SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

Kubernetes in the Cloud (Part 3): (Almost) Everything You Need to Know about Stateful Workloads

Live Webcast
August 20, 2019
10:00 am PT

Today's Presenters

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

Mike Jochimsen
Director of Alliances
Kaminario

Paul Burt
Technical Product Marketing Engineer
NetApp

Ingo Fuchs
Chief Technologist, Cloud and DevOps
NetApp

SNIA Legal Notice

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

- The material contained in this presentation is copyrighted by the SNIA unless otherwise noted.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

SNIA-At-A-Glance

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

185
industry leading organizations

2,000
active contributing members

50,000
IT end users & storage pros worldwide

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

What We Do

Educate vendors and users on cloud storage, data services and orchestration

Support & promote business models and architectures: OpenStack, Software Defined Storage, Kubernetes, Object Storage

Understand Hyperscaler requirements Incorporate them into standards and programs

Collaborate with other industry associations

Kubernetes in the Cloud Series

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

- Kubernetes in the Cloud (Part 1)
  - What is Kubernetes? Why would you want to use it?
  - How does Kubernetes help in a multi-cloud/private cloud environment?
  - How does Kubernetes orchestrate & manage storage? Can Kubernetes use Docker?
  - How do we provide persistence and data protection?
  - On demand at: http://bit.ly/KubeCloud1
- Kubernetes in the Cloud (Part 2)
  - Persistent storage and how to specify it
  - Ensuring application portability between Private and Public Clouds
  - Building a self-service infrastructure (Helm, Operators)
  - Selecting Block, File, Object (Traditional Storage, SDS)
  - On demand at: http://bit.ly/Kube2

Agenda

- Kubernetes is a Platform for mostly *state**less*** work
- Why *state**ful*** work is challenging
  - The lifecycle is more complicated
  - Container's learning curve + tools
  - Security is paramount
- 5 Ways to run *State**ful*** work on Kubernetes
- Questions
- Links & Resources


“Kubernetes is becoming the Linux of the cloud”
- Jim Zemlin,
Executive Director at the Linux Foundation

Architecture

Dockerizing MySQL at Uber Engineering

Joakim Recht

Kelsey Hightower
@kelseyhightower

Kubernetes has made huge improvements in the ability to run stateful workloads including databases and message queues, but I still prefer not to run them on Kubernetes.

9:04 AM · Feb 13, 2018 · Twitter Web Client

Secrets Management

How to commit code without leaking

Risks

- In the API server secret data is stored in etcd; therefore:
  - Administrators should enable encryption at rest for cluster data (requires v1.13 or later)
  - Administrators should limit access to etcd to admin users
  - Administrators may want to wipe/shred disks used by etcd when no longer in use
  - If running etcd in a cluster, administrators should make sure to use SSL/TLS for etcd peer-to-peer communication.
- If you configure the secret through a manifest (JSON or YAML) file which has the secret data encoded as base64, sharing this file or checking it in to a source repository means the secret is compromised. Base64

Encrypt secrets at rest,
Use RBAC,
and other best practices…

Use proven tools

Search or jump to… /

Pull requests Issues Marketplace Explore

godaddy / **kubernetes-external-secrets**

Watch 20 Star 379

Code Issues 33 Pull requests 10 Projects 0 Wiki Security Insights

Kubernetes External Secrets

kubernetes secret-management secrets-management aws aws-secrets-manager

92 commits **8** branches **7** releases **12** contributors

Branch: **master** New pull request Create new file Upload files Find File Clone

Flydiverny and **jeffpearce** feat: allow setting type in external secret to support other than Opa… ··· Latest commit 22669

bin	refactor: use watch api and instant poll new or modified secrets (#107)
charts/kubernetes-external-secrets	chore: remove events interval milliseconds references (#129)
config	refactor: use watch api and instant poll new or modified secrets (#107)
examples	feat: allow setting type in external secret to support other than Opa…

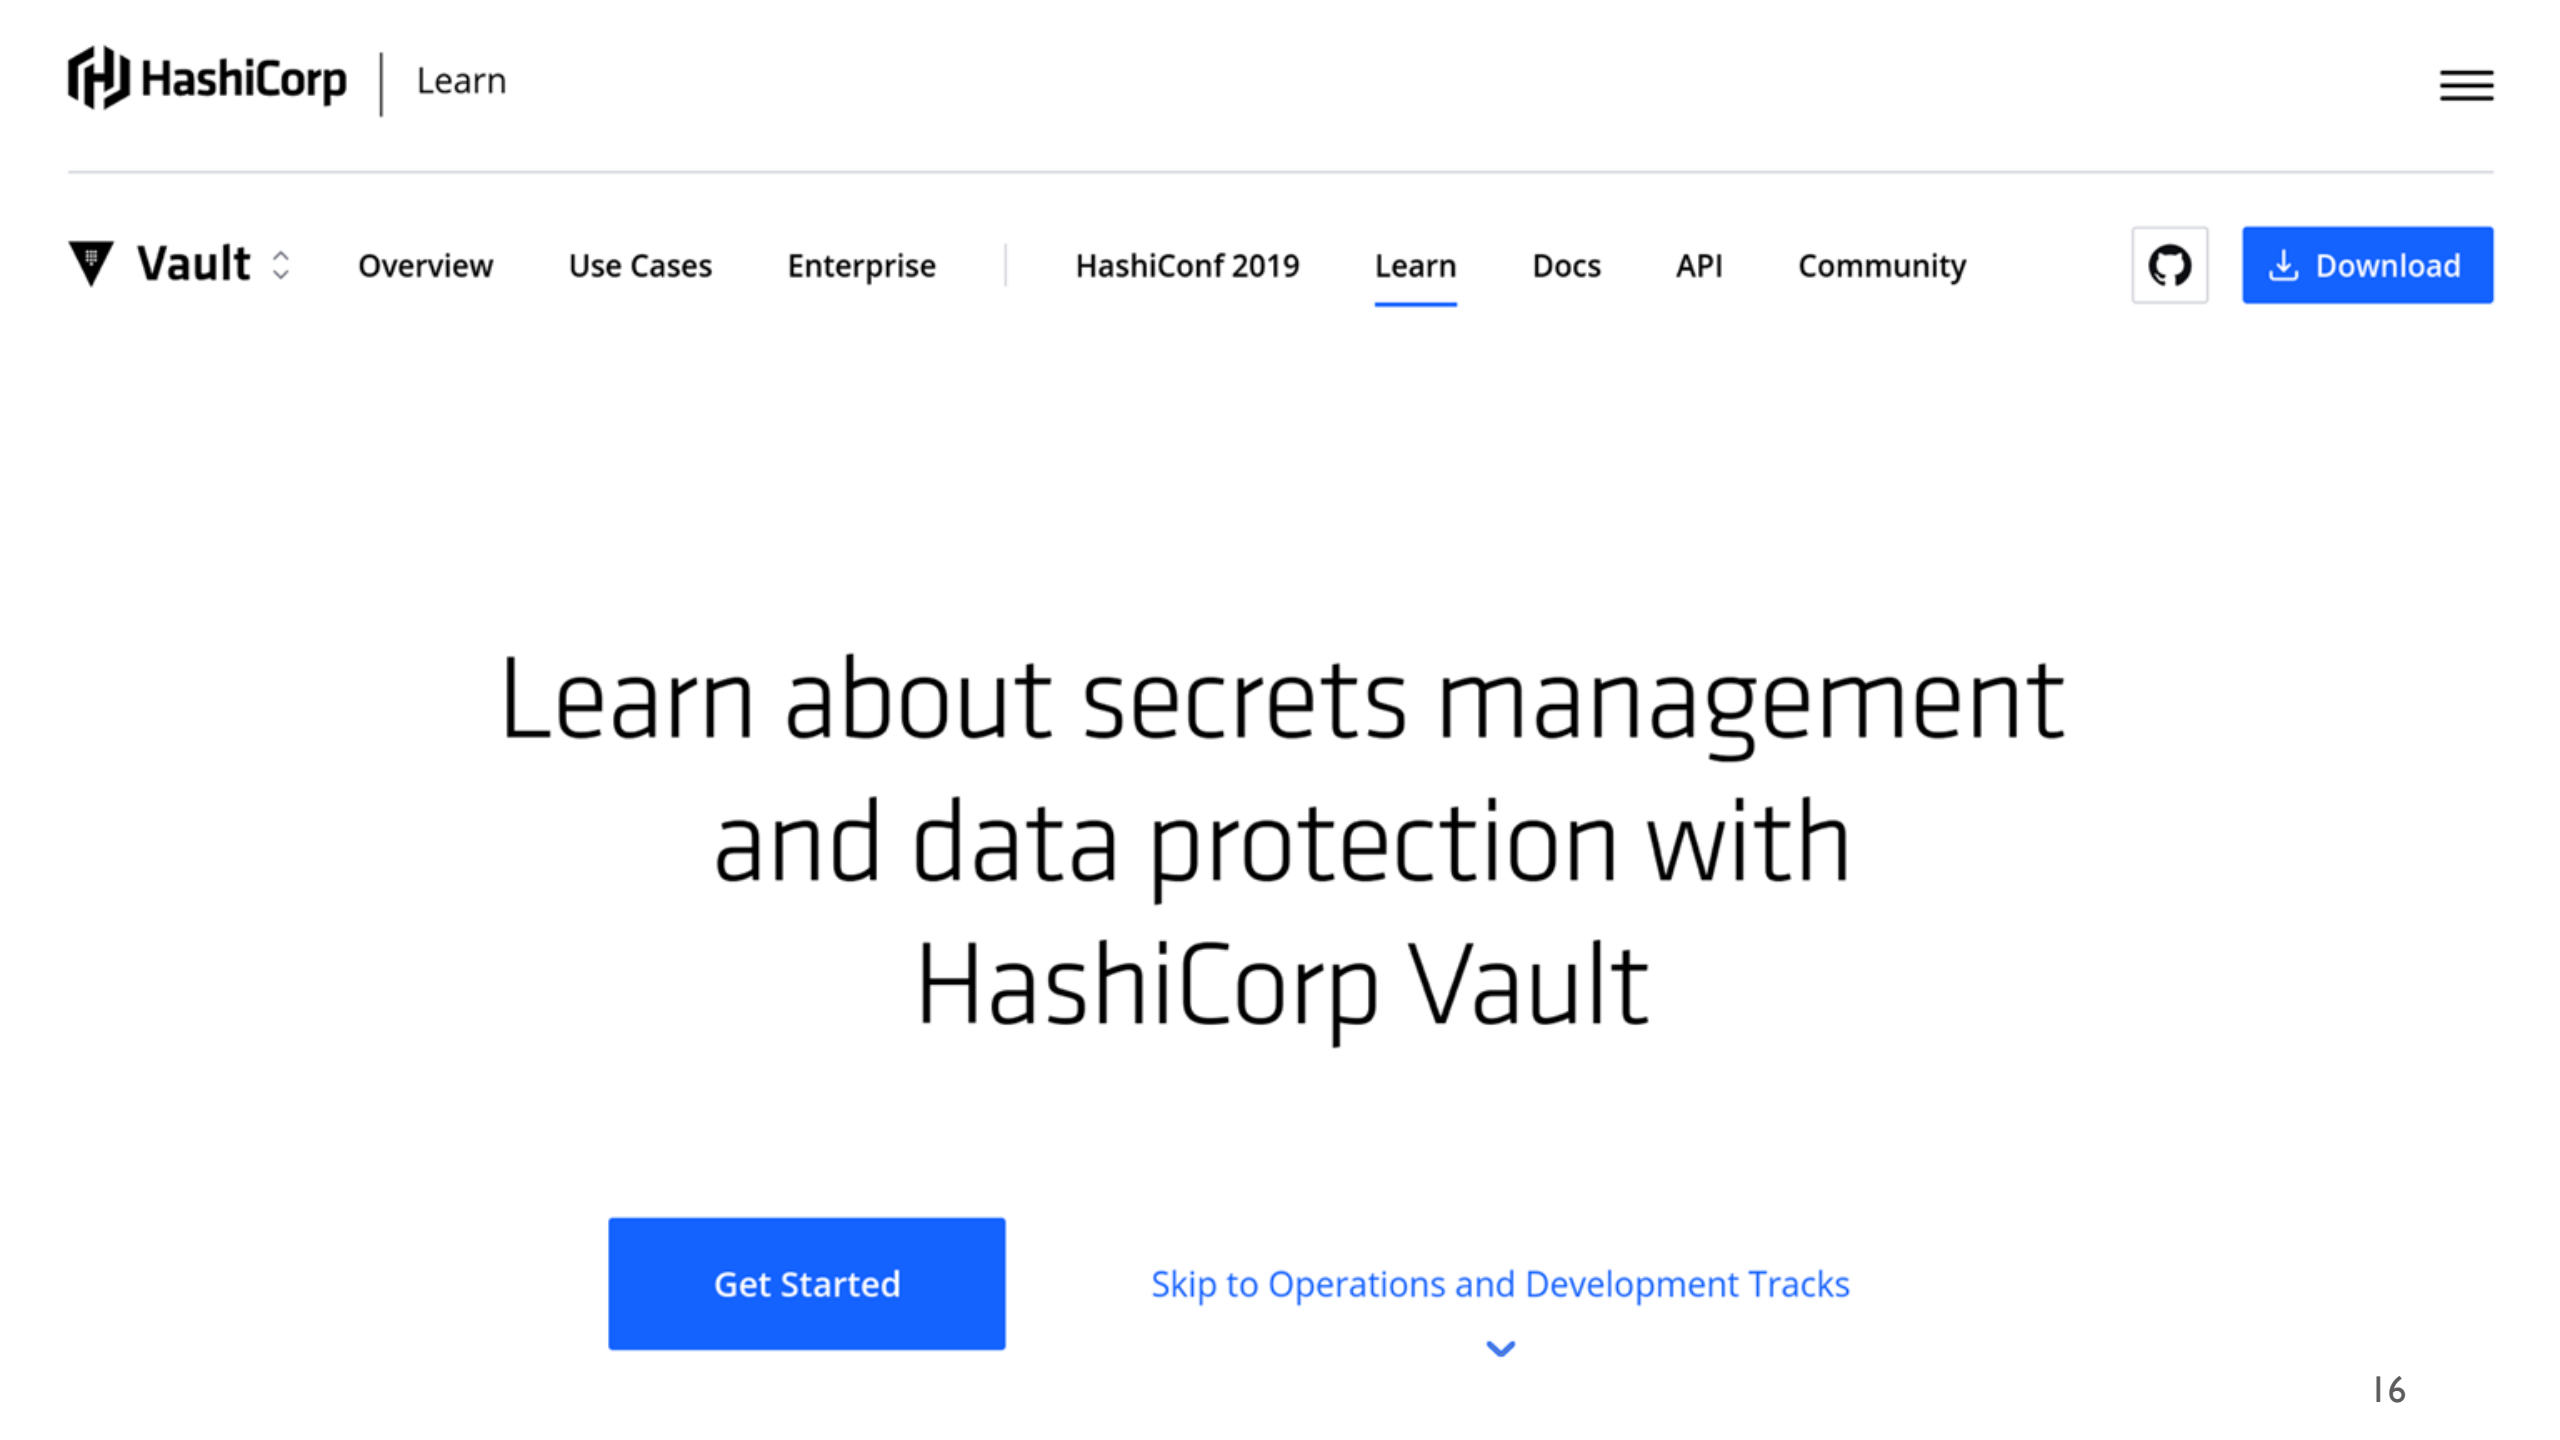
HashiCorp | Learn
Vault
Overview
Use Cases
Enterprise
HashiConf 2019
Learn
Docs
API
Community
Download
Learn about secrets management and data protection with HashiCorp Vault
Get Started
Skip to Operations and Development Tracks

Getting Started

12 TOPICS 64 MINS

Vault secures, stores, and tightly controls access to tokens, passwords, certificates, API keys, and other secrets in modern computing. Get started here.

START HERE

Install Vault →

2 MIN | The first step to using Vault is to get it installed.

Starting the Server →

5 MIN | After installing Vault, the next step is to start the server.

Your First Secret →

5 MIN | With the Vault server running, let's read and write our first secret.

Secrets Engines →

5 MIN | Secrets engines create, read, update, and delete secrets.

https://www.katacoda.com/courses/kubernetes

Start Scenario | Start Sc Networking Introduction | Start Scenario | Start Scenario

Getting Started With CRI-O and Kubeadm

Learn how to deploy a CRI-O based Kubeadm cluster

Start Scenario

BY RICHARD LI

Running Stateful Services on Kubernetes

Learn how to run stateful services on Kubernetes

Start Scenario

BY PETER BENJAMIN

Use Kubernetes To Manage Secrets And Passwords

Learn how Kubernetes can help keep secrets secure

Start Scenario

BY JONATHAN JOHNSO...

Deploy Docker Compose Files with Kompose

Learn how to use Kompose to deploy existing Docker Compose definitions

Start Scenario

BY JONATHAN JOHNSO...

Operating Kubernetes Clusters and Applications Safely

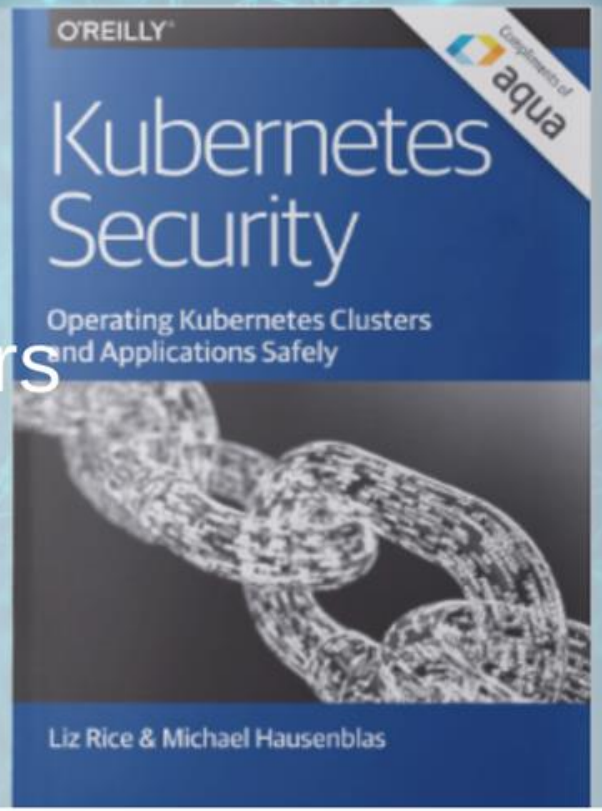

Download this eBook published by O'Reilly Media

Written by Liz Rice from Aqua Security and Michael Hausenblas from Red Hat

Kubernetes has fundamentally changed the way DevOps teams create, manage, and operate container-based

Get the O'Reilly Media Book

First name *

Las

DB on a VM

The best option, for most

Running a DB on a VM just needs some knowledge of …

- Services

GOOGLE CLOUD PLATFORM

Kubernetes best practices: mapping external services

Sandeep Dinesh
Developer Advocate

May 25, 2018

Editor's note: Today is the sixth installment in a seven-part video and blog series from Google Developer Advocate Sandeep Dinesh on how to get the most out of your Kubernetes environment.

If you're like most Kubernetes users, chances are you use services that live outside your cluster. For example, maybe you use the Twillio API to send text messages, or maybe the Google Cloud Vision API to do image analysis.

If your applications in your different environments connect to the same external endpoint, and have no plans to bring the external service into your Kubernetes cluster, it is perfectly fine to use the external service endpoint directly in your code. However, there are many scenarios where this is not the

```yaml
kind: Service
apiVersion: v1
metadata:
 name: mongo
spec:
 type: ExternalName
 externalName: ds149763.mlab.com
```

```yaml
kind: Service
apiVersion: v1
metadata:
 name: mongo
Spec:
 type: ClusterIP
 ports:
 - port: 27017
   targetPort: 27017

---

kind: Endpoints
apiVersion: v1
metadata:
 name: mongo
subsets:
 - addresses:
     - ip: 10.240.0.4
   ports:
     - port: 27017
```

That’s it! Super easy, and all of your old automation, monitoring, etc still work

DB in k8s via StatefulSet

Warning: can be problematic

StatefulSets need some knowledge of...

- Init Containers
- Persistent Volumes (PV)
- PV Claims (PVC)
- Storage Classes
- Services
- Pods
- ConfigMaps

StatefulSets need some knowledge of...

- Init Containers
- Persistent Volumes (PV)
- PV Claims (PVC)
- Storage Classes
- Services
- Pods
- ConfigMaps

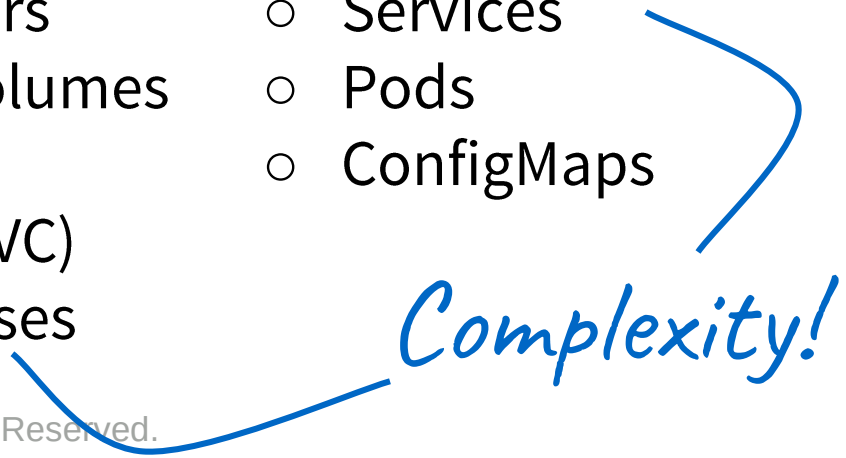

Kubernetes Storage Terminology

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

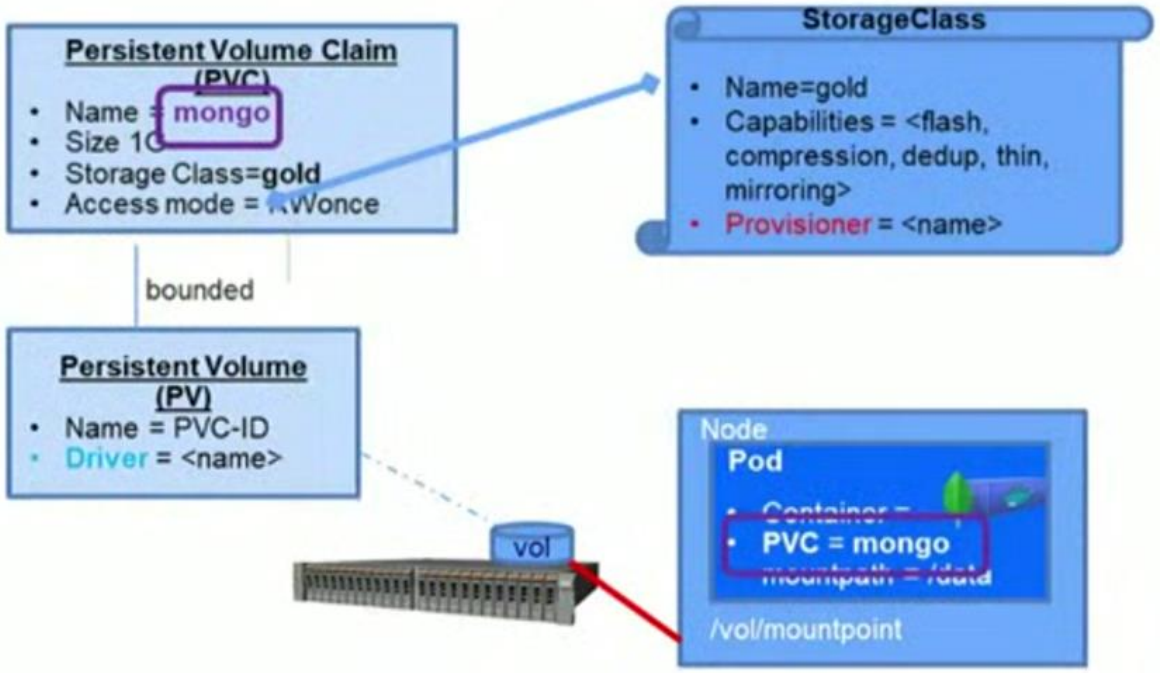

https://kubernetes.io/docs/concepts/storage/volumes/

Storage Class

To achieve dynamic volume creation, the admin must define a k8s **StorageClass** (e.g : gold, silver).

Provision a volume

1. The user creates a claim for volume (**PVC**).
2. The "Provisioner" (vendor specific) listens to new PVC requests, and dynamically creates the volume on the storage system (if no PV already matched)
   - The **PV** is created with "Driver" setting. The Driver(vendor specific) handles the volume attach\detach to the node.

Create a stateful POD

1. The user creates a **POD** with the new PVC.
2. K8s triggers the "Driver" in order to attach the PV to the node.
3. The volume is now mapped and mounted to the node.
4. k8s starts the POD with the PV mounted to /data inside the container.

✓ The stateful container is UP

Subscribe

Run a Single-Instance Stateful Application

This page shows you how to run a single-instance stateful application in using a PersistentVolume and a Deployment. The application is MySQL.

Be sure you’re thinking about Day 2 operations, not just the installation

Storing backups

Restoring from backup

Be sure you're thinking about Day 2 operations, not just the installation

Deleting the stateful app
(and reclaiming resources)

Upgrading the stateful app

Scaling the stateful app

You absolutely need failover and replication; containers fail for all sorts of silly reasons

DB in k8s via Operator

Introduces complexity, but sometimes worth it

“Application specific operational knowledge captured in software”

More specifically, an Operator is just:

CRDs + **automation**

More specifically, an Operator is just:

CRDs + **automation**

Custom Resource Definition

A native Kubernetes object, that gives you the power to customize the behavior of Kubernetes.

How can you create an Operator?

Operators, by their nature, are application-specific, so the hard work is going to be encoding all of the application operational domain knowledge into a reasonable configuration resource and control loop. There are some common patterns that we have found while building operators that we think are important for any application:

1. Operators should install as a single deployment e.g. `kubectl create -f https://coreos.com/operators/etcd/latest/deployment.yaml` and take no additional action once installed.
2. Operators should create a new third party type when installed into Kubernetes. A user will create new application instance using this type.
3. Operators should leverage built-in Kubernetes primitives like Services and Replica Sets when possible to leverage well-tested and well-understood code.
4. Operators should be backwards compatible and always understand previous versions of resources a user has created.
5. Operators should be designed so application instances continue to run unaffected if the Operator is stopped or removed.
6. Operators should give users the ability to declare a desired version and orchestrate application upgrades based on the desired version. Not upgrading software is a common source of operational bugs and security issues and Operators can help users more confidently address this burden.
7. Operators should be tested against a "Chaos Monkey" test suite that simulates potential failures of Pods, configuration, and networking.

An example of a complex application being started

```
etcd --name infra1 --listen-client-urls http://127.0.0.1:2379 \
--advertise-client-urls http://127.0.0.1:2379 --listen-peer-urls http://127.0.0.1:12380 \
--initial-advertise-peer-urls http://127.0.0.1:12380 --initial-cluster-token etcd-cluster-1 \
--initial-cluster
'infra1=http://127.0.0.1:12380,infra2=http://127.0.0.1:22380,infra3=http://127.0.0.1:32380' \
--initial-cluster-state new --enable-pprof
```

An example

Couchbase
START A REVOLUTION

Version Compatibility

You must run `cbbackupmgr` from a Couchbase Server installation with the same major and minor version as the host cluster. For example, to back up data from (or restore data to) a cluster running Couchbase Server 5.5, you must run `cbbackupmgr` from a Couchbase Server 5.5 node.

So, why not use an operator for everything?

The OperatorHub is a marketplace. Operators there should package everything you need.

OperatorHub.io

Search OperatorHub...

Contribute

Welcome to OperatorHub.io

OperatorHub.io is a new home for the Kubernetes community to share Operators. Find an existing Operator or list your own today.

CATEGORIES

- AI/Machine Learning
- Big Data
- Cloud Provider
- Database
- Integration & Delivery
- Logging & Tracing
- Monitoring
- Networking
- OpenShift Optional
- Security
- Storage
- Streaming & Messaging

PROVIDER

- [ ] Amazon Web Services (1)

39 ITEMS

VIEW SORT A-Z

Aqua Security Operator

provided by Aqua Security, Inc.

The Aqua Security Operator runs within Kubernetes cluster and provides a means to

AWS Service Operator

provided by Amazon Web Services, Inc.

The AWS Service Operator allows you to manage AWS

Camel K Operator

provided by The Apache Software Foundation

Apache Camel K (a.k.a. Kamel) is a lightweight integration

CockroachDB

provided by Helm Community

Community Jaeger Operator

provided by CNCF

Couchbase Operator

provided by Couchbase

Operator Framework

Report abuse

The Operator Framework is an open source toolkit to manage Kubernetes native applications, called Operators, in an effective, automated, and scalable way.

https://operatorhub.io/ Verified

Repositories 24 People 16

Pinned repositories

operator-sdk

SDK for building Kubernetes applications. Provides high level APIs, useful abstractions, and project scaffolding.

Go ★ 1.9k 434

operator-lifecycle-manager

A management framework for extending Kubernetes with Operators

Go ★ 339 143

operator-metering

Operator metering is responsible for collecting metrics and other information about what's happening in a Kubernetes cluster, and providing a way to create reports on the collected data.

Go ★ 166 37

Find a repository… Type: All Language: All

operator-metering

Operator metering is responsible for collecting metrics and other

Top languages

Go Java Shell JavaScript

Are Operators owned by Red Hat?

No, they're open source.



Should I use a configMap or a custom resource?

Use a ConfigMap if any of the following apply:

- There is an existing, well-documented config file format, such as a `mysql.cnf` or `pom.xml`.
- You want to put the entire config file into one key of a configMap.
- The main use of the config file is for a program running in a Pod on your cluster to consume the file to configure itself.

DB via cloud managed service

Leverage and expose managed services

AWS Open Source Blog

AWS Service Operator for Kubernetes Now Available 🚀

by Chris Hein | on 05 OCT 2018 | in Amazon Elastic Kubernetes Service, Open Source | Permalink | Comments | Share

The AWS Service Operator is an open source project in developer preview which allows to you manage your AWS resources directly from Kubernetes using the standard Kubernetes CLI, `kubectl`. It does so by modeling AWS Services as Custom Resource Definitions (CRDs) in Kubernetes and applying those definitions to your cluster. This means that a developer can model their entire application architecture from container to ingress to AWS services, backing it from a single YAML manifest. We anticipate that the AWS Service Operator will help reduce the time it takes to create new applications, and assist in keeping applications in the desired state.

Have you ever tried to integrate Amazon DynamoDB with an application running in Kubernetes? How about deploying an

Resources

Open Source at AW
Projects on GitHub

Follow

AWS Open Sourc
AWS Cloud
Facebook
LinkedIn

Search or jump to… /

Pull requests Issues Marketplace Explore

awslabs / **aws-service-operator**

Watch 55 Star 668

Code Issues 66 Pull requests 12 Security Insights

AWS Service Operator allows you to create AWS resources using kubectl.

kubernetes k8s aws operator

183 commits **3** branches **4** releases **13** contributors Ap

Branch: **master** New pull request Create new file Upload files Find File Clo

christopherhein Merge pull request #187 from pauldthomson/master ··· Latest commit

.github	Adding Governance Files for OSS
charts/aws-service-operator	helm chart: sync clusterrole.yaml; add events and patch
cloudformation	fix the error when deploy SQS with dead letter queue to a region does…
cmd/aws-service-operator	Add ability to namespace kubernetes API calls

```yaml
apiVersion: service-operator.aws/v1alpha1
kind: DynamoDB
metadata:
  name: example-table-name
spec:
  hashAttribute:
    name: user_id
    type: S
  rangeAttribute:
    name: created_at
    type: S
  readCapacityUnits: 5
  writeCapacityUnits: 5
```

Config Connector overview

Beta

This product or feature is in a pre-release state and might change or have limited support. For more information, see Product launch stages.

Config Connector is a Kubernetes addon that allows you to manage your Google Cloud Platform (GCP) resources through Kubernetes configuration.

Introduction

Operators make it easy to procure cloud resources, just like any other k8s resource

DB via Service Broker / Catalog

He's dead, Jim (Caution! Achtung!)

README.md

service-catalog

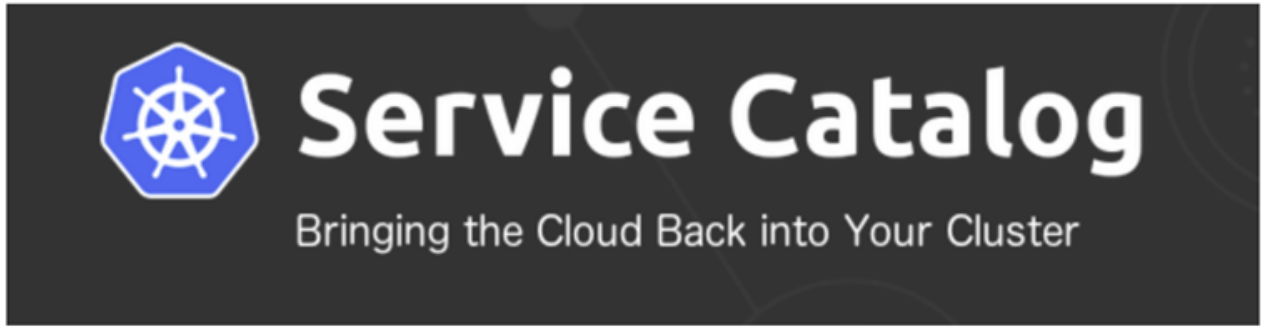

Service Catalog lets you provision cloud services directly from the comfort of native Kubernetes tooling. This project is in incubation to bring integration with service brokers to the Kubernetes ecosystem via the Open Service Broker API.

Documentation

Our goal is to have extensive use-case and functional documentation.

See the Service Catalog documentation on the main Kubernetes site, and svc-cat.io.

For details on broker servers that are compatible with this software, see the Open Service Broker API project's Getting Started guide.

Video links

- Service Catalog Intro

We're on v1.15

Project Status

We are currently working toward a beta-quality release to be used in conjunction with Kubernetes 1.9. See the milestones list for information about the issues and PRs in current and future milestones.

with Kubernetes 1.9.

Operators accomplish the same thing,

But they're easier

In conclusion

Summary

Storage

- We have **not** used the word "Storage" much
- A lot of our techniques are based on ideas from parts 1 & 2
- Data must be stored somewhere
- Persistent Volumes (PV), PVCs, and other concepts are hidden here, but they are an important part of everything covered today

Kubernetes Storage Terminology

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

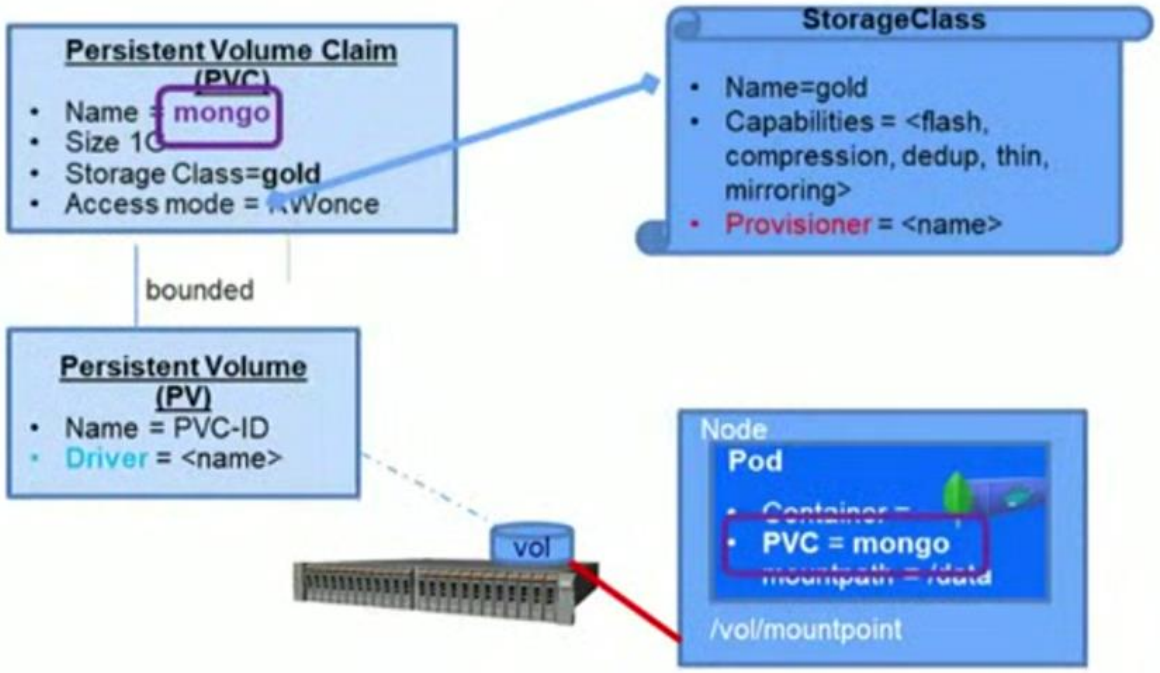

https://kubernetes.io/docs/concepts/storage/volumes/

Storage Class

To achieve dynamic volume creation, the admin must define a k8s **StorageClass** (e.g : gold, silver).

Provision a volume

1. The user creates a claim for volume (**PVC**).
2. The "Provisioner" (vendor specific) listens to new PVC requests, and dynamically creates the volume on the storage system (if no PV already matched)
   - The **PV** is created with "Driver" setting. The Driver(vendor specific) handles the volume attach\detach to the node.

Create a stateful POD

1. The user creates a **POD** with the new PVC.
2. K8s triggers the "Driver" in order to attach the PV to the node.
3. The volume is now mapped and mounted to the node.
4. k8s starts the POD with the PV mounted to /data inside the container.

✓ The stateful container is UP

Summary

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

Security is paramount

- There are a lot of wonderful online resource
- Read best practices on Kubernetes.io docs
- Use HashiCorp **Vault** or your cloud's **KMS** for secret
- Learn about Kubernetes and Security from:
  - KataCoda in-browser tutorials
  - The past SNIA webcasts in this series
- Buy or exchange your info for the ***Kubernetes Security*** book

Summary

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

Five ways to run *Stateful* workloads on Kubernetes

1. on VM (easier)
2. on k8s via StatefulSet (harder)
3. on k8s via Operator (harder)
4. via Cloud Managed Service (easier)
5. via Service Broker (harder)

Summary

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

Five ways to run *Stateful* workloads on Kubernetes

1. on VM (easier)
2. on k8s via StatefulSet (harder)
3. on k8s via Operator (harder)
4. via Cloud Managed Service (easier)
5. ~~via Service Broker (harder)~~

Kelsey Hightower 

@kelseyhightower

Kelsey's guide to running traditional databases on Kubernetes. Strongly consider using a managed service.

11:56 AM · Jan 20, 2017 · Twitter Web Client

Resources & Links, Part 1

1. https://twitter.com/kubernetesio/status/840257886202683392
2. https://www.youtube.com/watch?v=4x1r3Osu1Kg
3. https://twitter.com/kelseyhightower/status/963413508300812295?lang=en
4. https://kubernetes.io/docs/concepts/configuration/secret/
5. https://github.com/godaddy/kubernetes-external-secrets
6. https://learn.hashicorp.com/vault
7. https://www.katacoda.com/courses/kubernetes https://kubernetes-security.info/
8. https://info.aquasec.com/kubernetes-security https://cloud.google.com/blog/products/gcp/kubernetes-best-practices-mapping-external-services
9. https://cloud.google.com/blog/products/gcp/kubernetes-best-practices-mapping-external-services
10. https://kubernetes.io/docs/tasks/run-application/run-single-instance-stateful-application/
11. https://kubernetes.io/docs/tutorials/stateful-application/mysql-wordpress-persistent-volume/
12. https://cloud.google.com/blog/products/databases/to-run-or-not-to-run-a-database-on-kubernetes-what-to-consider

Resources & Links, Part 2

1. https://coreos.com/blog/introducing-operators.html
2. https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/
3. https://coreos.com/blog/introducing-operators.html
4. https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/
5. https://github.com/etcd-io/etcd/blob/master/Procfile.v2
6. https://docs.couchbase.com/server/6.0/backup-restore/enterprise-backup-restore.html
7. https://operatorhub.io
8. https://github.com/operator-framework
9. https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/#should-i-use-a-configmap-or-a-custom-resource
10. https://aws.amazon.com/blogs/opensource/aws-service-operator-kubernetes-available/
11. https://github.com/awslabs/aws-service-operator
12. https://cloud.google.com/config-connector/docs/overview
13. https://github.com/kubernetes-sigs/service-catalog
14. https://twitter.com/kelseyhightower/status/822488055709712384?lang=en

Find all these links in our blog at: http://bit.ly/KubeLinks

After This Webcast

SNIA CSTI | CLOUD STORAGE TECHNOLOGIES

- Please rate this webcast and provide us with feedback
- This webcast and a PDF of the slides will be posted to the SNIA Cloud Storage Technologies Initiative website and available on-demand at https://www.snia.org/forum/csti/knowledge/webcasts
- A full Q&A from this webcast will be posted to the SNIA Cloud blog: www.sniacloud.com/
- Follow us on Twitter @SNIACloud

Questions

Thank You