



DELIVERING OPENSTACK AND CEPH IN CONTAINERS

OpenStack Summit Sydney — November 7, 2017

The Three Amigos



**Andrew
Hatfield**

Practice Lead
Cloud Storage
Brisbane, Australia



**Federico
Lucifredi**

Product Management Director
Red Hat Ceph Storage
Boston, USA




**Sébastien
Han**

Principal Software Engineer
Storage Architect
Paris, France


Where we are today

Drivers


Why are we moving toward containers?




Packaging
format



Scalability




Speed




Upgrade
downgrade
flexibility



Immutable
infrastructure



Deployment
flexibility



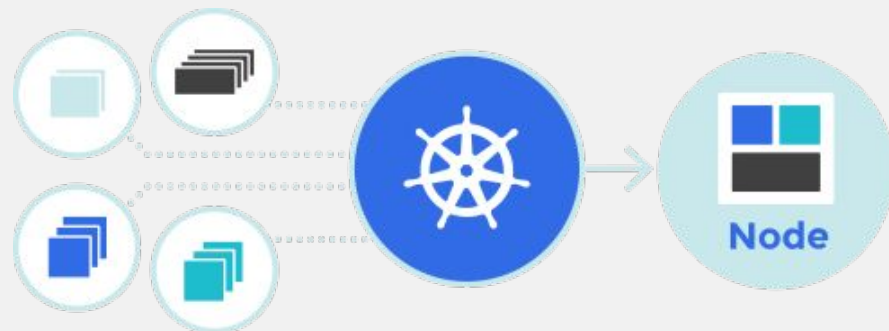
Resource
constraints

Entering Kubernetes with OpenStack

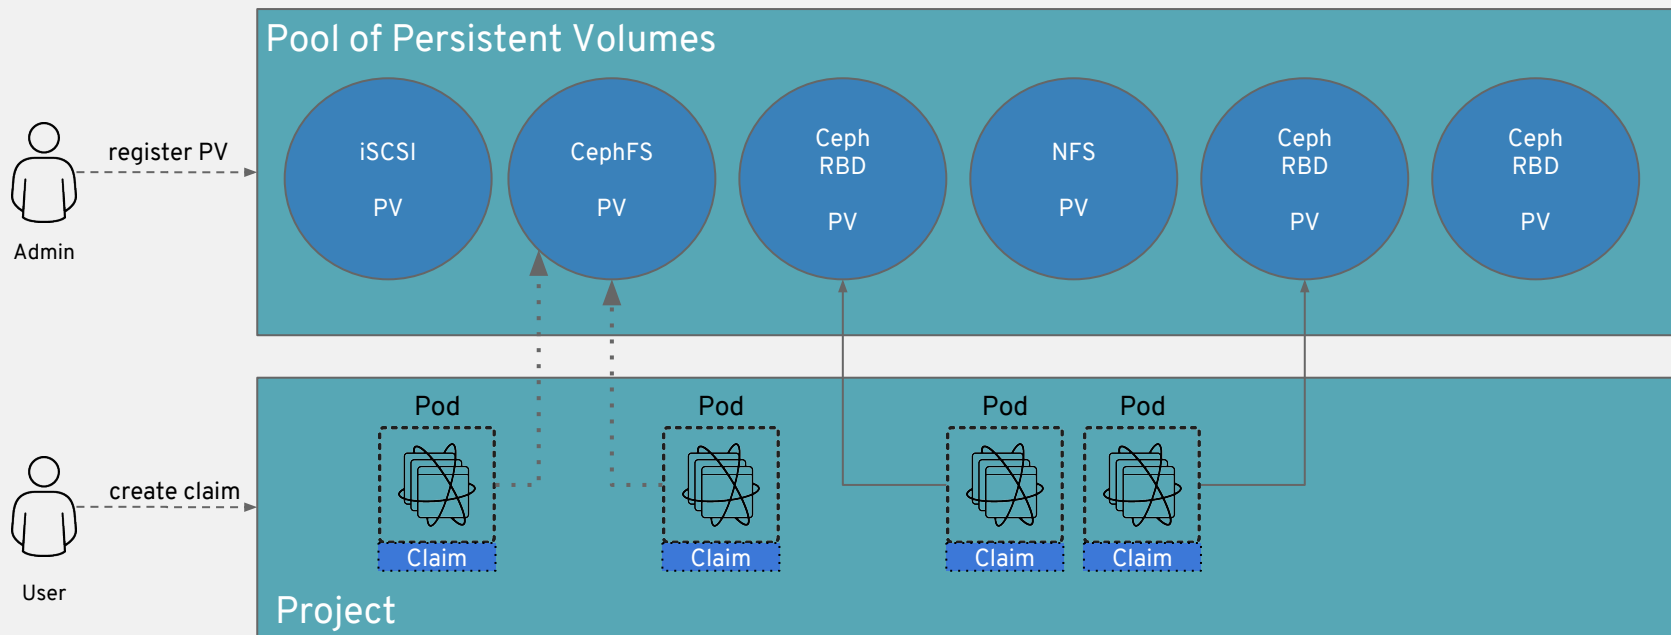
Kubernetes is an open-source orchestrator automating the deployment, scaling, and management of containerized applications.

Container management platform:

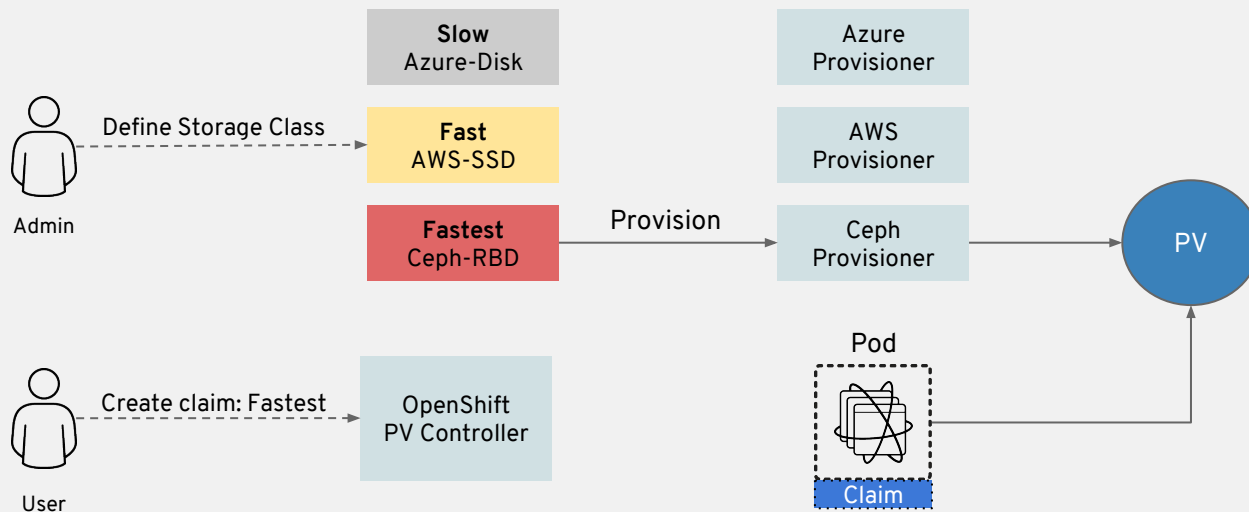
- Self-healing
- Load balancing
- Automated rollouts and rollbacks
- Pluggable architecture:
 - Storage
 - Network
 - Containers runtime
 - Scheduling



Persistent Storage



Dynamic Voluming Provisioning



OpenStack in Containers


Why OpenStack on Kubernetes?

Day 1: ease deployment

- OpenStack itself turned into microservice oriented architecture with its services
- Kubernetes manages applications
- Kubernetes load, driven by application needs, can vary dramatically over wide time scales (hours, days, weeks)
- Leverage the ease, scale and power of the Containers

Day 2:

- Share scheduling functions between Nova and Kubernetes



openstack®
CLOUD SOFTWARE

Kolla

Some background..

Let's go back in time again:

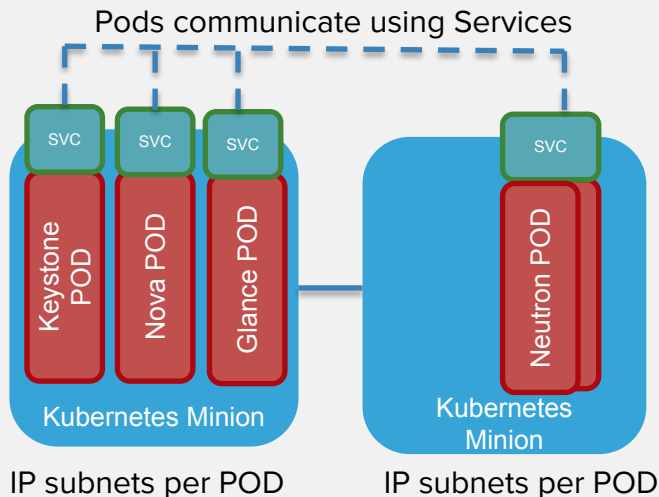
- Started in September 2014 by Red Hat
- Announced and advertised in November 2014 during Paris OpenStack Summit
- Kolla != Magnum
- Deployment methods
 - Mesos, started in October 2015 and abandoned as an orchestration scheduler by Kolla in April 2016
 - Ansible is the de facto tool
 - Kubernetes with the kolla-kubernetes project (started in May 2016)

Kolla

Installation of OpenStack services as set of containers

- Solves a **manageability** and **availability** problem with the current state of the art deployment systems in OpenStack

- Containerizing OpenStack is meant to **optimize image-based management** of OpenStack
- Uses Heat Templates and YAML to define services and pods



- Deploy OpenStack using container technology **for atomic** upgrades in seconds

<https://launchpad.net/kolla>

Kuryr and Fuxi

Exposing the power of the cloud to containers

- Taking advantage of the pluggable nature of Kubernetes
- Network and storage “interfaces” to interact with your Cloud
- Enhance the capabilities of Kubernetes by exposing more:
 - Network drivers from Neutron - Kuryr
 - Storage drivers from Cinder - Fuxi

Ceph in Containers

Deploying storage services with Kolla

Kolla support the containerization of the following storage services already today:

- Glance
- Cinder
- Swift
 - Kolla can deploy a full working Swift setup in either a **all-in-one** or **multinode** setup.
- Manila
 - Currently deploys manila-api, manila-scheduler and Manila-share
- Ceph
 - The out-of-the-box Ceph deployment requires 3 hosts with at least one block device each that can be dedicated for sole use by Ceph
 - Kolla also support external Ceph clusters by simply disabling Ceph deployment in `/etc/kolla/global.yml`
 - With tweaks to the Ceph cluster you can deploy a functional cluster with a single host and a single block device (this configuration provides no data resiliency).

ceph-container

Containerizing Ceph daemons

The project:

- Launched on Jan 18, 2015
- Upstream project: <https://github.com/ceph/ceph-container>
- Support from Jewel to the latest version of Ceph (currently Luminous)
- Distros: Ubuntu 16.04 and CentOS 7
- Automated builds on Docker Hub
- More than 5M+ pulls!



ceph/daemon
public | automated build

58
STARS

5M+
PULLS

>
DETAILS

Single daemon image

A single **ceph/daemon** container image is used to bootstrap a Ceph cluster with all its daemons running.

- To deploy MON, MGR, OSD, MDS, RGW, rbd-mirror and NFS
- Large variety of OSD scenarios:
 - Journal collocation
 - Dedicated device for journal (Filestore) and DB/WAL (Bluestore)
 - Encrypted
 - Filestore
 - Bluestore
 - Directory
- Deployment methods:
 - Ceph-ansible: containers are managed by systemd
 - Kubernetes: ceph-helm: <https://github.com/ceph/ceph-helm>

Persistent container storage

A Kubernetes cluster provides **persistent storage** to containers via the **persistent volume** abstraction.

- Persistent Volume
 - Abstracts underlying storage technology
- Persistent Volume Claim
 - Identifies a specific volume to container pods
- Storage class
 - The storage driver proper
 - Enables dynamic provisioning

Ceph


- | | |
|--|---|
| <ul style="list-style-type: none">• Today<ul style="list-style-type: none">◦ Ceph RBD◦ Cinder | <ul style="list-style-type: none">• Future<ul style="list-style-type: none">◦ CephFS◦ Manila |
|--|---|



ceph-nano

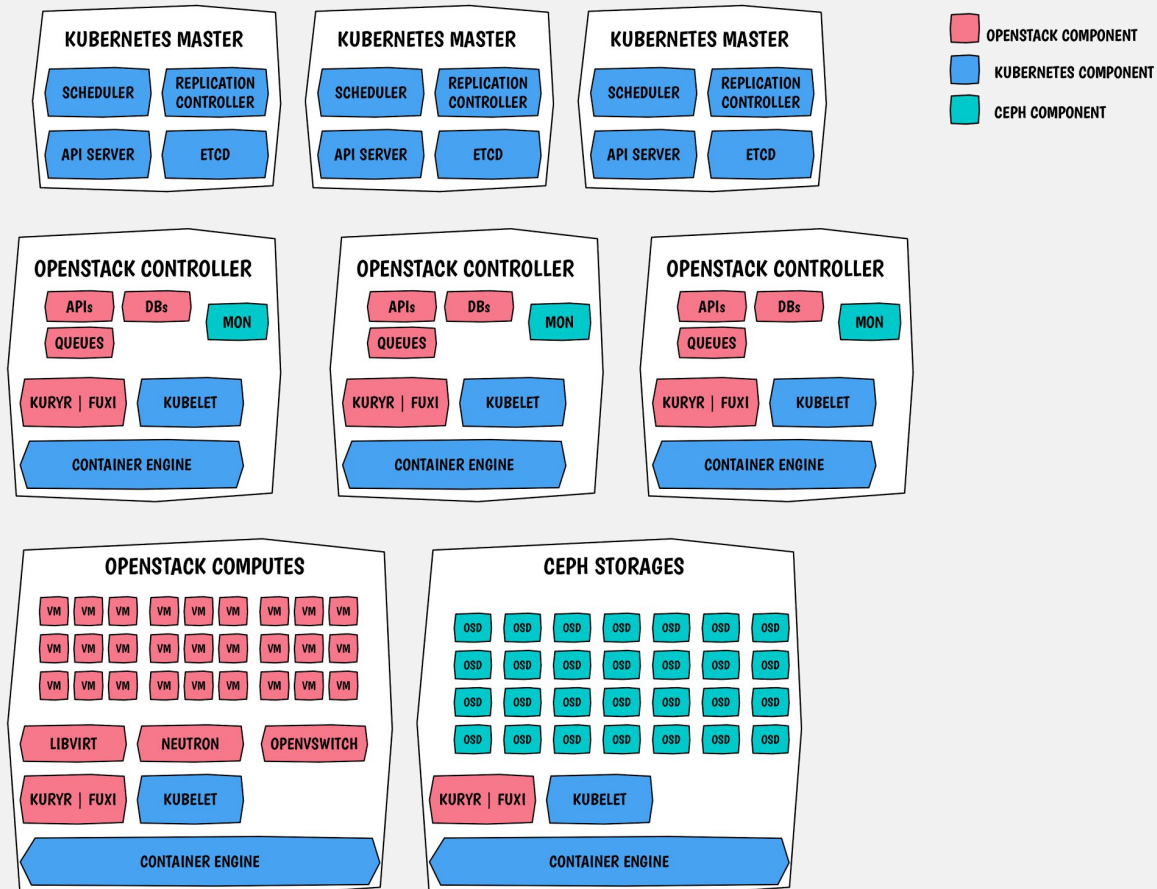
One step S3 in Container

- Runs all Ceph daemons into a single container
- Simple experience
- Currently exposes S3 capabilities
- Persistent (supports reboot)
- <https://ceph.github.io/ceph-container/>



Architecture use cases


CONTAINERIZED OPENSTACK CLOUD



Hyperconverged

Hypercon... what?

HYPERCONVERGED NODE IN-DEPTH



Deployment tools

TripleO


OpenStack standard with Heat



- Use Heat hook that allow Heat to orchestrate container deployment
- Result:
 - You can create a containerized OpenStack solution by using TripleO
 - Using the containers from the Kolla project
 - Ansible is available and ready today! TripleO uses it via Mistral workflows.

Ansible

The de-facto standard



Main purpose:

- Friendly learning curve
- “Flat” deployment, no lifecycle
- No placement intelligence or orchestration
- More traditional users, not kubernetes-ready
- Let systemd manage your containers (start/stop/restart/watchdog)

Kubernetes

The future

- Are you ready for containers?
- Nice orchestration / application's life cycle
- Ceph-helm
- Not fully compliant yet (see next slide)



Challenges - Gap analysis

Component	Requirement	Met
Network	Network Isolation	WIP
Network	Ability to disable network overlays and use host's network	Yes
Network	IPv6	Yes
Network	SSL terminaison	Yes
Storage	Data persistency (outside of containers)	Yes
HA	Cluster bootstrapping (boot order, node replicas, etc)	Yes
HA	Pods Monitoring	Yes
HA	Load balancing strategies	Yes
HA	Pod Fencing	WIP

What's next?

Ceph container roadmap

The road ahead

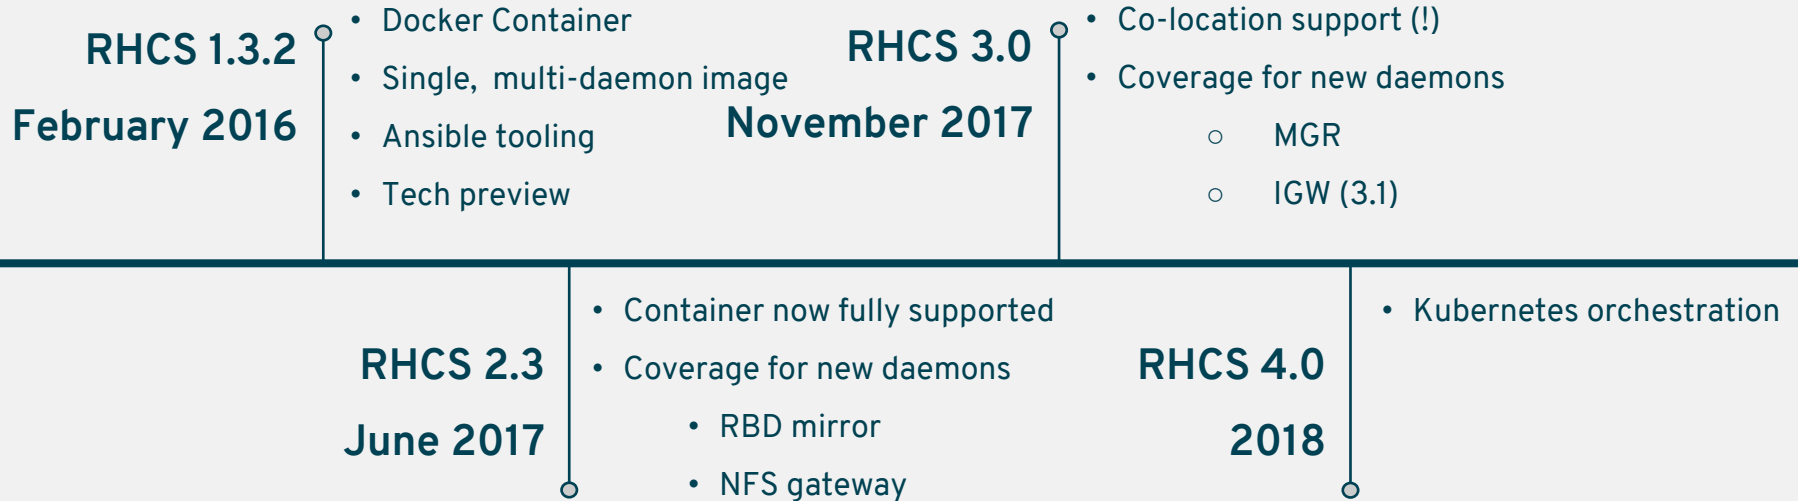
Done:

- Ceph-container
 - container images available for two years (Hammer)
- First stable release in October 2017 (Luminous, Jewel)
 - tag-stable-3.0-luminous-centos-7
 - tag-stable-3.0-luminous-ubuntu-16.04

Next:

- Kubernetes prototyping underway in the Mimic cycle
- Kolla to use Ceph container images

RHCS container roadmap



Summary

Takeaways:

- Ansible + Kolla are excellent candidates to deploy containerized infrastructures
- Support for OpenStack and Ceph in containers is here
- Kubernetes is the future but not fully there yet
- Getting the right networking model/setup is difficult but it's critical

Wanna hear more?

Resources

Resources:

- [Ceph Container](#) on Docker Hub
- [Ceph Container](#) on GitHub
- [Ceph Ansible](#)
- [Manila in Kolla](#)
- [Swift in Kolla](#)
- [Cinder in Kolla](#)
- [Containerized Ceph from Red hat and QCT](#)



@sebastien_han
@andrewhatfield
@0xF2

THANK YOU



plus.google.com/+RedHat



facebook.com/redhatinc



linkedin.com/company/red-hat



twitter.com/RedHatNews



youtube.com/user/RedHatVideos

Backup Slides (mwahahaha!)

BREAKOUT SESSIONS - Monday, November 6th

Panel: experiences scaling file storage with CephFS and OpenStack

Gregory Farnum, Sage Weil, Patrick Donnelly, and Arne Wiebalck (CERN)

1:30pm - 2:10pm

Multicloud requirements and implementations: from users, developers, service providers

Mark McLoughlin, Jay Pipes (Mirantis), Kurt Garloff (T-Systems International GmbH), Anni Lai (Huawei), and Tim Bell (CERN)

2:20pm - 3:00pm

BREAKOUT SESSIONS - Monday, November 6th

CephFS: now fully awesome (what is the impact of CephFS on the OpenStack cloud)

Andrew Hatfield, Ramana Raja, and Victoria Martinez de la Cruz

2:30pm - 2:40pm

Putting OpenStack on Kubernetes: what tools can we use?

Flavio Percoco

4:20pm - 5:00pm

Neutron-based networking in Kubernetes using Kuryr – a hands-on lab

Sudhir Kethamakka, Geetika Batra, and Amol Chobe (JP Morgan Chase)

10:50am - 12:20pm

BREAKOUT SESSIONS - Wednesday, November 8th

Questions to make your storage vendor squirm

Gregory Farnum

9:50am - 10:30am

Standing up and operating a container service on top of OpenStack using OpenShift

Dan McPherson, Ata Turk (MOC), and Robert Baron (Boston University)

1:50pm - 2:30pm

Glance image import is here...now it's time to start using it!

Erno Kuvaja and Brian Rosmaita (Verizon)

4:30pm - 5:10pm