



KubeCon

North America 2017

decco: Secure, Multitenant Cluster for Internet Facing Services

Bich Le, Chief Architect @Platform9

Background & Motivation



deploy, manage & maintain



as a service

on the infrastructure of your choice



Objectives

- Run control plane services on shared Kubernetes cluster(s)
- Secure enough to host multiple customers
- Simplify Kubernetes without hiding its power and API
- Useful to community

Deployment Cluster Config & Ops

- Public repo: <https://github.com/platform9/decco>
- Follows the Kubernetes controller/operator pattern
- Runs inside or outside cluster
- Automates provisioning and teardown of
 - DNS records
 - K8s resources: namespaces, deployments, services, ingress, netpolicies
- Introduces 2 custom resources: Space, App

DNS conventions for service endpoints

For hypothetical customer “foo”

- Customer-wide fqdn: `foo.platform9.net`
- Region-specific fqdn: `foo-regionname.platform9.net`
- Services using reverse proxy (HTTP only):
`foo-regionname.platform9.net/serviceName`
- Services not using reverse proxy (TCP and HTTP):
`serviceName.foo-regionname.platform9.net`

Network security

General observations about control plane services:

- TLS encryption for all traffic
- HTTP services
 - don't use client certificate authentication
 - use app-level authentication (Keystone)
- TCP services
 - use client certificate auth
 - optionally use app-level auth (e.g. rabbitmq / mysql password)

Network security (cont'd)

Internally, decco secures the network using a combo of


- TLS
 - encryption
 - mutual certificate verification
- Network Policy resources

decco concepts

- Space
- Project
- App

Space

- A naming and isolation boundary for a related set of Internet-facing services



Space – example yaml

```
1  apiVersion: "decco.platform9.com/v1beta2"  
2  kind: "Space"  
3  metadata:  
4    name: "example-space-1"  
5  spec:  
6    domainName: platform9.net  
7    httpCertSecretName: pf9-net-secret  
8    tcpCertAndCaSecretName: example-space-1  
9    encryptHttp: true  
10 project: dev
```

Space – what happens upon creation

Decco creates:

- Namespace
- DNS record
- Optional network policy resource (see *Project* slide)
- Resources needed for network routing and TLS processing

DNS update details

Decco uses same DNS library as Kubernetes Federation

1. Compute FQDN
 - Space: `${spaceName}.${domainName}`
 - TCP Service: `${svcName}.${spaceName}.${domainName}`
2. Lookup external endpoint (IP/fqdn) of cluster's ingress service
3. Using configured DNS provider
 - Lookup hosted zone corresponding to the space's domain name
 - If found, create A or CNAME record for the FQDN in the hosted zone, otherwise fail

Project

- An optional property to restrict incoming connections to those from spaces with the same project name
- Internally enforced using a Network Policy with namespaceSelector
- Tested on GKE “Alpha” 1.8.x cluster with Calico network plugin

Example

global.platform9.net

consul

project: (none)

project: foo

foo.platform9.net

mysql

keystone

project: bar

bar.platform9.net

mysql

keystone

foo-east.platform9.net

mysql

nova

foo-west.platform9.net

mysql

nova

bar-west.platform9.net

mysql

nova

App

- A thin wrapper around a PodSpec, specifying a service listening on a given port (*Planned: support for multiple service ports*)
- Service can be externally exposed via
 - reverse proxy (`httpUrlPath` not empty)
 - direct TCP (`httpUrlPath` empty)
- TCP specific options
 - `tlsCaAndCertSecretName` (use space default secret if not specified)
 - `verifyTcpClientCert`
 - `createDnsRecord`
- Other options: `initialReplicas`, `tlsEgresses`

App – http example

```
1  apiVersion: "decco.platform9.com/v1beta2"
2  kind: "App"
3  metadata:
4    name: "keystone"
5  spec:
6    initialReplicas: 1
7    httpUrlPath: "/keystone"
8    pod:
9      containers:
10     - name: keystone
11       image: monasca/keystone
12     ports:
13     - containerPort: 5000
```

App – tcp example

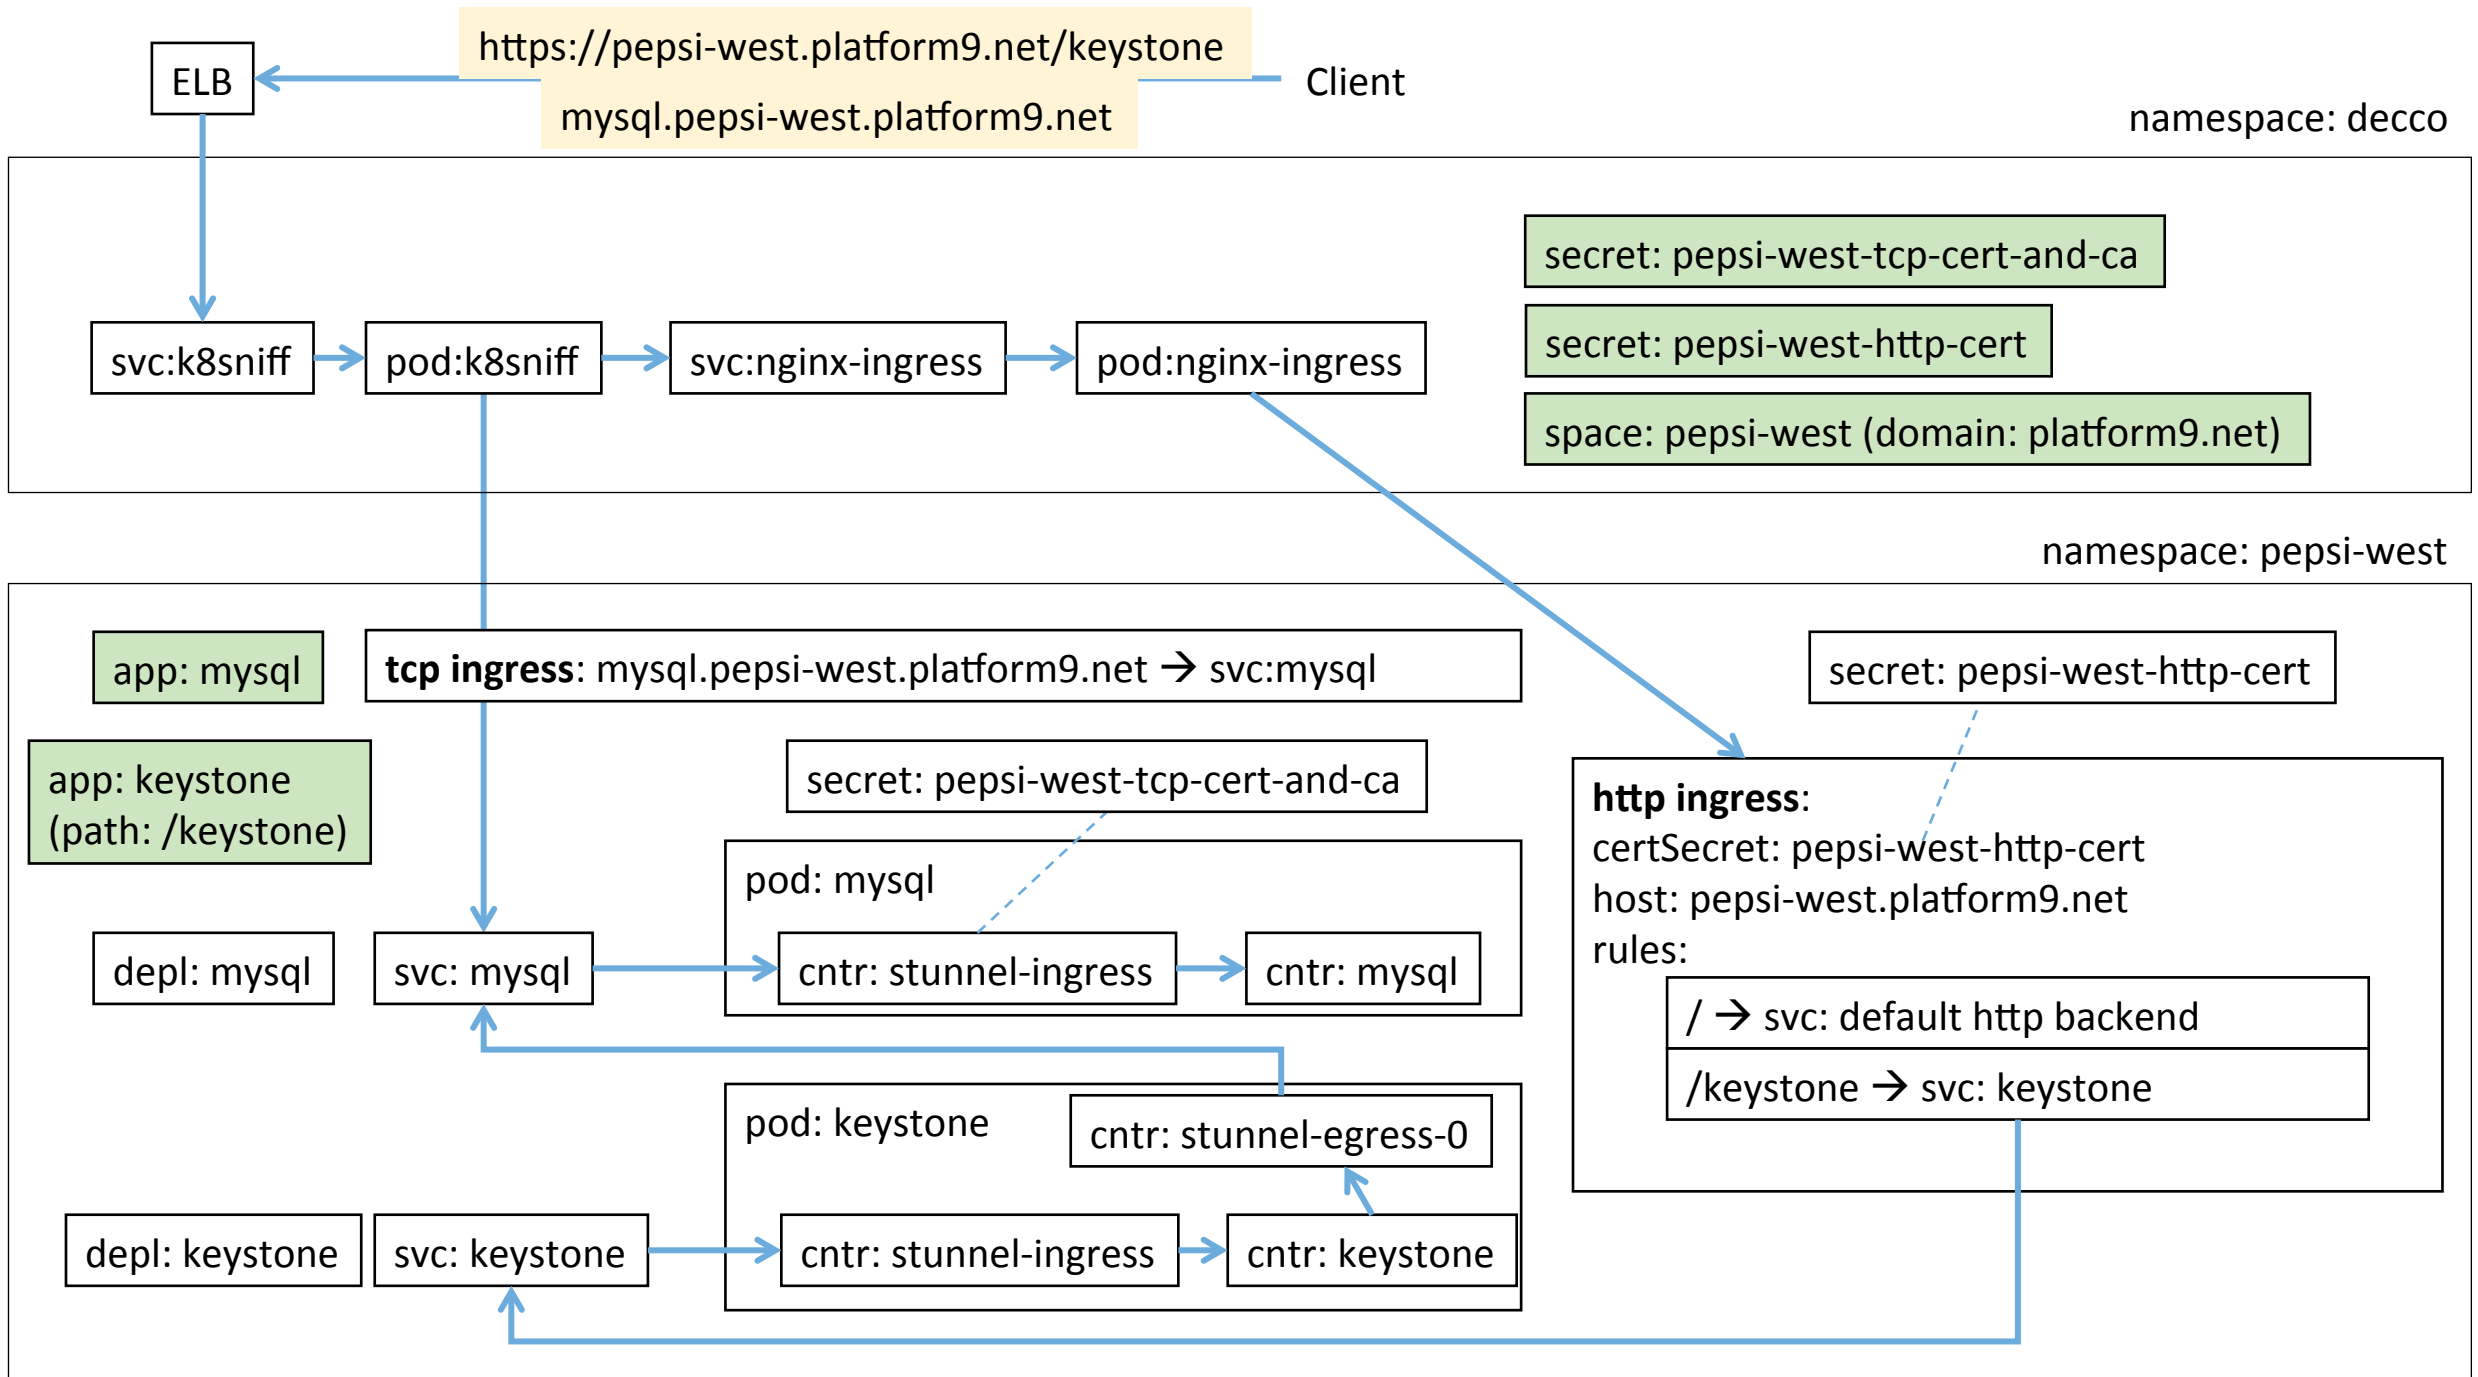
```
1  apiVersion: "decco.platform9.com/v1beta2"
2  kind: "App"
3  metadata:
4    name: "timeserver-tcp"
5  spec:
6    initialReplicas: 1
7    verifyTcpClientCert: true
8    pod:
9      containers:
10     - name: timeserver
11       image: platform9systems/current-time-standalone
12     ports:
13     - containerPort: 80
```

App – what happens at creation time

- Create service resource(s)
- Create DNS record if `createDnsRecord` is true
(example fqdn: `myservice.myspace.platform9.net`)
- Create resources for network routing and TLS
 - Ingress resources / path rules
 - Service mesh rules
- Create deployment resource with modified pod spec
 - Inject side cars needed for network routing

Network routing and TLS

- Planned: service mesh integration (e.g. Istio)
- Out of the box: cascaded ingress controllers
 - [k8sniff](#) inspects SNI headers and routes based on SNI server name
 - If no match, forwards to [nginx](#) ingress controller
 - Decco manages ingress resources for both (TCP and HTTP)
 - TLS
 - Certs need to be externally signed and supplied via secrets
 - Stunnel side cars



Roadmap

- Service mesh integration
- Integration with Availability Zones and Kubernetes Federation
- Support multiple service endpoints per app
- Open source project
 - Docs
 - Automated test
- Explore better container isolation via VM-based solutions (e.g. Frakti / runV)

Demo and Thanks!

- For more info
 - <https://github.com/platform9/decco>
 - <http://www.platform9.com>
- Other talks
 - Fission (serverless on Kubernetes)
(Thursday 2:45pm)
 - Cost-effective Compute Clusters using Spot and Preemptible Instances
(Friday 2pm)