IT2240 - Red Hat OpenShift and Splunk - Better Together!

Mattia Mascia | Red Hat

Matthew Modestino | Splunk

Mattia
Principal Consultant | Red Hat

Matt
IT Practitioner | Splunk

splunk> .conf19

Forward-Looking Statements

During the course of this presentation, we may make forward-looking statements regarding future events or plans of the company. We caution you that such statements reflect our current expectations and estimates based on factors currently known to us and that actual events or results may differ materially. The forward-looking statements made in the this presentation are being made as of the time and date of its live presentation. If reviewed after its live presentation, it may not contain current or accurate information. We do not assume any obligation to update any forward-looking statements made herein.

In addition, any information about our roadmap outlines our general product direction and is subject to change at any time without notice. It is for informational purposes only, and shall not be incorporated into any contract or other commitment. Splunk undertakes no obligation either to develop the features or functionalities described or to include any such feature or functionality in a future release.

Agenda

Timestamps

1. Intro & Overview 2 min
2. OpenShift Container Platform 5 min
3. Red Hat Universal Base Image 5 min
4. Splunk 2min
5. Splunk Operator 3 min
6. Splunk Connect for Kubernetes, App for Infrastructure & Business Flow 10min
7. Use Cases 5 min
8. Q&A 5 min

“Last Page First!”

A great guiding principal for us Techies ☺

Red Hat & Splunk

Red Hat + Splunk partnership delivers key outcomes for our customers!

- Enterprise Platforms
- Simple, Secure & Scalable
- Hybrid - Any Data Center. Any Data.
- Opensource & Community

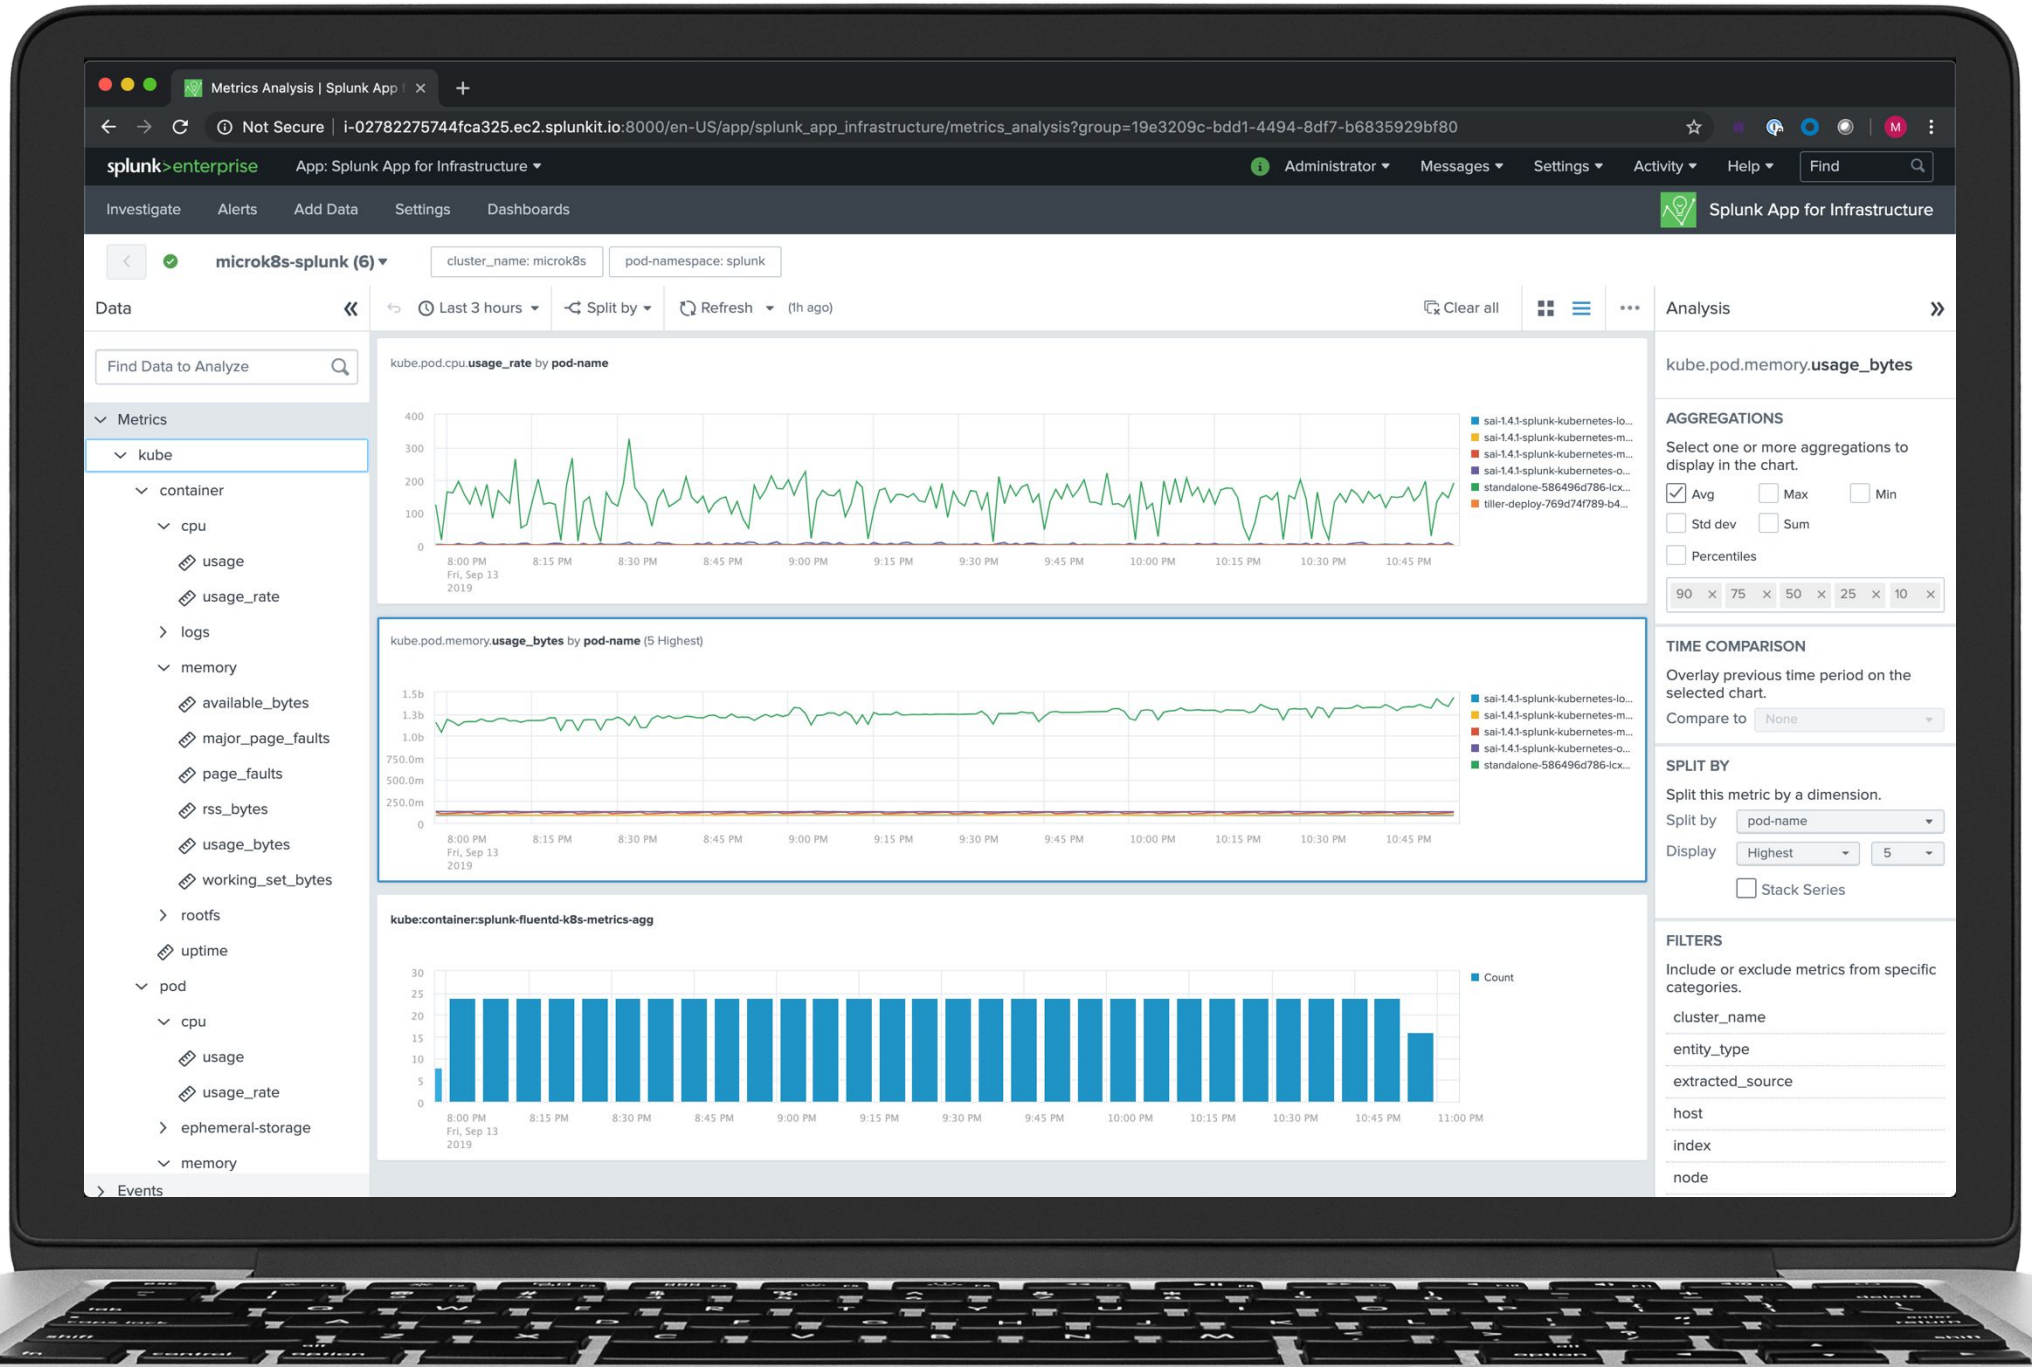

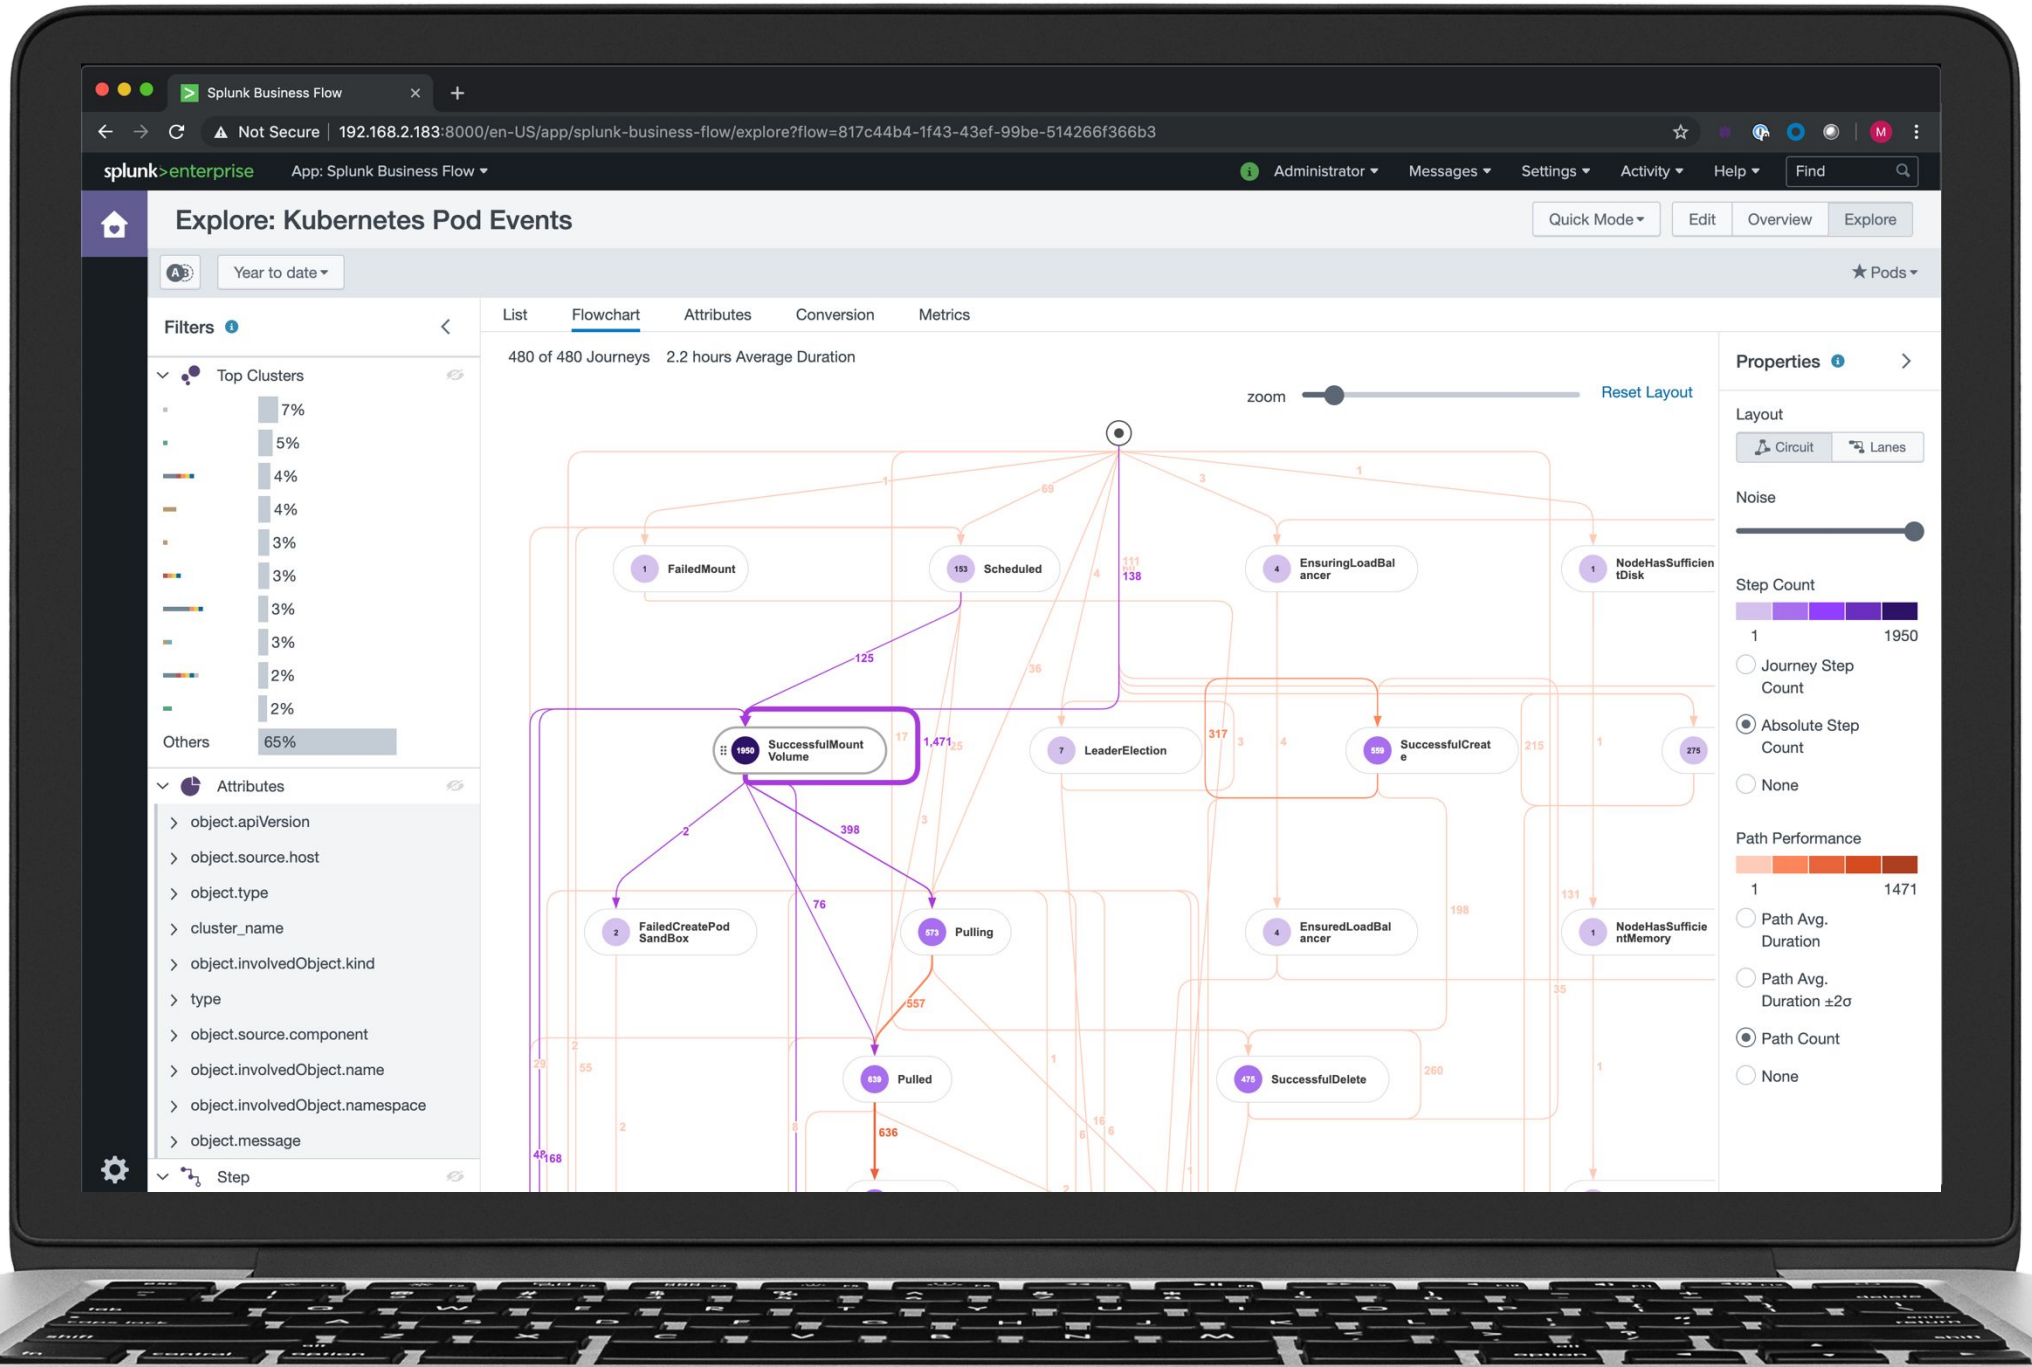

OpenShift Container Platform

What is it?

splunk> .conf19

“OpenShift is an Integrated Platform”

It’s not Just Kubernetes Certified!

OpenShift Container Platform

Everything you need, out of the box

Red Hat OpenShift 4

- Fully integrated and automated architecture
- Seamless Kubernetes deployment on any cloud or on-premises environment

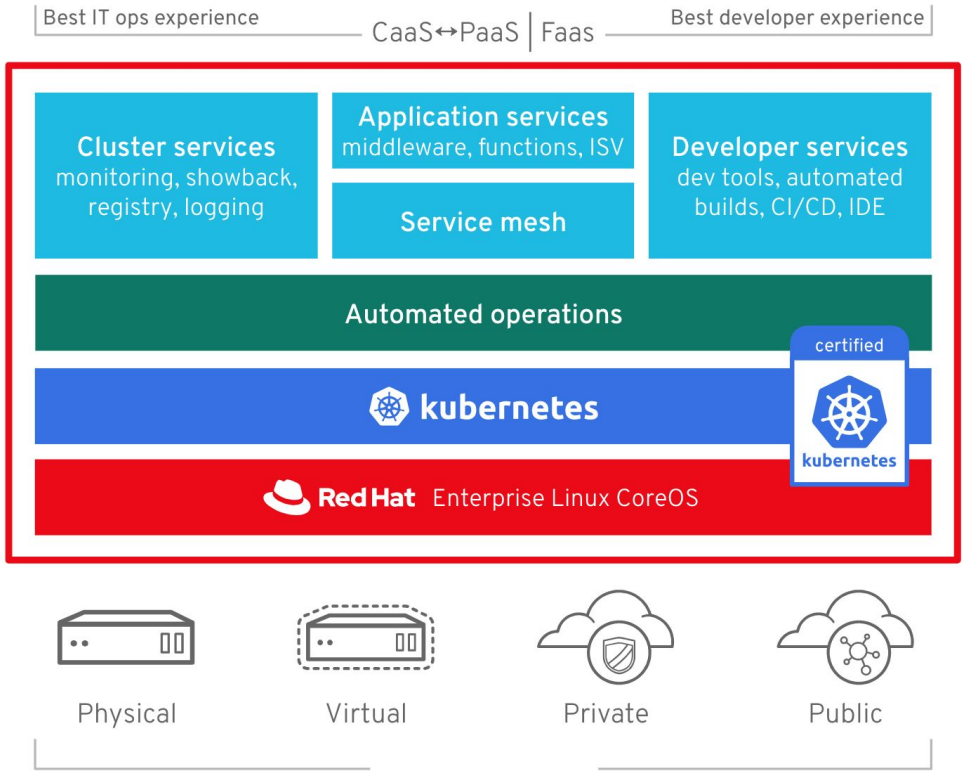

OpenShift Container Platform

Everything you need, out of the box

Red Hat OpenShift 4

- Fully automated installation, from cloud infrastructure to OS to application services
- One click platform and application updates
- Auto-scaling of cloud resources

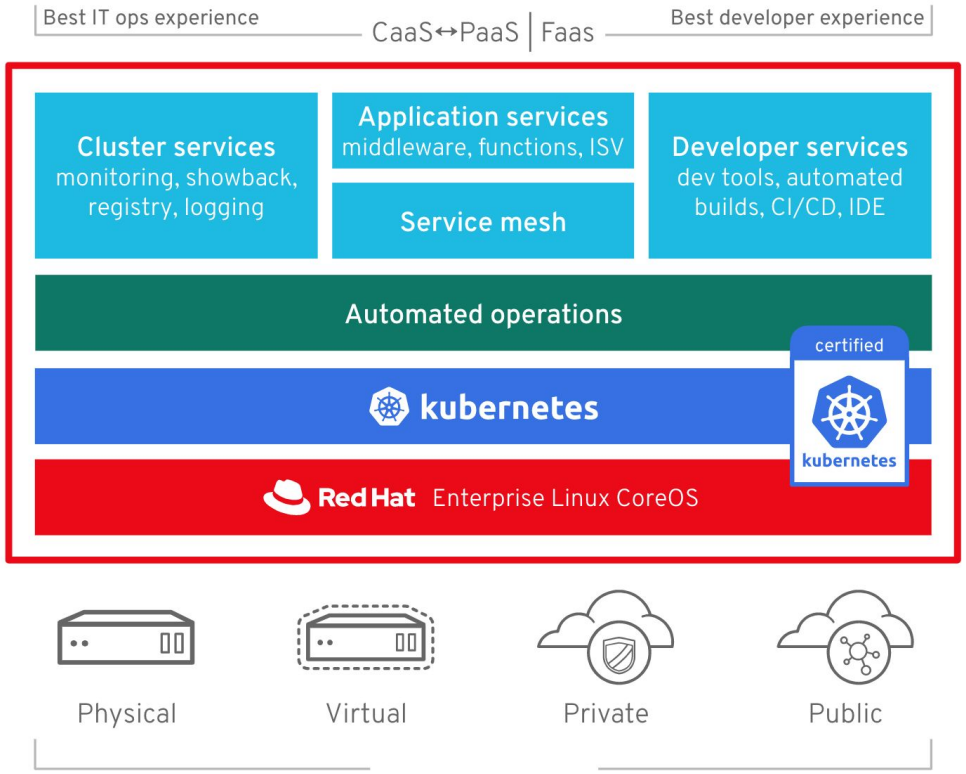

Red Hat Universal Base Image

A bold change that will help us expand the ecosystem

“Red Hat Universal Base Image”

The aim is to be the highest quality and most flexible base container image available

Can Be Build & Deployed Anywhere

Building on UBI is the first step

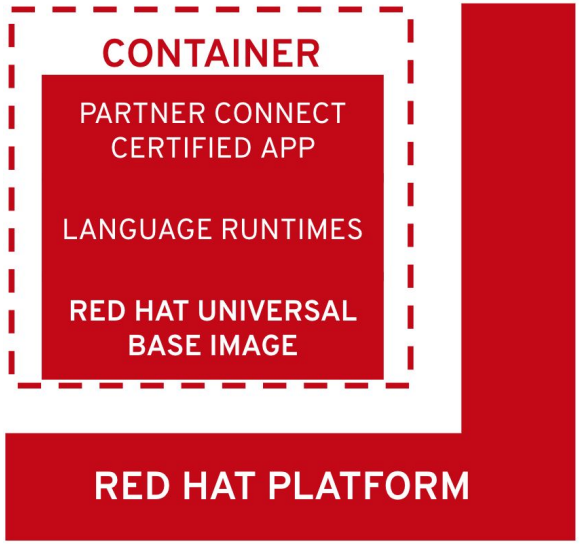

"Certification provides the highest level of support."

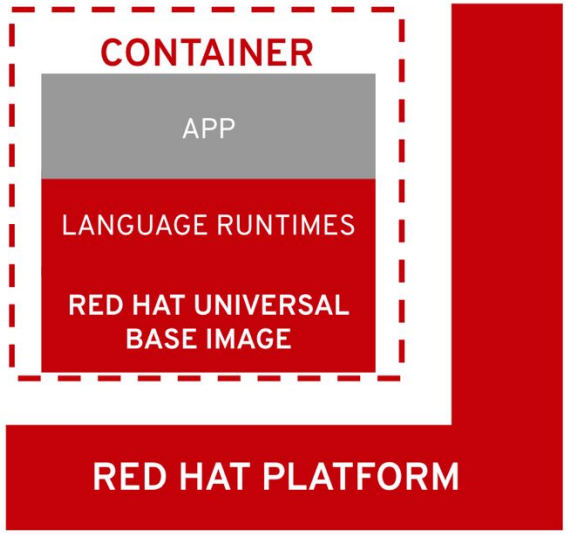

"Enterprise support when run on Red Hat platforms"

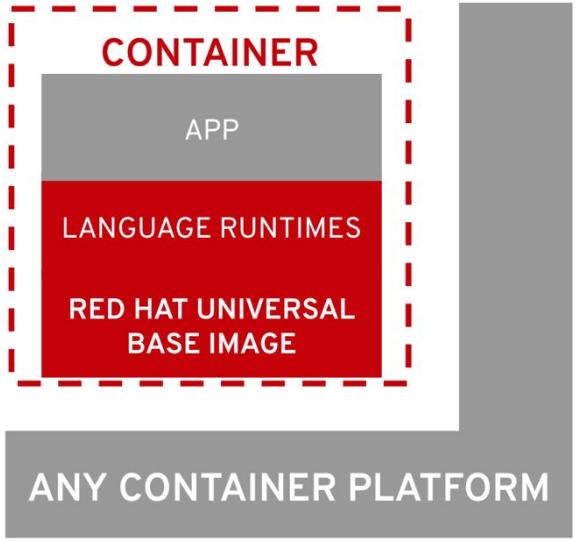

"Trusted base for any environment"

“On The Shoulders of Giants!”

Leaving the the OS & Infra to the Pros so we can Splunk!

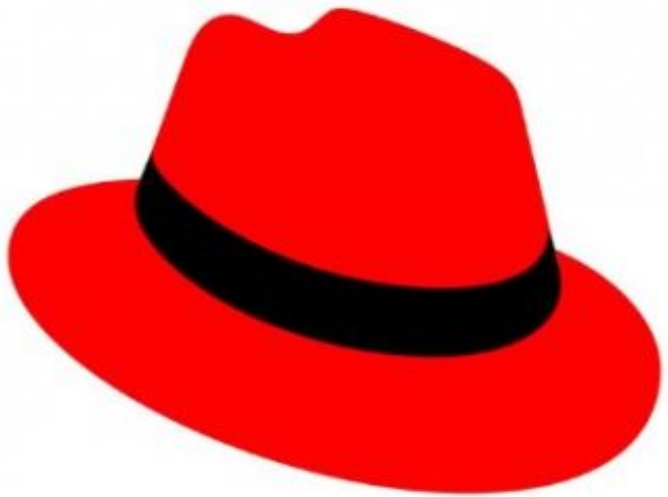

splunk>

Splunk

Any Question. Any Data. One Splunk.

splunk> .conf19

Splunk Operator

Deploying the Splunk Platform on Openshift with Red Hat Universal Base Image

splunk> .conf19

Splunk Operator

Encoding Operational Knowledge

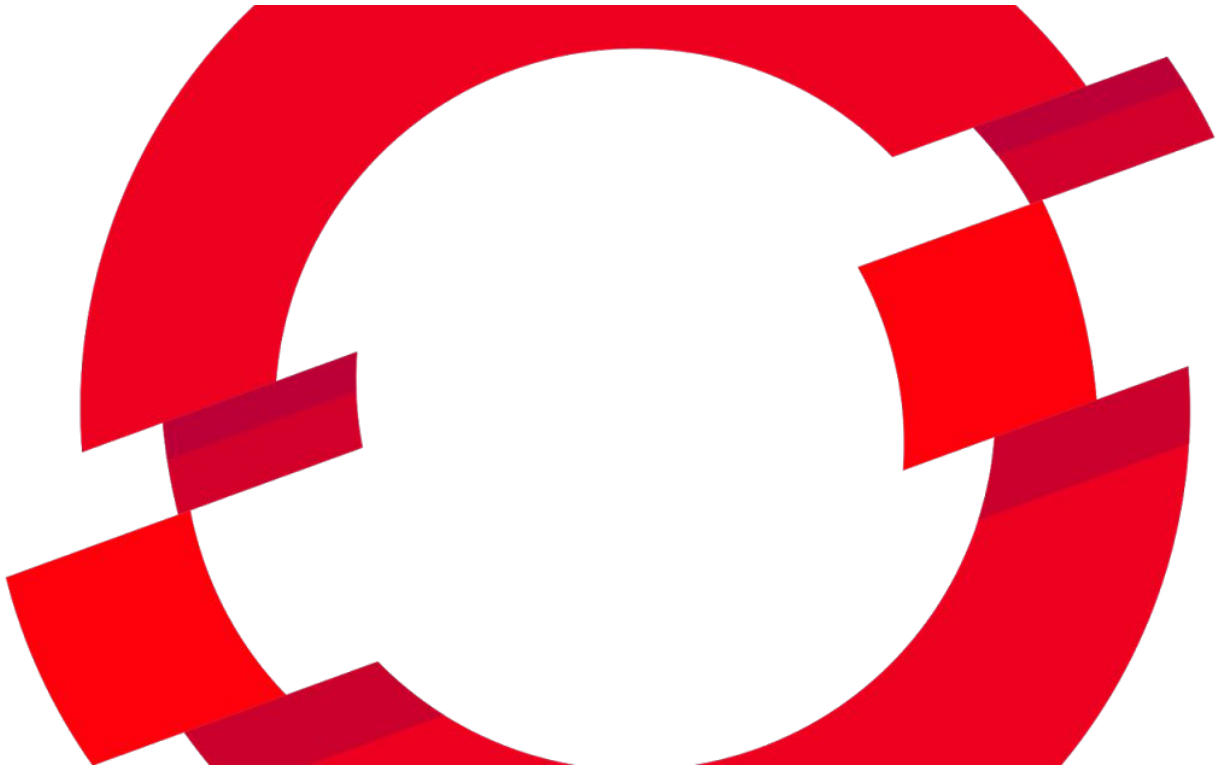

```
apiVersion: enterprise.splunk.com/v1alpha1
kind: SplunkEnterprise
metadata:
  name: cluster
  finalizers:
  - enterprise.splunk.com/delete-pvc
spec:
  splunkVolumes:
  - name: licenses
    configMap:
      name: splunk-licenses
  licenseUrl: /mnt/licenses/enterprise.lic
  resources:
    splunkVarStorage: 10Gi
    splunkIndexerStorage: 50Gi
  topology:
    indexers: 3
    searchHeads: 3
```

Splunk App for Infrastructure

The easy button for OpenShift Integration

splunk> .conf19

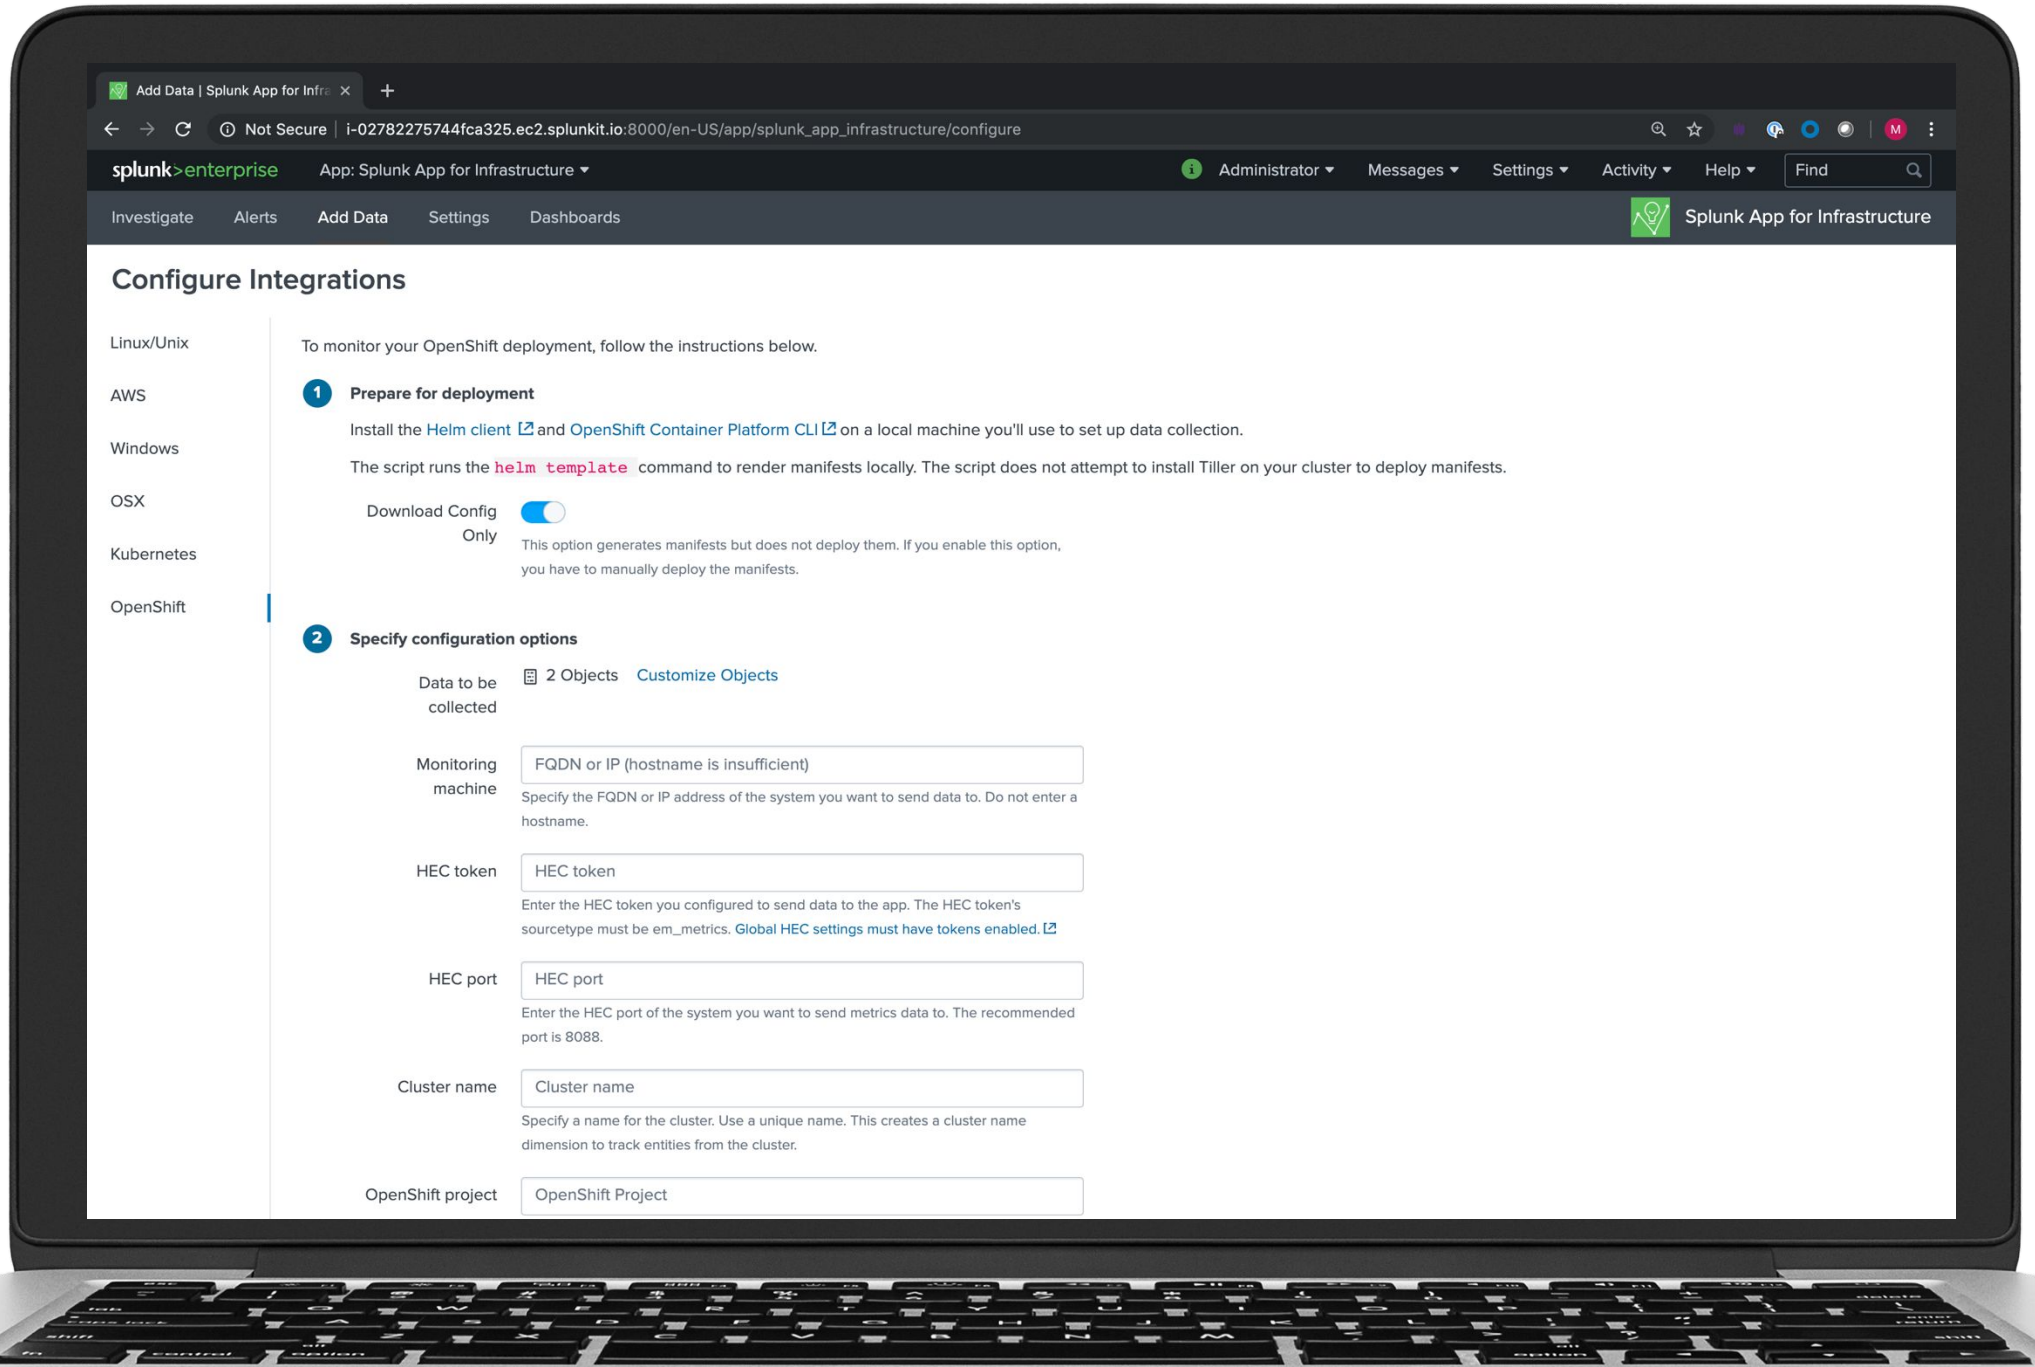

Splunk Connect for Kubernetes

Kubernetes Data Collection by Splunk & the open source Community!

splunk> .conf19

Kubernetes Logging Architecture

Node Agent

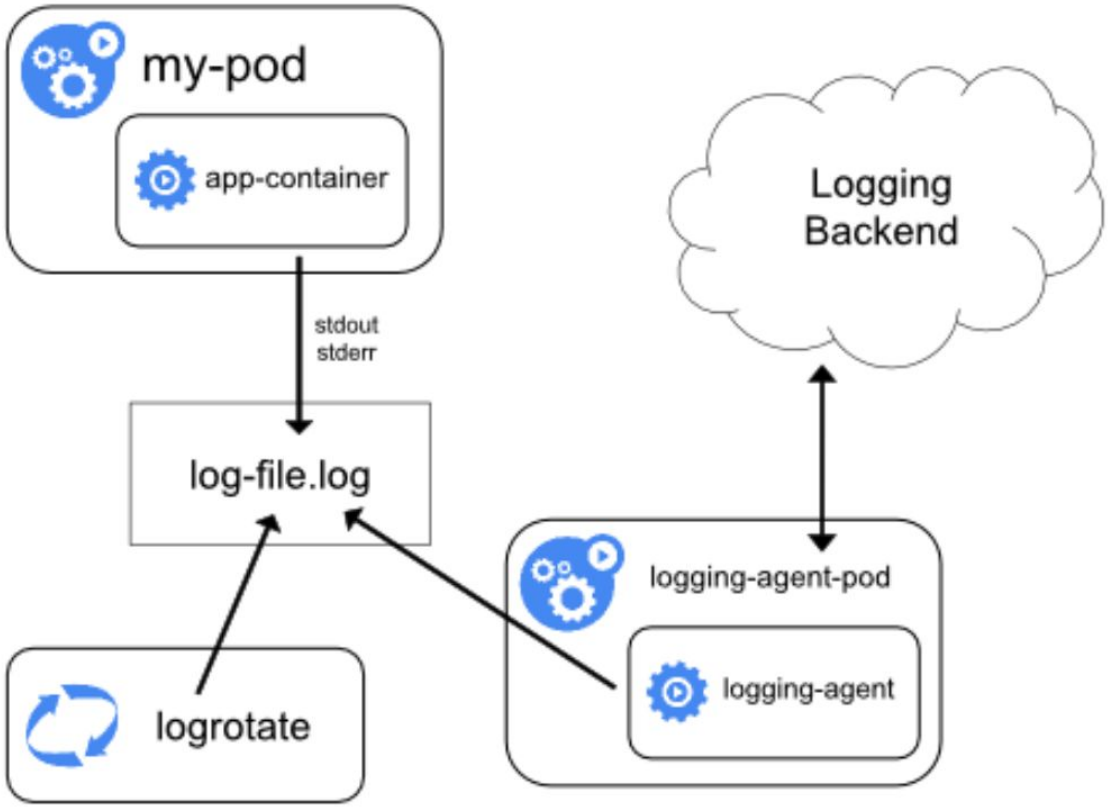

Kubernetes Logging Architecture

Sidecar Agent

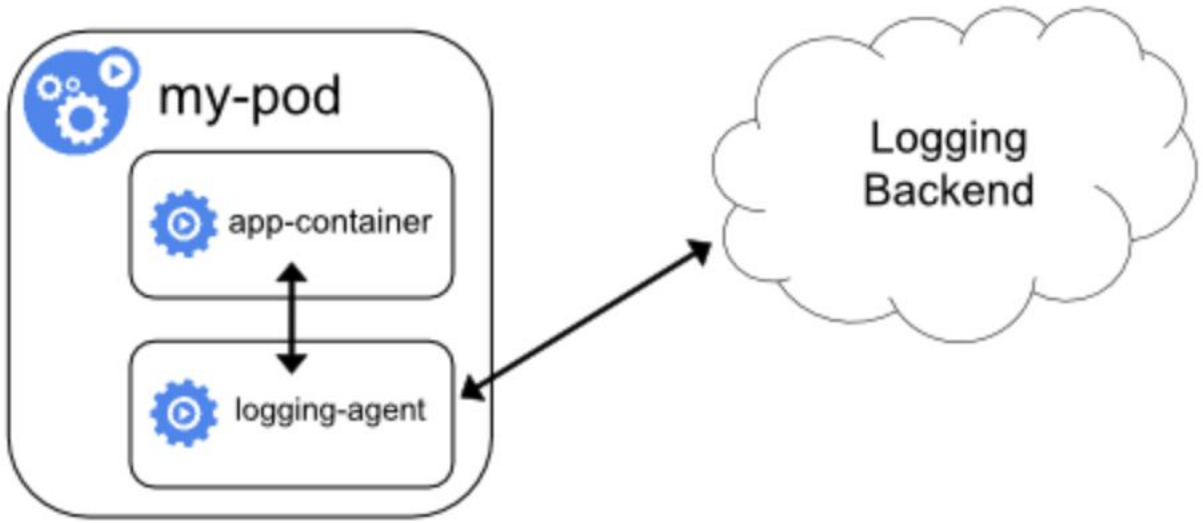

Splunk Connect for Kubernetes

Kubernetes Data Sources

Splunk Kubernetes Logging

Splunk Kubernetes Metrics

Splunk Kubernetes Objects

Architecture

Logging, Metrics, Objects

- Logging DaemonSet
- Metrics DaemonSet
- Metrics-Agg Deployment
- Objects Deployment
- OpenShift Web Console

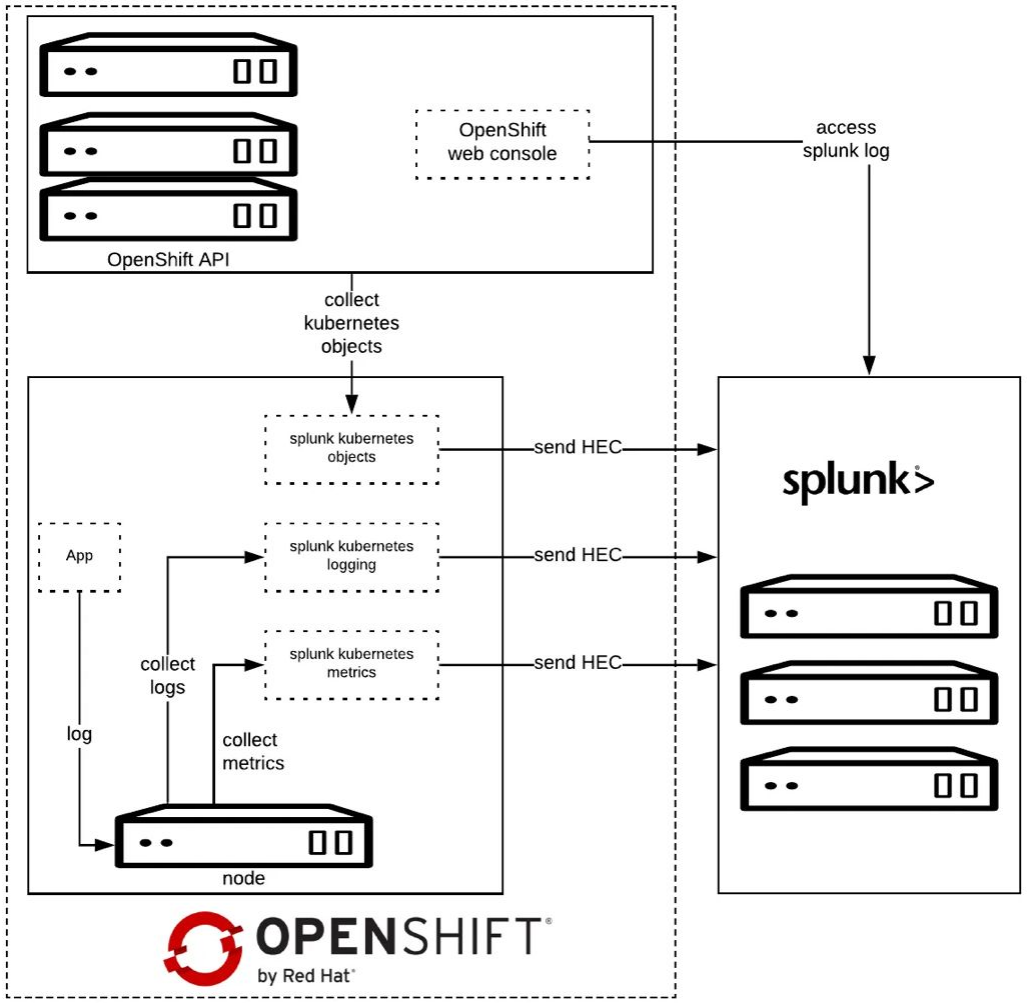

Splunk Kubernetes Logging

Application & OpenShift Cluster Logging

Docker/CRI-O Logs
Journald
Custom Log Sources

in_tail
systemd
jq_transformer
fluentd-hec

HTTP Event Collector
Event Index

Splunk Kubernetes Objects

OpenShift Metadata Collection

Kubernetes Objects
OAPI Objects

Kubernetes_objects
jq_transformer
fluentd-hec

HTTP Event Collector
Event Index

Splunk Kubernetes Metrics

OpenShift Platform Metrics

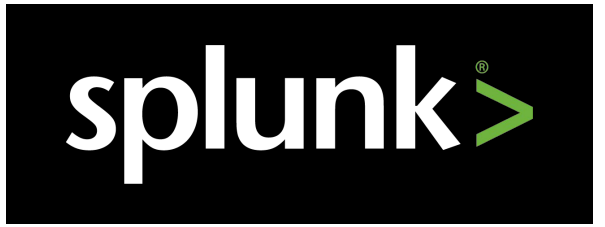

Kubernetes API Metrics
Kubelet Metrics

kubernetes_metrics
kubernetes_metrics_aggregator
record_modifier
fluentd-hec

HTTP Event Collector
Metrics Index

OpenShift Web Console

Extending OpenShift Console

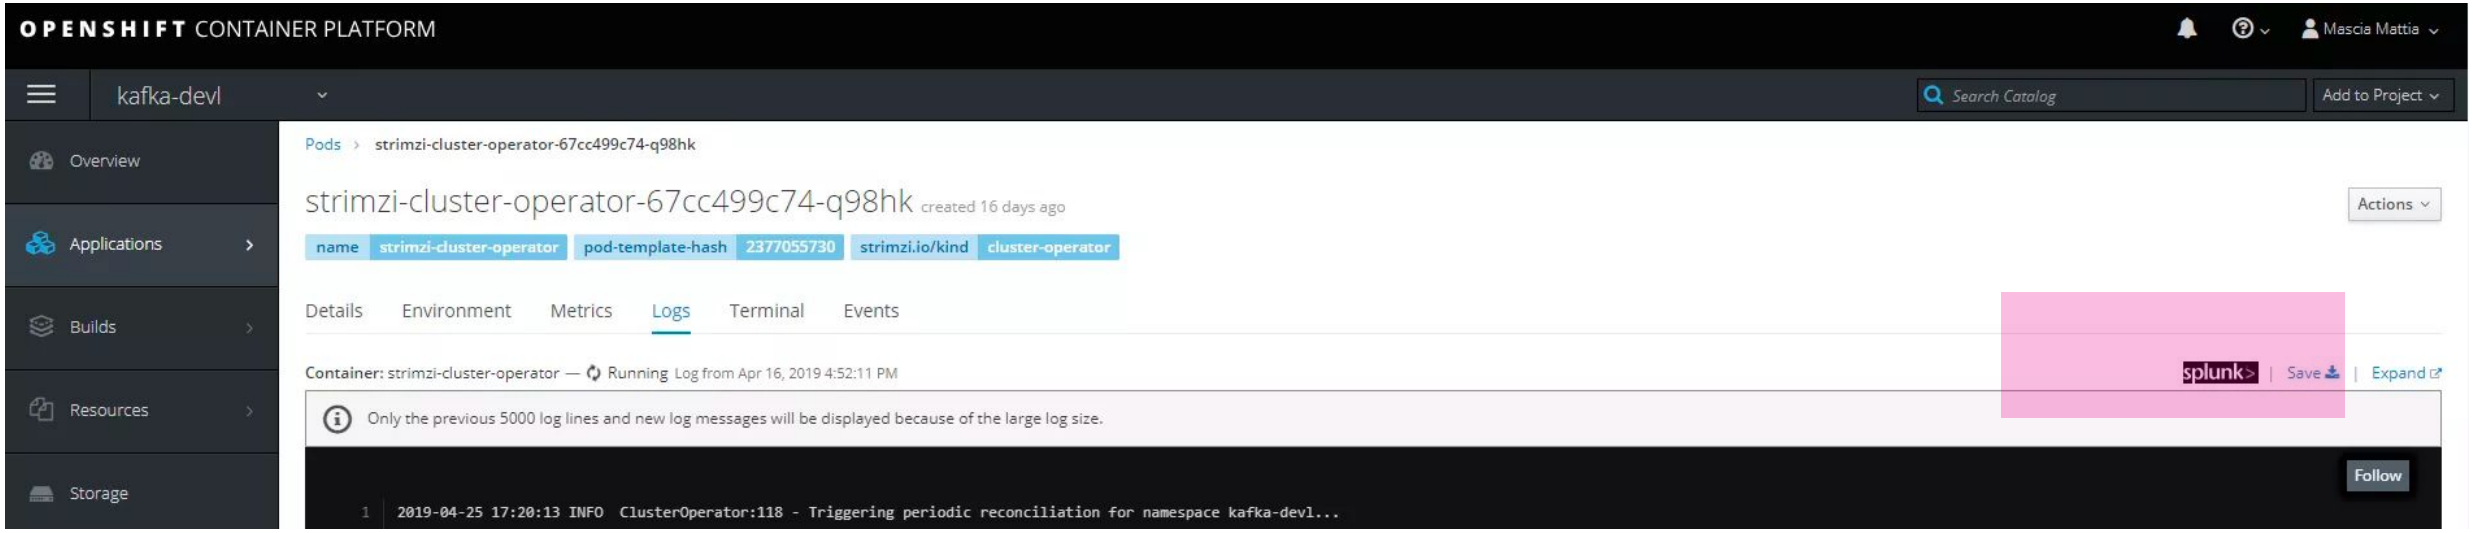

“Make the Data Dance!”

OpenShift Use Case Demos

splunk> .conf19

Demo

Red Hat & Splunk

Red Hat + Splunk partnership delivers key outcomes for our customers!

- Enterprise Platforms
- Simple, Secure & Scalable
- Hybrid – Any Data Center. Any Data.
- Open source & Community

Q&A

Mattia Mascia | Red Hat
Matthew Modestino | Splunk

splunk> .conf19

https://www.openshift.com

https://docs.splunk.com/Documentation/InfraApp/latest/Admin/AddDataOpenShift

https://github.com/splunk/splunk-connect-for-kubernetes

https://github.com/splunk/docker-splunk

https://github.com/splunk/splunk-operator

.conf19

splunk>

Thank You !

Go to the .conf19 mobile app to

RATE THIS SESSION