



Containerized VNFs in ONAP

Cloud-Native Lifecycle Management in Kubernetes

Tal Liron
Principal Software Engineer
NFV Partner Engineering

ONAP Casablanca Release Developer Forum
Beijing, June 19-22, 2018

Arguments

1. Kubernetes is not a VIM for containers

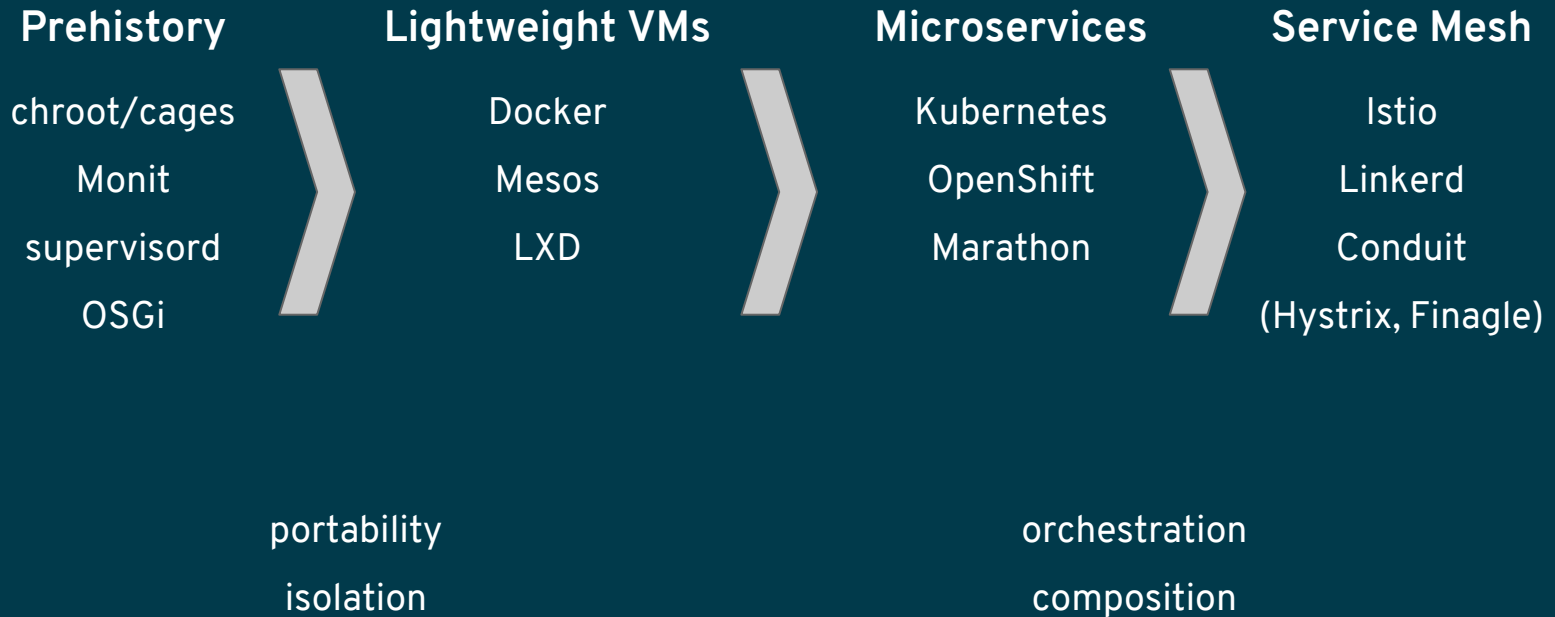
If that's what you want, check out Canonical's LXD

2. VNFs must manage their own lifecycle

Because:

1. they know best, and
2. clustering is very often domain-specific

Containers != Containers



Clustering Is Tricky

1. Cluster != Replicas

You can easily create 10 instances of a pod, but that doesn't mean they will magically work together

2. Control plane != data plane

Every pod has a single IP interface

Sometimes that's fine

("piggybacking" on the control plane)

Cloud-Native Lifecycle Management

1. Dynamic Kubernetes

The running application itself will modify its Kubernetes resources (Deployment, Service, etc.)

2. Triggered by Events

Some can be internal: e.g. high load, must scale

Some can be external: e.g. policy has changed

Introducing: “Operators”

1. Dedicated pod

Listens for events and modifies Kubernetes resources

2. Custom Resource

Domain-specific representation of the cluster

Stateful Example: Cassandra

<https://github.com/operator-framework/awesome-operators>

1. Clustering

Resize, duplicate, recover, load balancing

2. Admin

Backup, rolling upgrade, caching

VNF Example: Clearwater IMS

1. Clustering

Several “Sprout” nodes (SIP router)

Plus a single “Bono” node (SIP edge proxy)

2. Data plane != control plane

Each of these nodes needs its own IP address and needs to know about the others



THANK YOU



plus.google.com/+RedHat



facebook.com/redhatinc



linkedin.com/company/red-hat



twitter.com/RedHat



youtube.com/user/RedHatVideos