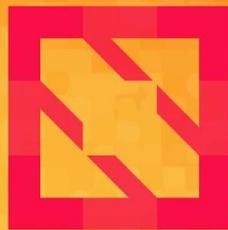




KubeCon



CloudNativeCon

North America 2019





KubeCon



CloudNativeCon

North America 2019

Solving Multi-Cluster Network connectivity with Submariner

Chris Kim <chris.kim@rancher.com>

Miguel Ángel Ajo <majopela@redhat.com>



SUBMARINER



What is Submariner?



KubeCon



CloudNativeCon

North America 2019

- Submariner is a tool that connects Kubernetes cluster networks together
- The Kubernetes Cluster Networking model provides specifications/provisions for the physical/virtual networks contained within a cluster, but leaves it up to the end user to determine how to connect multiple cluster networks together
- There are ways to do this specific to cloud providers, i.e. connect VPC's together, but this means you must be using only that providers cloud
- Submariner aims to connect the virtual networks together of multiple Kubernetes clusters through the utilization of IPsec tunnels

History of Submariner



KubeCon



CloudNativeCon

North America 2019

- Submariner was first conceptualized in 2017
- First prototypes were written in late 2018, utilizing bash scripts to automatically establish IPsec tunnels using StrongSwan
- Submariner 0.0.1 was released in March 2019, supporting IPsec for cross-cluster connection and local cluster host-gw routing
- Rancher and Red Hat (+ Community) have been collaborating on Submariner to enhance/increase reliability
- Submariner 0.0.2 (current release) was released in October 2019, supporting IPsec for cross-cluster connection and VXLAN for local cluster routing and src-ip preservation.



KubeCon



CloudNativeCon

North America 2019

Roadmap

Roadmap



KubeCon



CloudNativeCon

North America 2019

- Subctl (more about this when talking about deployment options)
- Support for overlapping IP CIDR ranges
- DNS service discovery across clusters
- NetworkPolicy support across clusters (experimental)
- Support more network plugins: ovn-kubernetes, “GKE”, etc...

Roadmap (Overlapping CIDRs)



KubeCon



CloudNativeCon

North America 2019

Handling clusters with Overlapping CIDRs

Handling overlapping CIDRs is being developed by the submariner team, the implementation is based on a global overlay CIDR which will be used for colliding clusters.

Roadmap (DNS Discovery)



KubeCon



CloudNativeCon

North America 2019

`submariner.io/lighthouse`

Lighthouse facilitates **DNS Service Discovery** in
multi-cluster connected environments

The idea is that the application does not need to know
where the service lives, it accesses local and remote
services in the same manner

Roadmap (DNS Discovery)



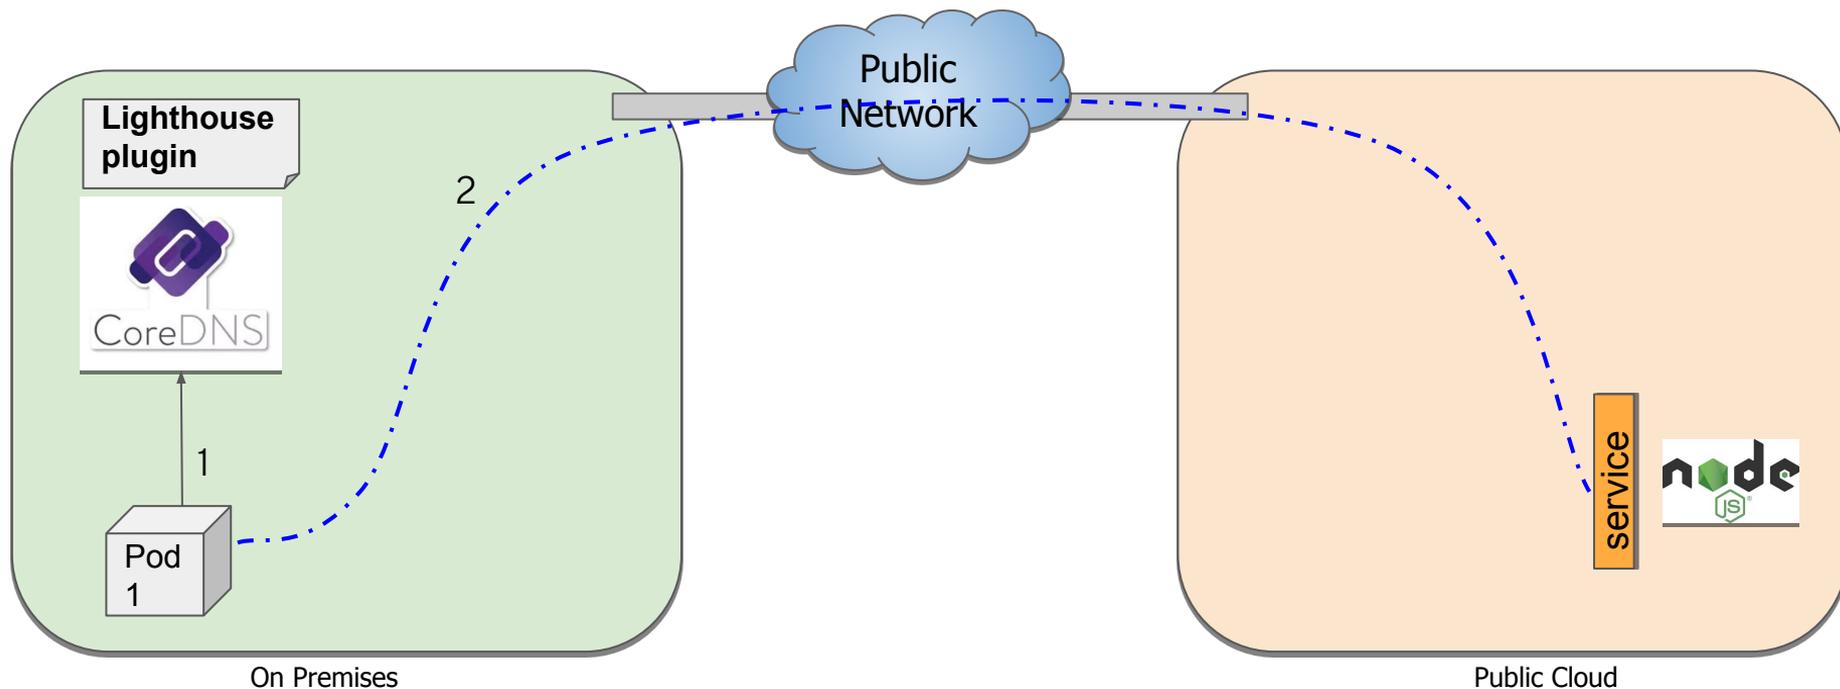
KubeCon



CloudNativeCon

North America 2019

submariner.io/lighthouse



Roadmap (DNS Discovery)



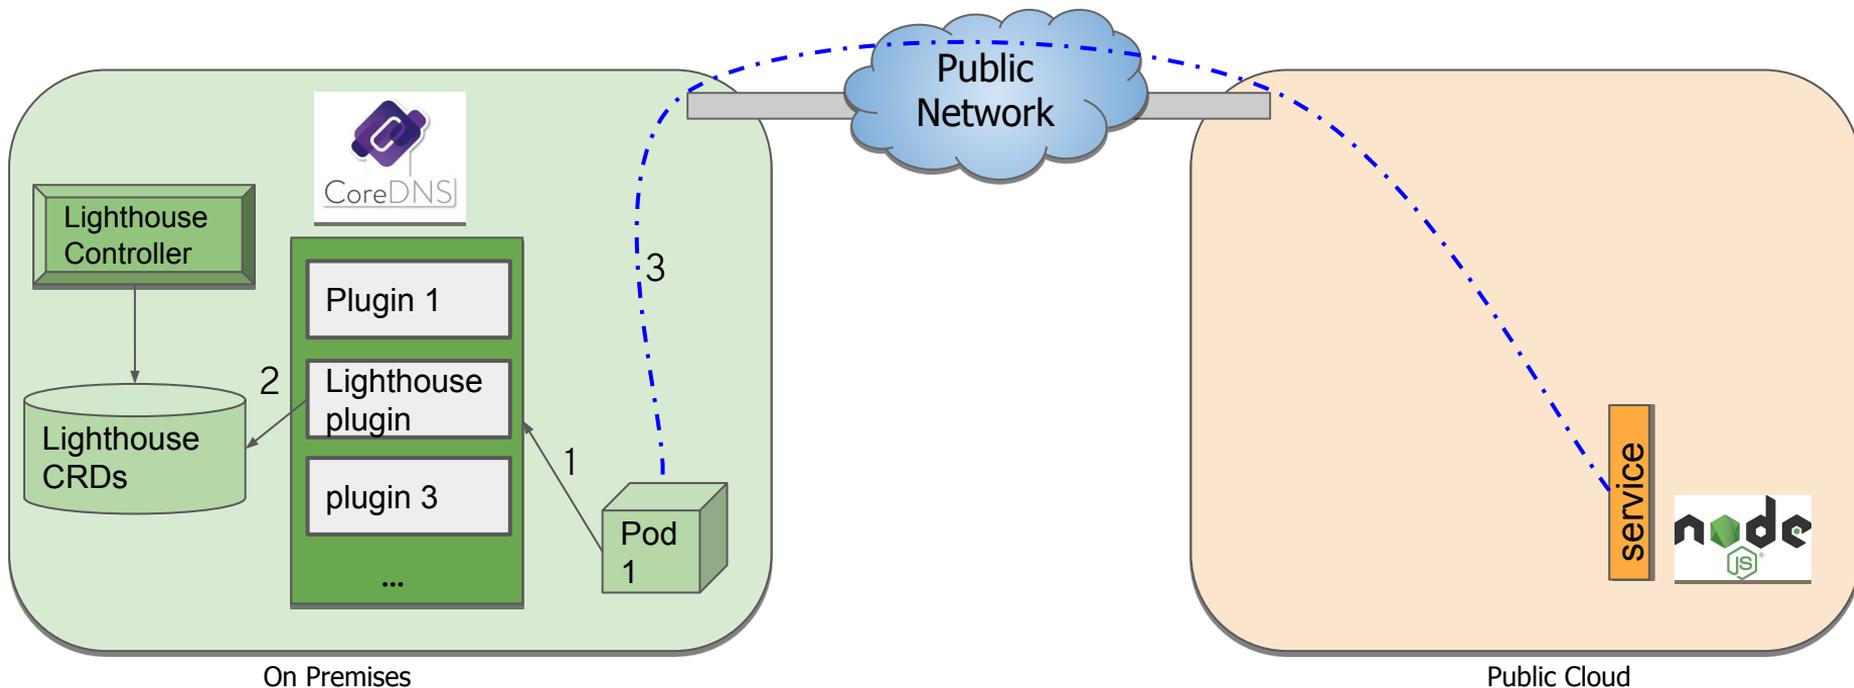
KubeCon



CloudNativeCon

North America 2019

submariner.io/lighthouse



Roadmap (Network Policies)



KubeCon



CloudNativeCon

North America 2019

submariner.io/**coastguard**

Coast Guard facilitates NetworkPolicy functionality in multi-cluster connected environments

Network plugins are unable to know about remote cluster pod IPs when processing NetworkPolicy objects within a cluster. Coastguard **analyzes the network policies and provides the IP addresses of remote cluster pods matching the network policy.**

<https://github.com/submariner-io/coastguard>

Roadmap (Network policies)



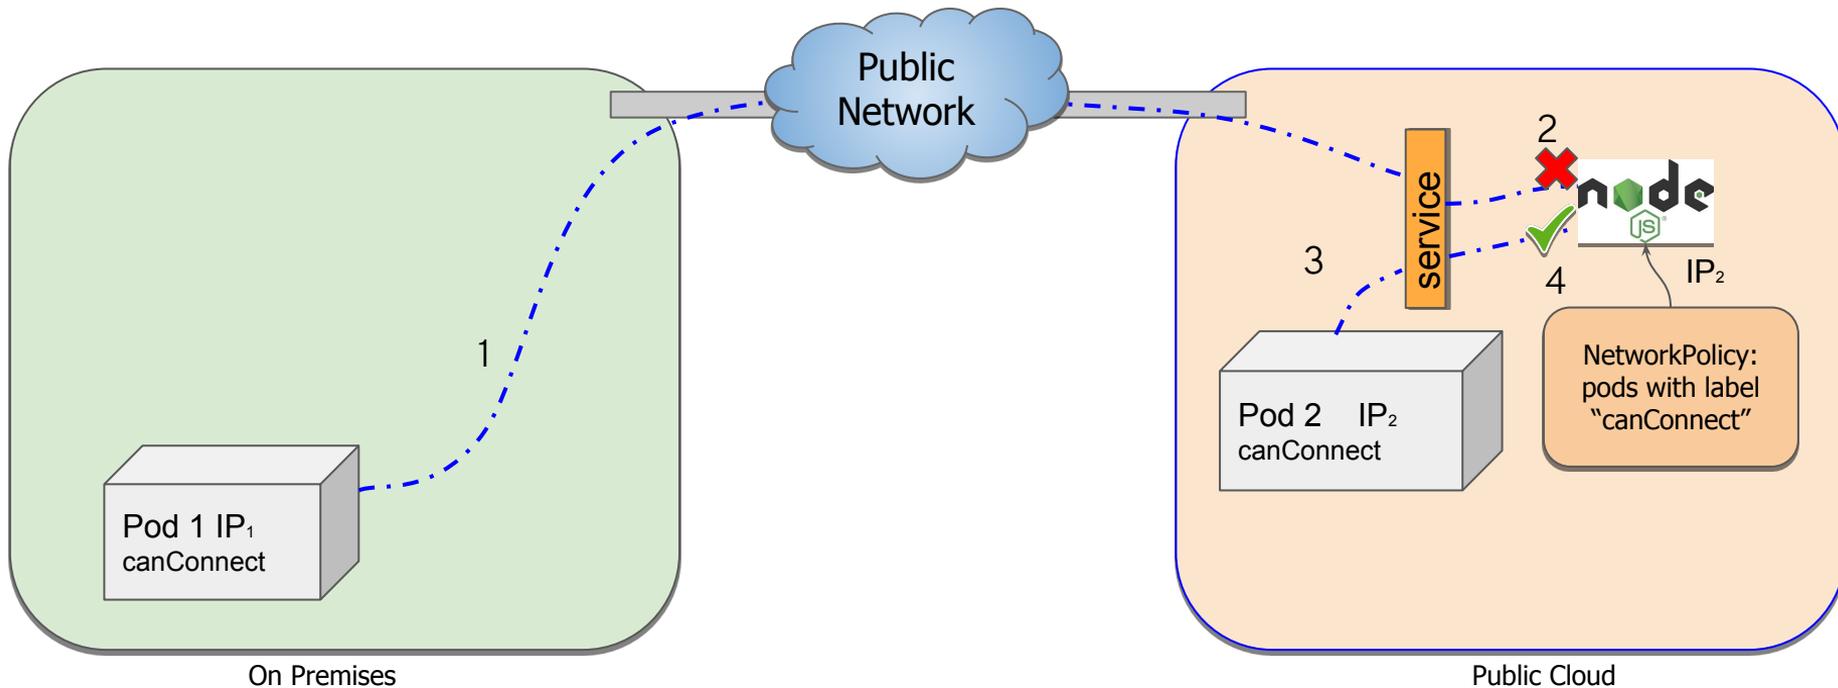
KubeCon



CloudNativeCon

North America 2019

submariner.io/coastguard (without coastguard running)



Roadmap (Network policies)



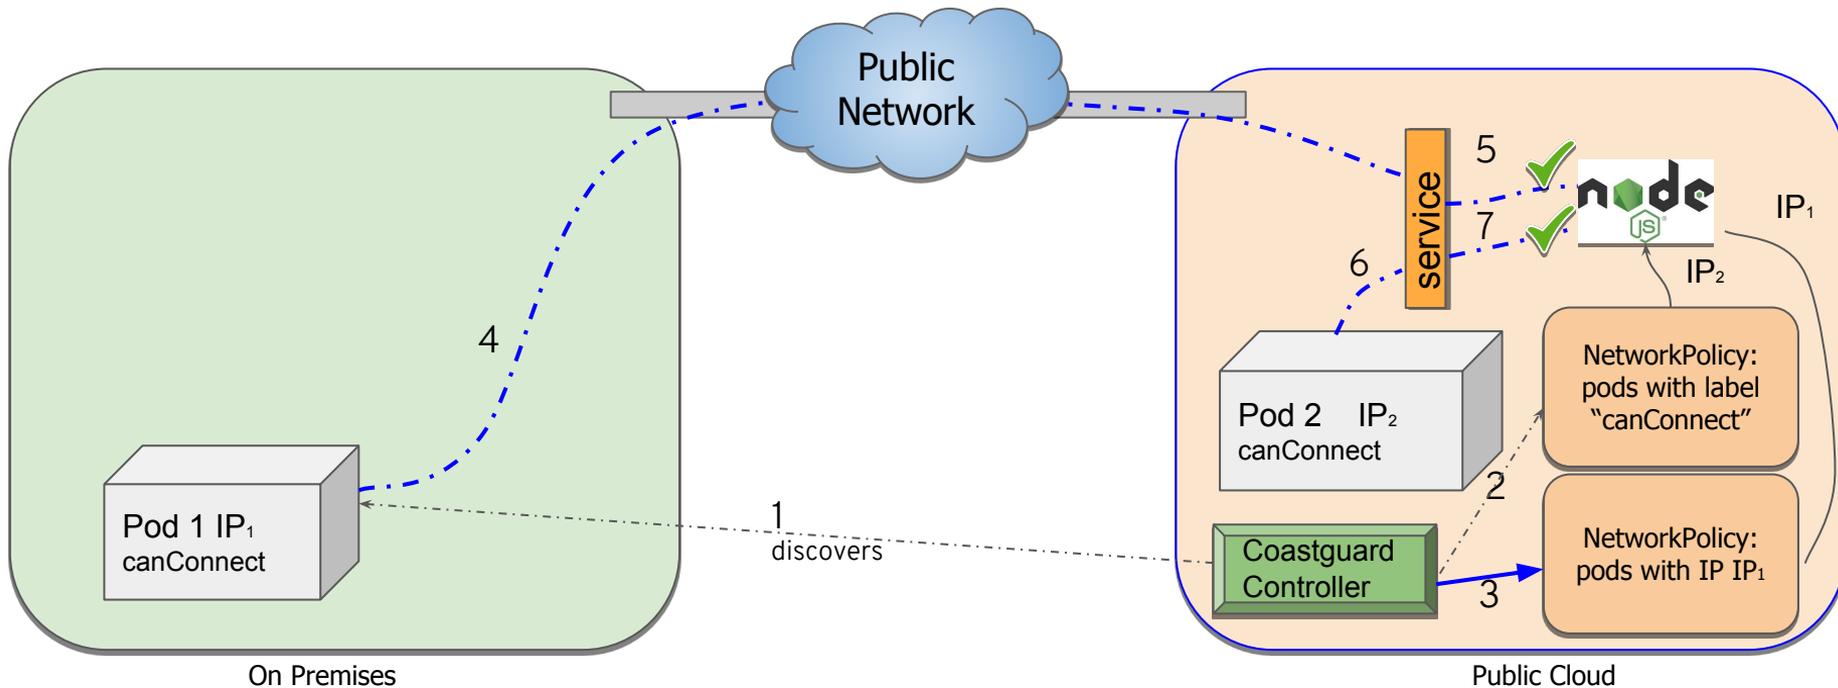
KubeCon



CloudNativeCon

North America 2019

submariner.io/coastguard



Roadmap (Network policies)



KubeCon



CloudNativeCon

North America 2019

submariner.io/**coastguard**

**In-depth proposal on how Network Policies could work
across clusters:**

https://docs.google.com/document/d/1_QzuasJPiQ-4t8tUGODoRS2E-Q3KLJajqMyCvLPH2tY



KubeCon



CloudNativeCon

North America 2019

Deployment Options

Deployment Options



KubeCon



CloudNativeCon

North America 2019

- **submariner-operator**
 - Installed on the cluster manages your submariner deployment and maintenance (upgrades, status, etc.)
- **helm**
 - <https://github.com/submariner-io/submariner-charts>
 - helm repo add submariner-latest https://submariner-io.github.io/submariner-charts/charts
- **subctl (+ subm-operator)**
 - cmdline interface to the operator (install, easy-configuration, status, upgrade...)

subctl Deployment



KubeCon



CloudNativeCon

North America 2019

subctl

```
$ export KUBECONFIG=my-broker-cluster/kubeconfig
```

```
$ subctl deploy broker
```

```
* Deploying submariner-operator.
```

```
* No labeled gateway node detected, please select one:
```

```
< list of nodes and IPs >, emphasize/suggest nodes with external IP.
```

```
* Deployed: SubmarinerBroker, Submariner...
```

```
* Created random IPSEC PSK stored in a k8s secret in the broker.
```

```
* Created submariner-xxx-yy-zz.json with signup details
```

```
* Waiting on submariner deployment to be ready: ... done.
```

Your broker is ready, use submariner-xxx-yy-zz.json to join other clusters.

subctl Deployment



KubeCon



CloudNativeCon

North America 2019

subctl

```
$ export KUBECONFIG=cluster-B/kubeconfig
```

```
$ subctl join submariner-xxx-yy-zz.json
```

```
* Deploying submariner-operator.
```

```
* No labeled gateway node detected, please select one:
```

```
< list of nodes and IPs >, emphasize/suggest nodes with external IP.
```

```
* Stored Secrets with broker access credentials and IPSEC details.
```

```
* Deployed: SubmarinerBroker, Submariner...
```

```
* Created submariner-xxx-yy-zz.json with signup details
```

```
* Waiting on submariner deployment to be ready: ... done.
```

```
Your cluster pods are ready to talk to other clusters.
```

Demo



KubeCon



CloudNativeCon

North America 2019

Demo time!



KubeCon



CloudNativeCon

North America 2019

Q&A

Thank you :-)

Time tracking:

- 2.0 min. What's submariner?
- 1.5min. History of submariner
- 0.5min. Roadmap list
 - 0.5min Overlapping IPs
 - 1.5min. Lighthouse
 - 2.0min. Coastguard
- 2.5min deployment options
- 7.0 min Demo

Total: 18min

Available for Q&A: $30-18 = 12$ min