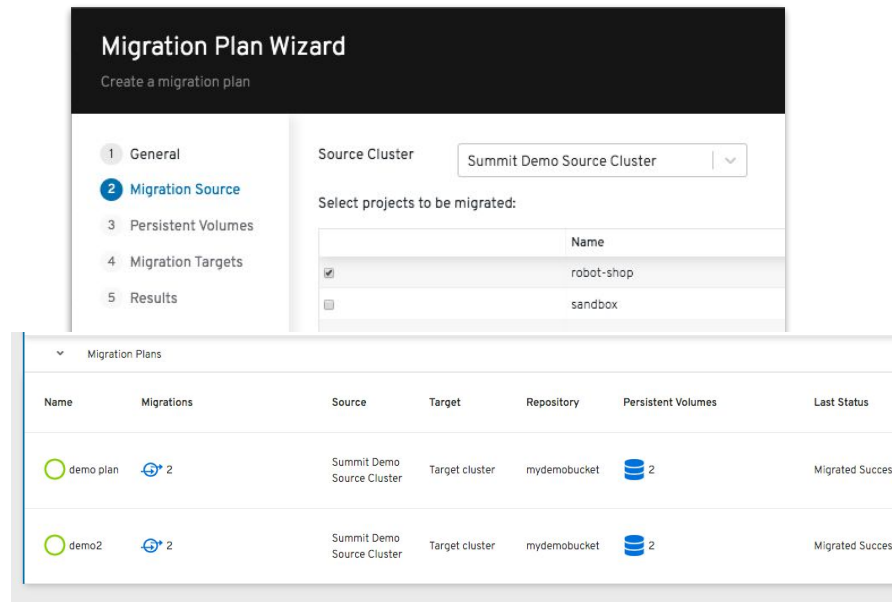
- Namespaces
- Persistent Volumes (move or copy)
- All important resource objects (Deployments, StatefulSets, etc)

Available in OpenShift 4.2

Customers are anxious to get their hands on this, but we want to get it right. We would love to receive sample application workloads to test.

Product Manager: Maria Bracho

Not Available Yet



OCS 4.2

Openshift Container Storage Based on CEPH Technology

- OVERVIEW
- META OPERATOR
- MACHINE API and MACHINE SETS
- ROOK CEPH OPERATOR and CEPH PODS
- MCG (NOOBAA) OPERATOR & PODS

WHY WE MOVE TO CEPH

- **MOTIVATION**

As cloud-native applications have evolved, we are noticing more customer requests for : native, easy to use and elastic storage (SDS) including S3/object interface (apps like registry, chargeback, metering, AI/ML) with persistent volumes (RWX & RWO) on the platform.

By leveraging Ceph, OCS can now provide a 3in1 Storage solution (Block, File & Object Storage) with production-grade S3 interface in addition to persistent volumes for stateful applications.

WHAT CHANGED

- **OPENSIFT**

OpenShift transitions from OCP 3.11 to OCP 4.2

- **OPENSIFT CONTAINER STORAGE**

also transitions from OCS 3.11 to OCS 4.2

- **OCS 4.2** will be based on **ROOK.IO**, which uses **Red Hat Ceph Storage** and the recently acquired **NooBaa** technology as the **Red Hat Multi Cloud Gateway**

- **Will OCS 3.11 work with OCP 4.2?**



NO. [Migration tooling](#) will be available to facilitate the move to OCS 4.2

- **MIGRATION PATH**

There will be a supported migration path offered for OCS 3.11 to OCS 4.2

WHAT DOES A TYPICAL OCS DEPLOYMENT LOOK LIKE?

CONFIDENTIAL Designator



Motivation

=

Get the best of OpenShift Customer experience without being on call



THE OCS 4.2 TECHNOLOGY STACK



+



+



ROOK & META OPERATOR

THE OPERATOR for Red Hat OpenShift Container Storage (OCS)

Functionalities delivered by the Meta operator:

- Facilitates the other OCS operators by performing administrative tasks
- Watches these other OCS operators
- Configures their CustomResources (CRs)

OCS 4.2 makes use of **MACHINE API**

- Performs **NODE** host provisioning after cluster installation
- Offers an Elastic, Dynamic provisioning method
- Manage primary resources : **MACHINES** Unit that describes the host for a Node, **MACHINE SETS** Groups of Machines, acting as replicas, can scale up- or down



ROOK CEPH STORAGE OPERATOR

- CLOUD NATIVE & CONTAINER STORAGE ORCHESTRATOR FOR AUTOMATED DEPLOYMENT & LIFE CYCLE MANAGEMENT
- FUNCTIONS :
BOOTSTRAP, CONFIGURE, PROVISION AND MONITOR THE STORAGE CLUSTER
FOR SCALING, UPGRADING, MIGRATION, DISASTER RECOVERY & RESOURCES MANAGEMENT
- **ROOK AND CEPH - CAN PROVIDE “3in1” APPs ACCESS TO STORAGE WITH :**
 - 1.BLOCK DEVICES
 - 2.SHARED FILE SYSTEM VOLUMES
 - 3.OBJECT STORAGE

MACHINE SET - OCS RUNS ON A MACHINE SET

MACHINE SET - MORE DETAIL

- **METADATA** - Includes info about MachineSet NAME and other labels
- **SELECTOR** - Defines which MACHINES are associated with the MACHINESET
- **TEMPLATE** - METADATA: MACHINES properties and details and location
 - SPEC: Details on how the MACHINE is created
 - PROVIDERSPEC: Details on provider specifics, e.g. AMI ID

OCS - RUNS ON AN A BY ADMINISTRATOR SELECTED MACHINESET USING TAINTS

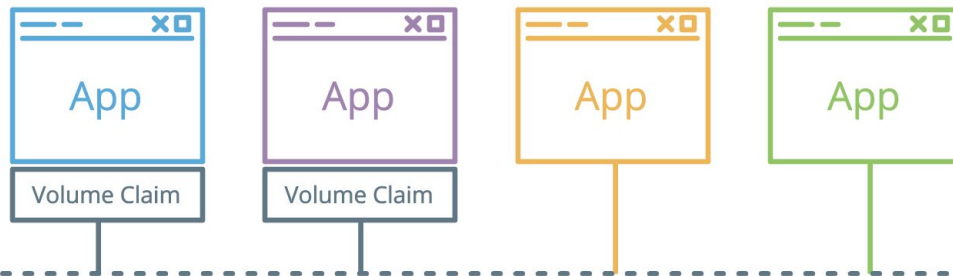
- **UI SELECTABLE**

The UI prompts the admin to select a number of nodes for OCS
Then tags and taints are set to those nodes

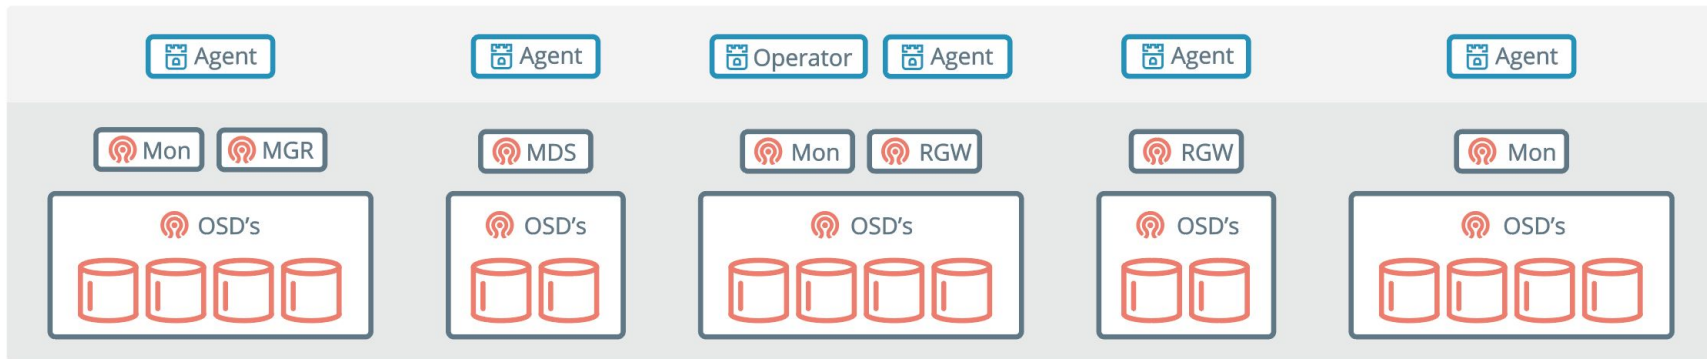
- **TAINTS**

Allow a Node to control which Pods should (or should not) be scheduled

Rook Architecture



ROOK pods



ROOK CEPH STORAGE OPERATOR

Rook Operator
for Ceph

ROOK FEATURES

- HIGH AVAILABILITY AND RESILIENCY
- DATA PROTECTION
- CONSISTENT STORAGE PLATFORM ACROSS HYBRID CLOUD
- BLOCK, FILE & OBJECT STORAGE SERVICE
- SCALE UP/DOWN
- DASHBOARD



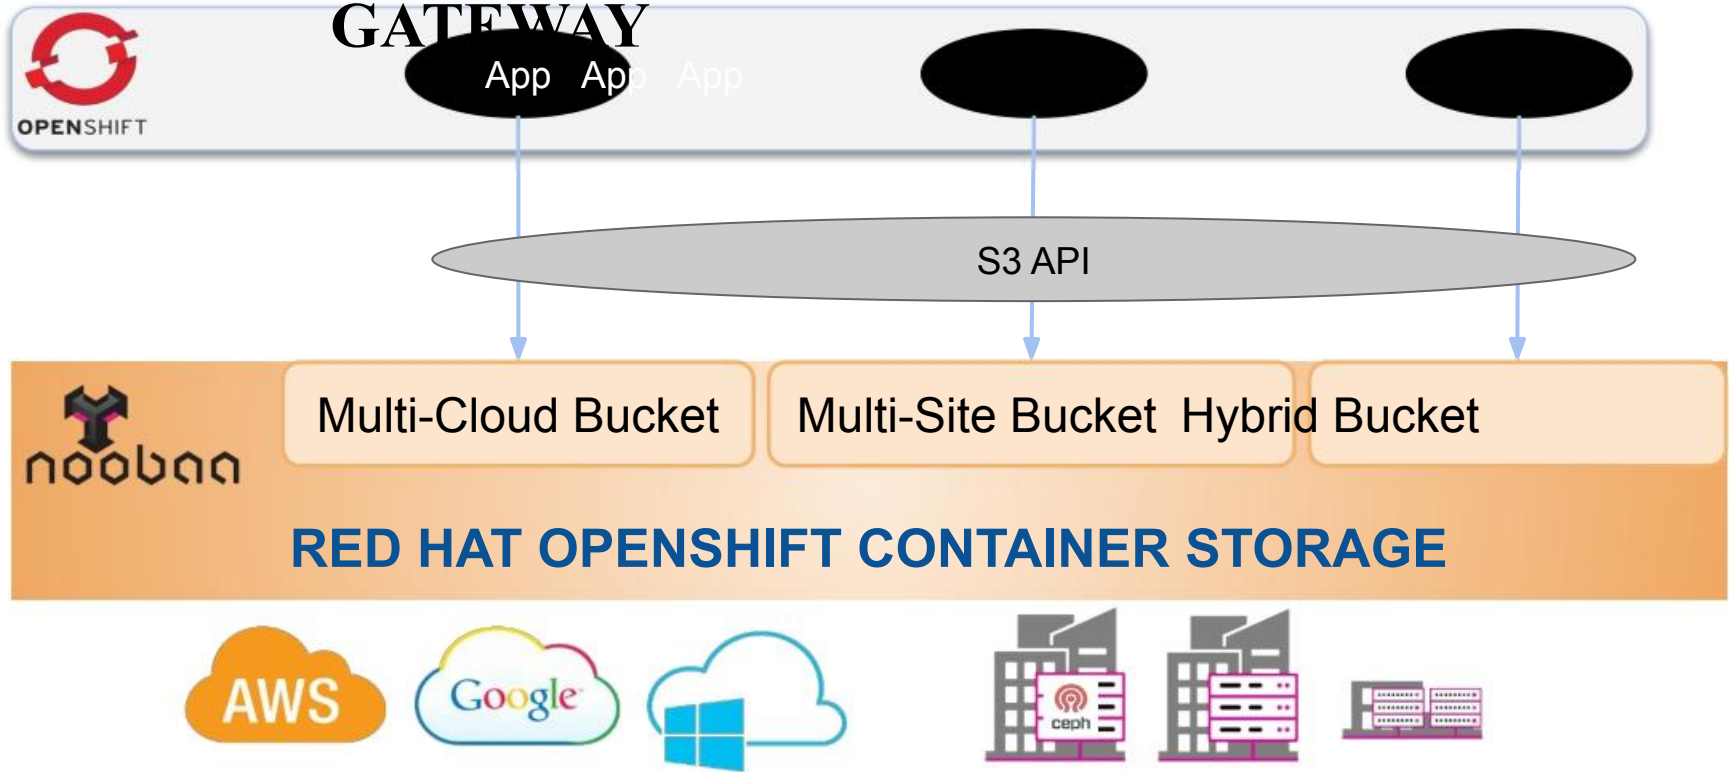
NOOBAA.IO

MULTI CLOUD GATEWAY (NOOBAA) OPERATOR



- **OCS MULTI CLOUD GATEWAY (NOOBAA) :** starting with OCS 4.2
NooBaa provides a consistent S3 endpoint across different infrastructures (AWS, Azure, GCP, Bare Metal, VMware)
- **MCG OPERATOR** WILL CREATE AND RECONCILE A NOOBAA SYSTEM
- **FUNCTIONALITY:** Multi Cloud Object Gateway: Active/Active read/write across different clouds
- **WHAT MCG PROVIDES :**
S3 OBJECT STORE SERVICE ABSTRACTION
DATA PLACEMENT POLICIES THAT ENABLE FOR HYBRID AND MULTI CLOUD
SOFTWARE-DRIVEN INFRASTRUCTURE
AGILITY

MULTI-CLOUD OBJECT GATEWAY



ACTIVE - ACTIVE MULTI CLOUD - READ/WRITE

OCS 4.2

OpenShift Container Storage

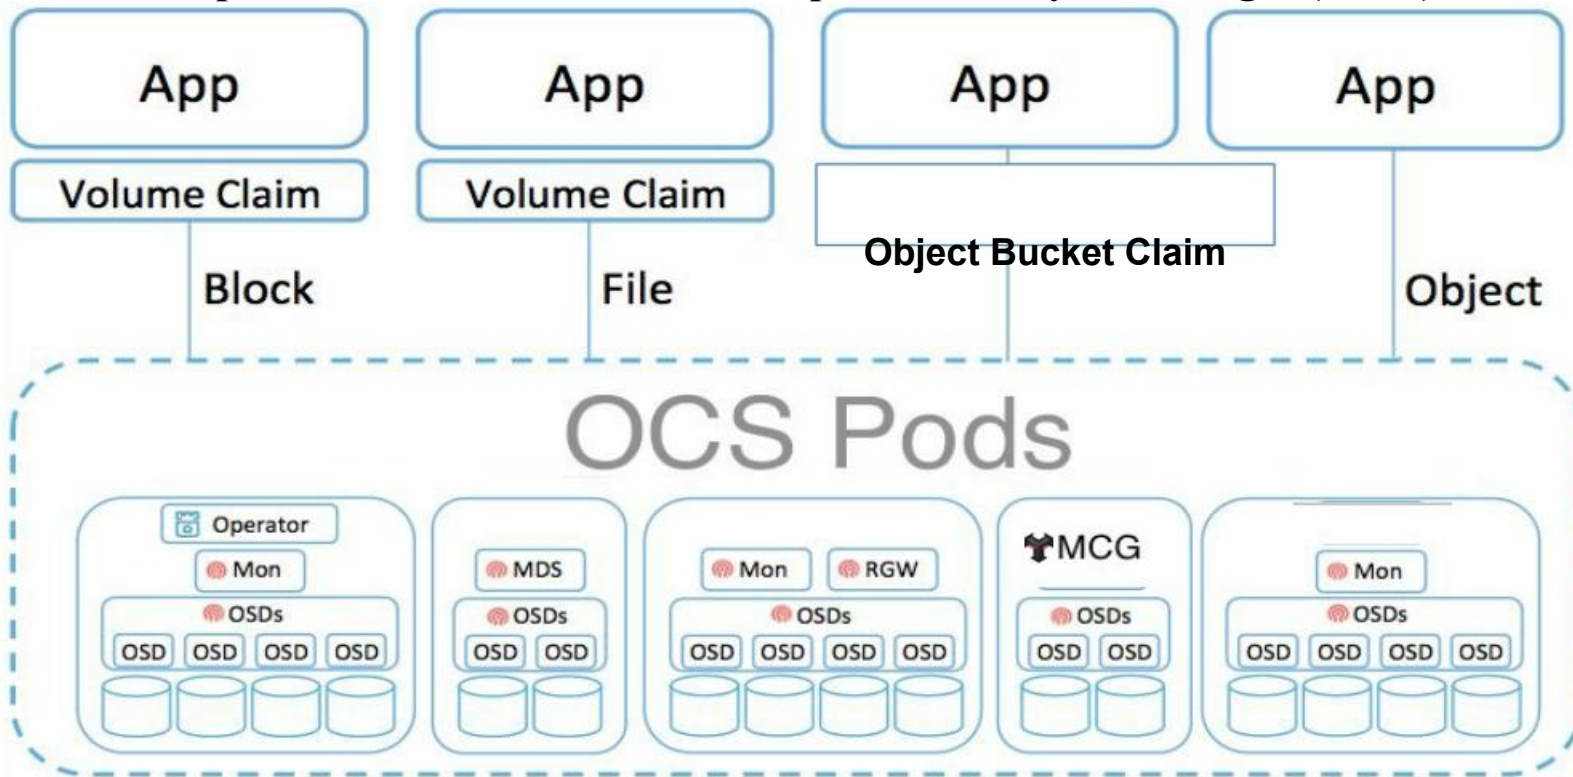
- Executive Summary -

THE FACTS - SUMMARY

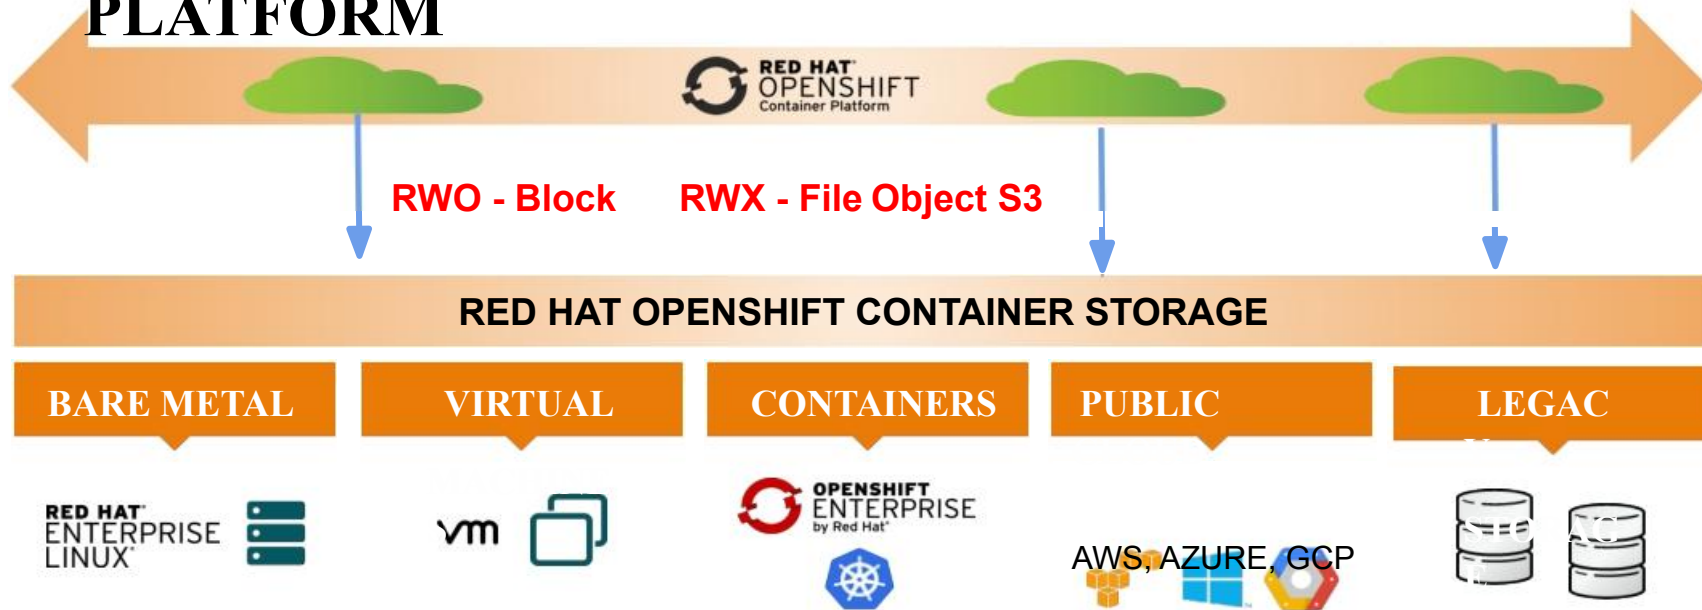
- CONTAINERS ARE SYSTEM PROCESSES AND ARE “VOLATILE” BY DEFAULT
- CONTAINERS THEREFORE NEED **PERSISTENT STORAGE & SMART ORCHESTRATION**
- **OCP 4.2** USES **OPERATORS** TO MANAGE THE ENTIRE OCP CLUSTER
- RHOCS NOW CHANGES AND WILL NOW USE **CEPH** AND **NOOBAA** ‘UNDER THE HOOD’
- **OCS 4.2** IS GA @ STARTING **OCP 4.2**
- **OCS 4.2** WILL OFFER **FILE, BLOCK AND OBJECT** STORAGE

OCS 4.x Operator Install, Upgrade, Expansion

OCS Operator based on Rook.io with Operator Lifecycle Manager (OLM)



COMPLETE STORAGE FOR CONTAINER PLATFORM

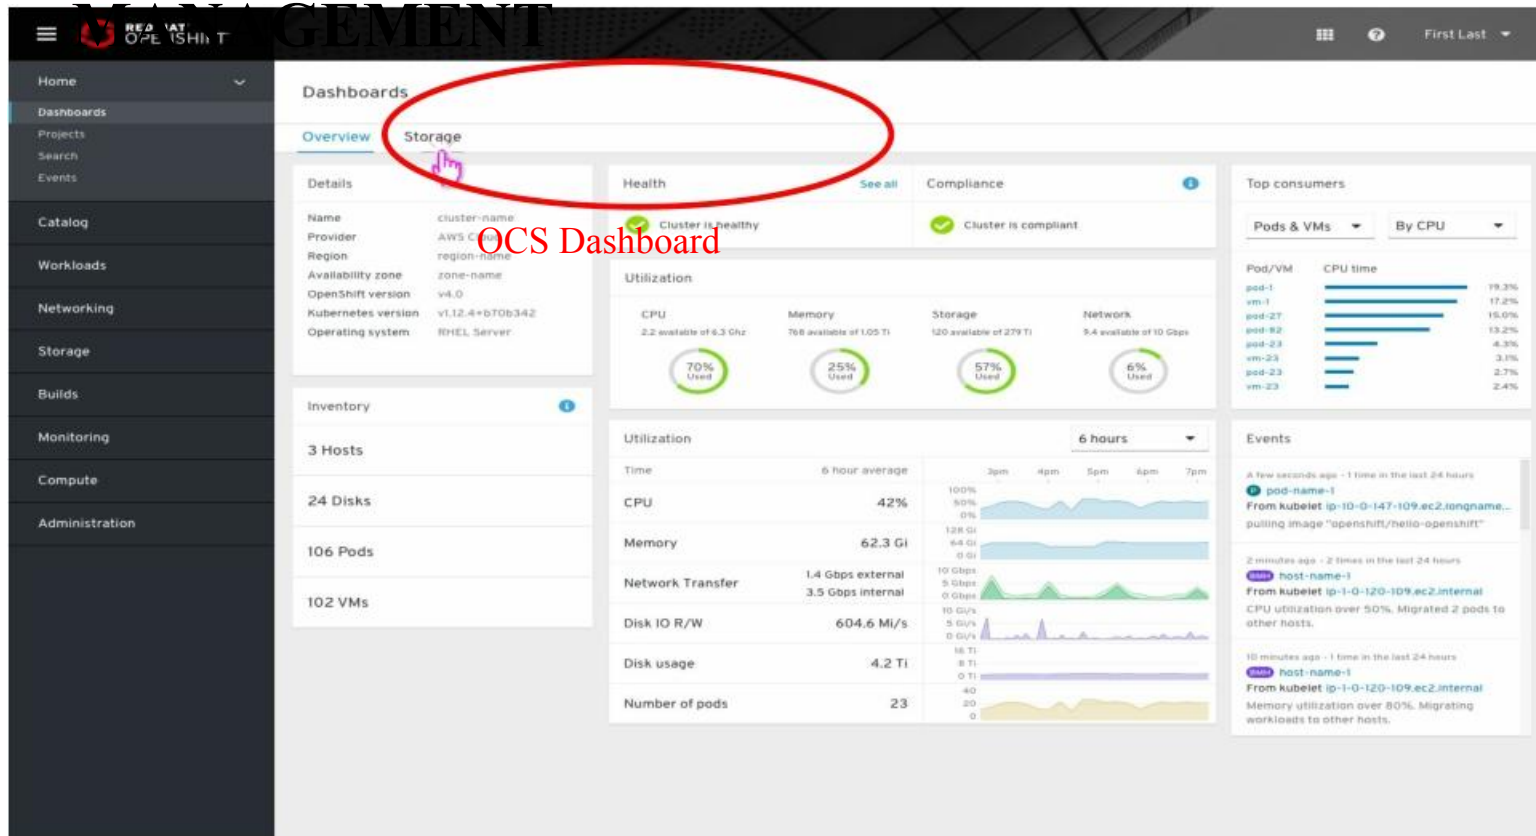


Provides Persistent Storage & Smart Orchestration for All Apps and Container infrastructure Services in their native interfaces

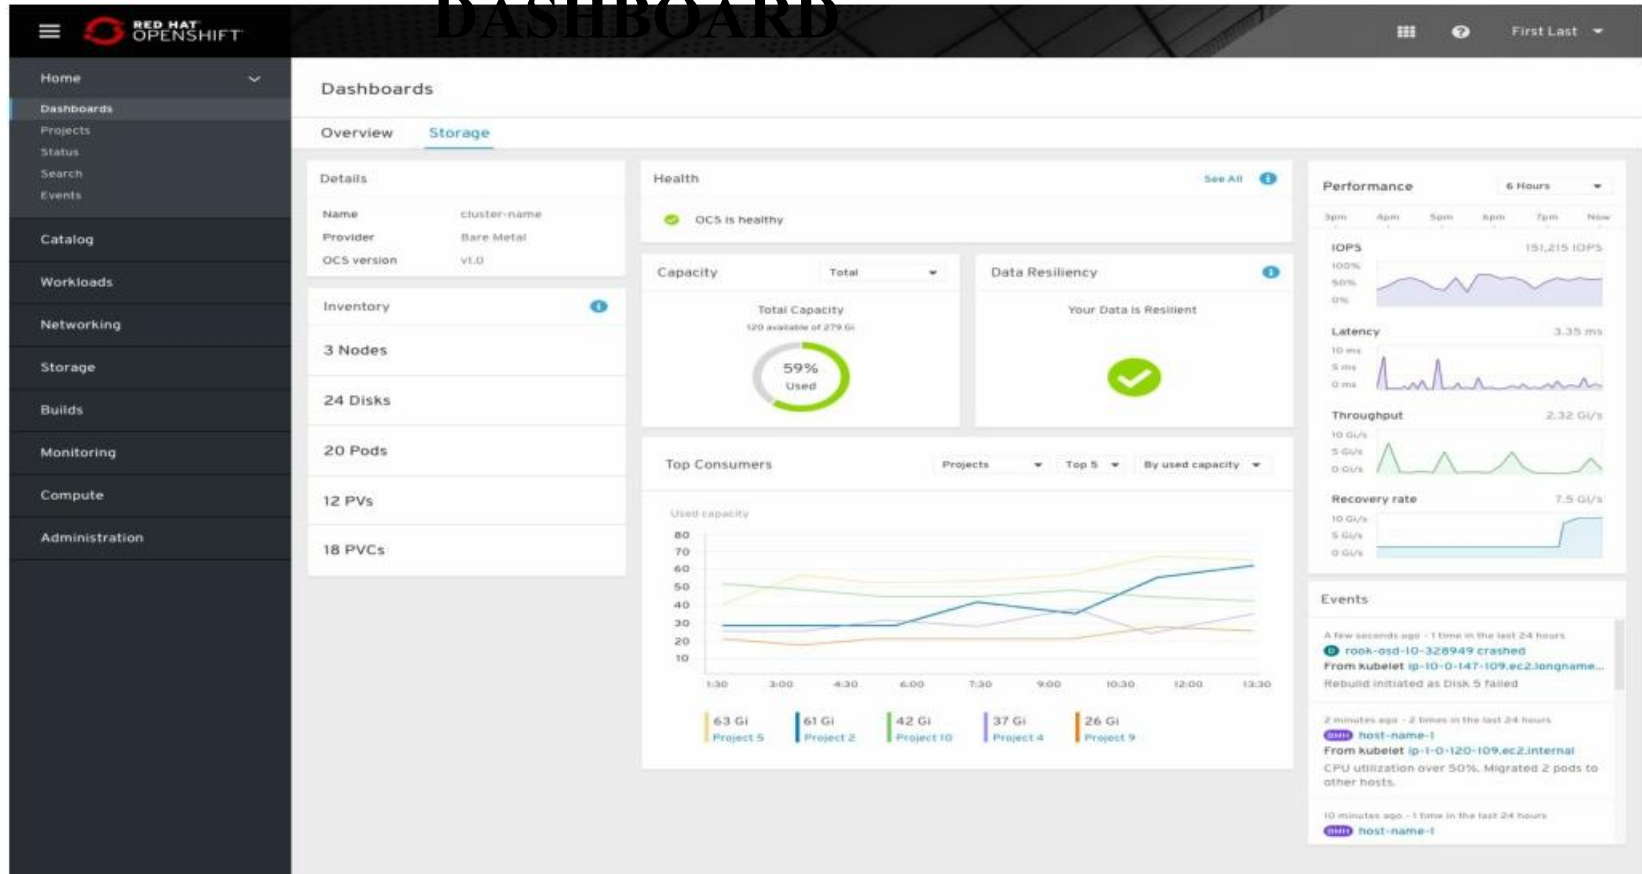
OCS 4.x Operator Driven Install from OperatorHub

The screenshot shows the Red Hat OpenShift console interface. On the left is a sidebar with navigation links: Home, Dashboards, Projects, Search, Events, Catalog, Workloads, Networking, Storage, Builds, Monitoring, Compute, and Administration. The main content area is titled 'Dashboards' and contains several tabs: Overview, Storage, Health, Compliance, and Top consumers. The 'Storage' tab is highlighted with a red circle and the text 'OCS Operator' in red. Below the tabs, there are sections for Details, Health, Compliance, Utilization, Inventory, and Events. The 'Details' section shows cluster information: Name (cluster-name), Provider (AWS Cloud), Region (region-name), Availability zone (zone-name), OpenShift version (v4.0), Kubernetes version (v1.12.4+b70b342), and Operating system (RHEL Server). The 'Health' section shows 'Cluster is healthy' and 'Cluster is compliant'. The 'Compliance' section shows 'Cluster is compliant'. The 'Utilization' section shows four circular gauges: CPU (70% Used), Memory (25% Used), Storage (57% Used), and Network (6% Used). The 'Inventory' section shows 3 Hosts, 24 Disks, 106 Pods, and 102 VMs. The 'Events' section shows a list of events, including 'pod-name-1' and 'host-name-1'.

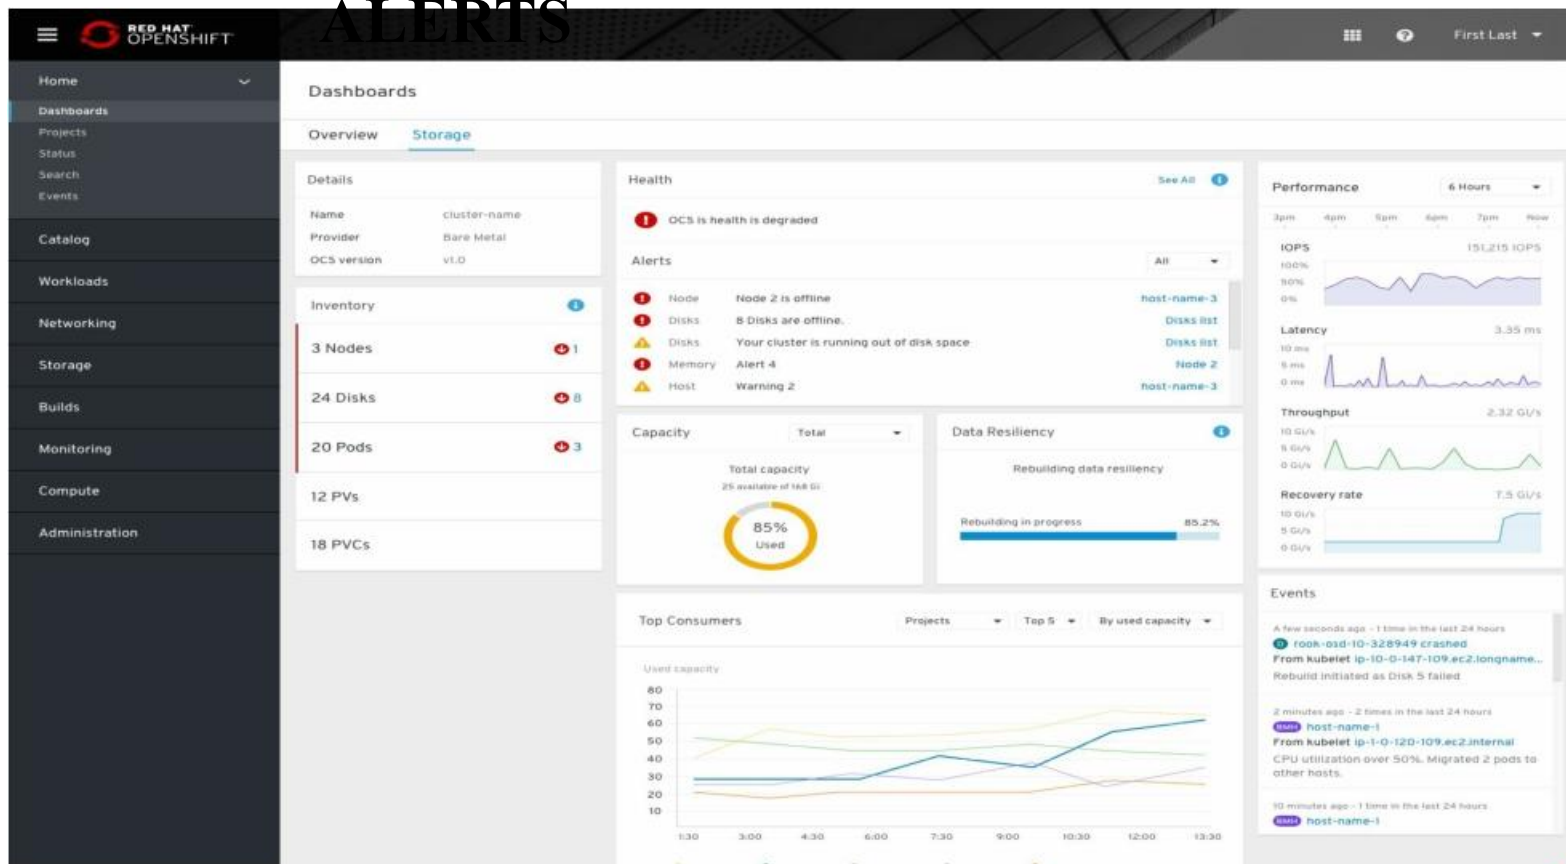
INTEGRATED MONITORING AND



OCS INTEGRATED DASHBOARD



OCS INTEGRATED DASHBOARD - ALERTS




OCS SIZING GUIDELINES

OCS SIZING - GUIDANCE

- **MINIMUM NODES #** - The MINIMUM amount of storage nodes is **3**
- **PV SUPPORTED #** - Out-of-the-Box OCS4.2 supports up to **1500** PVs
- **ADDITIONAL NODES** - Each additional node enables for **+500** PVs
- **MAXIMUM NODES #** - The MAXIMUM number of nodes in a cluster is **10**
- **MAXIMUM PV #** - The MAXIMUM number of PVs can scale to **5000** PVs

Cloud-like Simplicity, Everywhere

Full-stack automated operations across any on-premises,
cloud, or hybrid infrastructure



Thank you

Red Hat is the world's leading provider of enterprise open source software solutions. Award-winning support, training, and consulting services make Red Hat a trusted adviser to the Fortune 500.



linkedin.com/company/red-hat



youtube.com/user/RedHatVideos



facebook.com/redhatinc



twitter.com/RedHat