

worldpay
from FIS

OPENSIFT HIVE AT WORLDPAY





Bernd Malmqvist

Principal Container Platform Engineer



Matthew Simons

Container Platform Engineer

About Worldpay

Coming soon!

- Short introduction

- few details on the department/division that you currently work?
- something about the customers, if possible to mention a few
 - shortly about offering more side services, microservices, “all applications been build microservices / multiple small services. “
 -

How we create and manage OpenShift 4 clusters?

Our path to OpenShift 4

- The Red Hat collaboration and the early OpenShift 4 Beta engagement, and the continued feedback we received
- Red Hat works out in the open and this is how it let us find and use OpenShift Hive
- The Hive operator was Open Source but not known to us but it attracted our attention when we first heard of it
- Discover Hive capabilities and extended to fit our needs

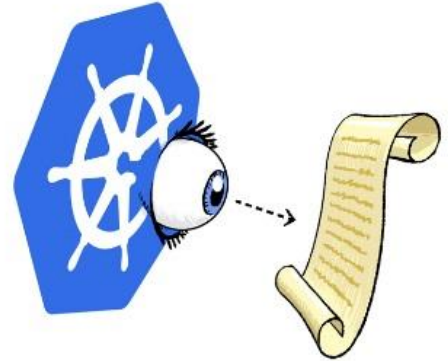
How we deploy cluster?

- Coming from OpenShift 3 we knew the problem of managing infrastructure and cluster provisioning separately
- The OpenShift 4 Installer was good for single cluster deployment but not great for multiple clusters by a team
- We wanted to treat our OpenShift 4 cluster as Cattle and not Pets
- We were looking for API driven cluster deployment and wanting to keep the cluster configuration in sync

OpenShift Hive Cluster Deployment

- A simple Kubernetes manifest defining base information to create the cluster
- We use default provisioning model with Installer Provisioned Infrastructure
- Complete cluster deployment and apply configuration takes around ~45mins and ready to use

```
ClusterDeployment
cluster.yaml
metadata:
  labels:
    wp.com/env: "eng"
  name: mycluster
spec:
  platform:
    aws:
      region: eu-west-2
  imageSet:
    name: "openshift-v4.2.7"
  clusterName: mycluster
  baseDomain: eng.wp.com
```



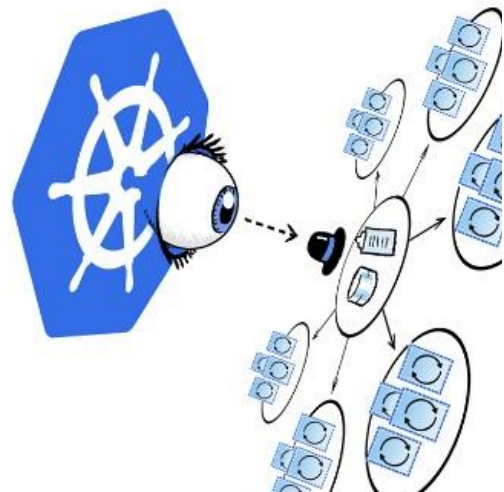
Cluster Deployment demo

How we configure cluster?

- Managing hundreds of configuration manifest and promotion across environments with multiple regional clusters whilst avoiding drift and ensuring idempotency
- OpenShift Hive has SyncSets to apply and push cluster configuration
- Regularly reconcile configuration to keep cluster configuration synchronized

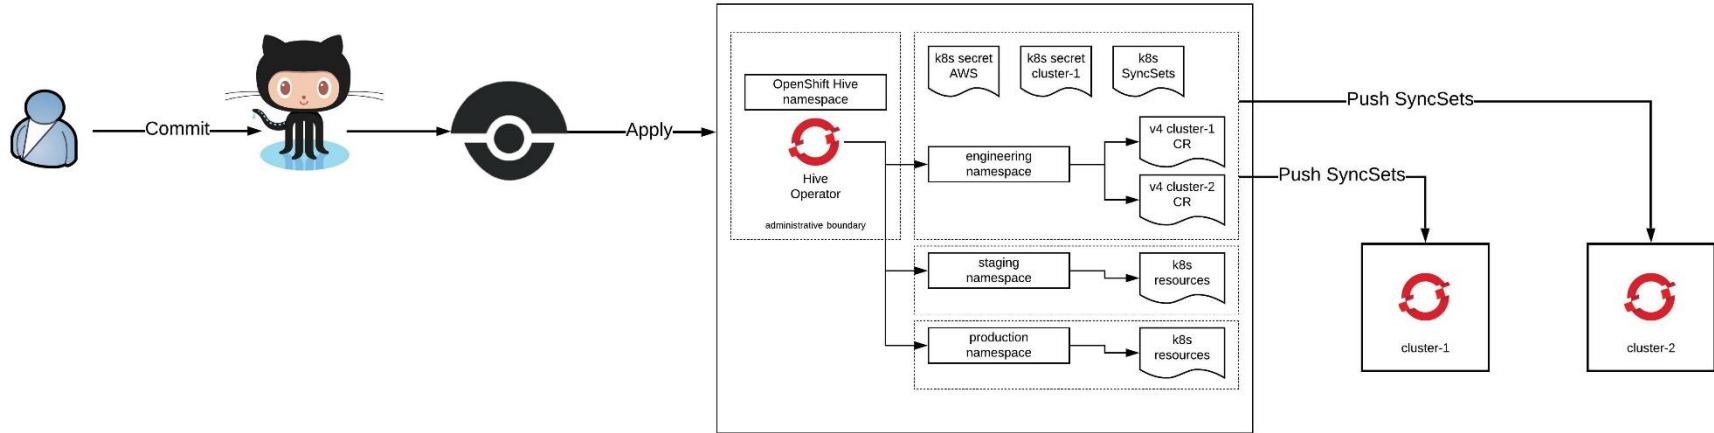
Subscription model using Hive SyncSets

- Just like in yum where you would subscribe to stable, unstable, or testing
- We use this model to promote versions to higher environments
 - Production = v0.3.1
 - Staging = v0.4.0
 - Engineering = master



- video link:
<https://youtu.be/KoDu8QLJ2E0>
Cluster SyncSet demo

GitOps



OpenShift Hive SyncSets

- It is difficult to just write a SyncSet manifest because of the size and scope
- To overcome this complexity we created a SyncSet generator by reusing the Red Hat Hive libraries to full-fill our needs of the cluster configuration
- The SyncSet generator takes normal Kubernetes manifests and generates a cluster SyncSet which can be applied using Hive

How we want to contribute back to the community

- We have ongoing talks with Red Hat and give feedback about how we use Hive and ways to improve
- Everyone benefit from Open Source so we are contributing our SyncSet generator back to the community

<-Github Worldpay Account->

Benefits for Worldpay

- OpenShift Hive operator enabled us to adopt GitOps delivery model
- Transformed the team to shift-left “software developer mindset”
- Closer working relationship with Red Hat and the community
- Improve the self-service capabilities for OpenShift 4
- Release channels and controlled promotions

**Please join us for the AMA panel in the afternoon for
Q&A,
thank you!**

THANK YOU