OpenShift e Container Storage

KUBERNETES ENTERPRISE PER LE GRANDI IDEE

PIERLUIGI SFORZA
Solution Architect
psforza@redhat.com

#RedHatOSD

… so you want to do containers and Kubernetes?

RED HAT HAS BEEN A KUBERNETES LEADER SINCE DAY 1

K8S 1.0

K8S 1.0

We were very lucky to be joined early on by the very capable OpenShift team … without their perspective and contributions, I don't think we would be standing here today

Brendan Burns, co-creator of Kubernetes

K8S 1.6

K8S 1.8

K8S 1.9

#RedHatOSD

redhat.

OPENSHIFT IS KUBERNETES
FOR THE ENTERPRISE

Security fixes
100s of defect and performance fixes
200+ validated integrations
Middleware integrations
(container images, storage, networking, cloud services, etc)
9 year enterprise lifecycle management
Certified Kubernetes

Kubernetes Workloads

MANAGE YOUR APPLICATIONS

PIERLUIGI SFORZA
Solution Architect
psforza@redhat.com

ALESSANDRO ARRICHIELLO
Solution Architect
ale@redhat.com

#RedHatOSD

CONTROLLERS MATTERS!

Different types of applications (stateful, stateless, batch, agent, ...) require different orchestrator behaviors

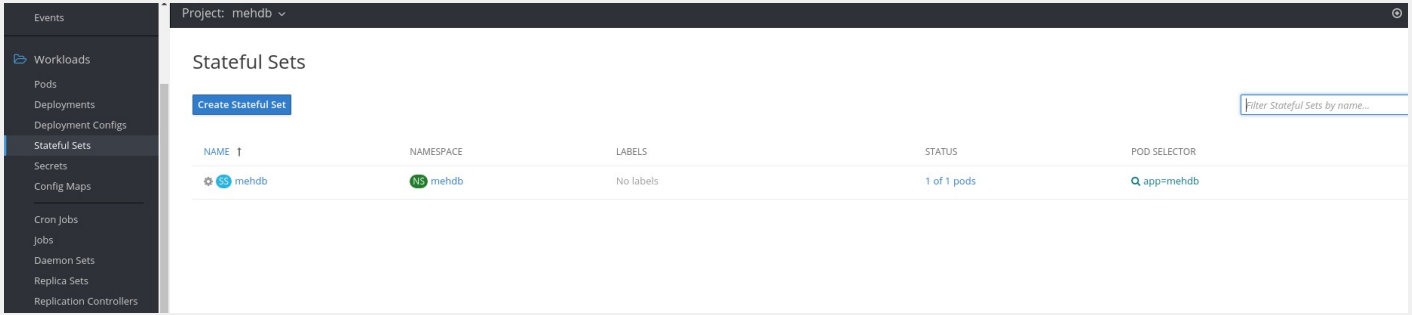


```
func NewControllerInitializers(loopMode ControllerLoopMode) map[strin
        controllers := map[string]InitFunc{}
        controllers["endpoint"] = startEndpointController
        controllers["replicationcontroller"] = startReplicationControl
```

Main controller types:

- Replica Sets
- Stateful Sets
- Daemon Sets
- Jobs (OneTime, Cron)

WHAT IS A POD?

CONTAINERS ARE WRAPPED IN PODS WHICH ARE UNITS OF DEPLOYMENT AND MANAGEMENT

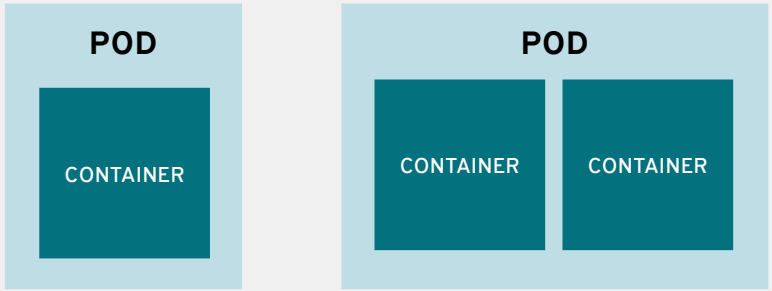

CONTROLLER & CONTROLLER-MANAGER

- The **controller-manager** is the Master's component that manage the controllers
- A **controller** is a loop that governs the status of kubernetes resources (such as pods) in order to bring it from the current state to the desired state
- Controllers react to **kubernetes events** and define **how resources should be orchestrated**

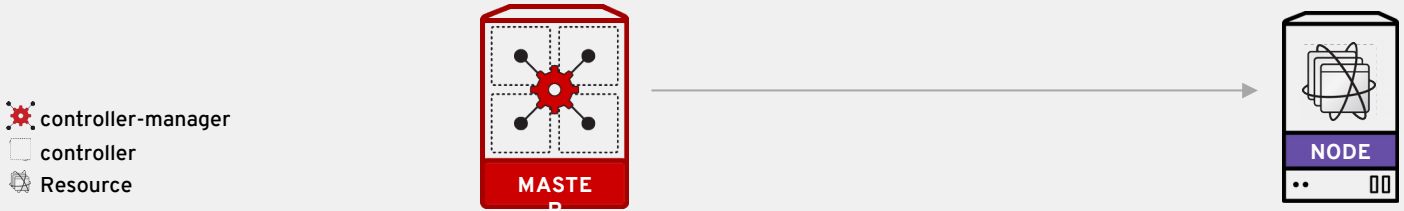

CONTROLLER & CONTROLLER-MANAGER

- The **controller-manager** is the Master's component that manage the controllers
- A **controller** is a loop that governs the status of kubernetes resources (such as pods) in order to bring it from the current state to the desired state
- Controllers react to **kubernetes events** and define **how resources should be orchestrated**

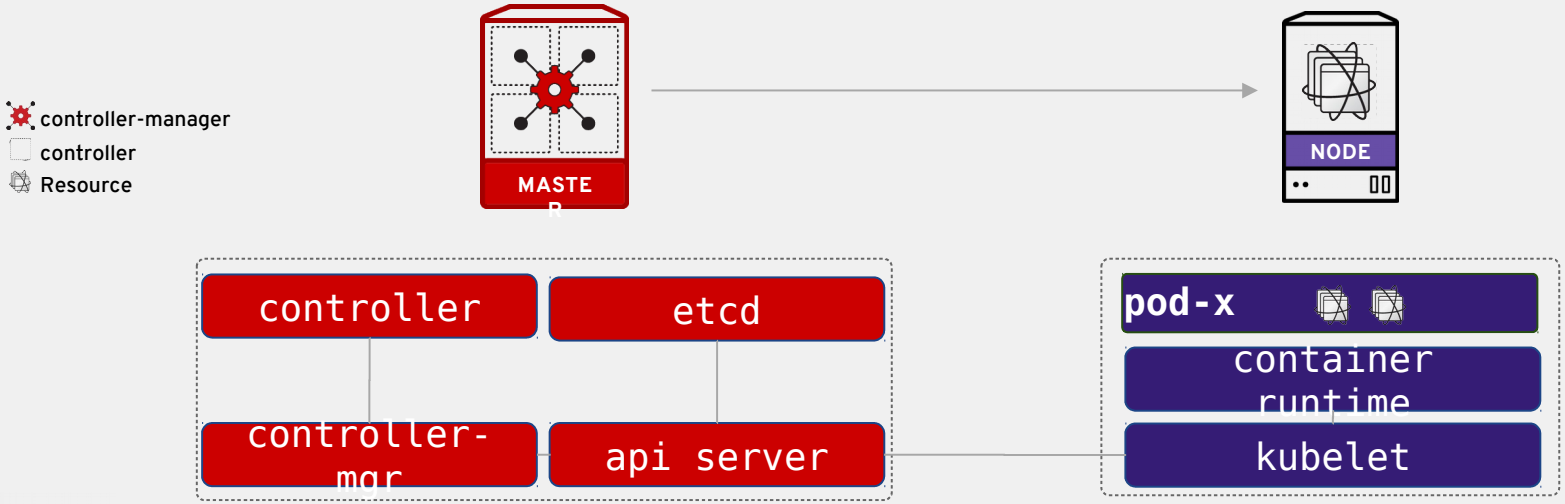

DEPLOYMENT AND REPLICASET

- A Deployment controller provides declarative updates for Pods and ReplicaSets
- ReplicaSet controller ensures that a specified number of pod replicas are running at any given time
- Recommended to run stateless application

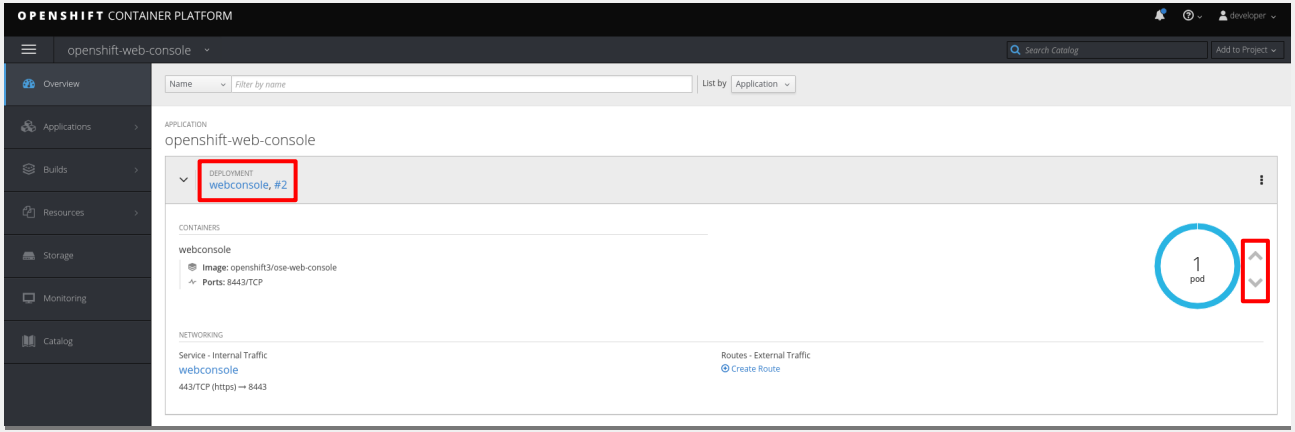

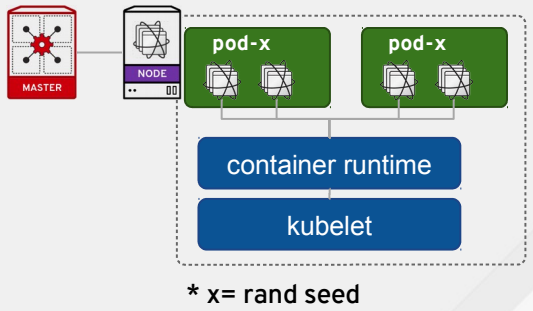

* x= rand seed

STATEFULSET

- A stateful set ensure
  - stabile resource allocation such as name and storage
  - ordered, graceful deployment, scaling up and termination

- ideal for highly available workloads in a “clustered mode”

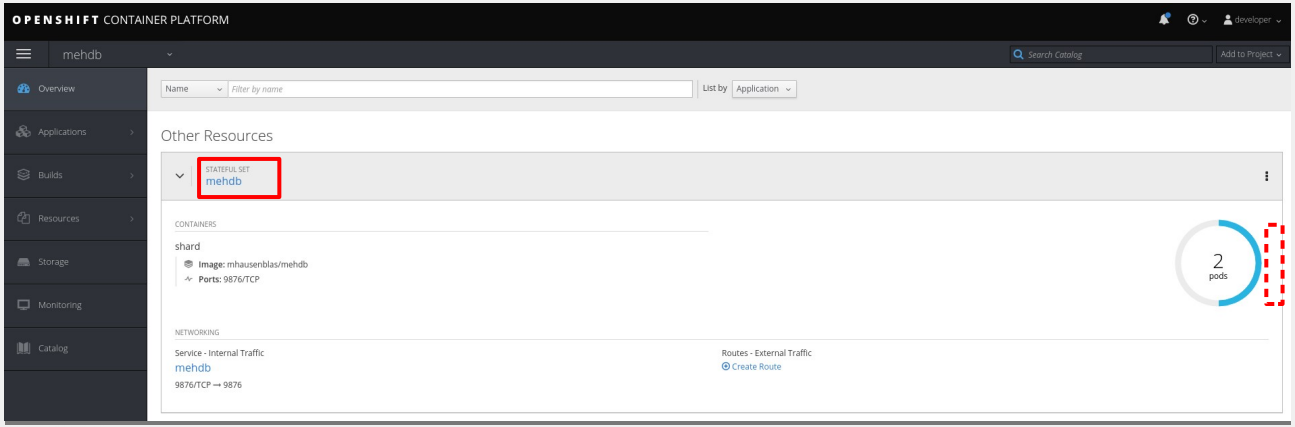

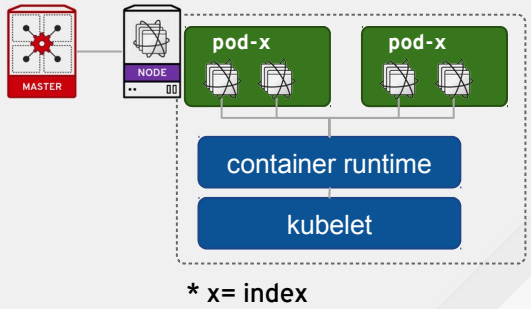

DAEMONSET

- A daemon set ensure to have **just 1 copy** of a pod on every node
- Daemon set is useful for: Logging Aggregators, Monitoring, Load Balancers / Reverse Proxies / API Gateways, single host batch…

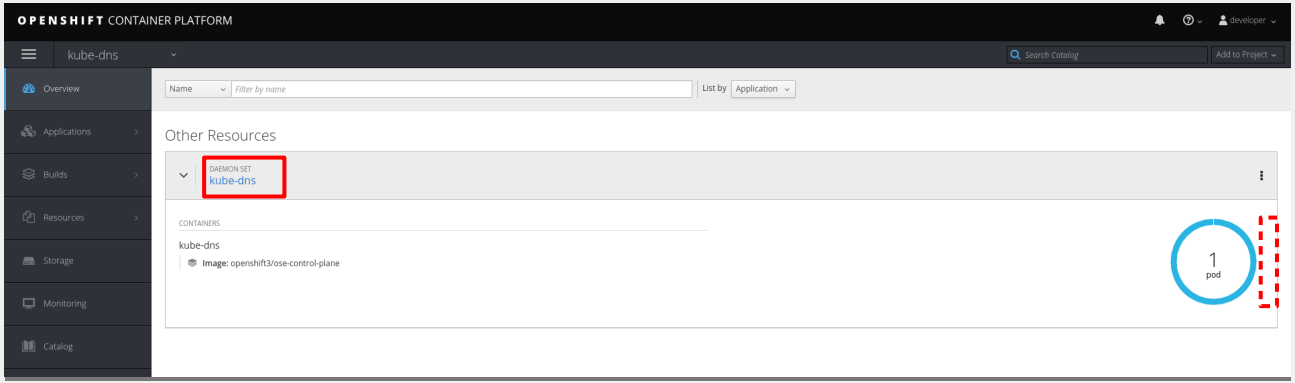

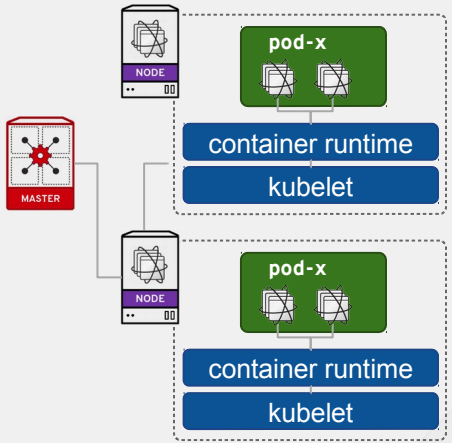

* x= available node count

REFERENCE ARCHITECTURE
FOR ENTERPRISE KUBERNETES

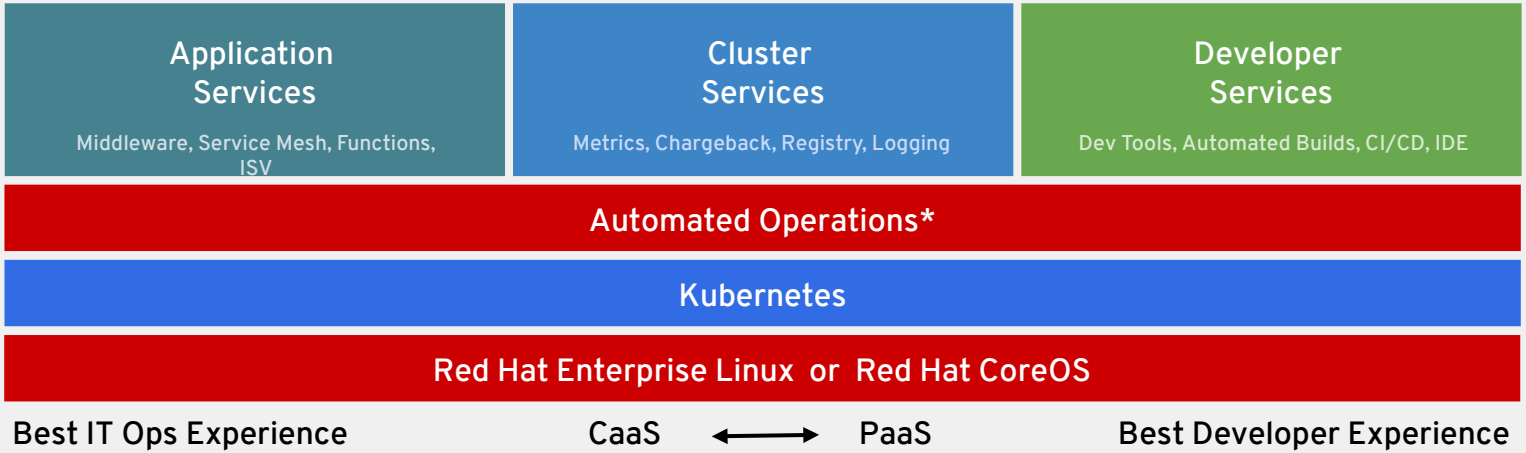

*coming soon with OCP 4.0
(targeted for GA Dec 2018)

Istio Service Mesh

FOR SERVICE-TO-SERVICE COMMUNICATIONS

NATALE VINTO
Specialist Solution Architect
nvinto@redhat.com

#RedHatOSD

OPENSHIFT SERVICE MESH: ISTIO*

Istio makes it easy to create a network of deployed services with load balancing, service-to-service authentication, monitoring, and more, helping to avoid operational nightmares.

POLICY

Grants the ability to write policy that applies to all applications and is not language specific

ROUTING

Allows for the control of routing flows

TELEMETRY

Provides the observability needed to manage microservices, such as how services are invoked, communication flows, and points of latency

* Technology Preview

ISTIO COMPANION: KIALI & JAEGER

Kiali and Jaeger make the perfect companion for Istio Service Mesh

VISUALIZATION

Kiali works with Istio to visualize the service mesh topology, features like circuit breakers or request rates.

TRACING

Kiali includes **Jaeger** Tracing, which provides distributed tracing out of the box.

SERVICE MESH ARCHITECTURE

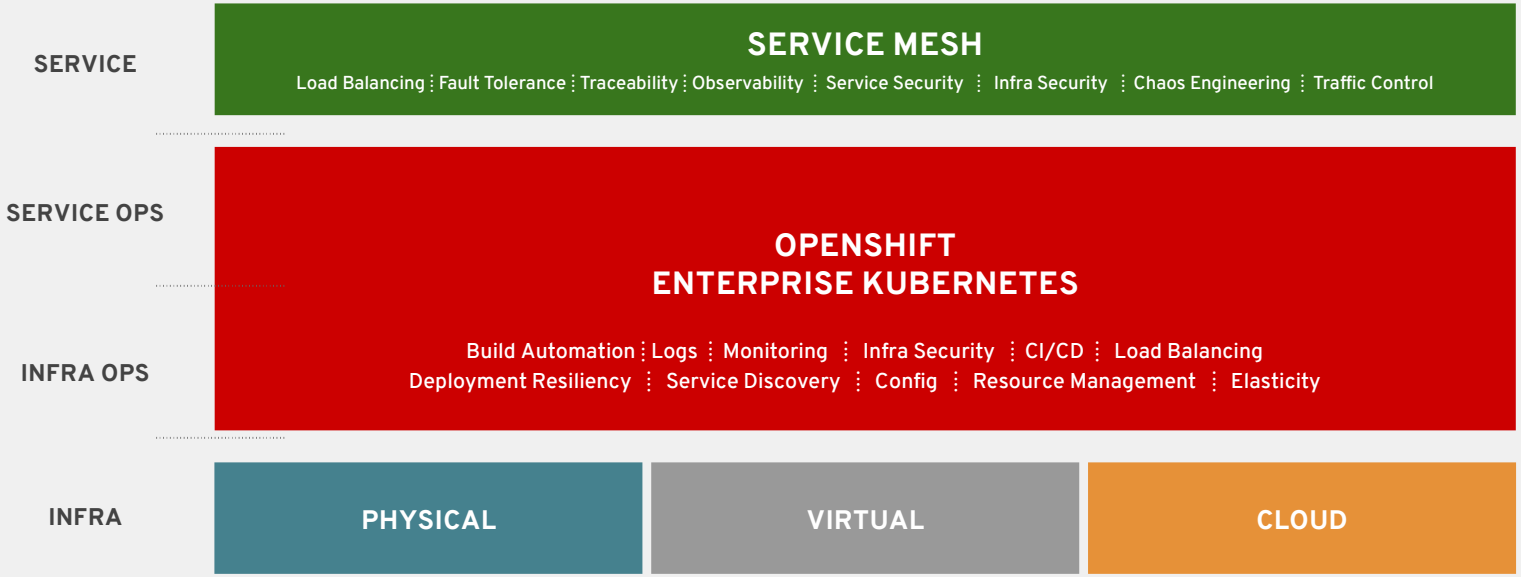

MICROSERVICES EVOLUTION

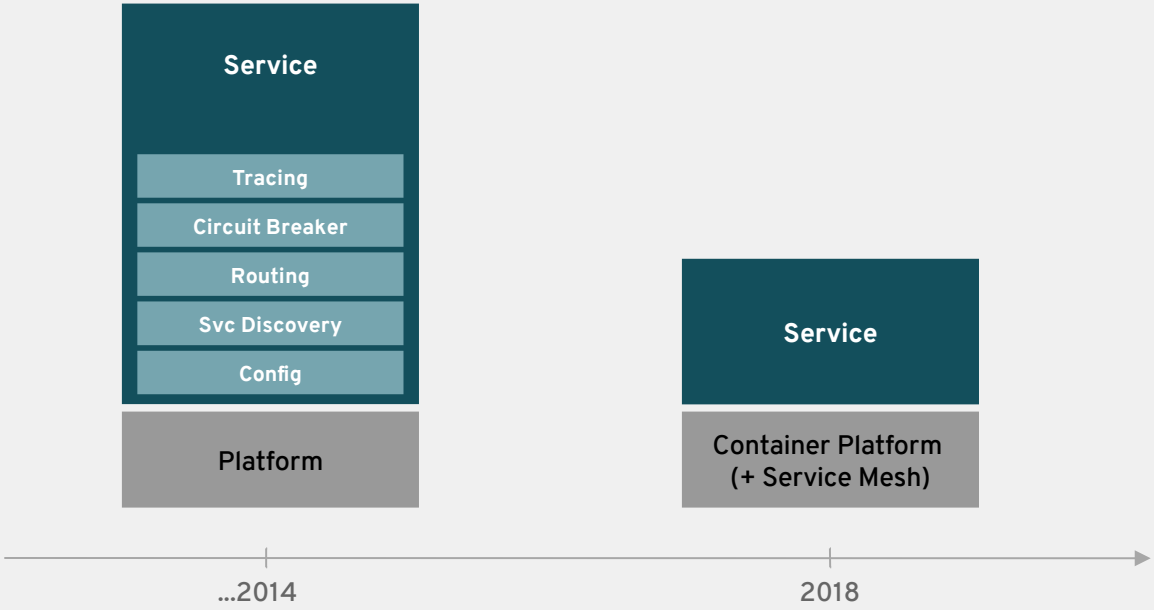

SERVICE MESH ARCHITECTURE

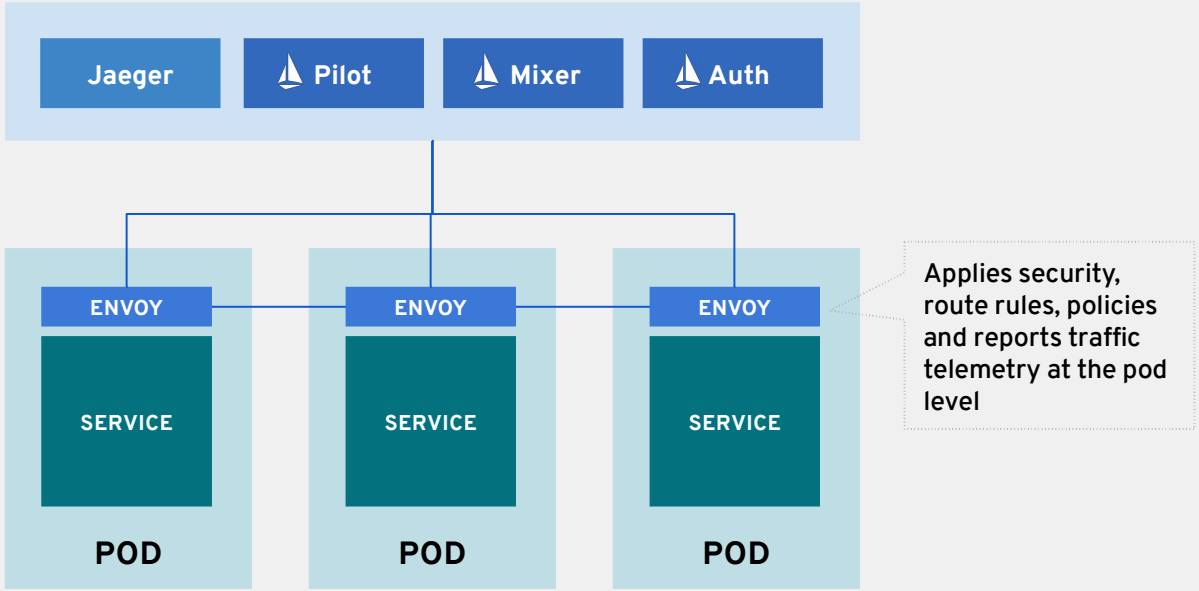

CIRCUIT BREAKERS WITH ISTIO

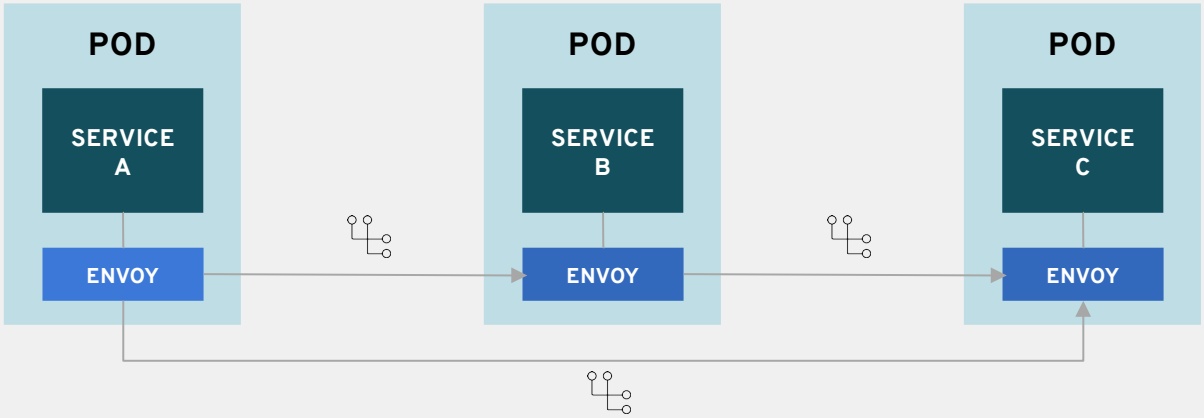

transparent to the services

SECURE COMMUNICATION WITH ISTIO

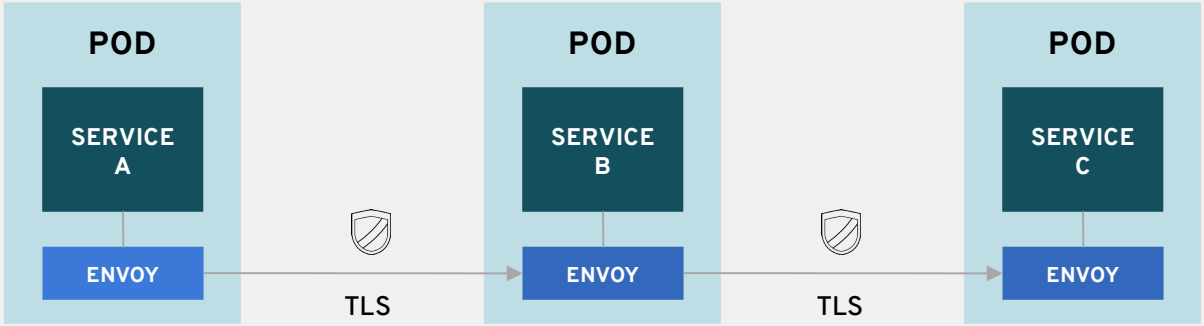

mutual TLS authentication, transparent to the services

DISTRIBUTED TRACING WITH ISTIO & JAEGER

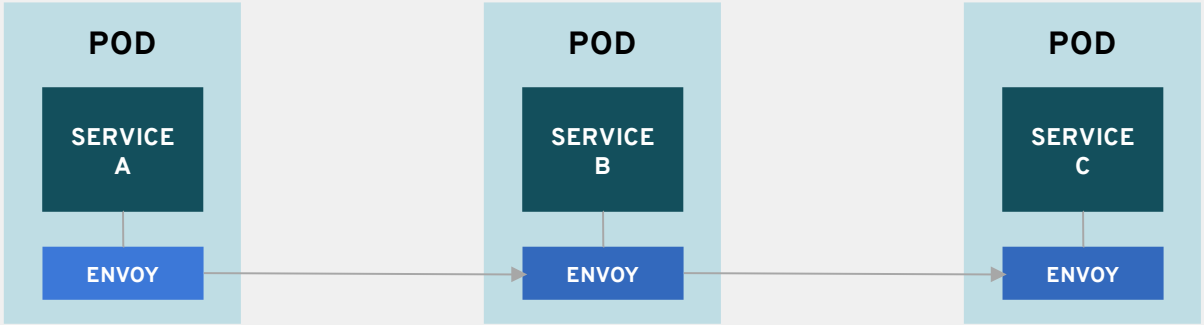

discovers service relationships and process times, transparent to the services

DEMO TIME:
Istio Internals

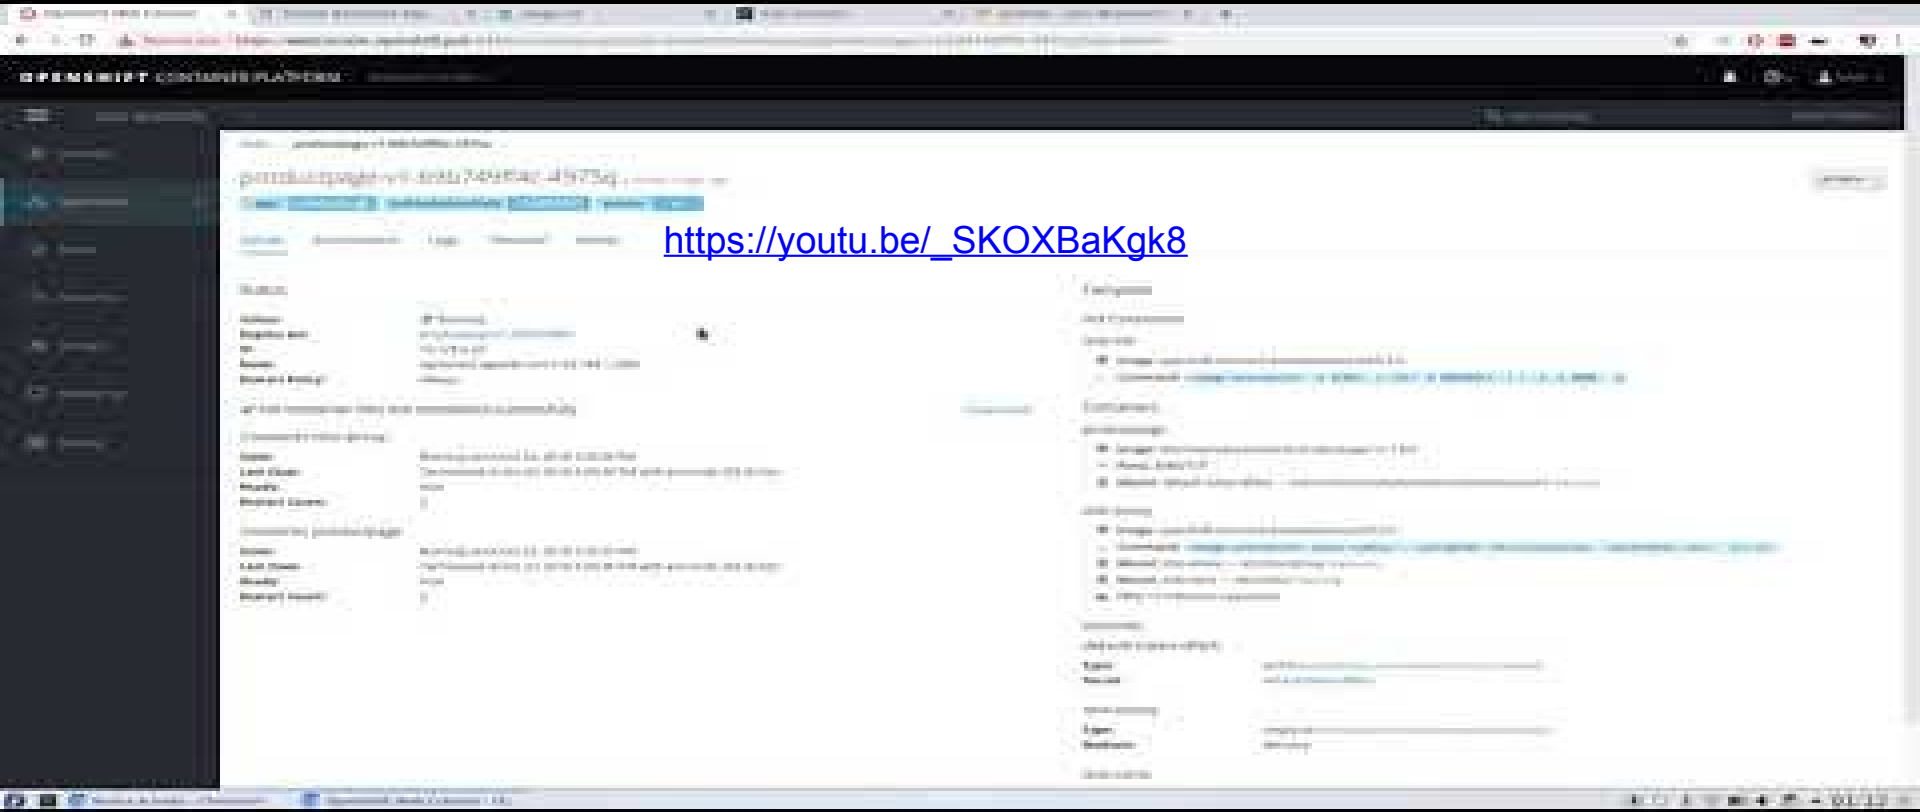
https://youtu.be/_SKOXBaKgk8

Prometheus Cluster Monitoring

PROVIDING ALERTS ALSO FOR OPENSHIFT CONTAINER STORAGE

NATALE VINTO
Specialist Solution Architect
nvinto@redhat.com

CARLOS TORRES
Specialist Solution Architect
ctorres@redhat.com

#RedHatOSD

COMPREHENSIVE MONITORING SUITE

The stack includes three distinct UIs:

- **Alertmanager** UI to manage alerts which have been fired
- **Prometheus** UI for querying and plotting any metrics
- **Grafana** to browse cluster-level dashboards

All UIs are accessible directly via the new admin console under the “Monitoring” menu.

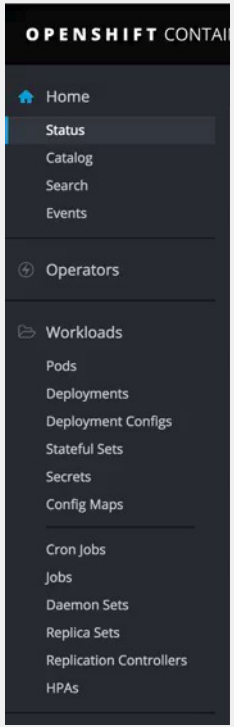

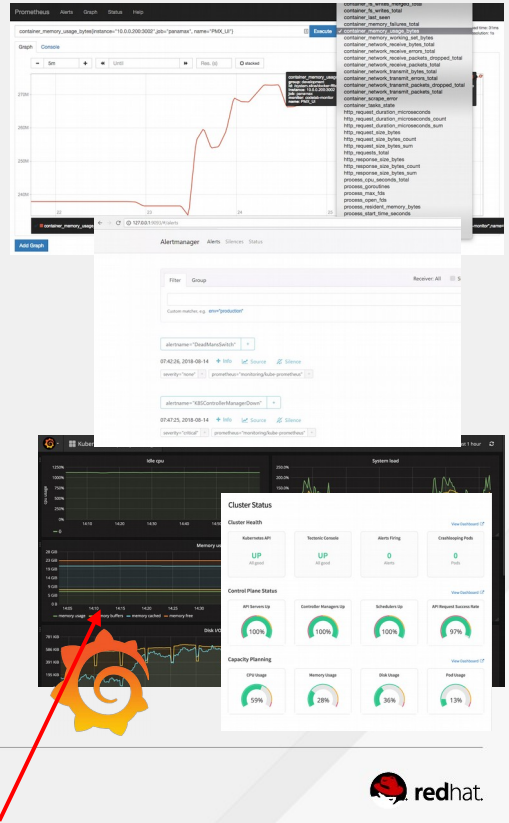

DEMO TIME:
Cluster Console - EventFeed

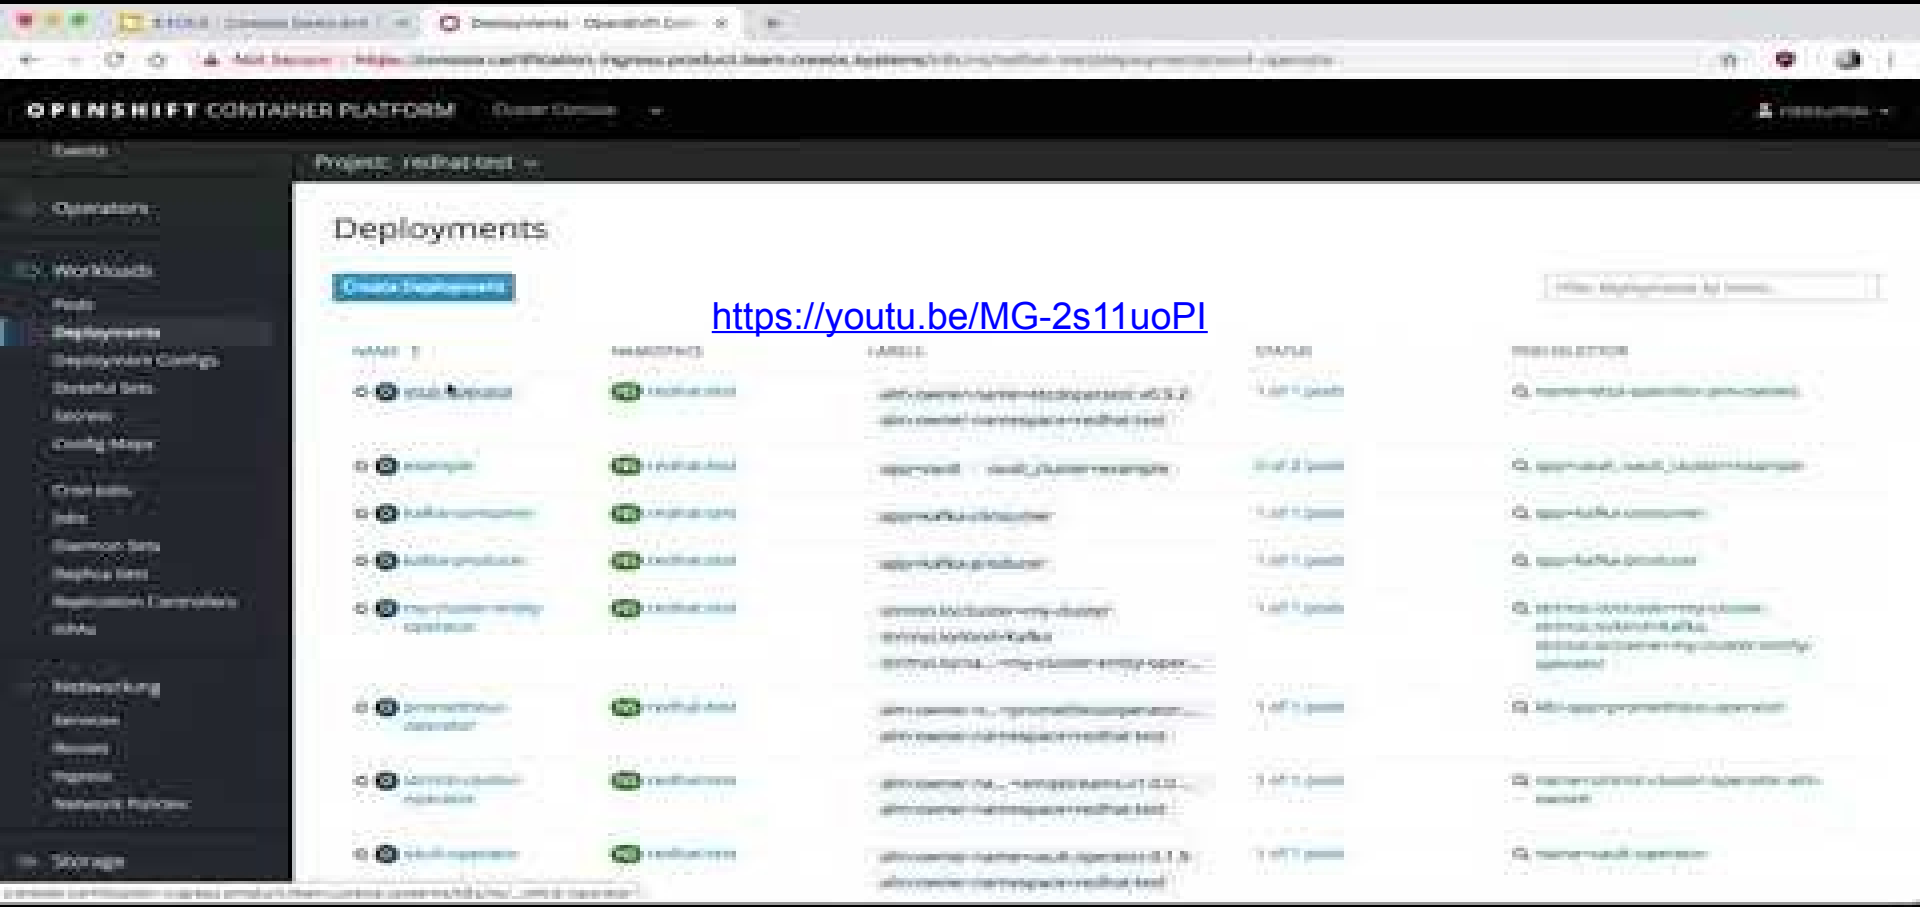
OPENSHIFT CONTAINER PLATFORM
Project: redhat-test
Deployments
https://youtu.be/MG-2s11uoPI

DEMO TIME:
Cluster Console - Monitoring

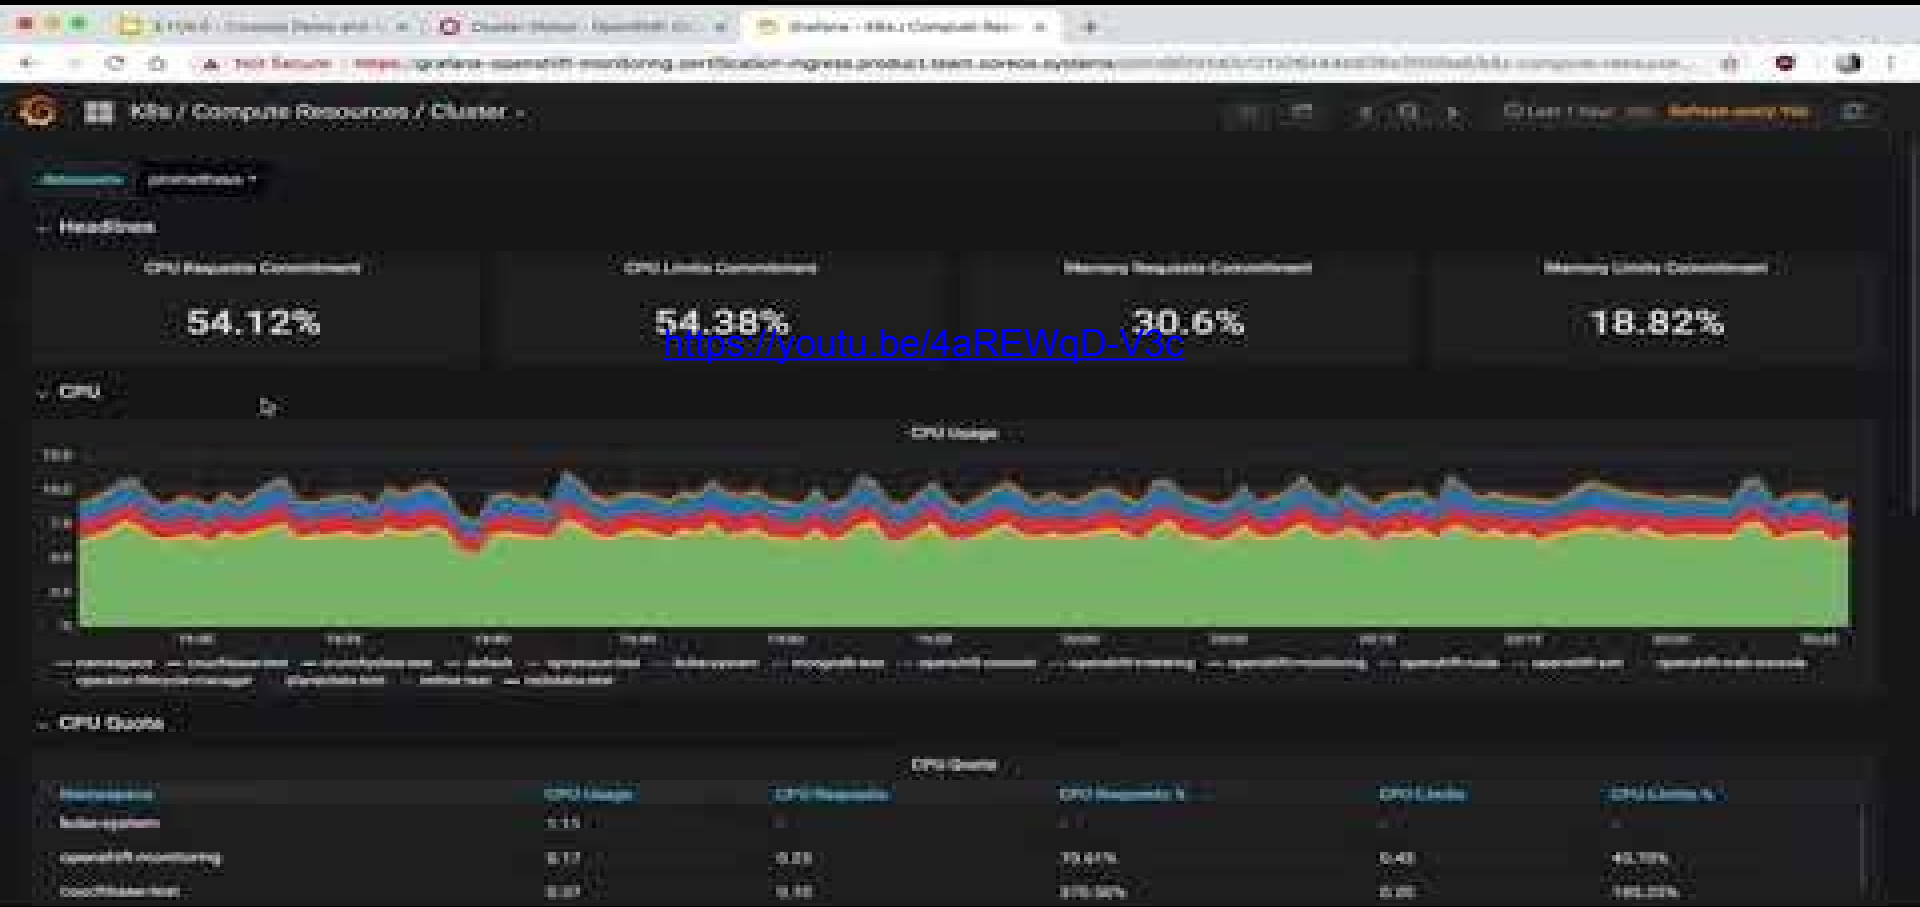
Kite / Compute Resources / Cluster -
Headlines
CPU Requests Commitment
54.12%
CPU Limits Commitment
54.38%
Memory Requests Commitment
30.6%
Memory Limits Commitment
18.82%
https://youtu.be/4aREWqD-V3c
CPU
CPU Usage
CPU Quota

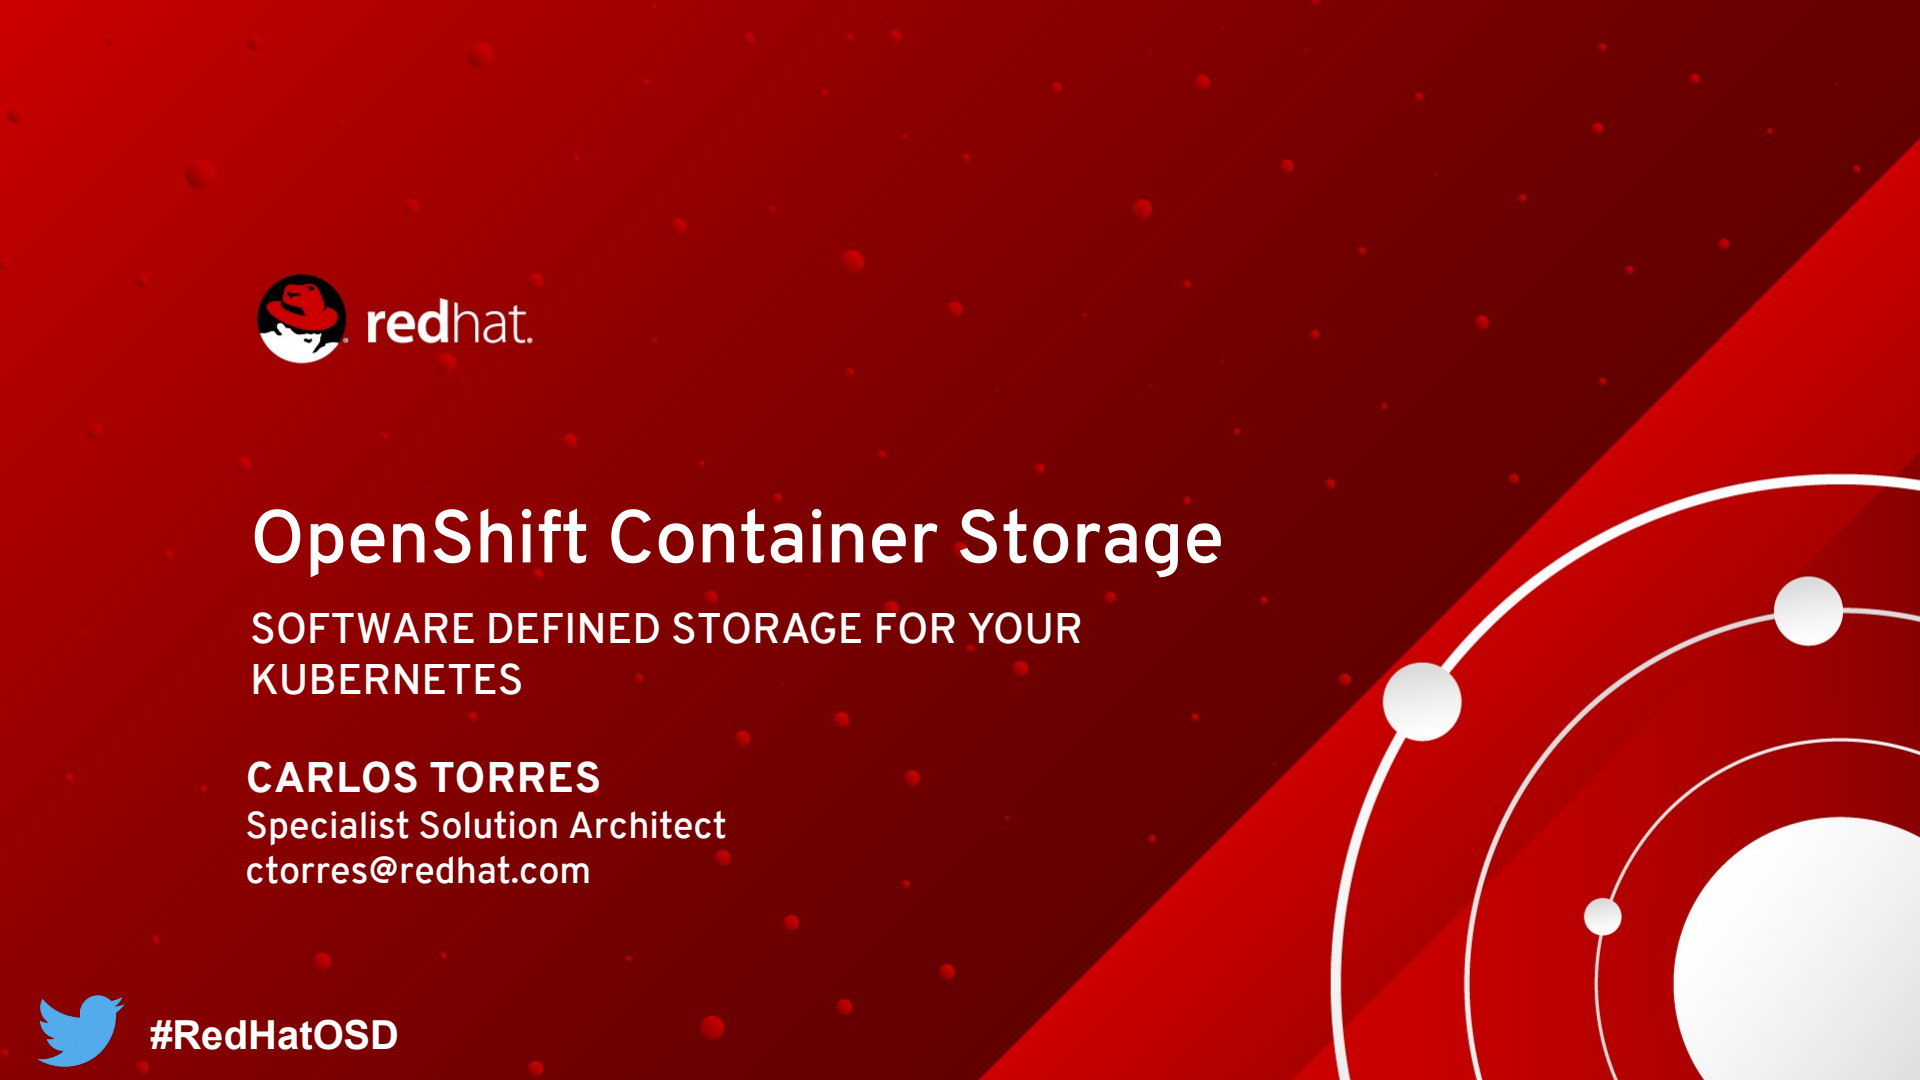
redhat.
OpenShift Container Storage
SOFTWARE DEFINED STORAGE FOR YOUR KUBERNETES
CARLOS TORRES
Specialist Solution Architect
ctorres@redhat.com
#RedHatOSD

RED HAT OPENSHIFT CONTAINER STORAGE

Flexible deployment with the same user experience and features

Converged = in containers

Persona: DevOps, App Architects

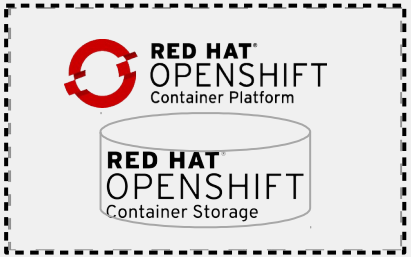

- Highly scalable, scale app+storage, start small and scale fast
- Storage life cycle managed by OCP

Independent = for containers

Persona: Storage Admins, Infrastructure Admins

- Highly scalable, independent scalability from OCP platform
- Adaptative to in-place BC/DR strategies

OPENSHIFT FULL INTEGRATION

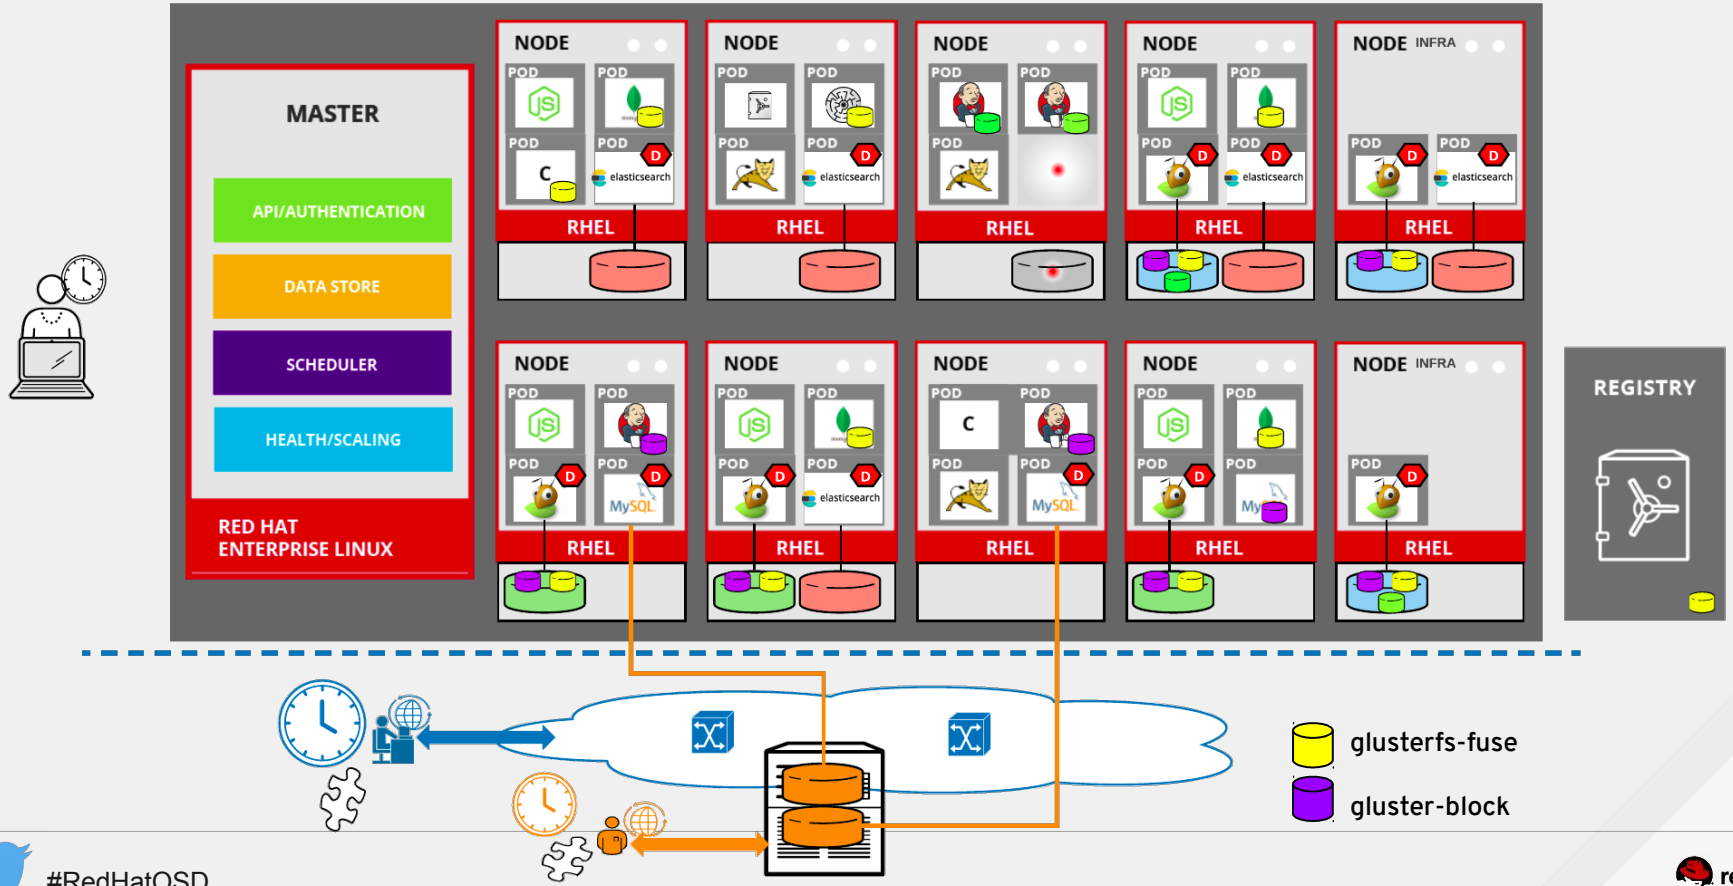

DEMO TIME:
Monitoring - OpenShift Container Storage

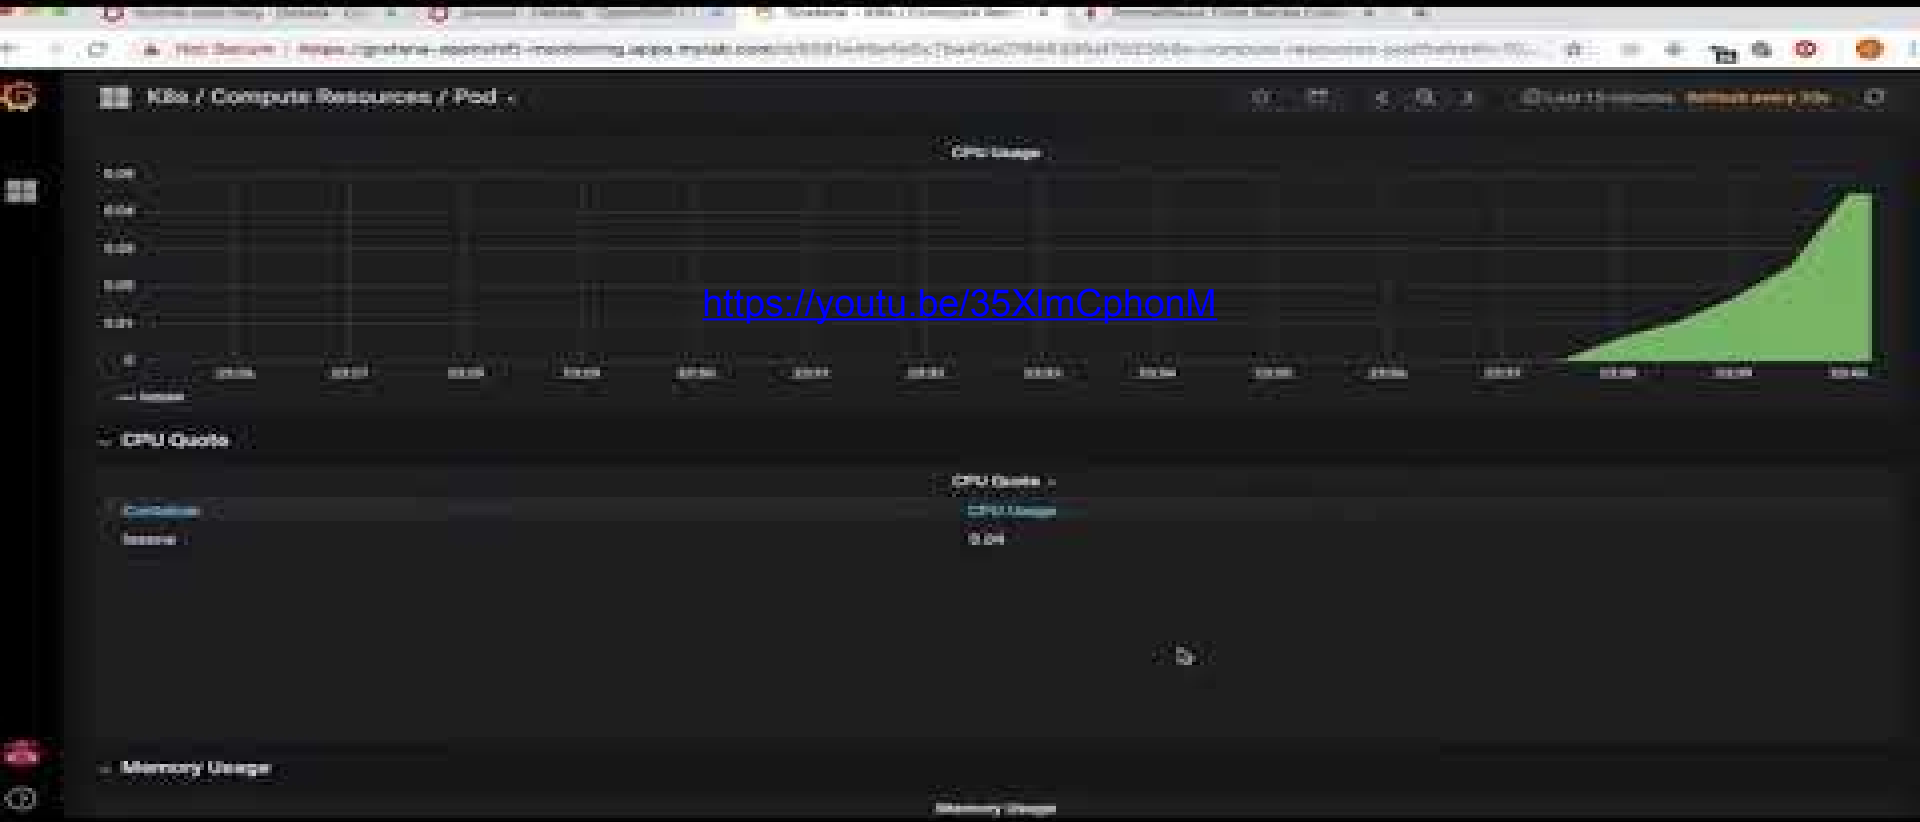
K8s / Compute Resources / Pod
CPU Usage
https://youtu.be/35XlmCphonM
CPU Quota
CPU Quota
Memory Usage

RHOCS: ANSIBLE ADVANCED DEPLOYMENT

Converged playbooks already available

Deployment workflow	Registry	Metrics	Logging	Applications
Deploying Red Hat Openshift Container Storage in Converged Mode				✔
Deploying Red Hat Openshift Container Storage in Converged Mode with Registry	✔			
Deploying Red Hat Openshift Container Storage in Converged Mode with Logging and Metrics		✔	✔	
Deploying Red Hat Openshift Container Storage in Converged mode for Applications with Registry, Logging, and Metrics	✔	✔	✔	✔

https://red.ht/2DaKPzg

Openshift Ansible Service Broker

And the road to Kubernetes Operators!

ALESSANDRO ARRICHIELLO
Solution Architect
ale@redhat.com

#RedHatOSD

WHAT IS A SERVICE BROKERAGE?

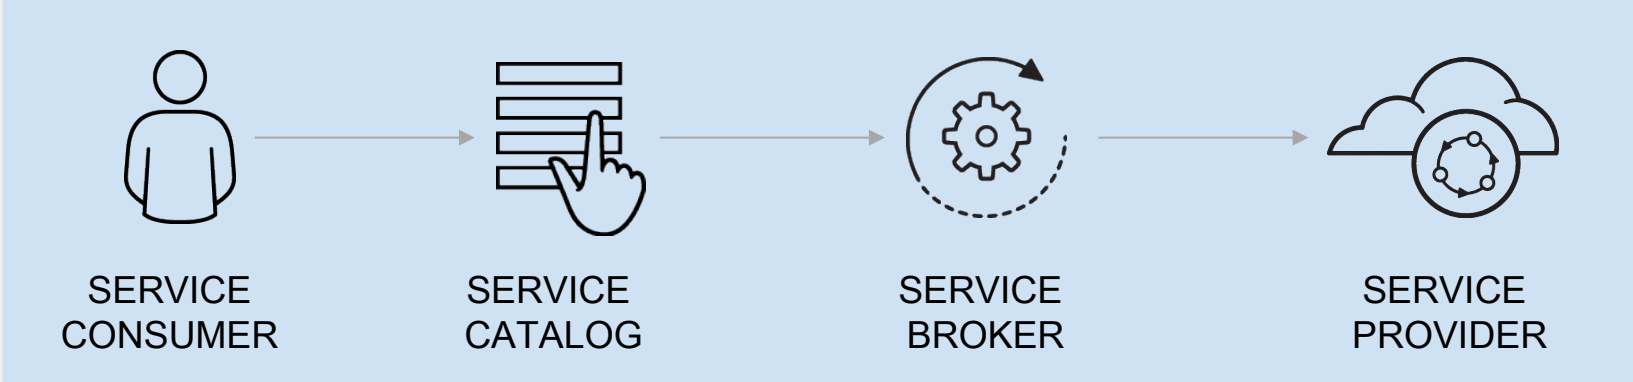

Automated, Standard and Consistent

OPEN SERVICE BROKER API™

A multi-vendor project to standardize how services are consumed on cloud-native platforms across service providers

FUJITSU Pivotal IBM redhat. Google SAP

BROKERAGE WITH OPENSHIFT

SERVICE CATALOG

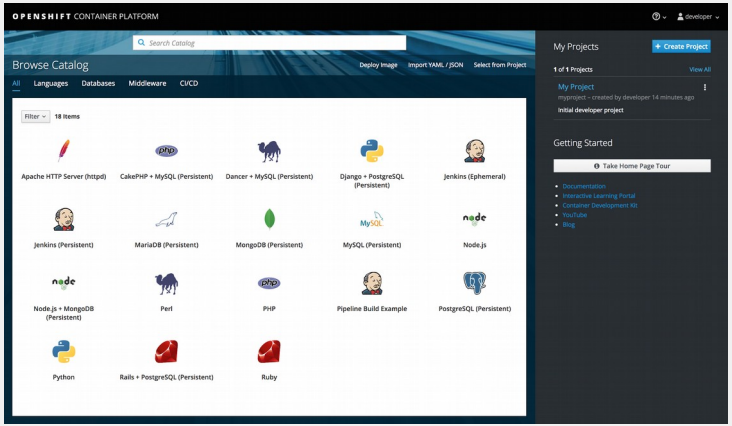

Broker

Service

OPEN SERVICE BROKER API™

Broker	Service	
OpenShift Template Broker	OPENSHIFT	*OpenShift Templates*
OpenShift Ansible Broker	ANSIBLE	*Ansible Playbook Bundles*
AWS Service Broker	AWS	*AWS Services*
Other Service Brokers	OTHER COMPATIBLE SERVICES	*Other Services*

OPENSHIFT ANSIBLE BROKER

Anything you can do with Ansible, you can do with the Ansible Broker

- Use Ansible on OpenShift to
  - Deploy containerized applications
  - Manage external components (e.g. Oracle database)
  - Provision cloud services (e.g. AWS RDS)
  - Orchestrate multi-service solutions
  - Manage dependencies or other logics on deployments (e.g. database initialization)

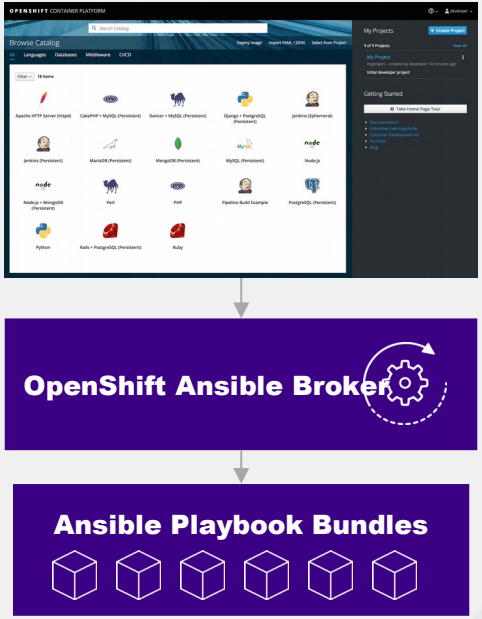

TP available from version 3.6
GA from 3.7

ANSIBLE PLAYBOOK BUNDLES (APB)

- Packaged as a container image
- Embed Ansible runtime
- Use named playbooks for actions
- Fulfill Service Catalog dynamically with services and parameters
- Provide a command line tool to manage APBs

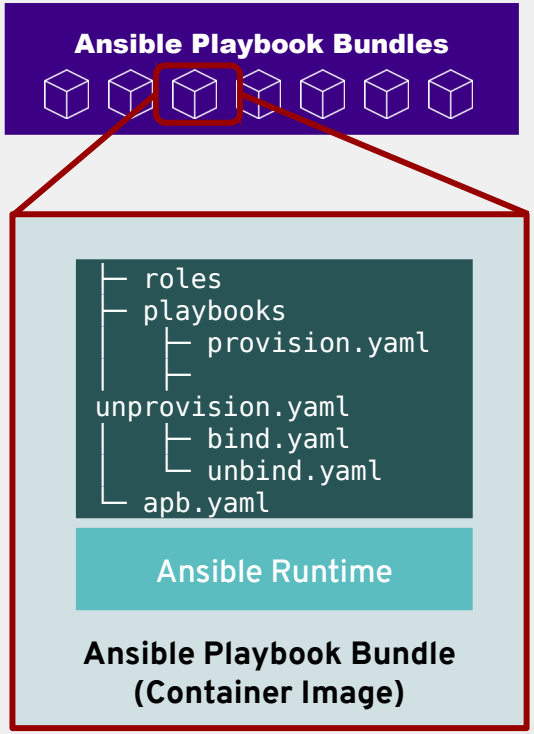

APB CREATION WORKFLOW

Site Reliability Engineer

playbooks and $vars

APB image

Service Catalog update

INIT → CUSTOMIZATION → PREPARE AND BUILD → PUSH

https://developers.redhat.com/blog/2018/05/23/customizing-an-openshift-ansible-playbook-bundle/

DEMO TIME:
MariaDB Provisioning on Remote RHEL

OPENSHIFT APB
MARIADB REMOTE PROVISIONING

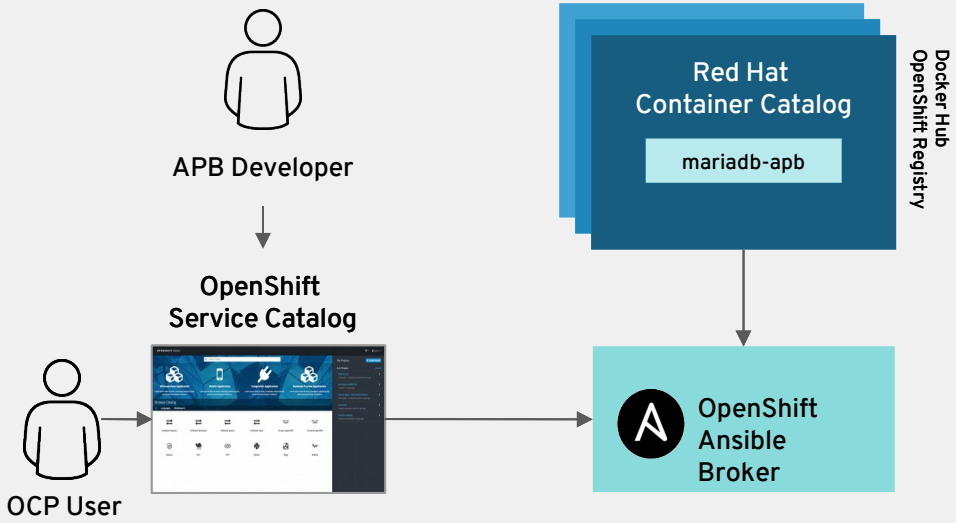

`oc run mariadb-apb provision $vars`

APB container runs `provision.yaml` playbook to install and configure MariaDB on external VM

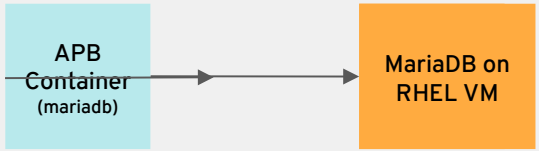

`ansible-playbook provision.yaml $vars`

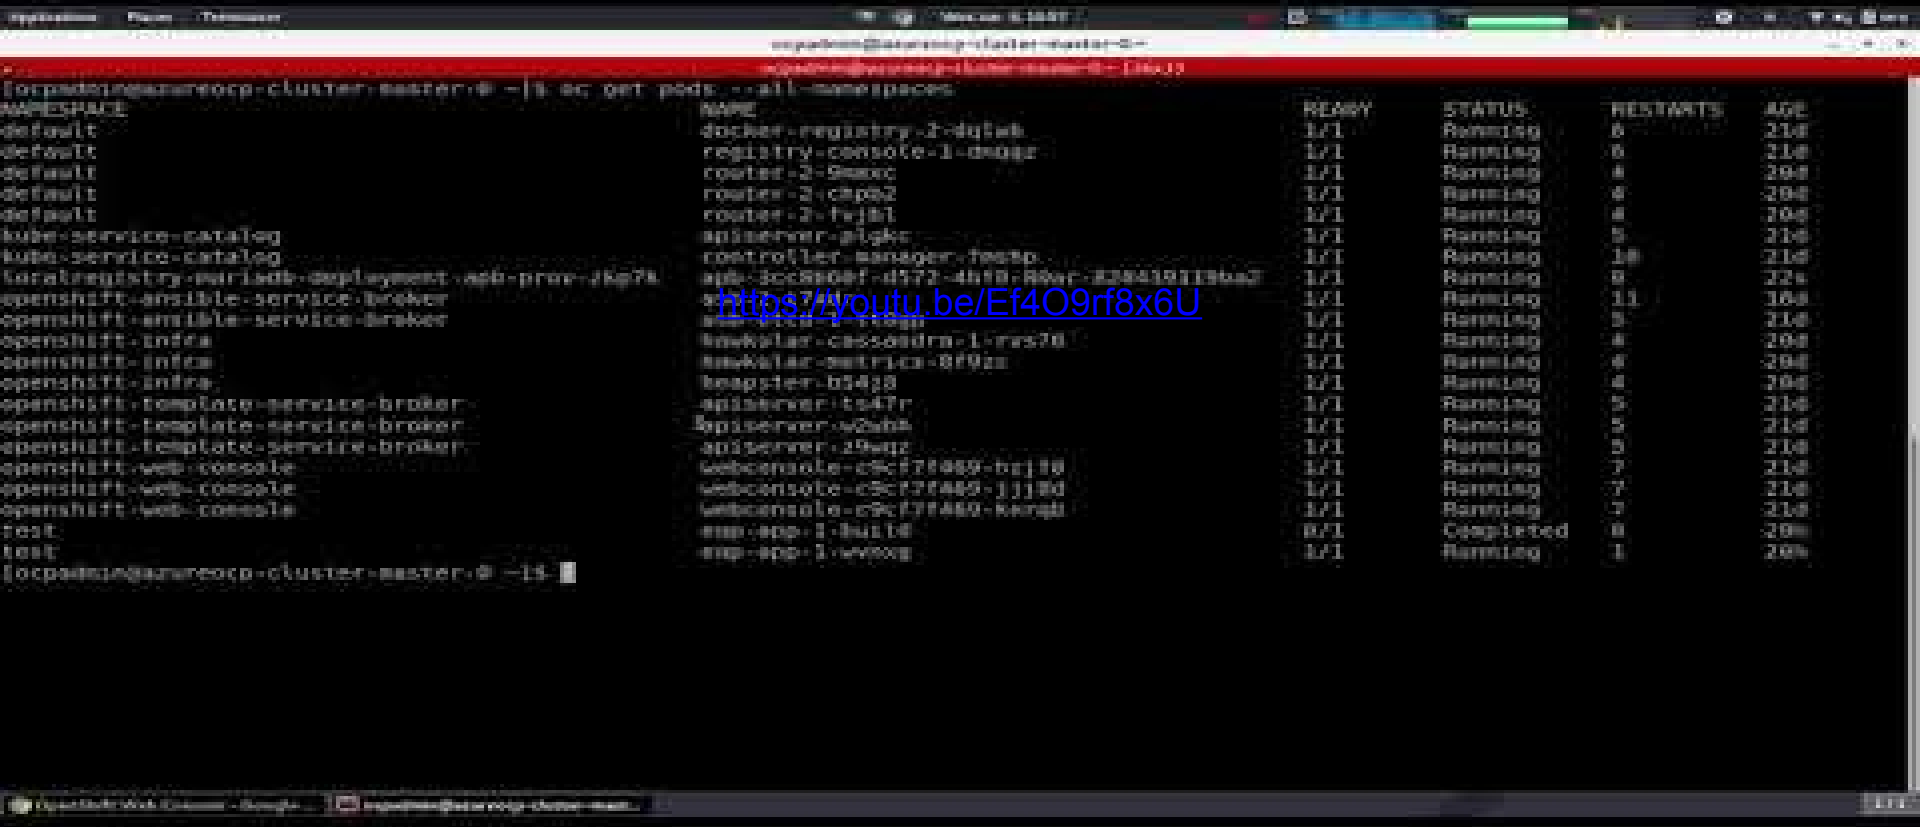

APB INTEGRATION WITH ANSIBLE GALAXY

Support discovering/running APB sources published to Ansible Galaxy from the OpenShift Ansible Service Broker.

How it works:

- APB's can be now be created right from `mazer` command line tool using the init command and then pushed to Ansible Galaxy.
- Broker should now be able to discover and provision APB-based services published to Ansible Galaxy and also make them available in the service catalog.

What's Next?
Operators!

KUBERNETES OPERATORS

THE EASE OF THE CLOUD EVERYWHERE

OPERATOR

- encode human operational knowledge
- **automatically patch, upgrade, recover, and tune apps and services**
- Kubernetes-native
- Purpose-built for a specific application or service

ENCODING AND AUTOMATING OPS KNOWLEDGE WITH OPERATORS

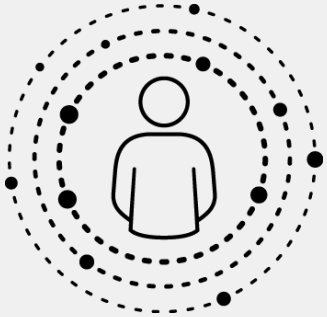

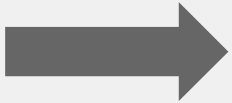

WITHOUT OPERATORS

REACTIVE

- Continually checks for anomalies
- Alert humans for response
- Requires manual change to fix

WITH OPERATORS

PROACTIVE

- Continually adjusts to optimal state
- Automatically acts in milliseconds

OPERATOR FRAMEWORK

An open source toolkit to manage application instances on Kubernetes in an automated, scalable way

Build Operators without specialized knowledge of the Kubernetes API

Install, update, and manage Operators and their dependencies

Enable usage reporting for Operators

https://github.com/operator-framework

OPERATOR IMPLEMENTATION PATHS

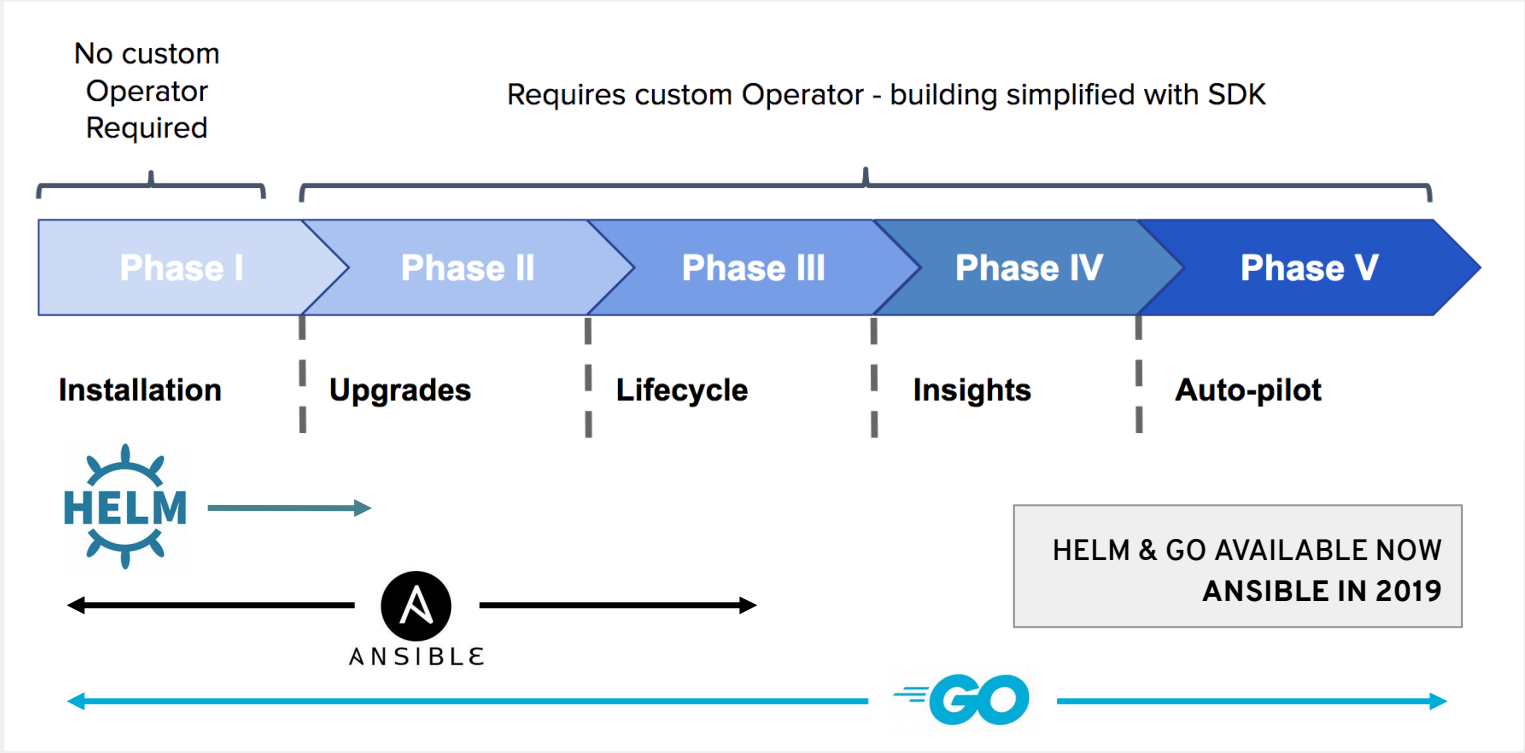

OPERATORS IN PREVIEW IN OCP 3.11

APPLICATION OPERATORS
DEVELOPER PREVIEW

OPERATOR LIFECYCLE MANAGER (OLM)
TECH PREVIEW

Install, manage, and upgrade Operators and their dependencies

RED HAT OPENSHIFT

Portable application services across any infrastructure

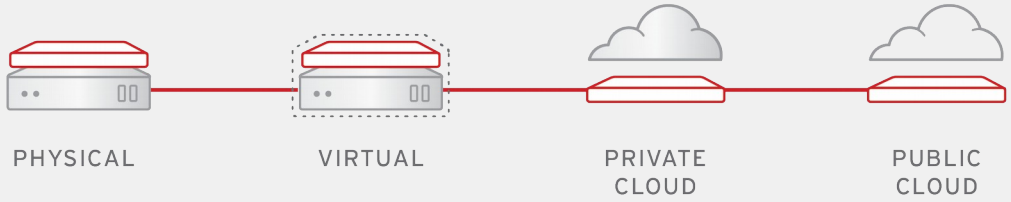

PORTABLE HYBRID CLOUD SERVICES WITH ISV OPERATORS

60+ Certified ISV Operators in Red Hat Early Access Program

Container-native Virtualization

THE FUTURE OF VIRTUALIZATION!

FEDERICO SIMONCELLI
CNV Engineering Manager
fsimonce@redhat.com

#RedHatOSD

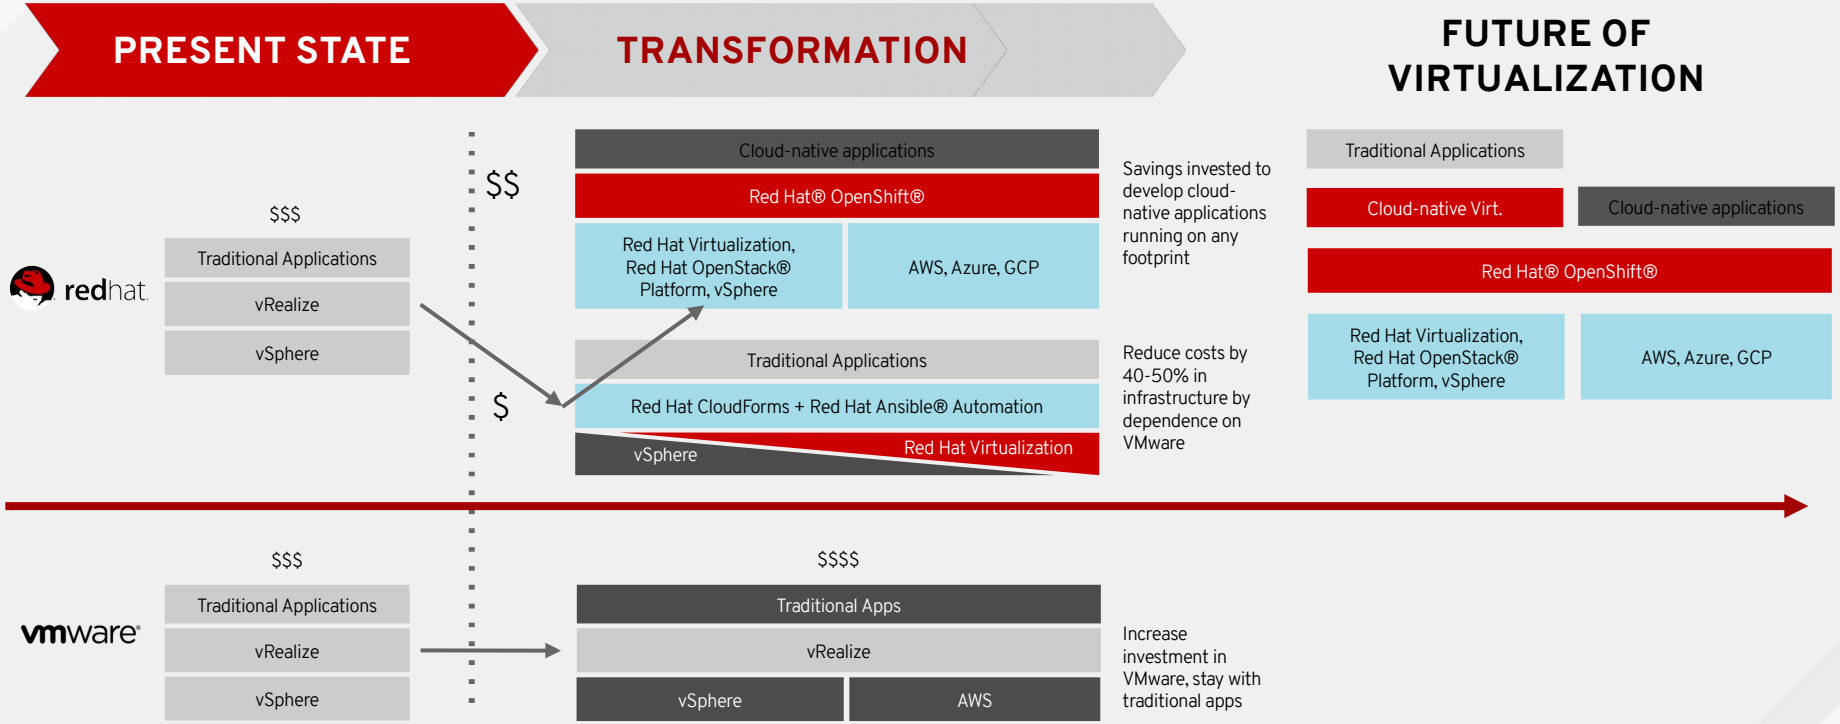
PRESENT STATE
TRANSFORMATION
FUTURE OF VIRTUALIZATION
redhat.
$$$
Traditional Applications
vRealize
vSphere
$$
$
Cloud-native applications
Red Hat® OpenShift®
Red Hat Virtualization, Red Hat OpenStack® Platform, vSphere
AWS, Azure, GCP
Traditional Applications
Red Hat CloudForms + Red Hat Ansible® Automation
vSphere
Red Hat Virtualization
Savings invested to develop cloud-native applications running on any footprint
Reduce costs by 40-50% in infrastructure by dependence on VMware
Traditional Applications
Cloud-native Virt.
Cloud-native applications
Red Hat® OpenShift®
Red Hat Virtualization, Red Hat OpenStack® Platform, vSphere
AWS, Azure, GCP
vmware
$$$
Traditional Applications
vRealize
vSphere
$$$$
Traditional Apps
vRealize
vSphere
AWS
Increase investment in VMware, stay with traditional apps

CONTAINERS AND VIRTUAL MACHINES

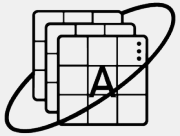

CONTAINER INFRASTRUCTURE AND ORCHESTRATION

Containerized applications and Kubernetes container orchestration as provided by OpenShift are becoming **the** standard for new applications.

VIRTUALIZED WORKLOADS

Virtualized workloads are not going anywhere fast! Business reasons (cost, time to market) and technical reasons (different or older operating system)

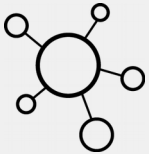

BARE-METAL RESURGENCE

Increasingly customers are pursuing bare-metal clusters for net new business functionality being built in containers.

As the technology mix changes, you will reach a tipping point where containers are the default but some workloads are still more suited to run as VMs

COMPONENTS OF CNV

- **KubeVirt**
  The virtual machine operator
  https://github.com/kubevirt/kubevirt/

- **Containerized Data Importer (CDI)**
  Importing disks
  https://github.com/kubevirt/containerized-data-importer

- **OpenShift Web Console**
  With UI extensions
  https://github.com/openshift/origin-web-console

- **Containerized Virt-v2v**
  Importing a whole virtual machine
  https://github.com/kubevirt/v2v-job

Leverages tried and trusted RHEL & RHV (KVM) virtualization capabilities.

Container-native Virtualization Demo

http://kubevirt.io/get_kubevirt/

Pre-requisites:

- kubectl
- minikube/minishift

Notes:

- Yes, we're running nested virt here - fine for getting started!
- Using upstream bits, for now, in product preview coming!

```
sgordon@:kubevirt-minishift-demo/ $> # Let's look at the new pods our KubeVirt CRDs are running in the kube-syste
m namespace.
sgordon@:kubevirt-minishift-demo/ $> oc get crds
NAME                                                        AGE
datavolumes.cdi.kubevirt.io                                 3h
openshiftwebconsoleconfigs.webconsole.operator.openshift.io 3h
virtualmachineinstancepresets.kubevirt.io                   3h
virtualmachineinstancereplicasets.kubevirt.io               3h
virtualmachineinstances.kubevirt.io                         3h
virtualmachines.kubevirt.io                                 3h
sgordon@:kubevirt-minishift-demo/ $> # The CDI controller runs in the default namespace.
sgordon@:kubevirt-minishift-demo/ $> oc get pods -n default
NAME                        READY     STATUS      RESTARTS   AGE
cdi-deployment-767b445c45-wp7pb   1/1       Running     0          3h
docker-registry-1-2gqht     1/1       Running     0          3h
persistent-volume-setup-658qq     0/1       Completed   0          3h
router-1-nn7gx              1/1       Running     0          3h
sgordon@:kubevirt-minishift-demo/ $> # Our own namespace is as expected empty right now.
sgordon@:kubevirt-minishift-demo/ $> oc get pods
No resources found.
sgordon@:kubevirt-minishift-demo/ $> oc get all
No resources found.
sgordon@:kubevirt-minishift-demo/ $> # Lets look at a VM definition
sgordon@:kubevirt-minishift-demo/ $> vim fedora-vm.yaml
sgordon@:kubevirt-minishift-demo/ $> # Let's now create the VM
sgordon@:kubevirt-minishift-demo/ $> kubectl create -f fedora-vm.yaml
virtualmachine.kubevirt.io "fedora-vm" created
sgordon@:kubevirt-minishift-demo/ $> # The VirtualMachinne object is the persistent representation of our virtual
 machine.
```

ROADMAP THEMES

(What's missing today?)

Supportability

- Simplify upgrade process
- Debug tooling support (sosreports, Insights)
- Broad provider support

Production Workloads

- Layer-2 Networking
- Live Migration
- Upload image as Template
- Guest agent introspection

Embrace the Platform

- Operators for all
- Integrated VM management
- Metrics and monitoring

Container-native Virtualization is **not** a drop-in replacement for traditional virtualization today.

Technology Preview access in an upcoming release of OpenShift.

OpenShift Container Platform 3.11

WHAT'S NEW?

ALESSANDRO ARRICHIELLO
Solution Architect
ale@redhat.com

PIERLUIGI SFORZA
Solution Architect
psforza@redhat.com

#RedHatOSD

NEW ADMIN-FOCUSED CONSOLE

Users have a choice of experience based on their role or technical abilities

- Admin/CaaS experience with heavy exposure to Kubernetes
- AppDev/PaaS experience with standard OpenShift UX
- Sessions are not shared across the Consoles but credentials are
- Both hosted on cluster, in `openshift-console` and `openshift-webconsole` namespaces

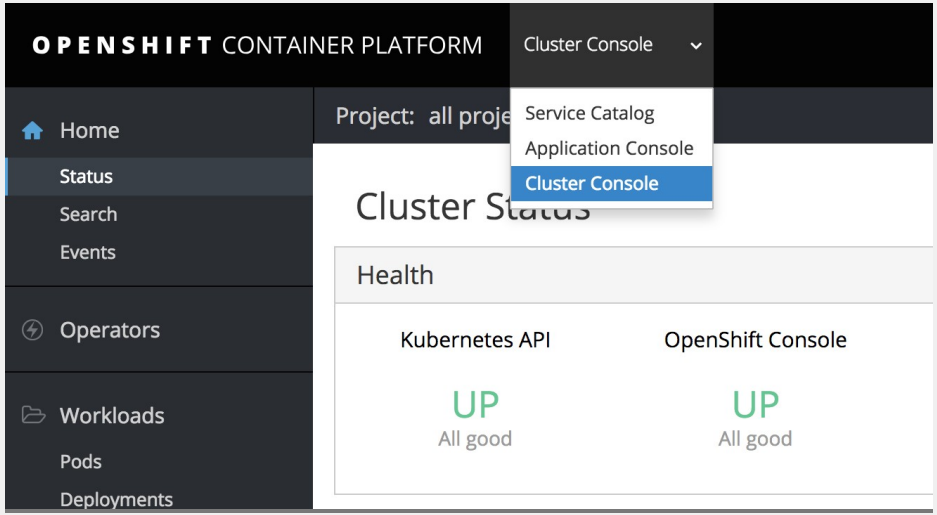

ACCESS CONTROL MANAGEMENT

Visual management of the cluster's RBAC Roles and RoleBindings

- Track down users and service accounts with a specific Role
- View cluster-wide or namespaced bindings
- Visually audit a Role's verbs and objects

Project admins can self-manage roles and bindings scoped to their namespace

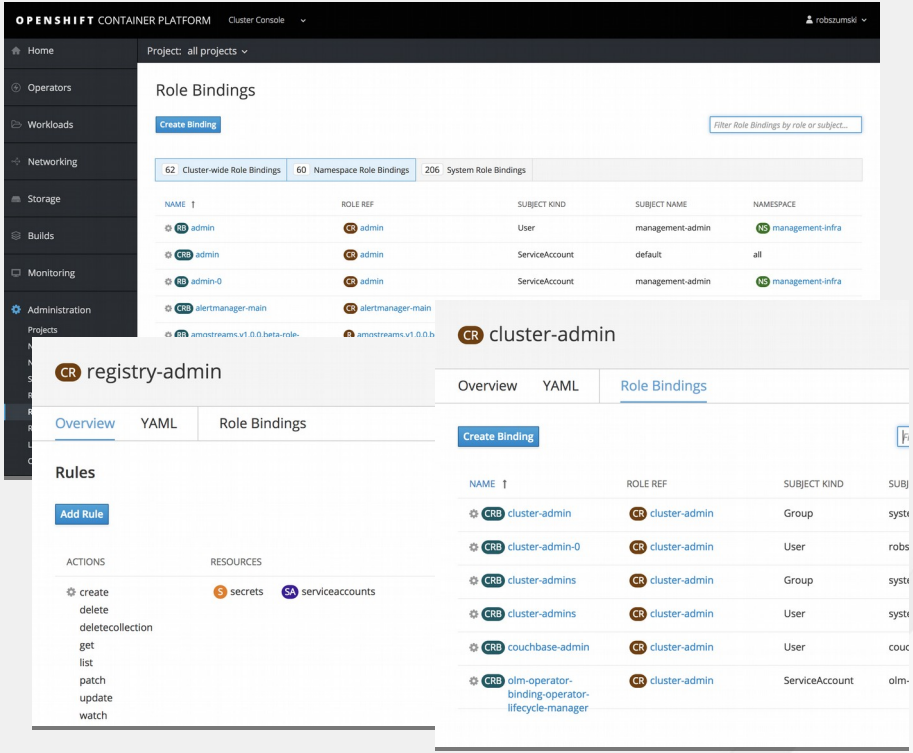

CRI-O / BUILDAH / PODMAN

cri-o

- Becoming the default for partners
- Crictl for node debugging and troubleshooting
- Podman for image tagging & management
- Continues to mature with OpenShift online, customer, and community deployments

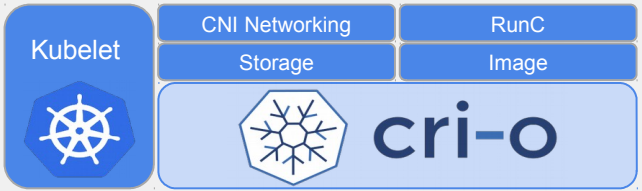

buildah

Start from an existing image or from scratch

Generate new layers and/or run commands on existing layers

Commit storage and generate the image manifest

Deliver image to a local store or remote OCI / docker registry

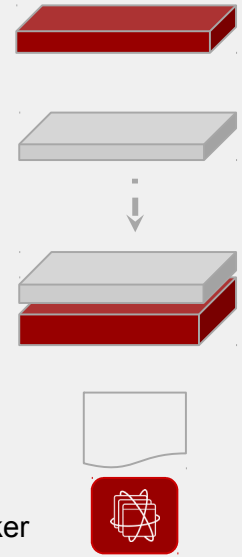

podman

Podman is planned to GA with RHEL 7.6.

A daemon-less CLI/API for running, managing, and debugging OCI containers and pods

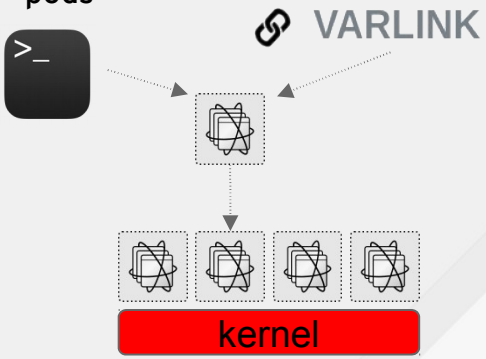

REFERENCE ARCHITECTURE GUIDES

Release: ocpsupplemental-3.11 (in 4-6 weeks after 3.11 GA)

Since 3.10, Reference Architecture Implementation guides are now part of the OpenShift product documentation (https://docs.openshift.com).

Documentation for deploying OCP 3.11 on: *(not live yet)*

- OpenShift 3.11 on Red Hat OpenStack Platform (RHOSP)
- OpenShift 3.11 on Amazon Web Services (AWS)
- OpenShift 3.11 on Microsoft Azure
- OpenShift 3.11 on VMware vSphere
- OpenShift 3.11 on Google Cloud Platform (GCP)
- OpenShift 3.9 on Red Hat Virtualization 4 (RHV) *(update in progress)*

LOCAL DEVELOPMENT

CDK 3.6

- OpenShift Container Platform v3.10.45 (and update to 3.11)
- Based on Minishift 1.24

Minishift 1.24

- Configuration used to start a profile is not saved
- Provide a way to modify the kube-apiserver config same as openshift-apiserver.
- Do not apply templates in xpaas addon one by one
- Local proxy server to handle proxy issues. (technology preview)

kubectl

- We always shipped kubectl for Linux on the master's file system, but now we will offer it in the oc client downloads

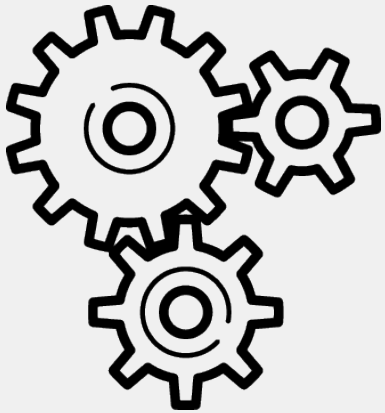

… so you want to do containers and Kubernetes?

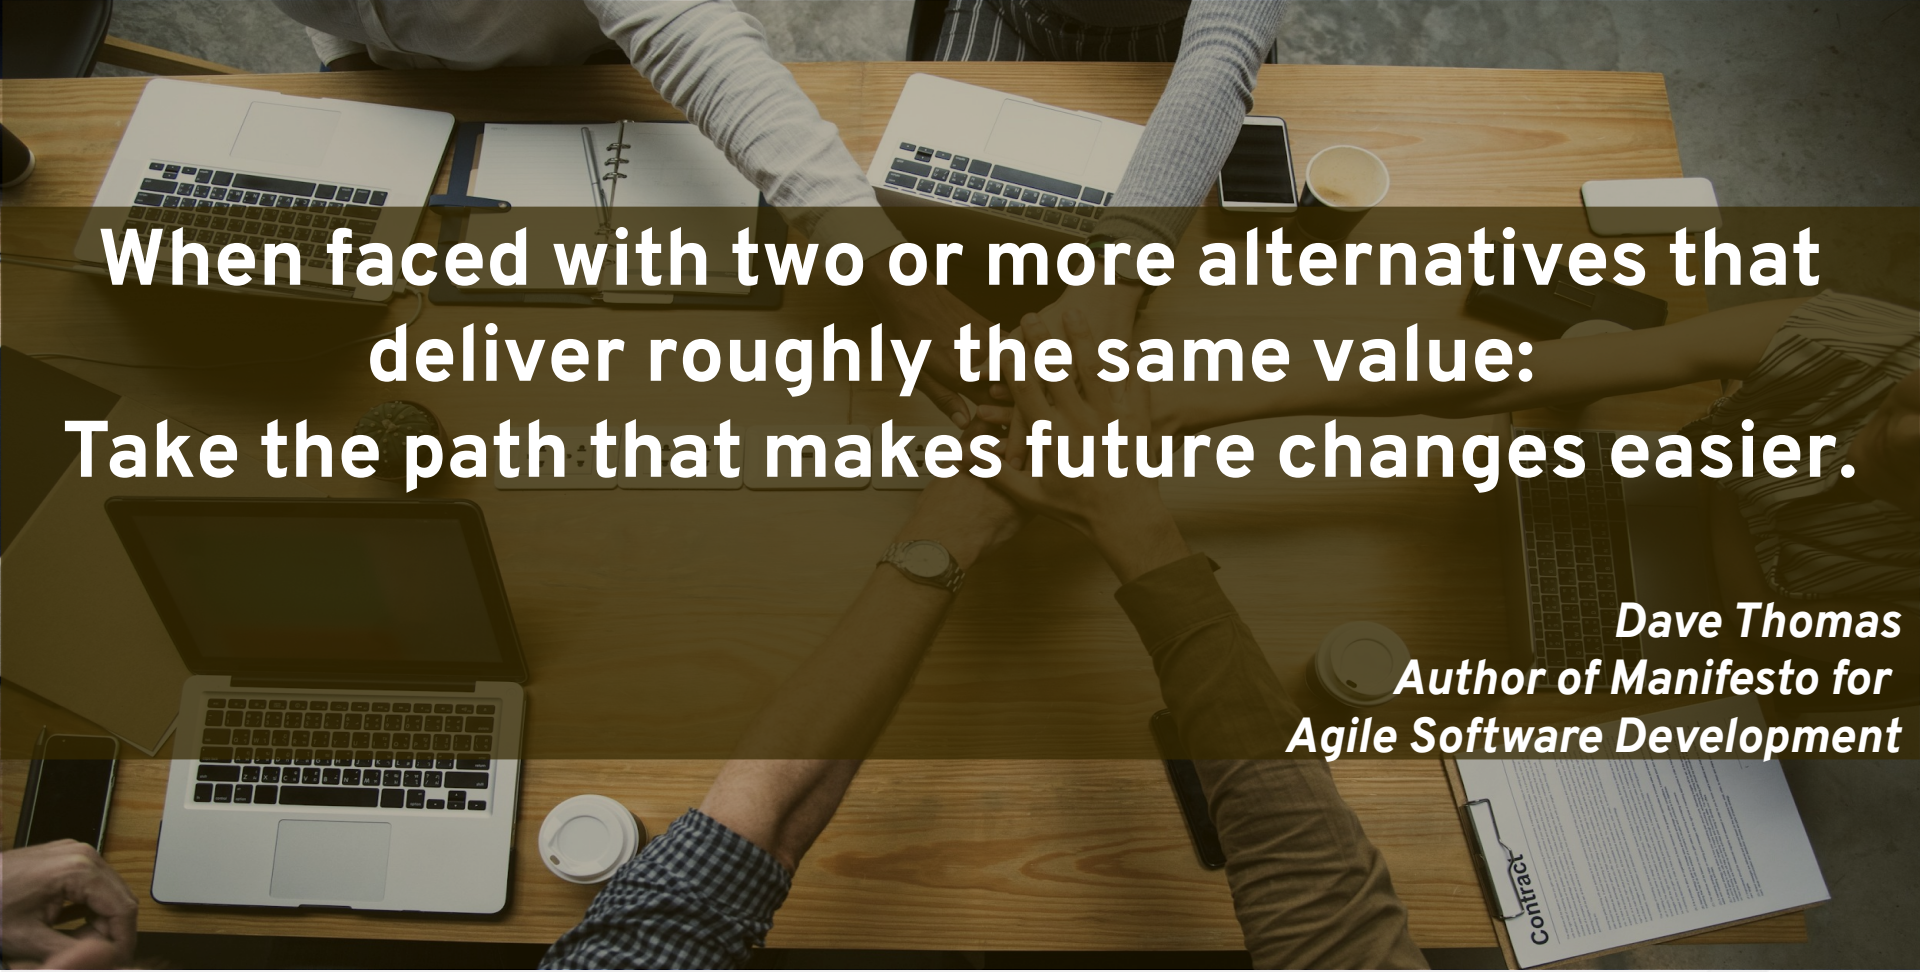
When faced with two or more alternatives that deliver roughly the same value:
Take the path that makes future changes easier.
Dave Thomas
Author of Manifesto for
Agile Software Development
Contract

GRAZIE PER L'ATTENZIONE

ALESSANDRO ARRICHIELLO
Solution Architect
ale@redhat.com

PIERLUIGI SFORZA
Solution Architect
psforza@redhat.com

#RedHatOSD