OPENSHIFT 4 & PCI-DSS

Kirsten Newcomer
Director, Cloud and DevSecOps Strategy
Red Hat, Cloud Platforms
December 2020

Red Hat

What makes an effective hybrid cloud platform?

BROAD ECOSYSTEM	BROADEST APPLICATION SUPPORT	DEVELOPER EXPERIENCE & ON-DEMAND
AUTOMATED OPERATIONS	STANDARDS, PORTABILITY & INTEROPERABILITY	SECURITY & COMPLIANCE

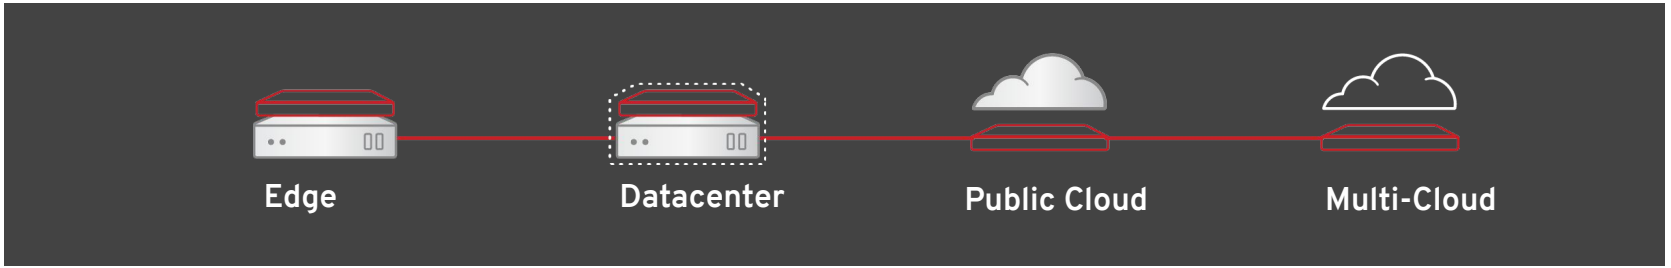

Red Hat

Security must be continuous

And integrated throughout the IT lifecycle

Identify security requirements & governance models

Built-in from the start; not bolted-on

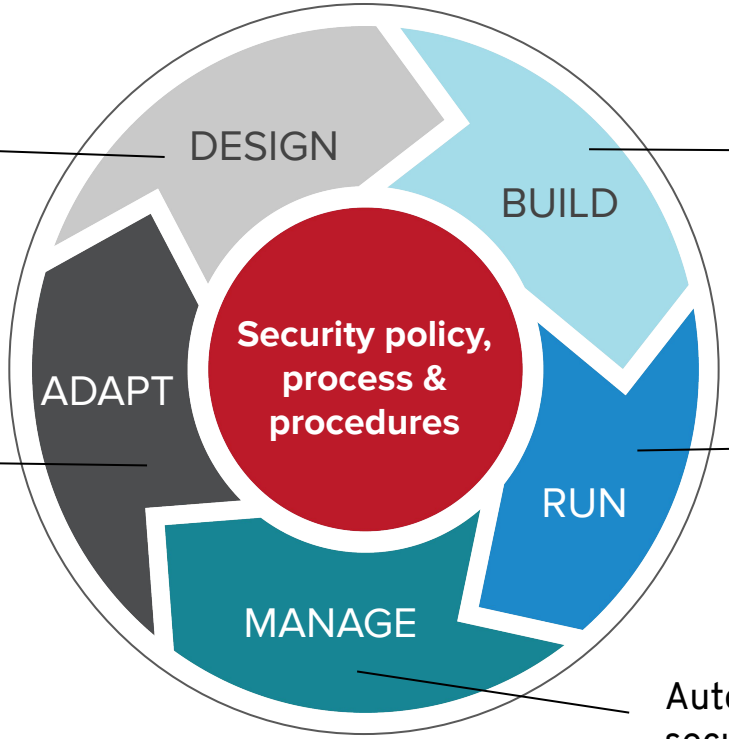

Revise, update, remediate as the landscape changes

Deploy to trusted platforms with enhanced security capabilities

Automate systems for security & compliance

Red Hat

Enterprise Kubernetes from Red Hat

Red Hat Advanced Cluster Management for Kubernetes

Multicluster management
Observability ⁝ Discovery ⁝ Policy ⁝ Compliance ⁝ Configuration ⁝ Workloads

Red Hat Advanced Cluster Security for Kubernetes

Advanced security
Declarative security ⁝ Vulnerability management ⁝ Network segmentation ⁝ Threat detection & response

OpenShift Container Platform

Manage workloads	Build cloud-native apps	Data-driven insights	Developer productivity
Platform services	**Application services**	**Data services**	**Developer services**
Service mesh ⁝ Serverless Builds ⁝ CI/CD pipelines Log management Cost management	Languages & runtimes API management Integration Messaging Process Automation	Databases ⁝ Cache Data ingest & prep Data analytics ⁝ AI/ML Data mgmt & resilience	Developer CLI ⁝ IDE Plugins & extensions CodeReady Workspaces CodeReady Containers

certified kubernetes

OpenShift Kubernetes Engine

Kubernetes cluster services
Automated Ops ⁝ Over-the-air updates ⁝ Monitoring ⁝ Logging ⁝ Registry ⁝ Networking ⁝ Router ⁝ Virtualization ⁝ OLM ⁝ Helm

Kubernetes (orchestration)

Linux (container host OS)

Physical | **Virtual** | **Private cloud** | **Public cloud** | **Managed cloud** (Azure, AWS, GCP, IBM, Red Hat) | **Edge**

Red Hat

Red Hat delivers continuous security for containers and Kubernetes

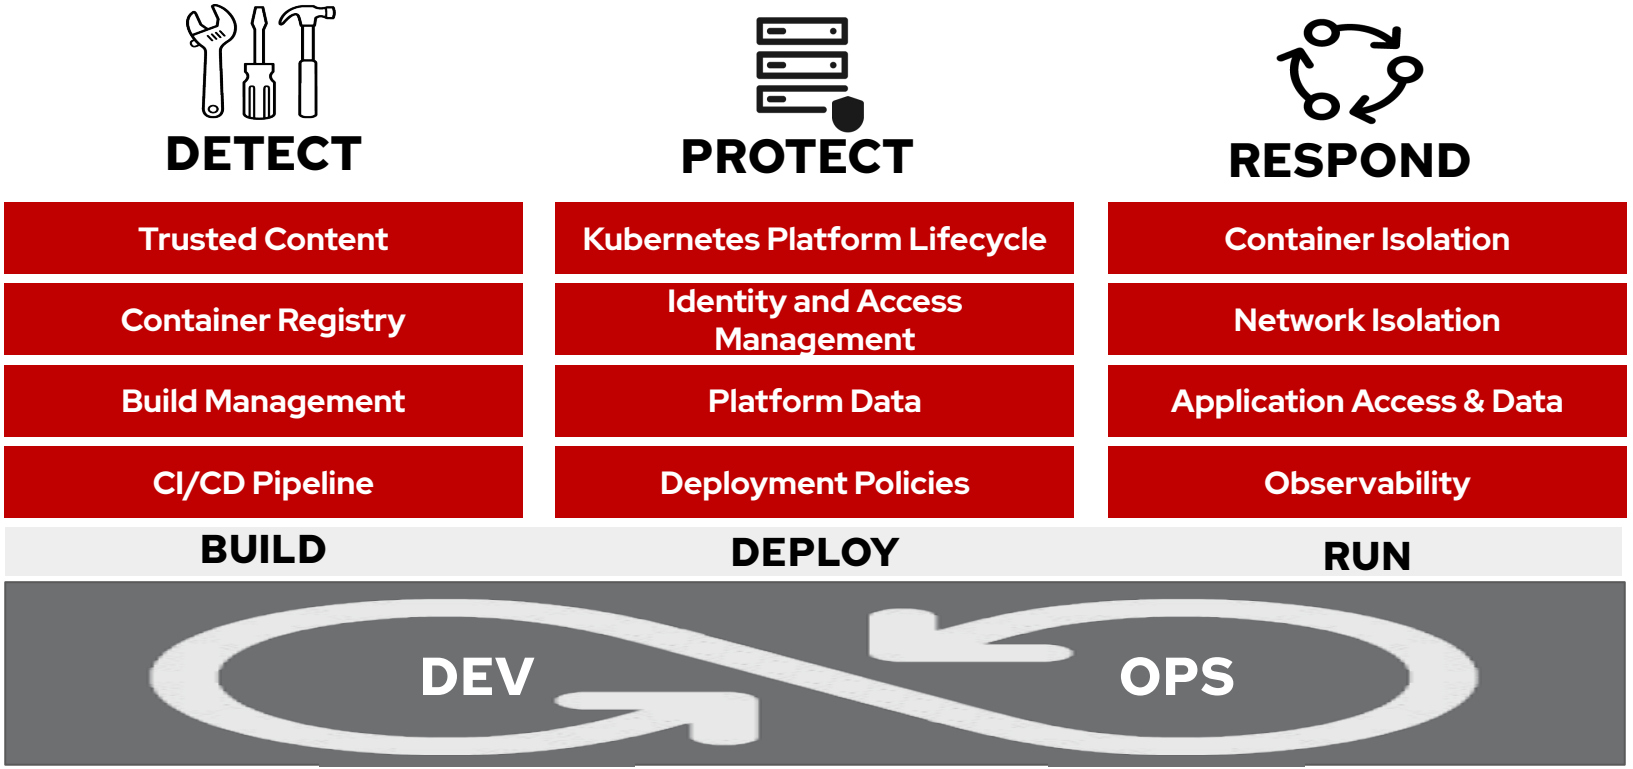

[A layered approach to container and Kubernetes security]

Red Hat

Hardening, applicability guides, certifications

OpenShift 4

- Available now
  - HIPAA
  - ISO 27001 (ask RH for a copy)
  - FISMA
  - The OpenShift Security Guide
  - OpenShift 4 Hardening Guide (ask RH for a copy)
  - PCI-DSS
- Target Q2 CY 2021
  - CIS OpenShift Benchmark
  - HITRUST

Managed Services certifications

- SOC2-type 1, SOC2-type 2
  - OpenShift Dedicated (OSD) on AWS
  - ARO, IBM ROKS, ROSA
  - In process for OSD on GCP
- ISO-27001
  - OSD on AWS, ARO
- PCI-DSS
  - ARO, IBM ROKS
  - In process for OSD on AWS and GCP, ROSA
- FedRAMP
  - ARO, IBM ROKS
  - In process for OSD on AWS, ROSA
- HIPAA and/or HITRUST
  - ARO, IBM ROKS
  - In discussion for OSD and ROSA

OpenShift Container Platform

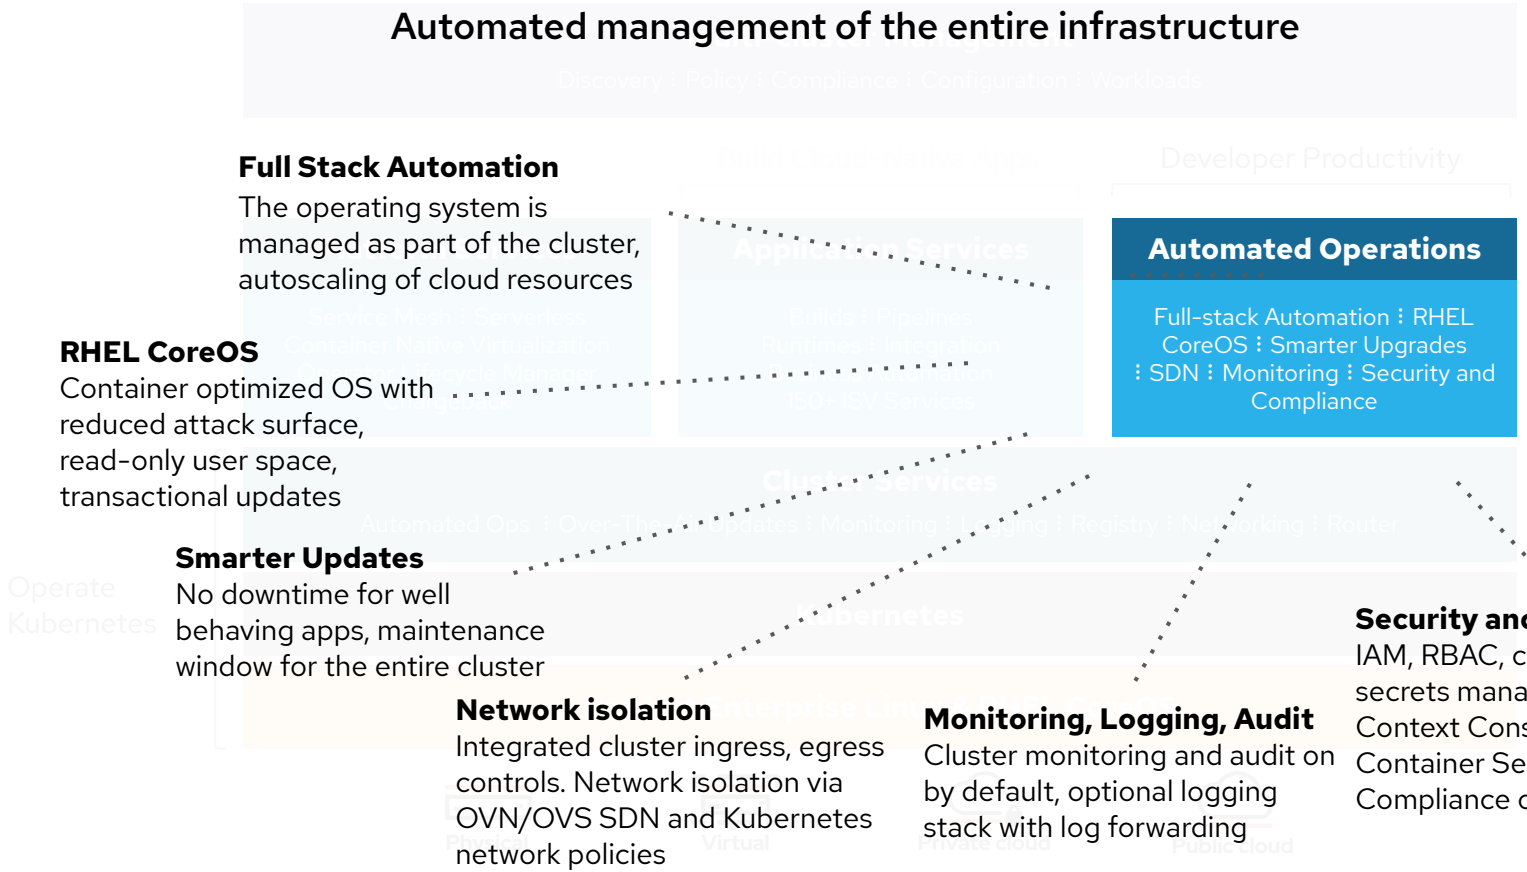

Full Stack Automation
The operating system is managed as part of the cluster, autoscaling of cloud resources

RHEL CoreOS
Container optimized OS with reduced attack surface, read-only user space, transactional updates

Smarter Updates
No downtime for well behaving apps, maintenance window for the entire cluster

Network isolation
Integrated cluster ingress, egress controls. Network isolation via OVN/OVS SDN and Kubernetes network policies

Monitoring, Logging, Audit
Cluster monitoring and audit on by default, optional logging stack with log forwarding

Security and Compliance
IAM, RBAC, certificate and secrets management, Security Context Constraints. Container Security operator, Compliance operator

Red Hat

OpenShift Platform Services

Platform Services to manage workloads and tie them into OpenShift infra capabilities

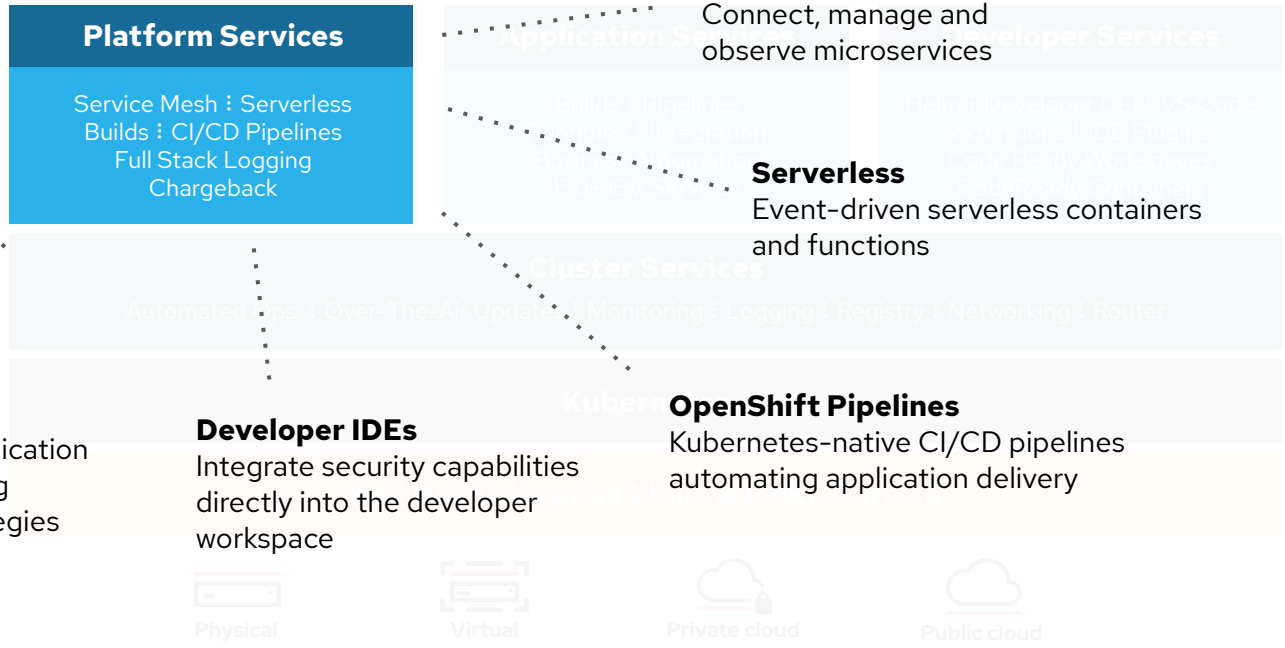

Service Mesh
Connect, manage and observe microservices

Serverless
Event-driven serverless containers and functions

OpenShift Build
Build images from application source and binary using Kubernetes build strategies

Developer IDEs
Integrate security capabilities directly into the developer workspace

OpenShift Pipelines
Kubernetes-native CI/CD pipelines automating application delivery

PCI-DSS RECOMMENDATIONS

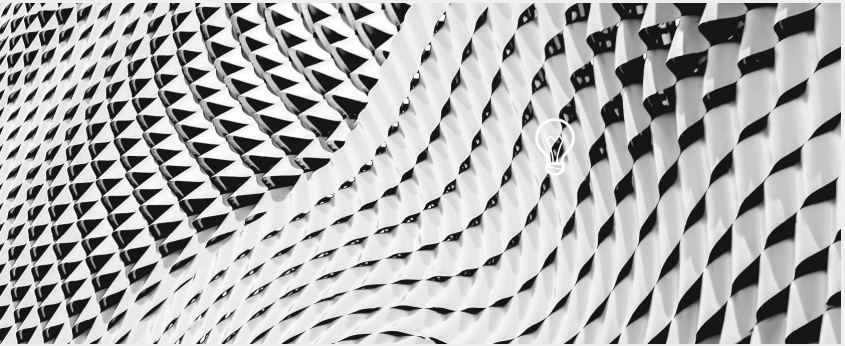

Red Hat

Coalfire Product Applicability Guide & Reference Architecture

- 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*
- Technical requirements 1, 2, 5, 6, 7, 8, 10, 11 are applicable, discussed

*Coalfire is working on a new version of this guide for OpenShift 4

Coalfire conclusion

“OpenShift hosted on Red Hat Enterprise Linux, as reviewed by Coalfire, can be effective in providing support for the outlined objectives and requirements of PCI DSS v3.2.1. Through proper implementation and integration into the organization’s overall technical infrastructure and information management systems, OpenShift may be useable in a PCI DSS v3.2.1 controlled environment.”

Securing the container platform

- Configuration and lifecycle management
- Host & runtime security
- Identity and Access Management
  - Project namespaces
  - Role Based Access Controls
- Data at rest, data in transit
  - Ingress & egress controls
  - Encryption
- Logging, Monitoring, Metrics
- Audit and Compliance

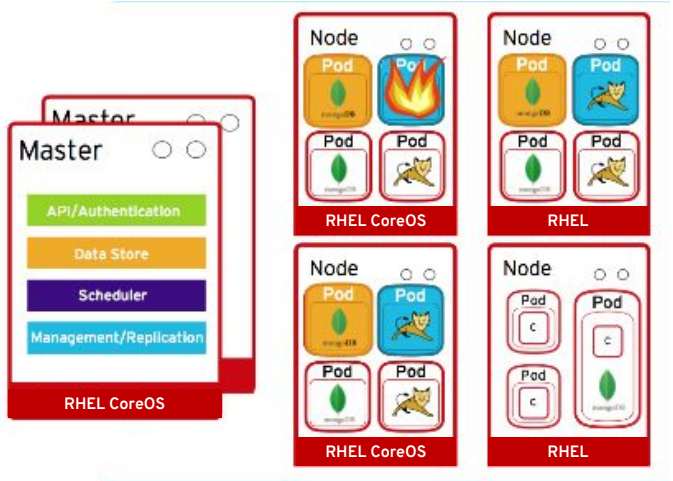

Red Hat

Automated Configuration and Lifecycle Management

Dramatically simplified for the Hybrid Cloud

Machines

Machines are complex for ops

Make machines easy
(like containers)

Configuration

Config change is risky

Make config management
and config change
easy and safe

Lifecycle

Software lifecycle is hard

Automate software
lifecycle on Kube

Automated Container Operations

FULLY AUTOMATED DAY-1 AND DAY-2 OPERATIONS

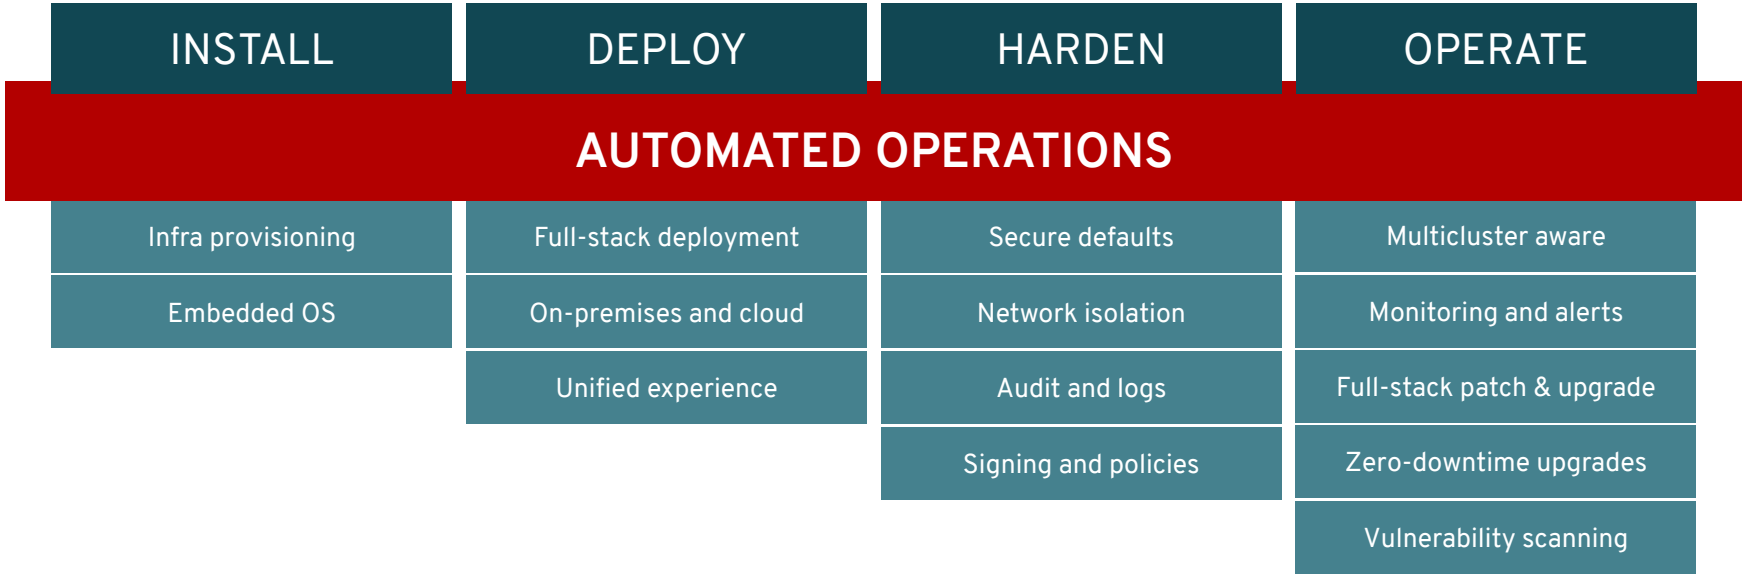

Red Hat

The Value Of Kubernetes Operators

No need for operator

Requires custom Operator built with SDK

Installation
Automated application provisioning and configuration management

Upgrades
Patch and minor version upgrades supported

Lifecycle
App lifecycle, storage lifecycle (backup, failure recovery)

Deep Insights
Metrics, alerts, log processing and workload analysis

Auto-pilot
Horizontal/vertical scaling, auto config tuning, abnormal detection, scheduling tuning...

Red Hat

REQUIREMENT 1

Install & maintain a firewall configuration

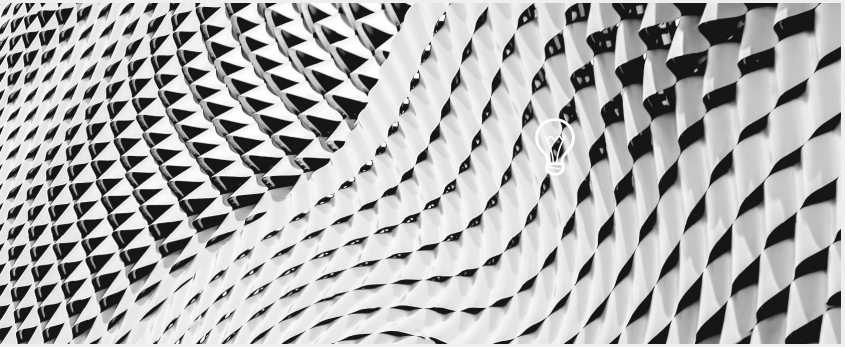

Red Hat

Requirement 1: Install & maintain a firewall configuration to protect cardholder data

"Requirement 1 is primarily concerned with traditional edge protections between the Internet or "untrusted networks" and internal networks. It is recommended that the OpenShift environment be placed in an internal controlled network that is protected with traditional edge protections provided by third-party solutions. ***As such, Coalfire determined that most, if not all, of these requirements were not pertinent to OpenShift's capabilities. However, assessors often look at implementation of firewalls and routers on internal networks used to isolate or segment workloads as a method of reducing assessment scope. With this in mind, many of the requirements may apply to the internal network elements that perform segmentation, including the SDN elements provided by OpenShift***."

Requirement 1 Applicability

- Protect with “traditional edge protections provided by third party solutions.”
- Use the available SDN network policies to provide micro-segmentation and isolation of workloads
- Use ingress and network policy objects to restrict inbound traffic
- OpenShift provides the ability to separate workloads onto different servers (nodes)
- Similarly, infrastructure pods (ingress and egress router) can be hosted on separate nodes from the master
- Use egress to restrict outbound traffic

OpenShift networking

- Built-in internal DNS to reach services by name
- Software Defined Networking (SDN) for a unified cluster network to enable pod-to-pod communication
- OpenShift follows the Kubernetes Container Networking Interface (CNI) plug-in model
- Isolate applications from other applications within a cluster
- Isolate environments (Dev / Test / Prod) from other environments within a cluster

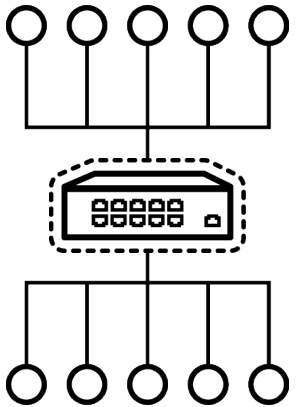

External Access to Cluster Resources

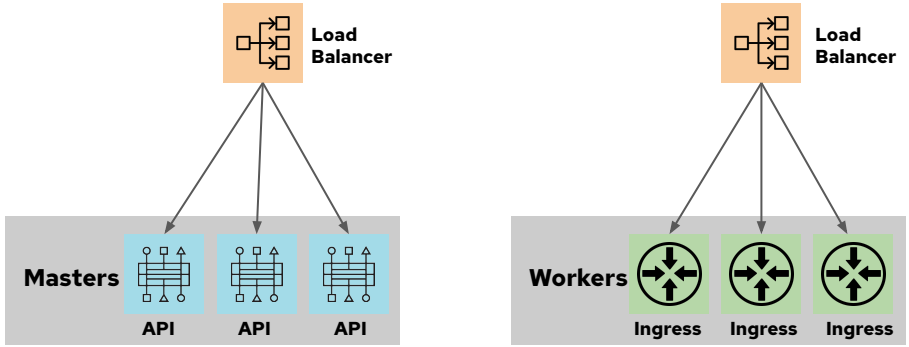

Ingress Traffic

- Two primary entry points into OpenShift
  - API
  - Ingress/Router
- Proper DNS entries must be configured
- Additional ingress types available
  - NodePort (requires additional port resources)
  - LoadBalancer

OpenShift cluster with multiple zones

Using multiple ingress controllers, network policies, multiple egress pods

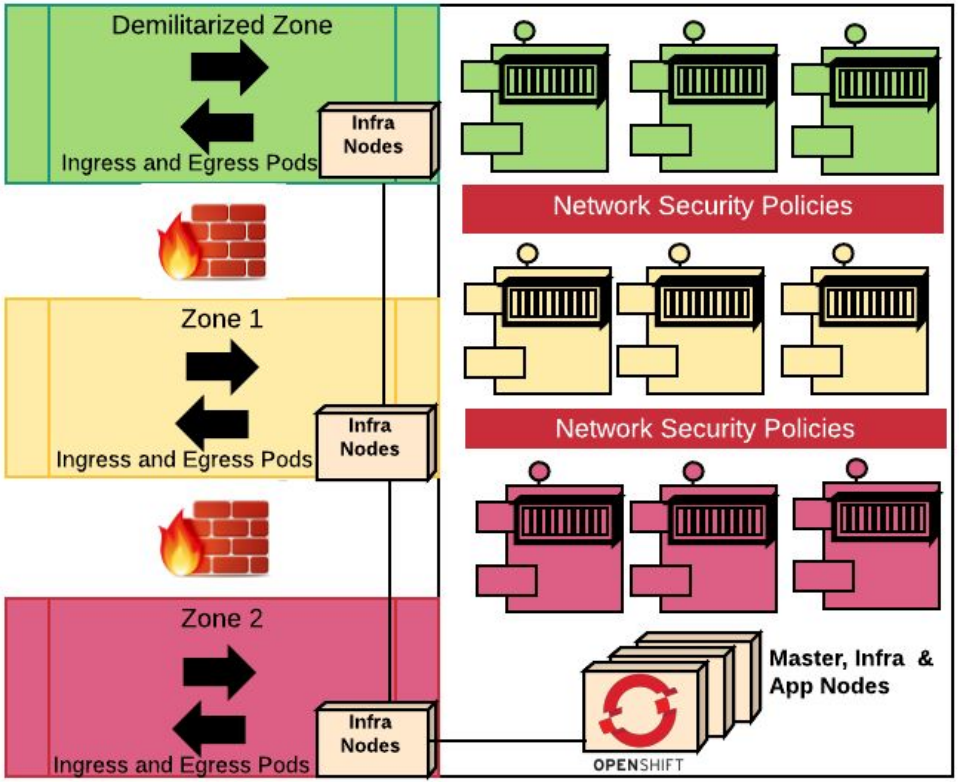

Application pods run on one OpenShift Cluster.

Microsegmented with Network Security policies.

Infra Nodes in each zone run Ingress and Egress pods for specific zones.

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

Red Hat

OpenShift SDN

Network policy enabled by default

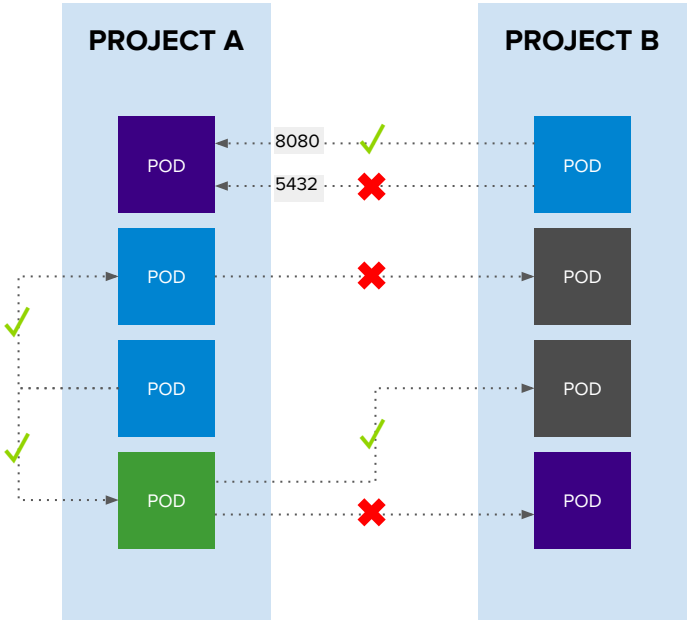

Example Policies

- 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
```


Red Hat

Controlling Egress Traffic

Egress IP high availability (multiple IPs)

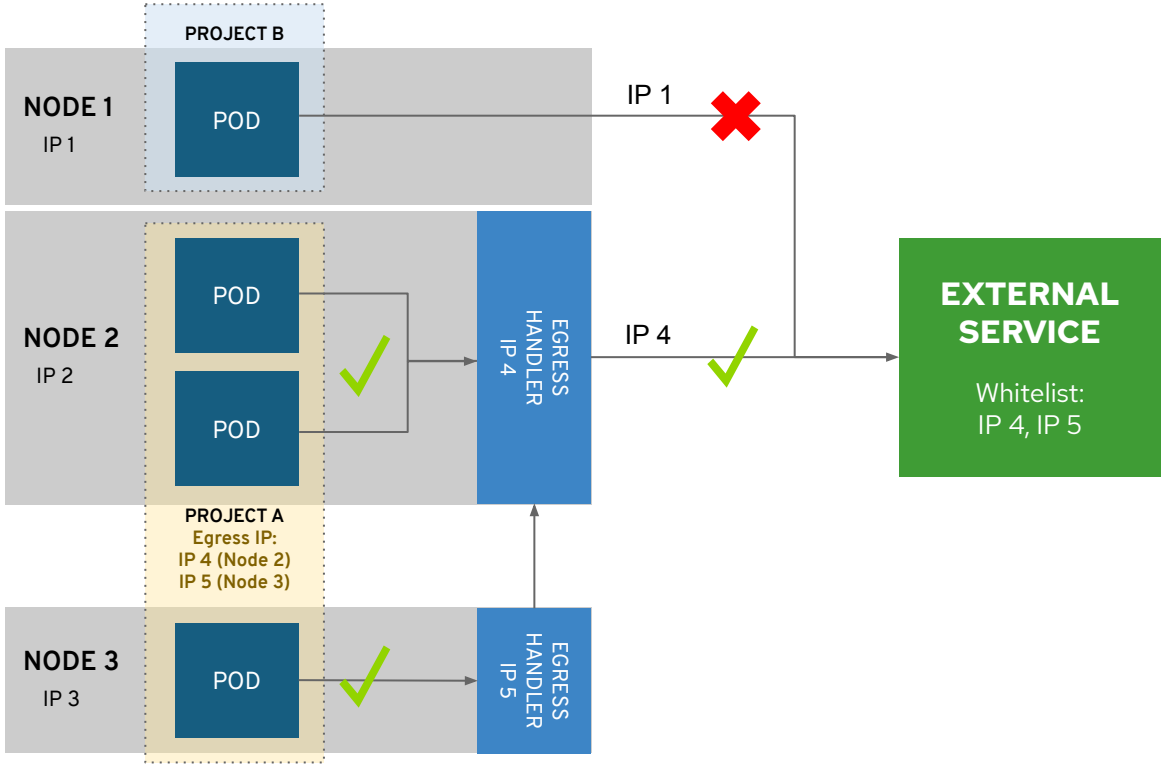

Red Hat

Controlling Egress Traffic

Egress IP high availability (multiple IPs)

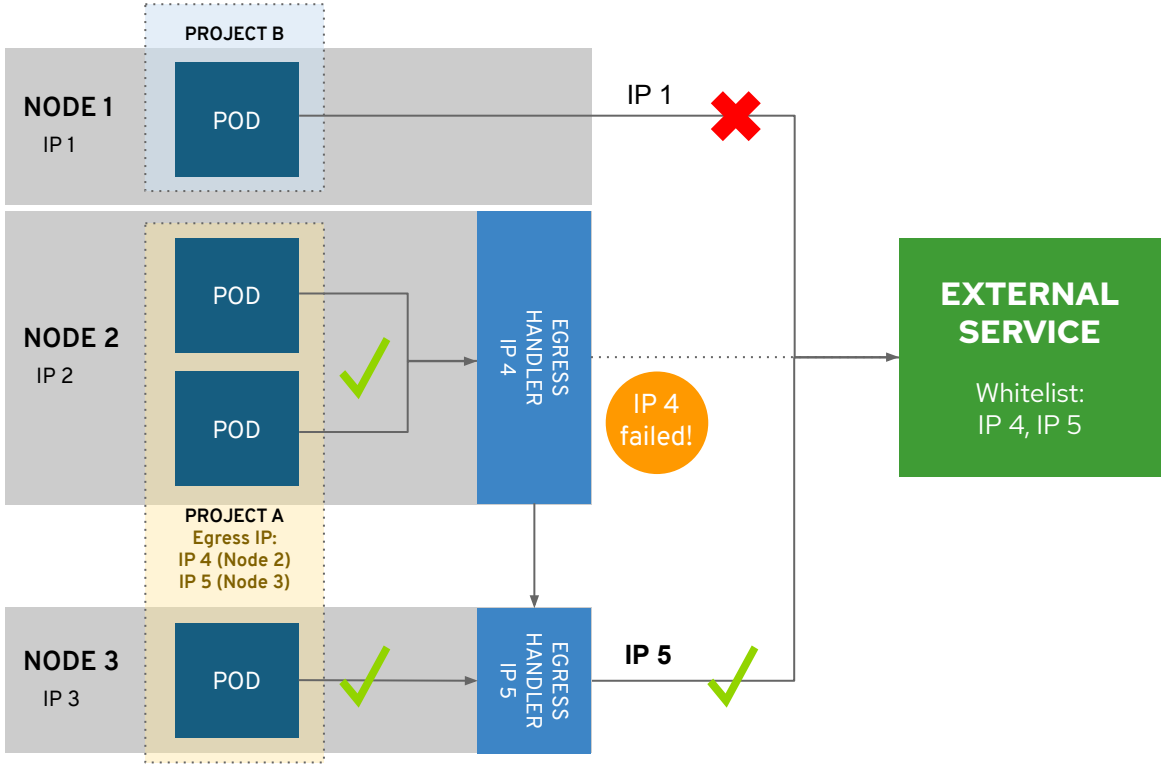

Red Hat

Egress firewall to limit access

to external addresses accessed by some or all pods from within the cluster

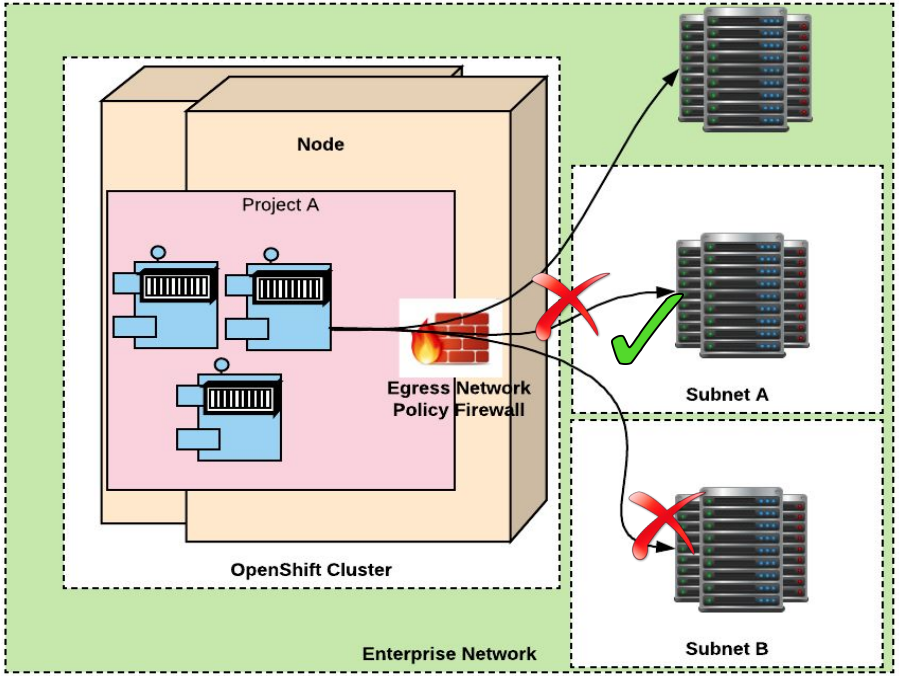

Examples:

A pod can talk to hosts (outside OpenShift cluster) but cannot connect to public internet

A pod can talk to public internet, but cannot connect to hosts (outside OpenShift cluster)

Public Internet

A pod cannot reach specific subnets/hosts

Red Hat

REQUIREMENT 2

Vendor Standards and Configs

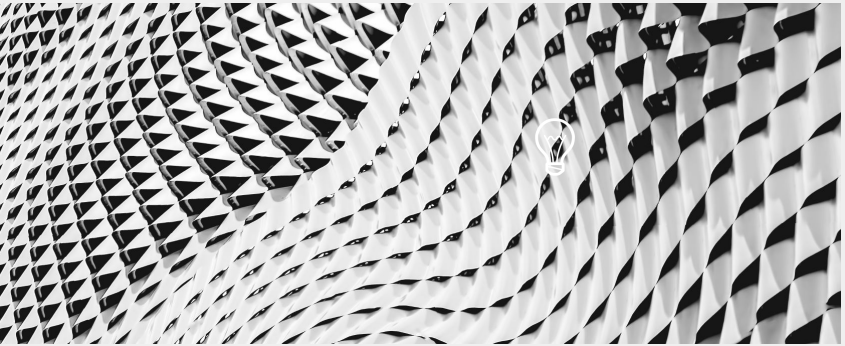

Red Hat

Requirement 2
Vendor defaults & configuration standards

2.1 Always change vendor-supplied defaults and remove or disable unnecessary default accounts

2.2 Develop config standards for all components

- 2.2.2 Enable only necessary services, protocols, daemons, etc., as required for the function of the system
- 2.2.3 Implement additional security features as required. For example, encryption
- 2.2.4 Configure system security parameters to prevent misuse
- 2.2.5 Remove all unnecessary functionality

2.3 Encrypt all non-console administrative access

Requirement 2 Applicability

- No vendor provided default passwords are in use
- Configure an external identity provider, create your own cluster admin, remove kubeadmin user
- By default, RHEL CoreOS is a container-optimized OS and includes only the components needed to run OpenShift 4
- Encrypt RHEL CoreOS volumes
- Encrypt the etcd datastore
- By default, all communication between the control plane and the data plane is encrypted
- Install the Compliance operator (requires 4.6)

Container host vision

An Ideal Container Host would be	RHEL CoreOS
Minimal	Only what's needed to run containers
Secure	Read-only & locked down
Immutable	Immutable image-based deployments & updates
Always up-to-date	OS updates are automated and transparent
Updates never break my apps	Isolates all applications as containers
Updates never break my cluster	OS components are compatible with the cluster
Supported on my infra of choice	Inherits majority of the RHEL ecosystem
Simple to configure	Installer generated configuration
Effortless to manage	Managed by Kubernetes Operators

Red Hat Enterprise Linux CoreOS

The Immutable Container Optimized Operating System

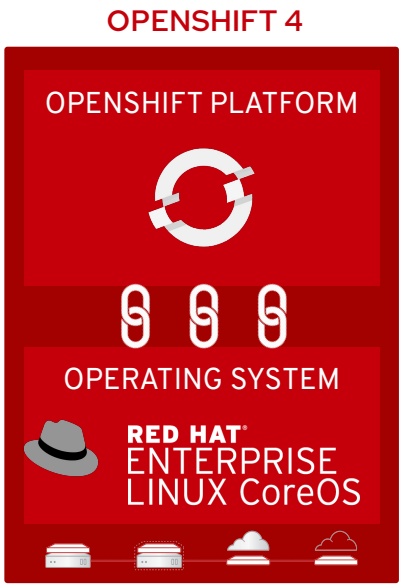

Role in OpenShift Ecosystem

- Versioned and validated for specific OpenShift version
- Required for masters. RHEL option for workers
- User space read-only

Managed by the OpenShift Cluster

- Considered a member of an OpenShift Deployment
- Configuration managed by the Machine Config Operator
  - Container runtime
  - Kubelet configuration
  - Authorized container registries
  - SSH Configuration

Red Hat

cri-o

A lightweight, OCI-compliant container runtime

Optimized for Kubernetes

Any OCI-compliant container from any OCI registry (including docker)

Improve Security and Performance at scale

[CRI - the Container Runtime Interface]
[OpenShift 4 defaults to CRI-O]
[Red Hat contributes CRI-O to the Cloud Native Computing Foundation]

Key characteristics of RHEL CoreOS

- **Transactional updates** - RHCOS is distributed as an image and each operating system update is versioned and distributed as containers. Major releases (and some z stream releases) provide new boot images. The OS always boots into a known-good version; this is similar in principle to how container images are managed and deployed.
- **Immutable management** - RHCOS is built to be managed in an immutable fashion by the Machine Config Operator and Kubernetes API. While certain parts of the OS are truly immutable, others are not. Immutable management enables us to spawn new nodes and ensure that the cluster is the single source of truth for provisioning configurations, OS versions, and run-time configuration. Apart from consistency, this also enables elastic clusters to spawn and destroy nodes.

Key characteristics of RHEL CoreOS (cont'd)

- **Applications need to run in containers** - Installing RPMs on RHCOS is not supported. The OS is built to run all processes outside the OS as a container. This allows us to guarantee successful upgrades and automation beyond what a traditional operating system can deliver.
- **rpm-ostree** - This is the technology used to assemble the operating system. RHEL RPMs are used to create the OS images, and versions can easily be queried using the rpm command.
  - **/usr** is where the operating system binaries and libraries are stored and is read-only.
  - **/etc**, **/boot**, **/var** are writable on the system but only intended to be altered by the Machine Config Operator.
  - **/var/lib/containers** is the graph storage location for storing container images.

RHEL CoreOS management

- **Regular management** of the underlying RHCOS cluster nodes is designed to be performed via the OpenShift API itself.
- **The only users** that exist on an RHCOS OpenShift node are *root* and *core*.
  - A user named core is created, with your ssh key assigned to that user. This allows you to log in to the cluster with that user name and your credentials.
  - The core user has permission to run privileged commands.
  - Adding additional users at the node level is highly discouraged.
  - **Updates are managed through the OpenShift MachineConfigOperator**
  - OS upgrades are delivered as an atomic unit.
  - The new OS deployment is staged during upgrades and goes into effect on the next reboot.
  - RHCOS upgrades in OpenShift Container Platform are performed during cluster updates.

OpenShift Cluster Management

- OpenShift Container Platform **creates the kubeadmin user** after the installation process completes.
- The **kubeadmin user has the cluster-admin role automatically applied** and is treated as the root user for the cluster. The password is dynamically generated and unique to your OpenShift Container Platform environment. After installation completes the password is provided in the installation program's output.
- **After you define an identity provider and create a new cluster-admin user, you can remove the kubeadmin to improve cluster security**.
- Cluster configuration changes are managed through cluster operators.

For more information see:
https://docs.openshift.com/container-platform/4.5/authentication/understanding-identity-provider.html

Encrypted control plane communication

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

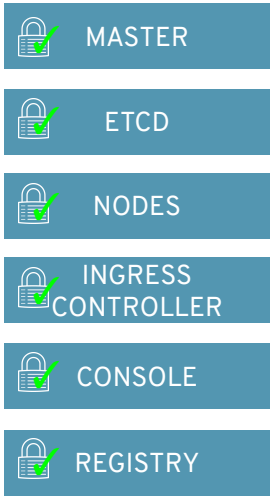

Red Hat

Encrypt secrets in transit and at rest

- Secure mechanism for holding sensitive data, such as
  - Passwords and credentials
  - SSH Keys
  - Certificates
- Secrets are made available as
  - Environment variables
  - Volume mounts
  - Interaction with external systems (e.g. vaults)
- Encrypted in transit and at rest
  - Encrypt the etcd datastore
  - Encrypt RHCOS volumes
- Never rest on the nodes

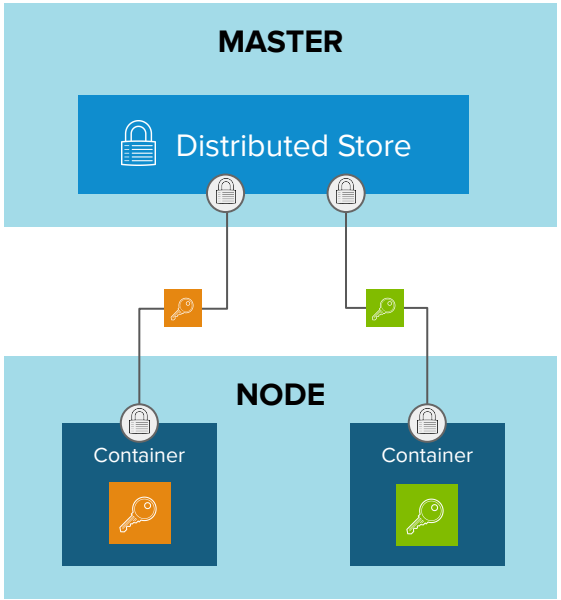

Red Hat

Volume Encryption

Network Bound Disk Encryption

- Provides encryption for local storage
- Addresses disk/image theft
- Platform/cloud agnostic implementation
- TPM/vTPM (v2) and Tang endpoints for automatic decryption

Openshift Compliance Operator: Declarative Security Compliance

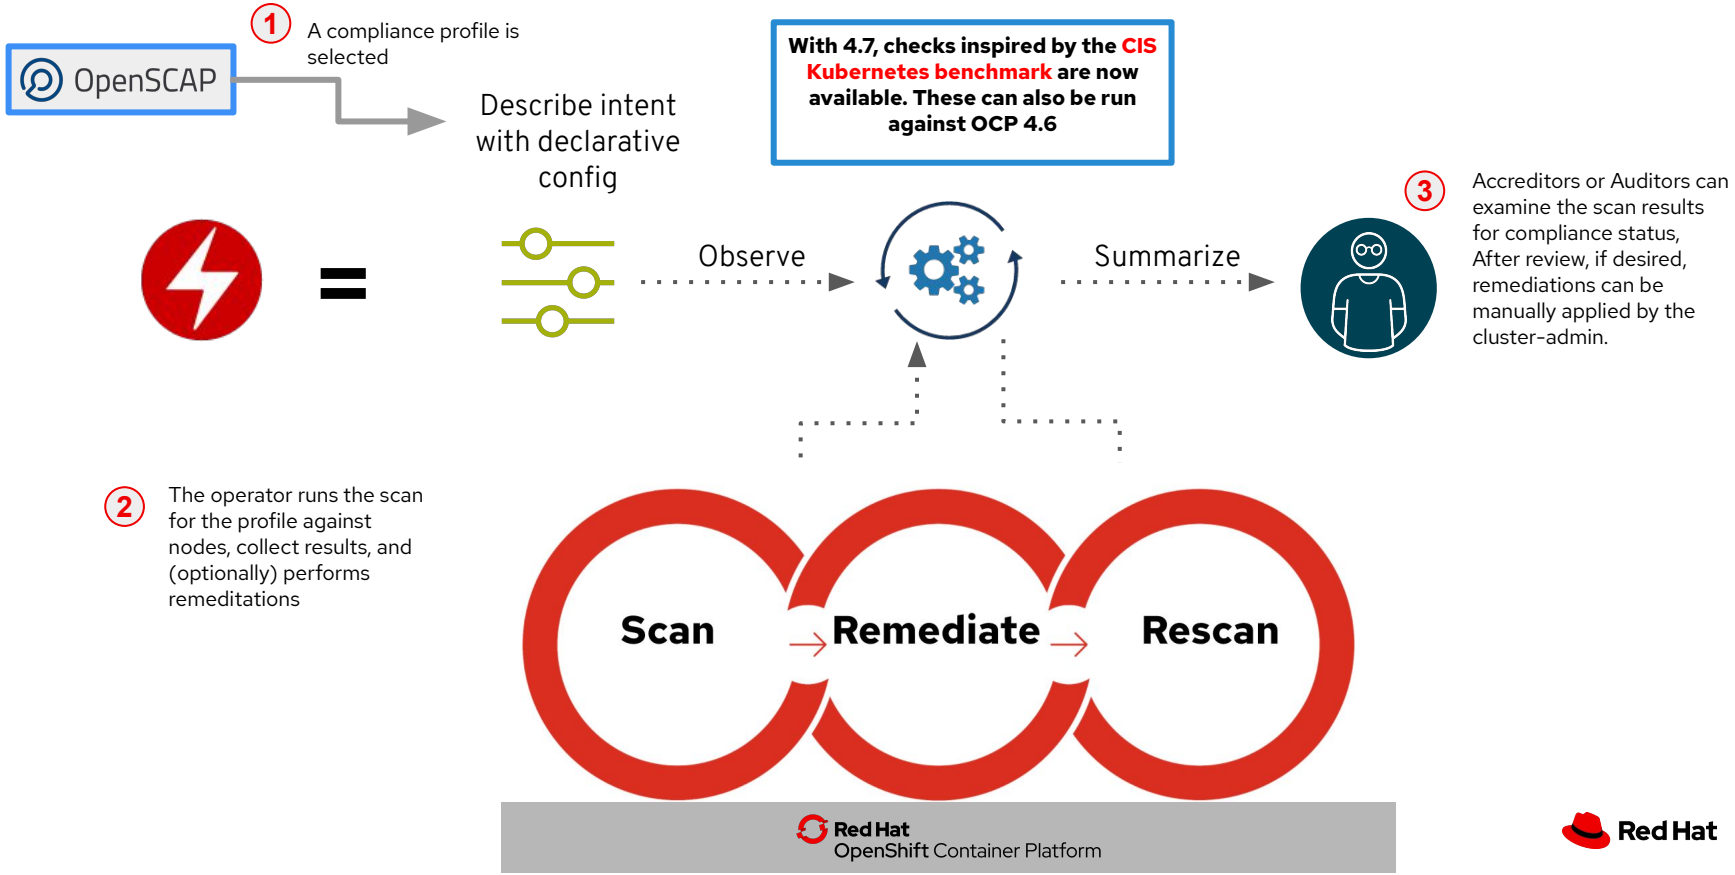

Red Hat

RECOMMENDATION 2

Protect against malware

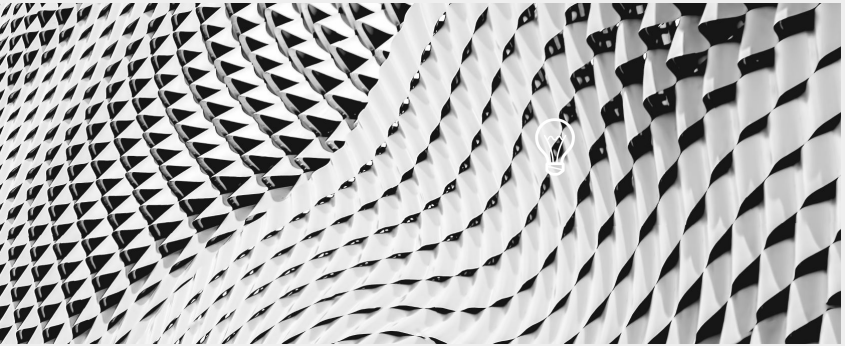

Red Hat

Requirement 5

Protect all systems against malware and regularly update anti-virus software or programs

Requirement 5: Considerations & Recommendations

- Look for a malware and/or anti-virus solution that is designed to work with containers or container images.
- Deploy the container security opertator
- Use Security Context Constraints
  - Containers cannot run as root by default
- Monitor for deprecated nodes
- Deploy the file integrity operator
- Container runtime security solutions (behavioral analysis) are available from Red Hat partners such as
  Aqua Security, Neuvector, Palo Alto (Twistlock), Sysdig

RHCOS & anti-virus scanners

- Solutions, such as anti-virus scanners, can be deployed as daemonsets or container images. However, few, if any anti-virus vendors are delivering their software in this form.
- We recommend talking with your anti-virus vendor and asking what solutions they have available for a container-optimized OS.
- It's still RHEL
  - To see the **packages installed**, run the same commands as on a RHEL system: `rpm -qa |sort`
  - To see which **ports are in use**: `ss -tulpn`
- However, in a 2019 paper from Gartner on Cloud Workload Protection Platforms, recommends that clients "Replace antivirus (AV)-centric strategies with a "zero-trust execution"/default deny/application control approach to workload protection where possible, even if used only in detection mode." [1]

1. Gartner: Market Guide for Cloud Workload Protection Platforms, ID G00356240, April 8, 2019

View Security Vulnerabilities with the Quay Operator

See all your Container Vulnerabilities right from the Console Dashboard

- Link out to **Red Hat Quay** for more in depth information
- The Quay Operator supports both **On-premise and External** Quay Registries
- Currently uses **Clair for Security Scan**; Planning to expand to other Vendors(TwistLock, Aqua, e.g.)
- *Only works for images managed by Quay*

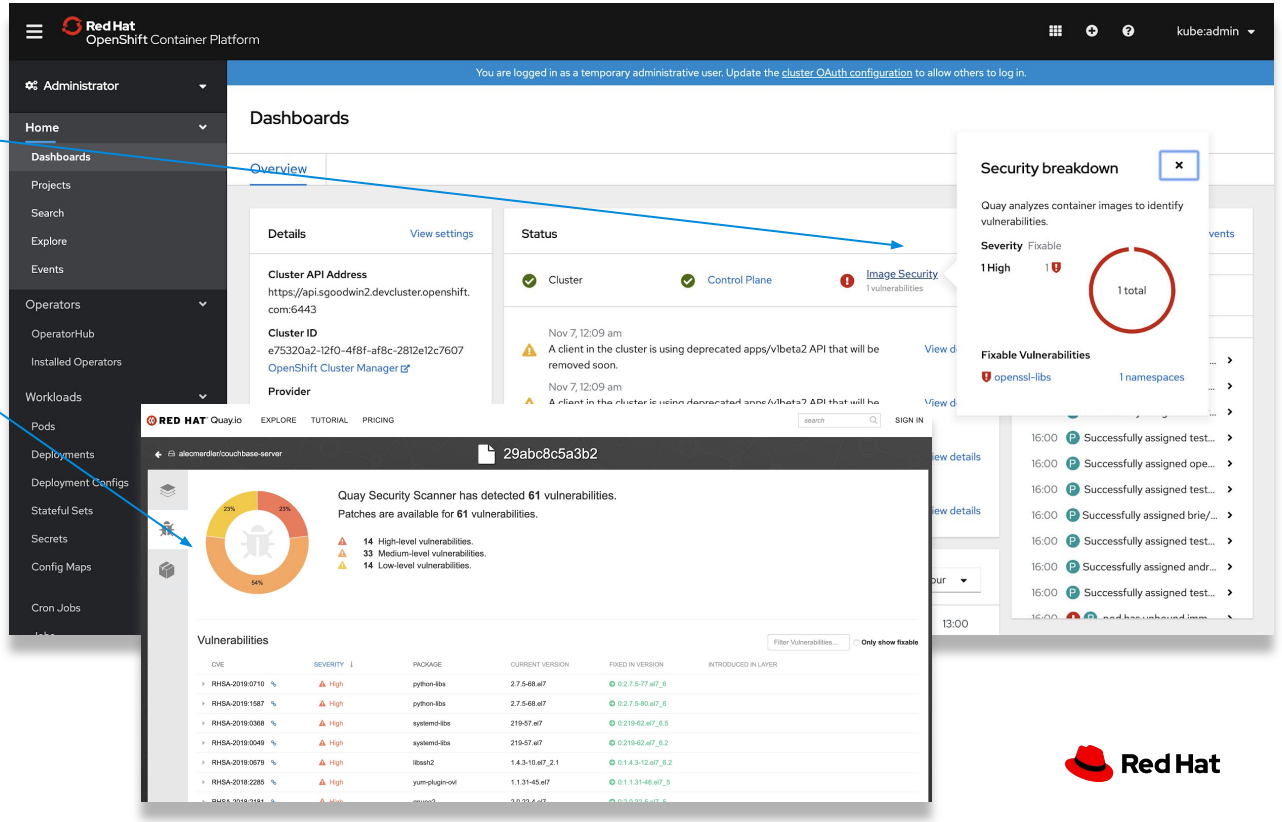

Red Hat

Container security starts with Linux security

- Security in the RHEL host applies to the container
- RHEL enables container multitenancy
- 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

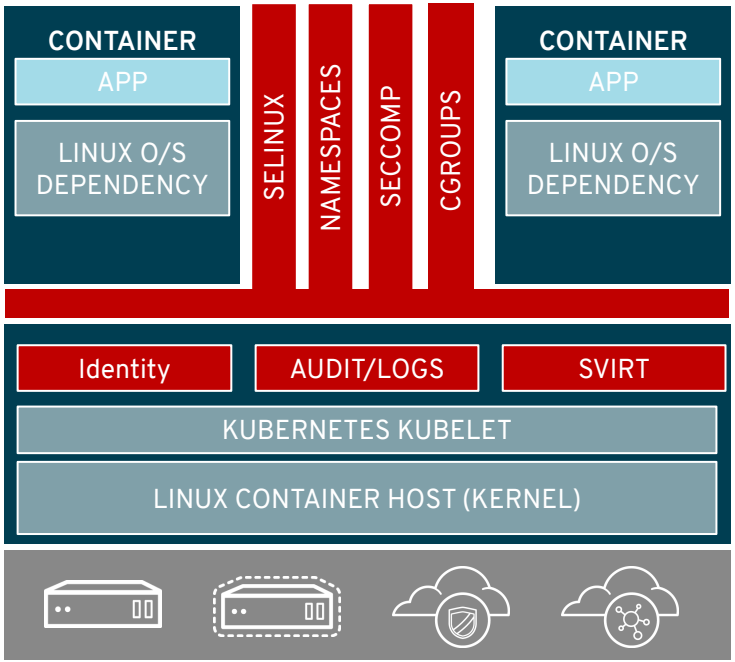

Red Hat

SELinux mitigates container runtime vulnerabilities

SELinux Mitigates container Vulnerability

January 13, 2017 | Joe Brockmeier

< Back to all posts Tags

A new CVE, (CVE-2016-9962), for the docker container runtime and runc were rec
released. Fixed packages are being prepared and shipped for RHEL as well as Fedo
CentOS. This CVE reports that if you `exec` d into a running container, the processe
the container could attack the process that just entered the container.

https://www.redhat.com/en/blog/selinux-mitigates-container-vulnerability

Latest container exploit (runc) can be blocked by SELinux

February 28, 2019 | Dan Walsh

< Back to all posts Tags: *Security*, *Containers*

A flaw in runc (CVE-2019-5736), announced last week, allows container processes to "escape" their containment and execute programs on the host operating system. The good news is that well-configured SELinux can stop it.

https://www.redhat.com/en/blog/latest-container-exploit-runc-can-be-blocked-selinux

Runtime security policies

(Pod Security Policies / Security Context Constraints)

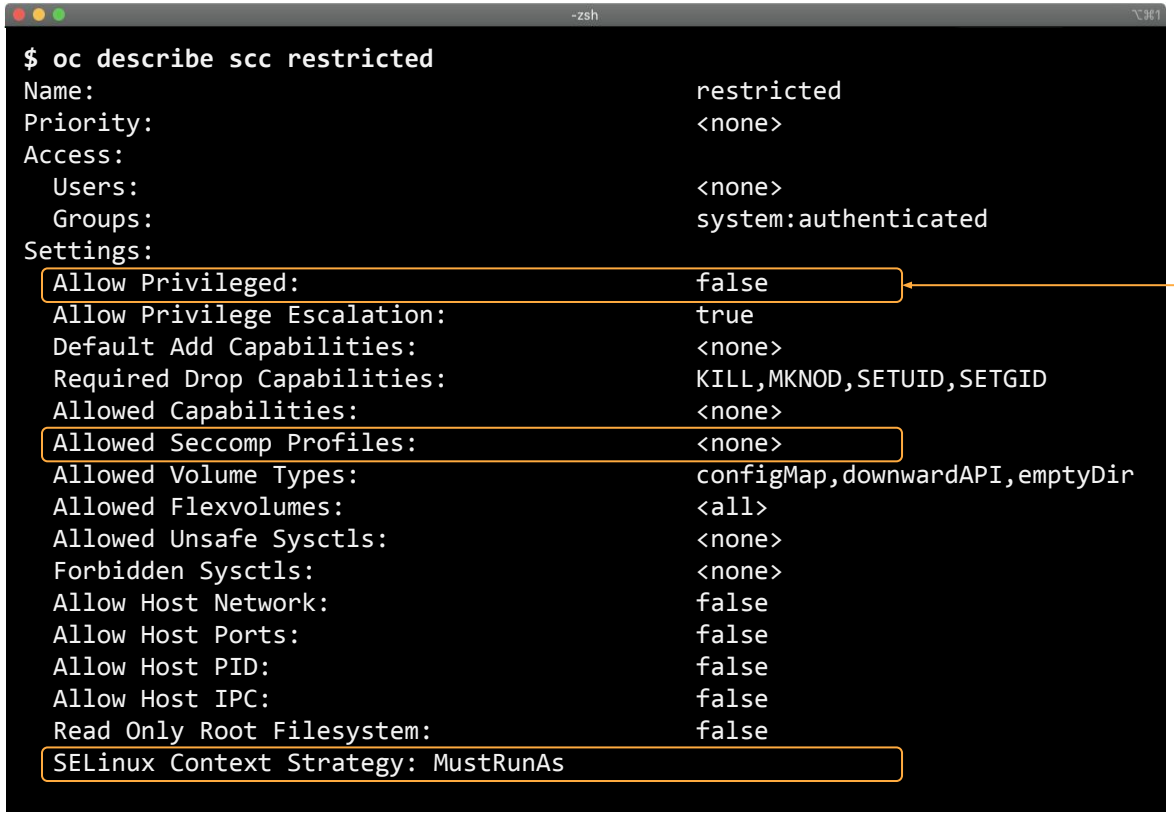

- Allow administrators to control permissions for pods
- By default, ensure no containers can run as root
- Admin can grant access to privileged PSP / SCC
- 8 included. Custom SCCs can be created

Red Hat

File Integrity

- Secure Boot - provides guarantee that a trusted, unmodified Kernel is loaded
- File integrity monitoring
  - /usr is read only
  - Machine Config Operator marks nodes with wrongly configured files as degraded
- Optional OpenShift File Integrity Operator using AIDE
  - Advanced Intrusion Detection Environment is a utility that creates a database of files on the system, and then uses that database to ensure file integrity and detect system intrusions

Red Hat OpenShift certified operators - Security

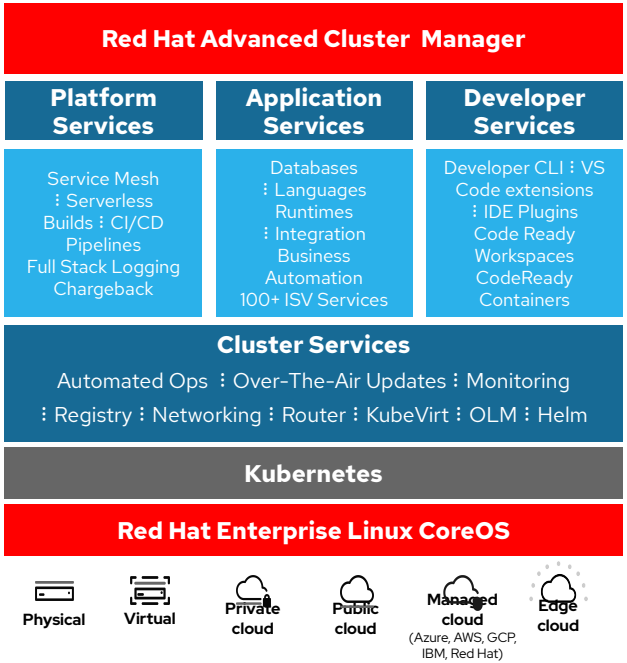

Red Hat

REQUIREMENT 6

Develop and maintain secure applications

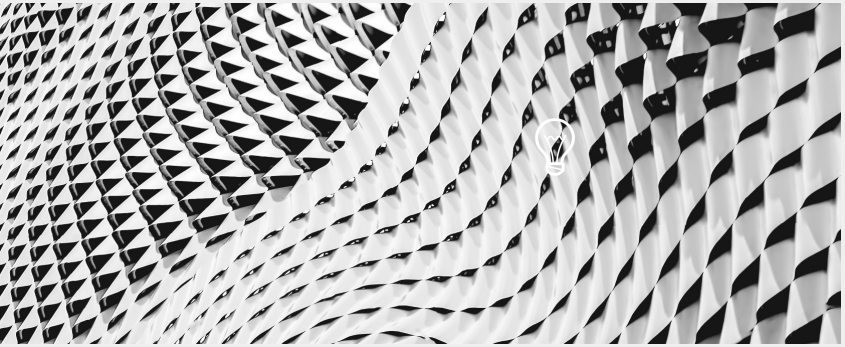

Red Hat

Requirement 6: Develop & Maintain Secure Applications

6.1 Establish a process to identify security vulnerabilities

6.2 Ensure that all system components and software are protected by applying applicable vendor-supplied patches

6.4 Follow change control policies for all changes to system components (GitOps)

Requirement 6 Applicability

- “Immutable containers are containers that will never be changed while running.”
- The best practice for patching containerized applications is to rebuild the container image and redeploy from the image.
- Separate development, test and production deployments with OpenShift projects and deploy to separate hosts if appropriate (node-selector)
- Deploy OpenShift Service Mesh to encrypt service to service communication and add protection for application ingress and egress
- Keep up with OpenShift z stream releases

Securing Containerized Applications

An opportunity to shift security left

Best practices

- Use trusted sources for external content
- Use a private registry to manage images
- CI/CD must have security gates
- Application secrets management
- Apply runtime security policies
- Rebuild and redeploy - never patch a running container
- Ensure application logging, monitoring

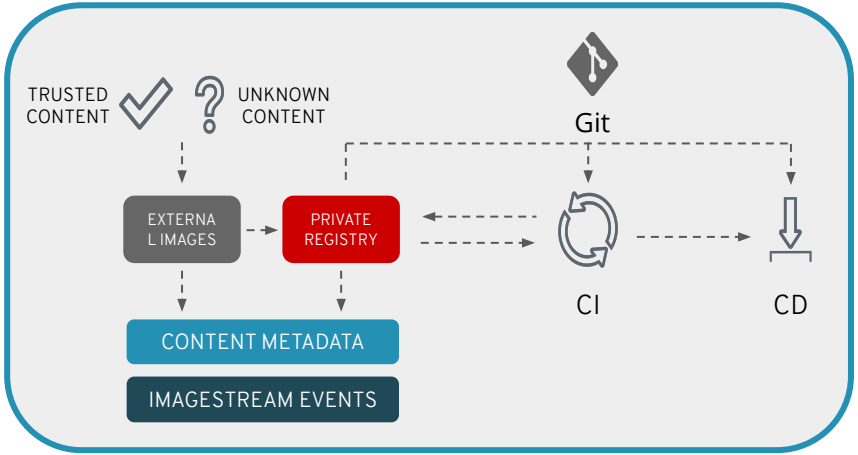

Red Hat

Secure & Automate The Content Lifecycle

Elements of the Openshift container pipeline

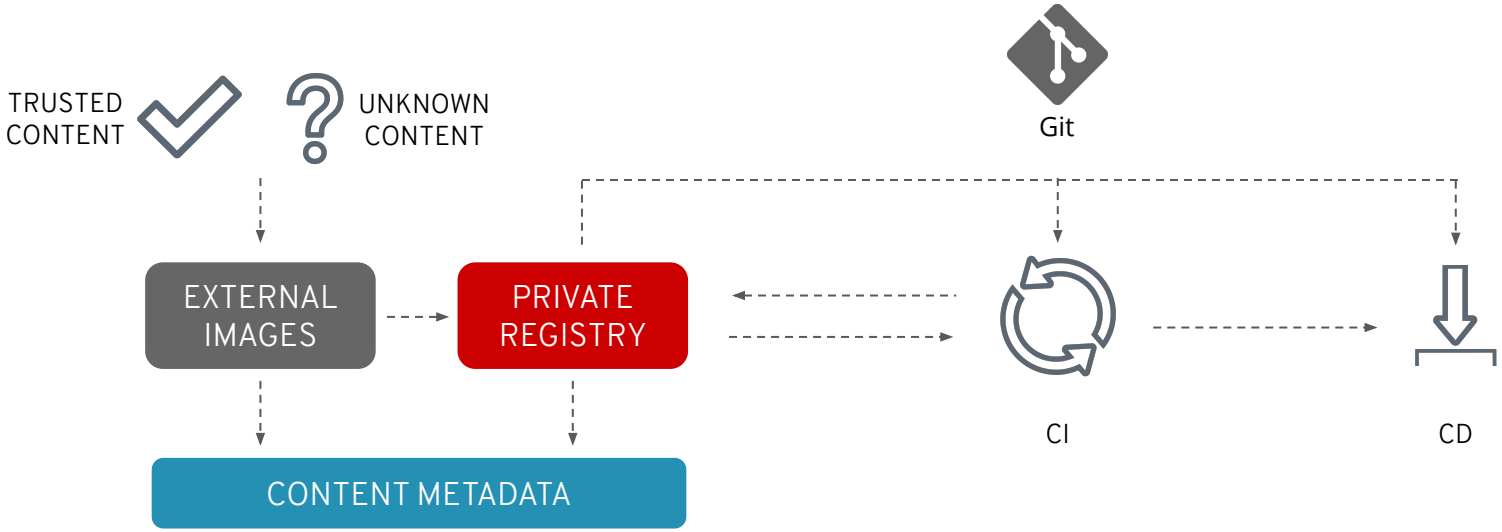

Red Hat

Trust is temporal: rebuild and redeploy as needed

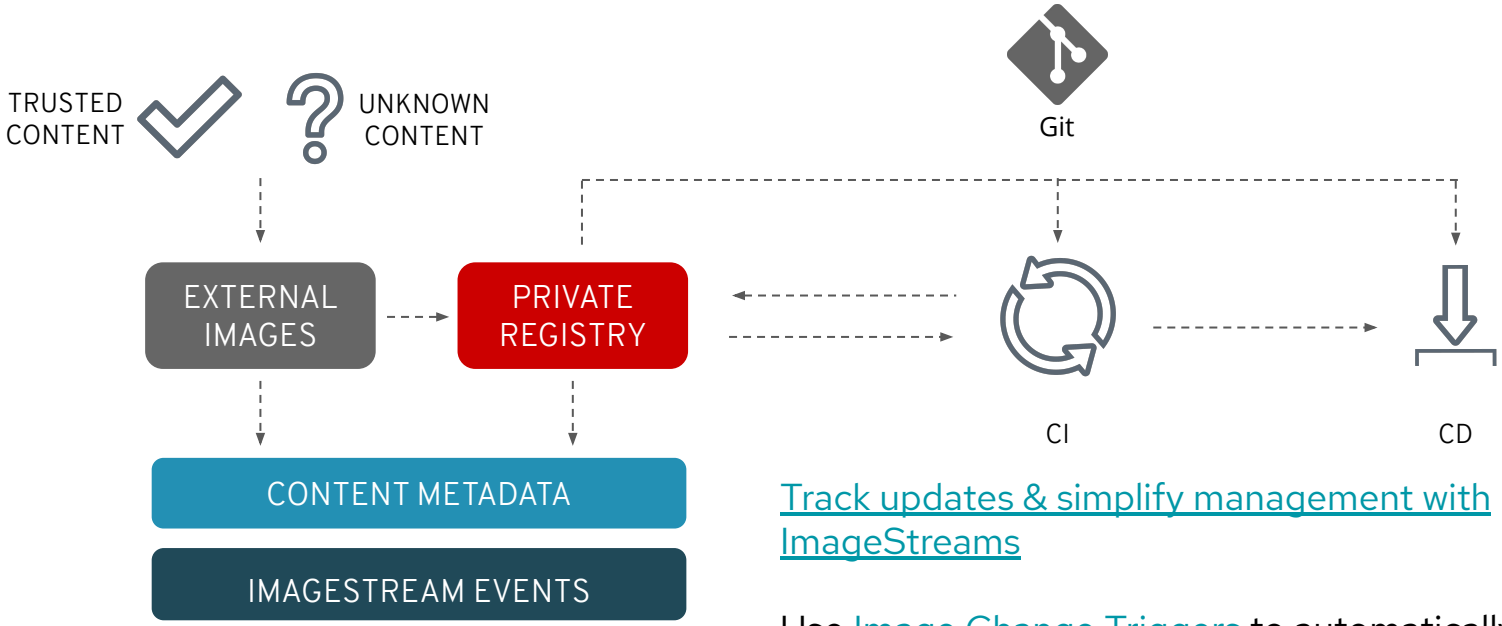

Track updates & simplify management with ImageStreams

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

Red Hat

External Content: Use Trusted Sources

Red Hat Container Images

- Signed Images
- Health Index
  (A to F grade)*
- Security advisories &
  errata (patches)

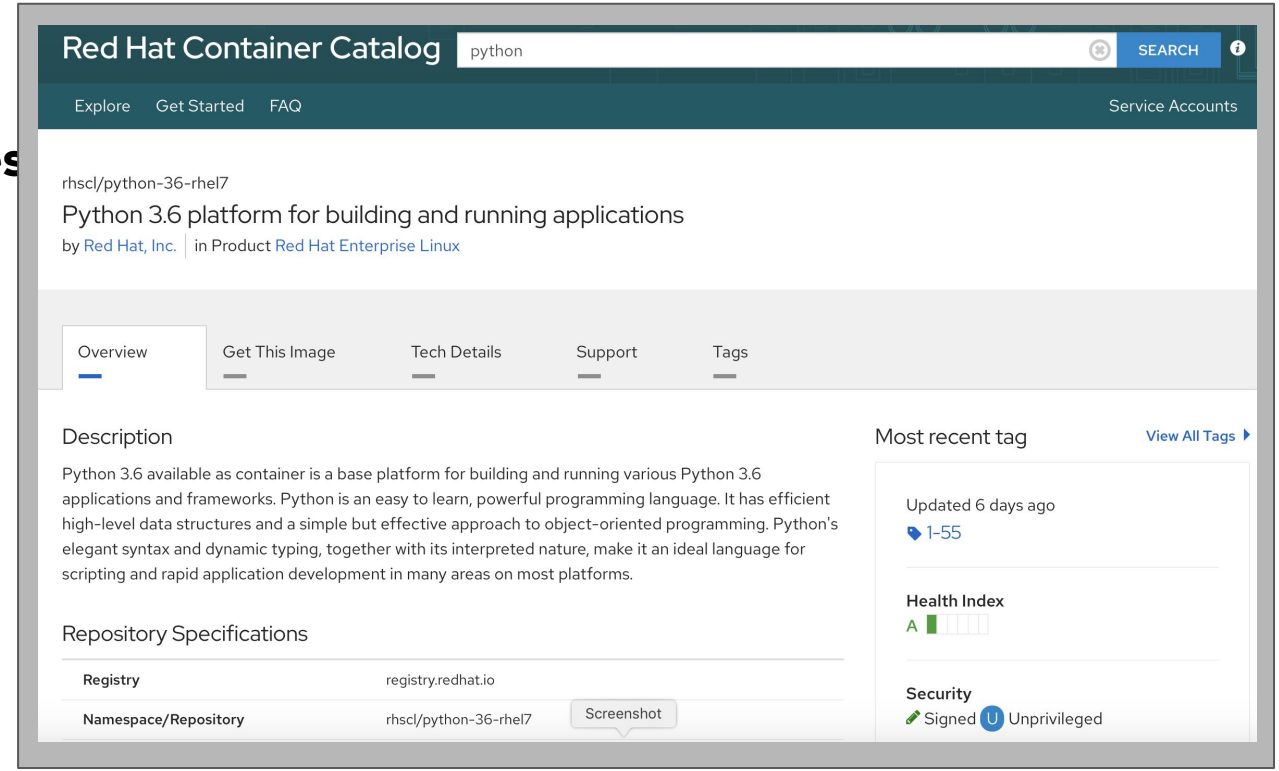

*More about the Health Index

Red Hat

The Red Hat Universal Base Image

The base image for all your needs -- enterprise architecture, security and performance

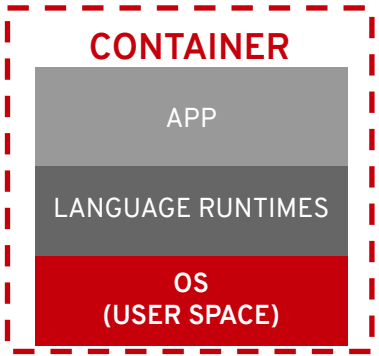

The Red Hat Universal Base Image is based on RHEL and made available at no charge by a new end user license agreement.

Development

- Minimal footprint (~90 to ~200MB)
- Programming languages (Modularity & AppStreams)
- Enables a single CI/CD chain

Production

- Supported as RHEL when running on RHEL
- Same Performance, Security & Life cycle as RHEL
- Can attach RHEL support subscriptions as RHEL

Red Hat

Supportability matrix

Red Hat Support and Community Support

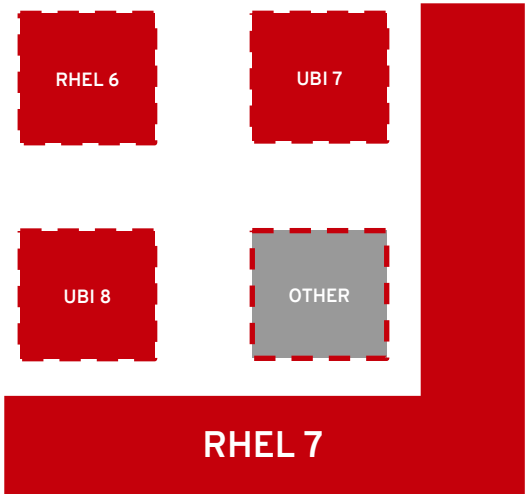

Red Hat Enterprise Linux 7

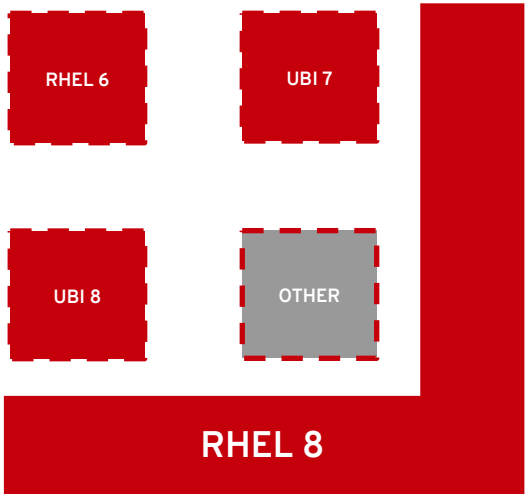

Red Hat Enterprise Linux 8

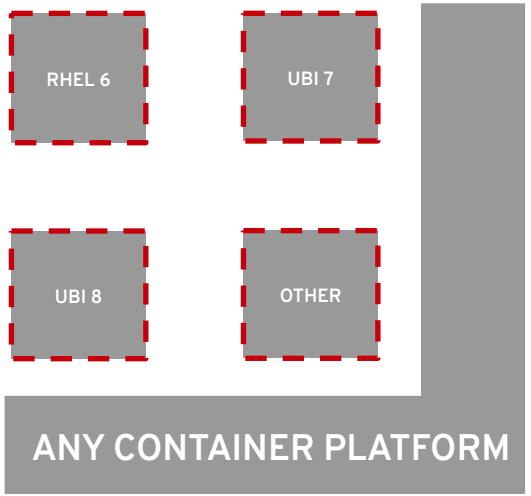

Like any upstream project

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

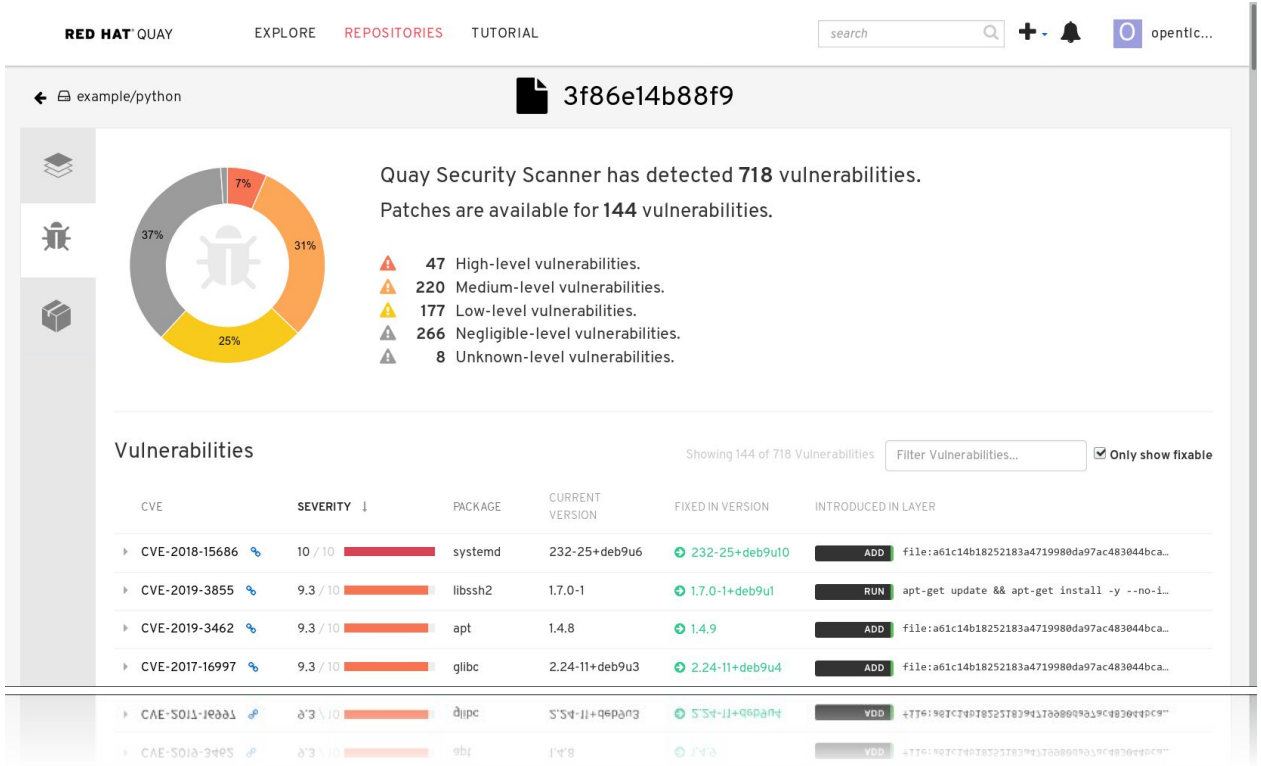

Integrate Security in your CI/CD Pipeline

Automated quality and security: because you can't inspect quality into a product

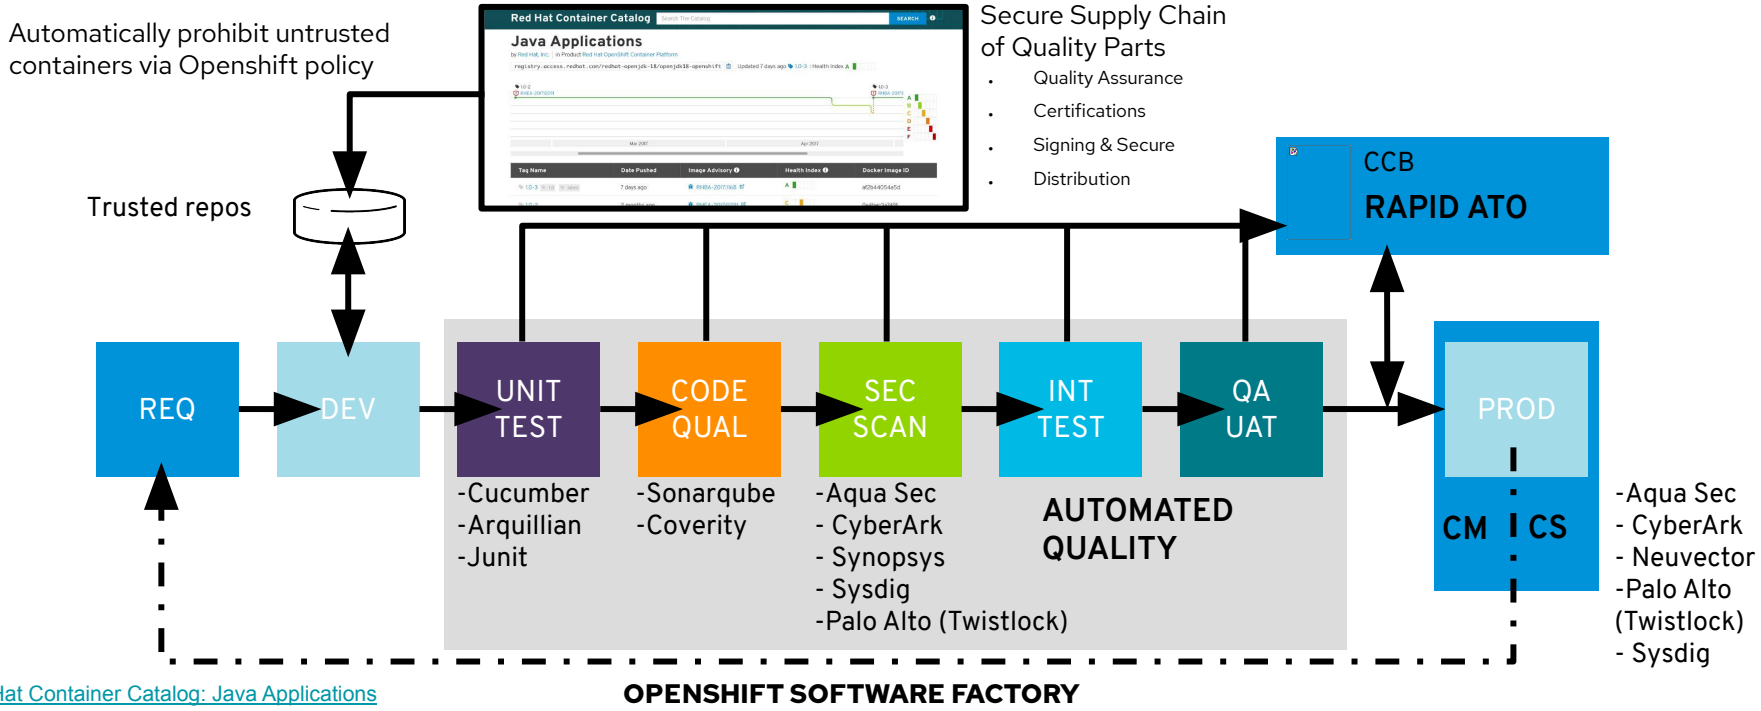

[Red Hat Container Catalog: Java Applications]

[More about the Container Health Index]

Red Hat

Plugins available for VSCode and Eclipse IDEs

Enhancing Secure Application Development and DevSecOps

"Shift Left" - find CVEs and license issues during development

Red Hat Dependency Analytics IDE plugins provide security and license warnings for any project dependency:

- Be notified of CVEs in any package or sub-package
- Remediation advice (upgrade / downgrade)
- Uses open source and Snyk CVE databases
- Supported for Java, Node, Python

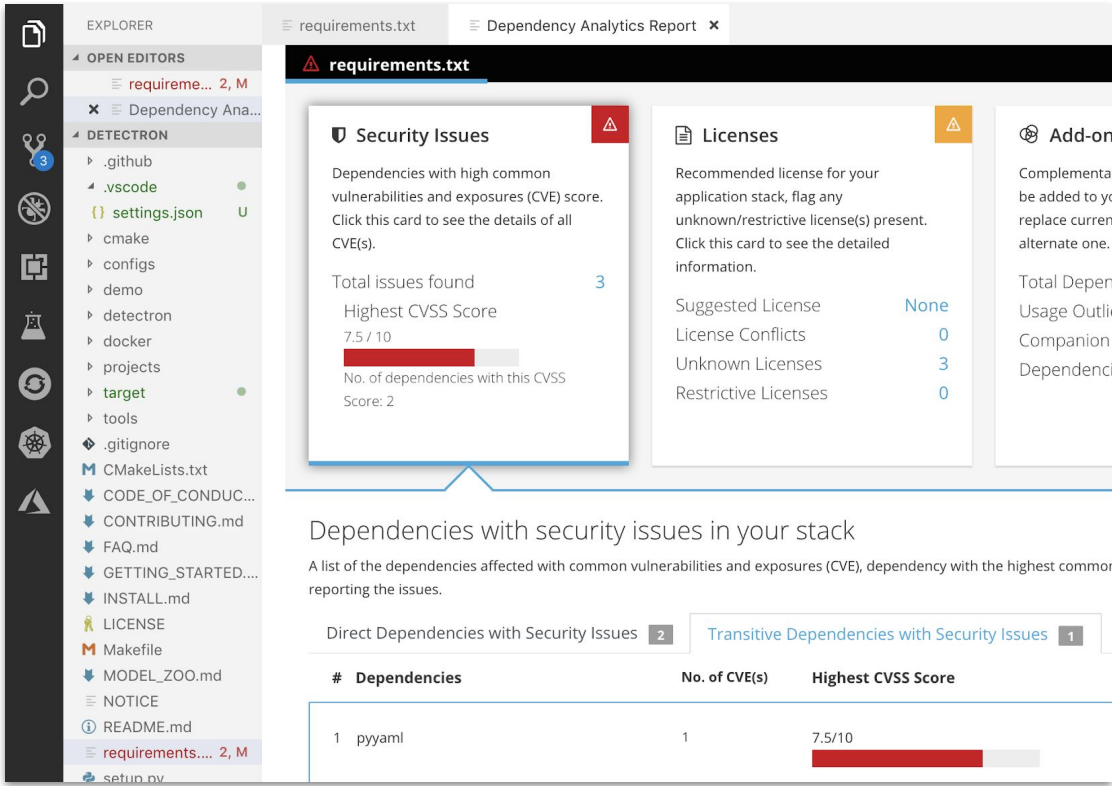

Jenkins CI/CD, run in OpenShift and deploy to OpenShift

Jenkins is still the most used CI/CD platform in enterprises and can be used from inside OpenShift.

An intuitive pipeline visualization makes it simple for users to see how builds are progressing.

The full Jenkins UI is also available.

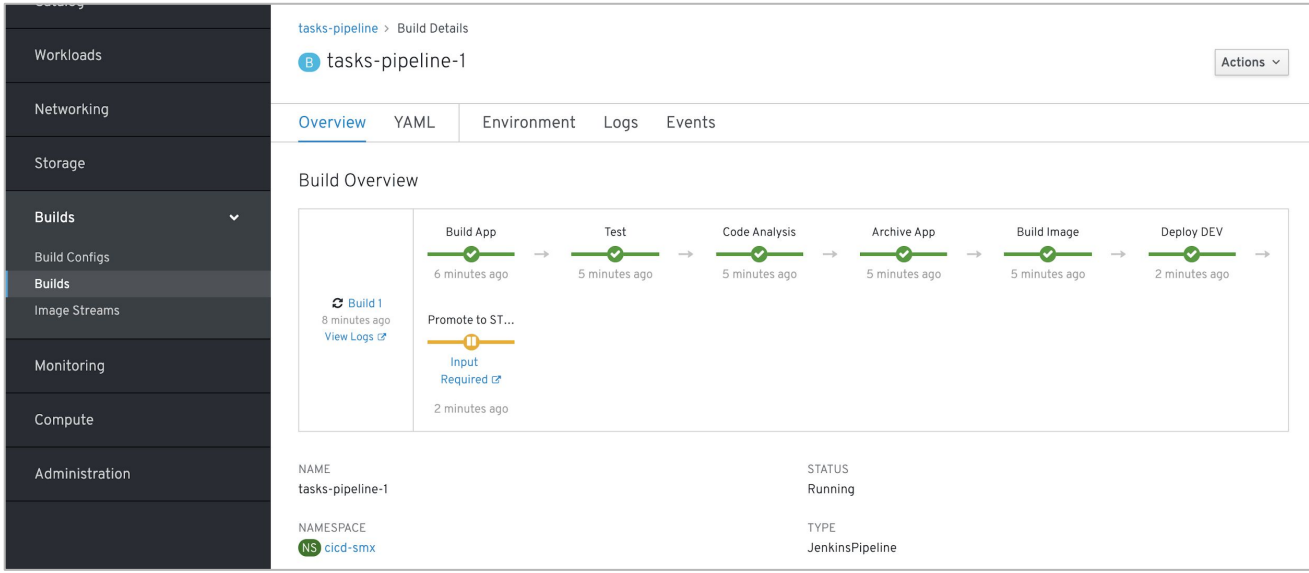

Why? Build in, or for, OpenShift from your enterprise CI/CD system.

Red Hat

OpenShift Pipelines: A Kubernetes-native CI/CD platform

Provides a next-gen Kubernetes CI/CD pipeline that works for containers (including serverless).

Based on the Tekton project (which was spun out of the Knative Pipelines project) started by Google, Red Hat and others.

Target general availability in OpenShift 4.7.

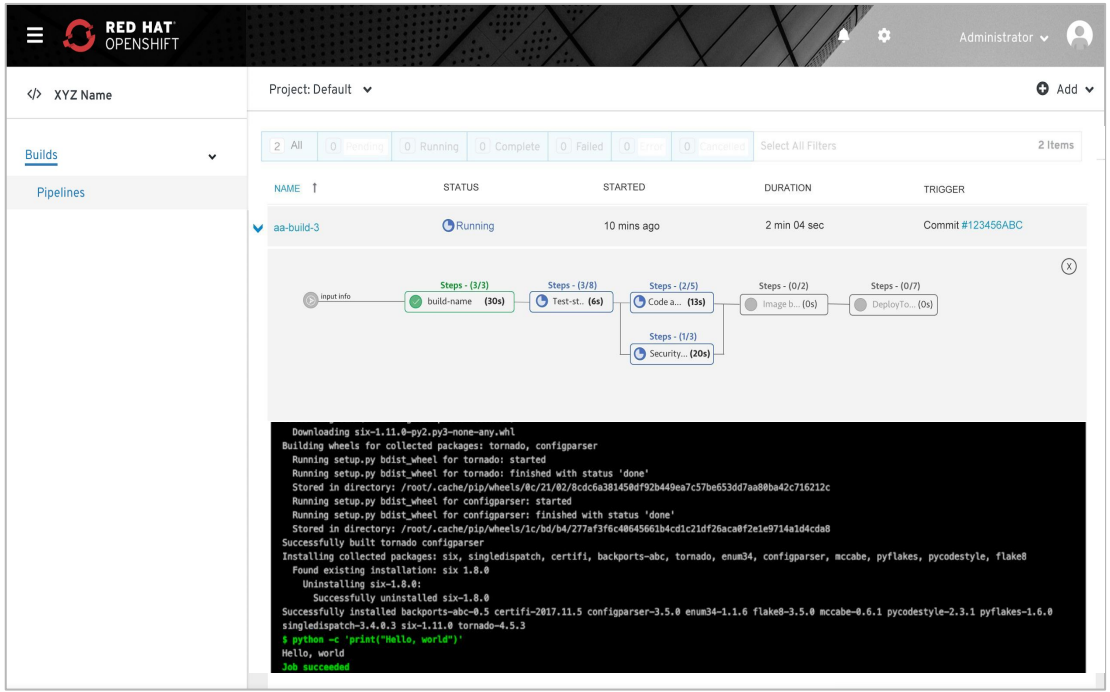

Why? A faster, less resource-intensive CI/CD platform that's Kubernetes-native.

Container Signing

A simplified and automated approach to signing container images

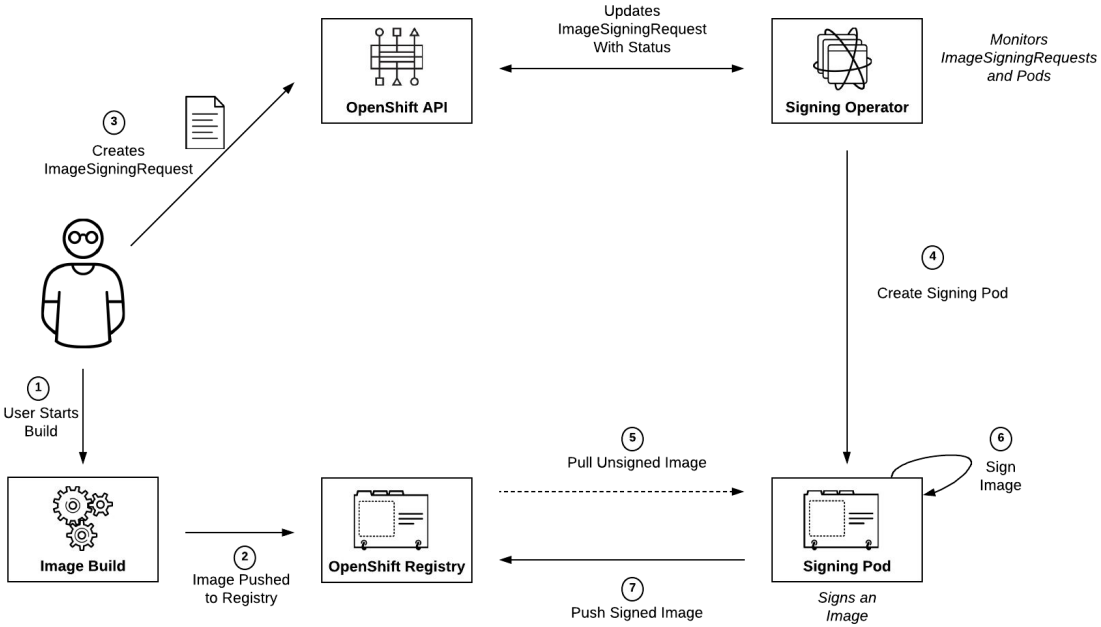

Red Hat

Managing Container Deployment

- Deployments: Containerized App Configuration as Code
- Whitelist / Blacklist external repos
- Apply runtime security policies
- Validate image signatures
- Monitor for new vulnerabilities
- Trust is temporal:
  rebuild & redeploy as needed

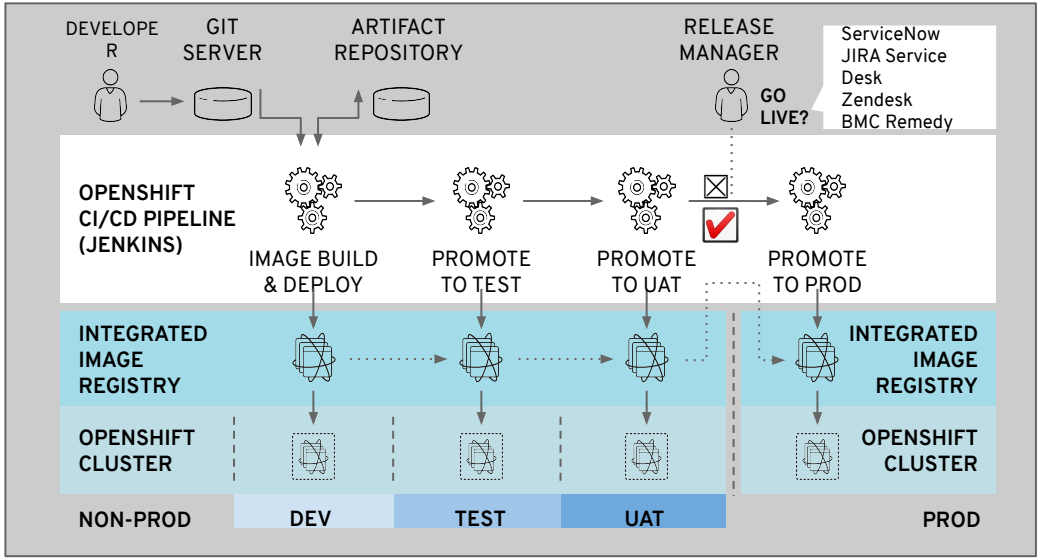

Red Hat

Projects isolate applications across teams, groups and departments

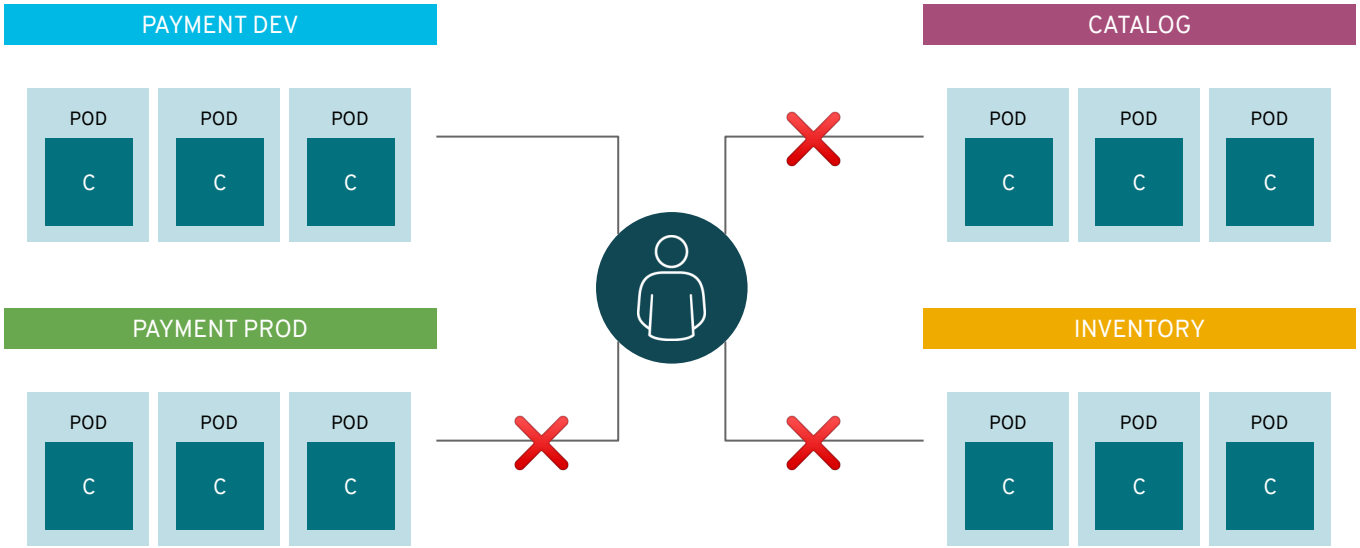

Red Hat

Secure microservices with 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

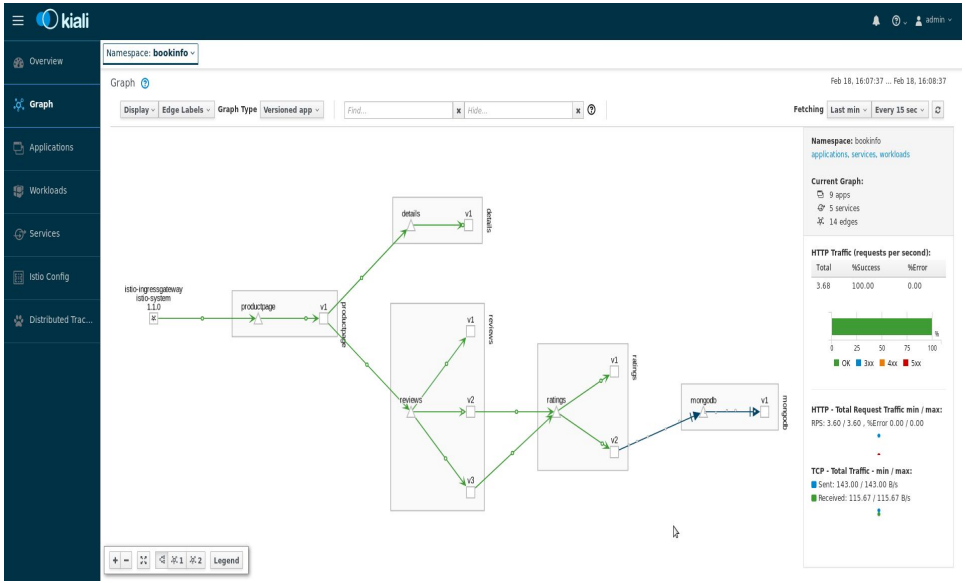

Red Hat

Observability with Kiali

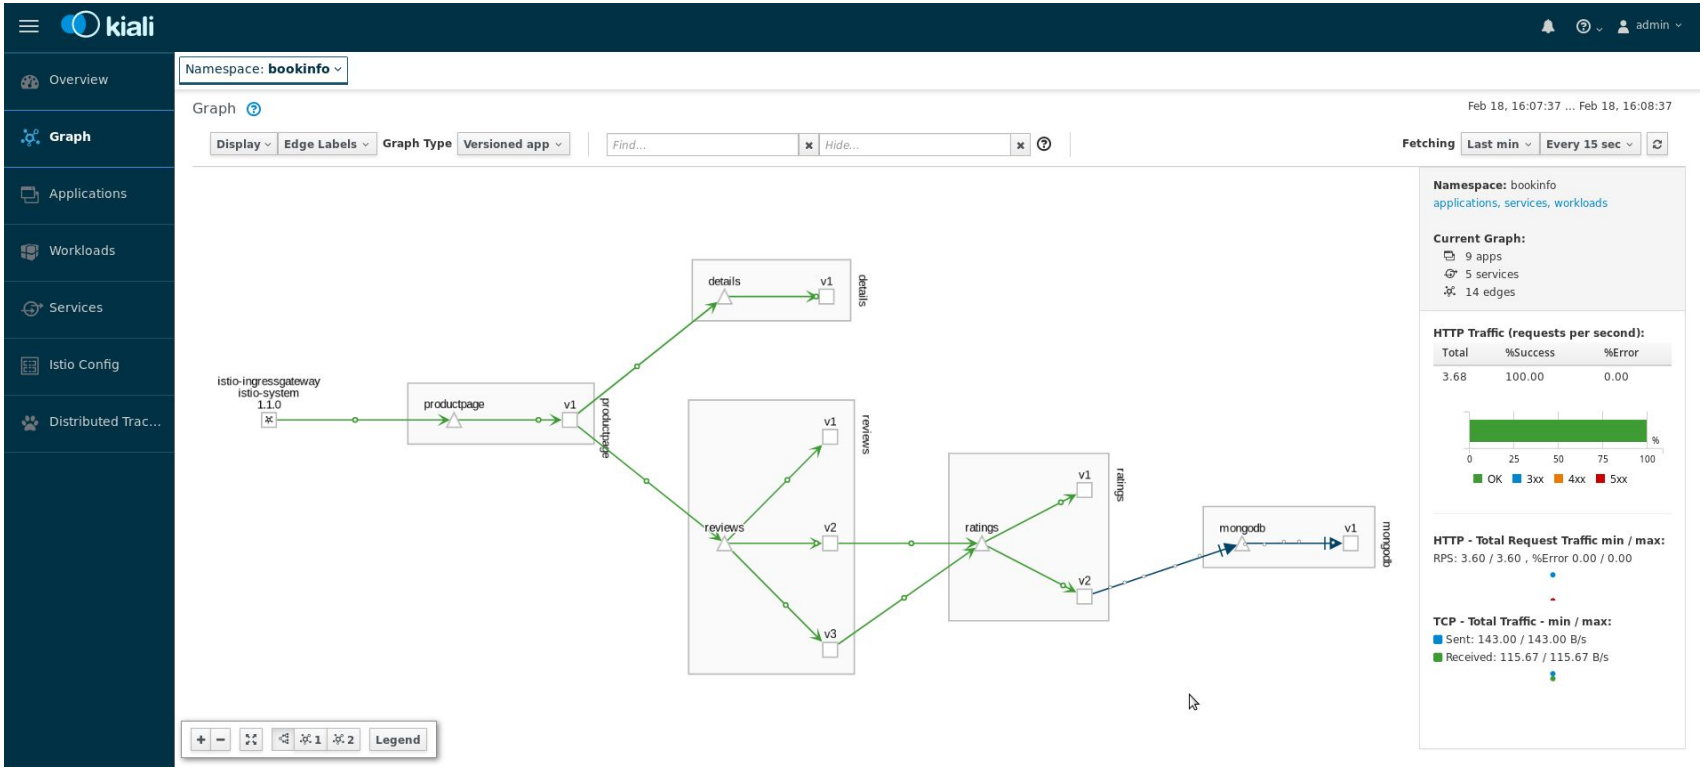

Red Hat

Application API management

Consider configuring an API gateway for container platform & application APIs

- Authentication and authorization
- LDAP integration
- End-point access controls
- Rate limiting

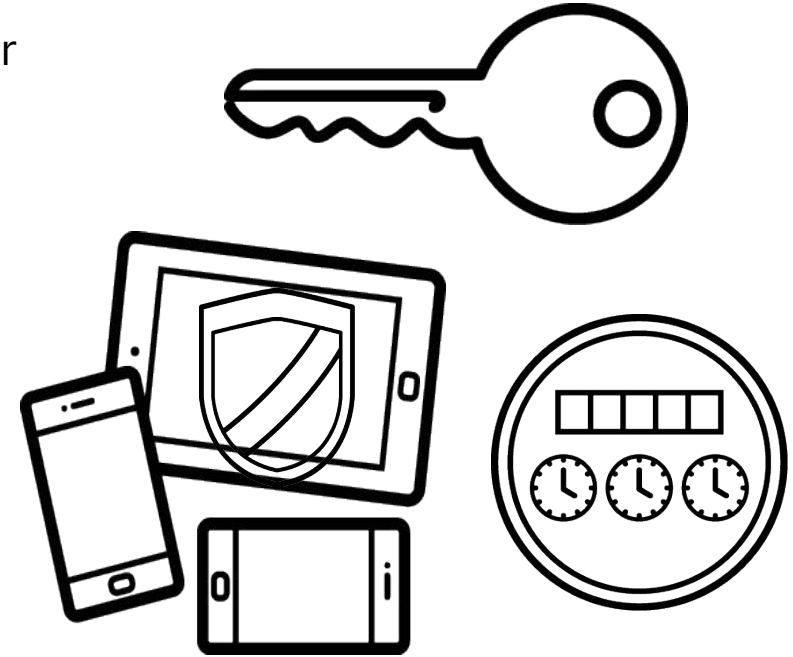

Smarter Software Updates

No downtime for well behaving apps

Applications with multiple replicas, using liveness probes, health checks and taints/tolerations

Node Pools with more than one worker and slack resources

Maintenance window for entire cluster

No need for separate windows for each component

Upgrade runs completely on the cluster

No more long running processes on a workstation

Constant health checking from each Operator

Operators are constantly looking for incompatibilities and issues that might arise

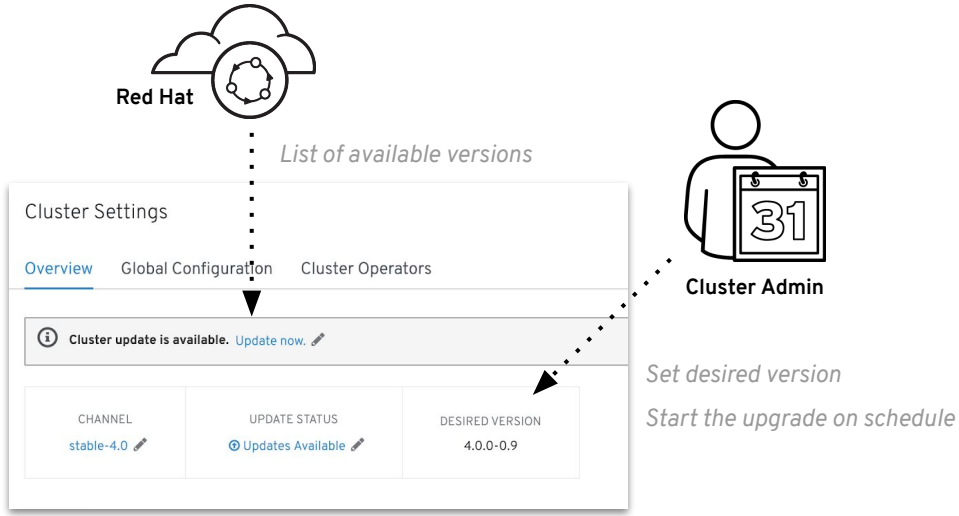

GitOps with ArgoCD Guide

- Guide published to GitHub
  github.com/openshift/openshift-gitops-examples
- Topics
  - Install and configuration of ArgoCD
  - Cluster configs with ArgoCD
  - Operator installation
  - Multi-cluster configs

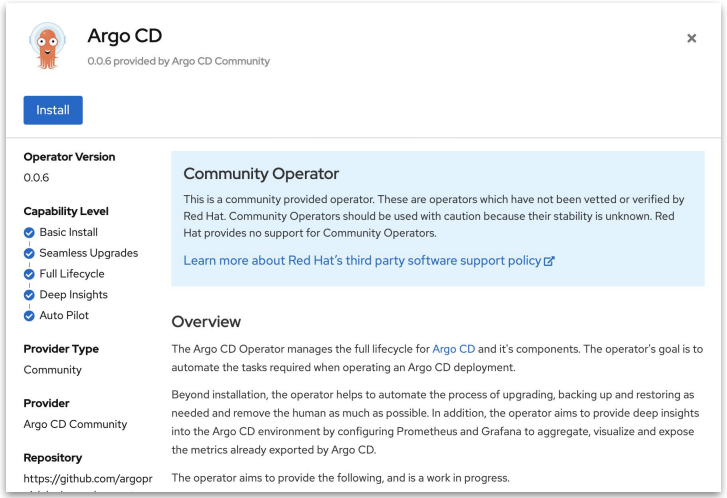

Red Hat

REQUIREMENT 7 & 8

Restrict access by need to know
Identity and access management

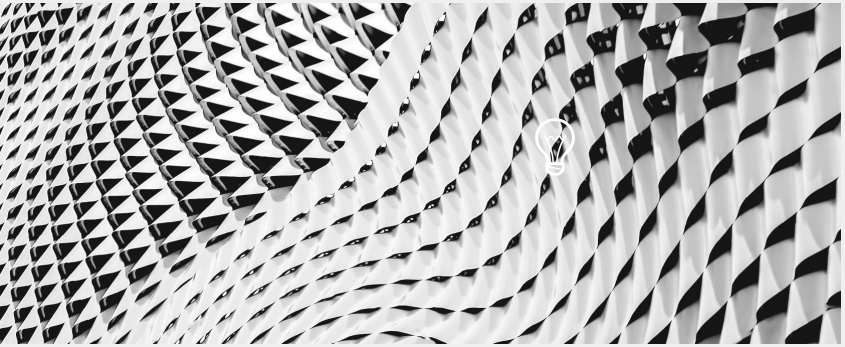

Red Hat

Requirement 7 & 8: Restrict access by need to know, IAM

7. Restrict access to cardholder data by business need to know

    7.1 Limit access to system components and cardholder data to only those who need to know

    7.2 Establish an access control system for system components; default to deny all

    7.3 Ensure that security policies and procedures are documented and in use (Git Ops, Compliance operator)

8. Identify and authenticate access to system components

Requirement 7 & 8 Applicability

- Use OpenShifts built-in RBAC, projects and network policies to isolate users, teams and applications from each other
- Use a 3rd party, external identity provider (integrate via LDAP or OAuth)
- Manage account lockouts, session timeouts, via external IdP (can manage session timeouts in OpenShift)

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)

Restrict access by need to know

Role based authorization

- 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

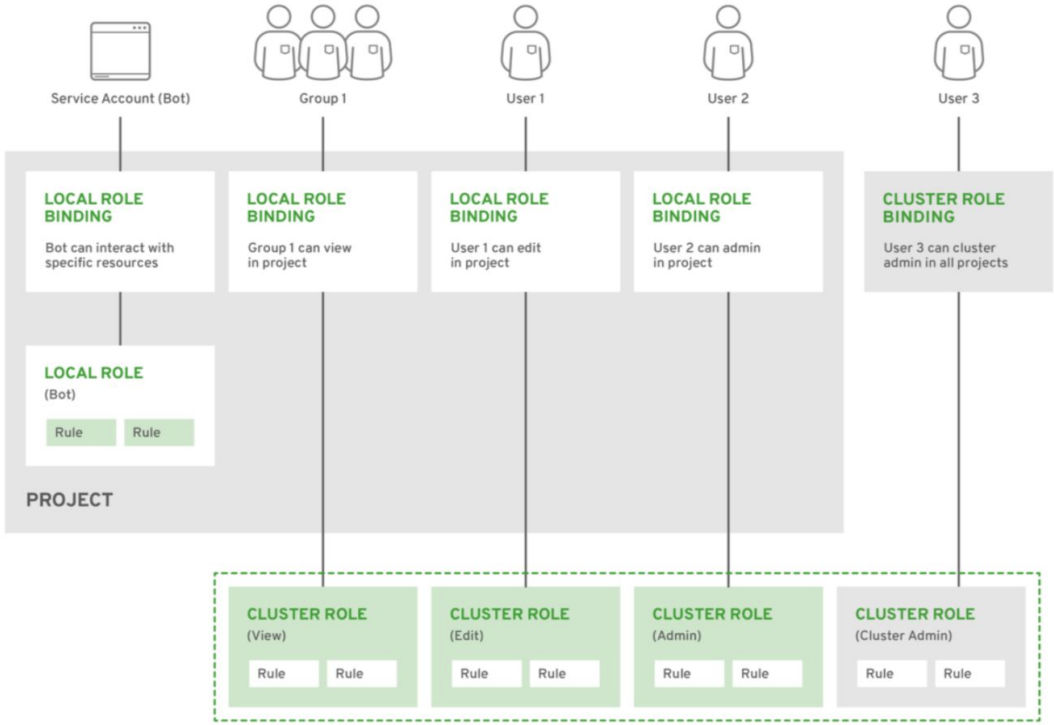

Figure 12 - Authorization Relationships

For more information see: [Managing RBAC in OpenShift]

GitOps with ArgoCD Guide

- Guide published to GitHub
  github.com/openshift/openshift-gitops-examples
- Topics
  - Install and configuration of ArgoCD
  - Cluster configs with ArgoCD
  - Operator installation
  - Multi-cluster configs

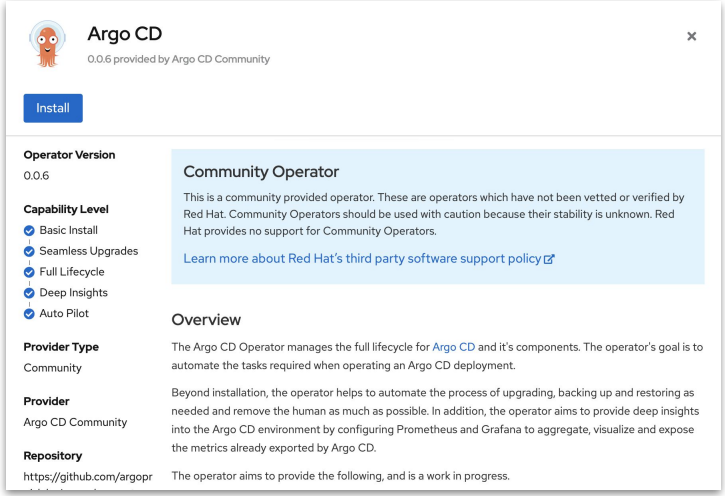

Red Hat

REQUIREMENT 10

Track and monitor all access

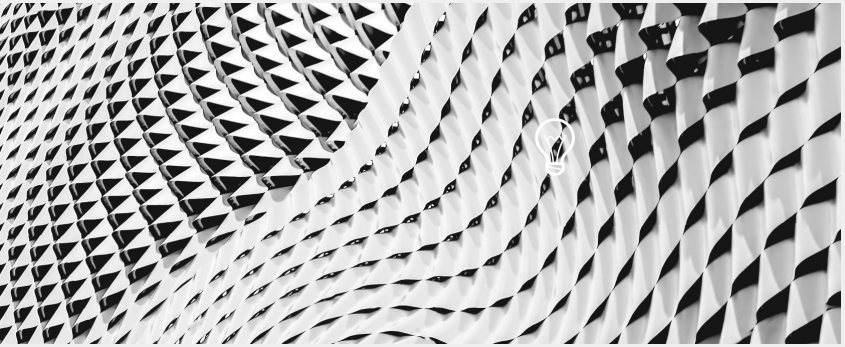

Red Hat

Requirement 10: Track and Monitor all access

10.1 Implement audit trails to link system actions to users

10.2 Implement automated audit trails to monito
- 10.2.1 All individual user access to card data
- 10.2.2 All actions taken by privileged users
- 10.2.3 Access to audit trails
- 10.2.4 Invalid logical access attempts
- 10.2.5 Use of and changes to identity and authentication methods
- 10.2.6 Initializing, stopping, pausing of audit logs
- 10.2.7 Creation and deletion of system level objects

10.3 Record user identity, type of event, date and time of event, success or failure, origin of event, identity or name of target of event

10.4 Use time synchronization technology

10.5 Secure audit trails so they cannot be altered

Requirement 10 Applicability

- Host level and API server audit is on by default
- Evaluate default audit levels
- Forward all logs, including audit logs, to SIEM

Cluster log and audit management

Install the Elasticsearch and Cluster Logging Operators

- 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
- API server events are automatically audited
- Logging pipelines collect API server and host audit logs as well as cluster and application logs for forwarding to the SIEM of your choice

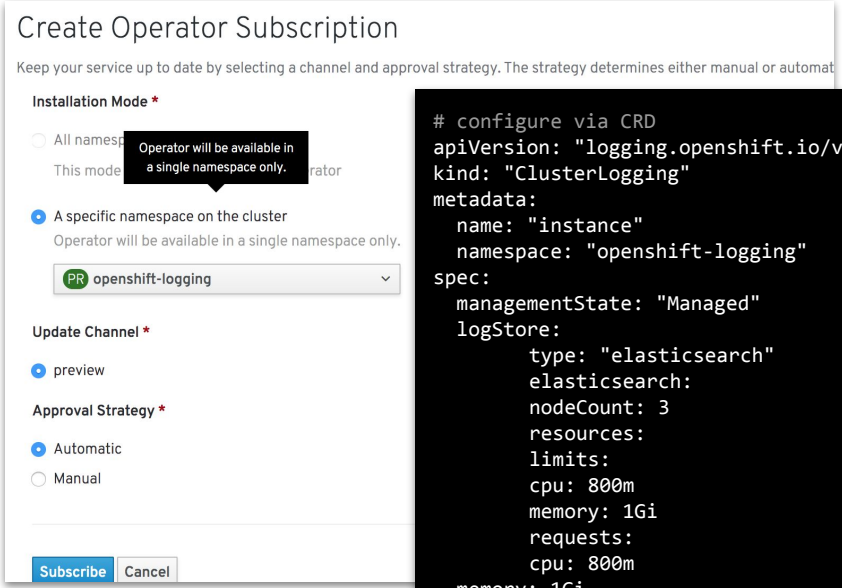


```
# configure via CRD
apiVersion: "logging.openshift.io/v1"
kind: "ClusterLogging"
metadata:
  name: "instance"
  namespace: "openshift-logging"
spec:
  managementState: "Managed"
  logStore:
        type: "elasticsearch"
        elasticsearch:
        nodeCount: 3
        resources:
        limits: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

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

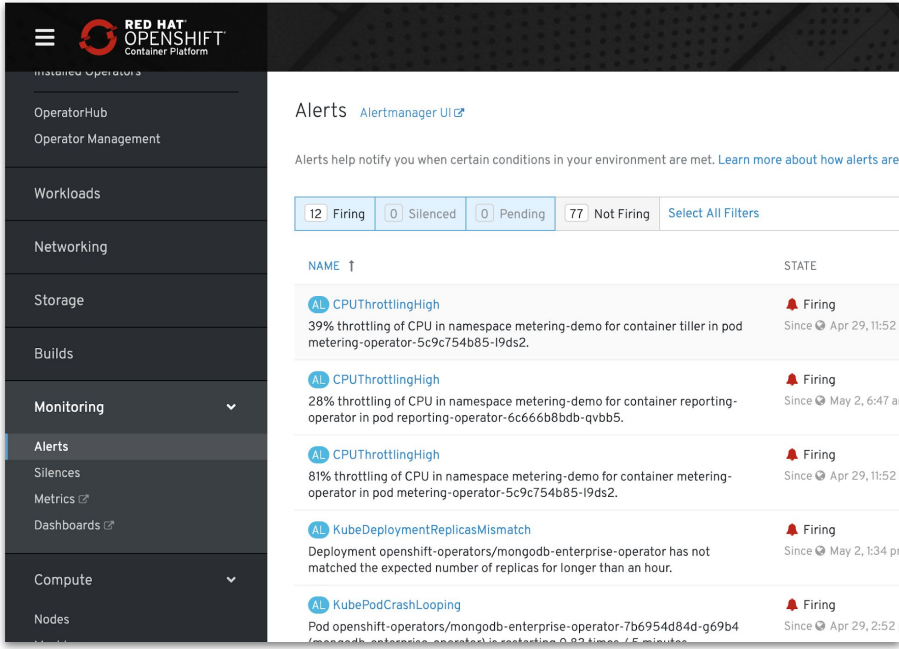

Red Hat

Ingress Access Logging

- There is a new API field on the IngressController resource to configure it:
  - Ability to enable access logs
  - Choice of logging to a pod container or to a Syslog server
  - Options to configure HTTP log format and Syslog facility
  - *Limitation:* Syslog endpoint must be UDP
- Log the hostname of a node from which the log message originated (send-log-hostname) enabled

Log to Sidecar Container

```
$ oc -n openshift-ingress-operator patch ingresscontroller/default --type=merge
--patch='{"spec":{"logging":{"access":{"destination":{"type":"Container"}}}}}'
```

Log to a "facility" on a Syslog server

```
$ oc -n openshift-ingress-operator patch ingresscontroller/default --type=merge
--patch='{"spec":{"logging":{"access":{"destination":{"type":"Syslog","syslog":{"addr
ess":"1.2.3.4","port":10514,"facility":"audit"}}}}}}'
```

View the Logs

```
$ oc -n openshift-ingress logs deploy/router-default -c logs --tail=10 --follow
```

Auditd

- Low level system wide auditing system
- Integrated in Kernel and userspace - no security event escapes!
- Very detailed feed that meets all existing compliance standards
- Actively used by customers that need to adhere to tight security practices
- Auditd is included in RHEL CoreOS
- Host level audit logs are collected for forwarding by the OpenShift Logging Pipelines feature

SUMMARY

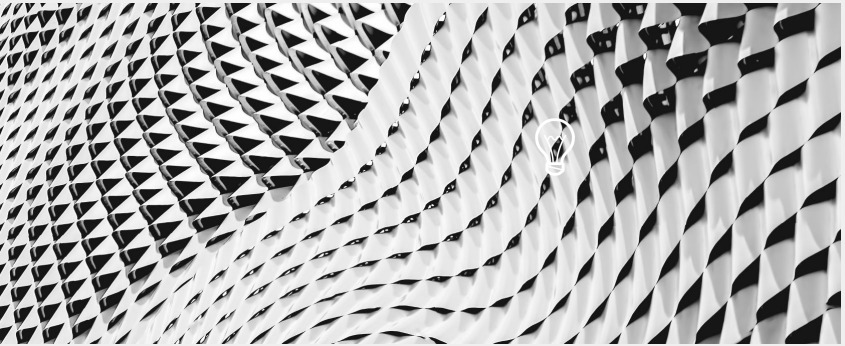

Red Hat

A comprehensive approach to securing containers and Kubernetes

Detect

Trusted Content
- RH supply chain (backport fixes)
- RH Trusted Content with Health Index
- Universal Base Images
- Runtime images

Private Registry
- Integrated registry
- Quay with Clair for image scanning

Build Management
- Source2Image
- ImageStreams track changes to external images

Pipelines & developer tools
- IDE plugins for dependency analysis
- Code Ready Workspaces
- Jenkins / Tekton Pipelines

Protect

Configuration & Lifecycle Management
- OpenShift operators manage drift
- OLM manages operator privileges
- One maintenance window for the full stack
- Upgrades with zero application downtime
- Automate Compliance

Identity and Access Management
- Built-in token based authentication
- Supports 9 Identity Providers including AD/LDAP
- RBAC with Multi-Level Access Control

Platform Data Protection
- Encrypt secrets at rest (etcd datastore)
- All traffic to master nodes is encrypted by default; x.509 certificates for authentication
- Configure cipher suites

Deployment Policies
- SCC (Security Context Constraints)
- No privileged containers by default

Respond

Container isolation
- RHCOS Immutable user space
- SELinux+
- Secure boot
- LUKS volume encryption / FIPS mode
- Non-root containers

Network Isolation
- Ingress / Egress control
- Multus CNI plugin
- Network microsegmentation

Application access and data
- Projects with SELinux annotations control Access to Resources
- Encrypt east / west traffic (Service Mesh)

Observability
- Host and K8s event audit on by default
- Monitoring on by default
- Applications can use cluster monitoring
- Service Mesh traceability
- Container Security Operator

A layered approach to container and Kubernetes security

Red Hat

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

Red Hat

DEFEND INFRASTRUCTURE

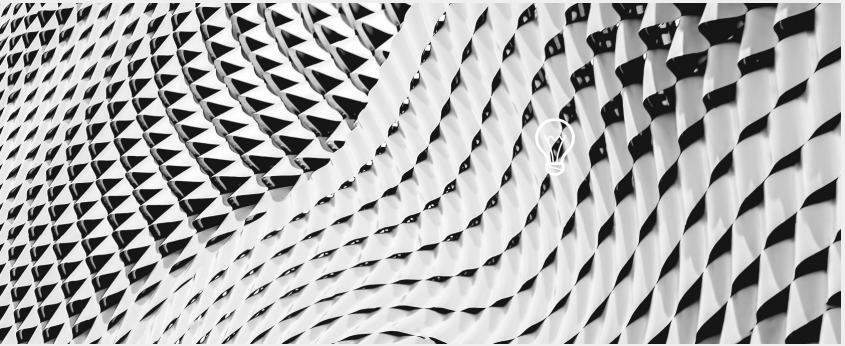

Red Hat

Day 2 Configuration

Global Configuration
You complete most of the cluster configuration and customization after you deploy your OpenShift Container Platform cluster.

Change via Cluster Settings screen

Once you have discovered your desired settings (prev. slide), changes can be made via Console or CLI.

Operators apply these updates

One or more Operators are responsible for propagating these settings through the infrastructure

- Identity Provider
- Ingress Controller
- Logging, Metrics

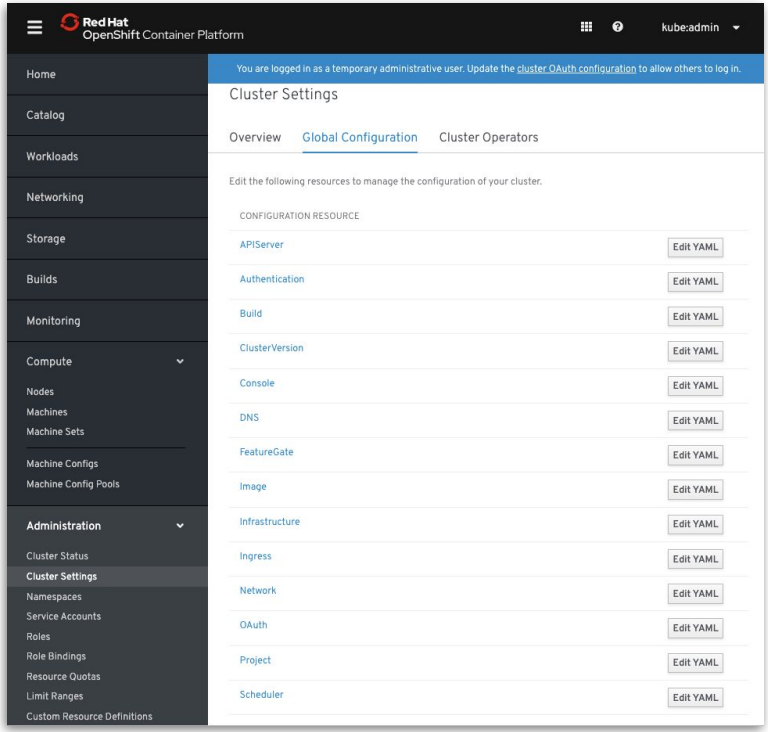

Attached storage

Secure storage by using

- SELinux access controls
- Secure mounts
- Supplemental group IDs for shared storage
- Network bound disk encryption

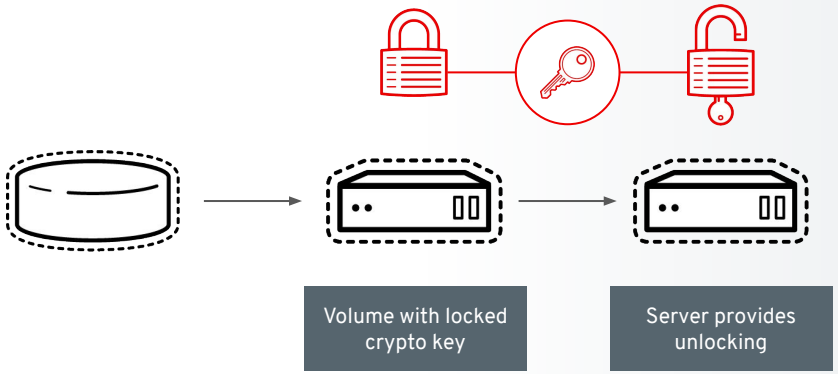

Red Hat

EXTEND SECURITY

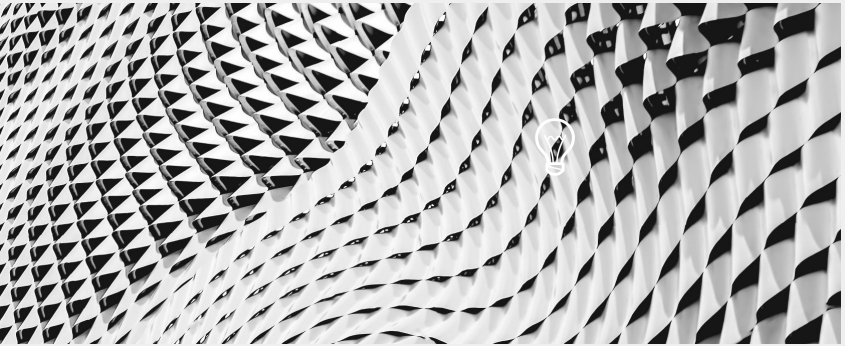

Red Hat

The Security Ecosystem

For enhanced security, or to meet existing policies, you may choose to integrate with enterprise security tools, such as

- Identity and Access management / Privileged Access Management
- External Certificate Authorities
- External Vaults / Key Management solutions
- Filesystem encryption tools
- Container content scanners & vulnerability management tools
- Container runtime analysis tools
- Security Information and Event Monitoring (SIEM)

Red Hat Advanced Cluster Management

Red Hat Advanced Cluster Management for Kubernetes

Robust, proven, award-winning

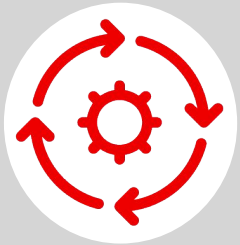

Multicluster life-cycle management

Policy-driven governance, risk, and compliance

Advanced application life-cycle management

Advanced Cluster Management

Application-centric Management
Deploy, upgrade, and manage applications with consistency across multiple clouds

Policy-Based Governance
Enforce policies and ensure compliance across clusters, applications and infrastructures

Cluster Lifecycle Management
Centrally, create, update, delete clusters across the enterprise

Multicluster Management

Infrastructure Management

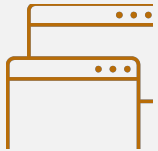

Application Management

Event Management

Existing Tools & Processes

Security & Compliance Management

Multi-Cluster Management and Security with Red Hat Advanced Cluster Management for Kubernetes

- Centrally set & enforce policies for security, applications, & infrastructure
- Quickly visualize detailed auditing on configuration of apps and clusters
- Built-in CIS compliance policies and audit checks
- Immediate visibility into your compliance posture based on your defined standards

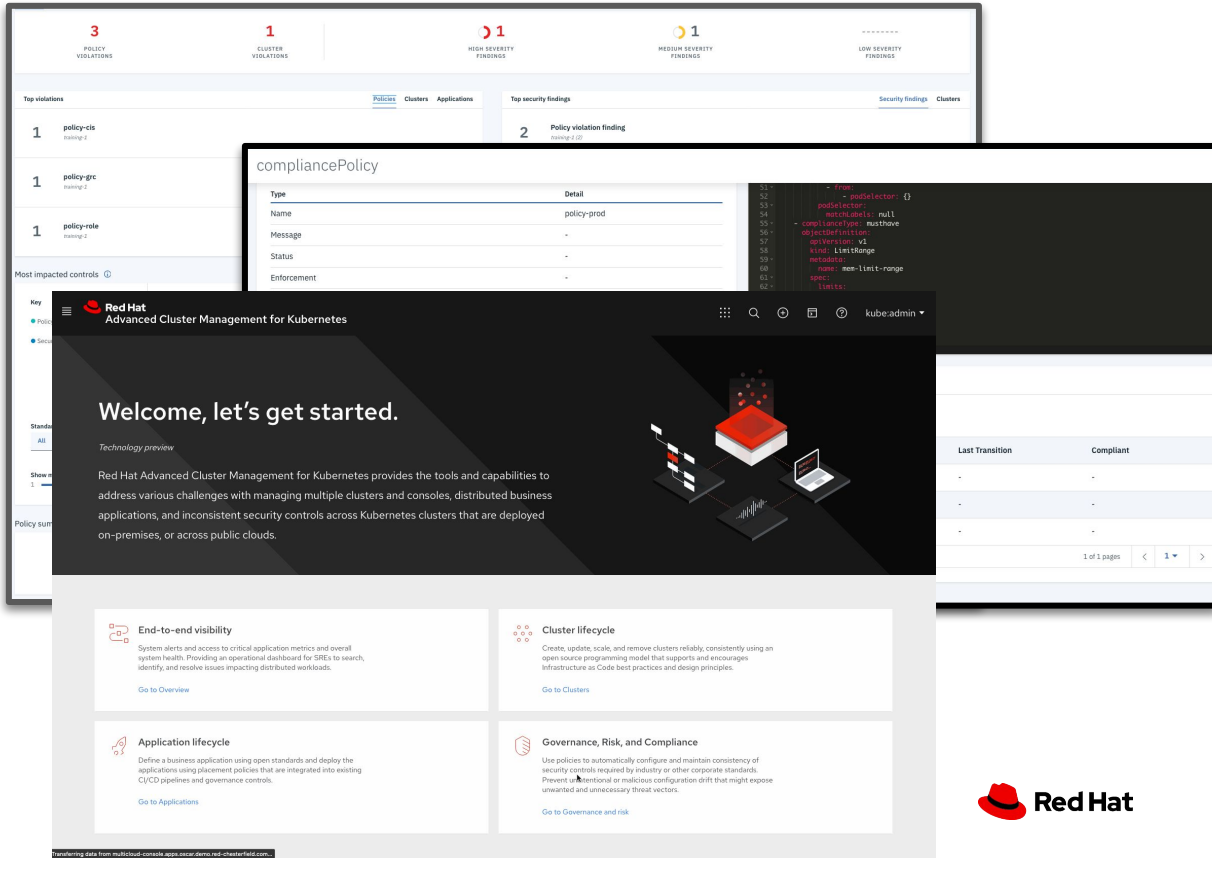

Red Hat

Appendix: External hybrid cloud security guidance

Securing Kubernetes

Guidance from the CNCF Kubernetes Security Audit

" While Kubernetes facilitates high-availability workload deployments, the underlying hosts, components, and environment of a Kubernetes cluster must be configured and managed. This management has a direct impact on the capabilities of the cluster, and affects the behavior of an operator's composed objects.

With this in mind, the options available for configuring components of Kubernetes often fluctuate significantly in supported versions, and vary in their approach to default settings. This leads to a non-trivial amount of configuration required by an administrator to stand-up a functional cluster for a given workload.

More effort must then be spent maintaining the cluster to abide by these settings, especially when planning and executing upgrades of Kubernetes components."

Kubernetes Security Whitepaper, Trail of Bits, May 31, 2019

Securing the container host

Guidance from NIST

Use container-specific host OSs instead of general-purpose ones to reduce attack surfaces.

A container-specific host OS is a minimalist OS explicitly designed to only run containers, with all other services and functionality disabled, and with read-only file systems and other hardening practices employed. When using a container-specific host OS, attack surfaces are typically much smaller than they would be with a general-purpose host OS, so there are fewer opportunities to attack and compromise a container-specific host OS. Accordingly, whenever possible, organizations should use container-specific host OSs to reduce their risk.

NIST Special Publication 800-190
Application Container Security Guide

Securing cloud native workloads

Guidance from Gartner

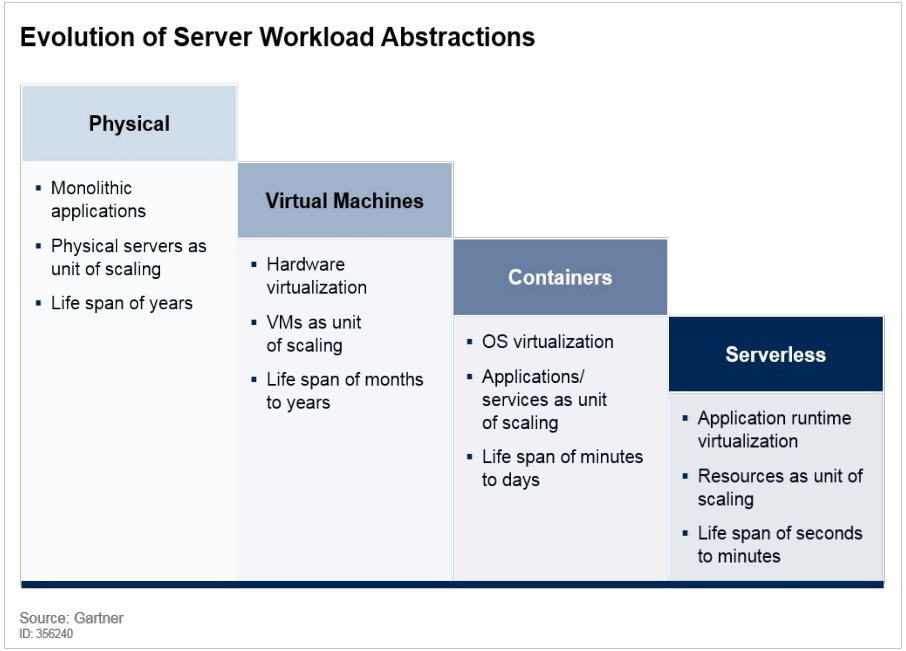

"The best way to secure these rapidly changing and short-lived workloads is to start their protection proactively in the development phase … so that when a workload is instantiated in production, it is "born" protected."

"Replace antivirus (AV)-centric strategies with a "zero-trust execution"/default deny/application control approach to workload protection where possible…." [1]

1. Gartner: Market Guide for Cloud Workload Protection Platforms, ID G00356240, April 8, 2019

Appendix: Resource management

Resource & Cluster Capacity Management

- Manage compute resources, object counts, storage resources
  - Resource quotas per project
  - Resource quotas across multiple projects
- OpenShift Cluster Capacity Tool
  - Simulate a sequence of scheduling decisions to determine how many instances of an input pod can be scheduled on the cluster before it is exhausted of resources

Table 1. Compute resources managed by quota

Resource Name	Description
`cpu`	The sum of CPU requests across all pods in a non-terminal state cannot exceed this value. `cpu` and `requests.cpu` are the same value and can be used interchangeably.
`memory`	The sum of memory requests across all pods in a non-terminal state cannot exceed this value. `memory` and `requests.memory` are the same value and can be used interchangeably.
`ephemeral-storage`	The sum of local ephemeral storage requests across all pods in a non-terminal state cannot exceed this value. `ephemeral-storage` and `requests.ephemeral-storage` are the same value and can be used interchangeably. This resource is available only if you enabled the ephemeral storage technology preview. This feature is disabled by default.
`requests.cpu`	The sum of CPU requests across all pods in a non-terminal state cannot exceed this value. `cpu` and `requests.cpu` are the same value and can be used interchangeably.
`requests.memory`	The sum of memory requests across all pods in a non-terminal state cannot exceed this value. `memory` and `requests.memory` are the same value and can be used interchangeably.
`requests.ephemeral-storage`	The sum of ephemeral storage requests across all pods in a non-terminal state cannot exceed this value. `ephemeral-storage` and `requests.ephemeral-storage` are the same value and can be used interchangeably. This resource is available only if you enabled the ephemeral storage technology preview. This feature is disabled by default.
`limits.cpu`	The sum of CPU limits across all pods in a non-terminal state cannot exceed this value.
`limits.memory`	The sum of memory limits across all pods in a non-terminal state cannot exceed this value.

Tech Preview

Descheduler

Evict a running Pod so that the Pod can be **rescheduled** onto a more suitable node.

Situations when to use:

Nodes are underutilized or overutilized

Pod and node affinity requirements, such as taints or labels, have changed and the original scheduling decisions are no longer appropriate for certain nodes.

Node failure requires Pods to be moved.

New nodes are added to clusters.

https://docs.openshift.com/container-platform/4.4/nodes/scheduling/nodes-descheduler.html#nodes-descheduler

Install the descheduler via OperatorHub to the openshift-kube-descheduler-operator namespace and Create a descheduler instance

```
apiVersion: operator.openshift.io/v1beta1
kind: KubeDescheduler
metadata:
  name: cluster
  namespace: openshift-kube-descheduler-operator
spec:
  deschedulingIntervalSeconds: 3600
  strategies:
    - name: "LowNodeUtilization"
      params:
       - name: "cputhreshold"
         value: "10"
       - name: "memorythreshold"
         value: "20"
       - name: "podsthreshold"
         value: "30"
       - name: "memorytargetthreshold"
         value: "40"
       - name: "cputargetthreshold"
         value: "50"
       - name: "podstargetthreshold"
         value: "60"
       - name: "nodes"
         value: "3"
            - name: "RemoveDuplicates"
```

Nodes Resource views

At-a-glance views for your nodes right from the OpenShift Console.

- Key node data surfaced in the **List view**
- Offers a **new Overview** to provide **insights into critical data** back to you
  - Role/Type/Zone/Address
  - Status/Health Checks
  - Resource Utilizations
    - CPU/Memory/Filesystem
- Directly access to your node with a new **Terminal** view right in the console
  - Act as **root** on the node
  - Access Node Logs (journalctl)

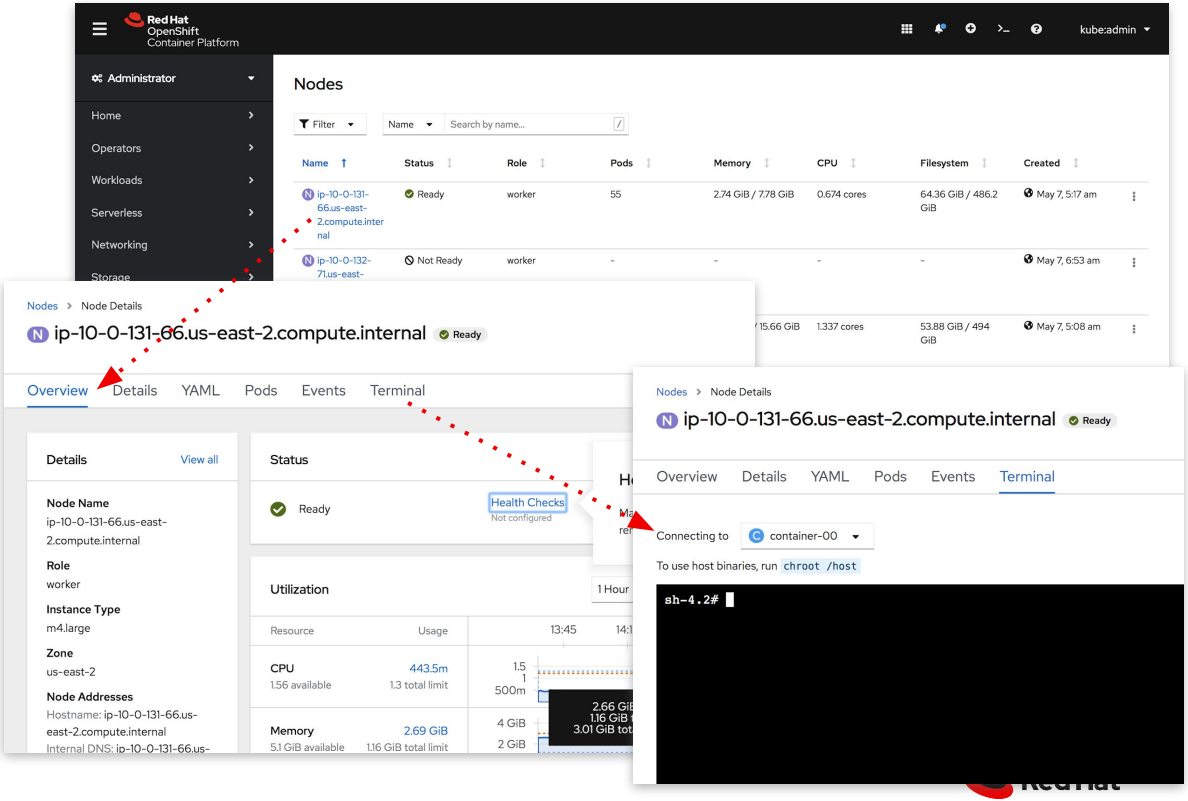

Vertical Pod Autoscaler

Tech Preview

The **VPA** can determine the the **right size for pods** and **frees** the user from having to set **pod resource requests and limits.**

Three controllers:

Recommender - Recommends values for cpu and memory requests based on past consumption

Updater - Kills pods where VPA recommendations do not match the current settings so that they can be recreated by their controllers with the updated requests.

Admission Plugin - Sets the correct resource requests on new pods (due to Updater's activity).

The Vertical Pod Autoscaler Operator is managed by the Cluster Version Operator (CVO) and creates the openshift-vertical-pod-autoscaler namespace

```
apiVersion: autoscaling.openshift.io/v1
kind: VerticalPodAutoscalerController
metadata:
 name: default
spec:
 safetyMarginFraction: 0.15
 podMinCPUMillicores: 25
 podMinMemoryMb: 250
```

Appendix: Etcd encryption & cipher suites

OpenShift 4 etcd Encryption

Encrypt secrets, config maps...

- Encryption of the etcd datastore is optional. Once enabled, encryption cannot be disabled.
- The aes-cbc cipher is used.
- Keys are created and automatically rotated by an operator and stored on the master node's file system.
- Keys are available as a secret via the kube API to a cluster admin.
- Assuming a healthy cluster: after enabling encryption, within a day, all relevant items in etcd are encrypted
- Backup: The etcd data store should be backed up separately from the file system with the key.
- Disaster recovery: a backup of both the encrypted etcd data and encryption keys must be available.

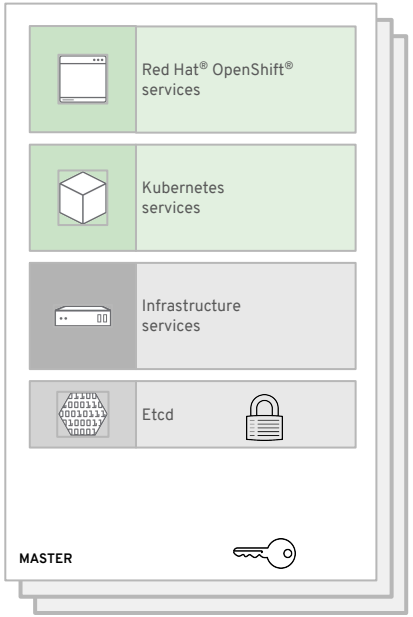

Ingress & API Cipher Suite Configuration

- Allow customers to meet policies requiring them to use specific cipher suites and/or to ensure that disallowed ciphers are not available.
- The TLSSecurityProfile defines the schema for a TLS security profile that will be used by Ingress and the API server.
- Type is one of Old, Intermediate, or Custom. The Modern profile is currently not supported because it is not yet well adopted by common software libraries.

```
// custom is a user-defined TLS security profile. Be extremely careful using a custom
// profile as invalid configurations can be catastrophic. An example custom profile
// looks like this:
//
//   ciphers:
//     - ECDHE-ECDSA-CHACHA20-POLY1305
//     - ECDHE-RSA-CHACHA20-POLY1305
//     - ECDHE-RSA-AES128-GCM-SHA256
//     - ECDHE-ECDSA-AES128-GCM-SHA256
//   minTLSVersion: TLSv1.1
//
// +optional
// +nullable
Custom *CustomTLSProfile `json:"custom,omitempty"`
```