



THE LINUX FOUNDATION
OPEN SOURCE SUMMIT



THE LINUX FOUNDATION
OPEN SOURCE SUMMIT

Setting up Kubernetes with Day 2 in Mind

Angela Chin, Senior Software Engineer, Pivotal
Urvashi Reddy, Senior Software Engineer, Pivotal

About Us



Angela

Software Engineer @ Pivotal
Based in Santa Monica, CA
Tell me the best hikes and
restaurants!



Urvashi

Software Engineer @ Pivotal
Based in San Francisco, CA
Tell me where the best
desserts are at!

Agenda

- Power of Kubernetes
- Managing Kubernetes
- BOSH
- CFCR
- Recent Work
- Summary
- What's Next?
- Resources



THE LINUX FOUNDATION
OPEN SOURCE SUMMIT

Power of Kubernetes

What is k8s?

- Kubernetes
- Container orchestrator built around Docker

Docker

- Allows standardization of OS and binaries
- Easily reproducible due to Dockerfiles
- Anyone can now harness the power of containers!



THE LINUX FOUNDATION
OPEN SOURCE SUMMIT

But...

Managing many containers can be a challenge

What is k8s?

- Kubernetes
- **Container orchestrator** built around Docker

What is k8s?

- Kubernetes
- **Container orchestrator** built around Docker
- “Have a datacenter on their laptop”

What is k8s?

- Kubernetes
- **Container orchestrator** built around Docker
- “Have a datacenter on their laptop”
- Dev environment == production environment

Benefits of k8s

- Multi-cloud: AWS, GCP, Azure, vSphere, etc.
- Community convergence & 3rd party integrations

3rd Party Integrations

- Prometheus: monitoring system
- Zipkin: distributed tracing system
- CNI plugins (calico, flannel, rkt): networking
- Istio & Linkerd: service meshes

Benefits of k8s

- Multi-cloud: AWS, GCP, Azure, vSphere, etc.
- Community convergence & 3rd party integrations
- Flexibility



THE LINUX FOUNDATION
OPEN SOURCE SUMMIT

Managing Kubernetes

Tracking features is a full time job

- Minors shipped every ~100 days
- Alpha vs Beta vs Stable features

Alpha, Beta, and Stable

	Availability	Upgradeability	Audience
Alpha	Disabled by default	Can be backwards incompatible	Early feedback
Beta			
Stable			

Alpha, Beta, and Stable

	Availability	Upgradeability	Audience
Alpha	Disabled by default	Can be backwards incompatible	Early feedback
Beta	Enabled by default	May be downtime	Early feedback
Stable			

Alpha, Beta, and Stable

	Availability	Upgradeability	Audience
Alpha	Disabled by default	Can be backwards incompatible	Early feedback
Beta	Enabled by default	May be downtime	Early feedback
Stable	Enabled by default	Possible	Everyone

Tracking features is a full time job

- A lot of features to keep track of
 - Over 31 notable changes in the last 1.11 release
- Highly developed project results in unforeseen impacts

Security

- Many entry points for things to go wrong
- Intra cluster TLS is still a work in progress
- CVE Fixes

Upgrading Clusters

- Address vulnerabilities
- Manage diverse workloads
- Upgrade Checklist
 - Gracefully drain nodes
 - Consider shutting down the control plane
 - Deploy the upgrade
 - Is the cluster operational?
 - Plan for the worst case scenario



THE LINUX FOUNDATION
OPEN SOURCE SUMMIT

BOSH

oh my BOSH!

BOSH is a project that unifies release engineering, deployment, and lifecycle management of small and large-scale cloud software. BOSH can provision and deploy software over hundreds of VMs. It also performs monitoring, failure recovery, and software updates with zero-to-minimal downtime.

oh my BOSH!

BOSH is a project that **unifies release engineering, deployment, and lifecycle management** of small and large-scale cloud software. BOSH can provision and deploy software over hundreds of VMs. It also performs monitoring, failure recovery, and software updates with zero-to-minimal downtime.

oh my BOSH!

BOSH is a project that **unifies release engineering, deployment, and lifecycle management** of small and large-scale cloud software. BOSH can provision and deploy software over hundreds of VMs. It also performs **monitoring, failure recovery, and software updates with zero-to-minimal downtime.**

BOSH

- Day 2 Operations
- Large-scale cloud software

IAAS Support

- vSphere
- AWS
- GCP
- Azure
- OpenStack

BOSH

- Day 2 Operations
- Large-scale cloud software
- Security fixes within 48 hours

BOSH

- Day 2 Operations
- Large-scale cloud software
- Security fixes within 48 hours
- Security aware

BOSH

- Day 2 Operations
- Large-scale cloud software
- Security fixes within 48 hours
- Security aware
- Monitors and resurrects VMs

BOSH

- Day 2 Operations
- Large-scale cloud software
- Security fixes within 48 hours
- Security aware
- Monitors and resurrects VMs
- Reproducible deployments

BOSH

- Day 2 Operations
- Large-scale cloud software
- Security fixes within 48 hours
- Security aware
- Monitors and resurrects VMs
- Reproducible deployments

and...

8 years old!

Proven with Cloud Foundry!!



THE LINUX FOUNDATION
OPEN SOURCE SUMMIT

Cloud Foundry Container Runtime (CFCR)

What is CFCR?

"CFCR is a BOSH release that offers a uniform way to instantiate, deploy, and manage highly available Kubernetes clusters on a cloud platform using BOSH"

What is CFCR?

- Manages Kubernetes packages & dependencies
 - Up to date with the k8s features
- Configurable
- Tested

CFCR: Leverages BOSH

- Reliable upgrades
- Self-healing VMs
- Certificate and Credential Management
- BOSH DNS

CFCR: Tooling

- Smoke Tests
- Tooling for targeting Clusters



THE LINUX FOUNDATION
OPEN SOURCE SUMMIT

Recent Work

What we're working on now?

- Highly available clusters with multiple masters
- The ability to backup and restore a cluster



THE LINUX FOUNDATION
OPEN SOURCE SUMMIT

What's Next?

Future Plans

- Support for Azure Cloud Provider
- Hardening Security
- Pluggable CNI



THE LINUX FOUNDATION
OPEN SOURCE SUMMIT

But...we really want to hear
what you think should be
next



THE LINUX FOUNDATION
OPEN SOURCE SUMMIT

Summary

Summary

- k8s is awesome! ...but hard to manage
- BOSH + CF CR provides a way to ensure Day 2 operations for your k8s cluster

Resources

[BOSH Docs](#)

[CFCR Docs](#)

[CFCR Repo](#)

#cfcf @ cloudfoundry.slack.com



THE LINUX FOUNDATION
OPEN SOURCE SUMMIT

Thank you!
achin@pivotal.io
ureddy@pivotal.io