

From Zero to Hero with OpenShift

Red Hat Forum 2019-10-03, Stockholm

Jimmy Falkbjer, Principal Architect Software

Introduction

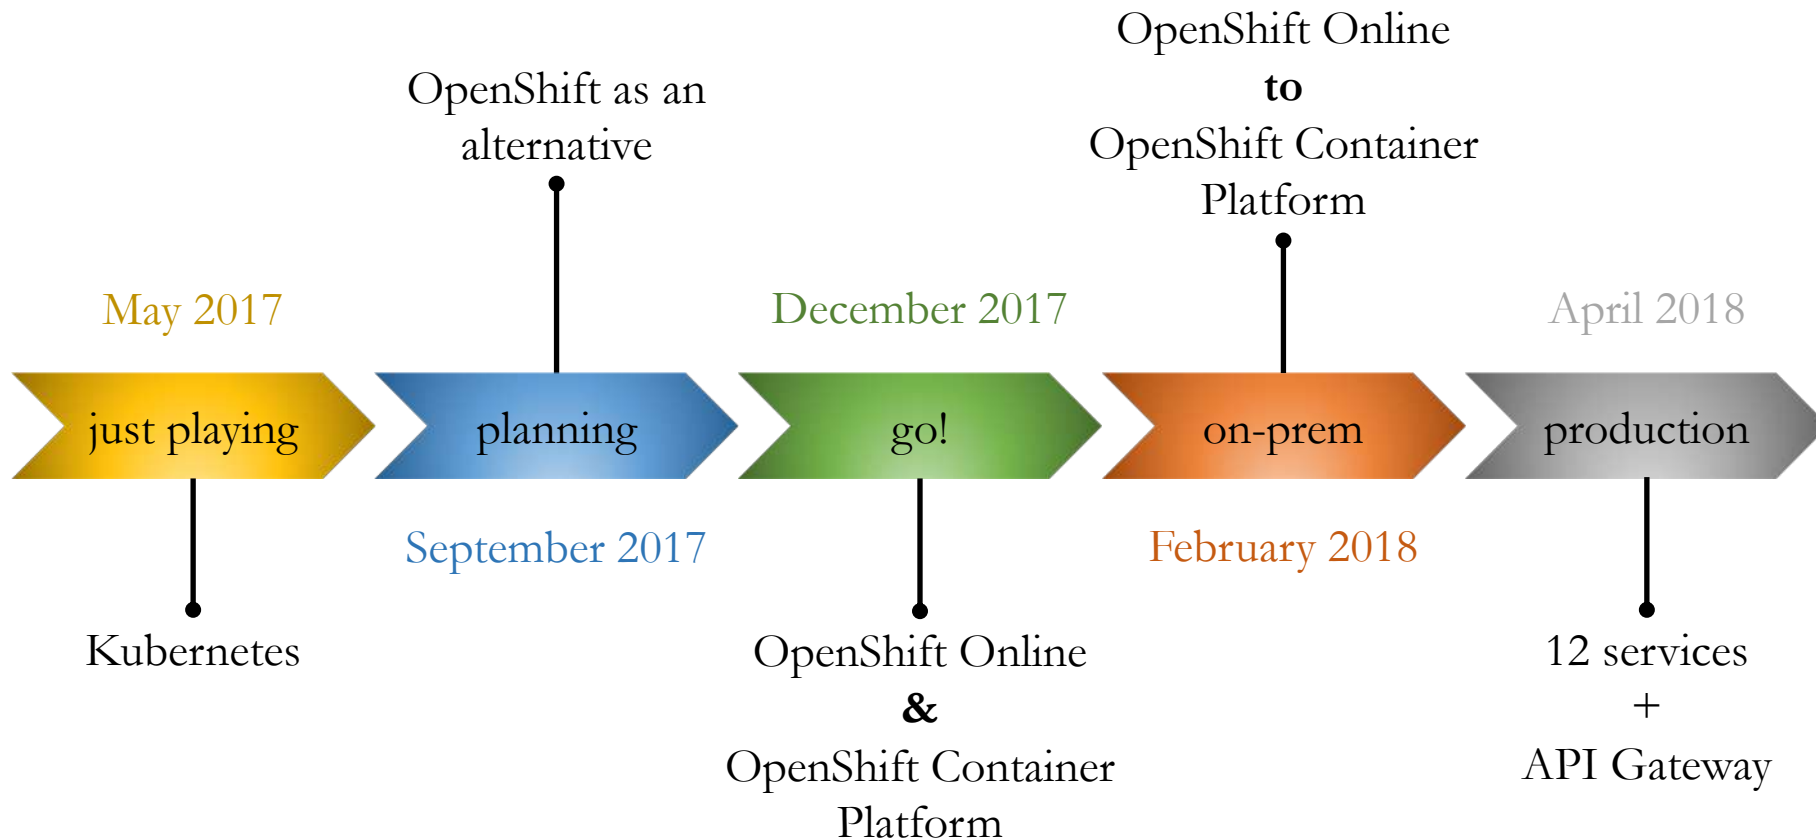
- Jonas Björk, Operations
- Kristian Ejvind, Operations
- Mathias Åhsberg, System Developer

- Pernilla Lundqvist, Project Manager
- Niclas Tegnér, System Developer

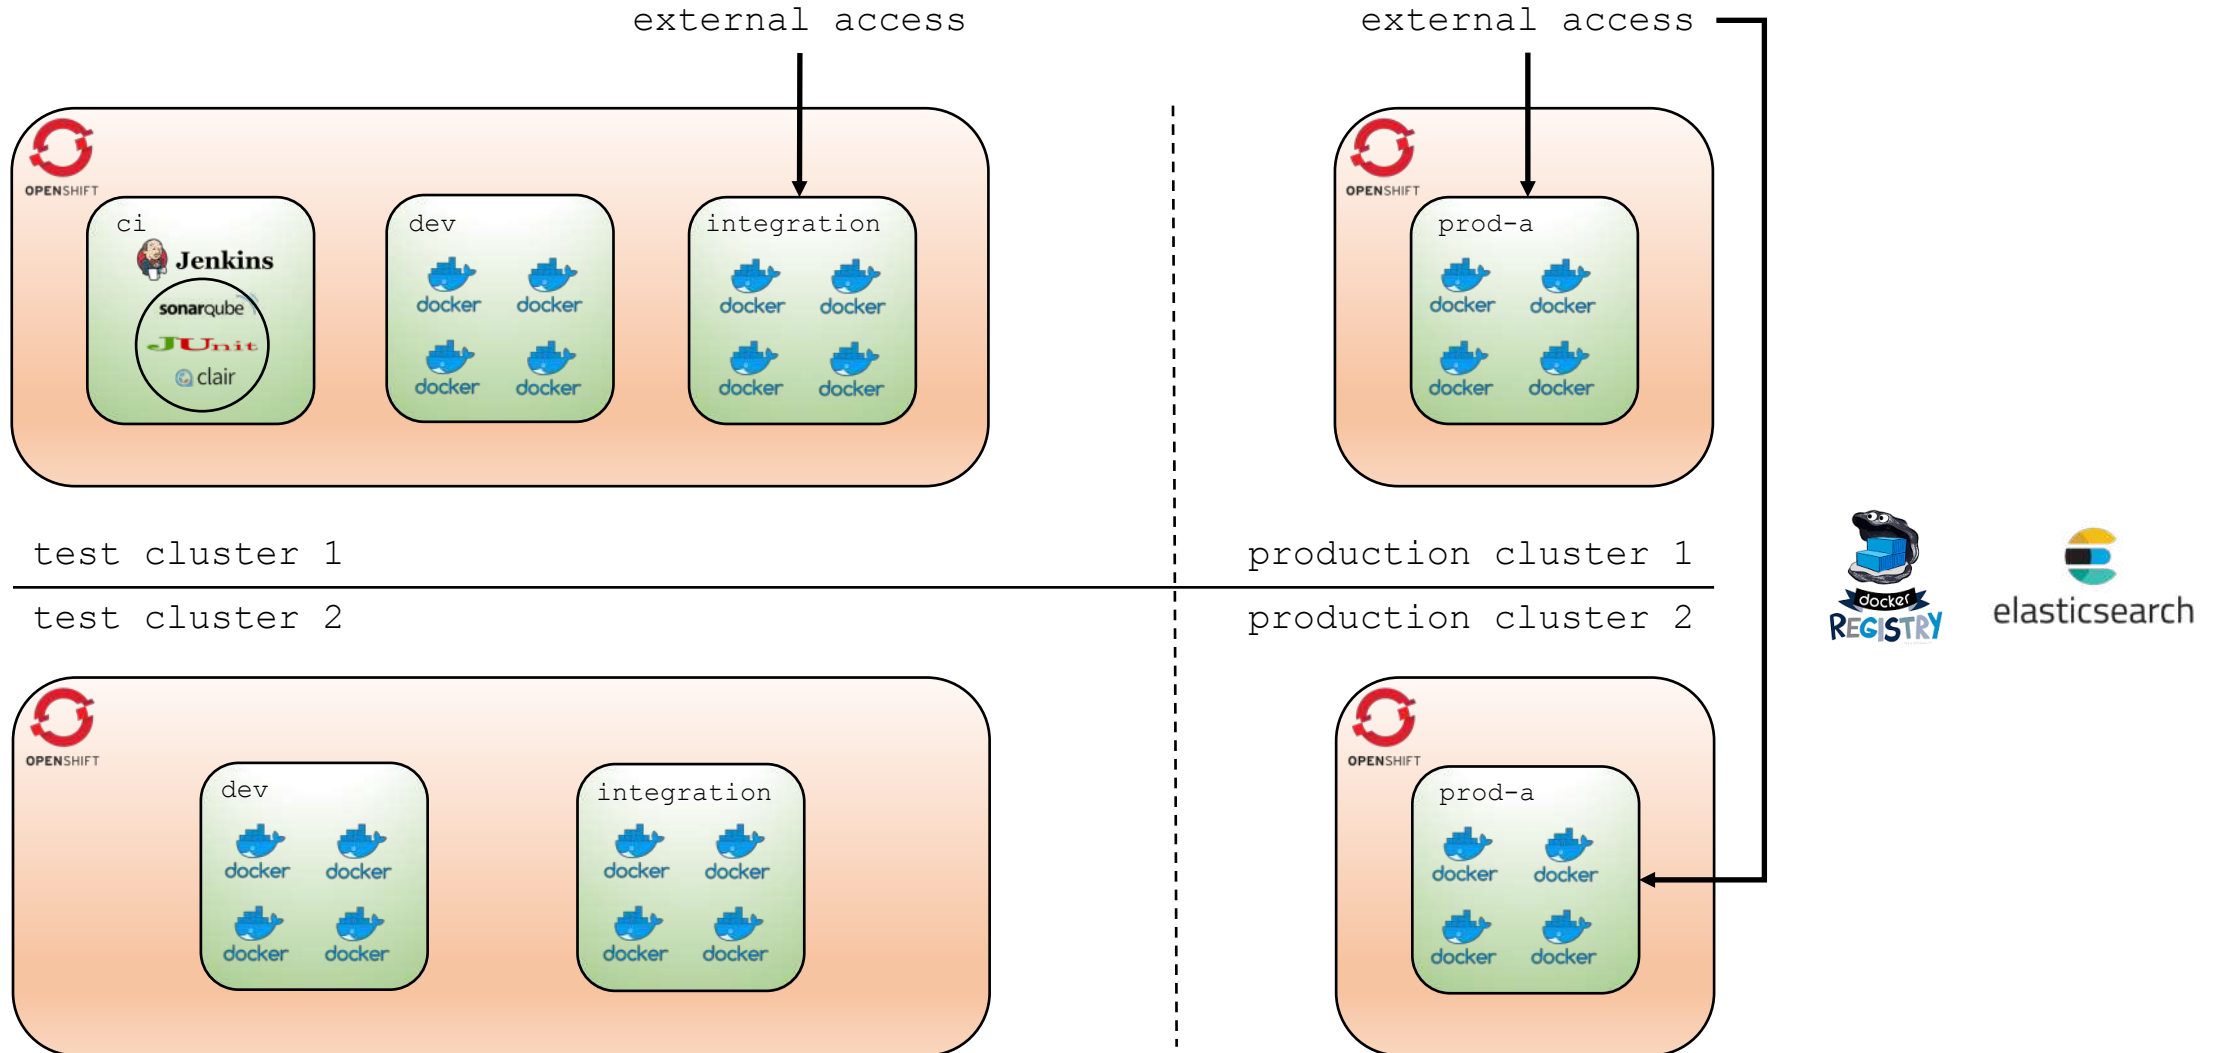
Goals for the new Environment

- Containers, of course
- Secure
 - Services
 - Data
- Clear separation of concerns
 - Secrets
 - Keep up software/frameworks upgrades
- No more monoliths
- Automate as much as possible
 - Build and deploy pipeline
 - Deploy to production during working hours

Road to production



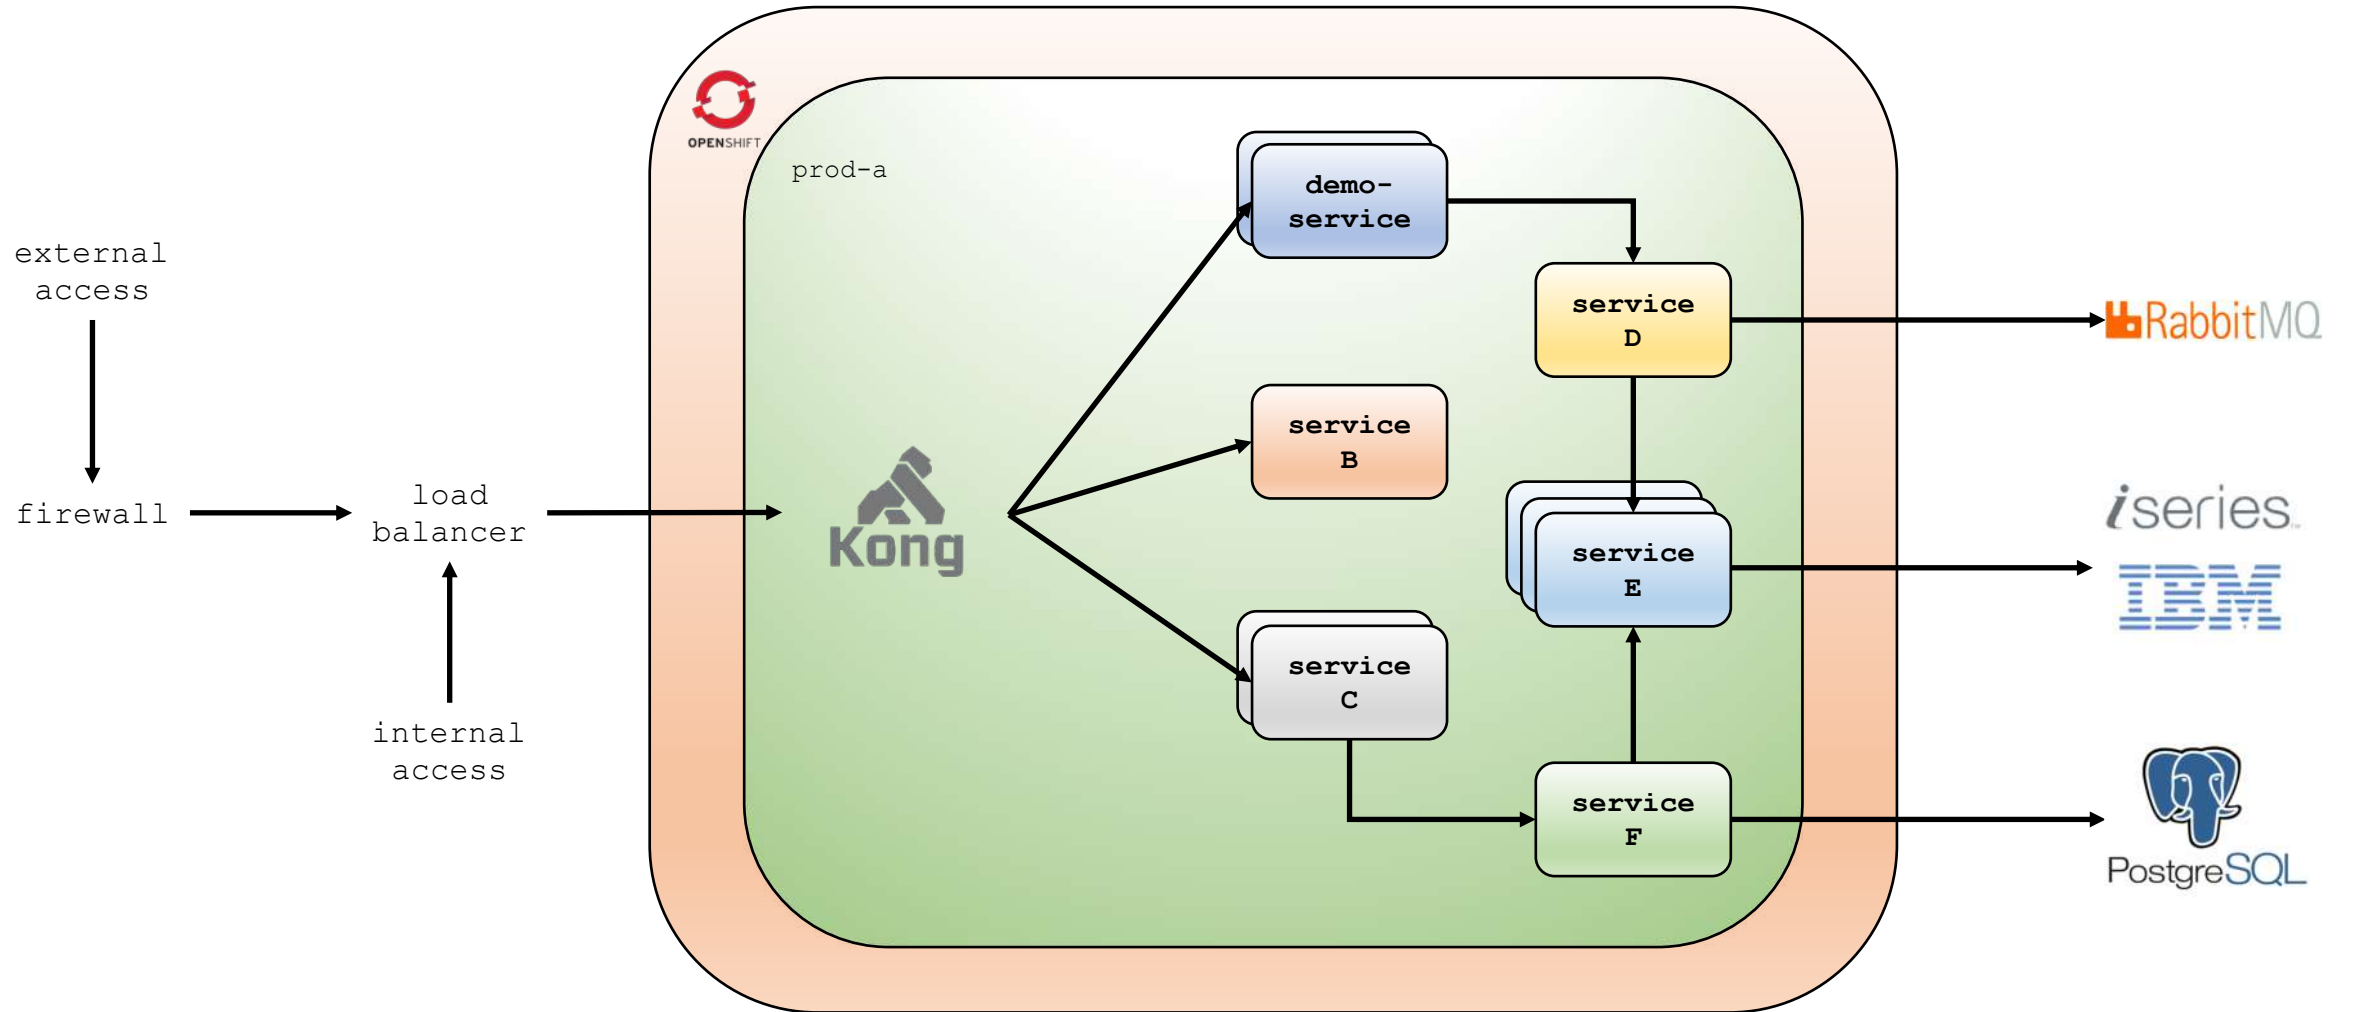
What did we have



Environment

- OpenShift 3.7 on VMware, with two compute/storage tiers per cluster
- Persistent storage using vSphere cloud provider
- One common Elasticsearch instance for all clusters
- One common Openshift Standalone Registry for all clusters

Production



Kong

- Øredev, 2016
 - Jeremy Seitz - Domain-Driven Desire: API architecture, cross-functional teams, and war stories around the topic <https://vimeo.com/191051851>
 - . . . *based on NGINX, it is Open Source* . . .
- Plugin based
 - Key Authentication; API key for consumers
 - IP Restriction; configure service for internal use
 - Prometheus; expose metrics
 - Zipkin; distributed tracing
 - Pre/Post-function; exchange client certificate
 - and some more . . .

Kong configuration

- Configuration via curl commands

```
curl -X POST "http://localhost:8001/services/" \  
  -d "name=demo-service" \  
  -d "url=http://demo-service:8080/api"
```

```
curl -X POST "http://localhost:8001/services/demo-service/routes" \  
  -d "paths[]=api/demo_service" \  
  -d "strip_path=true"
```

```
curl -X POST "http://localhost:8001/services/demo-service/plugins" \  
  -d "name=prometheus"
```

- From version 0.12.2 to 1.3.0, via 9 upgrades, with no downtime

Kong future, perhaps

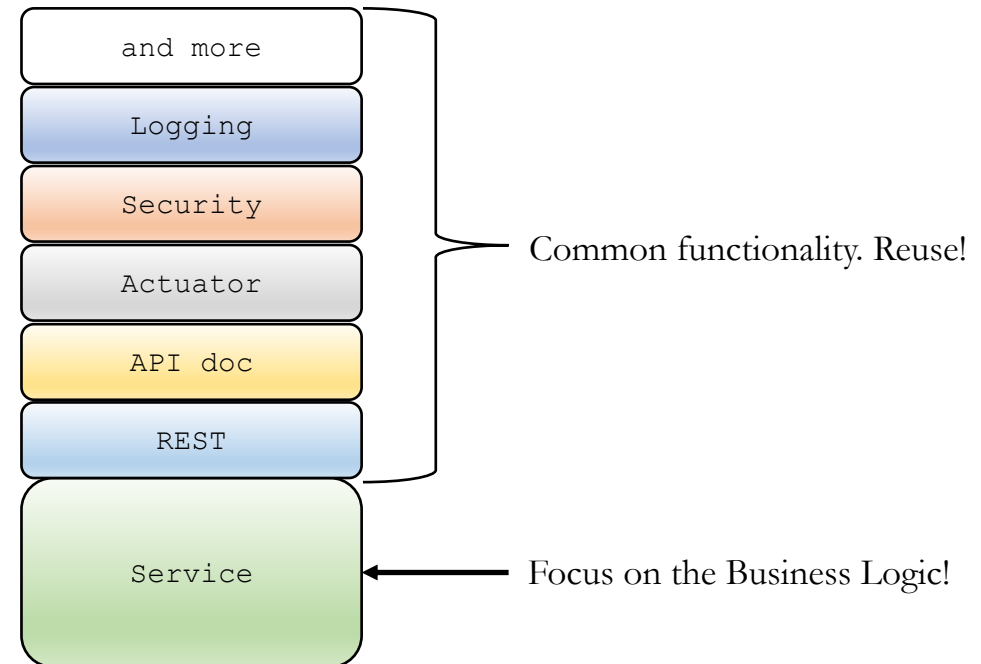
- Remove curl commands
- Use declarative configuration
 - All configuration in yaml format
 - `oc apply/delete`
- Replace HAProxy with Kong Ingress Controller

Service definition

- Endpoints under /api
- OpenAPI spec under /api_docs
- Distributed tracing conforming to OpenTracing standard
- Aggregated application logging to Resurs Bank central log in UTC timestamps
- Validation of JWT authorization header when applicable
- Check required authorities for endpoints when applicable
- Health endpoint (200 OK)
- Metrics endpoint (prometheus format)
- Secrets configured by Operations in Production
- Follow endpoint/naming/coding conventions
- Common error response structure with traceId
- README.md
- Jenkinsfile
- Dockerfile

Spring Boot Service with Starters

- REST
 - /api
- API documentation
 - /api_docs
- Actuator
 - /actuator/*
- Security
 - JWT
 - Authorities
- Logging
- Error handling
- and more . . .



Base images by Operations

- RHEL base image
 - resurs-minimal:latest

```
FROM registry.access.redhat.com/rhel7-minimal:latest
RUN yum -y upgrade
COPY Resurs-RootCA.crt /etc/pki/ca-trust/source/anchors
```

- Java base image
 - resurs-minimal-java11:latest

```
FROM base/resurs-minimal:latest
RUN microdnf update && \
    microdnf --enablerepo=rhel-7-server-rpms install java-11-openjdk-headless --nodocs && \
    microdnf clean all
```

Building service image

- demo-service:b141d0b

```
FROM base/resurs-minimal-java11:latest
ADD demo-service/build/libs/demo-service.jar /app.jar

RUN sh -c 'touch /app.jar'

EXPOSE 8080

ENV JAVA_OPTS="-Djava.security.egd=file:/dev/./urandom"
ENV JAVA_MEM_OPTS=""

ENTRYPOINT ["sh", "-c", "java $JAVA_OPTS $JAVA_MEM_OPTS -jar /app.jar"]
```

Configuration, take 1

- All configuration in git, including secrets
- Developers create required .yaml files

application.yml, in the source project

```
spring.application.name: demo-service
client:
  connect-timeout: 3000
  read-timeout: 15000

---
spring.profiles: openshift

another-service.url: http://another-service:8080/api
datasource.host: postgres
resurs.cloud.secret.directories: /etc/demo-service/secrets
```


postgres.yaml

```
apiVersion: v1
kind: Service
metadata:
  name: postgres
  namespace: prod-a
spec:
  type: ExternalName
  externalName: postgres.resurs.loc
```

configmap.yaml

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: demo-service
  namespace: prod-a
  labels:
    app: demo-service
data:
  application.yaml: |
    spring.profiles: openshift
  client:
    connect-timeout: 2000
    read-timeout: 10000
```

deployment.yaml

```
apiVersion: apps/v1beta1
kind: Deployment
  - env:
    - name: SPRING_PROFILES_ACTIVE
      value: openshift
    - name: SPRING_CONFIG_LOCATION
      value: classpath:application.yml,file:/etc/demo-service/config/application.yml
  volumeMounts:
    - mountPath: /etc/demo-service/config
      name: demo-service
  volumes:
  - configMap:
      items:
        - key: application.yml
          path: application.yml
      name: demo-service
  name: demo-service
```

service.yaml

```
apiVersion: v1
kind: Service
metadata:
  name: demo-service
  namespace: prod-a
  labels:
    app: demo-service
spec:
  selector:
    app: demo-service
  ports:
    - name: http
      protocol: TCP
      port: 8080
      targetPort: 8080
  type: ClusterIP
```

Secret

- Operations
 - Can see and reveal Secret
- Developers
 - Cannot see Secret

```
apiVersion: v1
kind: Secret
metadata:
  name: demo-service-secrets
  namespace: prod-a
data:
  db.username: UmVhbGx5Pz8/Cg==
  db.password: Tm90aGluZ0hlcmUuCg==
```

- But we wanted version control of our secrets!

SealedSecret

- SealedSecrets by Bitnami
 - Kubernetes controller
 - Anyone can encrypt
 - Only the controller can decrypt
- Operations
 - Create SealedSecret
 - Add SealedSecret to Git
- Developers
 - Can see the SealedSecret
- We have version control of our secrets!

```
apiVersion: bitnami.com/v1alpha1
kind: SealedSecret
metadata:
  name: demo-service-secrets
  namespace: prod-a
spec:
  encryptedData:
    db.username: AgBE7kx1fs0gbXZZ...
    db.password: AgAZ8Qz3J2krU2V8H...
```

Configuration, take 2

demo-service

└─ kubernetes

└─ openshift

└─ application-**dev**.yaml

└─ application-**integration**.yaml

└─ application-**prod**.yaml

└─ **deployment-config.json**

deployment-config.json

```
{
  "template": "springboot",
  "environments": {
    "base": {
      "labels": {
        "team": "softwaresolutions"
      },
      "kong": {
        "expose": "public"
      },
      "secrets": [{
        "name": "demo-service-secrets",
        "mountPath": "/etc/demo-service/secrets"
      }]
    },
    "dev": {
      "replicas": 1,
      "cpuLimit": 1
    },
    "integration": {},
    "prod": {}
  }
}
```


generated-deployment-list.yaml

- ConfigMap
 - application-<environment>.yaml
- Service
 - labels
 - selector
 - namespace
 - type
 - Ports
- *API Gateway configuration*
 - *curl commands*
- Deployment
 - labels
 - namespace
 - affinity
 - image
 - env
 - ports
 - livenessProbe/readinessProbe
 - resources
 - volumeMounts

Generated configuration

<namespace>

└─ apps

└─ demo-service

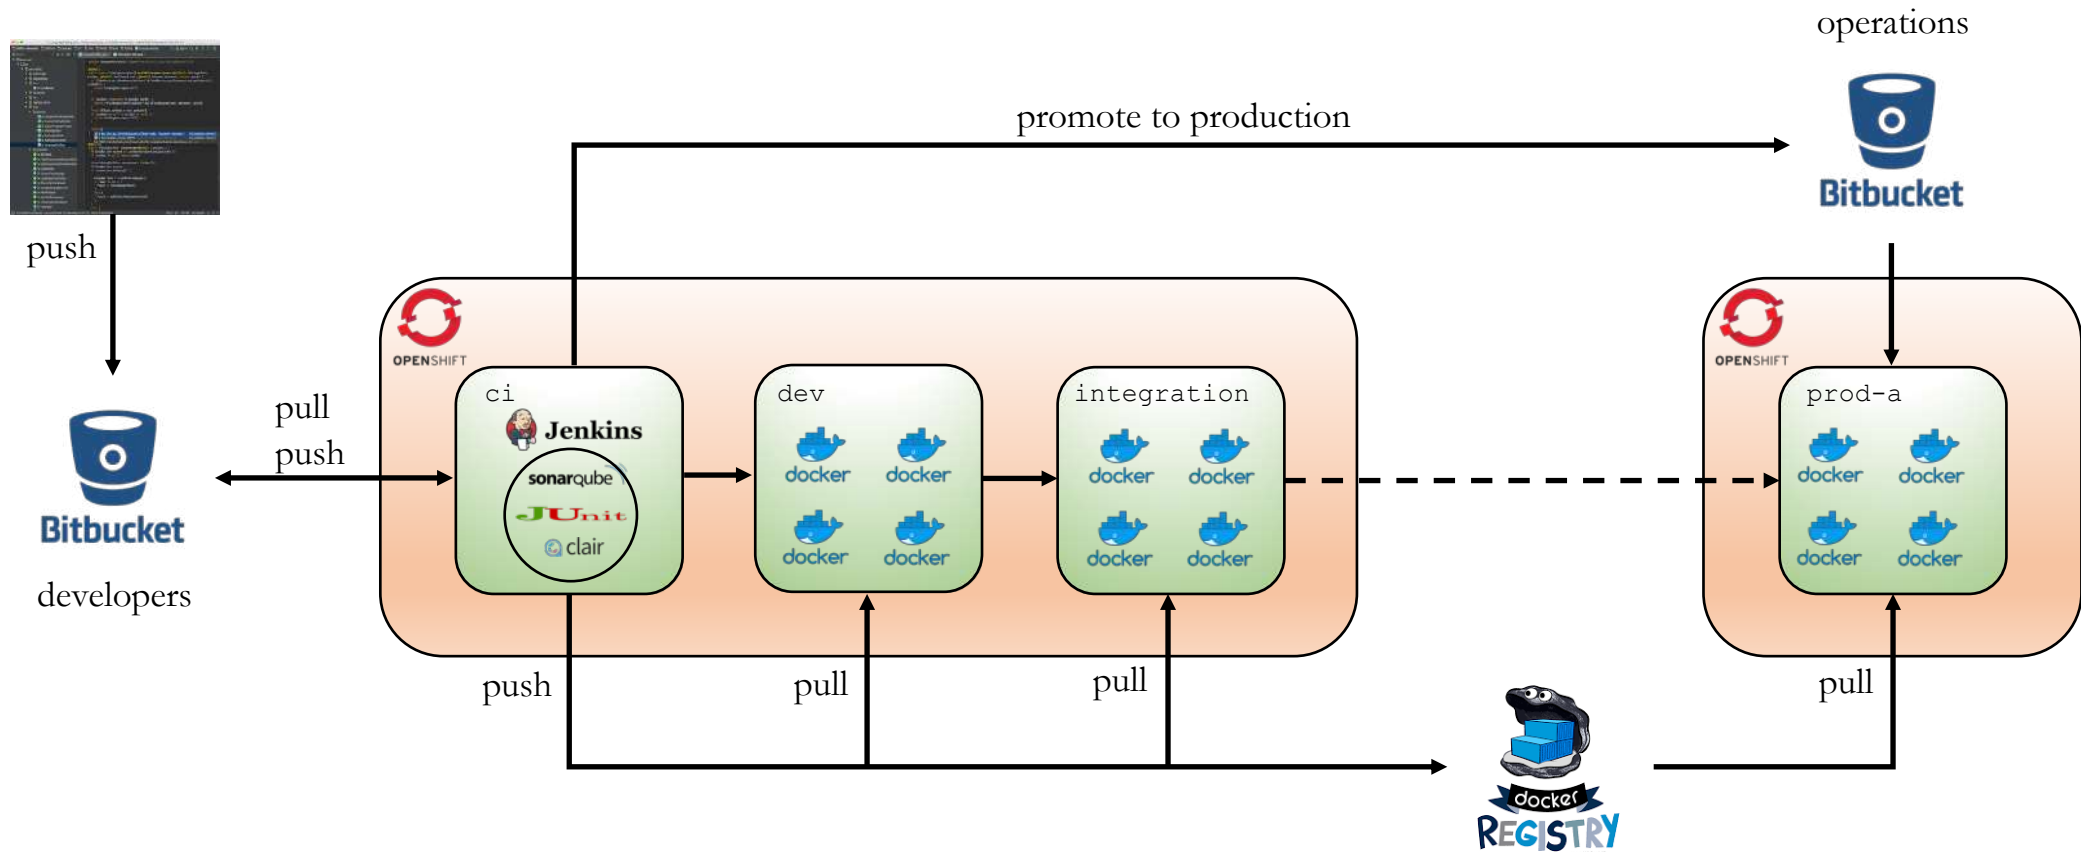
└─ demo-service-secrets.yaml

└─ **generated-deployment-list.yaml**

└─ **kong-configuration.txt**

- apiVersion: apps/**v1beta1** → apps/**v1beta2** → apps/**v1**

Pipeline



Jenkins create PR to Operations

Bitbucket Projects Repositories

Search for code, commits or repositories...

Openshift / openshift-production

Create pull request


Pull requests

FILTER BY: Merged S-Bitbucket-Jenkins Target branch I'm reviewing

Summary

| | Reviewers |
|---|-----------|
| Promote compass-client to prod → master S-Bitbucket-Jenkins - #520, last updated 45 minutes ago | Q 1 +6 |
| Promote notification-service to prod → master S-Bitbucket-Jenkins - #519, last updated 2 hours ago | Q 1 +6 |
| [NEW] Promote customer-card-service to prod → master S-Bitbucket-Jenkins - #508, last updated 2 hours ago | Q 1 +7 |
| Promote omni-account-payments-service to prod → master S-Bitbucket-Jenkins - #513, last updated 2 hours ago | Q 1 +7 |
| Promote notification-service to prod → master S-Bitbucket-Jenkins - #514, last updated 2 hours ago | Q 1 +6 |

New service with PR template


Source view Diff to previous History ▾
Show source Edit Blame Raw file

Service info

Please replace all the placeholders below with information about the new service.
Note: This form must be filled out 8 workdays prior to the requested release date.

Service name

`${serviceName}`

Contact person(s)

`<Firstname Lastname PhoneNumber>`

Service type

`<API | Web application | Other>`

System owner

`<Firstname Lastname Teamname>`

Dependencies (external and/or internal) to services not running in Openshift:

`<None | List of dependencies>`

Database setup :

`<None | Instructions>`

ServiceDesk Plus integration

- Standard change
- Change management gets what they want
- Deploy to Production, any time, almost . . .



S-Bitbucket-Jenkins

Change created in Servicedesk+ with id [2754](#)

[Reply](#) · [Delete](#) · [Create task](#) · [Create Jira issue](#) · [Like](#) · 14 Aug 2019



S-Bitbucket-Jenkins

OPENED

the pull request 14 Aug 2019

Faster deploy to production

- Legacy environment
 - Announce intent to deploy on Thursday, deployed on Tuesday 06.00
- OpenShift
 - Approve your PR by 09.00 on Tue/Thu, deployed immediately
- ~800 deployments to Production in 18 months
 - < 10 of them have been rolled back

Deploy all services

- `oc apply -recursive -f .`
- Unfortunately startup order of services matters
 - Configure `PriorityClass`
 - Use `priorityClassName` in `Deployment` for service
- Flux*
 - The state in the cluster matches the configuration in git, automatically

Upgrading OpenShift

- 3.7 → 3.10 was ok, but . . .

Upgrading OpenShift

- 3.7 → 3.10 was ok, but . . .
- 2018-11-20

PV/PVC failed using vSphere-managed storage

```
E1123 09:44:27.960669      1 vsphere.go:1077]
```

```
Failed to get shared datastore:
```

```
No shared datastores found in the Kubernetes cluster for nodeVmDetails:
```

```
[{NodeName:node1 vm:0xc42501bf20 VMUUID:}  
 {NodeName:node2 vm:0xc42501bf60 VMUUID:}  
 {NodeName:node3 vm:0xc42501bfa0 VMUUID:}  
 {NodeName:nodeX vm:0xc42501bfe0 VMUUID:}]
```

vSphere Storage will **NOT** work

- 2018-12-19

This configuration will **NOT** work.

I(eng) also think that, this **never worked** in 3.7 too.

- 2019-01-03

In Openshift-3.7 we(I am using we but it is vmware that changed this code) **used to allow** provisioning of volumes even if datastores being used was not shared between all VMs in the cluster.

Ceph

- Decided to move persistent storage from vSphere cloud provider to Ceph.
- New Ceph cluster installed, spanning both datacenters and all three compute/storage tiers. Block based (RBD) and S3 style access.
- OpenShift uses RBD storage
- Services use S3 storage

Continuing the upgrade

- Migrate from vSphere cloud provider to Ceph storage
- Upgrade test clusters to 3.11
- Upgrade production clusters to 3.11
- Install and configure new OpenShift 4.2 clusters

Why not 4.1?

- You can disable Telemetry, but ...
 - You cannot perform subscription management
 - No disconnected subscription management
- Waiting for 4.2
 - Can we do disconnected subscription management
 - Disable Telemetry

| Q3 CY2019 OpenShift 4.2 | |
|----------------------------|--|
| DEV | <ul style="list-style-type: none">• Developer Console GA• Serverless w/ Knative Tech Preview• OpenShift Pipelines (Tekton) Tech Preview• CodeReady Containers GA• Developer CLI (odo) GA |
| APP | <ul style="list-style-type: none">• GPU metering• OperatorHub Enhancements• Operator Deployment Field Forms• Application Binding with Operators• Application Migration Console |
| PLATFORM | <ul style="list-style-type: none">• Kubernetes 1.14 w/ CRI-O runtime• Disconnected Install and Update• Automated Installer for Azure, OSP, GCP• OVN Tech Preview• FIPS• Federation Workload API• Automated App cert rotation• OpenShift Container Storage 4.2 |
| HOSTED | <ul style="list-style-type: none">• UHC Multi-Cluster deployment• Proactive Support Operator |

Monitoring

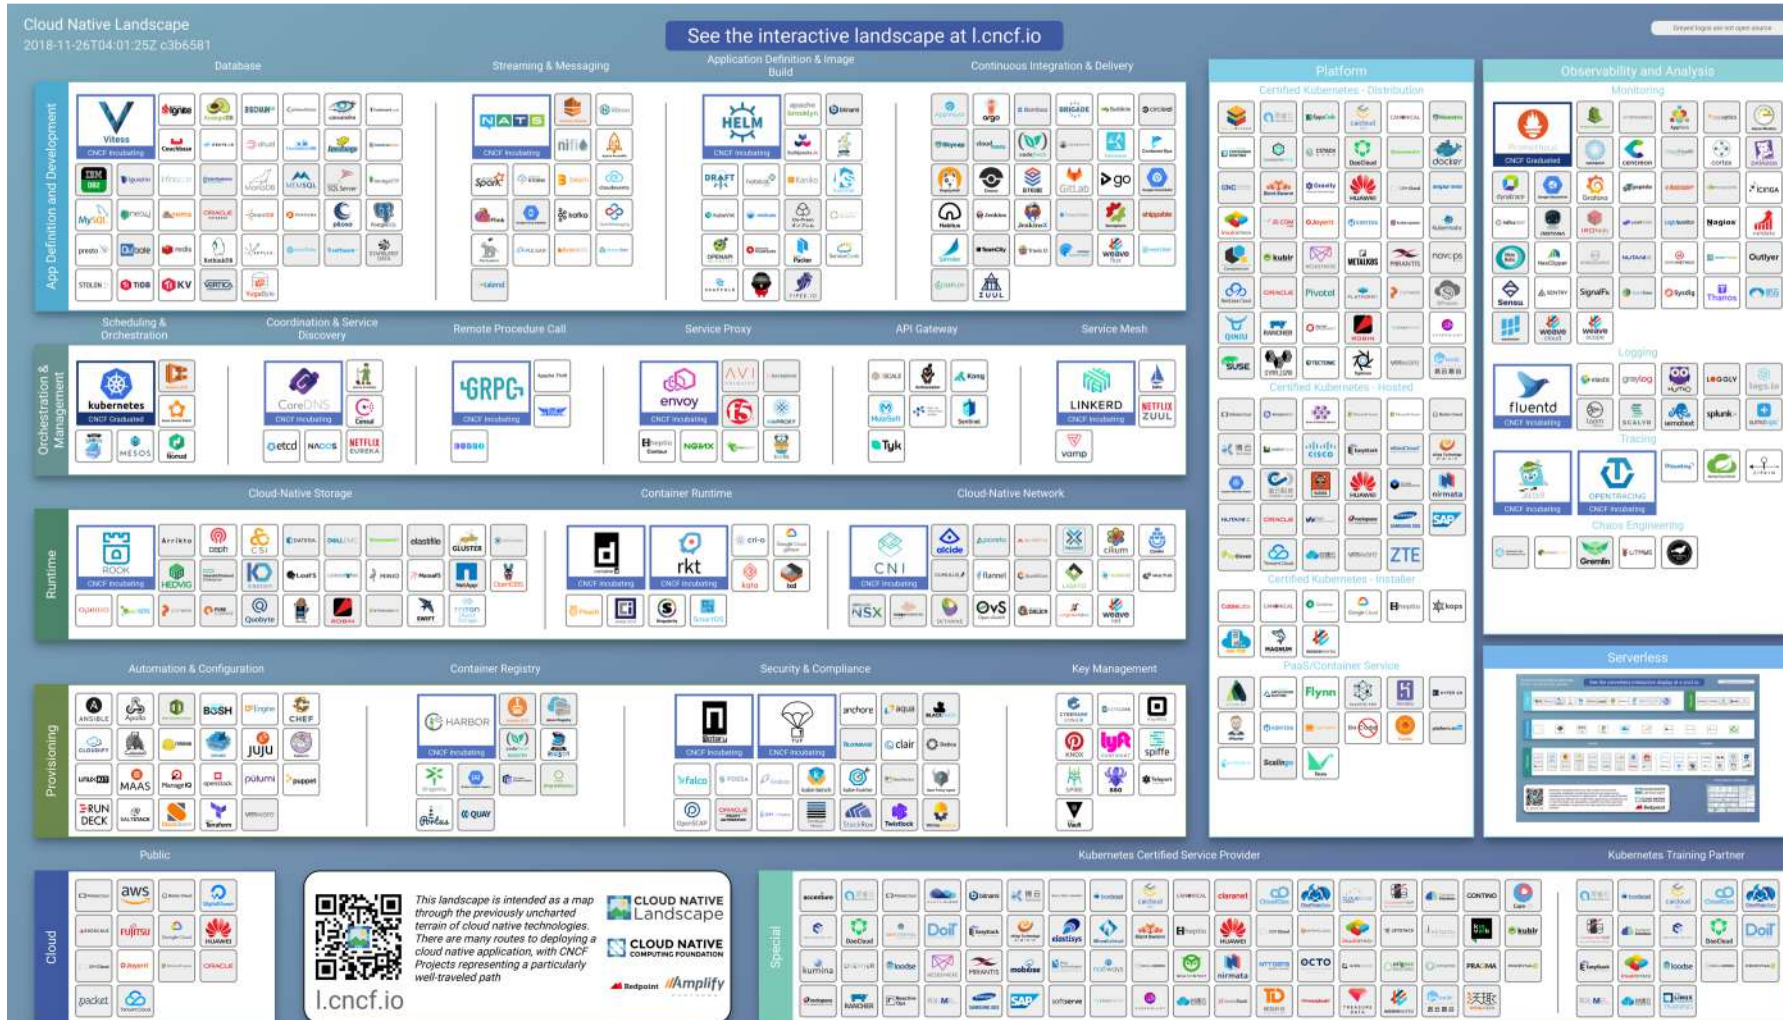
- AppDynamics is used for legacy services
- In OpenShift, just a new base image
 - FROM base/**resurs-minimal-java8-appdyn**:latest
 - And two AppDynamics environment variables
- But . . .

AppDynamics issues in OpenShift

- Needs agent
- License count not managed properly
- Problems with picking up some traffic
- Dependent of specific technology and versions
 - No Java 11 support, yet
 - No project Reactor support, yet




CNCF Landscape

<https://landscape.cncf.io>


















Graduated projects

Graduated CNCF Projects (3)

| | | |
|---|---|--|
|  envoy Envoy Cloud Native Computing Foundation (CNCF) ★ 7,406 |  kubernetes Kubernetes Cloud Native Computing Foundation (CNCF) ★ 44,828 |  Prometheus Prometheus Cloud Native Computing Foundation (CNCF) ★ 20,488 |
|---|---|--|

Incubating CNCF Projects (16)

| | | | | |
|---|---|---|---|---|
|  CNI Container Network Interface (CNI) Cloud Native Computing Foundation (CNCF) ★ 1,657 |  containerd containerd Cloud Native Computing Foundation (CNCF) ★ 3,202 |  CoreDNS CoreDNS Cloud Native Computing Foundation (CNCF) ★ 2,930 |  fluentd Fluentd Cloud Native Computing Foundation (CNCF) ★ 7,071 |  gRPC gRPC Cloud Native Computing Foundation (CNCF) ★ 18,282 |
|  HARBOR Harbor Cloud Native Computing Foundation (CNCF) ★ 5,151 |  HELM Helm Cloud Native Computing Foundation (CNCF) ★ 8,786 |  JAEGER Jaeger Cloud Native Computing Foundation (CNCF) ★ 6,551 |  LINKERD Linkerd Cloud Native Computing Foundation (CNCF) ★ 4,765 |  NATS NATS Cloud Native Computing Foundation (CNCF) ★ 4,906 |
|  |  |  |  |  |

Prometheus in OpenShift

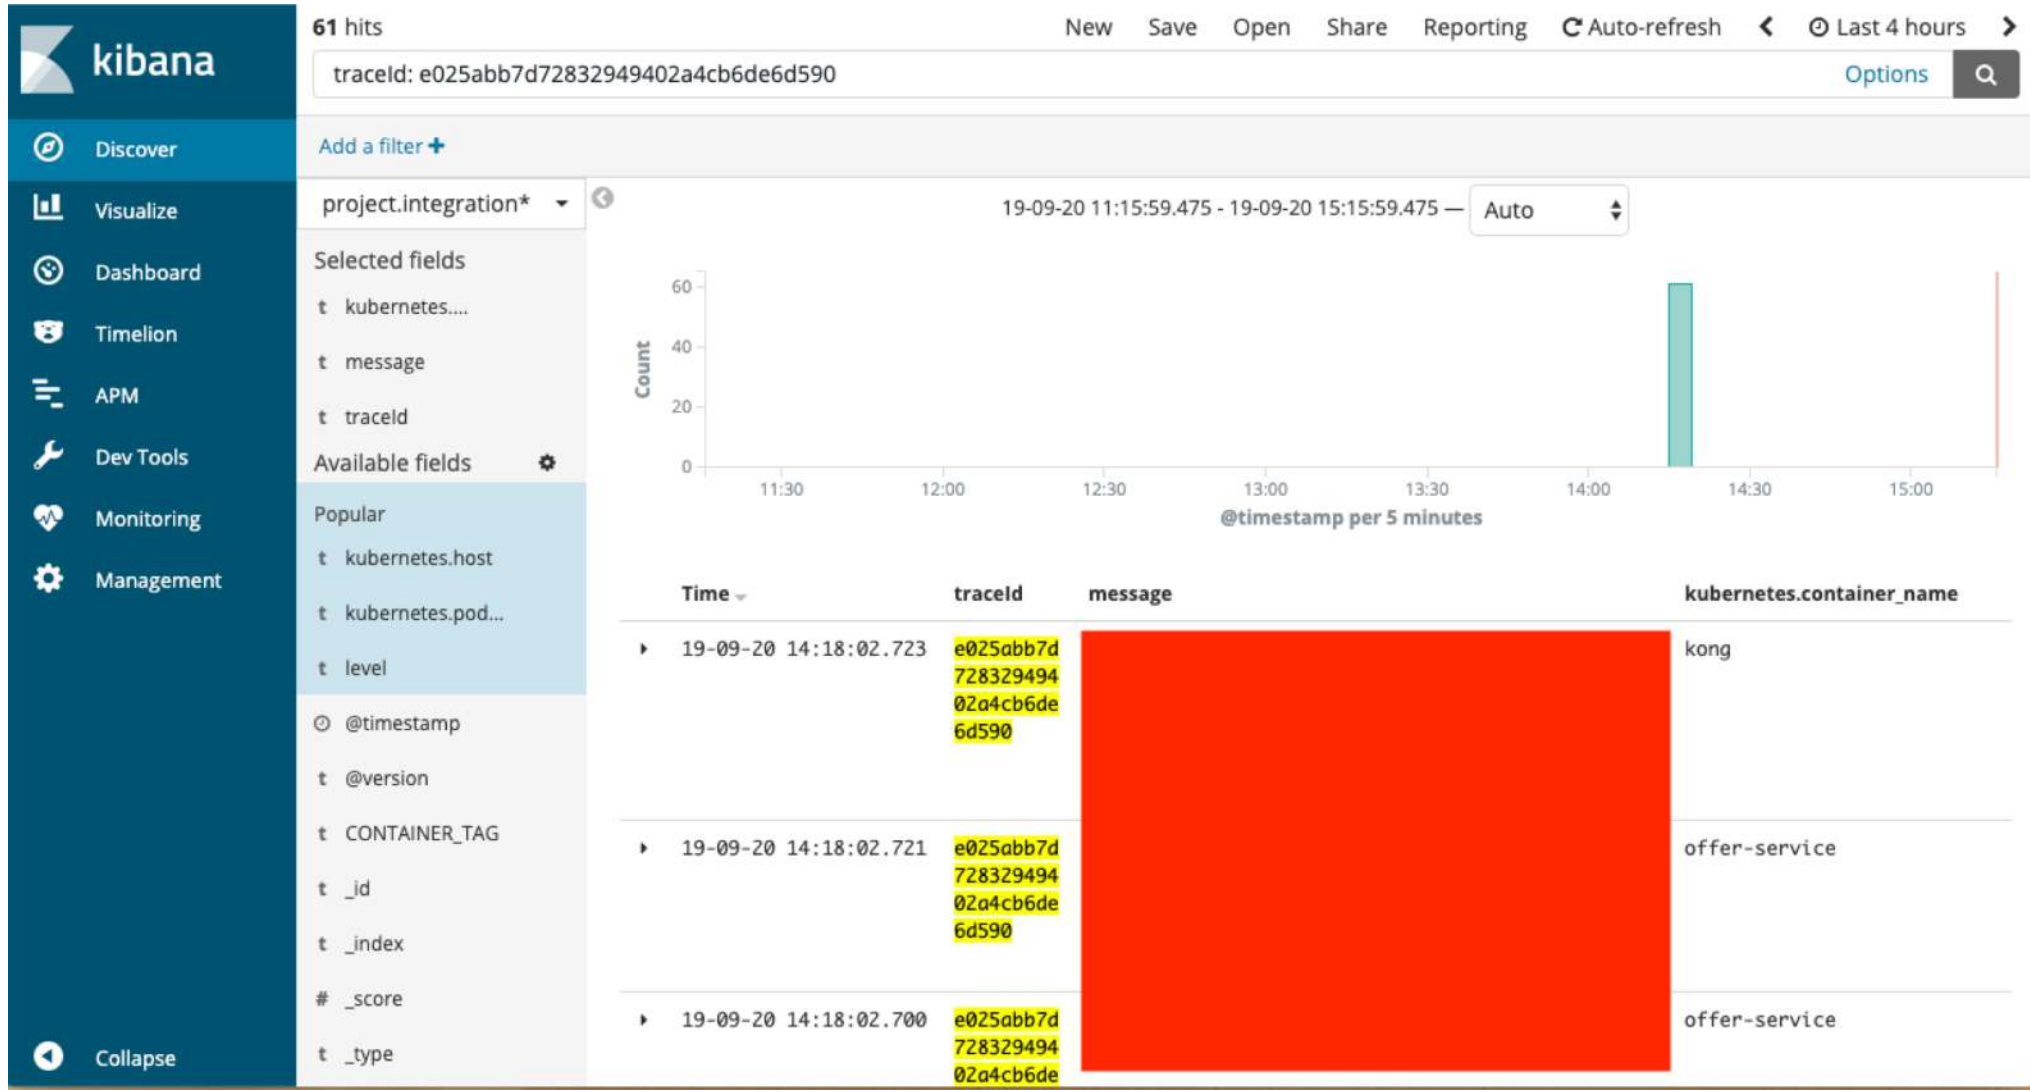
- Prometheus is the de-facto standard for Cloud Native applications
- All services are already Prometheus ready

| Feature | OCP 3.9 | OCP 3.10 | OCP 3.11 |
|-------------------------------|---------|----------|----------|
| Prometheus Cluster Monitoring | TP | TP | GA |

*Prometheus, a CNCF project that collects time-series data as a source for triggering alerts, has emerged as a **leading standard for cloud-native monitoring** within Kubernetes.**

*<https://www.redhat.com/en/blog/generally-available-today-red-hat-openshift-container-platform-311-ready-power-enterprise-kubernetes-deployments>

Kibana

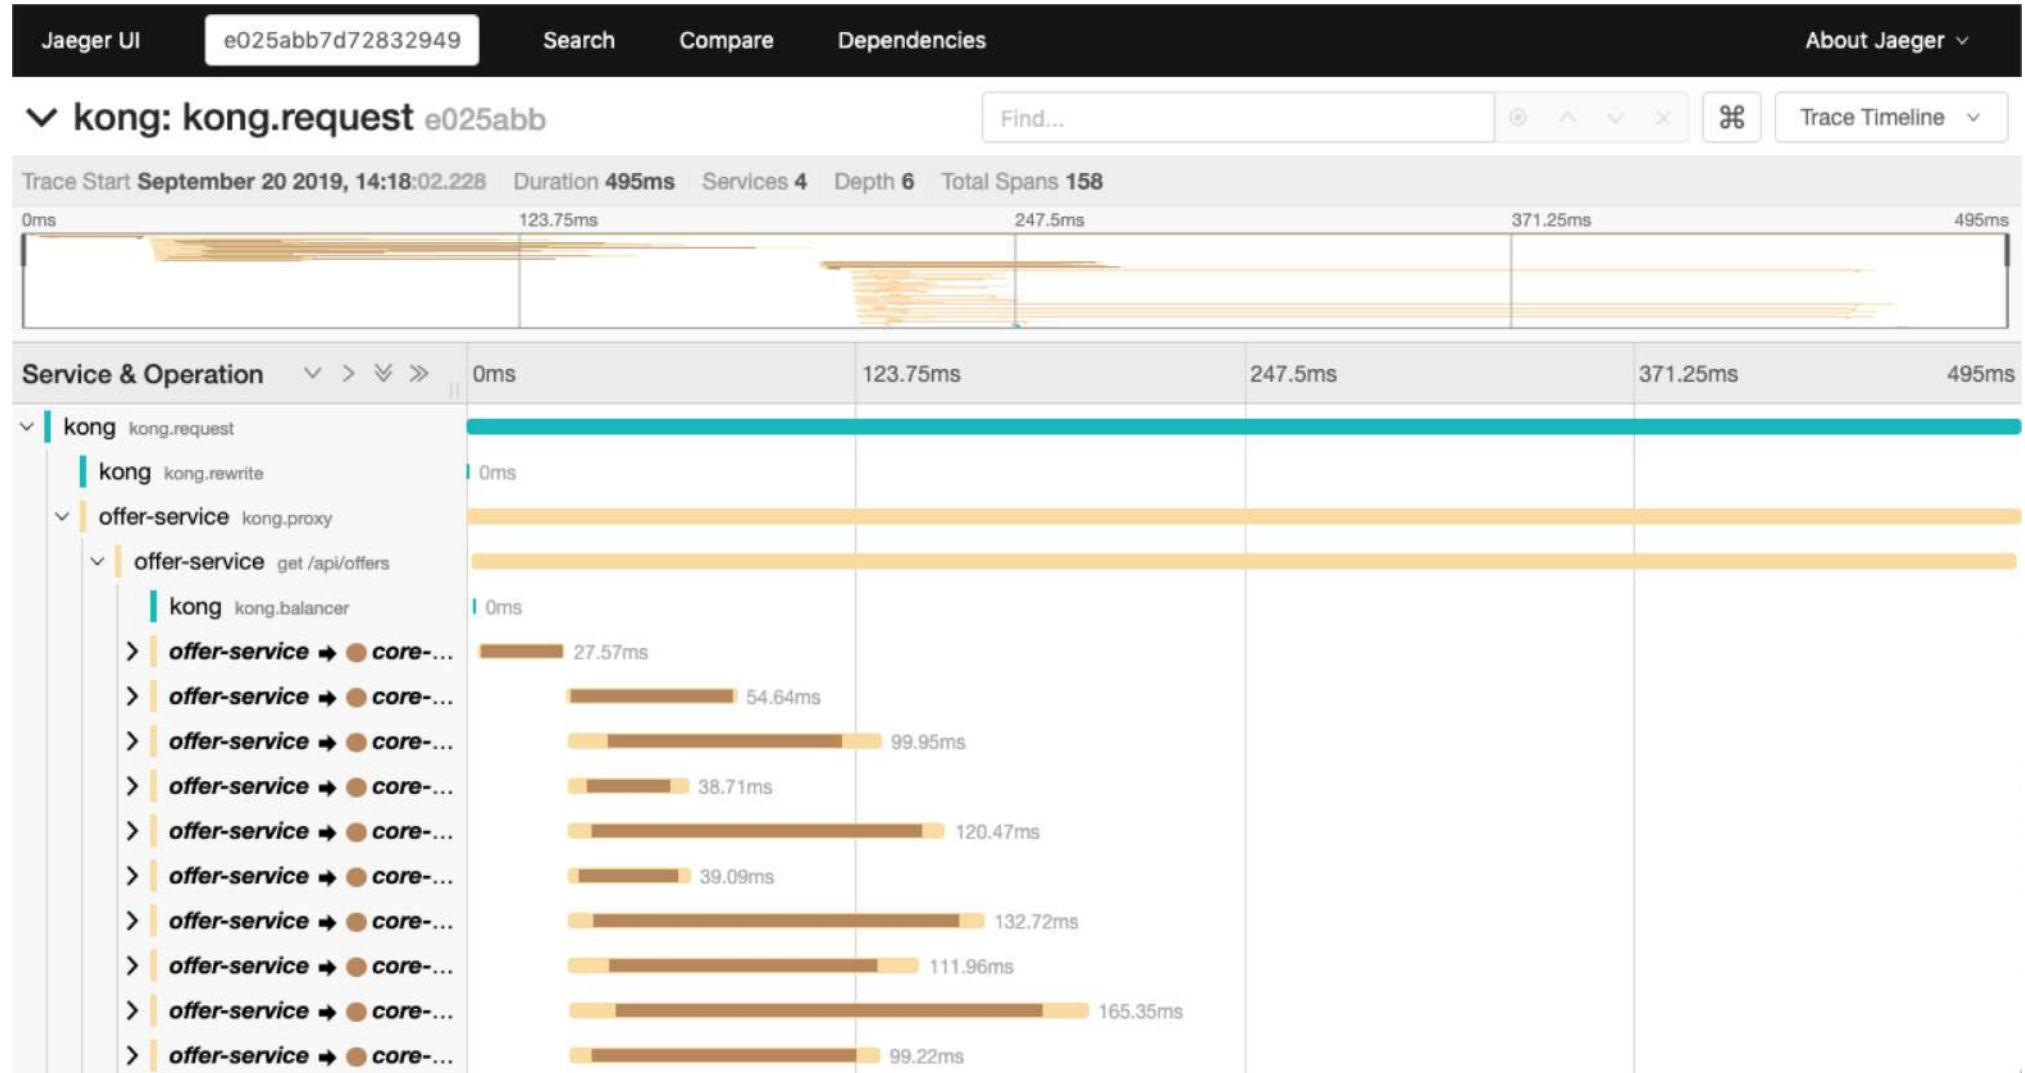


log_format for Kong

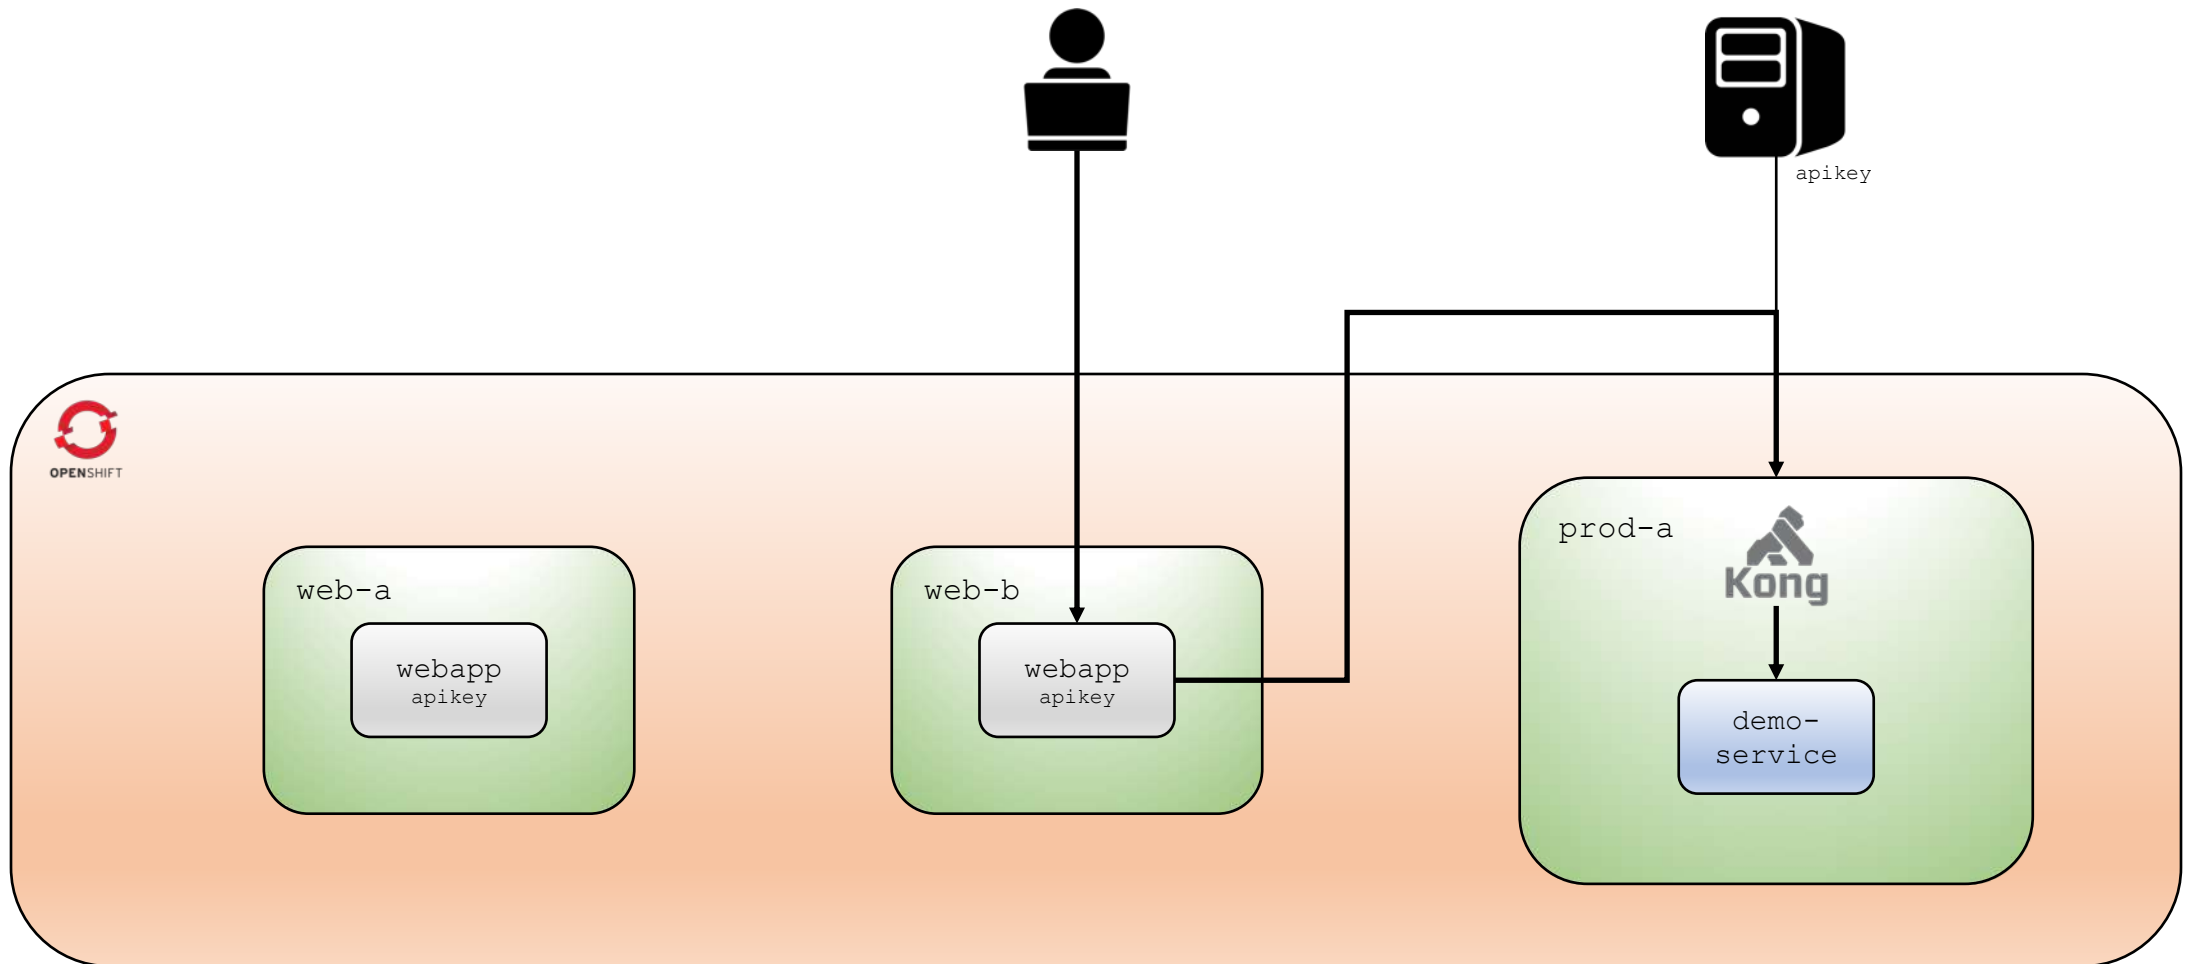
```
{  
    "traceId":    "$http_x_b3_traceid",  
    "spanId":     "$http_x_b3_spanid",  
    "consumer":   "$http_x_consumer_username",  
    "status":     $status,  
    "duration":   $request_time,  
    "message":    <standard access log format>  
}
```

- duration > 1.5 AND status: 200
- consumer: THE_CONSUMER AND status: 500 AND message: "/api/demo_service/xyz"

Jaeger



New namespace for web application



Status of today

- 85 services
- 11 web applications
- All teams have something running in OpenShift
 - Even departments not part of IT

The future, a never ending story

- Legacy application namespace
- Quarkus
- Operators
- OpenShift Pipelines (Tekton)
- Service Catalog
- Serverless with Knative
- Service Mesh
- NetworkPolicy
- Image signing

And finally

- Succeed with a small, focused team
- Make OpenShift the place where the developers want to be
- It's hard to satisfy everyone
- Automate everything

