Securing A Multitenant Kubernetes Cluster

Victoria, BC

Kirsten Newcomer
Senior Principal Product Manager

Red Hat

CONTAINERS ARE THE NEW WAY TO DELIVER APPLICATIONS

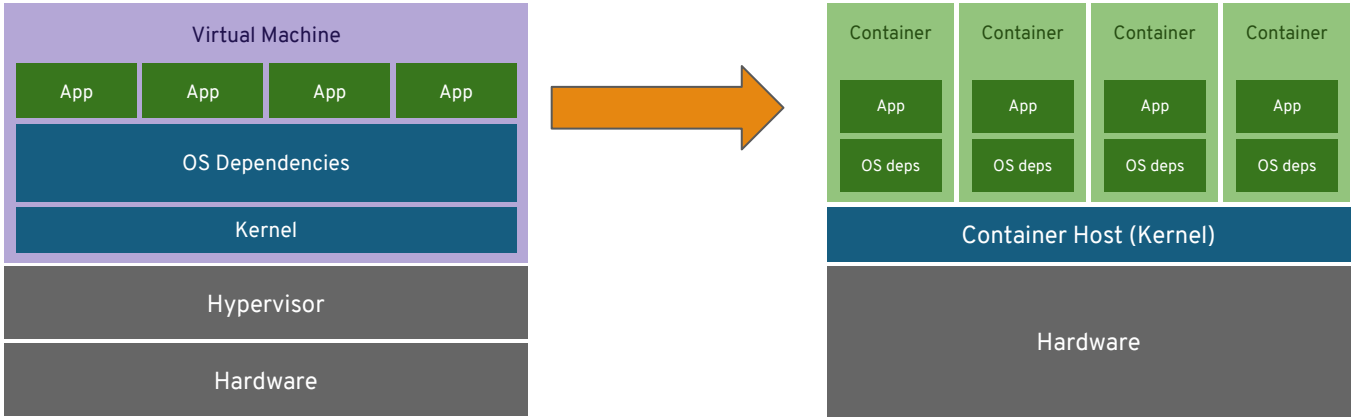

VMs virtualize the hardware

Containers virtualize the process

CONTAINER DEPLOYMENTS ARE INCREASING

Red Hat

KUBERNETES IS THE NEW WAY OF AUTOMATING APPLICATION RESILIENCY

- Auto scale
- Health checks
- Networking (CNI) & Routing
- Platform HA
- Application HA

Red Hat

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

Red Hat

OPENSHIFT HELPS YOU
DELIVER APPLICATIONS FASTER

Cloud-native Applications

AI & Machine Learning

Blockchain

Internet of Things

Innovation Culture

CONTAINERS, KUBERNETES, MICROSERVICES & DEVOPS ARE KEY INGREDIENTS

RED HAT OPENSHIFT BUSINESS VALUE

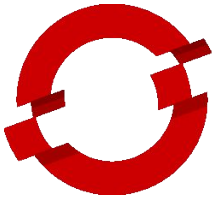

RED HAT® OPENSHIFT

531%
5 Year ROI

66%	36%	8 MONTHS	$1.29M
Faster development lifecycle	More applications per year	Payback period	Average annual benefits per 100 developers

The Business Value of Red Hat OpenShift, IDC #US41845816, October 2017
https://www.redhat.com/en/resources/The-Business-Value-of-Red-Hat-OpenShift

OPENSHIFT IS ENTERPRISE KUBERNETES
BUSINESS CRITICAL APPLICATIONS RUN ON OPENSHIFT

"Red Hat OpenShift allows us to go to market faster. We can move microservices and applications on OpenShift in a few seconds. That's the impact this has on our business." -- Luis Uguina, Chief Digital Officer, Macquarie Bank

- Digital-first bank, reshaping the Australian banking market
- Rethinking their mobile customer experience.
- Using RHEL, OpenShift and JBoss Fuse
- More than 60 business critical applications on OpenShift

This new model is helping us hire and retain top talent.

View the Macquarie Bank keynote

RED HAT PARTNER OROCK ACHIEVES FEDRAMP MODERATE ATO

RESTON, Va., July 23, 2019 /PRNewswire/ — ORock® Technologies, Inc. today announced that it received authorization from the Federal Risk and Authorization Management Program (FedRAMP) to offer Red Hat OpenShift Container Platform within its FedRAMP Moderate cloud environment.

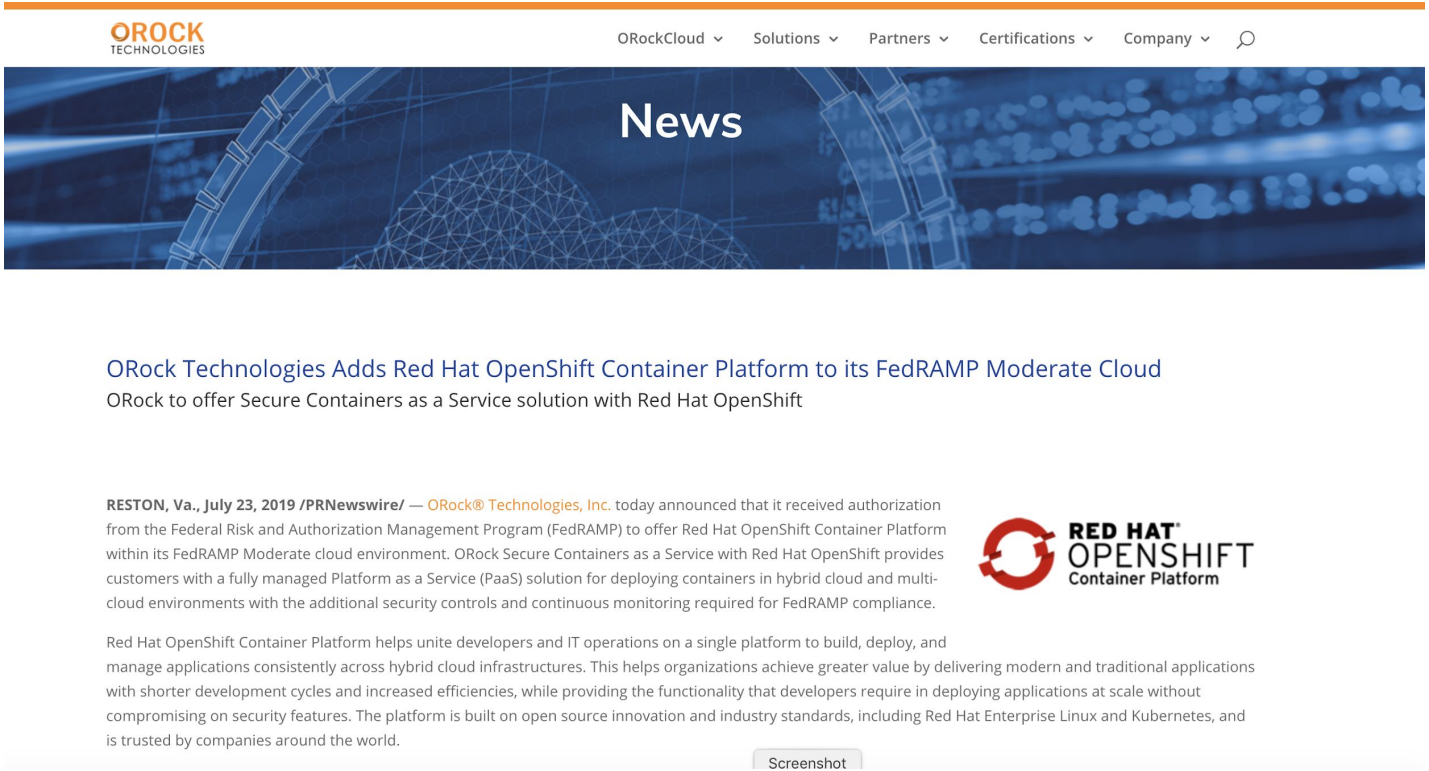

https://orocktech.com/company/news/2019-07-23/

Red Hat

SECURING A MULTI-TENANT CLUSTER

Requires security throughout the stack and the IT lifecycle

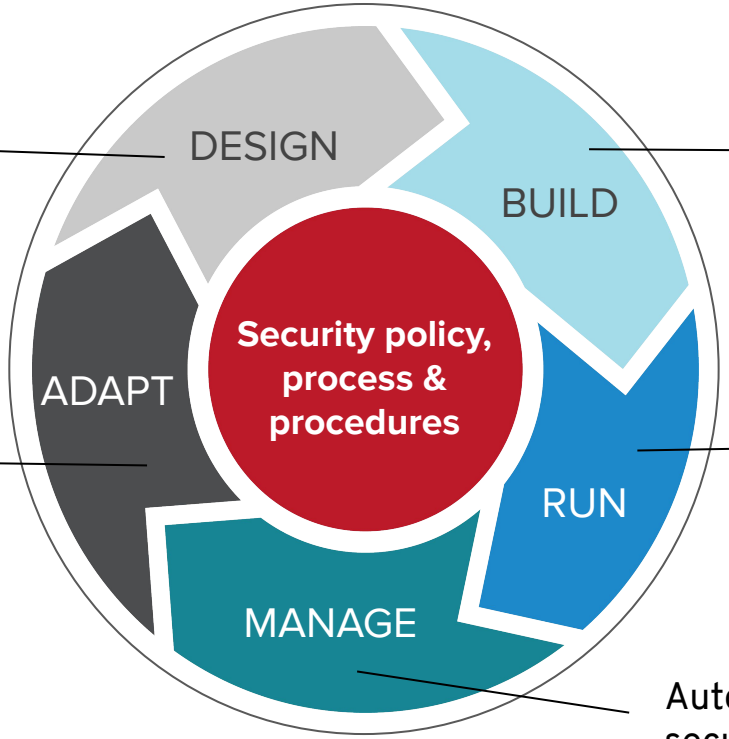

Identify security requirements & governance models

Built-in from the start; not bolted-on

Revise, update, remediate as the landscape changes

Deploy to trusted platforms with enhanced security capabilities

Automate systems for security & compliance

Red Hat

DEVSECOPS

THROUGH THE ADOPTION OF CONTAINERS

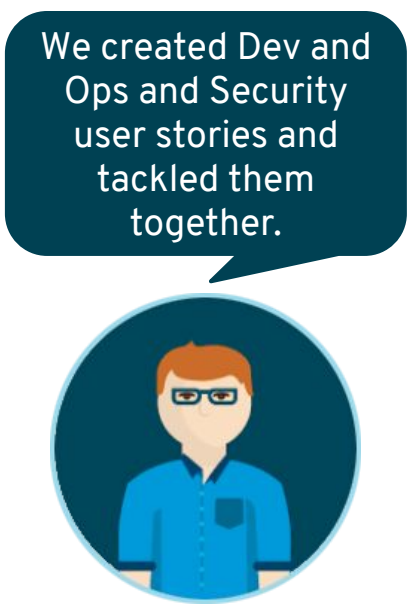

DEVELOPER

SECURITY

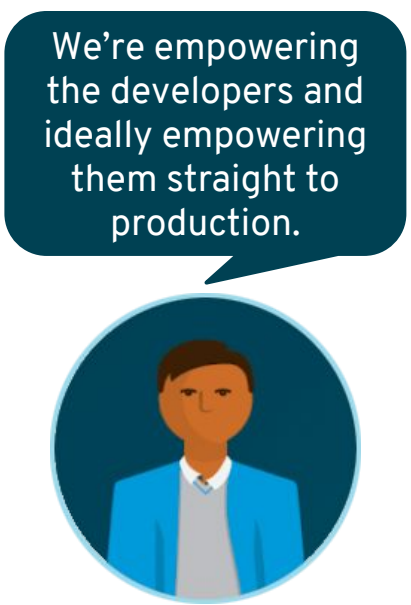

OPERATIONS

[DevNation Federal 2017 - The Journey to DevSecOps]

OPENSHIFT ENABLES MULTI-TENANCY
Layers and Lifecycle

1. Host OS
2. Container platform
3. Network
4. Containerized applications

1. HOST OS CONTAINER MULTI-TENANCY

Container Security starts with Linux Security

- Security in the RHEL host applies to the container
- SELINUX and Kernel Namespaces are the one-two punch no one can beat
- Protects not only the host, but containers from each other
- RHEL CoreOS provides minimized attack surface
- Common Criteria cert - including container framework

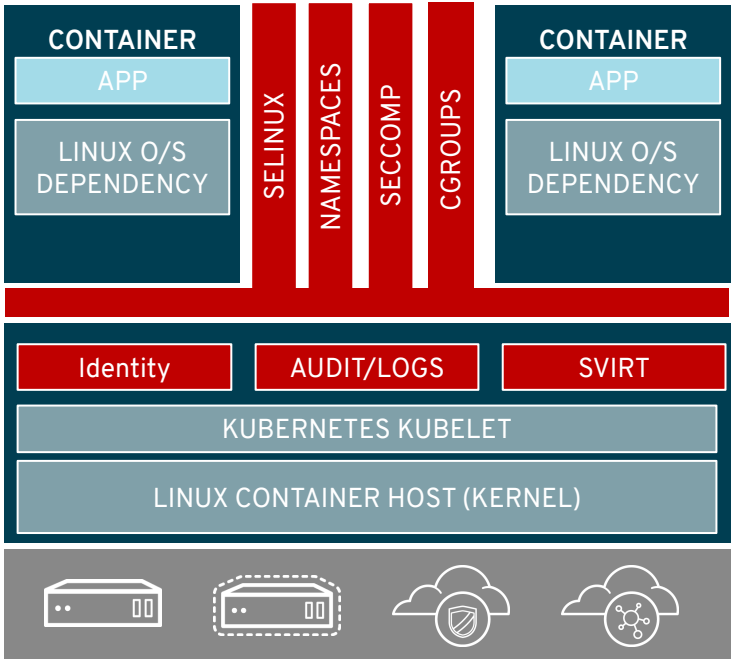

Red Hat

IMMUTABLE OPERATING SYSTEM

Red Hat Enterprise Linux CoreOS is versioned with OpenShift

Only what's needed to run containers
Immutable image-based deployments & updates
Read-only & locked down
Managed by Kubernetes Operators

Red Hat Enterprise Linux CoreOS is managed by the cluster

The Operating system is operated as part of the cluster, with the config for components managed by Machine Config Operator:

- CRI-O config
- Kubelet config
- Authorized registries
- SSH config

Control plane runs on RHEL CoreOS

Worker nodes can run RHEL CoreOS or RHEL

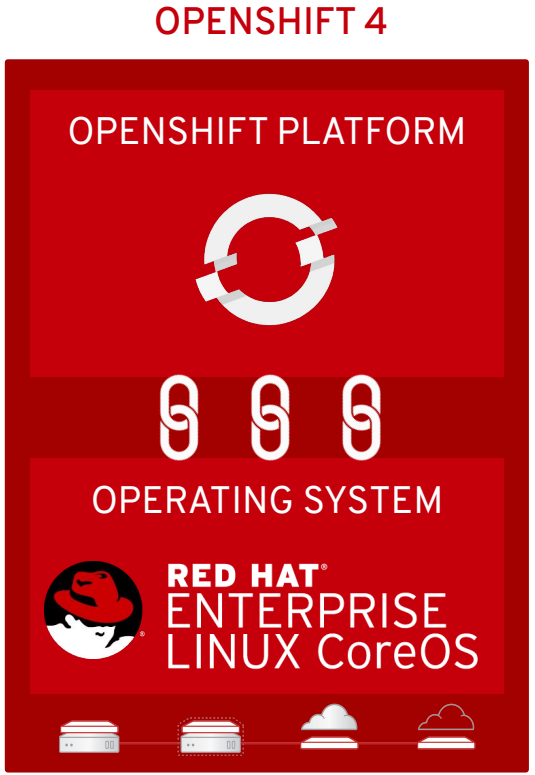

Red Hat

2. THE CONTAINER PLATFORM

OpenShift Security & Multitenancy Features Include

- Host & Runtime security
- Identity and Access Management
- Project namespaces
- Integrated & extensible secrets management
- Service CA
- Logging, Monitoring, Metrics

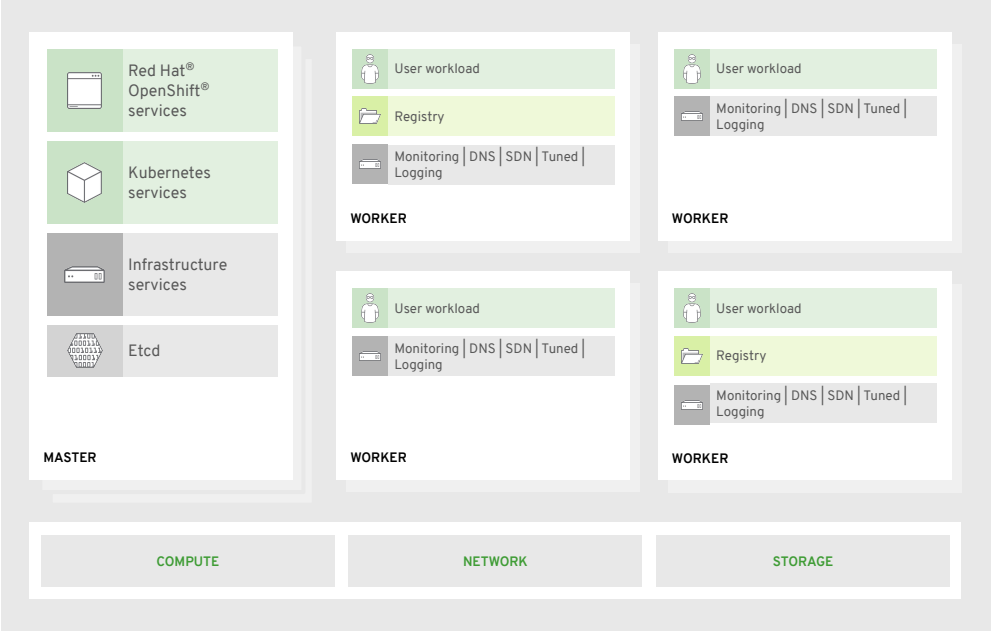

Red Hat

RUNTIME SECURITY POLICIES

SCC (Security Context Constraints)

Allow administrators to control permissions for pods

Restricted SCC is granted to all users

By default, no containers can run as root

Admin can grant access to privileged SCC

Custom SCCs can be created

```
$ oc describe scc restricted
Name:                                   restricted
Priority:                               <none>
Access:
  Users:                                <none>
  Groups:                               system:authenticated
Settings:
  Allow Privileged:                     false
  Default Add Capabilities:             <none>
  Required Drop Capabilities:           KILL,MKNOD,SYS_CHROOT,SETUID,SETGID
  Allowed Capabilities:                 <none>
  Allowed Seccomp Profiles:             <none>
  Allowed Volume Types:                 configMap,downwardAPI,emptyDir,persistentVolumeClaim,projected,
  Allow Host Network:                   false
  Allow Host Ports:                     false
  Allow Host PID:                       false
  Allow Host IPC:                       false
  Read Only Root Filesystem:            false
  Run As User Strategy: MustRunAsRange
```

IDENTITY AND ACCESS MANAGEMENT

OpenShift includes an OAuth server, which does three things:

- Identifies the person requesting a token, using a configured identity provider
- Determines a mapping from that identity to an OpenShift user
- Issues an OAuth access token which authenticates that user to the API

Managing Users and Groups in OpenShift
Configuring Identity Providers

Supported Identity Providers include

- Keystone
- LDAP
- GitHub
- GitLab
- GitHub Enterprise (new with 3.11)
- Google
- OpenID Connect
- Security Support Provider Interface (SSPI) to support SSO flows on Windows (Kerberos)

PROJECTS ISOLATE APPLICATIONS

across teams, groups and departments

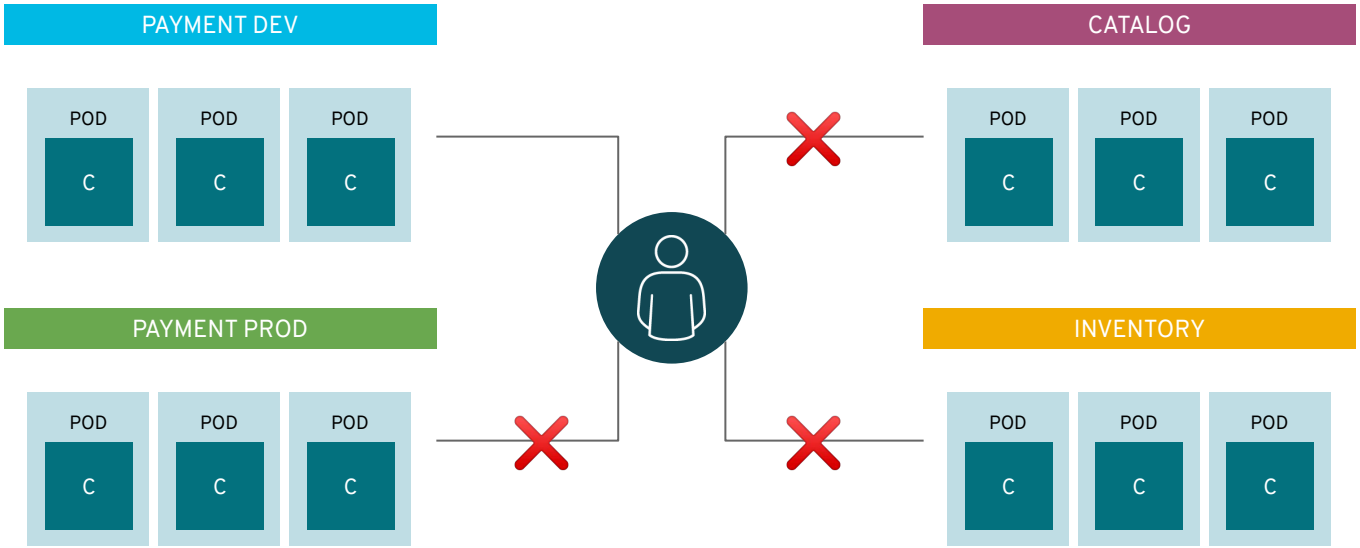

Red Hat

RESTRICT ACCESS BY NEED TO KNOW

Role based authorization (RBAC)

- Project scope & cluster scope available
- Matches request attributes (verb,object,etc)
- If no roles match, request is denied (deny by default)
- Operator- and user-level roles are defined by default
- Custom roles are supported

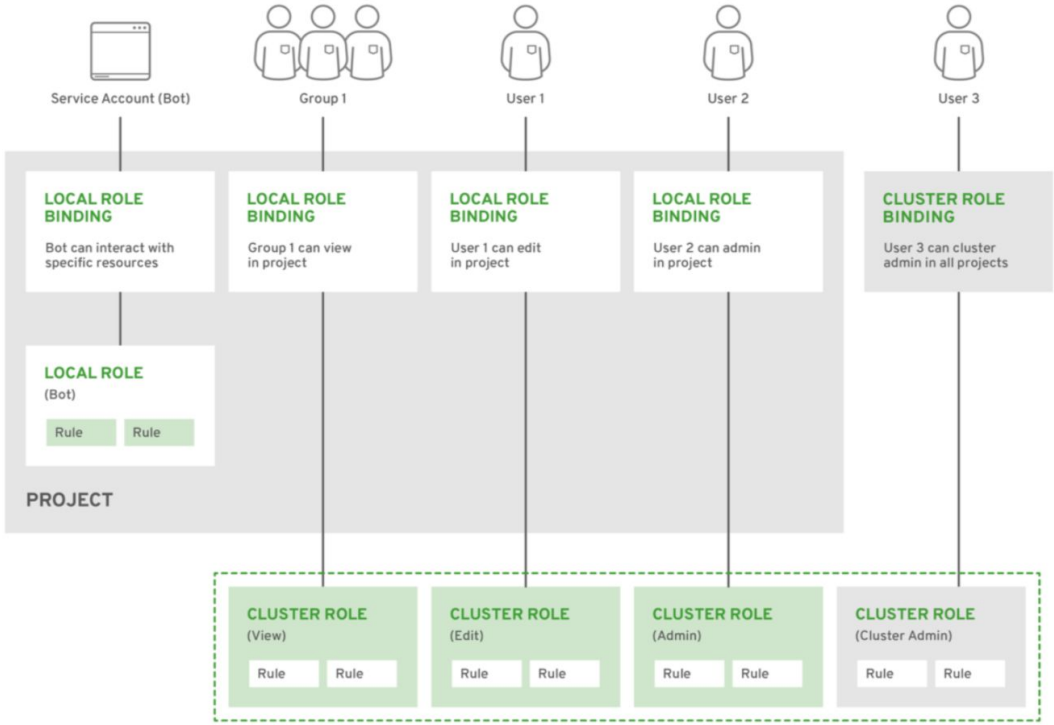

Figure 12 - Authorization Relationships

For more information see: [Managing RBAC in OpenShift]

SECRETS MANAGEMENT

- Platform secrets are stored in etcd
  - Passwords and credentials
  - SSH Keys
  - Certificates
- Application secrets can be stored in etcd or external vault
- Secrets are made available as
  - Environment variables
  - Volume mounts
  - Interaction with external systems (e.g. vaults)
- Encrypted in transit and at rest*
- Never rest on the nodes

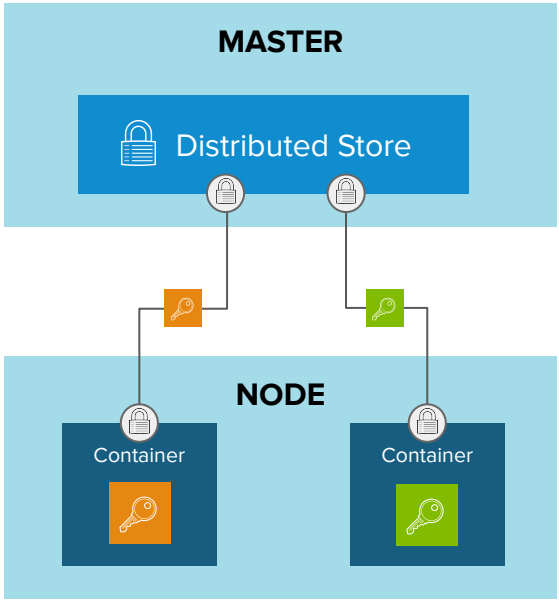

CLUSTER CERTIFICATE MANAGEMENT

- Certificates are used to provide secure connections to
  - master and nodes
  - Ingress controller and registry
  - etcd
- Certificate rotation is automated
- Configure external endpoints to use custom certificates
- For example:
  Requesting and Installing Let's Encrypt Certificates for OpenShift 4

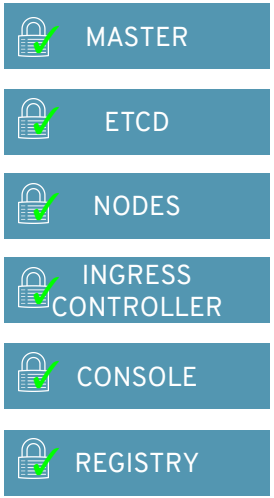

Red Hat

CLUSTER MONITORING

Cluster monitoring is installed by default

- Exposes resource metrics for Horizontal Pod Autoscaling (HPA) by default
  - HPA based on custom metric is tech preview
- No manual etcd monitoring configuration anymore
- New screens for managing Alerts & Silences
- More metrics available for troubleshooting purposes (e.g. HAproxy)
- Configuration via ConfigMaps and Secrets

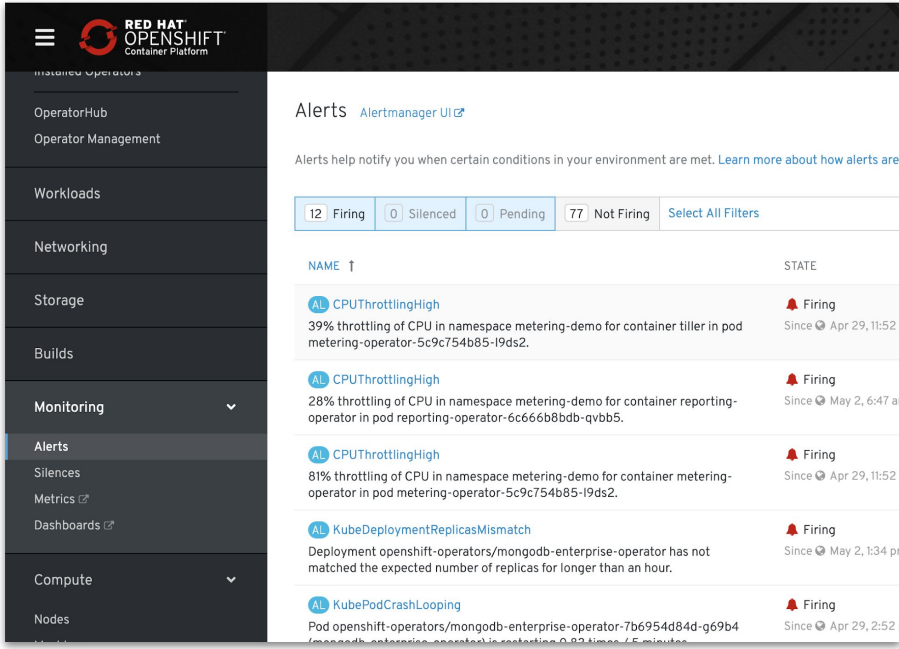

Red Hat

CLUSTER LOG MANAGEMENT

Install the Elasticsearch and Cluster Logging Operators from OperatorHub

- EFK stack aggregates logs for hosts and applications
  - Elasticsearch: a search and analytics engine to store logs
  - Fluentd: gathers logs and sends to Elasticsearch.
  - Kibana: A web UI for Elasticsearch.
- Access control
  - Cluster administrators can view all logs
  - Users can only view logs for their projects
  - Central Audit policy configuration
- Ability to send logs elsewhere
  - External elasticsearch, Splunk, etc

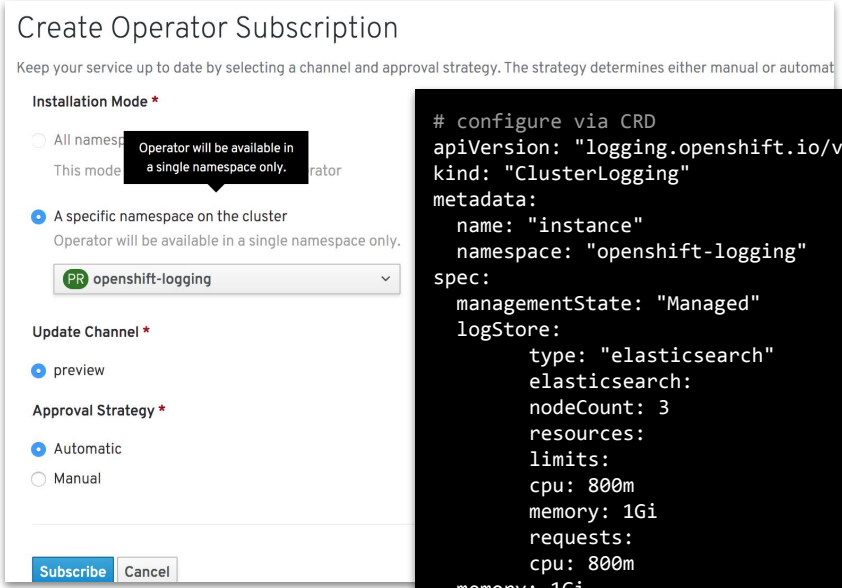


```
# configure via CRD
apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance"
  namespace: "openshift-logging"
spec:
  managementState: "Managed"
  logStore:
        type: "elasticsearch"
        elasticsearch:
        nodeCount: 3
        resources:
        limits:
        cpu: 800m
        memory: 1Gi
        requests:
        cpu: 800m
  memory: 1Gi
  storage:
        storageClassName: gp2
        size: 100G
        redundancyPolicy: "SingleRedunda
  visualization:
        type: "kibana"
        kibana:
        replicas: 1
  curation:
```

3. NETWORK MULTI-TENANCY

Fine Grained Control with Network Policy

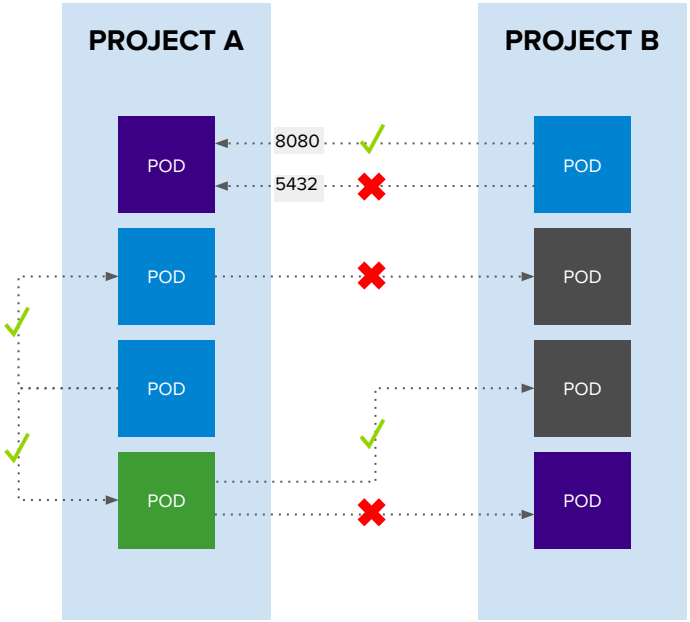

Example Policies

- Allow all traffic inside the project
- Allow traffic from green to gray
- Allow traffic to purple on 8080

```
apiVersion: extensions/v1beta1
kind: NetworkPolicy
metadata:
  name: allow-to-purple-on-8080
spec:
  podSelector:
    matchLabels:
      color: purple
  ingress:
  - ports:
    - protocol: tcp
      port: 8080
```

Enabled by default in OpenShift 4

Red Hat

MULTI-TENANT INGRESS & EGRESS CONTROL

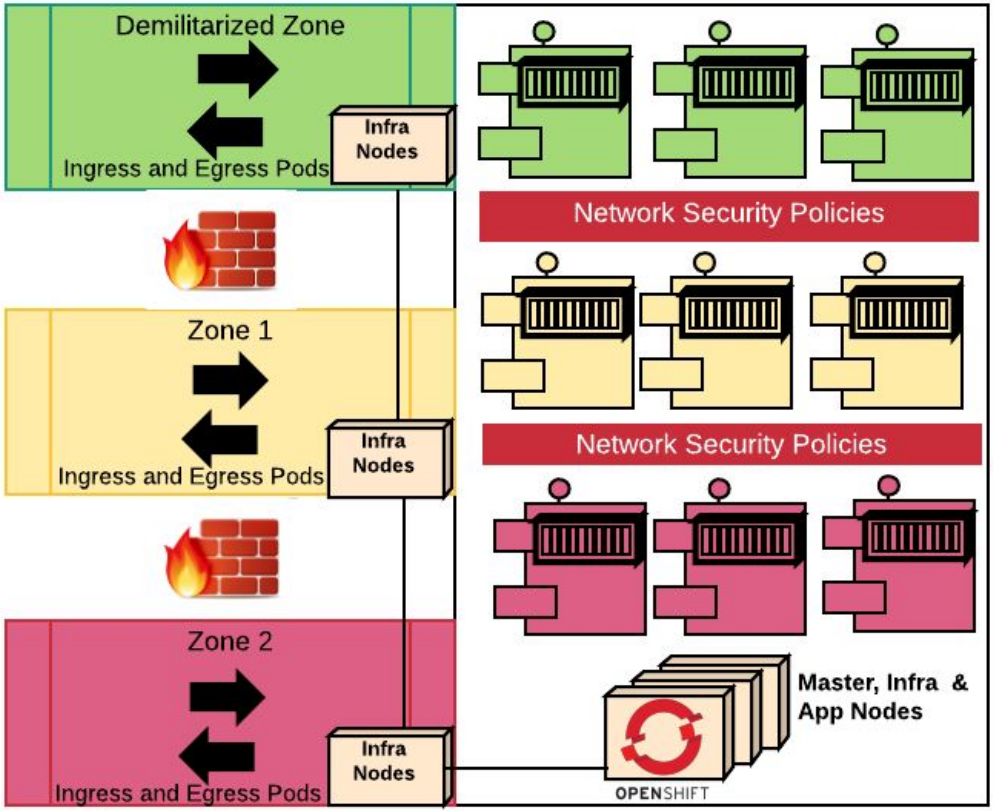

Application pods run on one OpenShift Cluster. Microsegmented with Network Security policies.

Infra Nodes in each zone run Ingress and Egress pods for specific zones. Egress firewall to limit external addresses accessed.

If required, physical isolation of pods to specific nodes is possible with node-selectors. But that can reduce worker node density.

There may be cases where a single tenant cluster is preferred.

Red Hat

OPENSHIFT MULTUS

Optionally Separate Control Plane and Data Plane

Multus Enables Multiple Networks & New Functionality to Existing Networking

The Multus CNI “meta plugin” for Kubernetes enables one to create multiple network interfaces per pod, and assign a CNI plugin to each interface created.

1. Create pod annotation(s) to call out a list of intended network attachments…
2. …each pointing to CNI network configurations packed inside CRD objects

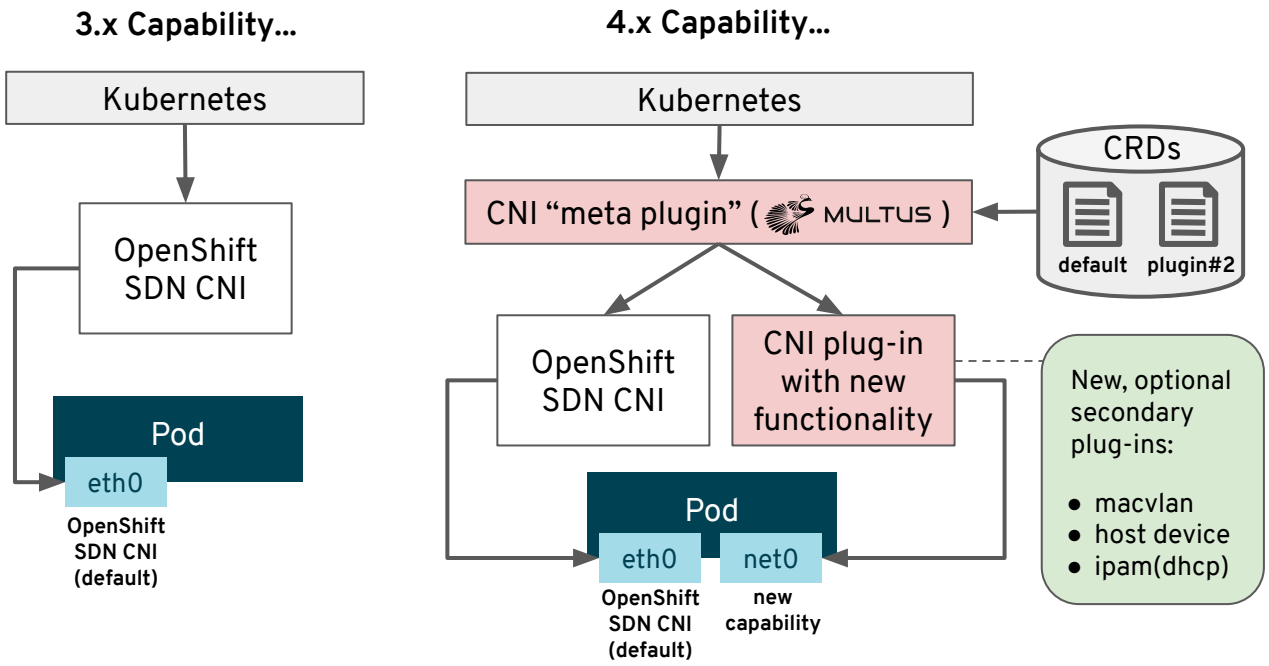

Red Hat

4. SECURING CONTAINERIZED APPLICATIONS

Secure and Automate the Content Lifecycle
Trust is temporal; rebuild and redeploy as needed

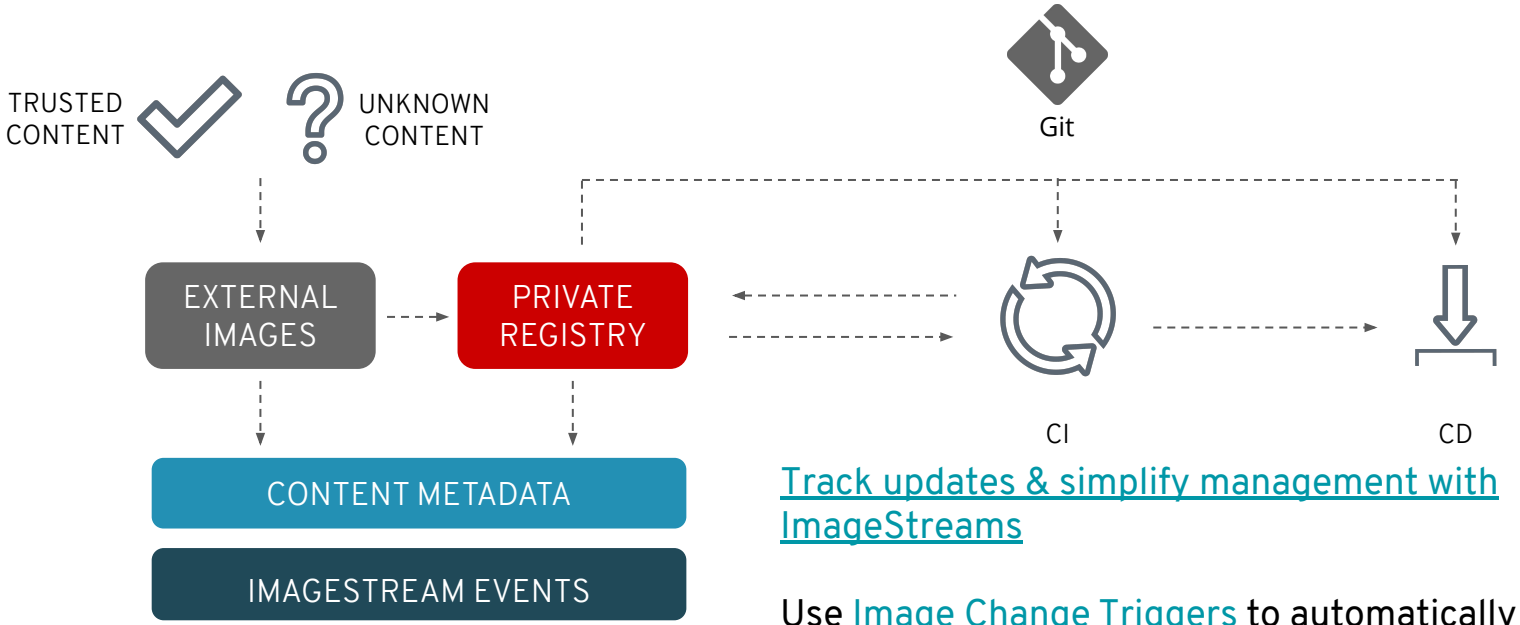

Track updates & simplify management with ImageStreams

Use Image Change Triggers to automatically rebuild custom images with updated (patched) external images

Red Hat

RED HAT QUAY
ENTERPRISE CONTAINER REGISTRY

- Offered as self-managed and as-a-service
- Vulnerability Scanning (Clair)
- Geographic Replication
- Build Image Triggers
- Image Rollback with Time Machine

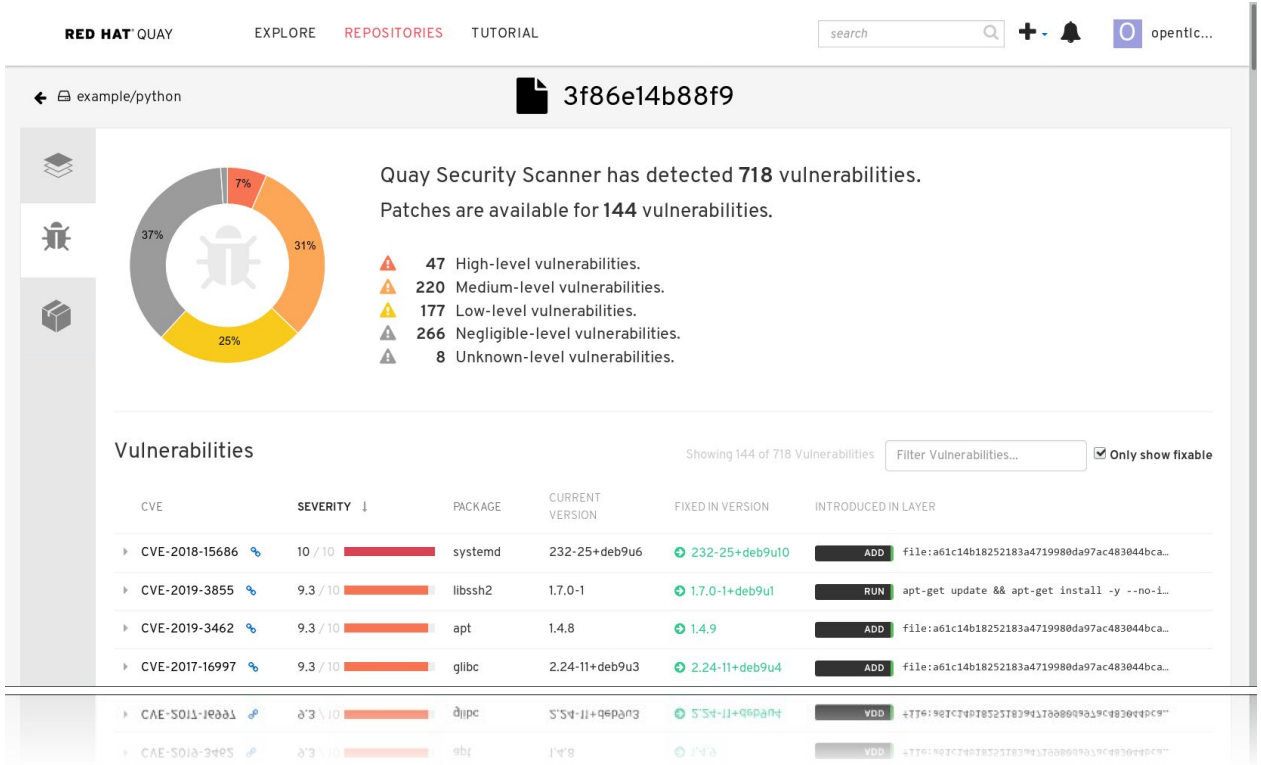

Red Hat

CI/CD MUST INCLUDE SECURITY GATES

- Integrate security testing into your build / CI process
- Use automated policies to flag builds with issues
- Sign your custom container images

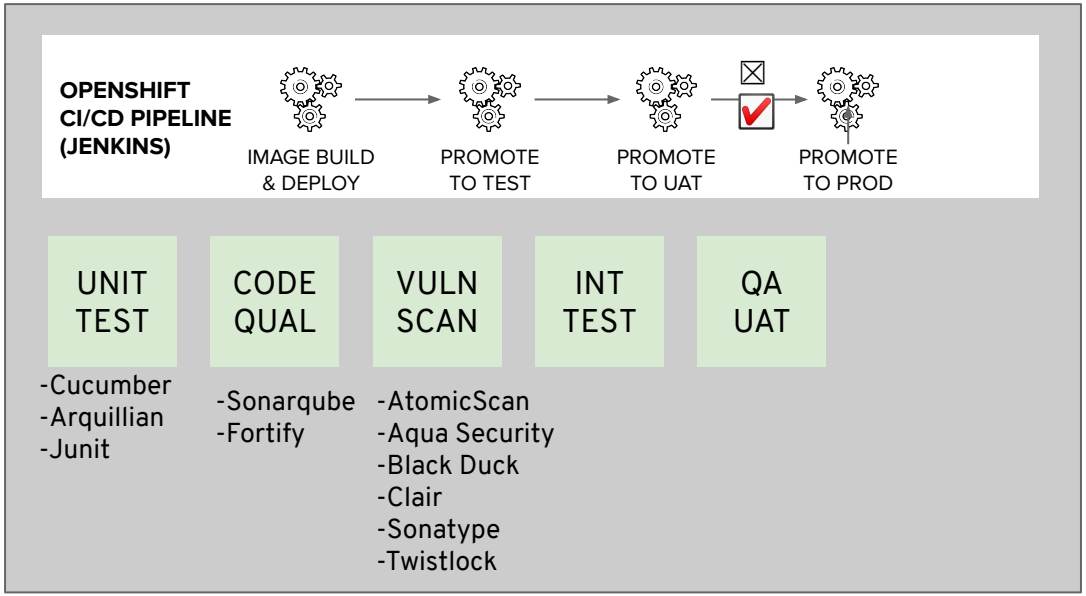

Red Hat

RED HAT SERVICE MESH

Key Features

- A dedicated network for service to service communications
- Observability and distributed tracing
- Policy-driven security
- Routing rules & chaos engineering
- Powerful visualization & monitoring
- Will be available via OperatorHub
- **Working on multi-tenancy for GA** (e.g. Kiali to use OpenShift RBAC)

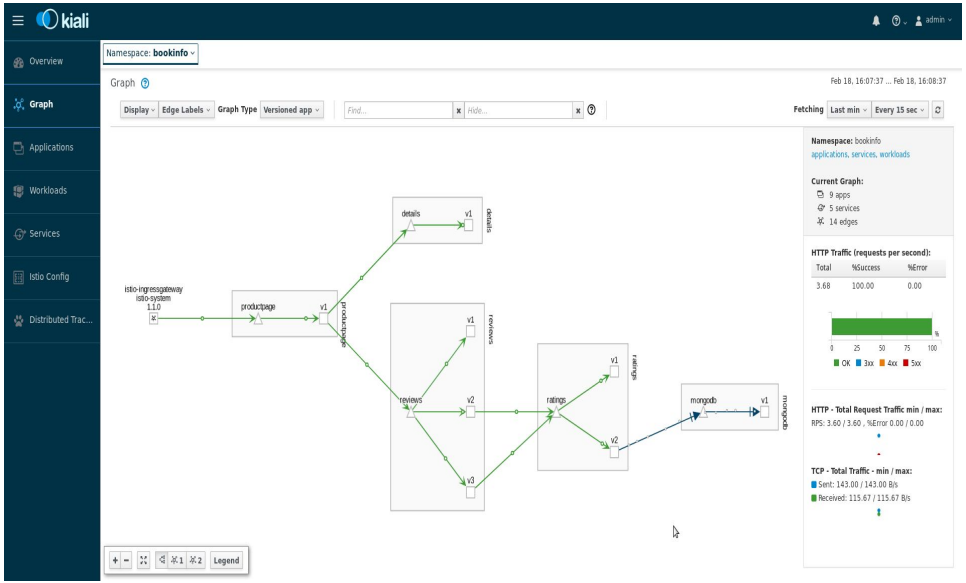

Generally Available in July

Red Hat

COMPREHENSIVE CONTAINER SECURITY

CONTROL
Application Security

- Container Content
- CI/CD Pipeline
- Container Registry
- Operators

DEFEND
Infrastructure

- Container Host Multi-tenancy
- Container Platform
- Network Isolation
- Storage
- Audit & Logging
- Service Mesh

EXTEND

- Security Ecosystem

THE BROAD SECURITY ECOSYSTEM









Red Hat

Next Steps

- Speak with a Red Hat expert here at Security Symposium
- Look for the slides in a “Thank You” email from us in the next few days
- Stay up to date with Red Hat at redhat.com/security
- Visit redhat.com/events to find out about workshops and other events like this one coming to your area

Thank you for coming.

Feedback or questions?

infrastructure@redhat.com

Thank you to our partner

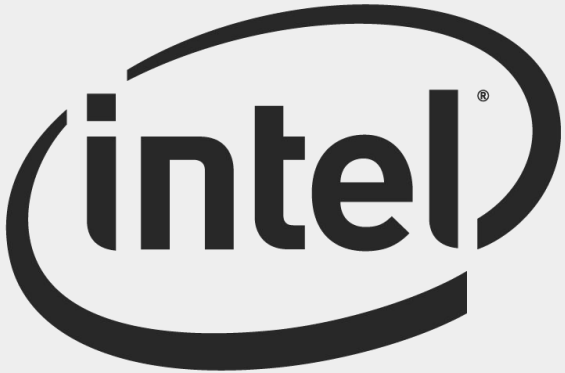

REGULATORY COMPLIANCE WITH OPENSHIFT

- Red Hat contracted with Coalfire to provide a PCI-DSS technical controls product applicability guide (PCI-DSS 3.2) and reference architecture (PCI-DSS 3.2.1) for OpenShift
- Guides also available for ISO 27001, FISMA (NIST) & FISMA
- OpenShift Hardening Guide for 3.10 & 3.11 - OpenShift 4 Guide planned for Fall 2019

OPERATORS SIMPLIFY MANAGEMENT OF COMPLEX APPLICATIONS ON KUBERNETES

OPERATOR SDK

DEVOPS	dynatrace
APM	APPDYNAMICS, INSTANA, New Relic., Sysdig
DATA SERVICES	GIGASPACES, hazelcast, PlanetScale
DATABASE	Couchbase, mongoDB., NUODB, PingCAP
SECURITY	aqua, anchore, BLACKDUCK BY SYNOPSYS, TREMOLO SECURITY, tufin
STORAGE	ROBIN, STORAGEOS
AND MANY MORE TO COME...	

Red Hat

- Containerized

- Containerized
- Cloud storage ready
- Replicated
- Backup
- Automated updates

- Containerized
- Container storage ready
- Replicated
- Backup
- Automated updates
- Enhanced observability
- Customization
- Local development
- Fully Open Source
- Any Kubernetes
- Certified on OpenShift

EVERYTHING RUNS IN PODS IN OPENSHIFT 4

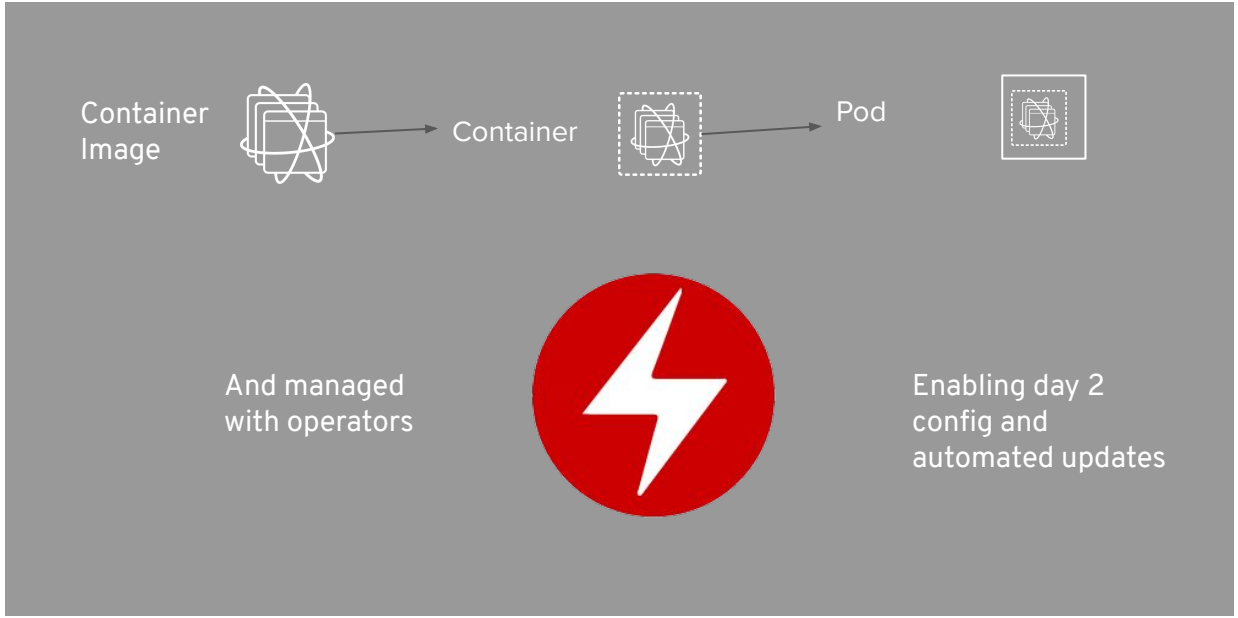

Red Hat

BUILD OPERATORS FOR YOUR APPS

OPERATOR SDK

Use OLM to Manager your Application Lifecycle

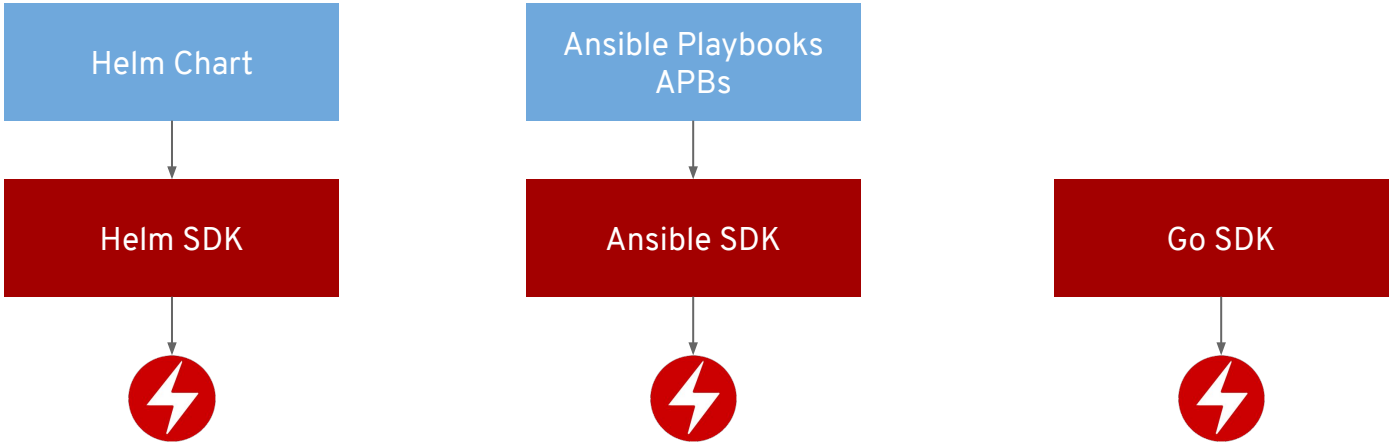

Build operators from Helm chart, without any coding

Build operators from Ansible playbooks and APBs

Build advanced operators for full lifecycle management

Red Hat

ATTACHED STORAGE

Secure storage by using

- SELinux access controls
- Secure mounts
- Supplemental group IDs for shared storage

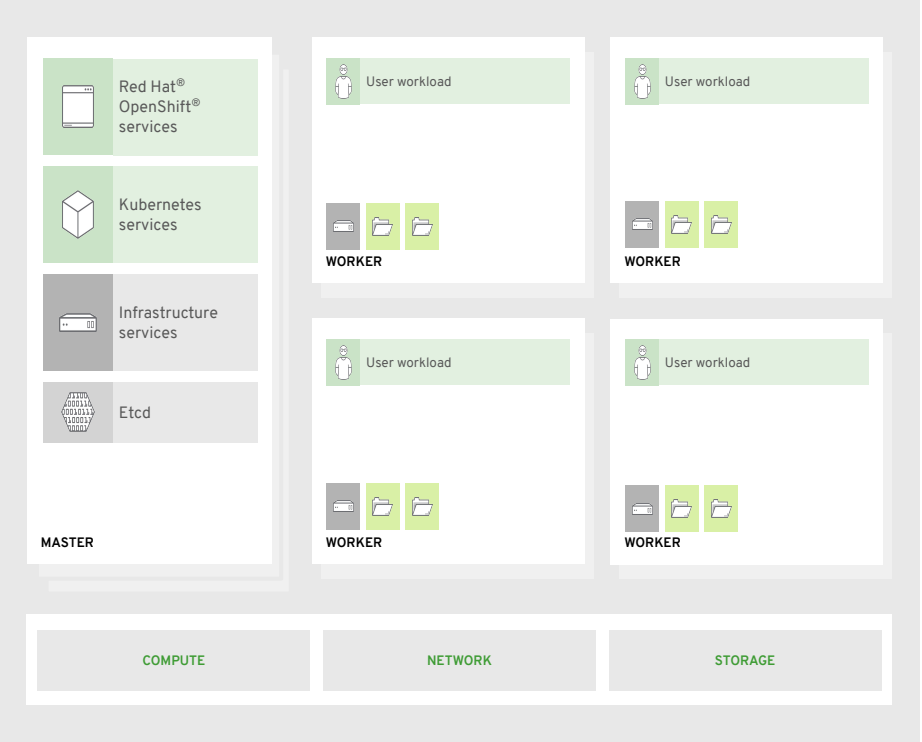

Red Hat