From a cluster to the Cloud

Jean-Frederic Clere
Tomcat
@jfclere

APACHECON Europe

Oct. 22-24, 2019

Agenda

- Who I am
- A cluster:
  - Session replication and application.
- The cloud:
  - Looking at the different cloud providers
  - KUBEPing and DNSPing
  - Modify the tomcat cluster
    - Allow a dynamic list of nodes
    - Only TCP. (8888 port exported via deployment.yml)
  - Operator and S2I
  - Demos
- What next? Questions / Suggestions

Who am I?

Jean-Frederic Clere

- Red Hat
- Years writing JAVA code and server software
- Tomcat committer since 2001
- Doing OpenSource since 1999
- Cyclist/Runner etc
- Lived 15 years in Spain (Barcelona)
- Now in Neuchâtel (CH)

Session replication in a cluster

- HTTP/1.1
  - No transaction
  - No persistent connection
- Web App:
  - Using cookies to carry session ID
  - Store information in the session:
    - Shopping cart etc.
- Multi nodes and dynamic
  - Route request to right node
  - Replicate information

A cluster

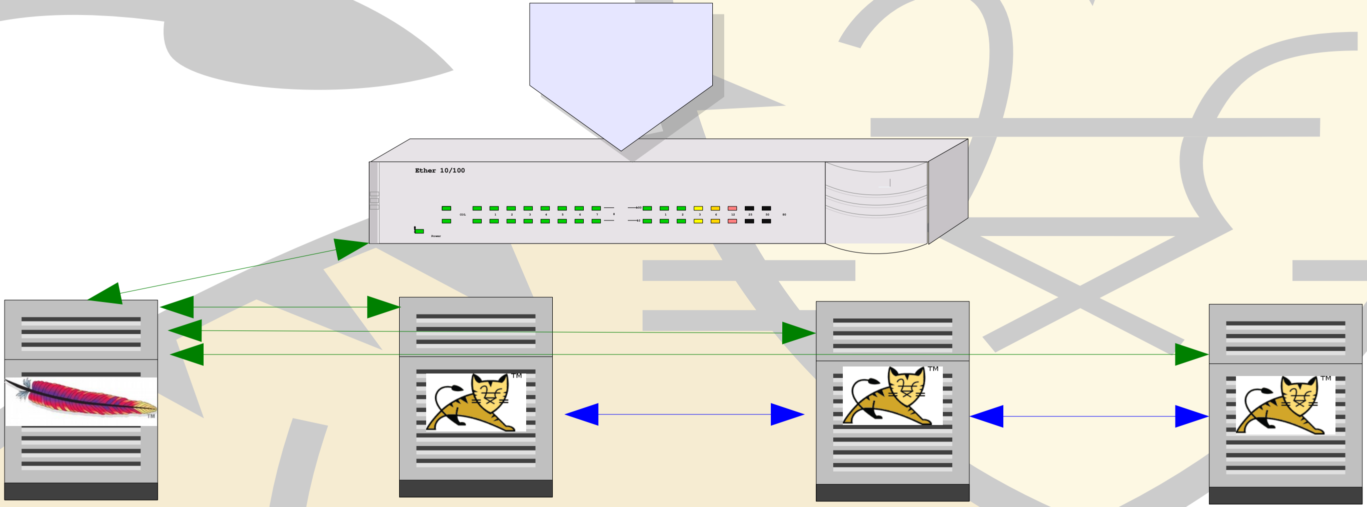

Load balancer
HTTPD
mod_proxy

Tomcat node

Tomcat node

Tomcat node

How to replicate sessions

In cluster:

<distributable/> in web.xml

<Cluster className="org.apache.catalina.ha.tcp.SimpleTcpCluster"/>

Port upd 45564

Ports tcp range 4000:4100

Kubernetes

Kubernetes is an open-source system for automating deployment, scaling, and management of containerized applications. ***https://kubernetes.io/***

Cloud providers

Most of the major cloud providers rely on Kubernetes as a container management solution.

Cloud providers

We worked on adding support for Kubernetes so that our solution would be available on all of these providers.

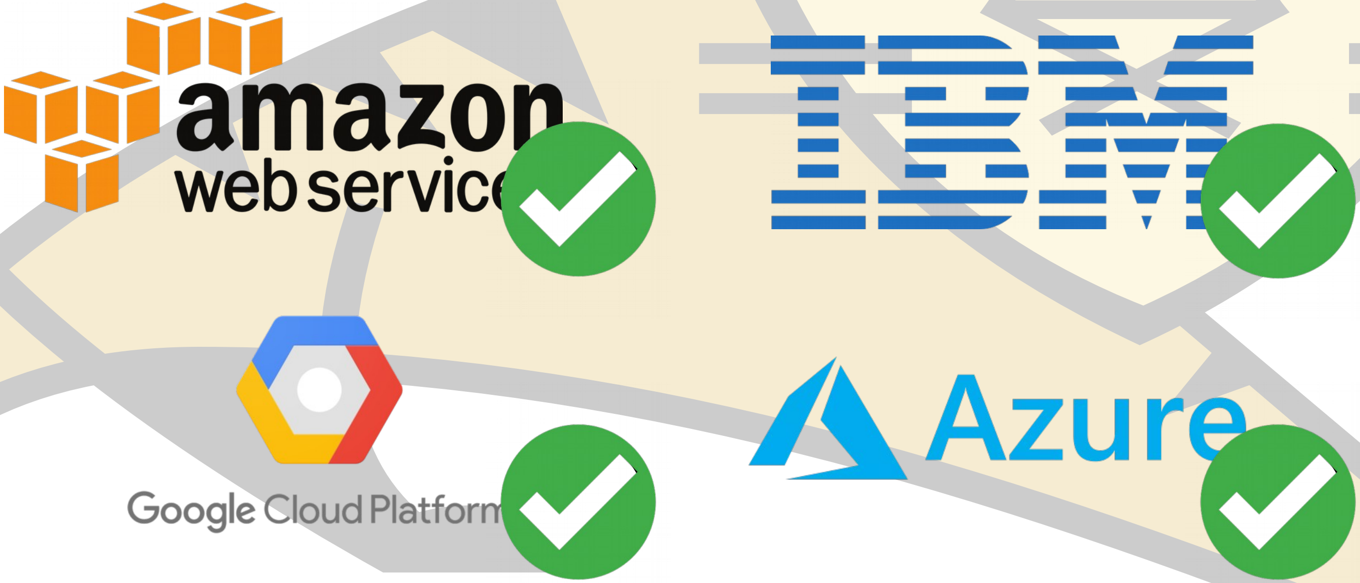

Deployment

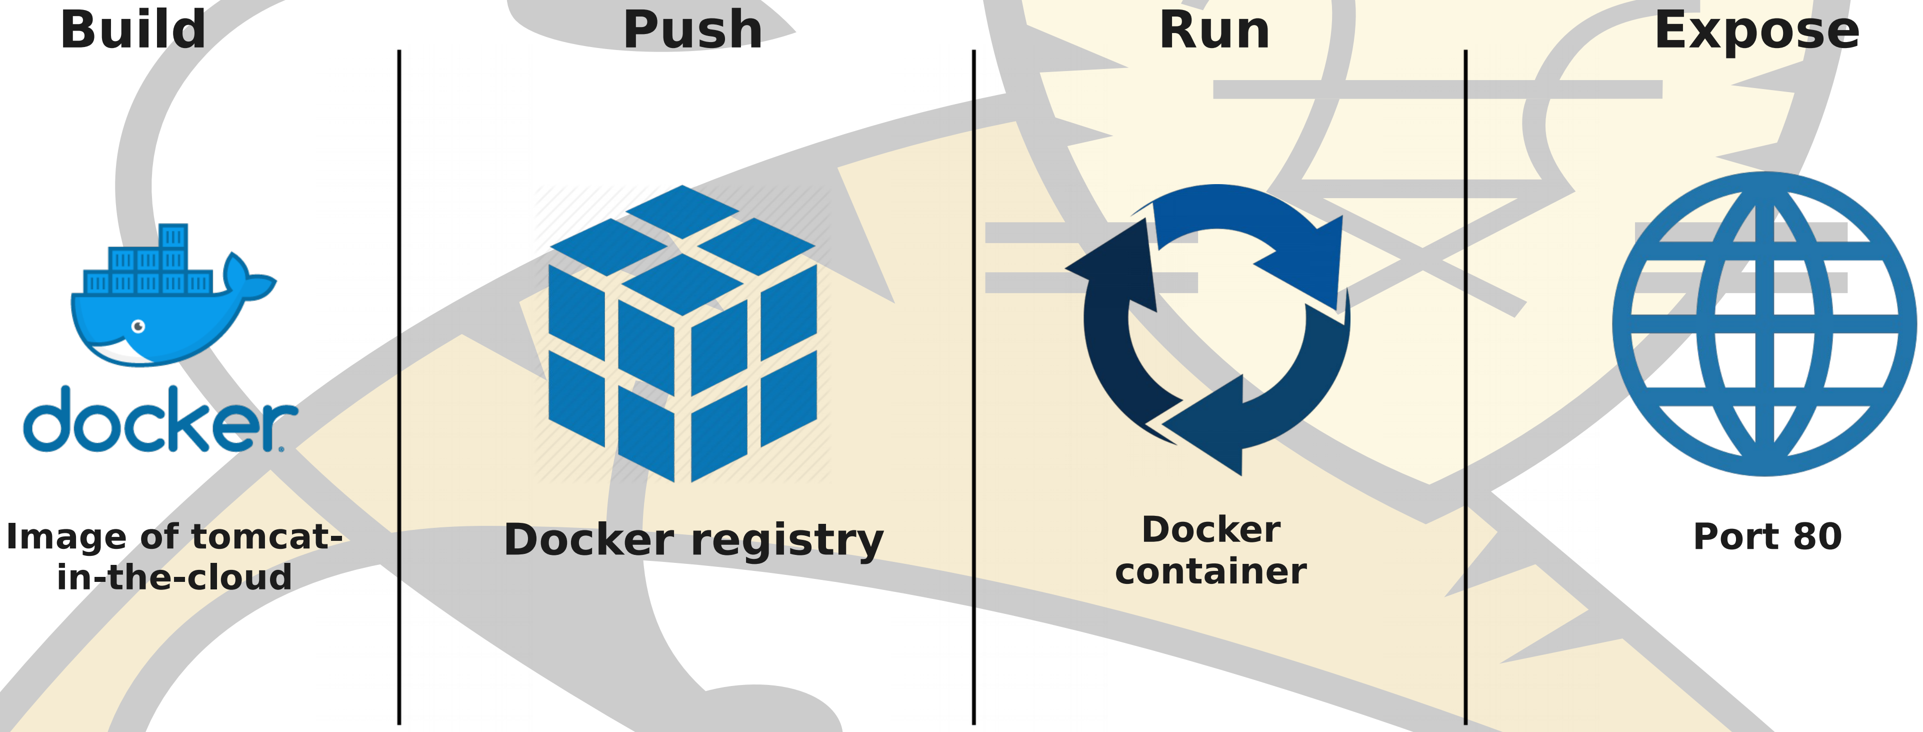

Build

Image of tomcat-in-the-cloud

Push

Docker registry

Run

Docker container

Expose

Port 80

Automation

Because the deployment can be time consuming and slightly different for each of the cloud providers (in terms of permission management). We’re currently working on automating the process preparing an operator.

AWS:

- **awscli /IAM console / docker / kops / kubectl**

Azure:

- **azure-cli /Azure console / docker / kubectl**

Google:

- **google-cloud-sdk / google cloud console / docker / kubectl**

IBM:

- **OpenShift / kubectl or oc (origin-clients)**

OPENSHIFT

A Red Hat project / product

See OpenShift

https://www.openshift.com/

Can use AWS (public cloud) or Private on premise.

A layer on top of kubernetes to make developpers life more easy.

Tomcat in OpenShift/Kubernetes

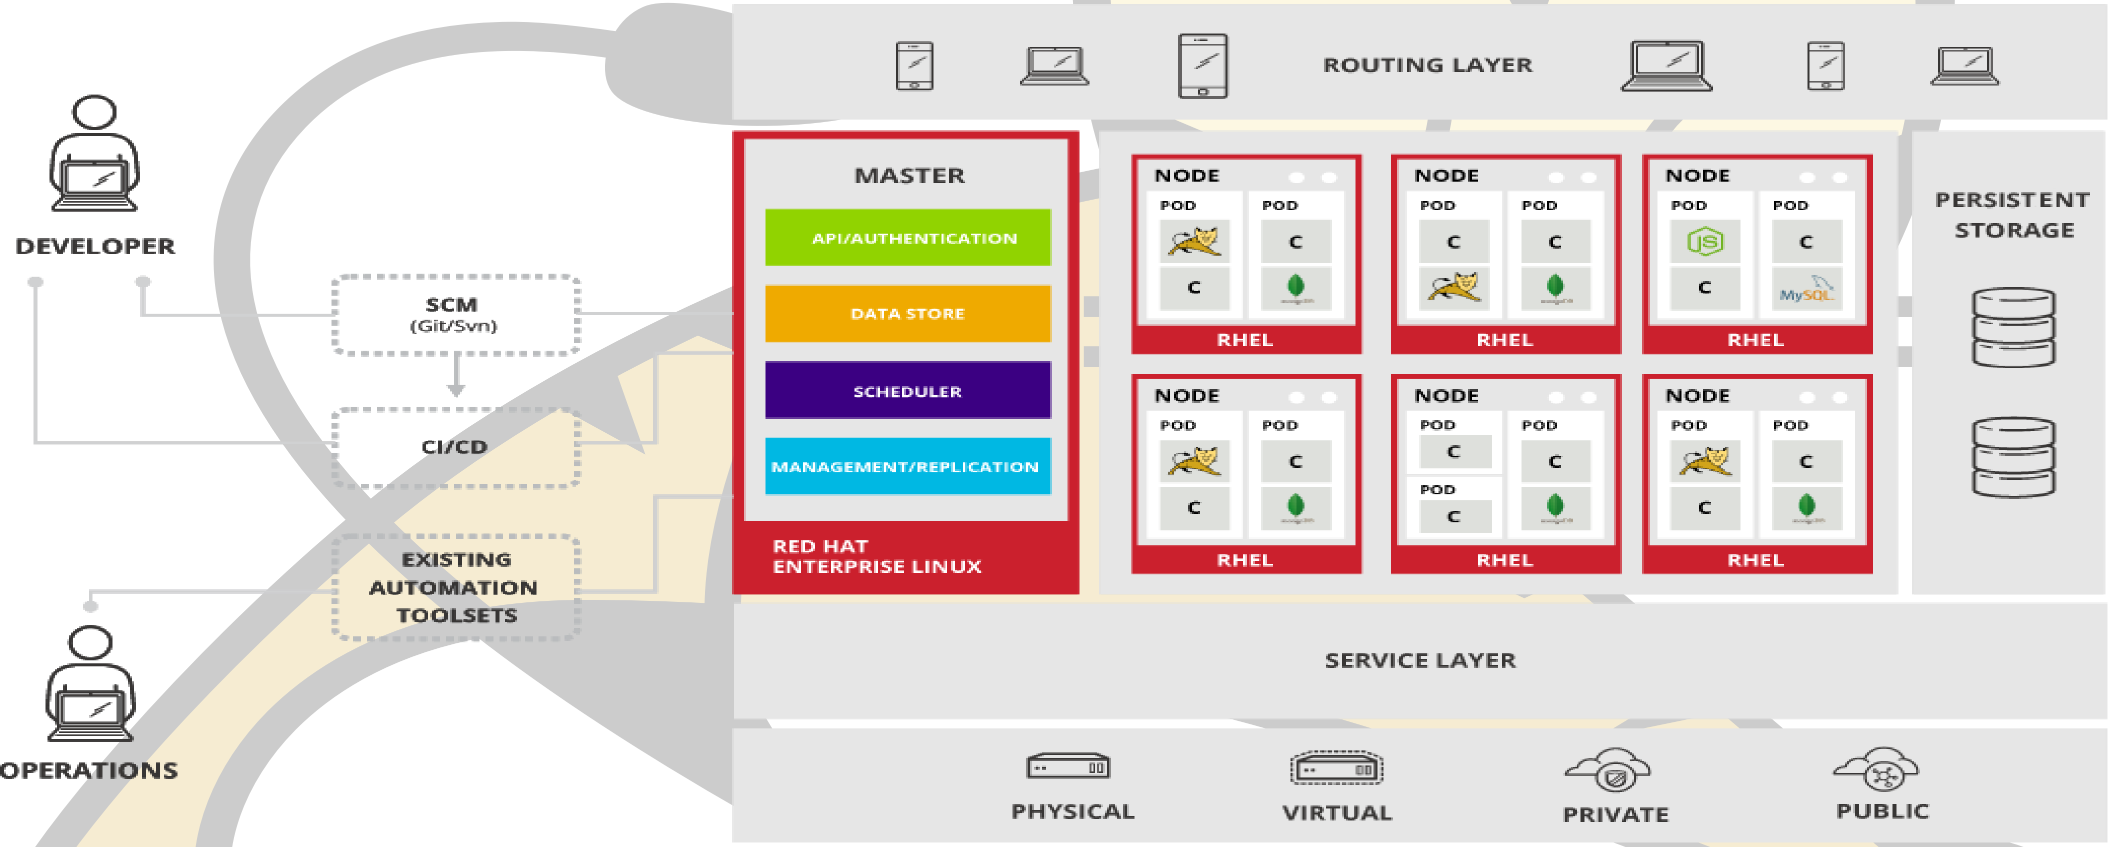

Developping Tomcat App in OpenShift/Kubernetes

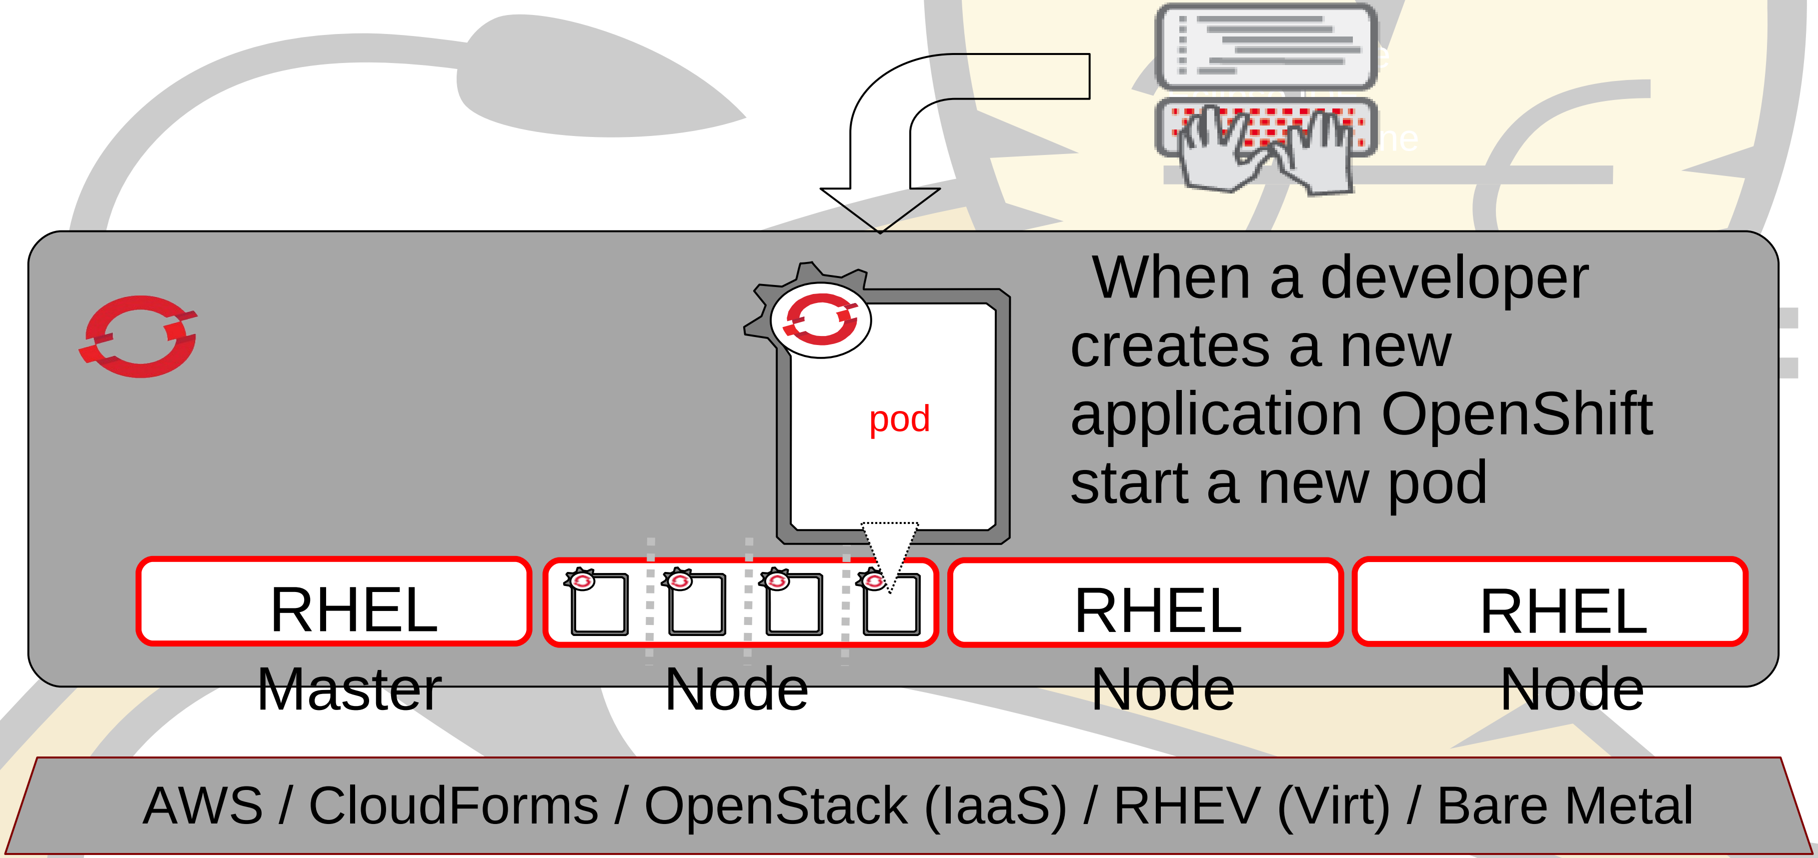

When a developer creates a new application OpenShift start a new pod

pod

RHEL

Master

Node

RHEL

Node

RHEL

Node

AWS / CloudForms / OpenStack (IaaS) / RHEV (Virt) / Bare Metal

Getting started

- minishift:
  - Allows a demo on a single box.
  - Easy to setup
  - Small demo

- Online:
  - We have prepared wiki to help you to start:
    - https://github.com/web-servers/tomcat-in-the-cloud/wiki
  - We have a katacoda tutorial:
    - https://katacoda.com/jfclere/courses/tomcat-in-the-cloud

- Bare metal / VM:
  - Use ansible to install (for openshift)
  - 2 nodes + master + infra minimal
- Tomcat webapp with sessions
  - Rest Counter demo.

From a cluster to the Cloud

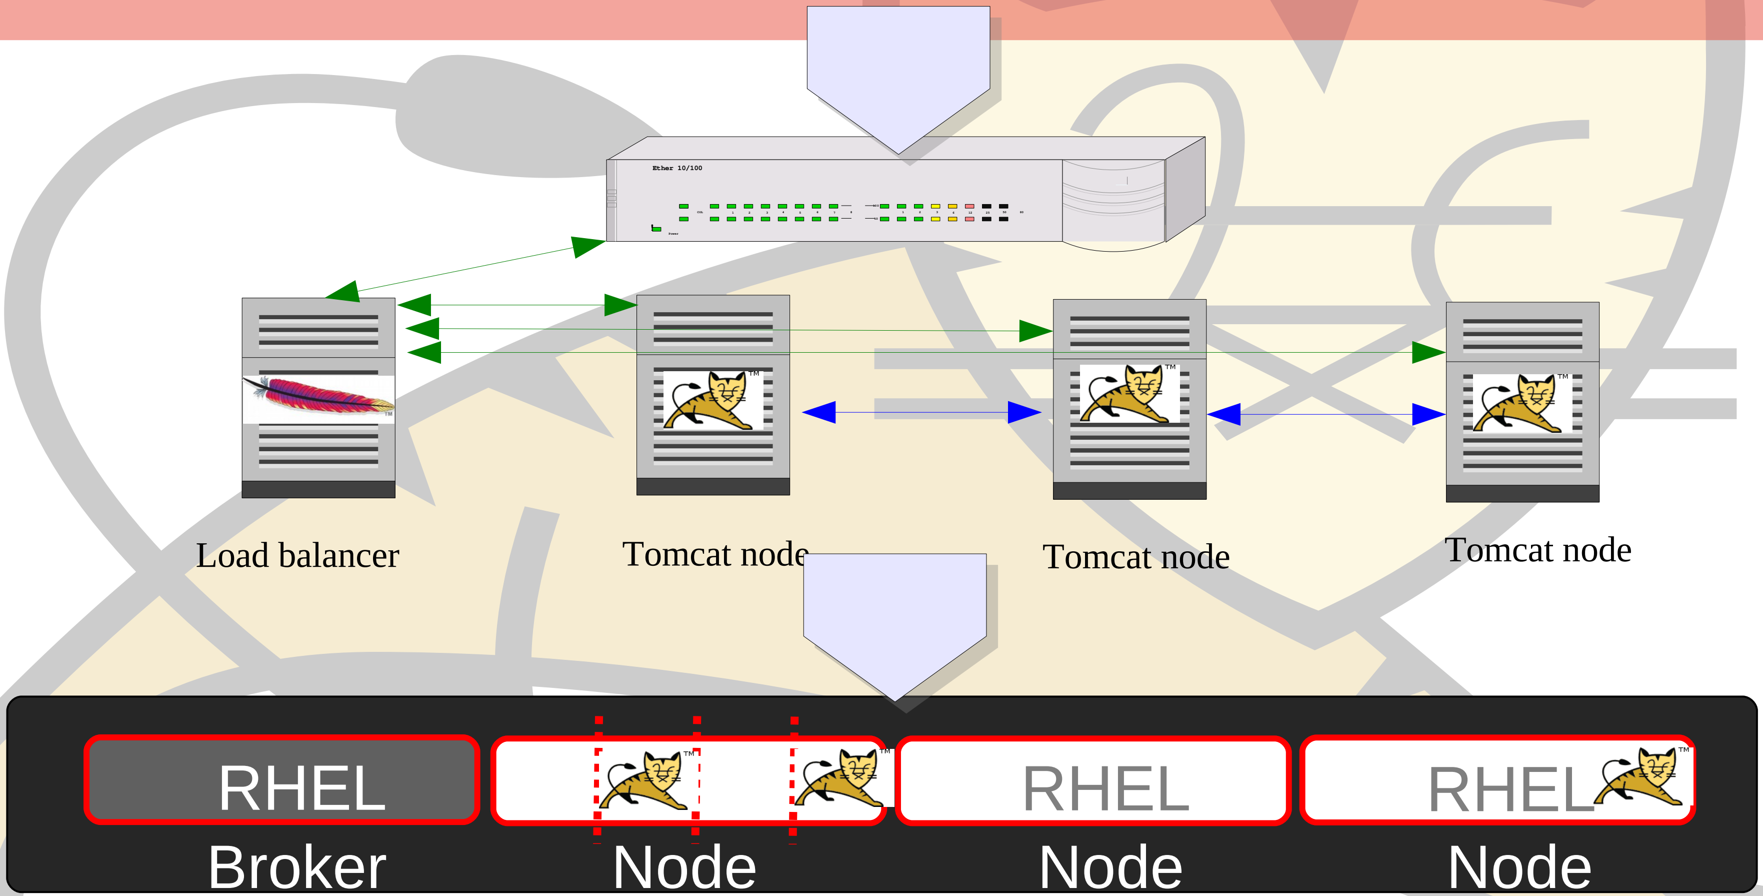

Problems for a cluster to cloud...

- Many ways to solve:
  - Embed tomcat with SpringBoot
  - Create a docker image
  - Extend an existing docker image
  - Fabric8 / S2I
- Tomcat session replication:
  - No multicast in the cloud.
  - Need a “ping” to find the other nodes (KubePing/DNSPing)
  - KubePing needs “view nodes” permission to the system account of the project.
  - DNSping uses DNS service in Kubernetes.

Solutions: KUBEPing

Tomcat cluster built-in solution
Peer discovery through multicast heartbeat messages
Does not work in a cloud environment

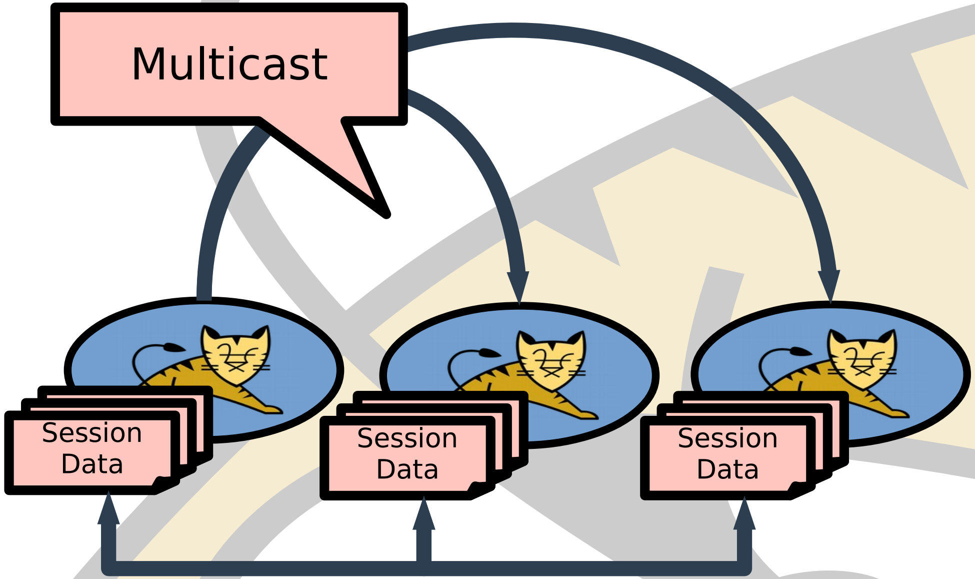

solution
Peer discovery through Kubernetes Downward API
Works in all kubernetes clouds

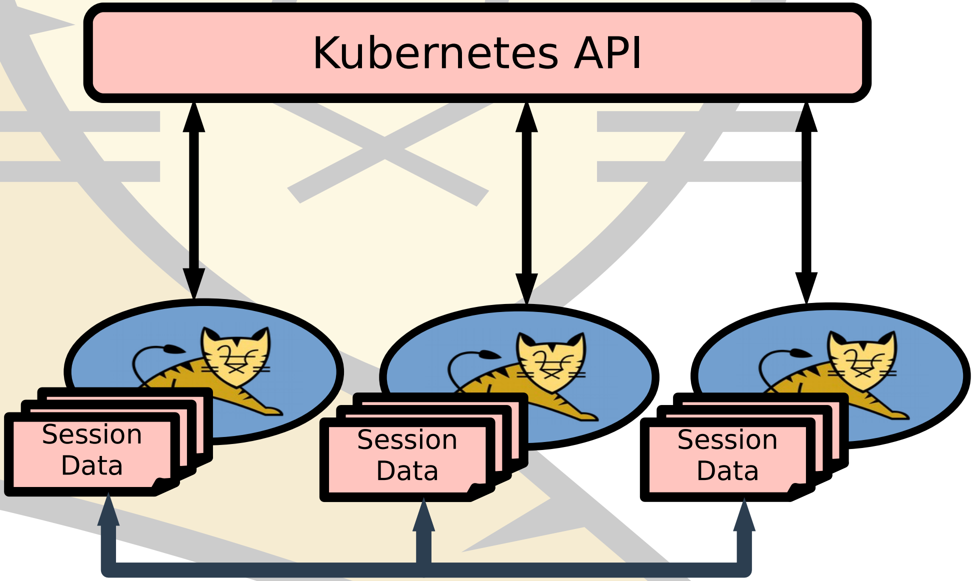

Kubernetes API

Tools for managing a Kubernetes cluster

Accessible from the pods within the cluster

GET /api/v1/namespaces/tomcat-in-the-cloud/pods

➔ Return a JSON representation of all the pods in the cluster

➔Requires permissions

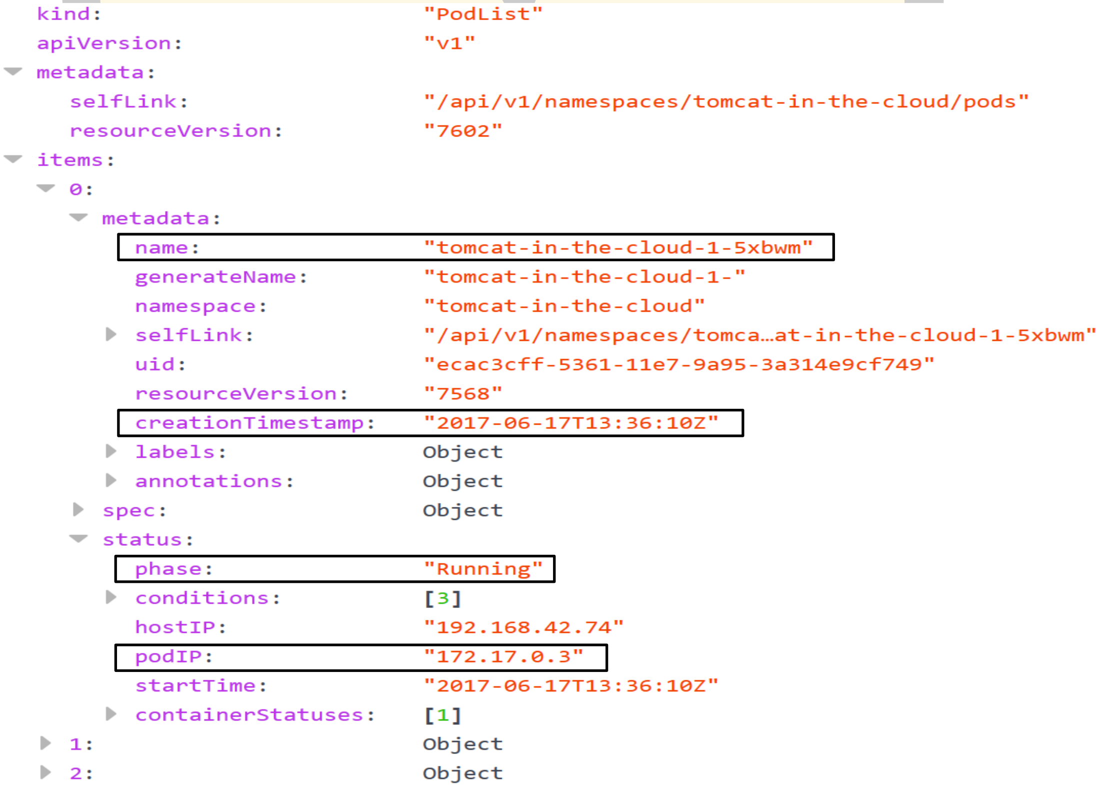

```
kind:                       "PodList"
apiVersion:                 "v1"
▾ metadata:
    selfLink:               "/api/v1/namespaces/tomcat-in-the-cloud/pods"
    resourceVersion:        "7602"
▾ items:
  ▾ 0:
    ▾ metadata:
        name:               "tomcat-in-the-cloud-1-5xbwm"
        generateName:       "tomcat-in-the-cloud-1-"
        namespace:          "tomcat-in-the-cloud"
      ▸ selfLink:           "/api/v1/namespaces/tomca…at-in-the-cloud-1-5xbwm"
        uid:                "ecac3cff-5361-11e7-9a95-3a314e9cf749"
        resourceVersion:    "7568"
        creationTimestamp:  "2017-06-17T13:36:10Z"
      ▸ labels:             Object
      ▸ annotations:        Object
    ▸ spec:                 Object
    ▾ status:
        phase:              "Running"
      ▸ conditions:         [3]
        hostIP:             "192.168.42.74"
        podIP:              "172.17.0.3"
        startTime:          "2017-06-17T13:36:10Z"
      ▸ containerStatuses:  [1]
  ▸ 1:                      Object
  ▸ 2:                      Object
```

Architecture KUBEPing case

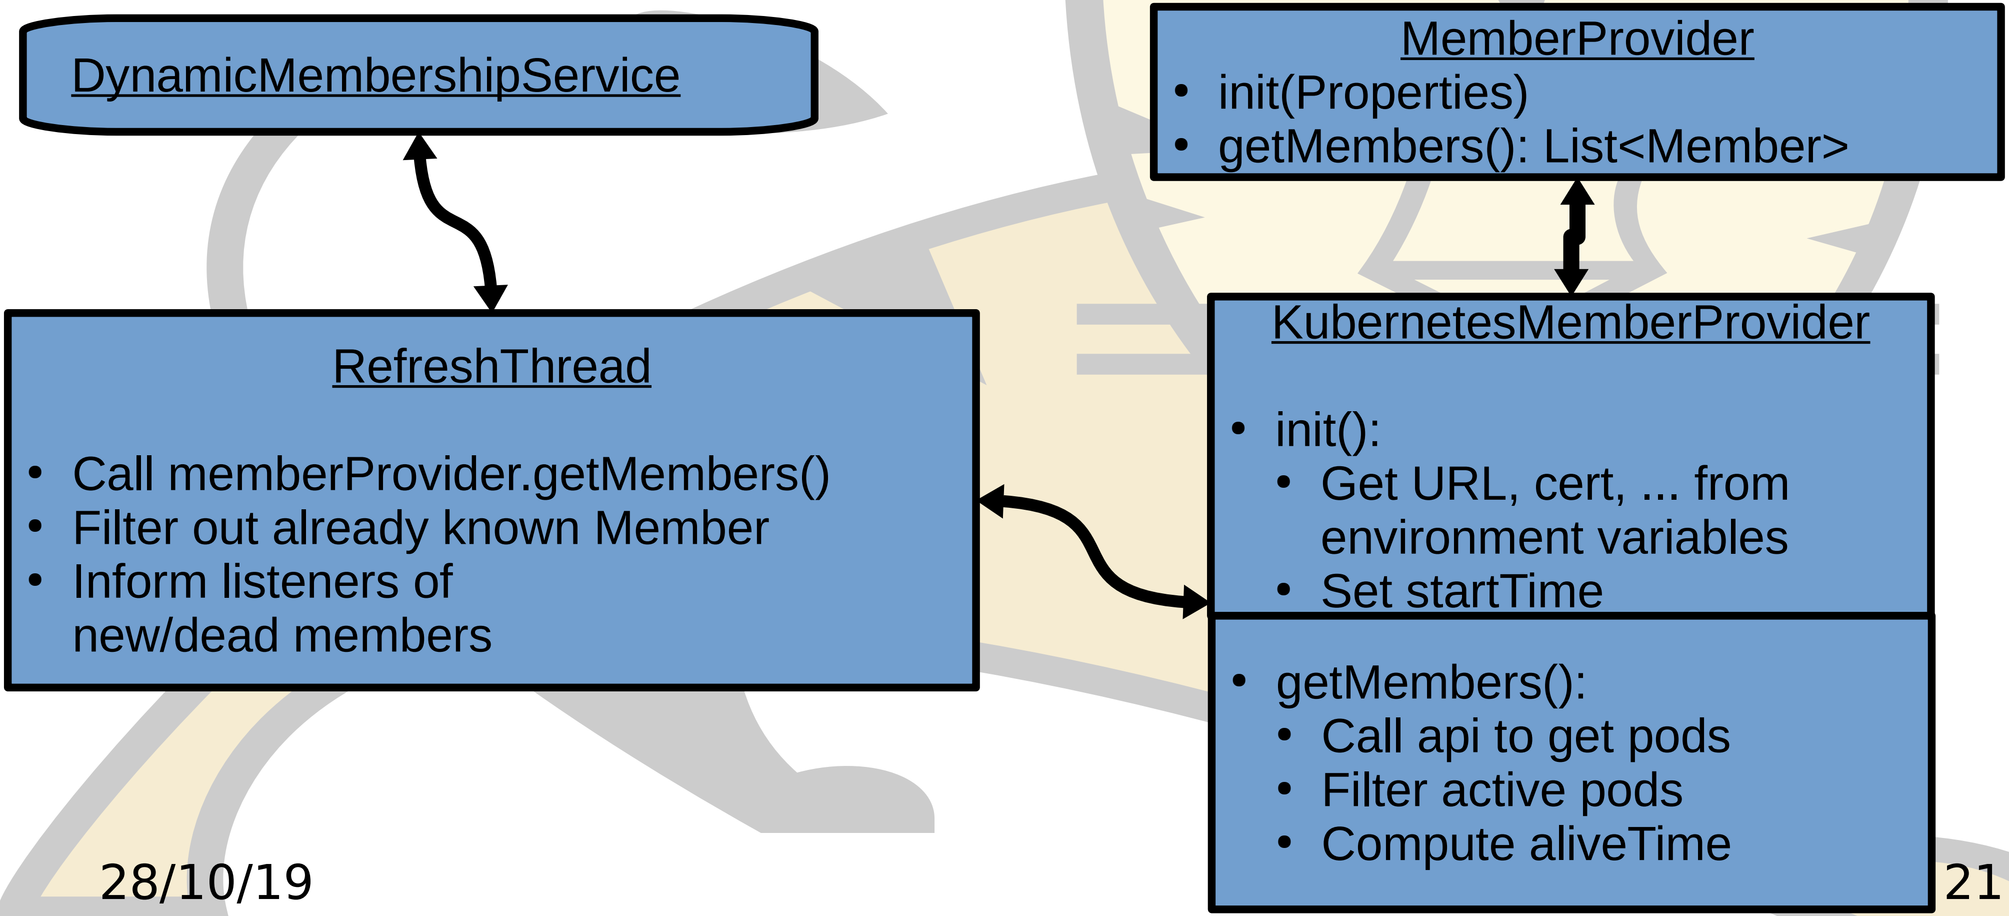

DynamicMembershipService

MemberProvider

- init(Properties)
- getMembers(): List<Member>

RefreshThread

- Call memberProvider.getMembers()
- Filter out already known Member
- Inform listeners of new/dead members

KubernetesMemberProvider

- init():
  - Get URL, cert, ... from environment variables
  - Set startTime
- getMembers():
  - Call api to get pods
  - Filter active pods
  - Compute aliveTime

Solutions: DNSPing

Tomcat cluster built-in solution
Peer discovery through multicast heartbeat messages
Does not work in a cloud environment

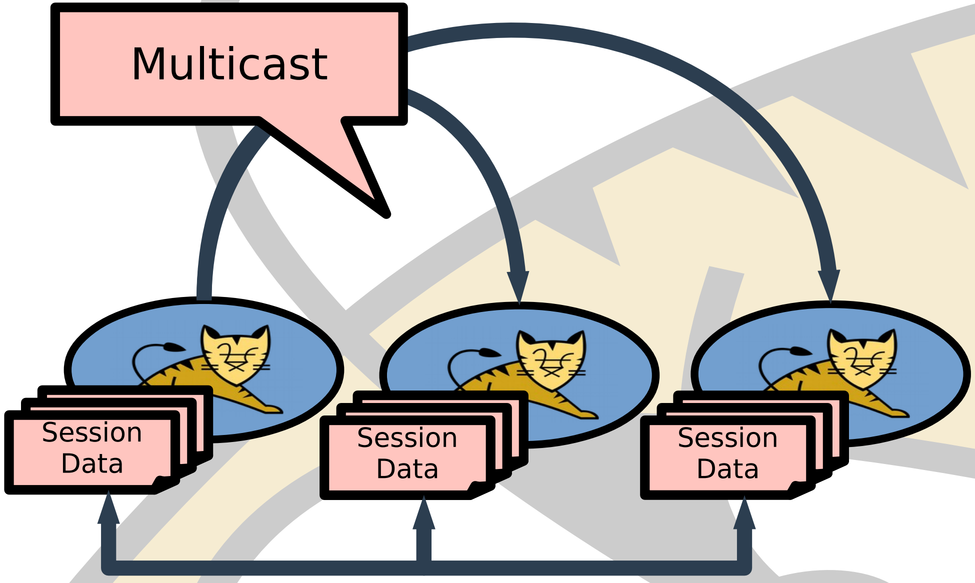

solution
Peer discovery through DNS lookup
Works in all kubernetes clouds

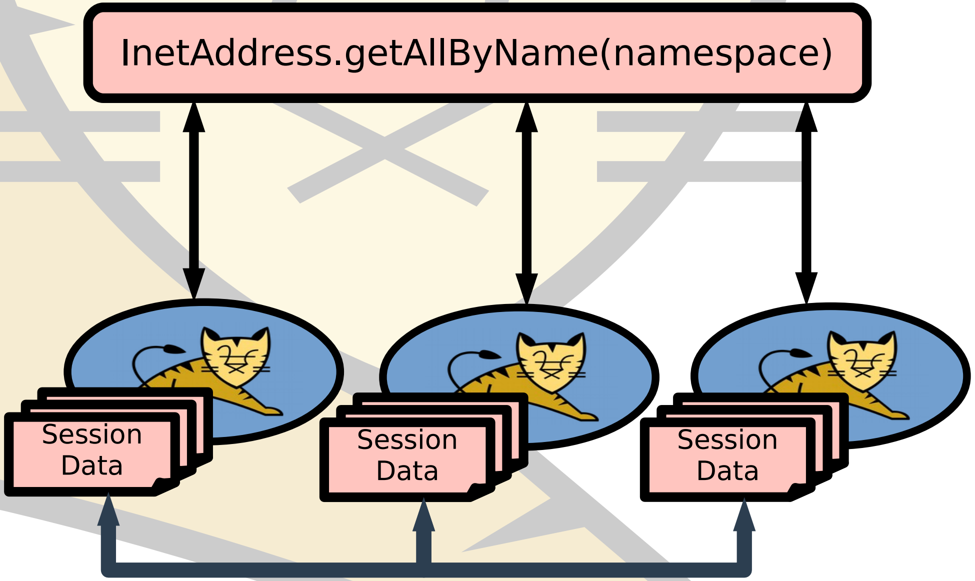

`InetAddress.getAllByName(namespace)`

DNS lookup

nslookup name-space

Accessible from the pods within the cluster

InetAddress.getAllByName()

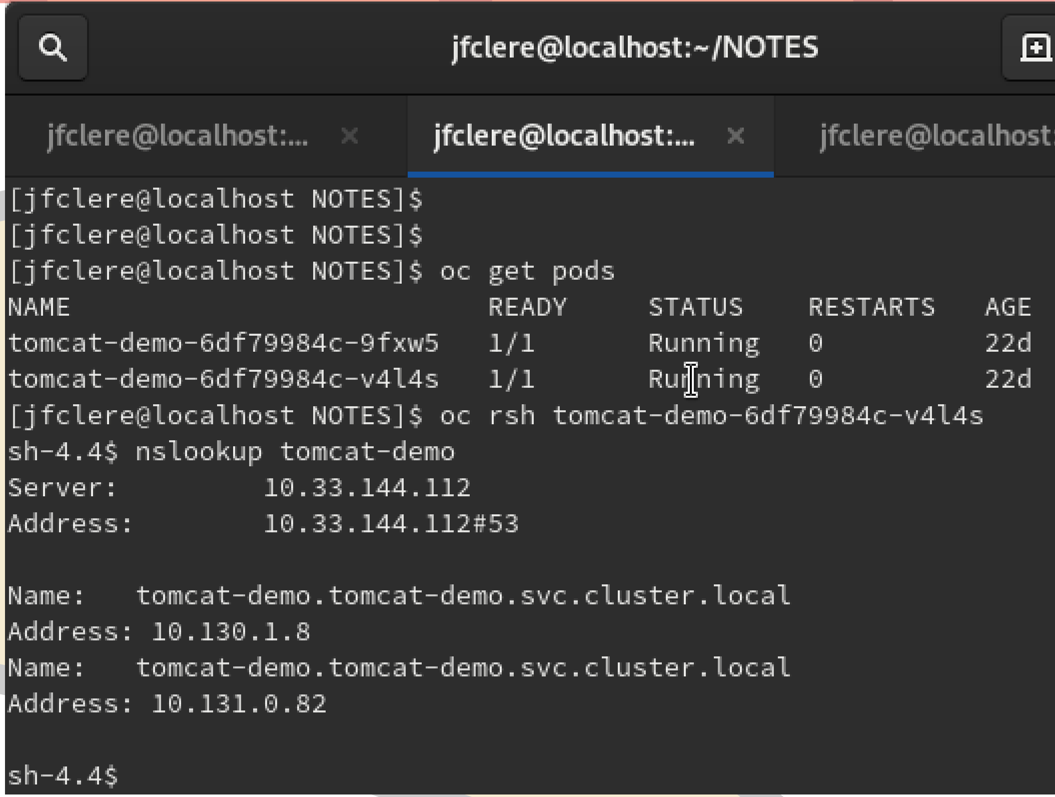

```
[jfclere@localhost NOTES]$
[jfclere@localhost NOTES]$
[jfclere@localhost NOTES]$ oc get pods
NAME                          READY     STATUS    RESTARTS   AGE
tomcat-demo-6df79984c-9fxw5   1/1       Running   0          22d
tomcat-demo-6df79984c-v4l4s   1/1       Running   0          22d
[jfclere@localhost NOTES]$ oc rsh tomcat-demo-6df79984c-v4l4s
sh-4.4$ nslookup tomcat-demo
Server:         10.33.144.112
Address:        10.33.144.112#53

Name:   tomcat-demo.tomcat-demo.svc.cluster.local
Address: 10.130.1.8
Name:   tomcat-demo.tomcat-demo.svc.cluster.local
Address: 10.131.0.82

sh-4.4$
```

What is done

Demos contents:

- **We use Katacoda for the demos.**

Where is the code:

- **In Tomcat tomcat-maven for the Docker Image**
- **In Tomcat for the DNS/KUBEPing cluster code**
- **In github for the https://github.com/web-servers/tomcat-s2i**
- **In github for the https://github.com/web-servers/tomcat-operator**
- **More documentation / tests are welcome**

Katacoda demo using KUBEPing

https://katacoda.com/jfclere/courses/tomcat-in-the-cloud

And the sources:

https://github.com/jfclere/intro-katacoda/tree/master/tomcat-in-the-cloud/deploy-titc-using-cli

Requires permission to read pods configuration (use it in private cloud)

Katacoda demo using DNSPing

https://katacoda.com/jfclere/scenarios/dnsping-tomcat

And the sources:

https://github.com/jfclere/intro-katacoda/tree/master/DNSPing-tomcat

Runs everywhere, but requires a service for DNS discovering.

Operator

- What is a Kubernetes operator
  - kubernetes definition
  - Basically it automates the services, routes and build (S2I) process.
- What do we have now
  - We have one written in GO (prototype)
  - S2I (source to image) just tooling (PR # 188) vetoed but doesn’t need to be in Tomcat.

Katacoda demo using operator

Operator demo in Katacoda

And the sources:

https://github.com/jfclere/intro-katacoda/tree/master/war-katacoda

Every thing is created by the operator: pods, services etc.

Where we are

Main sites:

https://github.com/web-servers/tomcat-in-the-cloud

https://github.com/jfclere/tomcatPI

https://docs.openshift.com

https://github.com/apache/tomcat

tomcat : res/tomcat-maven

DNSMembershipProvider / KubernetesMembershipProvider

Tomcat operator and S2I PR#188

Questions ?
Suggestions?

THANK YOU

JEAN-FREDERIC CLERE

@jfclere

jfclere@gmail.com