
CERN

One Tool to Rule Them All: How CERN Runs Application Servers on Kubernetes

Antonio Nappi

Few words about the speaker

Antonio Nappi
DevOps Engineer at CERN

Linkedin
https://www.linkedin.com/in/nappiantonio/

What is CERN?

Member States:
Austria, Belgium, Bulgaria, Czech Republic, Denmark, Finland, France, Germany, Greece, Hungary, Israel, Italy, Netherlands, Norway, Poland, Portugal, Romania, Serbia, Slovakia, Spain, Sweden, Switzerland and United Kingdom

Associate Members in the Pre-Stage to Membership:
Cyprus, Slovenia

Associate Members:
India, Pakistan, Turkey, Ukraine

Observers to Council:
Japan, Russia, United States of America, the European Commission, Joint Institute for Nuclear Research and UNESCO

~ 3 000 members of personnel
Users from 100 different countries

Budget ~ 1 100 MCHF per year

CERN: a unique environment

Study fundamental particles

- How they interact
- Understand the fundamental laws of nature

Large Hadron Collider (LHC)

- Largest particle collider in the world
- 27 km in circumference
- Thousand of magnets

Place where the Web was born

Science for peace

- Melting pot

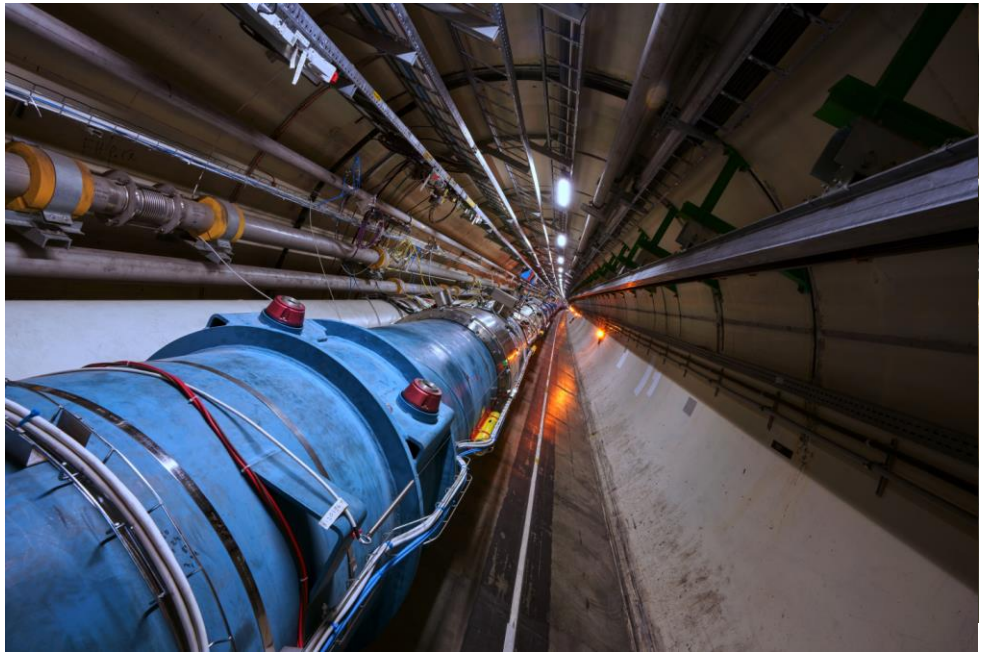

CERN Openlab

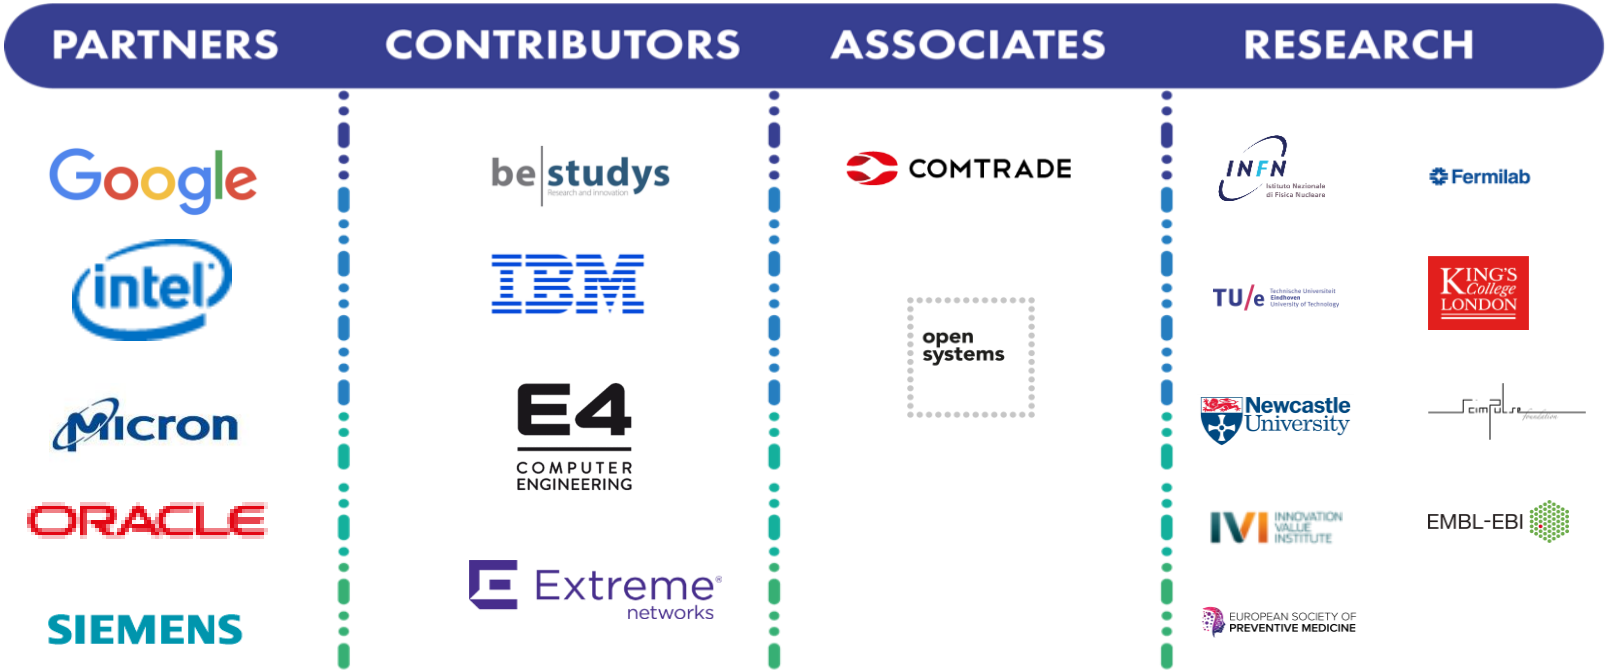

Applications servers at CERN

WebLogic 12.1.3

- ~ 358 Clusters
- ~ 100 Applications
  - DEV / TEST / PREPROD / PROD
- Web Profile application Stateful
  - No Java Message Service
  - No Enterprise JavaBeans
- Users
  - Engineers
  - Administration
  - IT

Tomcat

- Third party application
- 38 VMs
  - JDK7
  - JDK8
- Single instance applications
  - License problem

Applications servers on K8S

WebLogic on Kubernetes

- Different versions
  - 12.1.3
  - 12.2.1
- ~ 30 deployment
  - Across DEV/TEST/PREPROD
- 1 application prod

Tomcat

- Prototype
- Run legacy applications
  - Work just with JDK7

Why Kubernetes

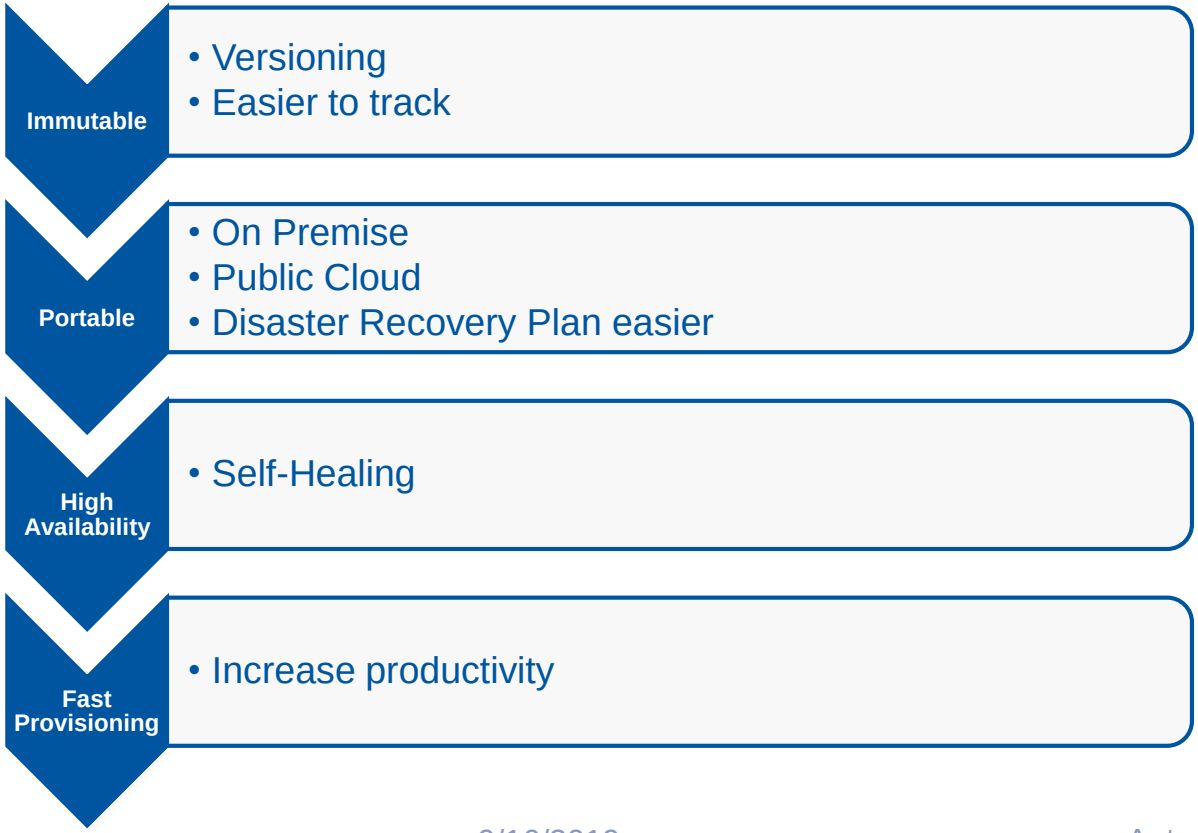

Constraints

Transparent for developers

Replicate VM model

Logging

Monitoring

CERN Infrastructure

OpenStack Private Cloud

Built for physics workloads

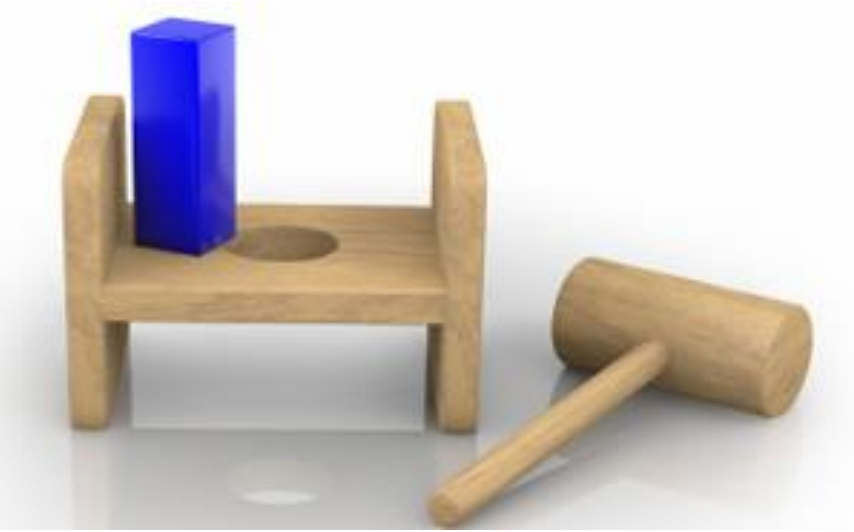

Current architecture

Kubernetes provisioning

Private Cloud/OpenStack Magnum

HA Production scenario

Application spread across 3 AZS

Weighted Round Robin

Canary deployments

Everything ephemeral

One application per K8S namespace

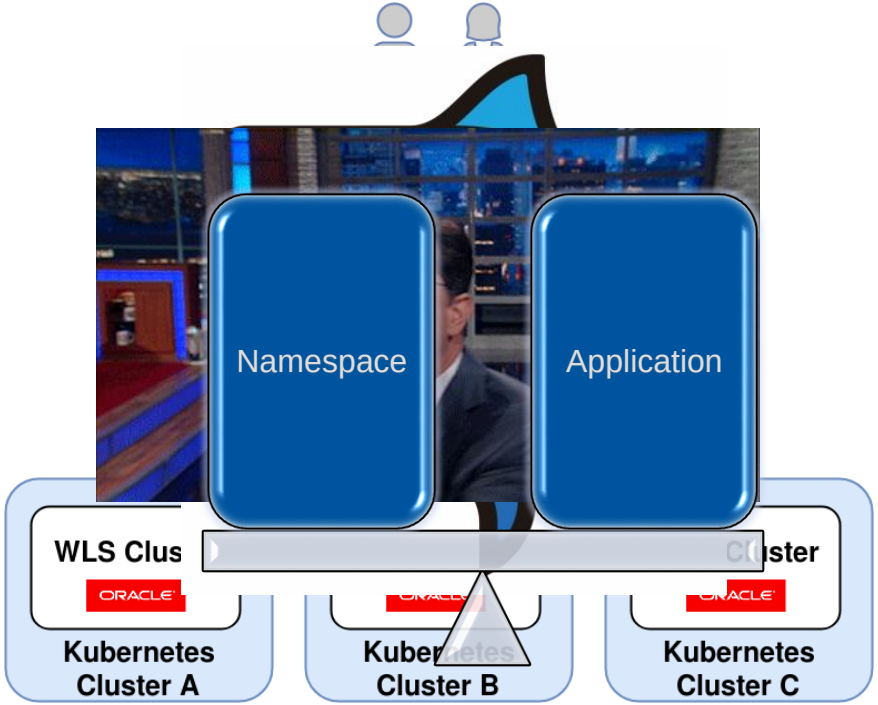

Current architecture

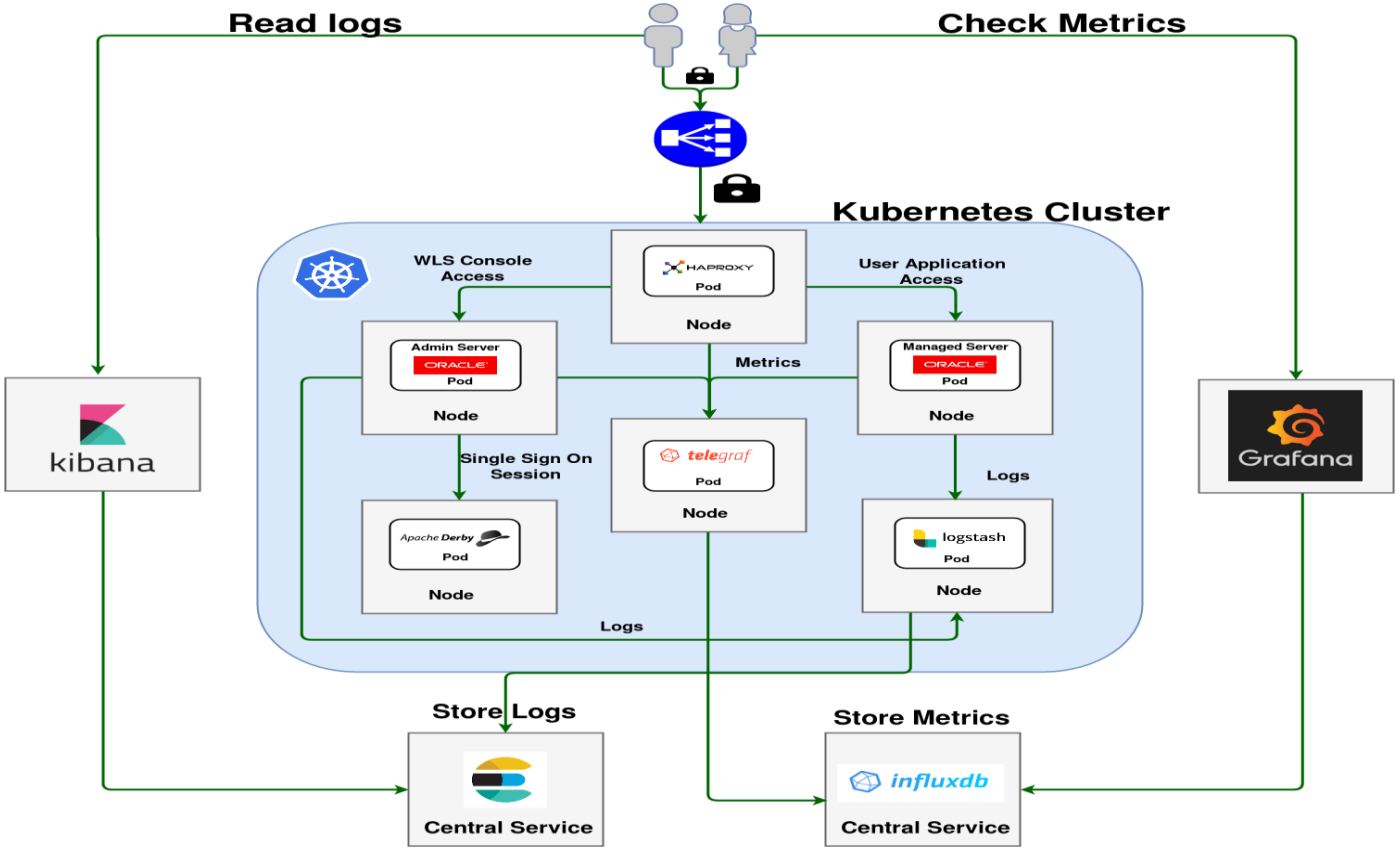

Load balancer

HA External Load Balancer

Pacemaker & Corosync

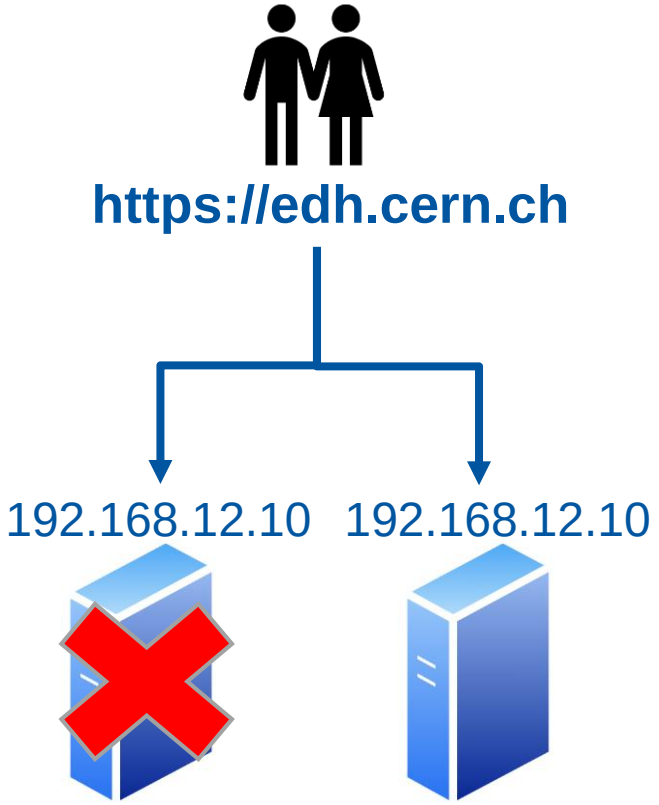

Load balancer

HA External Load Balancer

- Pacemaker & Corosync
- Discriminate traffic based on url

https://edh.cern.ch

URL	Cluster	Node:Port
edh.cern.ch	A	Node1:32000, Node2:32000
apt.cern.ch	B	Node1:32000, Node2:32000
e-groups.cern.ch	C	Node1:32000, Node2:32000
vip-events.cern.ch	A,B,C	A: Node1:32001, Node2:32001 B: Node1:32001, Node2:32001 C: Node1:32001, Node2:32001

Load balancer

HA External Load Balancer

- Pacemaker & Corosync
- Discriminate traffic based on url

Ingress HAProxy

- The only point of access
- No LoadBalancer service type

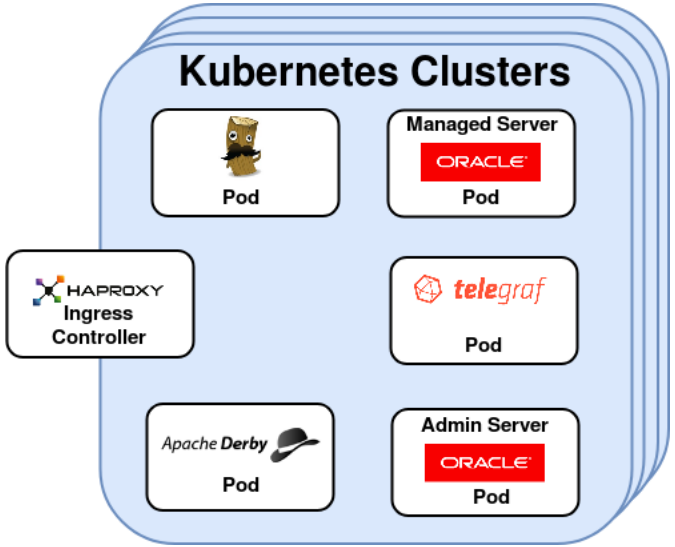

Load balancer

HA External Load Balancer

- Pacemaker & Corosync
- Discriminate traffic based on url

Ingress HAProxy

- The only point of access
- No LoadBalancer service type

Prevent unauthorized access

- No firewall control
- Shared secret

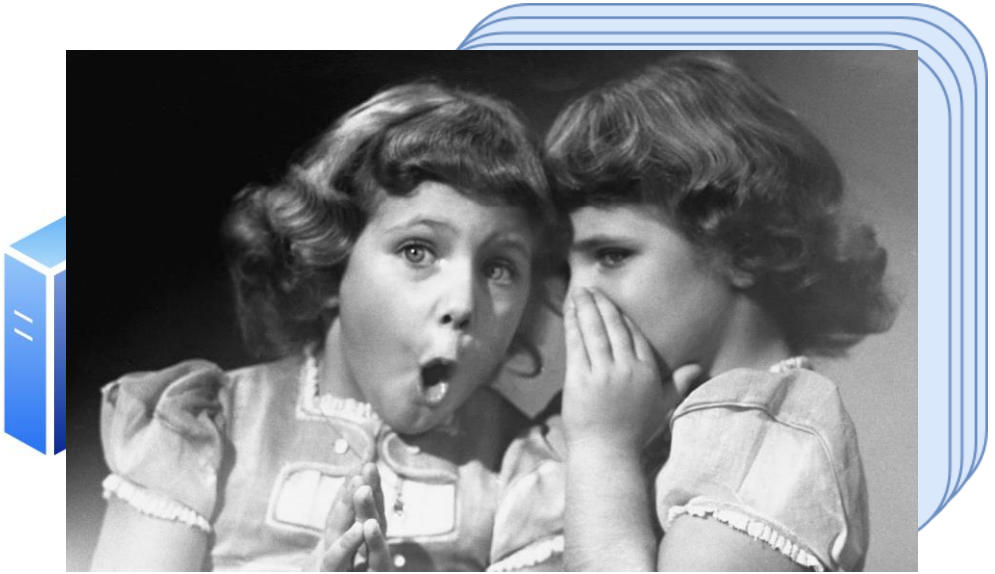

Load balancer

HA External Load Balancer

- Pacemaker & Corosync
- Discriminate traffic based on url

Ingress HAProxy

- The only point of access
- No LoadBalancer service type

Prevent unauthorized access

- Firewall no customizable
- Share a secret

Logging

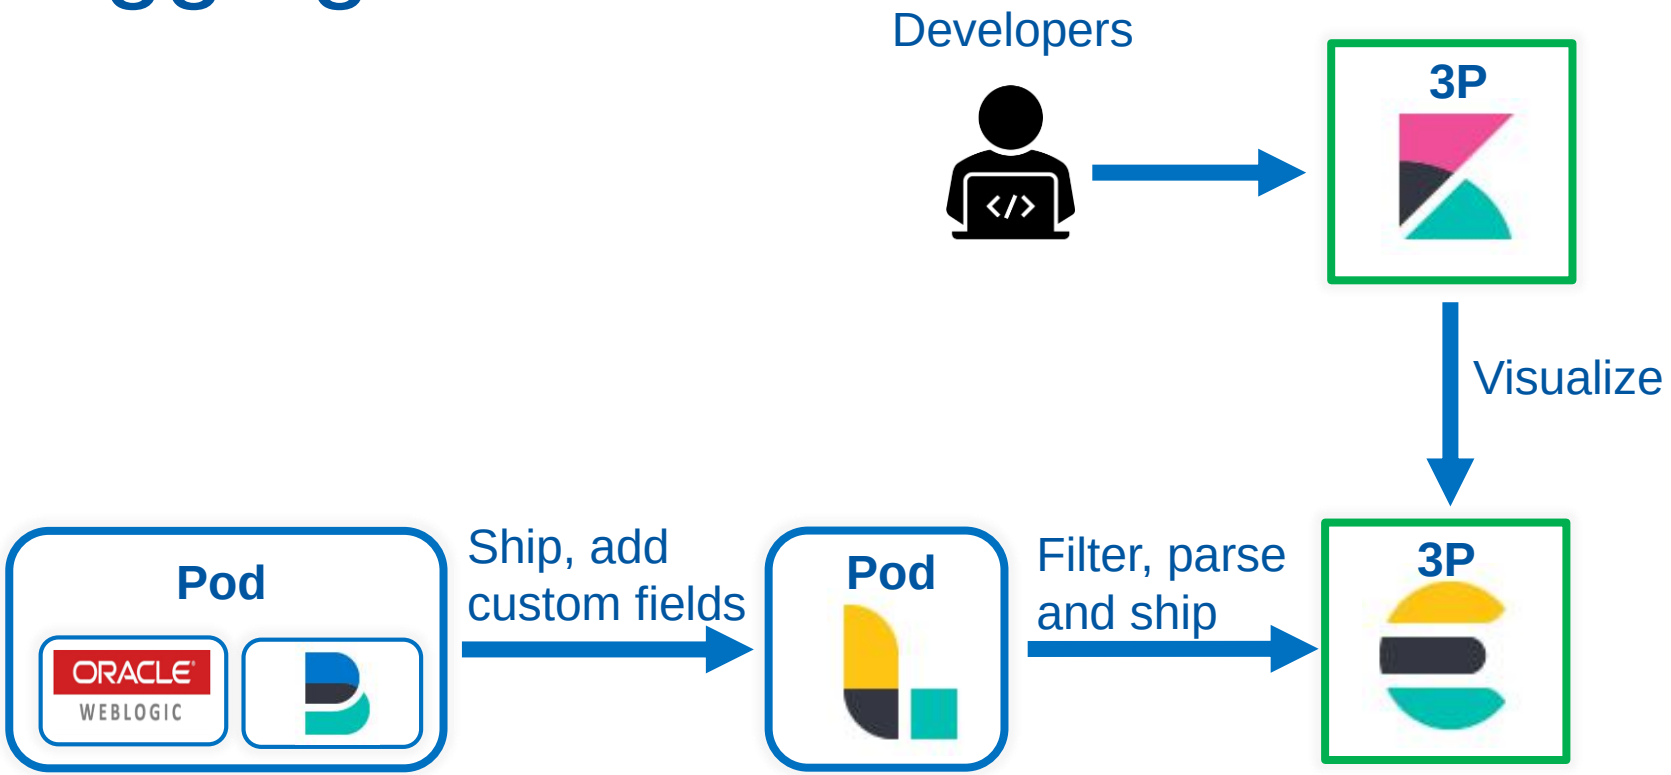
Developers
3P
Visualize
Pod
ORACLE
WEBLOGIC
Ship, add
custom fields
Pod
Filter, parse
and ship
3P

Metrics

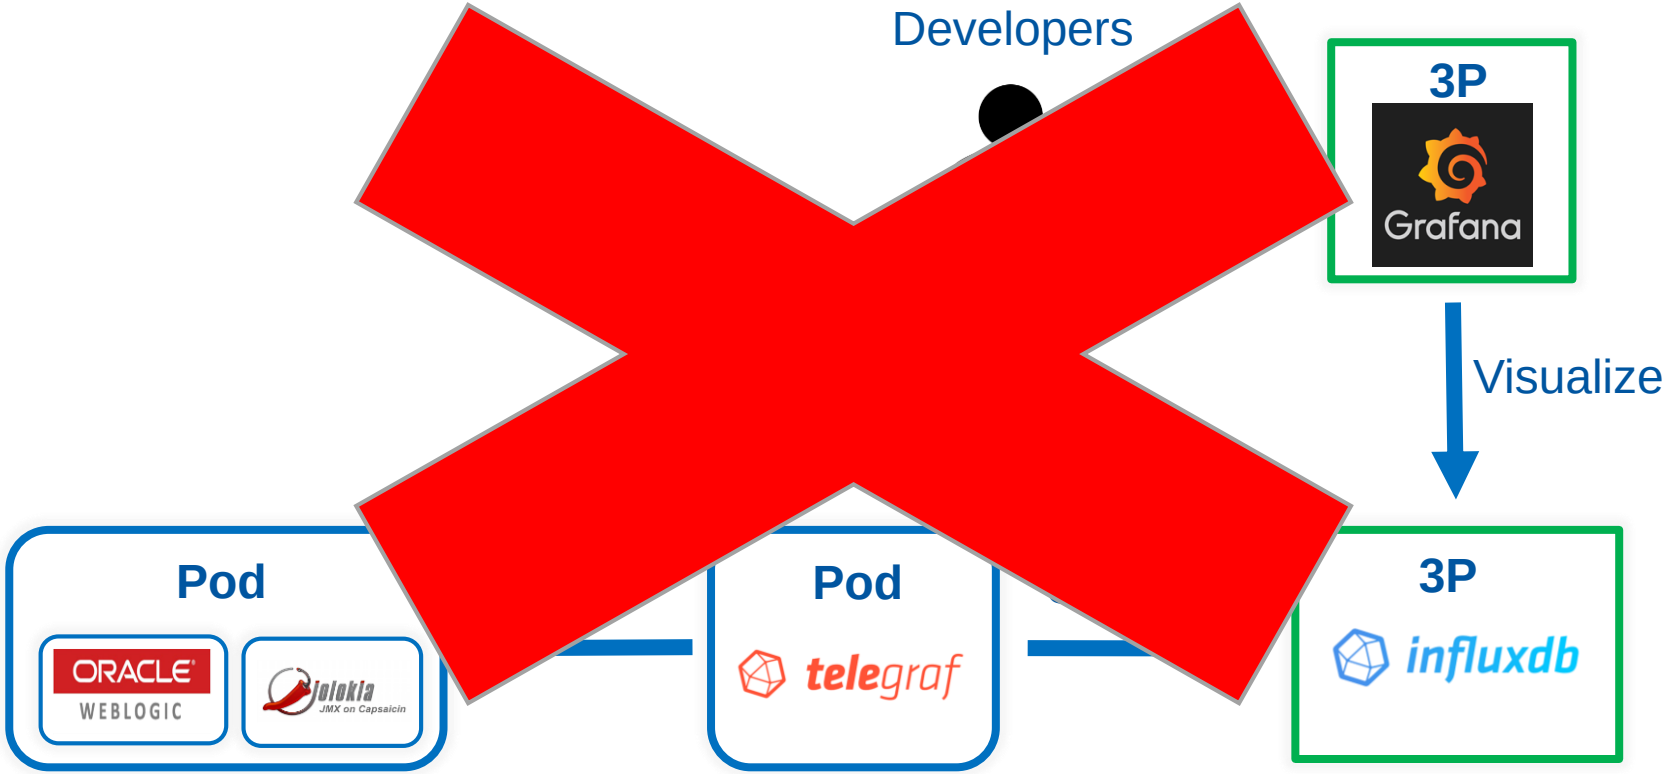
Developers
3P
Grafana
Visualize
Pod
ORACLE
WEBLOGIC
Jolokia
JMX on Capsaicin
Pod
telegraf
3P
influxdb

Metrics

Developers

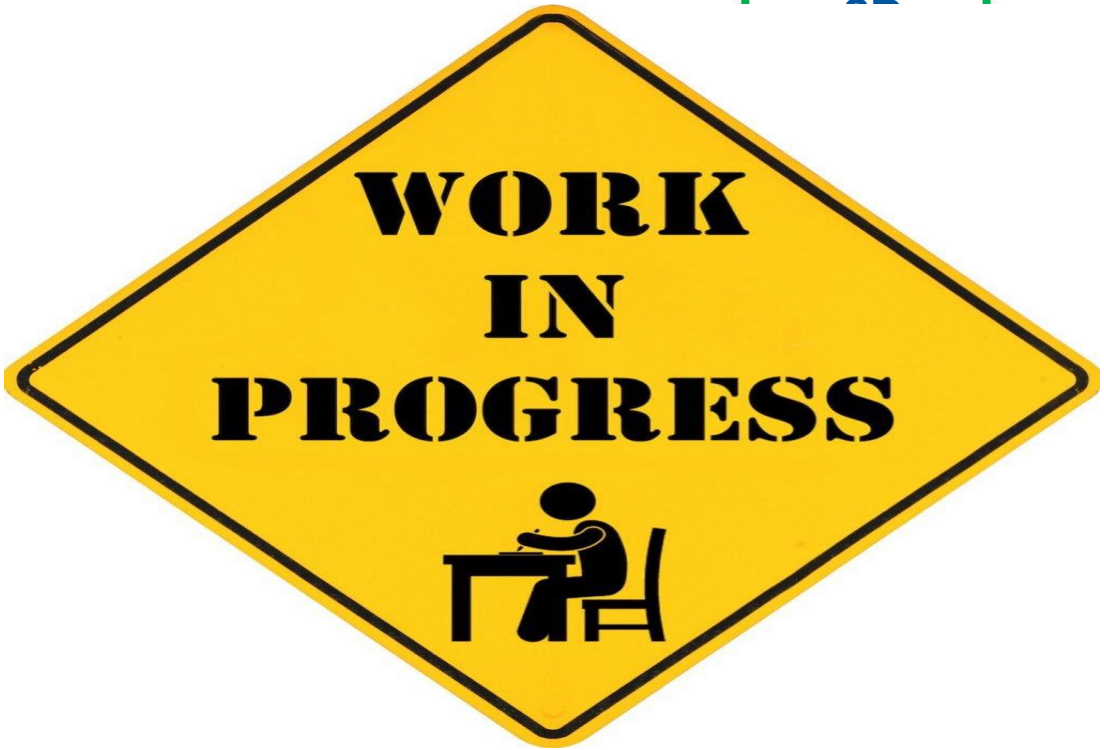

WebLogic

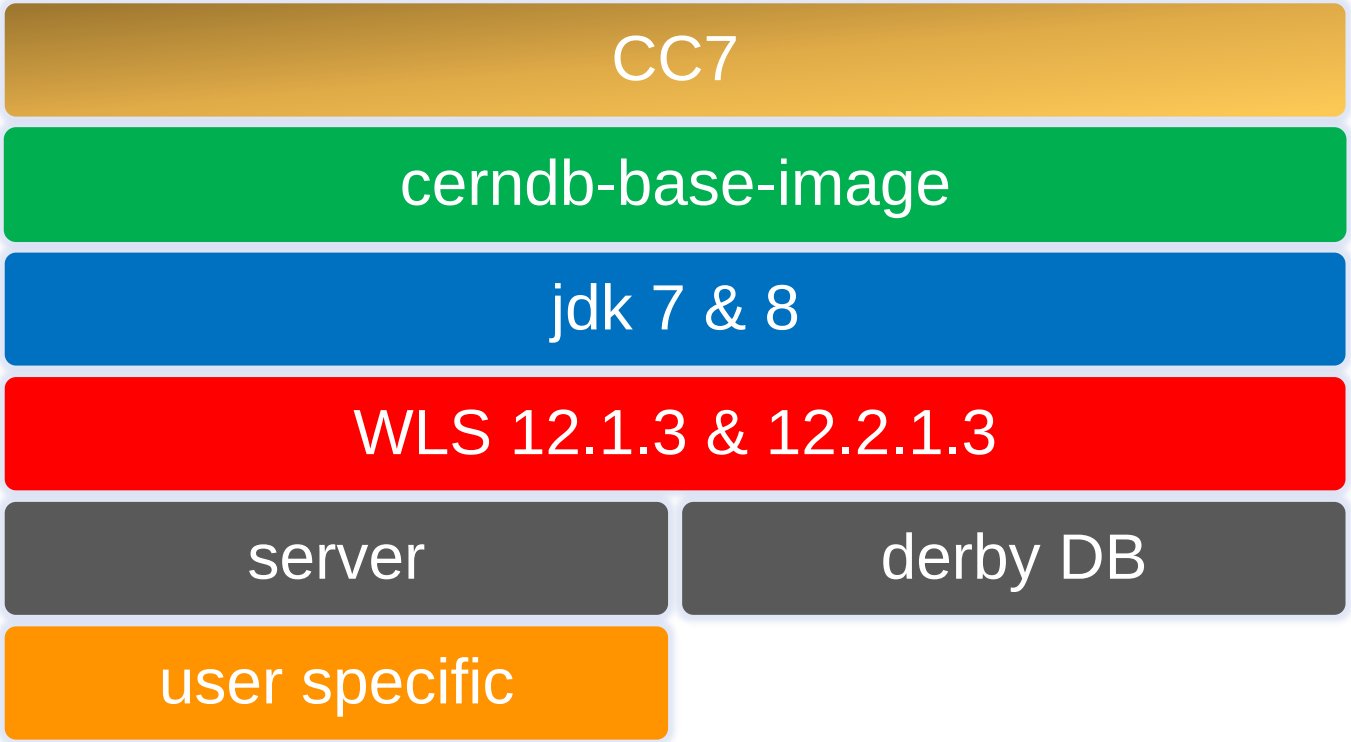
CC7
cerndb-base-image
jdk 7 & 8
WLS 12.1.3 & 12.2.1.3
server
derby DB
user specific

How we build them

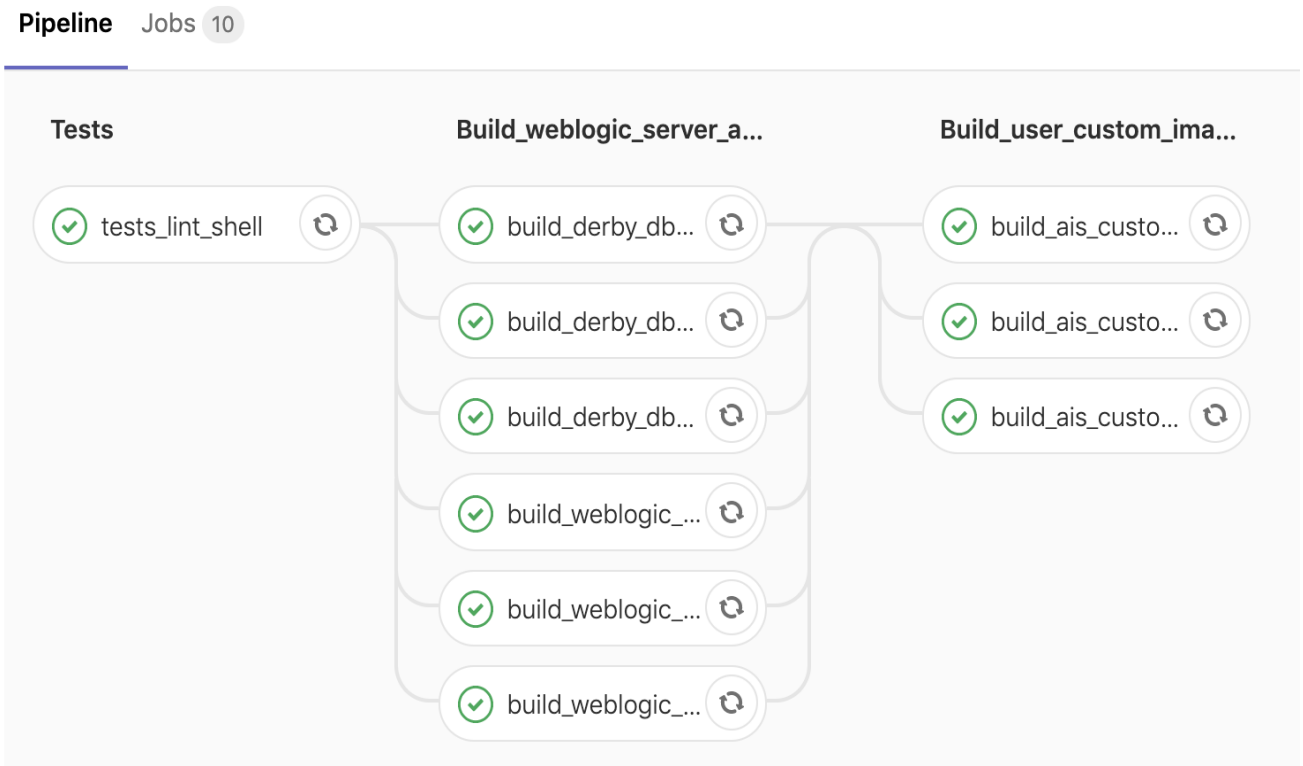

WebLogic domain configuration

WebLogic Deploy Tooling (WDT)

- Faster
  - Decreases sensitively our configuration time
- Domain config in YAML file.
- Configuration at runtime
- Works with different WebLogic versions
  - 12.1.3
    - Problem related to Single Sign On configuration
      - Bug in WLST

JSON file

- Portable
- Needs to be adapted to WDT

Secret Management

The application servers need:

- Passwords
- Certificates
- Etc. Etc.

How inject in Kubernetes

- Kubernetes secrets
  - File
  - Env variables

Secret Management

Problem:

- How avoid file or env variable to be accessible

What we do

- Unset env variables

Other solutions

- Java Security Manager
  - Use whitelist
  - With a trick, you can invert the behavior
- init container
  - Be careful with the restart
  - Health check should be reliable

Secret Management

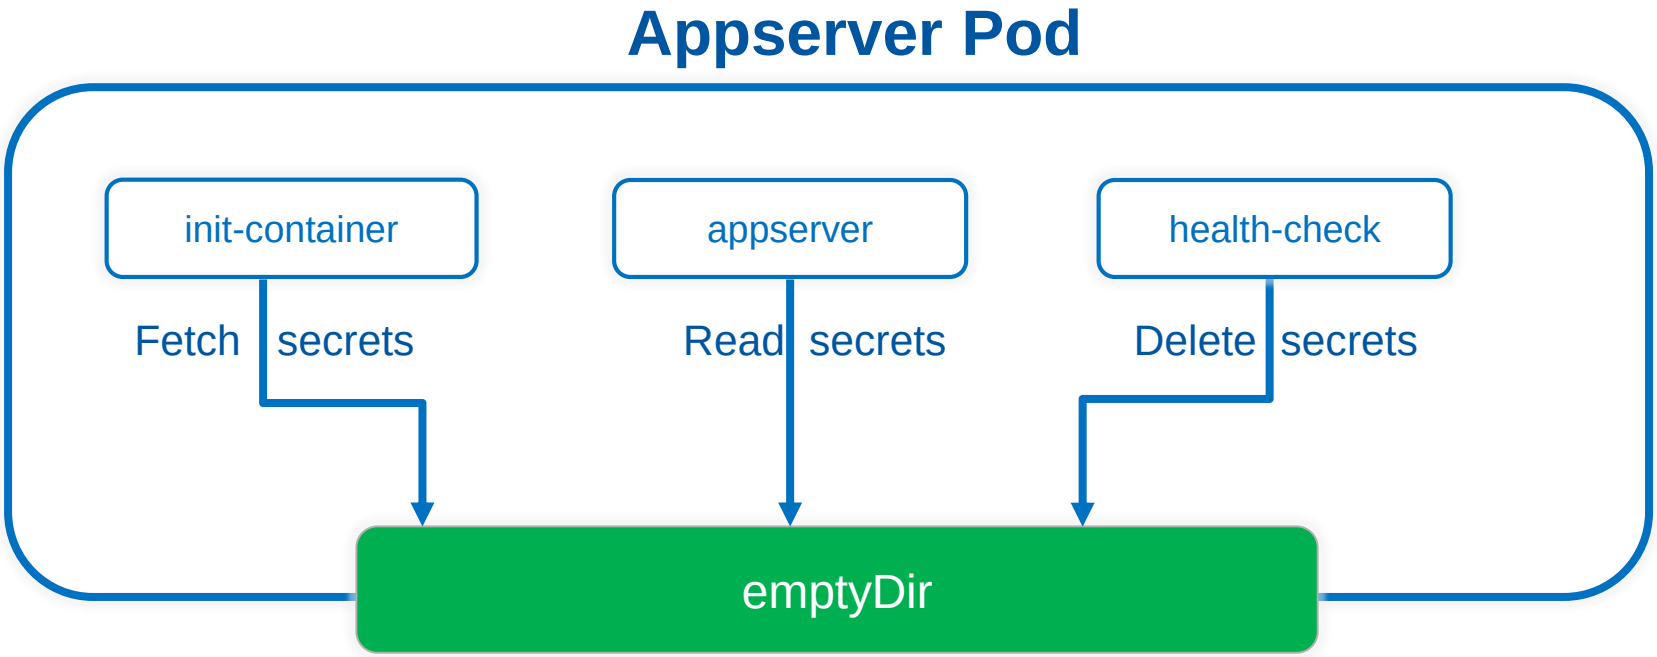

Deployment

Us

- Deployment script
  - Written in python
  - Fetch input and fill value files use by helm charts
  - Install Helm charts
- Built image with the deployment script
  - Run everywhere

Developers

- private Rundeck instance
  - container base on deployment image
- REST API call

Deployment

Developers

Accept: application/json’
Content-Type: application/json
X-Rundeck-Auth-Token: ${RUNDECK_USER_TOKEN}"

HTTP

RED HAT OPENSHIFT

RUNDECK

```
"options": {
      "JEEDY_APP_ENTITY_NAME":"CHOSEN_ENTITY",
      "JEEDY_ADM_ACTION":"CHOSEN_ACTION”,
      "JEEDY_DEPLOYMENT_TAG":"CHOSEN_TAG",
      "REFRESH_JSON":"{true,false}",
      "NOTIFICATIONS_EXTRA_RECIPIENTS":"CHOSEN_ADDRESS@cern.ch"
  }
```

Deployment

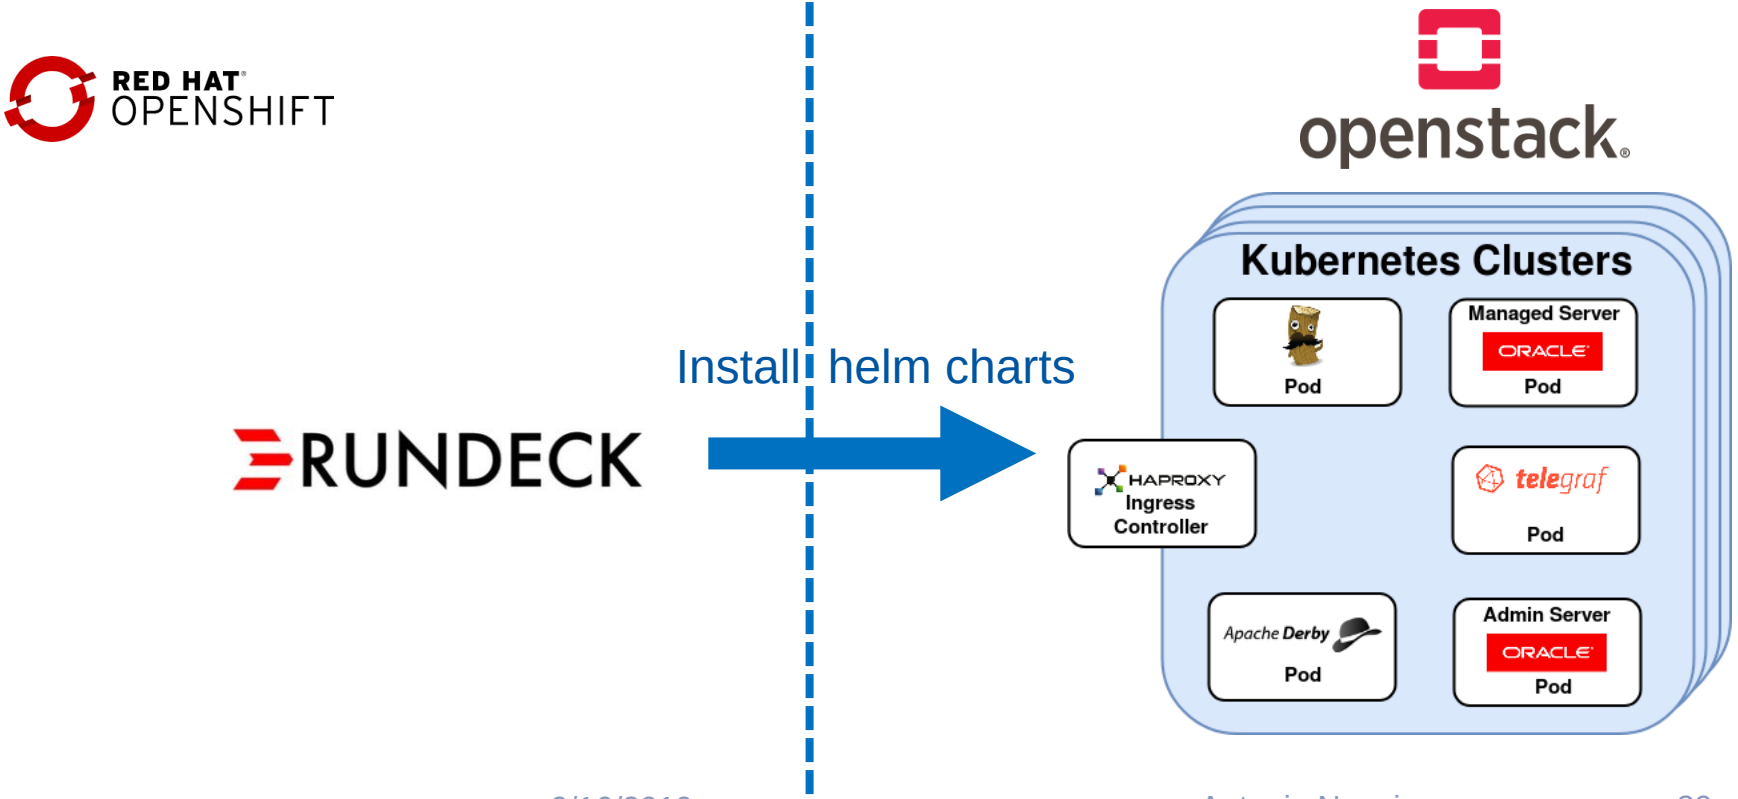
RED HAT OPENSHIFT
openstack.
RUNDECK
Install helm charts
Kubernetes Clusters
Pod
Managed Server
ORACLE
Pod
HAPROXY
Ingress
Controller
telegraf
Pod
Apache Derby
Pod
Admin Server
ORACLE
Pod

Security

Use trusted images or build our owns

Use tools to dynamically/statically check vulnerabilities

Run containers as non root user

Automatize

Helps to avoid errors

Demo

Next Steps

Complete the migration of environments to Kubernetes (ongoing)

EOL WLS 12.1.3 in November 2019

- 12.2.1.3 WebLogic only in Kubernetes

Disaster Recovery on the Cloud (OCI)

- Container Engine for Kubernetes

Delegation of actions

- Allow users to interact with K8S clusters

Prometheus

- WebLogic Exporter

Define upgrading procedures for Kubernetes Cluster

Move forward Tomcat on Kubernetes

Lessons learned

Production environments can run on containers

- Kubernetes standard de facto
  - Huge support from community

Increased portability and dynamicity of the system

- Open new doors
- Let us focus on developers needs
- No vendor lock-in

Infrastructure flexibility

- Each piece can be changed with one of same logic

All that glitters is not gold

- Everyday new tools
  - Not jump too fast on them

Images credits

- Slide 3
  - https://www.identitainsorgenti.com/a-santantonio-napoli-celebra-i-pizzaioli-patrimonio-mondiale-con-una-giornata-di-festa
- Slide 10
  - https://plus.maths.org/content/round-peg-square-hole-or-square-peg-round-hole
- Slide 11
  - https://www.sccpre.cat/show/TiTbiiT_lets-say-you-wish-to-deploy-and-manage/
  - https://giphy.com/gifs/colbertlateshow-magic-stephen-colbert-late-show-l46CnlK71u7CazMUE
- Slide 15
  - https://blog.converse.ai/the-secret-chat-weapon-every-marketer-needs-to-have-c4b41e4e7771
- Slide 23
  - https://giphy.com/gifs/aumanimation-8BkrxepXlKaWMbt5pK
- Slide 30
  - https://xebialabs.com/technology/rundeck/

THANK YOU

QUESTIONS?

antonio.nappi@cern.ch