



KubeCon



CloudNativeCon

Europe 2018

# Cloud native container networking in AWS using CNI plugins

Anirudh Aithal – Sr SDE, AWS



# \$ whoami



KubeCon



CloudNativeCon

Europe 2018



**Anirudh Aithal**

aithal

I'm a developer with Amazon  
Elastic Container Services (ECS)

@aws ECS

Seattle, WA

Overview

Repositories 15

Stars 42

Foll

## Popular repositories

**amazon-ecs-agent**

Forked from [aws/amazon-ecs-agent](#)

Amazon EC2 Container Service Agent

Go 9 1

**cni**

Forked from [containernetworking/cni](#)

Container Network Interface - networking for Linux  
containers

Go 1



# Topics covered in today's talk



KubeCon



CloudNativeCon

Europe 2018



- Networking requirements for containerized applications
- Container networking primitives
- Developing (Amazon VPC) CNI plugins

# What do applications need?

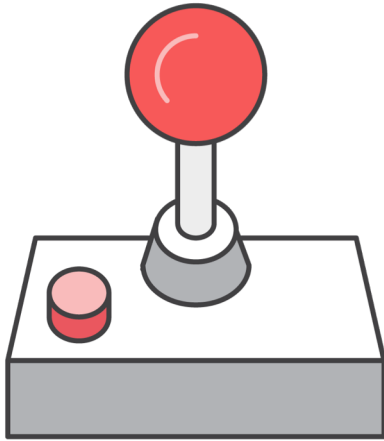


KubeCon



CloudNativeCon

Europe 2018



## Usability

- Simple abstractions
- Discovery



## Security

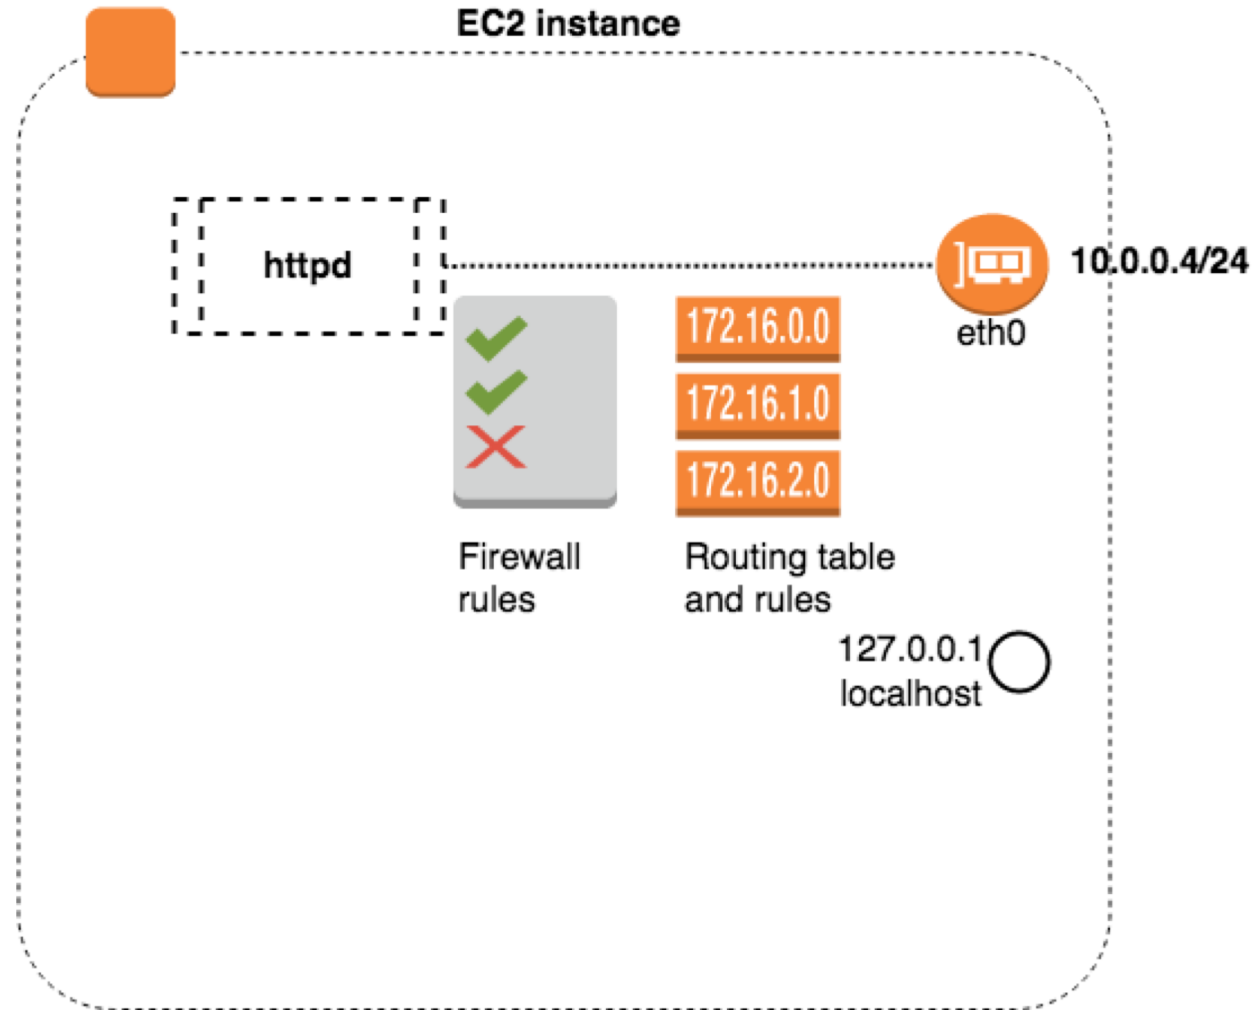
- Network isolation
- Access control
- Auditability



## Maintainability

- Scalability
- Performance
- Monitoring

# One listener : one host



# One listener : one host

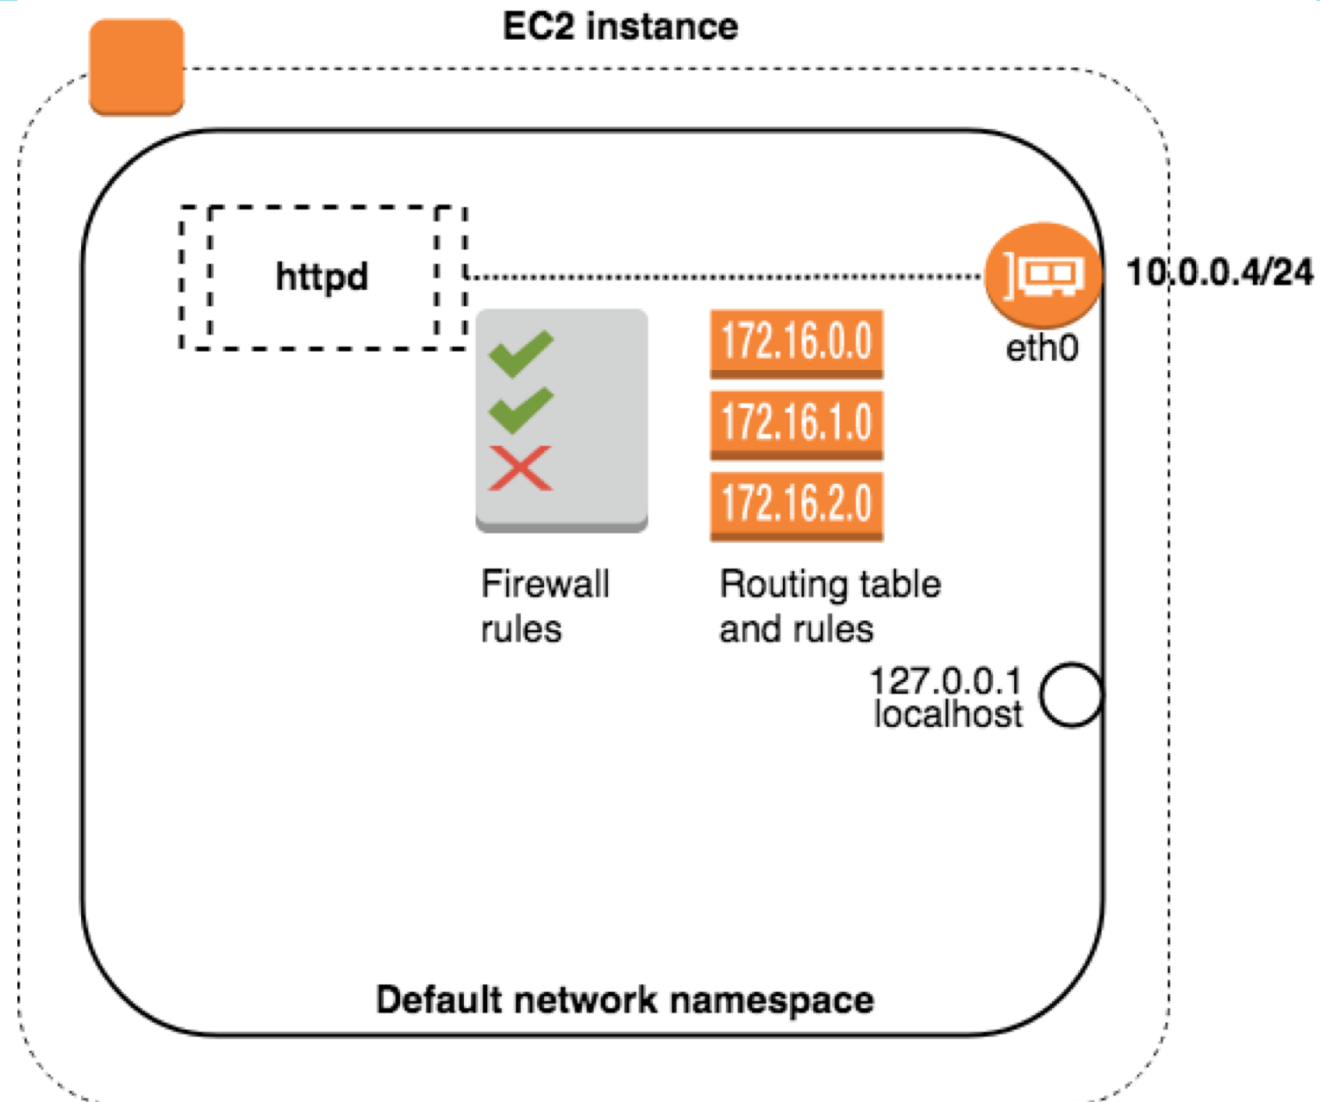


KubeCon



CloudNativeCon

Europe 2018



# Two listeners : one host !

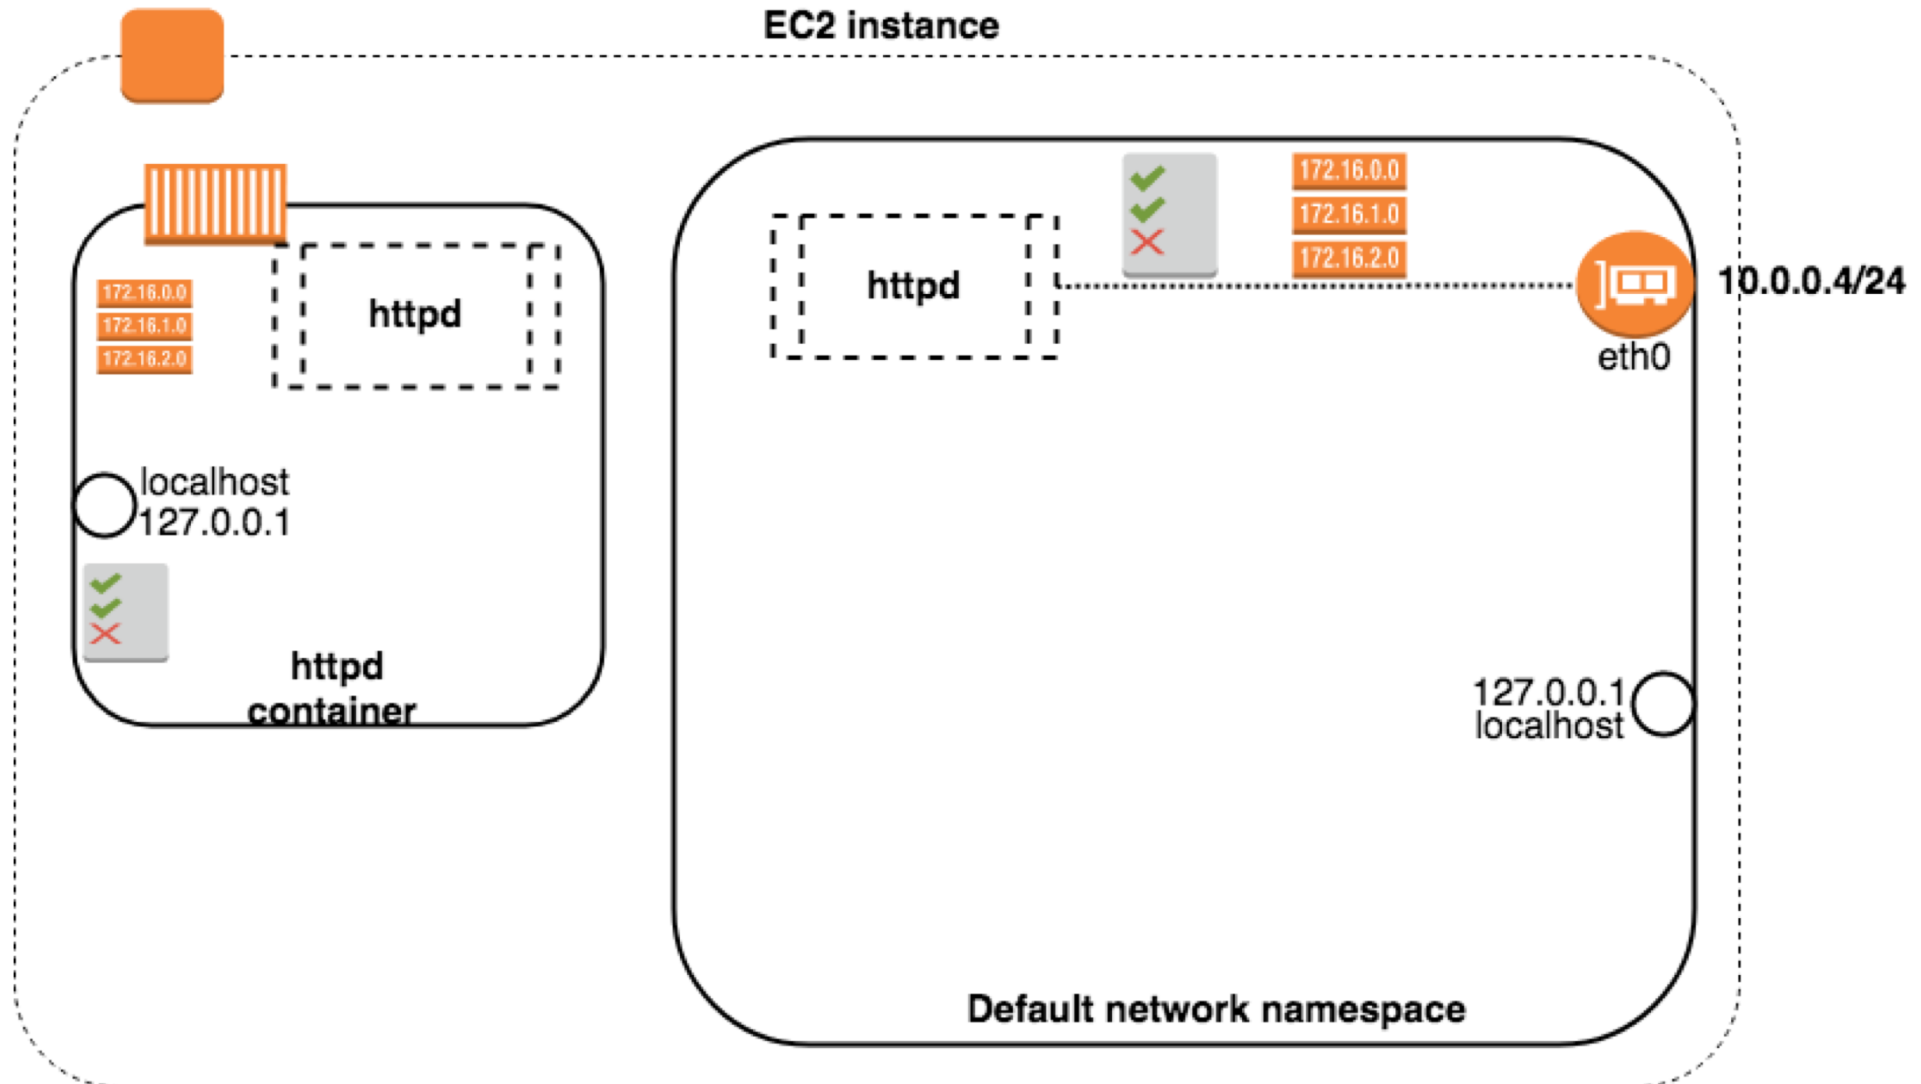


KubeCon



CloudNativeCon

Europe 2018



# Two listeners : one host !

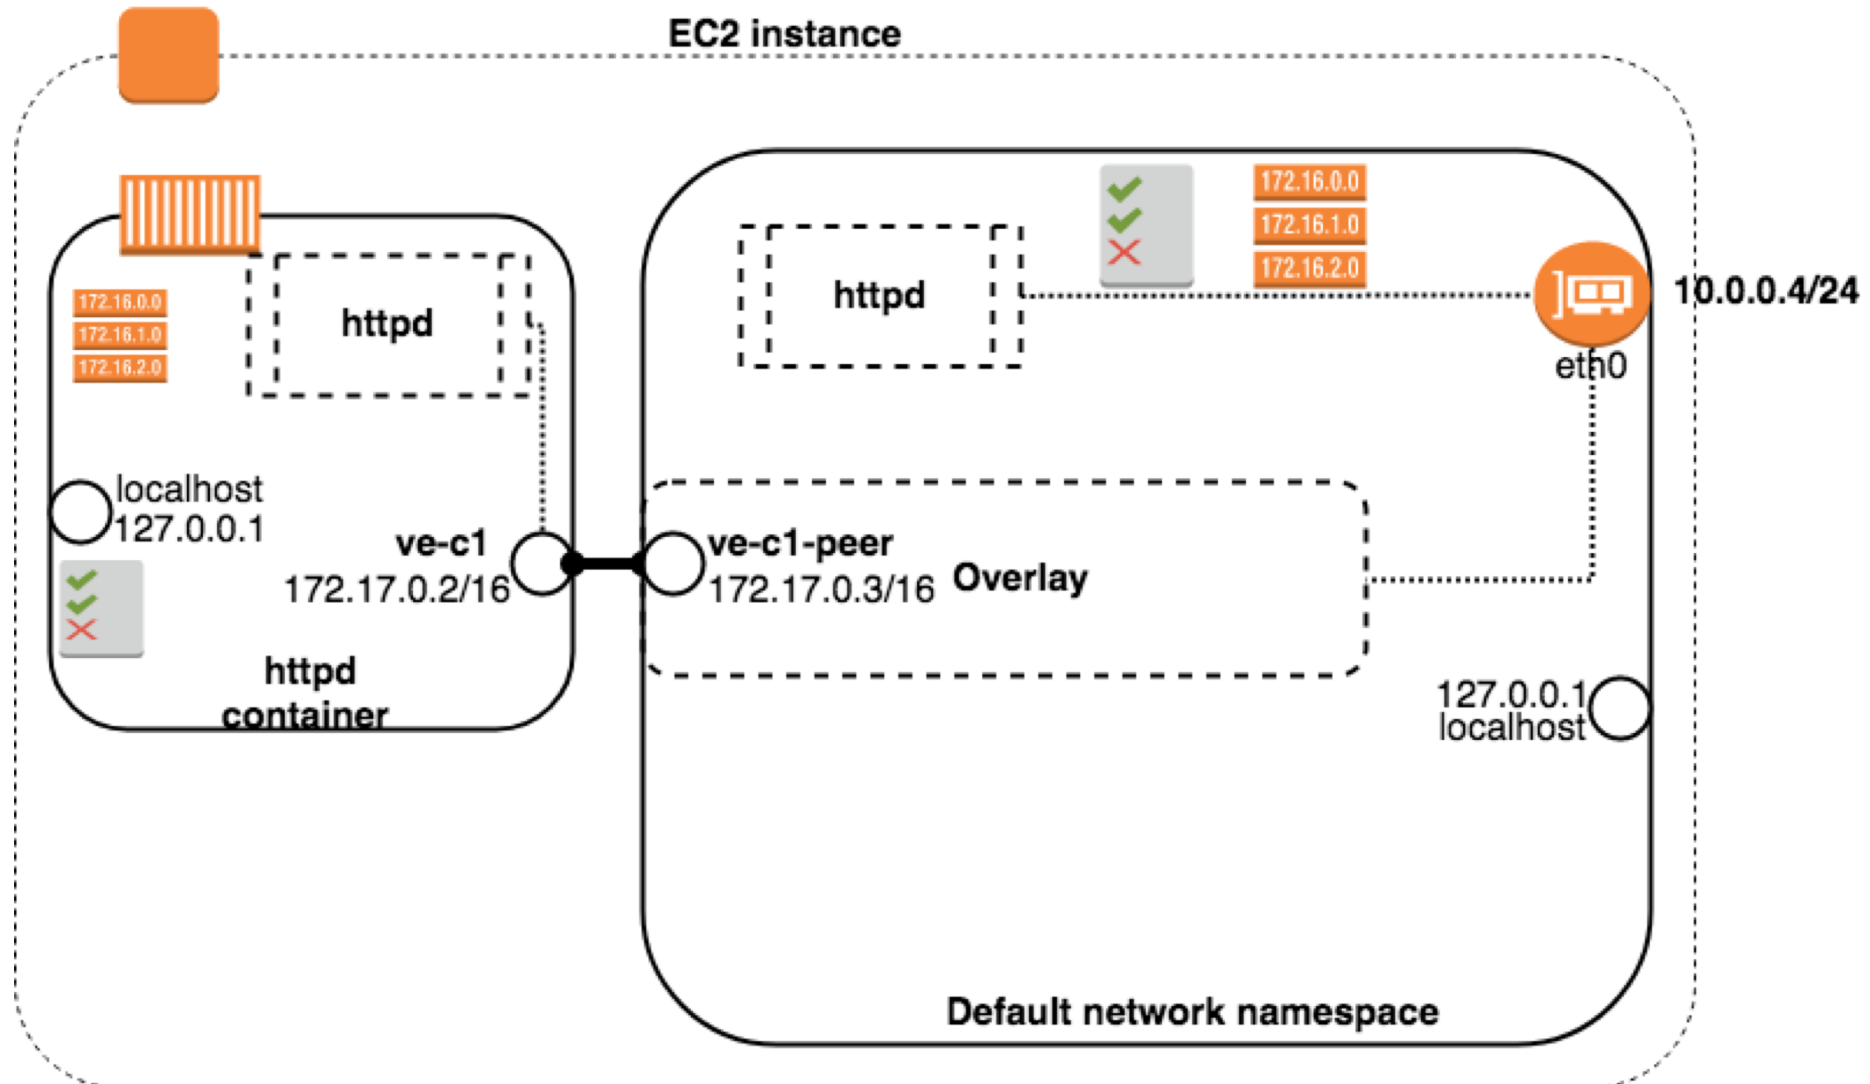


KubeCon



CloudNativeCon

Europe 2018





# Connecting containers across hosts

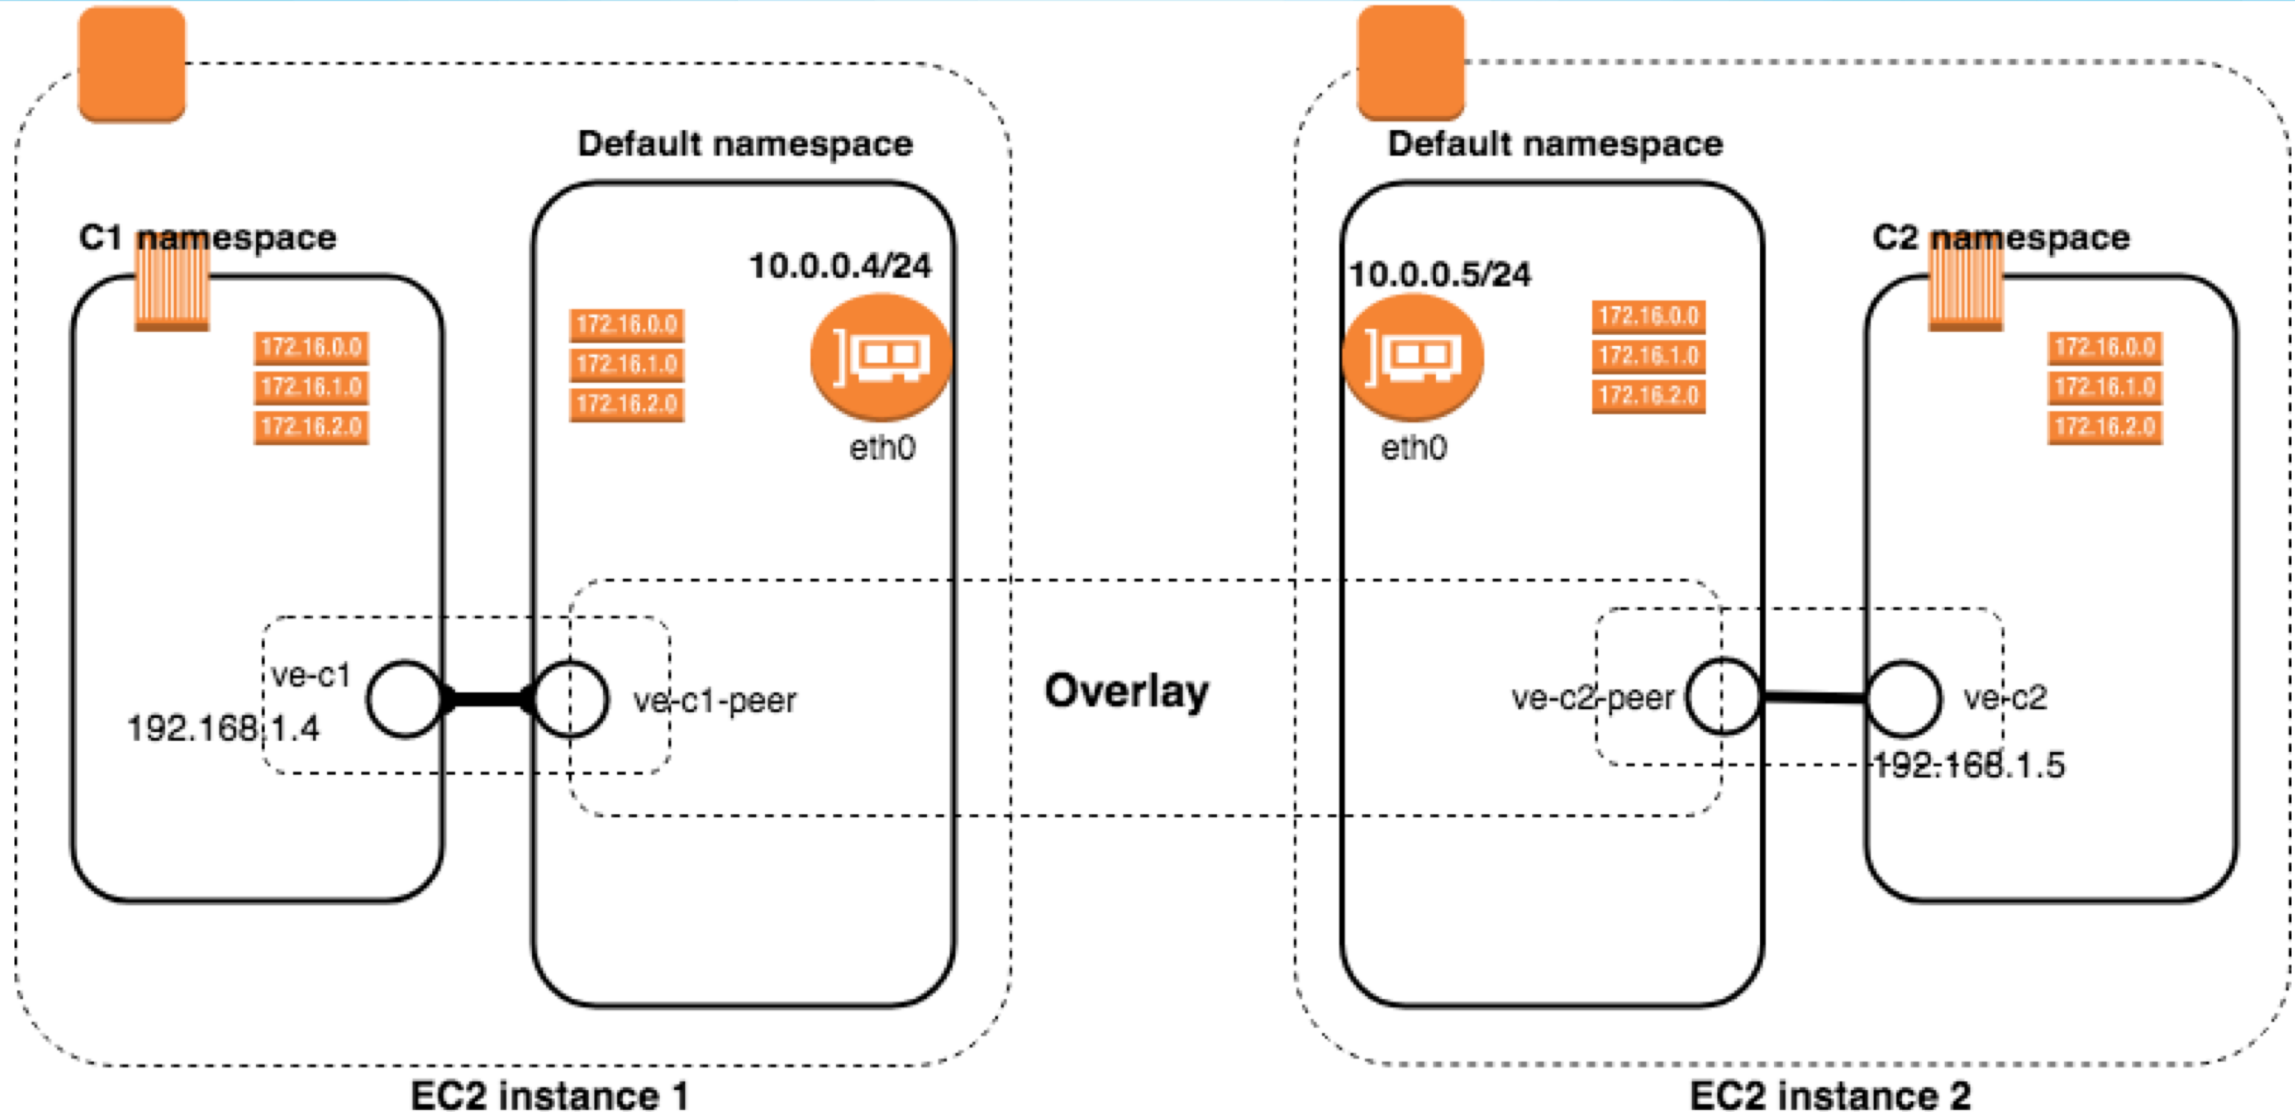


KubeCon



CloudNativeCon

Europe 2018



# Container communication – 0

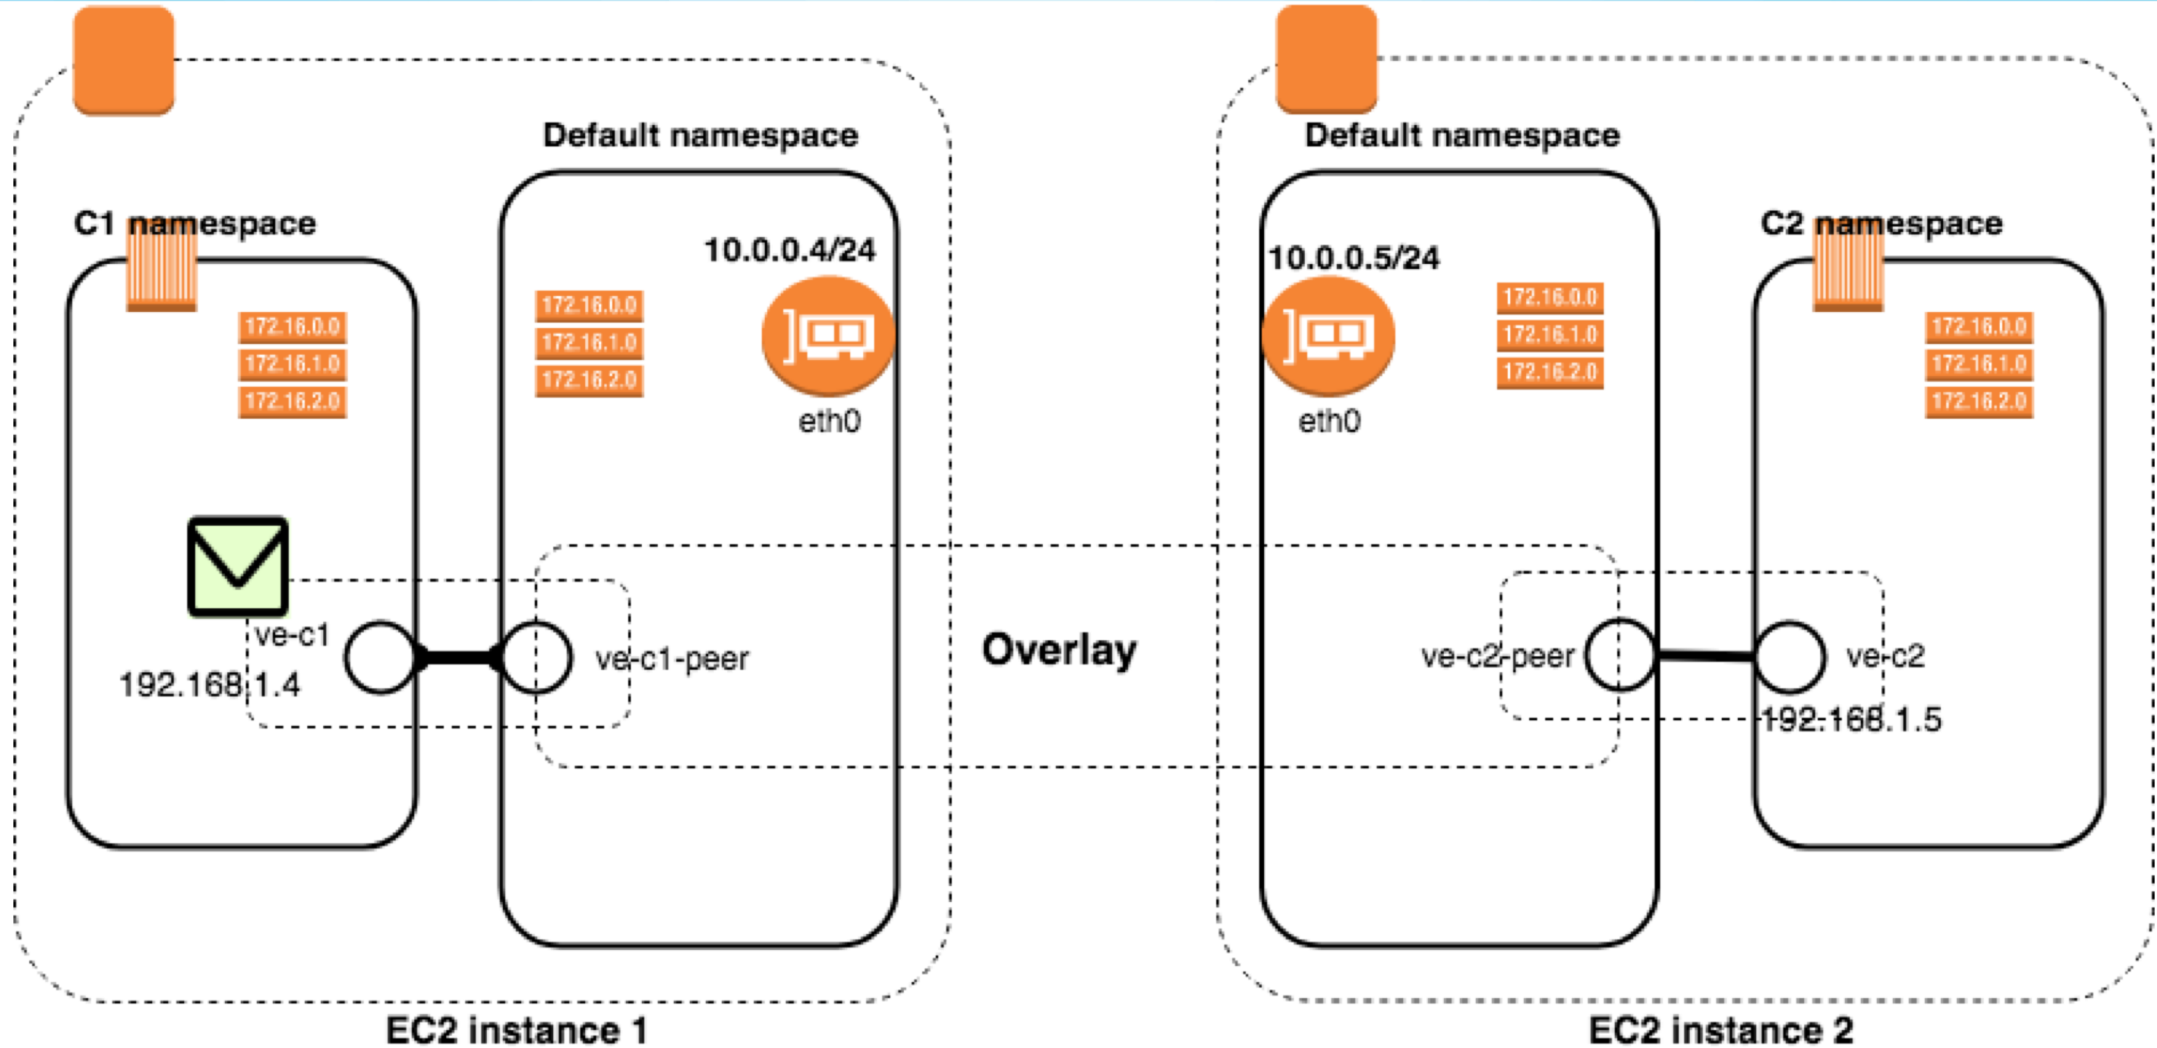


KubeCon



CloudNativeCon

Europe 2018



# Container communication – 1

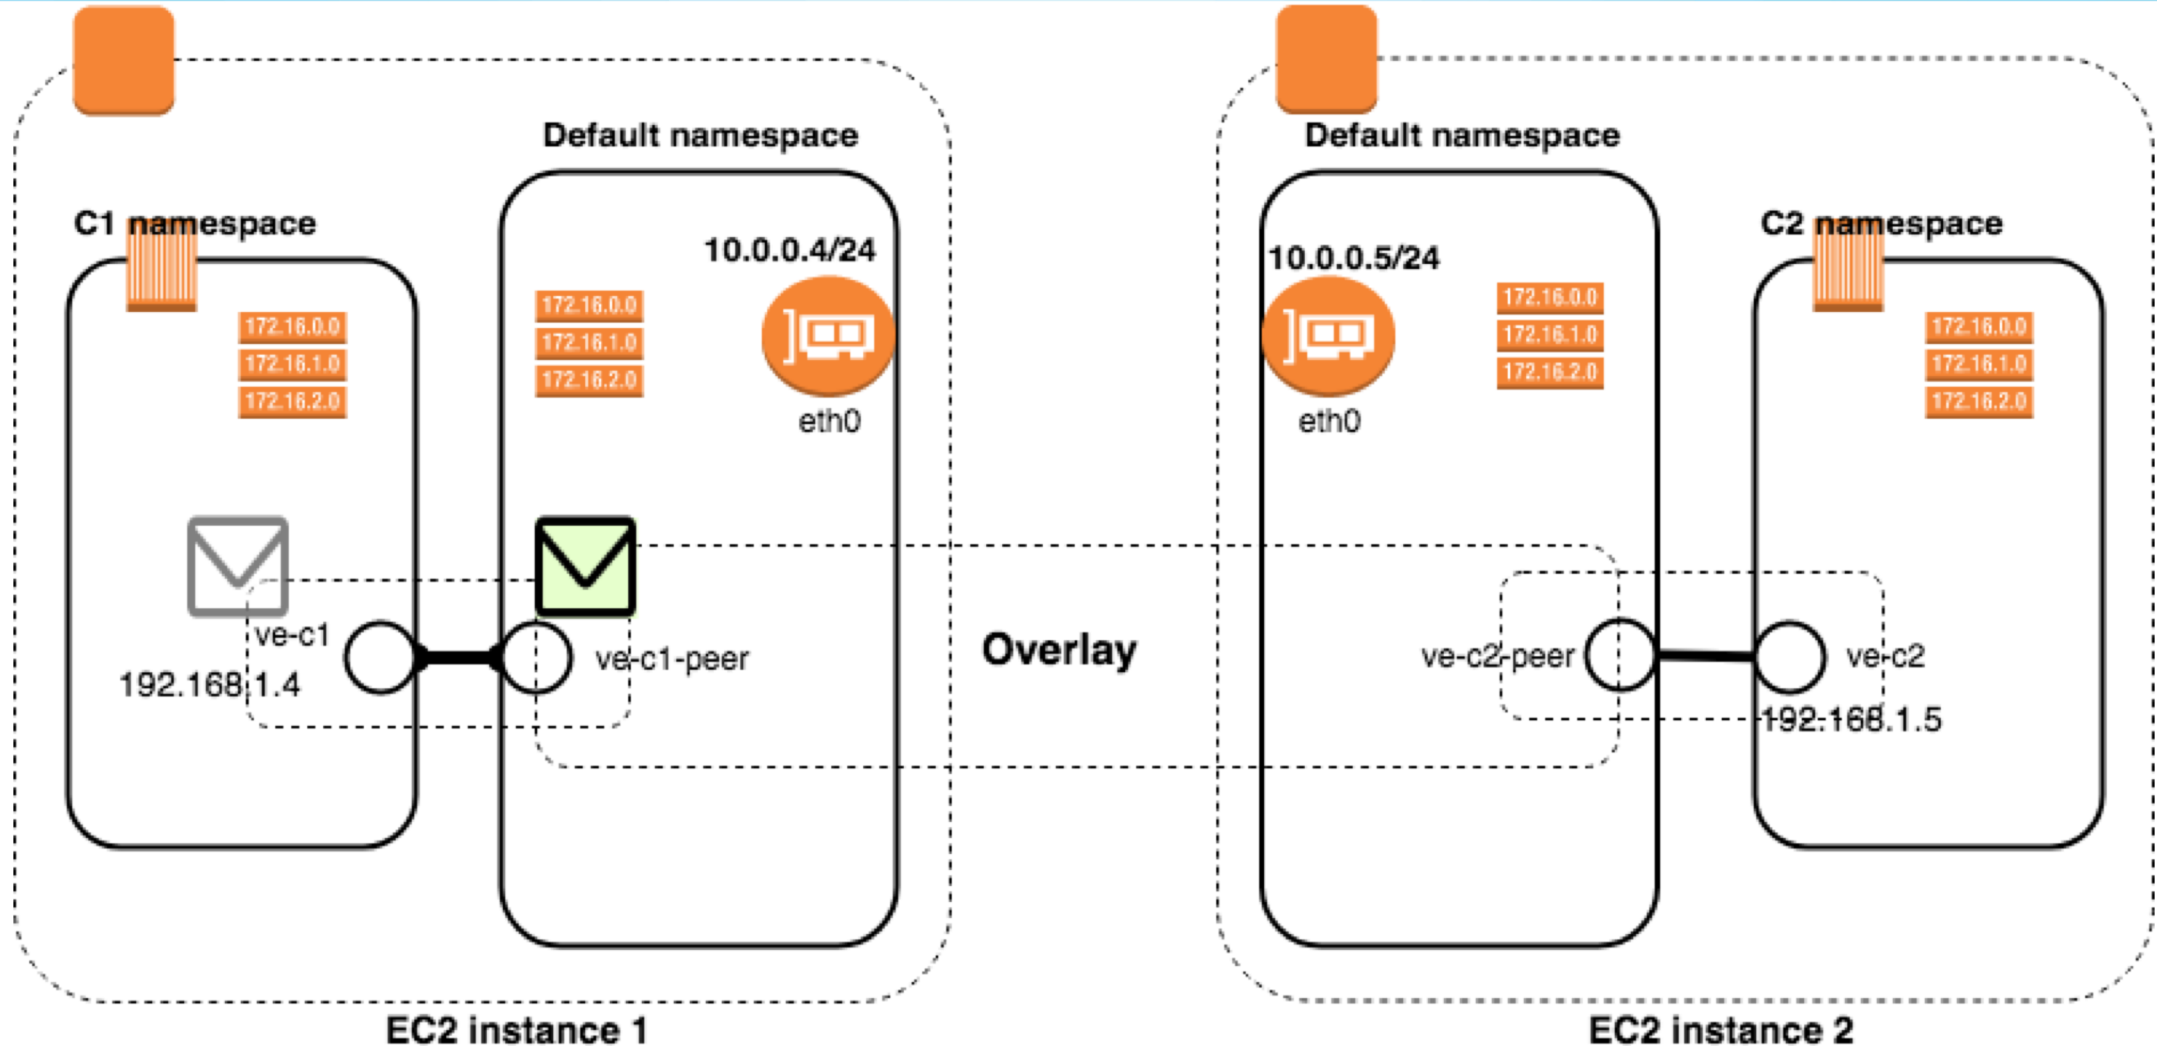


KubeCon



CloudNativeCon

Europe 2018



# Container communication – 2

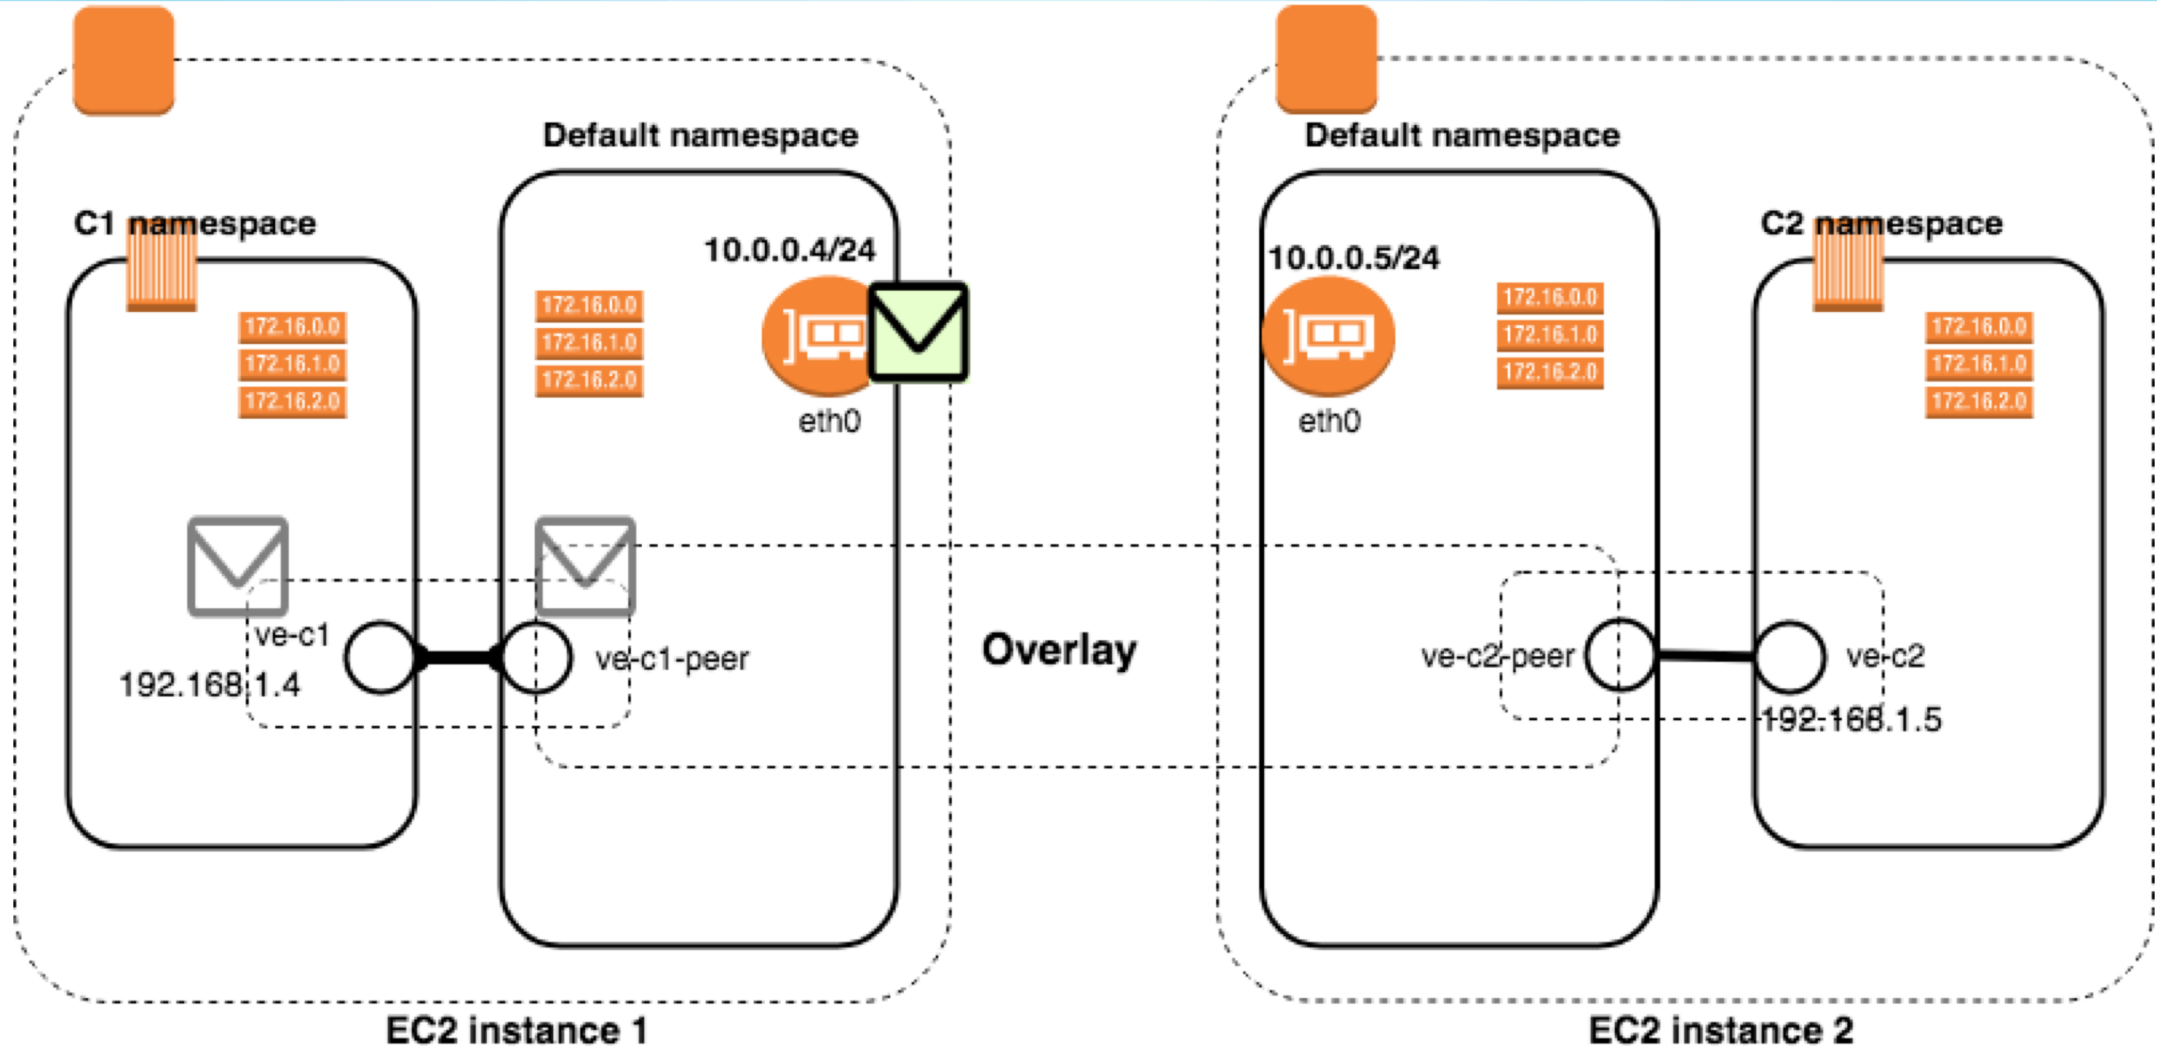


KubeCon



CloudNativeCon

Europe 2018



# Container communication – 3

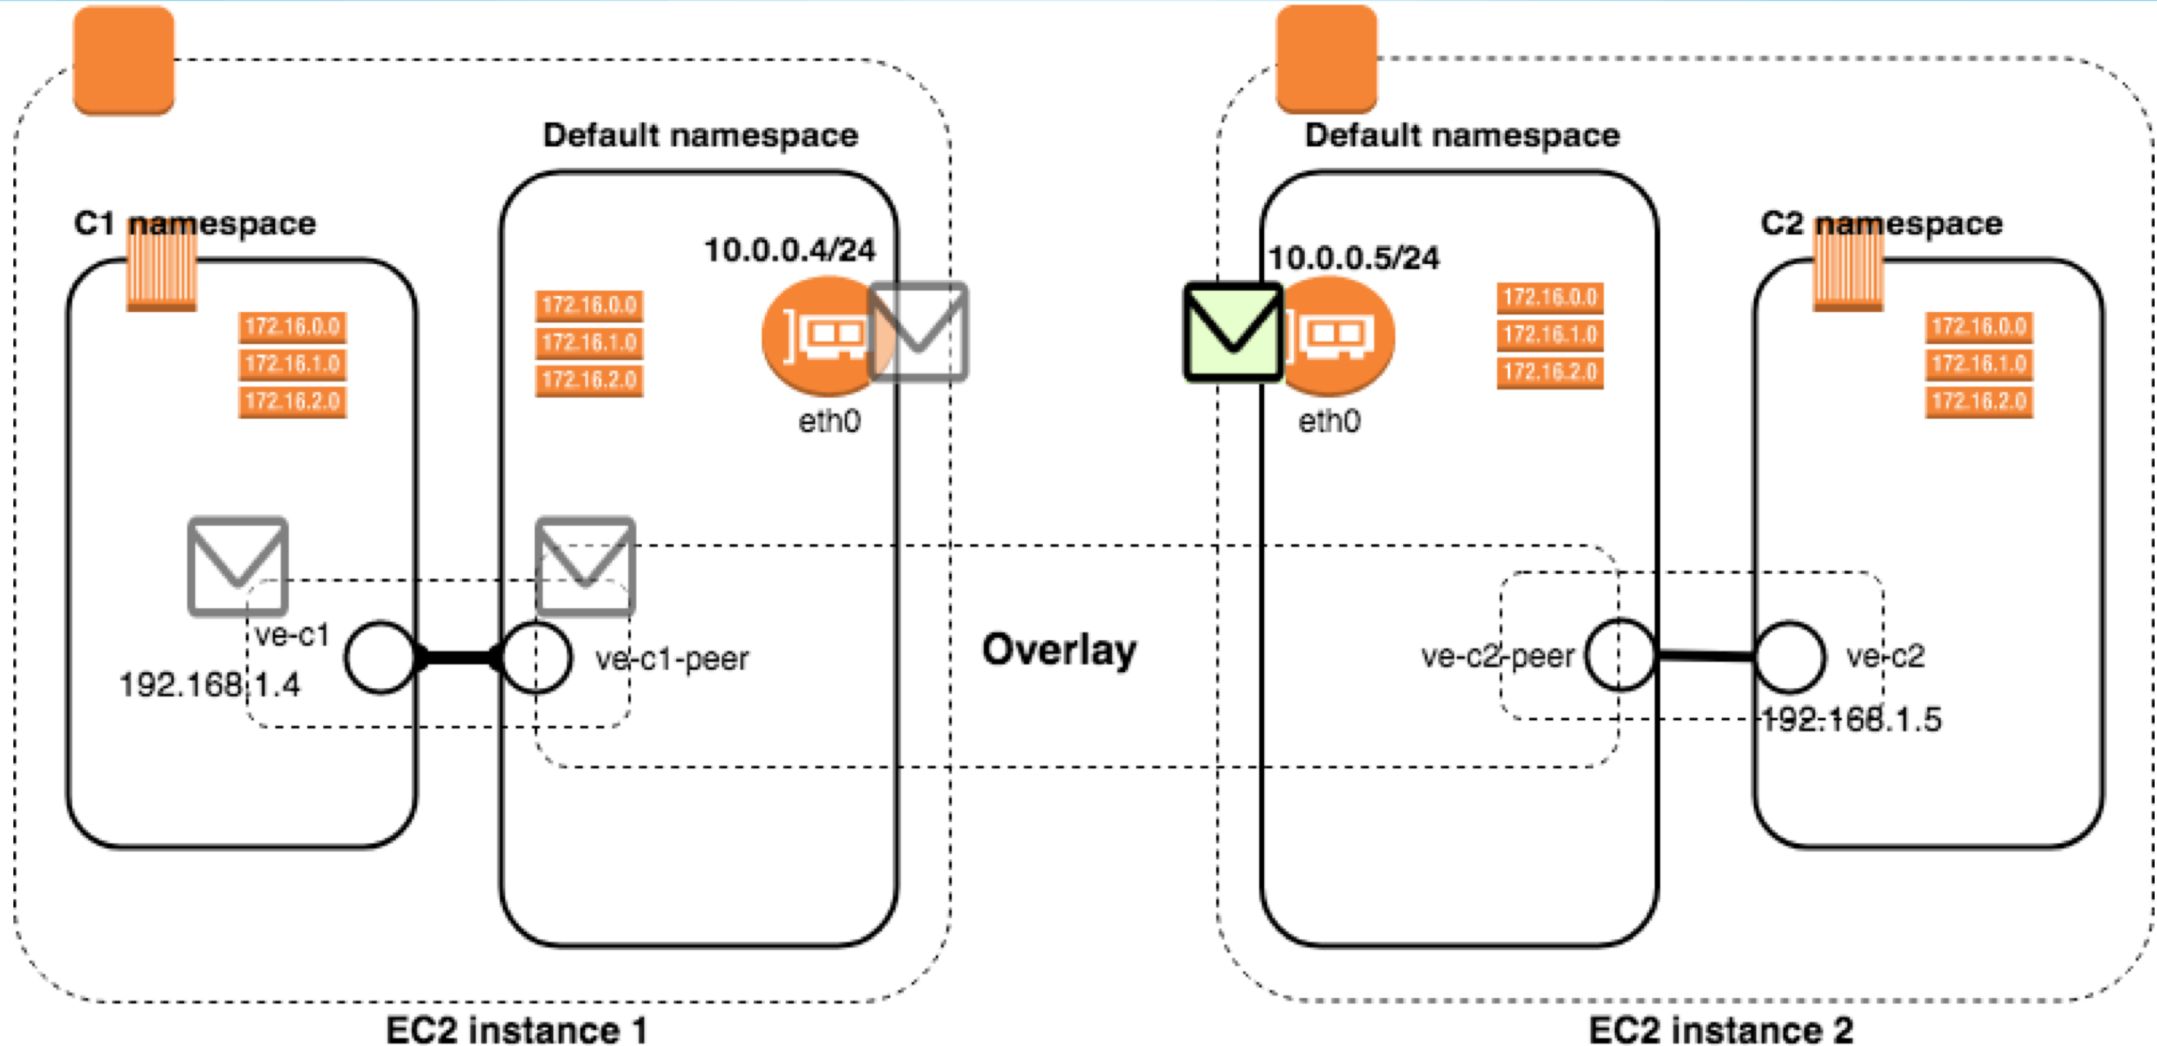


KubeCon



CloudNativeCon

Europe 2018



# Container communication – 4

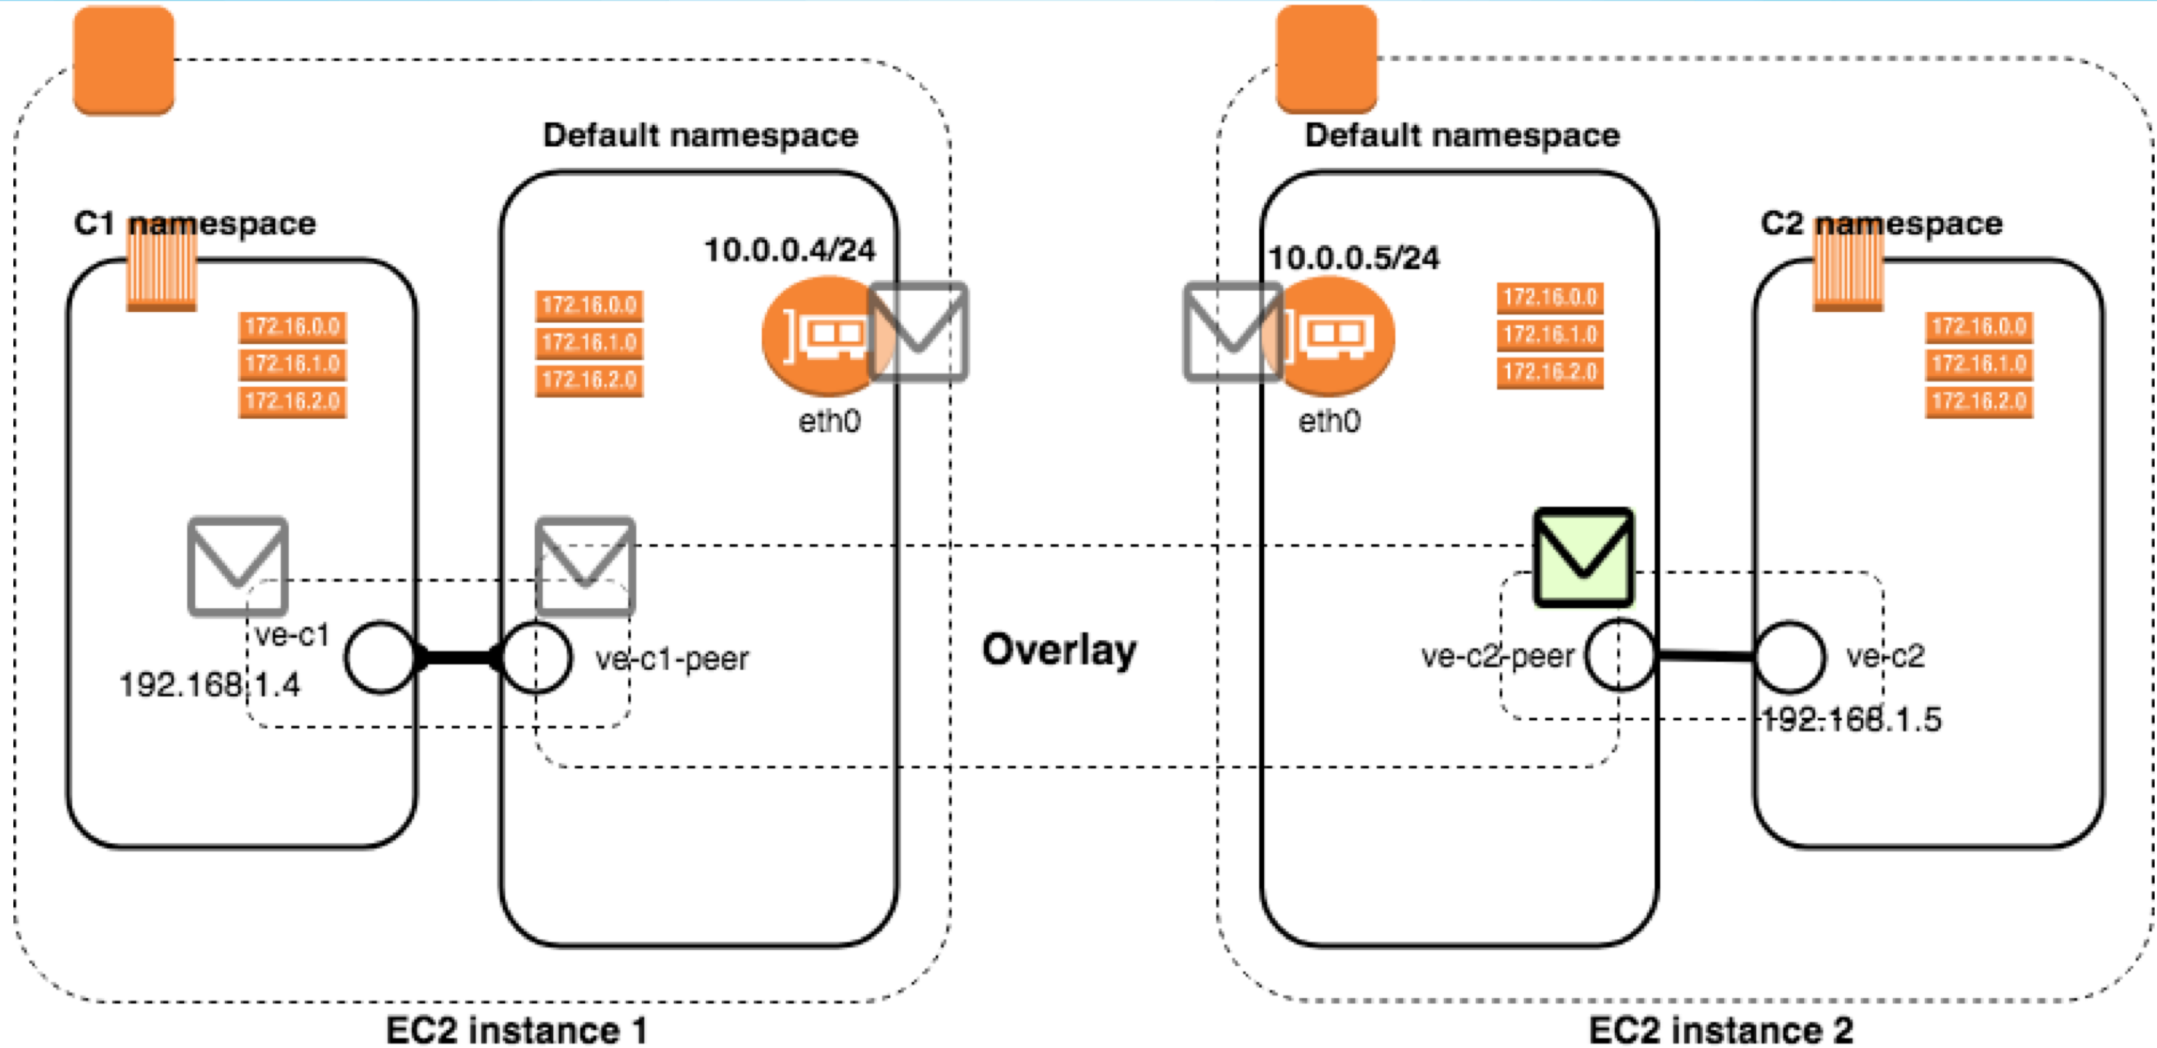


KubeCon



CloudNativeCon

Europe 2018





# Container communication – 5

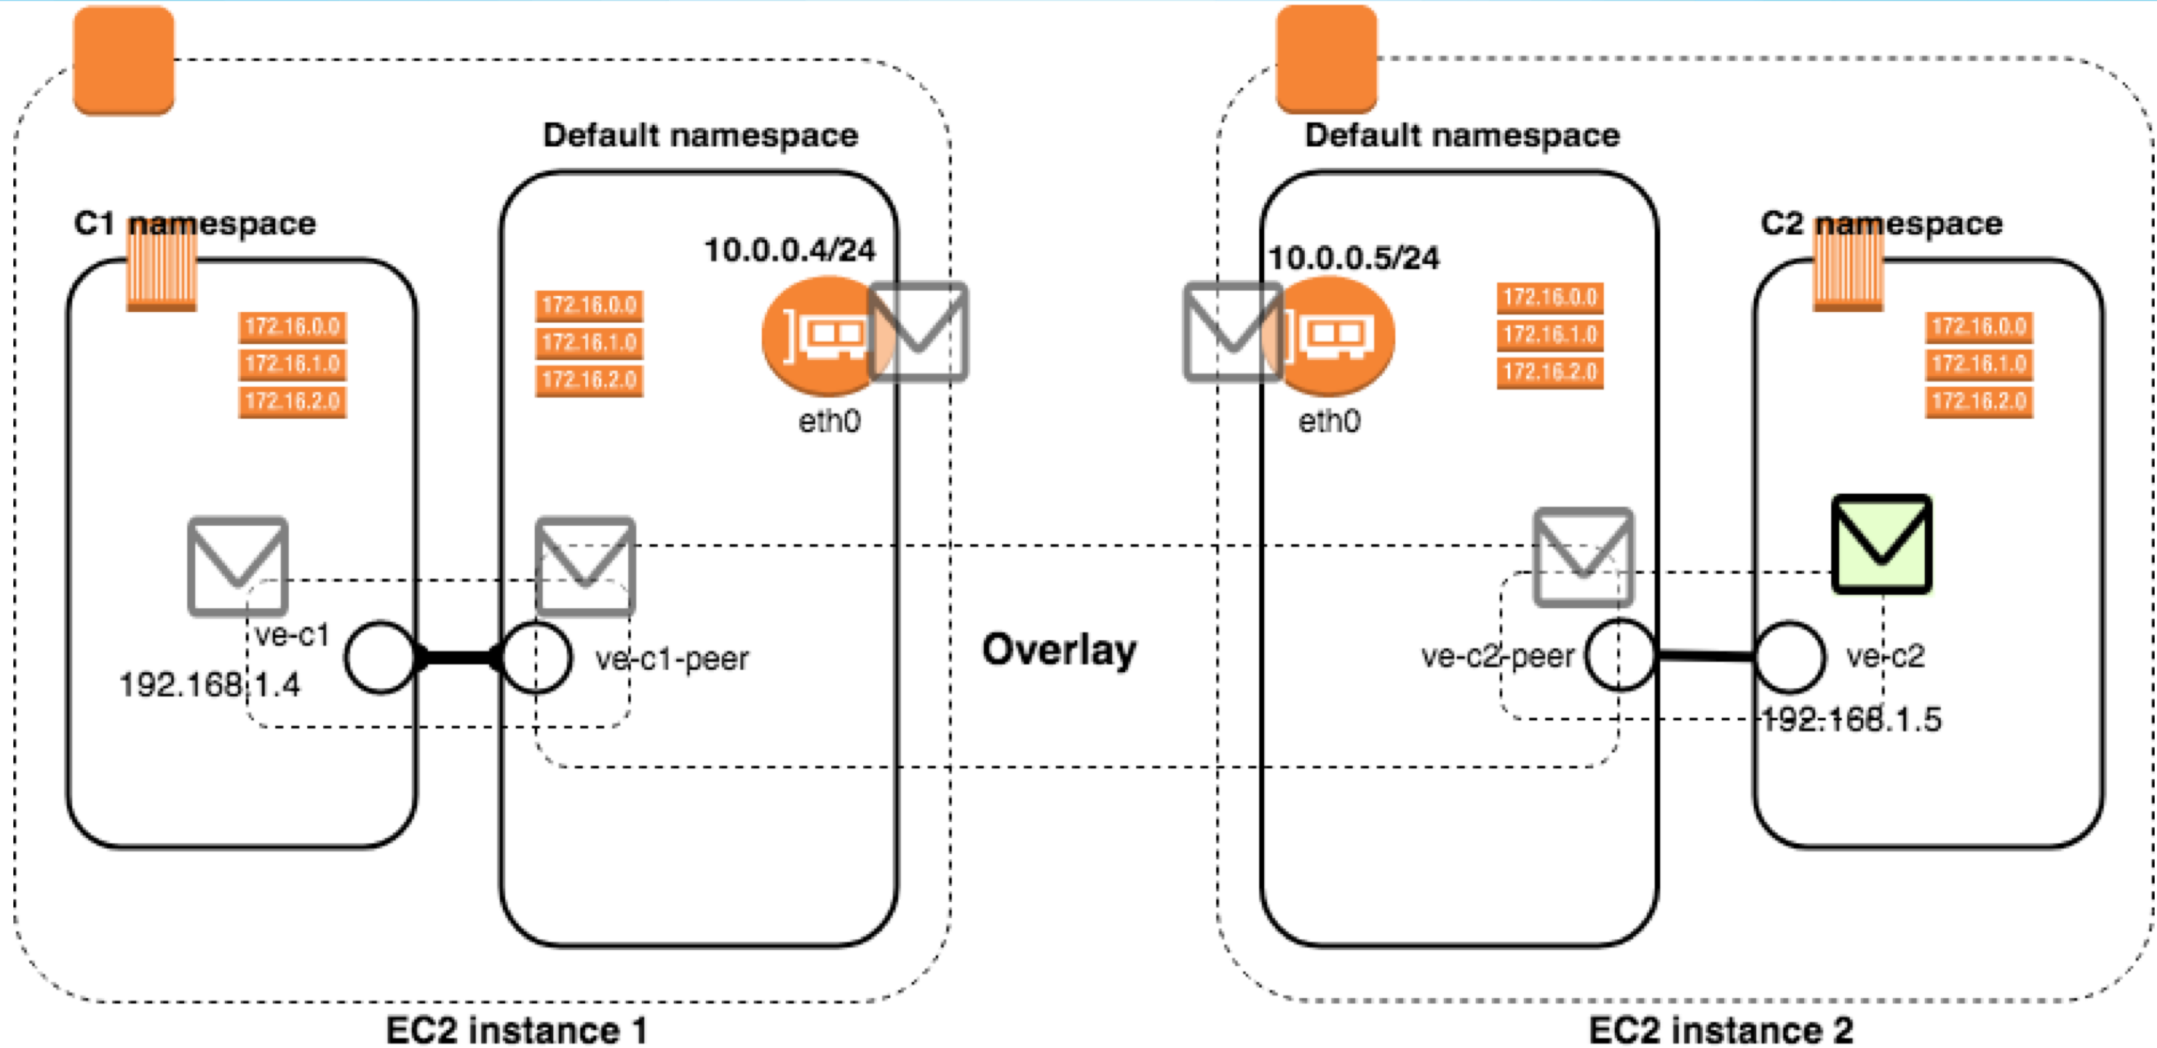


KubeCon



CloudNativeCon

Europe 2018



# How did we do?




KubeCon



CloudNativeCon

Europe 2018



## Usability

- Simple abstractions
- Discovery



## Security

- ~~Network isolation~~
- ~~Access control~~
- ~~Auditability~~



## Maintainability

- ~~Scalability~~
- ~~Performance~~
- Monitoring

# VPC networking – EC2 instances

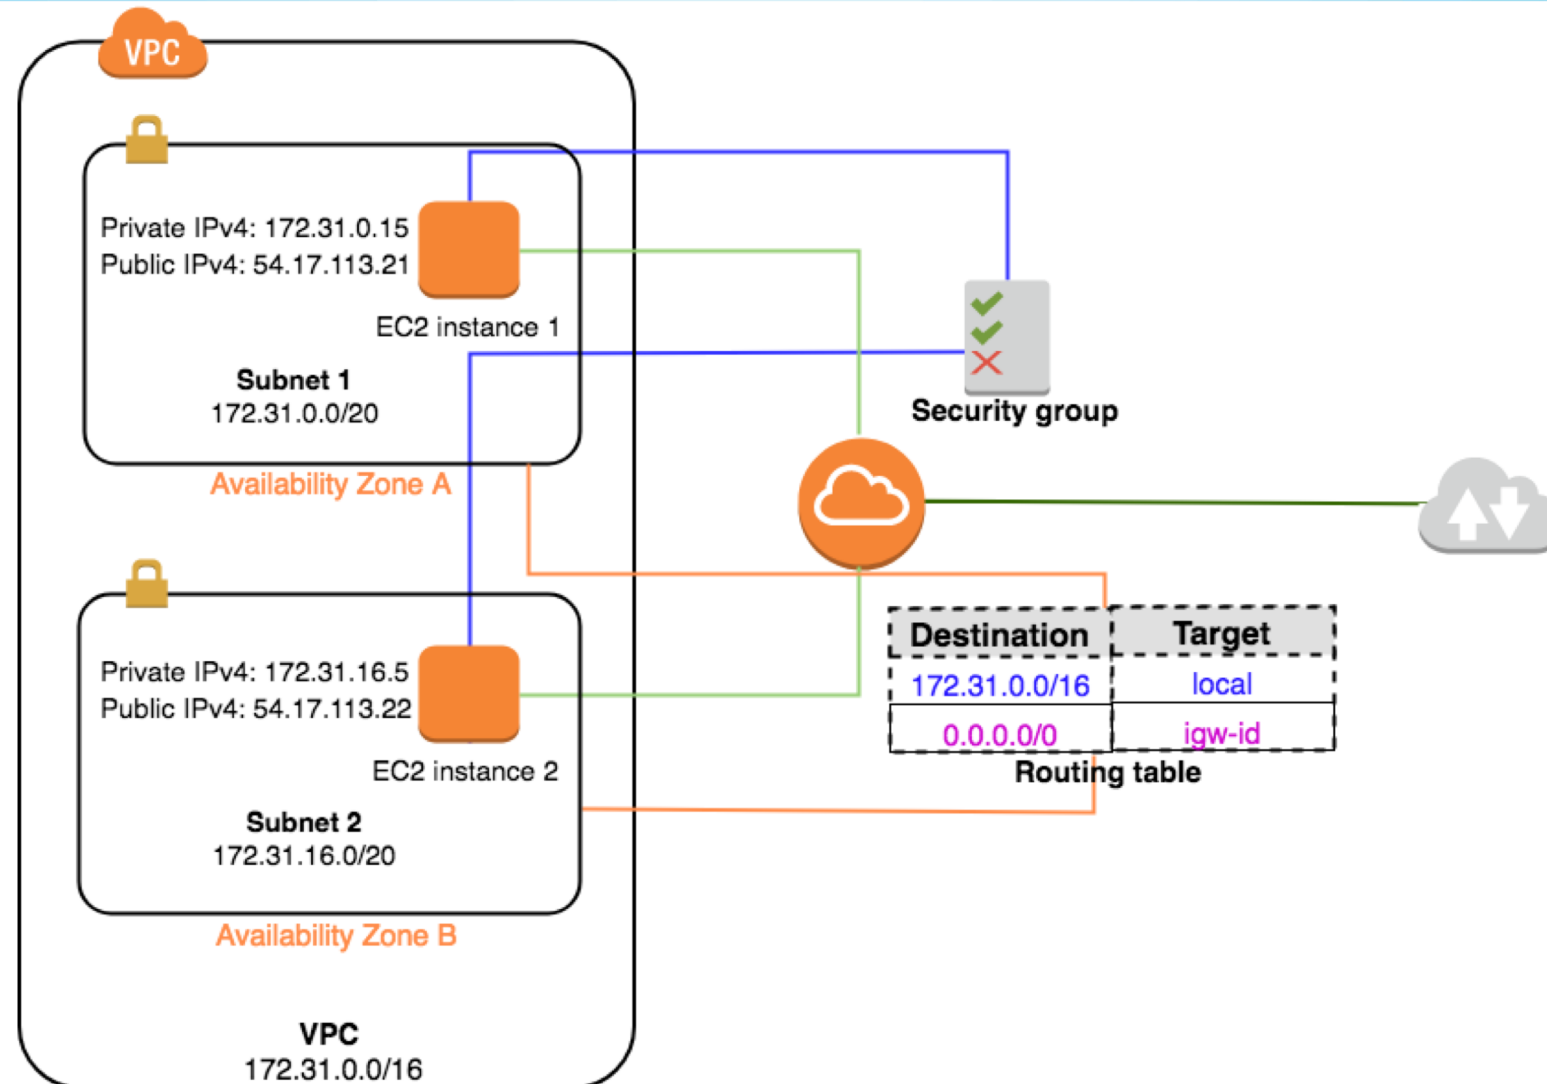


KubeCon



CloudNativeCon

Europe 2018



# Packet flow – EC2 instances

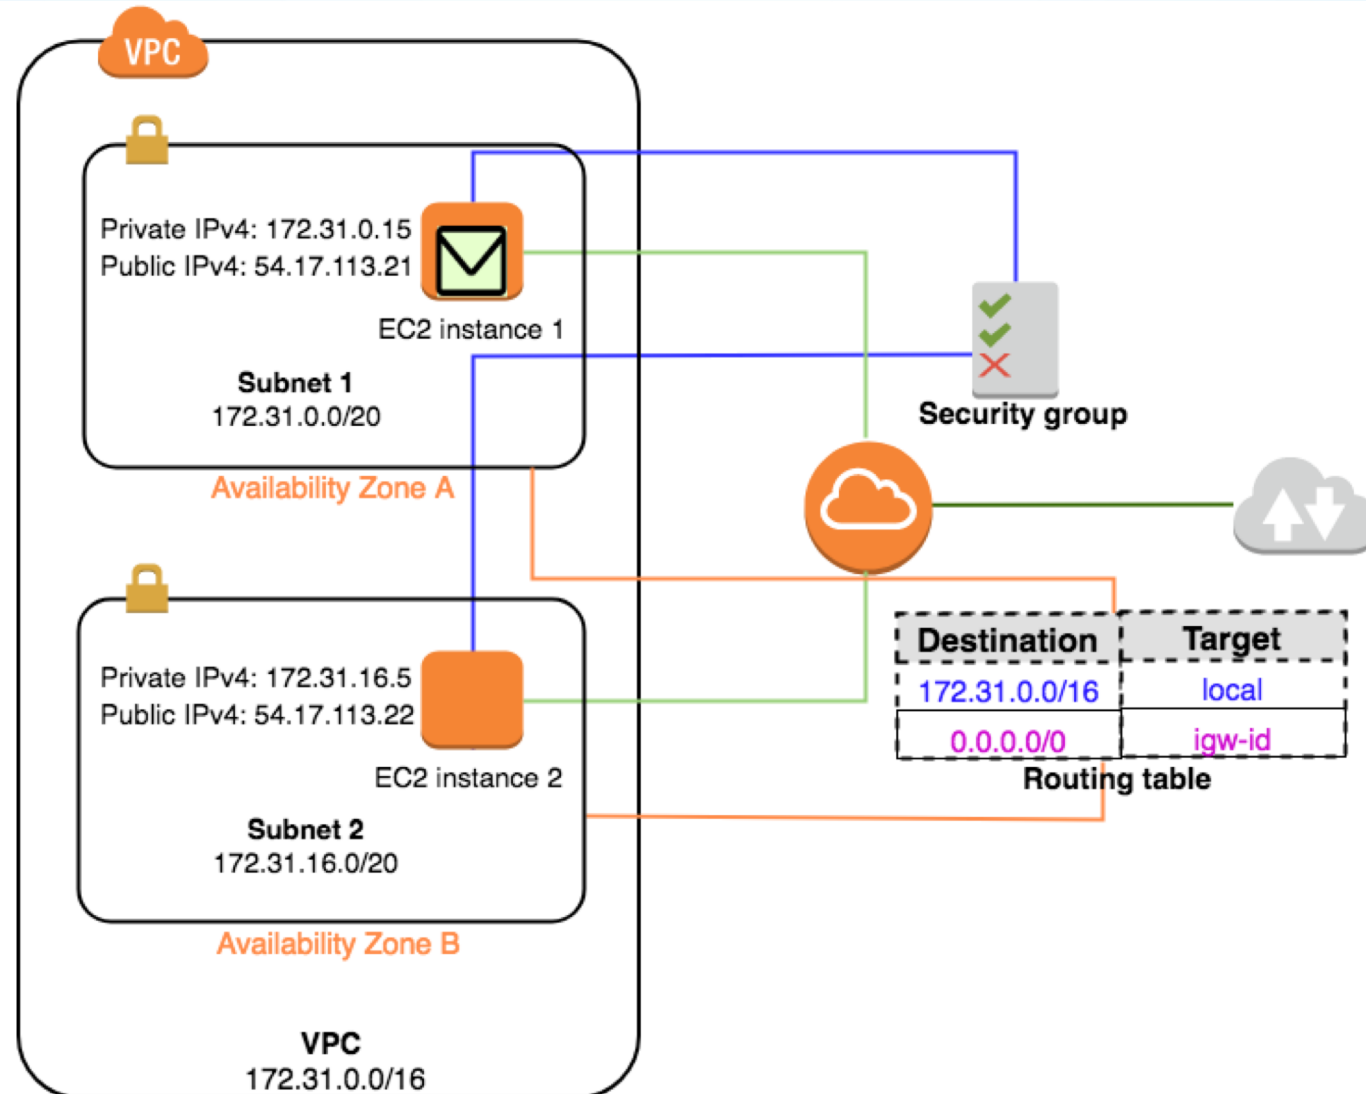


KubeCon



CloudNativeCon

Europe 2018



# Packet flow – EC2 instances

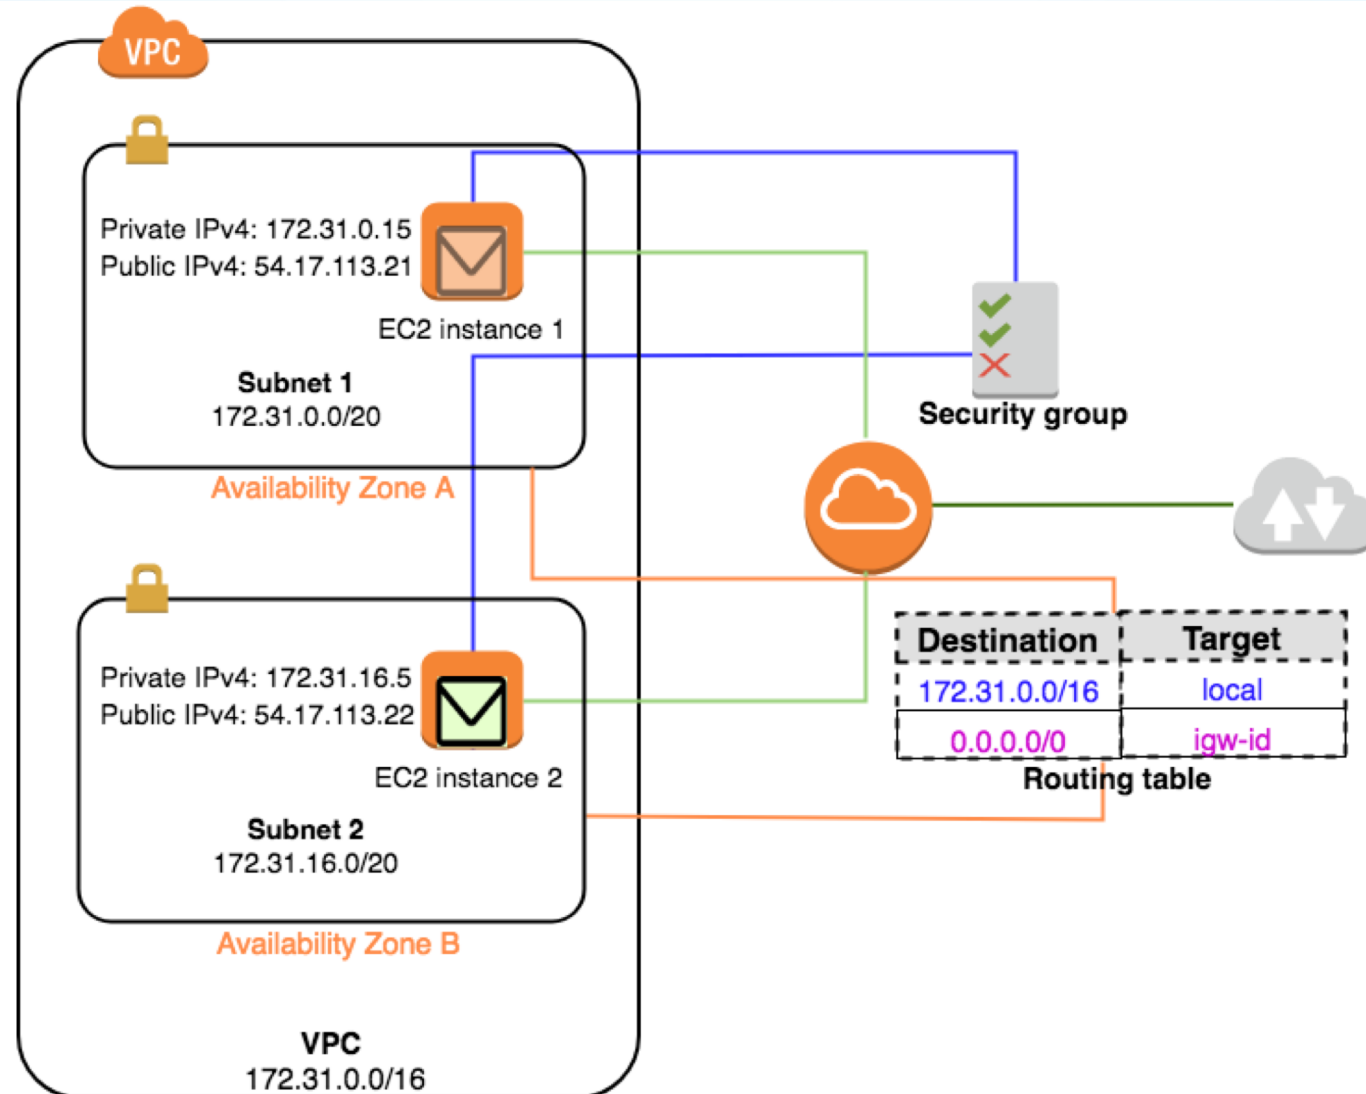


KubeCon



CloudNativeCon

Europe 2018



# VPC networking - ENIs

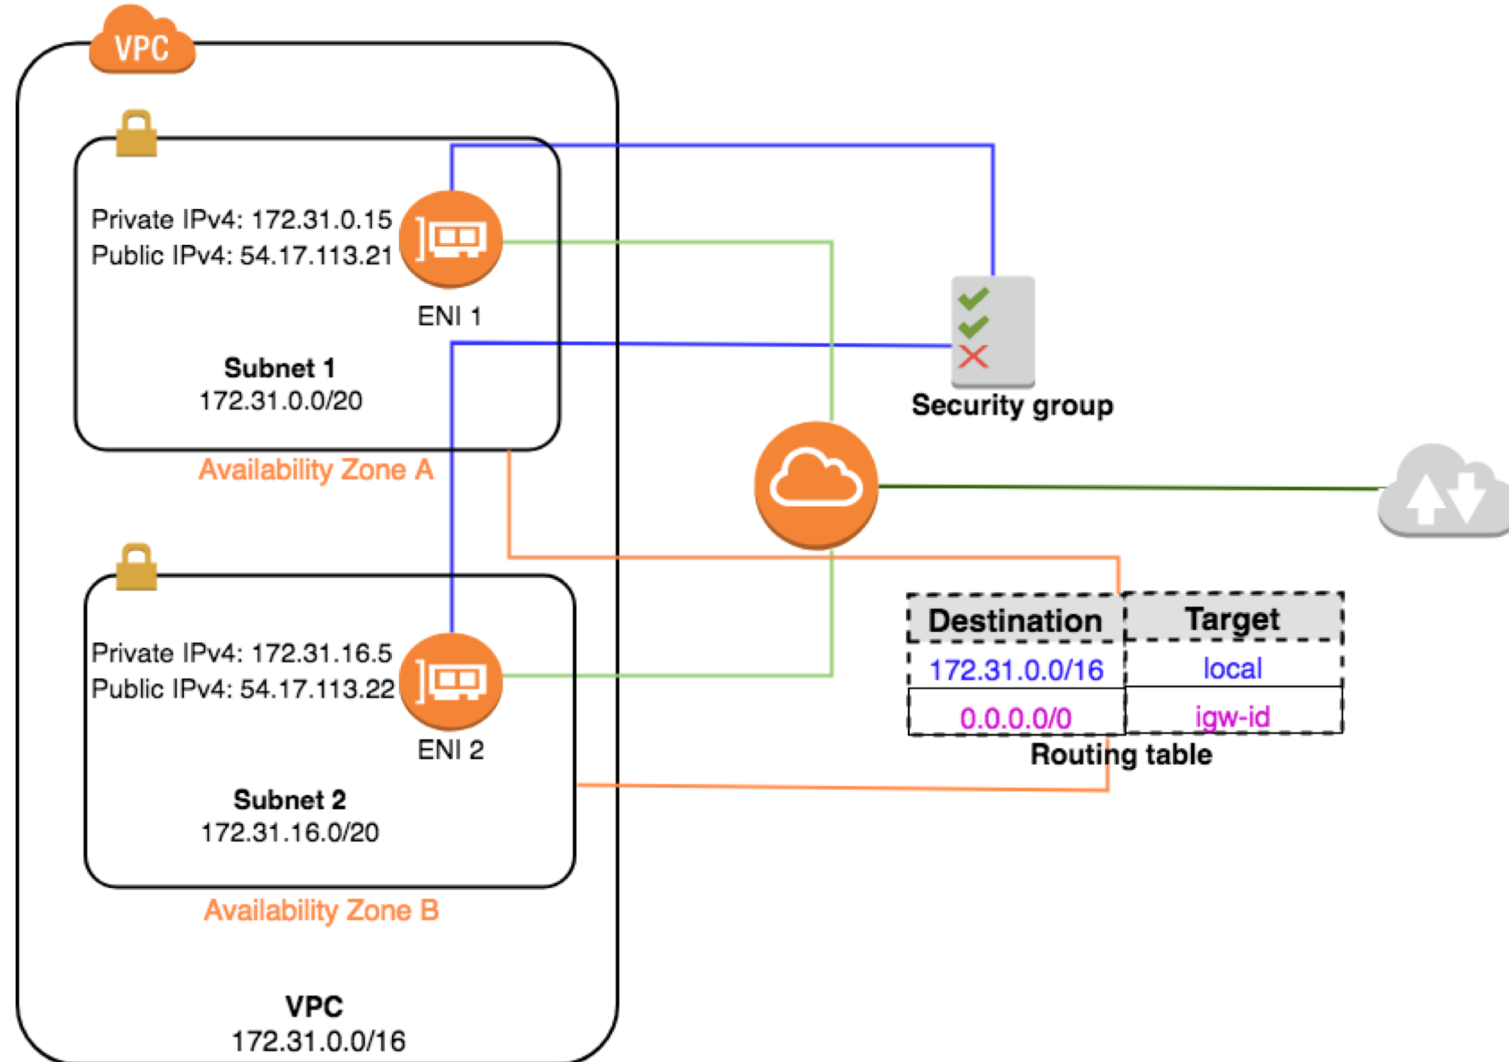


KubeCon



CloudNativeCon

Europe 2018





# ENIs for tasks & pods




KubeCon



CloudNativeCon

Europe 2018



## Amazon ECS Introduces Task Networking for Containers

Posted On: Nov 14, 2017

Tasks running on Amazon EC2 Container Service (Amazon ECS) can now take advantage of *awsvpc* mode for container networking. This new mode allocates an [elastic networking interface](#) to each running task, providing a dynamic private IP address and internal DNS name. This simplifies container networking operations, allowing tasks to run with full networking features on AWS, just like EC2 instances.

[aws / amazon-ecs-cni-plugins](#)

[Code](#) [Issues 9](#) [Pull requests 2](#) [Projects 0](#)

Networking Plugins repository for ECS Task Networking

[cni-plugin](#) [Manage topics](#)

[aws / amazon-vpc-cni-k8s](#)

[Code](#) [Issues 25](#) [Pull requests 8](#)

Networking plugin repository for pod networking i

# VPC networking – containers

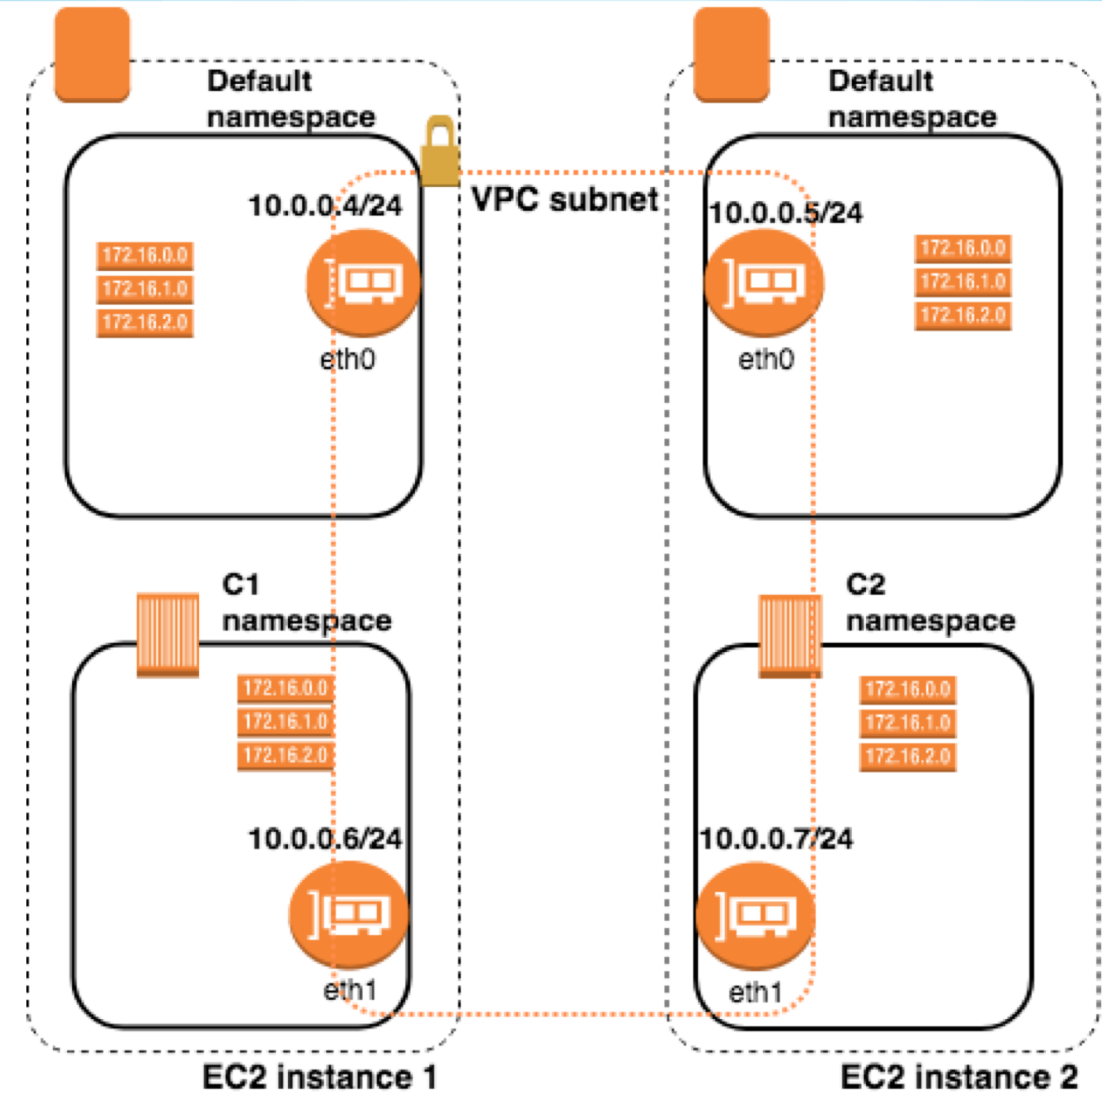


KubeCon



CloudNativeCon

Europe 2018



# Packet flow

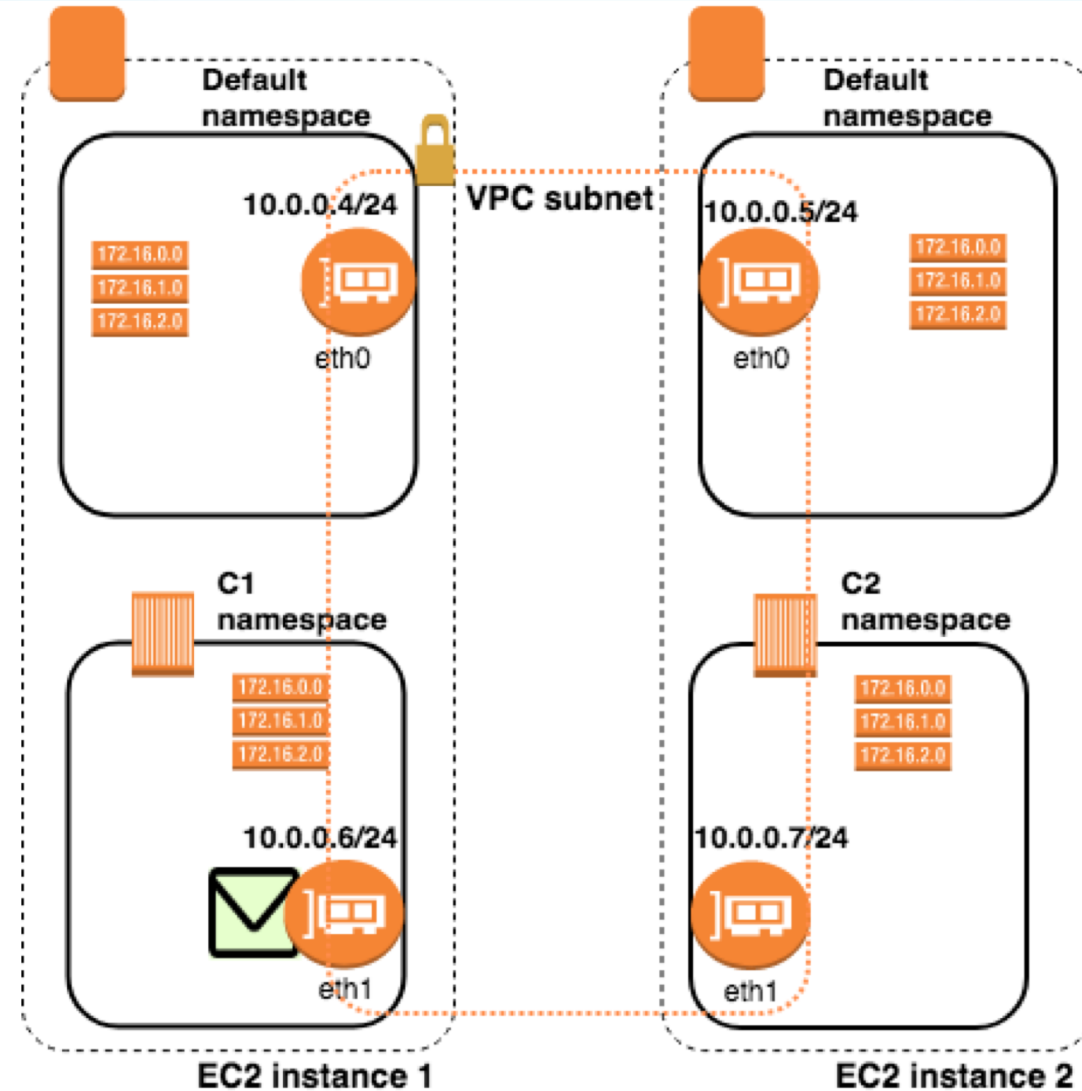


KubeCon



CloudNativeCon

Europe 2018



# Packet flow

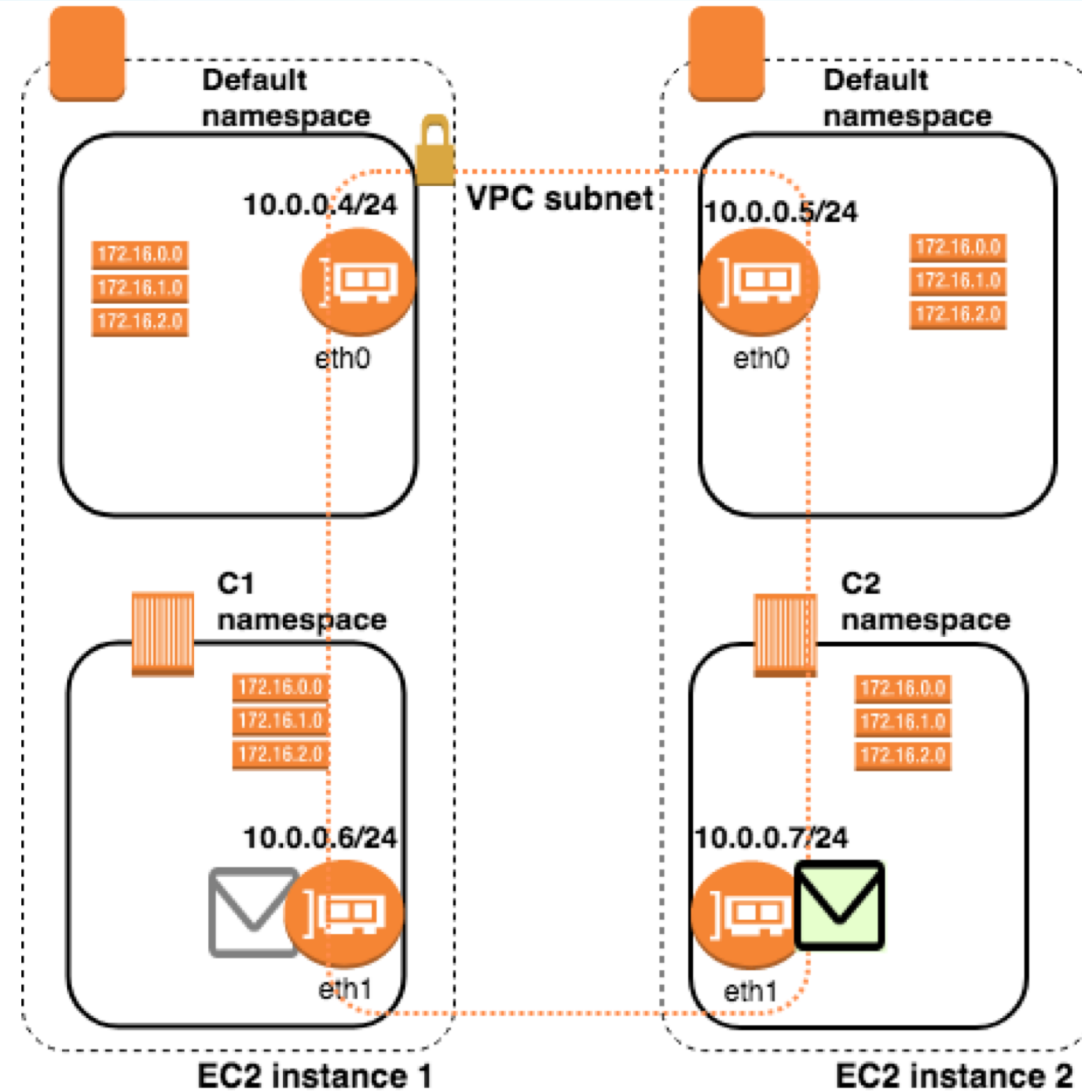


KubeCon



CloudNativeCon

Europe 2018






KubeCon



CloudNativeCon

Europe 2018



# CNI



# CNI or CNM ?



KubeCon



CloudNativeCon

Europe 2018

- Maintainability
  - Minimal intrusion to container life-cycle
  - Life-cycle management
  - Rolling out updates
- Simplicity
  - Consistent & reliable interface
  - Testability
- Extensibility
- Ecosystem support



# CNI plugins 101



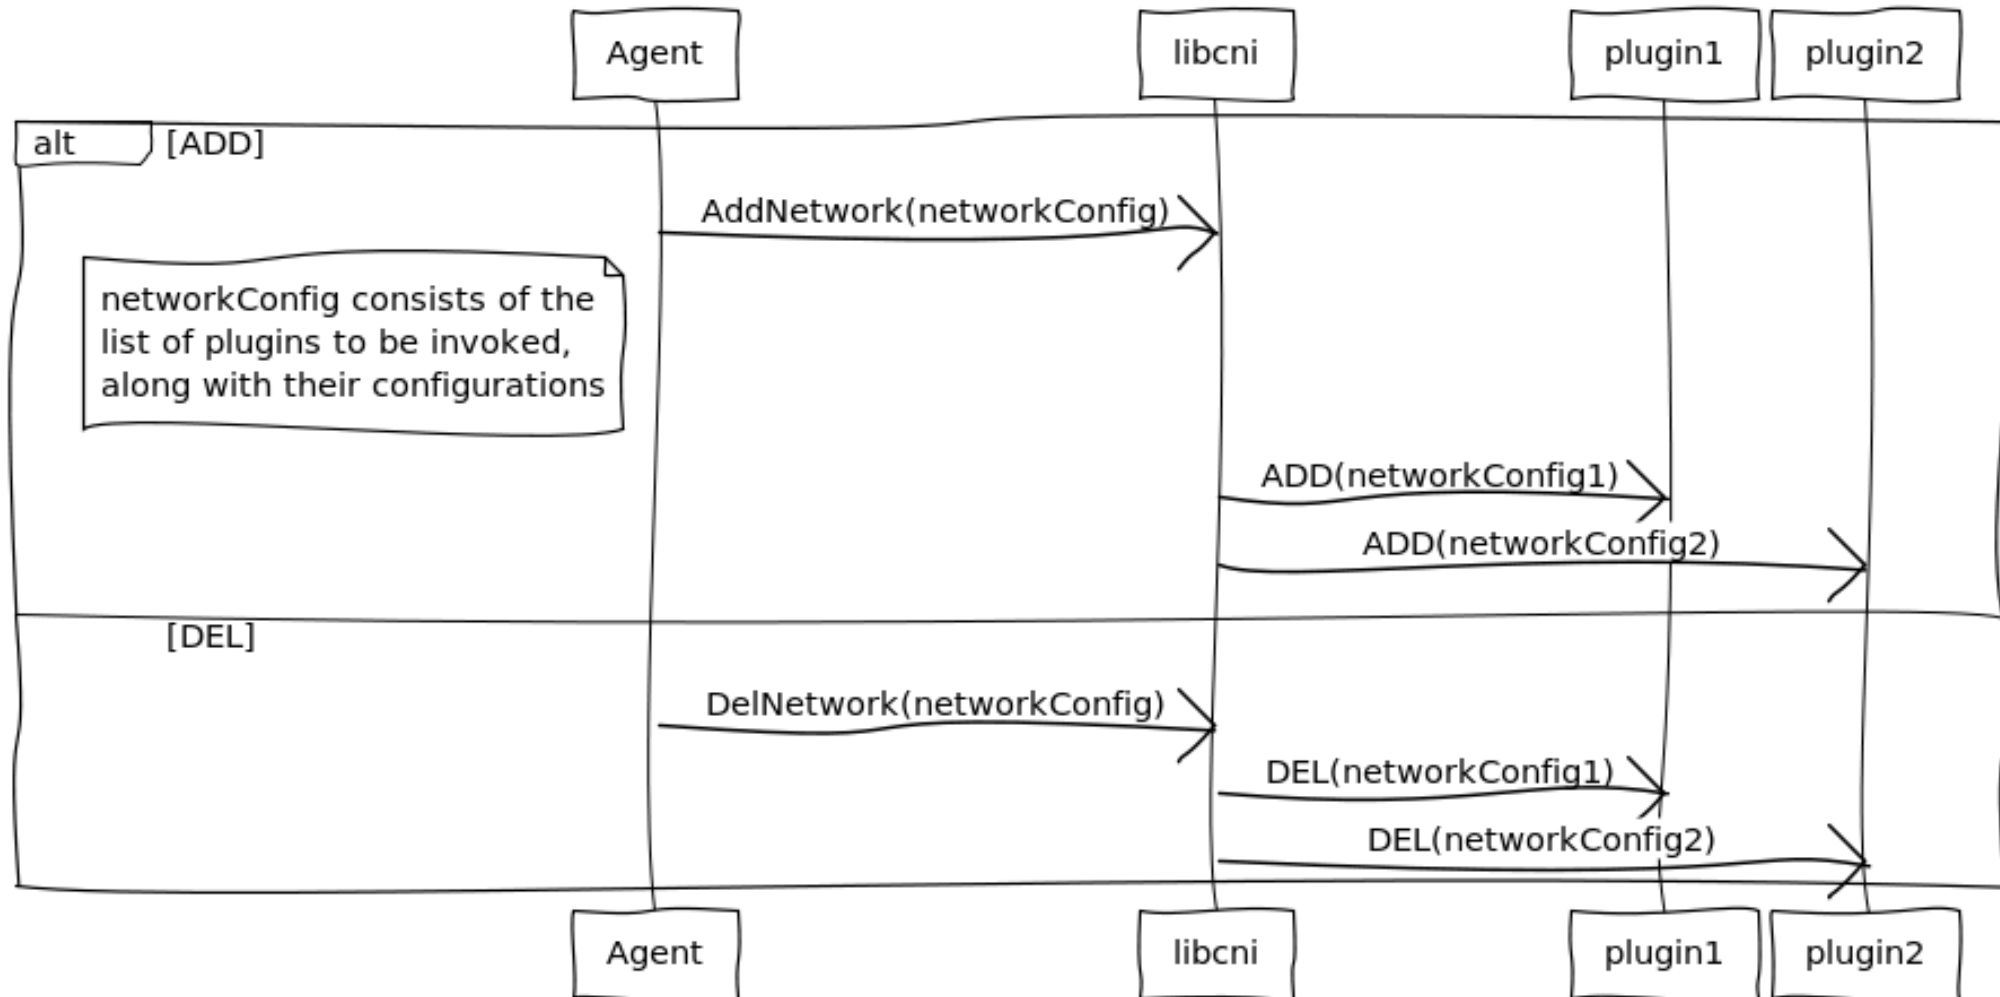
KubeCon



CloudNativeCon

Europe 2018

## CNI plugin invocation



# golang or ...?




KubeCon



CloudNativeCon

Europe 2018

- Static binary
- Ecosystem support



# Packaging & distribution




KubeCon



CloudNativeCon

Europe 2018

Branch: master [amazon-ecs-agent](#) / [scripts](#) / [dockerfiles](#) / [Dockerfile.release](#)

 vsiddharth CNI Plugins Packaging: Address review comments

4 contributors 

32 lines (25 sloc) | 1.13 KB

```
1 # Copyright 2014-2017 Amazon.com, Inc. or its affiliates. All Rights Reserved.
2 #
3 # Licensed under the Apache License, Version 2.0 (the "License"). You may
4 # not use this file except in compliance with the License. A copy of the
5 # License is located at
6 #
7 #     http://aws.amazon.com/apache2.0/
8 #
9 # or in the "license" file accompanying this file. This file is distributed
10 # on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either
11 # express or implied. See the License for the specific language governing
12 # permissions and limitations under the License.
13
14 # Not from scratch because we also want a little directory structure,
15 # specifically /tmp
16 FROM amazon/amazon-ecs-scratch:make
17
18 COPY out/amazon-ecs-agent /agent
19 COPY ["LICENSE", "NOTICE", "/"]
20
21 COPY out/amazon-ecs-pause.tar /images/amazon-ecs-pause.tar
22
23 # Copy our cni plugins ecs-eni, ecs-ipam and ecs-bridge
24 COPY out/cni-plugins /amazon-ecs-cni-plugins
25
26 # Copy our bundled certs to the first place go will check: see
27 # https://golang.org/src/pkg/crypto/x509/root_unix.go
28 COPY misc/certs/ca-certificates.crt /etc/ssl/certs/ca-certificates.crt
29
30 EXPOSE 51678 51679
31 ENTRYPOINT ["/agent"]
```

# Versioning



KubeCon



CloudNativeCon

Europe 2018

```
{  
  "version": "2017.06.0",  
  "dirty": false,  
  "gitShortHash": "226db3"  
}
```

# Automating version info



KubeCon



CloudNativeCon

Europe 2018

```
.PHONY: plugins
```

```
plugins: $(LOCAL_ENI_PLUGIN_BINARY) $(LOCAL_IPAM_PLUGIN_BINARY) $(LOCAL_BRIDGE_PLUGIN_BINARY)
```

```
$(LOCAL_ENI_PLUGIN_BINARY): $(SOURCES)
```

```
GOOS=linux CGO_ENABLED=0 go build -installsuffix cgo -a -ldflags "\
```

```
  -X github.com/aws/amazon-ecs-cni-plugins/pkg/version.GitShortHash=$(GIT_SHORT_HASH) \
```

```
  -X github.com/aws/amazon-ecs-cni-plugins/pkg/version.GitPorcelain=$(GIT_PORCELAIN) \
```

```
  -X github.com/aws/amazon-ecs-cni-plugins/pkg/version.Version=$(VERSION) -s" \
```

```
  -o ${ROOT}/${LOCAL_ENI_PLUGIN_BINARY} github.com/aws/amazon-ecs-cni-plugins/plugins/eni
```

```
@echo "Built eni plugin"
```

# Testing the plugin



KubeCon



CloudNativeCon

Europe 2018

```
.PHONY: unit-test integration-test e2e-test
```

```
unit-test: $(SOURCES)
```

```
    go test -v -cover -race -timeout 10s ./pkg/... ./plugins/...
```

```
integration-test: $(SOURCE)
```

```
    go test -v -tags integration -race -timeout 10s ./pkg/... ./plugins/...
```

```
e2e-test: $(SOURCE) plugins
```

```
    sudo -E CNI_PATH=${ROOT}/bin/plugins ${GO_EXECUTABLE} test -v -tags e2e -race -timeout 120s ./plugins/...
```

# In conclusion ...



KubeCon



CloudNativeCon

Europe 2018

- Avoid/minimize feature envy, especially for networking
- CNI plugins ftw!
- Version everything (with git SHAs)

# Related links



KubeCon



CloudNativeCon

Europe 2018

<https://github.com/aws/amazon-ecs-cni-plugins>

<https://github.com/aws/amazon-vpc-cni-k8s/>

<https://github.com/vishvananda/netns>

<https://github.com/vishvananda/netlink>

<https://github.com/containernetworking/cni>

<https://github.com/containernetworking/plugins>





KubeCon



CloudNativeCon

Europe 2018

# Thank you!

<https://github.com/aws/amazon-ecs-cni-plugins/>

<https://github.com/aws/amazon-vpc-cni-k8s/>

@aithal

