企业级容器应用平台

OpenShift最佳实践

陈沙克 招银云创金融PaaS研究中心总监
张 雷 九州云社区总监

OpenShift

OpenShift Origin is a computer software product from Red Hat for container-based software deployment and management. It is a supported distribution of Kubernetes using Docker containers and DevOps tools for accelerated application development.

- Openshift v2 / v3
  - v2 Cartridges / Gears / Broker
  - V3 Docker / Pods / Master
- openshift v3.9, kubernetes 1.9

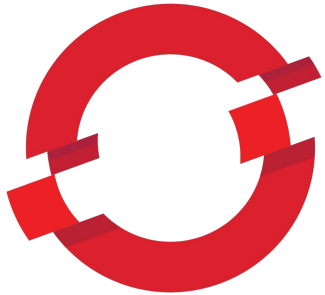

OPENSHIFT

Why OpenShift

FEATURE	KUBERNETES	OPENSHIFT ORIGIN	OPENSHIFT CONTAINER PLATFORM
Multi-host container scheduling	✔	✔	✔
Self-service provisioning	✔	✔	✔
Service-discovery	✔	✔	✔
Persistent storage	✔	✔	✔
Multi-tenancy		✔	✔
Collaboration		✔	✔
Networking		✔	✔
Image registry		✔	✔
Monitoring		✔	✔
Log aggregation		✔	✔
CI/CD and DevOps		✔	✔
Certified application services (databases, runtimes, ...)			✔
Certified middleware services			✔
Built-in operational management			✔
Enterprise-grade operating system			✔
100% Open Source	✔	✔	✔
Community support	✔	✔	✔
Enterprise 24/7 Support			✔
Security response team			✔
Stable Lifecycle (7 years)			✔

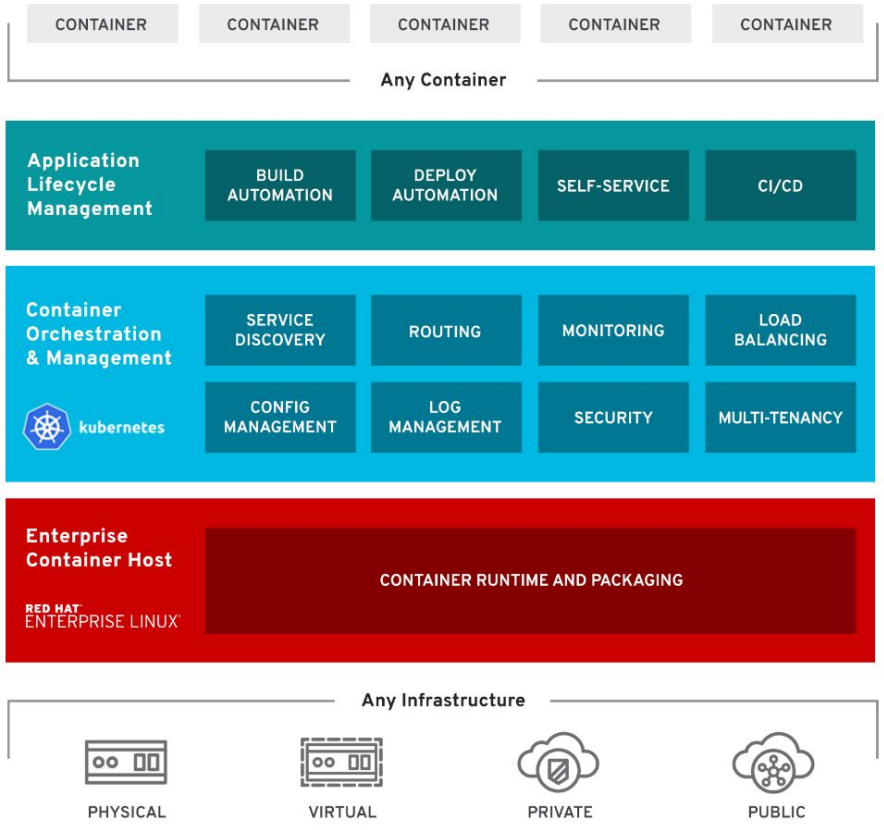

OpenShift Features

- Deployment
  - openshift-ansible / minishift / oc cluster up
- Image
  - Build, BuildConfig / source to image
  - ImageStream
  - Docker Registry
- Network
  - OpenShift SDN
- Service publish
  - Router vs Ingress
- Security
  - SCC vs Pod Security Policy
- EFK
- Monitoring
- Service Catalog
  - ansible-service-broker
  - template-broker
- CI / CD & Devops

OpenShift - Deployment

- minishift
  - minishift start
  - works like minikube
  - launch a vm and run openshift in it
- oc cluster up
  - run openshift in container
- openshift-ansible
  - production-ready
  - EFK / Metrics / Storage / Public & Private Cloud & Baremetal

Image & Registry

- Integrate with docker registry
  - Multi Tenancy
  - ImageStream / ImageStreamTag
- Support Building Image
  - BuildConfig
  - Dockerfile / Source to Image / Jenkins pipeline

Network

- OpenShift SDN
  - A CNI Implementation
  - OpenVSwitch + Vxlan
- Flannel SDN
- Contiv SDN
- Nuage SDN
- Kuryr SDN
- Calico SDN

Network - Isolation

- Network Policy
  - Complicated
- OpenShift SDN support multi tenant mode
  - ovs-multitenant / ovs-networkpolicy

Network - ovs-multitenant

- Each projects gets a unique VNID – identifies pods in that project
- Default projects – VNID=0 communicate with all others (Shared services)
- Pods' traffic inspected according to its project membership

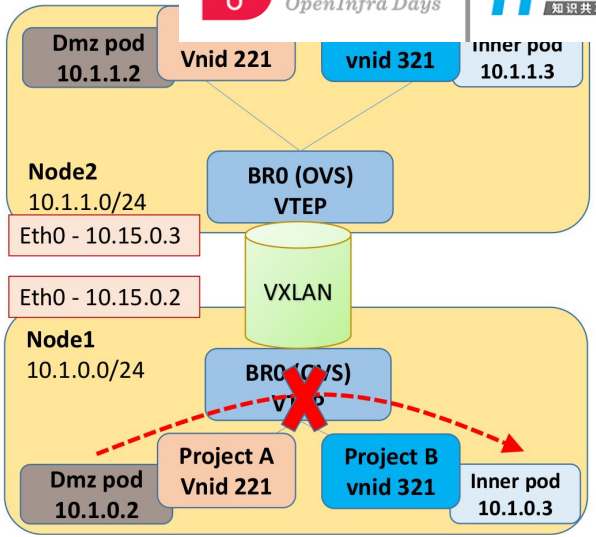


```
oc adm pod-network join-projects --to=<project1> <project2> <project3>
oc adm pod-network isolate-projects <project1> <project2>
```

Network - NetworkPolicy

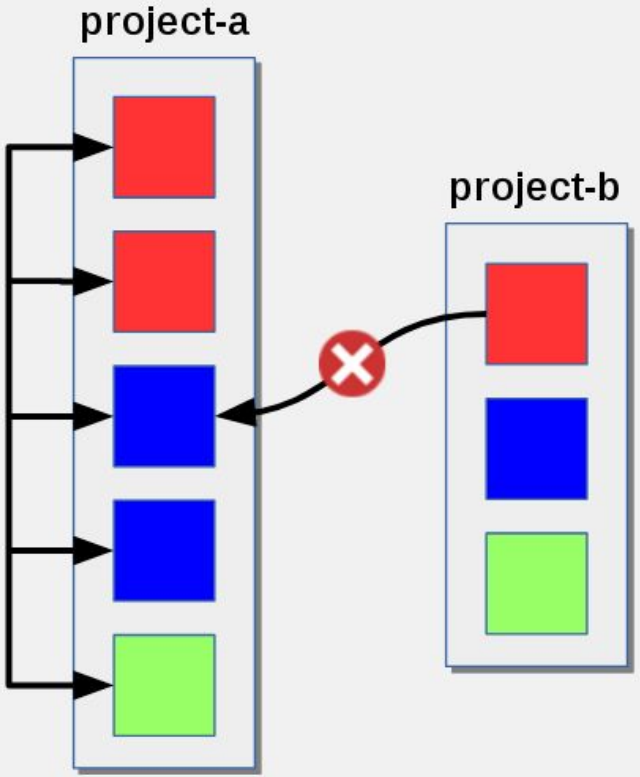

Policy applied to namespace: **project-a**

```
kind: NetworkPolicy
apiVersion: extensions/v1beta1
metadata:
  name: allow-from-same-namespace
spec:
  podSelector:
  ingress:
  - from:
    - podSelector: {}
```

Network - Egress

- Egress Network Policy
- Egress Router
  - Redirect (iptables)
  - HTTP Proxy (squid)
  - DNS Proxy (haproxy)

```
Kind: EgressNetworkPolicy
apiVersion: v1
metadata:
  name: default
spec:
  egress:
    - type: Allow
      to:
        cidrSelector: "192.168.10.0/24"
    - type: Allow
      to:
        dnsName: "www.g.cn"
    - type: Deny
      to:
        cidrSelector: "0.0.0.0/0"
```

Network - Egress

- Egress Network Policy
- Egress Router
  - Redirect (iptables)
  - HTTP Proxy (squid)
  - DNS Proxy (haproxy)

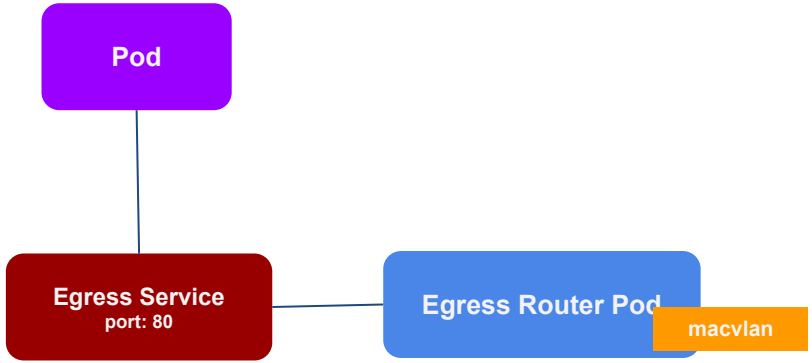

Network - QoS

```
kind: Pod
apiVersion: v1
metadata:
  name: iperf-slow
  annotations:
    kubernetes.io/ingress-bandwidth: 10M
    kubernetes.io/egress-bandwidth: 10M
spec:
  containers:
  - image: nginx
    name: nginx
```

Router vs Ingress

- Router support multi tenancy
- Router support shard
- Better TLS support
  - edge / reencrypt / passthrough

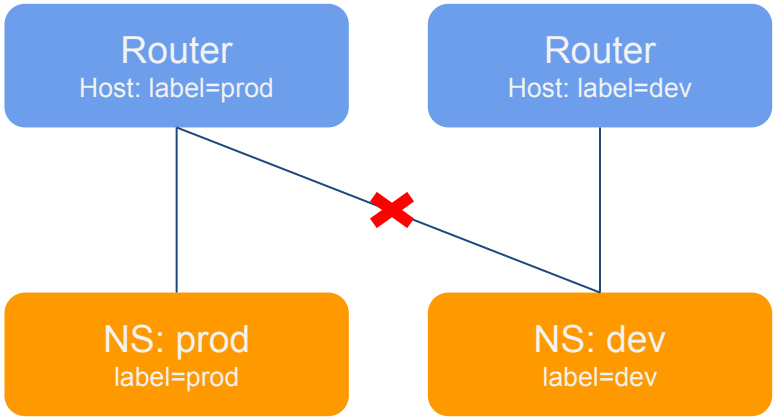

Security

- TLS
  - API / controllers / etcd / nodes / registry / router
- Network Isolation
  - ovs-multitenant / ovs-networkpolicy
- Service Serving Certificate
  - `service.alpha.openshift.io/serving-cert-secret-name: app1-tls`
- Security Context Constraints (SCC) vs Pod Security Context

Security Context Constraints (SCC) vs Pod Securit


```
NAME         PRIV   CAPS   SELINUX    RUNASUSER          FSGROUP    SUPGROUP   PRIORITY   READONLYROOTFS   VOLUMES
anyuid       false  []     MustRunAs  RunAsAny           RunAsAny   RunAsAny   10         false            [configMap ... ]
hostaccess   false  []     MustRunAs  MustRunAsRange     MustRunAs  RunAsAny   <none>     false            [configMap ... ]
hostnetwork  false  []     MustRunAs  MustRunAsRange     MustRunAs  MustRunAs  <none>     false            [configMap ... ]
nonroot      false  []     MustRunAs  MustRunAsNonRoot   RunAsAny   RunAsAny   <none>     false            [configMap ... ]
privileged   true   [*]    RunAsAny   RunAsAny           RunAsAny   RunAsAny   <none>     false            [*]
restricted   false  []     MustRunAs  MustRunAsRange     MustRunAs  RunAsAny   <none>     false            [configMap ... ]
```

```
$ oc adm policy add-scc-to-user anyuid -z useroot
```

EFK

- ElasticSearch & Curator & Fluentd & Kibana
- Collect logs from nodes and projects
- Authentication

Monitoring

- heapster & Hawkular & Cassandra
- Integrate with web console
- pod horizon auto scaling

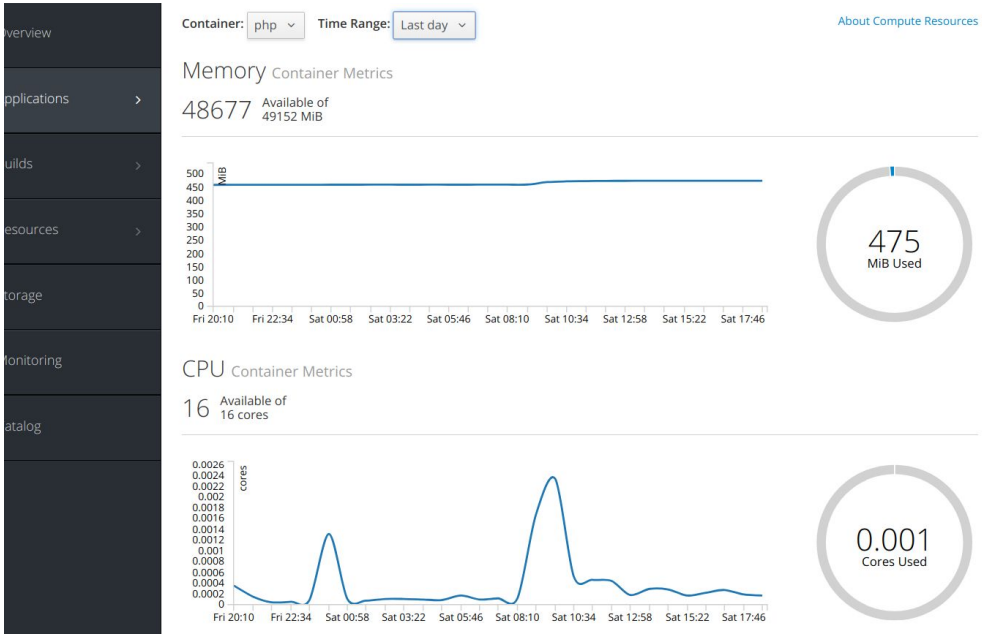

Monitoring

- heapster & Hawkular & Cassandra
- Integrate with web console
- pod horizon auto scaling

```
apiVersion: extensions/v1beta1
kind: HorizontalPodAutoscaler
metadata:
  name: frontend
spec:
  scaleRef:
    kind: DeploymentConfig
    name: frontend
    apiVersion: v1
    subresource: scale
  minReplicas: 1
  maxReplicas: 10
  cpuUtilization:
    targetPercentage: 80
```

Service Catalog

- Template-Broker
- Ansible-Service-Broker

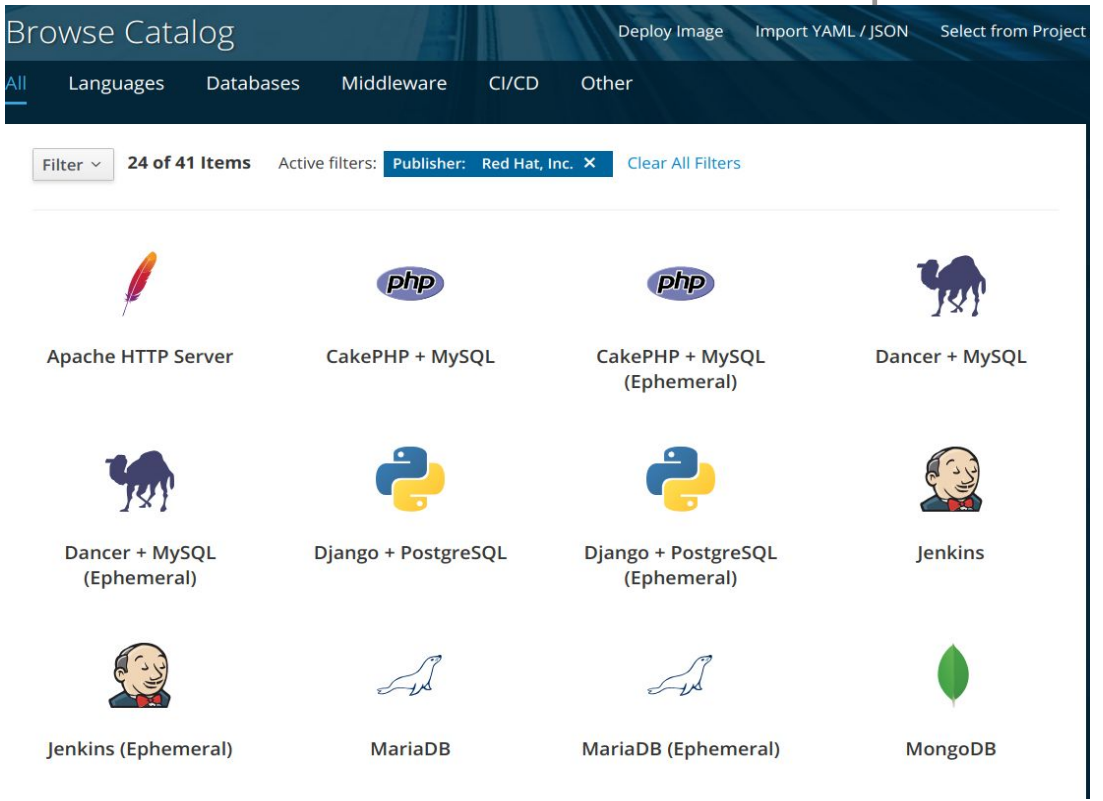

CICD & Devops

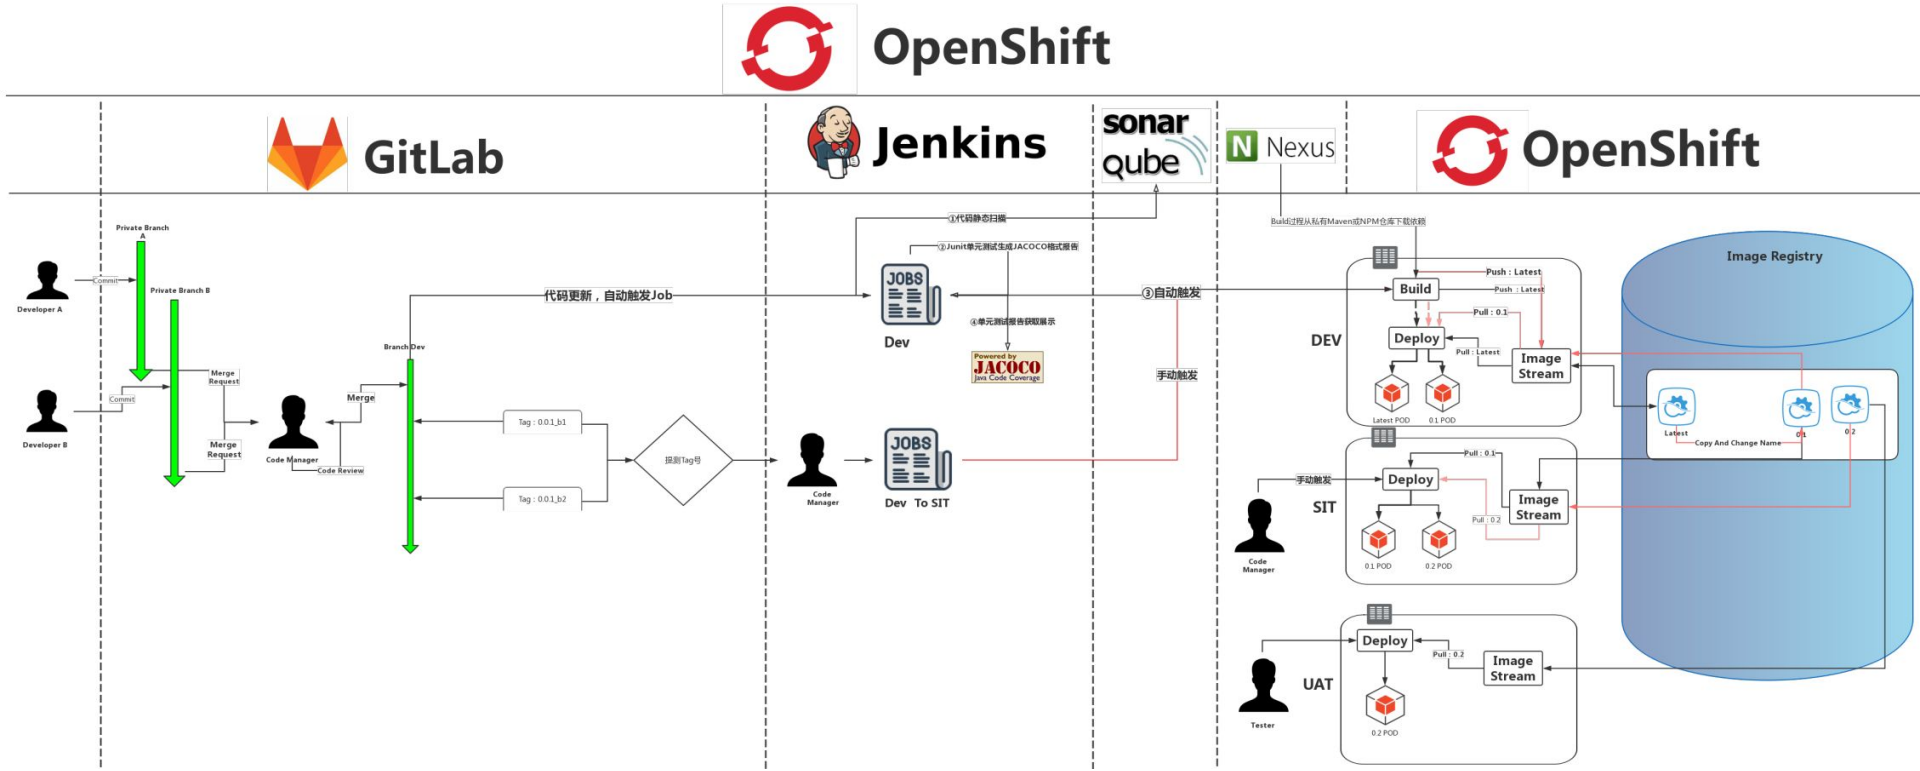
OpenShift
GitLab
Jenkins
sonar qube
Nexus
OpenShift
Private Branch A
Private Branch B
Developer A
Developer B
Commit
Merge Request
Code Manager
Code Review
Merge
Branch:Dev
代码更新，自动触发Job
Tag : 0.0.1_b1
Tag : 0.0.1_b2
提测Tag号
①代码静态扫描
②Junit单元测试生成JACOCO格式报告
④单元测试报告获取展示
JOBS
Dev
Powered by JACOCO Java Code Coverage
③自动触发
手动触发
Code Manager
JOBS
Dev To SIT
Build过程从私有Maven或NPM仓库下载依赖
DEV
Push : Latest
Build
Deploy
Pull : 0.1
Pull : Latest
Image Stream
Latest POD
0.1 POD
Image Registry
Latest
Copy And Change Name
0.1
0.2
手动触发
SIT
Pull : 0.1
Deploy
Image Stream
Pull : 0.2
0.1 POD
0.2 POD
Code Manager
UAT
Deploy
Pull : 0.2
Image Stream
0.2 POD
Tester

OpenShift Roadmap

OpenShift Container Platform 3.10 (July)

- Kubernetes 1.10 and CRI-O option
- Smart Pruning
- Istio (Dev Preview)
- oc client for developers
- Golden Image Tooling and TLS bootstrapping
- Windows Server Containers (Dev Preview))
- Prometheus Metrics and Alerts (Tech Preview)
- S3 Svc Broker

OpenShift Online & Dedicated

- Dedicated self-service: RBAC, templates, LB, egress
- Dedicated encrypted storage, multi-AZ, Azure beta

Q2 CY2018

Q3 CY2018

OpenShift Container Platform 3.11 (Sept)

- Kubernetes 1.11 and CRI-O default
- Infra monitoring ,alerting with SRE intelligence, Node Problem Detector
- Etcd, Prometheus, and Vault Operators - Tech preview
- Operator Certification Program and JBoss Fuse Operator
- Autoscaler for AWS and P-SAP features
- HPA Custom Metric
- Tech preview of ALM
- New web console for developers and cluster admins
- Ansible Galaxy ASB support
- CNV (Tech Preview)
- OVN (Tech Preview for Windows)
- FIPS and other Security PAGs

OpenShift Online & Dedicated

- OpenShift Online automated updates for OS
- Chargeback for OpenShift Online Starter

Q4 CY2018

OpenShift Container Platform 3.12 (Dec/Jar

- Kubernetes 1.12 and CRI-O default
- Converged Platform
- Full Stack Automated Installer
  - AWS, RHEL, Azure, OSP
- Over the Air Updates
- RHCC integrated experience
- Windows Containers GA
- Easy/Trackable Evaluations
- Red Hat CoreOS Container Linux with Ignition Automations
- Cluster Registry
- HPA metrics from Prometheus
- Metering and Chargeback (Tech Preview)

OpenShift Online & Dedicated

- Cluster Operator driven installs
- Self-Service Dedicated User Experience

Q1 CY2019

OpenShift Container Platform 3.13 (March)

- Kubernetes 1.13 and CRI-O default
- Full Stack Automation
  - GCP, VMware
- Istio GA
- Mobile 5.x
- Serverless (Tech Preview)
- RHCC for non-container content
- Integrated Quay (Tech Preview)
- Idling Controller
- Federated Ingress and Workload Policy
- OVN GA
- Che (Tech Preview)

OpenShift Online & Dedicated

- OpenShift.io on Dedicated (Tech Preview)

CHINA
OpenInfra Days

Thank You