


OpenEBS

Architecture and Design

v0.5 - Oct 2017



Prerequisites

Docker and Kubernetes - Namespaces, RBAC, Storage Classes, CRD, Dashboard

Running Stateful Workloads with PV, PVC and Dynamic Provisioner

Introduction to OpenEBS - Container Native Storage

Prometheus, Node Exporter, Grafana

Opentracing - Jaeger

Design Goals and Constraints

Storage optimized for Containerized Applications


Stable, Secure and Scalable - Horizontally scalable to millions of Containers, Fault tolerant and Secure by default

Seamless integration into any private and public cloud environments.
Vendor independent.

Non-disruptive software upgrades

Easy to setup. Low entry barrier. Developer and Operators Friendly.


Stateful Apps using OpenEBS Volumes



OpenEBS Control Plane

OpenEBS Volume (ov) containers will be managed by:

- Container Orchestrators like K8s and
- OpenEBS control plane services that specialize in storage operations.




The source code for the volume containers -- target and replica also known as frontend and backend is located under **openebs/jiva**

OpenEBS Control Plane Architecture


(1) Integrators



(2) Control Plane (Cluster Services) (aka. K8s master)




(3) Control Plane (Node Agents) (aka K8s DaemonSets)



- CO integrated
- Other Frameworks
- OpenEBS Services


OpenEBS Control Plane converged K8s Cluster



The source code for the control plane components is located in mainly **openebs/maya** repository

OpenEBS Initialization

1. Load Configuration



2. Launch - Cluster Services

maya-apiserver

maya-provisioner

maya-mulebot

maya-connect

prometheus-maya

grafana-maya

kube-dashboard*

3. Launch - Node Services

node_exporter (n)


maya-agent (n)

jaegar-maya (n)

4. Verify Status via


mayactl

Storage Schema - Raw Storage



Raw Storage (RS).
logical representation
of the underlying
storage attached to a
node.

Storage Schema - Storage Backend



Storage Backend (SB) represents a carved/partitioned unit of storage that will be allocated to OpenEBS Volume Replica.

Raw Storage (RS), representing a single instance of the underlying storage type.




Storage Schema - Usage

1. DevOps admin creates the SBA and RSA yaml files and feeds into the K8s DB.

2. maya-agent will look for the RSA that are assigned to it and initializes the adaptor discovery logic.

3. For the discovered disks, maya-agent will create the RS and may also create some SBs (like nvmedimm namespaces) using the SBAs (zpool, lvm)

4. In K8s/maya-apiserver will convert SBs into PVs before attaching to the OV replica containers.



Storage Backend (SB)

- **Name**
- **Node** - where it is currently attached (local disks) or mounted (remote -disks).
- **Type** - local disks, iscsi-disk, gpd, aws, nvme, nvdimm
- **ParentSB** (none - or is this accessed as partition)
- **AccessLayerType** : **ext4, zvol, raw-block.**
- **Params** specific to type
- **UsedBy**

YAML

Storage Backend Adaptors (SBA)

- **Name**
- **Node Filter** - which nodes does this apply to
- **Type** - local disks, ext4, zpool, lvm
- **Params** specific to type:


Raw (RS) -- actual disk entity or a partition -- raw storage that needs managed per node

YAML


Raw Storage Adaptor (RSA)

- **Name**
- **Node** - where it is required
- **Type** - nvme, nvdimm, aws etc.,
- Parameters:
 - Filesystem (with directory like /var/openebs)
 - local disks
 - Block device filters
 - K8s - local disk manager
 - OpenEBS iSCSI - peer node disks (?)
 - Auto Provisioner-gpd, aws, external SAN
 - Credentials
 - Zonal


maya-agent




maya-apiserver



maya-provisioner



maya-connect (TBD)



maya-mulebot (TBD)


Workflow

State Diagrams

Sequence Diagrams



OpenEBS Volume - State Diagram



User Interface

Kubernetes Dashboard Extensions



☰ Overview
Cluster

- Namespaces
- Nodes
- Persistent Volumes
- Roles
- Storage Classes

Namespace

default ▾

Overview
Workloads

- Daemon Sets
- Deployments
- Jobs
- Pods
- Replica Sets
- Replication Controllers
- Stateful Sets

Discovery and Load Balancing

Pods

| Name | Node | Status | Restarts | Age | CPU (cores) | Memory (bytes) |
|---|--------------|---------|----------|---------|-------------|----------------|
| ✓ maya-apiserver-3416621614-c... | minikube-dev | Running | 0 | 6 hours | 0 | 7.863 Mi |
| ✓ openebs-provisioner-42306262 | minikube-dev | Running | 0 | 6 hours | 0 | 5.543 Mi |

Deployments

| Name | Labels | Pods | Age | Images |
|------------------------------------|---------------------------|-------|---------|---------------------------------|
| ✓ maya-apiserver | name: maya-apiserver | 1 / 1 | 6 hours | openebs/m-apiserver:0.4.0 |
| ✓ openebs-provisioner | name: openebs-provisioner | 1 / 1 | 6 hours | openebs/openebs-k8s-provisioner |

Replica Sets

| Name | Labels | Pods | Age | Images |
|---|--|-------|---------|---------------------------------|
| ✓ maya-apiserver-3416621614 | name: maya-apiserver pod-template-hash: 3416621614 | 1 / 1 | 6 hours | openebs/m-apiserver:0.4.0 |
| ✓ openebs-provisioner-4230626287 | name: openebs-provisioner pod-template-hash: 4230626287 | 1 / 1 | 6 hours | openebs/openebs-k8s-provisioner |

Cluster


- Namespaces
- Nodes
- Persistent Volumes
- Roles
- Storage Classes

Namespace

default

Overview
Workloads

- Daemon Sets
- Deployments
- Jobs
- Pods**
- Replica Sets
- Replication Controllers
- Stateful Sets
- Discovery and Load Balancing

CPU usage

Memory usage

Details

Name: maya-apiserver-3416621614-c2116

Namespace: default

Labels: name: maya-apiserver pod-template-hash: 3416621614

Annotations: Created by: ReplicaSet maya-apiserver-3416621614

Creation time: 2017-10-12T11:46

Status: Running

Network

Node: minikube-dev

IP: 172.17.0.6

Containers

☰ Shell

Cluster

Namespaces

Nodes

Persistent Volumes

Roles

Storage Classes

Namespace

default ▾

Overview

Workloads

Daemon Sets

Deployments

Jobs

Pods

Replica Sets

Replication Controllers

Shell in maya-apiserver ▾ in maya-apiserver-3416621614-c2116

```
root@maya-apiserver-3416621614-c2116:/# maya help
Usage: maya [--version] [--help] <command> [<args>]

Available commands are:
  snapshot  Provides operations related to snapshot of a Volume
  version   Prints version and other details relevant to maya
  volume    Provides operations related to a Volume

root@maya-apiserver-3416621614-c2116:/#
```

Cluster

- Namespaces
- Nodes
- Persistent Volumes
- Roles

Storage Classes

Namespace

default ▾
Overview
Workloads

- Daemon Sets
- Deployments
- Jobs
- Pods
- Replica Sets
- Replication Controllers
- Stateful Sets

Discovery and Load Balancing

Storage Classes

| Name | Labels | Provisioner | Parameters | Age |
|-----------------------------------|--------------------------------|------------------------------|---|---------|
| openebs-basic | - | openebs.io/provisioner-iscsi | pool: hostdir-var replica: 2 size: 5G | 6 hours |
| openebs-cassandra | - | openebs.io/provisioner-iscsi | pool: hostdir-var replica: 2 size: 5G | 6 hours |
| openebs-jupyter | - | openebs.io/provisioner-iscsi | pool: hostdir-var replica: 2 size: 5G | 6 hours |
| openebs-kafka | - | openebs.io/provisioner-iscsi | pool: hostdir-var replica: 2 size: 10G | 6 hours |
| openebs-mongodb | - | openebs.io/provisioner-iscsi | pool: hostdir-var replica: 2 size: 5G | 6 hours |
| openebs-percona | - | openebs.io/provisioner-iscsi | pool: hostdir-var replica: 2 size: 5G | 6 hours |
| openebs-redis | - | openebs.io/provisioner-iscsi | pool: hostdir-var replica: 2 size: 5G | 6 hours |
| openebs-zk | - | openebs.io/provisioner-iscsi | pool: hostdir-var replica: 2 size: 5G | 6 hours |
| standard | addonmanager.kubernetes.io/... | k8s.io/minikube-hostpath | - | a day |

Cluster[Namespaces](#)[Nodes](#)[Persistent Volumes](#)[Roles](#)[Storage Classes](#)

Namespace

[default](#)**Overview**

Workloads

[Daemon Sets](#)[Deployments](#)[Jobs](#)[Pods](#)[Replica Sets](#)[Replication Controllers](#)[Stateful Sets](#)

Discovery and Load Balancing

Details

Name: openebs-percona**Annotations:** `last applied configuration`**Creation time:** 2017-10-12T11:46**Labels:** -**Provisioner:** openebs.io/provisioner-iscsi**Parameters:** `pool: hostdir-var` `replica: 2` `size: 5G`

TODO - Could list OpenEBS Volumes created using this Storage class

[Daemon Sets](#)
[Deployments](#)
[Jobs](#)
[Pods](#)
[Replica Sets](#)
[Replication Controllers](#)
[Stateful Sets](#)
[Discovery and Load Balancing](#)
[Ingresses](#)
[Services](#)
[Config and Storage](#)
[Config Maps](#)
[Persistent Volume Claims](#)
[Secrets](#)

Details

Name: percona

Namespace: default

Labels: name: percona

Annotations: last applied configuration

Creation time: 2017-09-14T07:43

Status: Running

Network

Node: kubeminion-02

IP: 10.36.0.3

Persistent Volume Claims

| Name | Volume | Labels | Age |
|--|--|--------|-----------|
|  openebs-vol1-claim | pvc-741c37cc-9920-11e7-8fdc-021c6f7dbe9d | - | 7 minutes |

default ▾

Overview

Workloads

Daemon Sets

Deployments

Jobs

Pods

Replica Sets

Replication Controllers

Stateful Sets

Discovery and Load Balancing

Ingresses

Services

Config and Storage

Config Maps

Persistent Volume Claims

Secrets

About

Details

Name: demo-vol1-claim**Namespace:** default**Annotations:** control-plane.alpha.kubernetes.io/leader: {"holderidentity":"77a9d034-af43-11e7-90a0-0242ac110007","leaseDurationSeconds":15,"acquireTime":"2017-10-12T18:34:00Z","renewTime":"2017-10-12T18:34:00Z","remainingTime":15} [last applied configuration](#) pv.kubernetes.io/bind-completed: yes pv.kubernetes.io/bound-by-controller: yes volume.beta.kubernetes.io/storage-provisioner: openebs.io/provisioner-iscsi**Creation time:** 2017-10-12T18:34**Status:** Bound**Volume:** pvc-f9fc9b42-af7b-11e7-9adc-080027b55b5e**Capacity:** {"storage": "5G"}**Access modes:** ReadWriteOnce**Storage class:** openebs-percona

Cluster

Namespaces

Nodes

Persistent Volumes

Roles

Storage Classes

Namespace

default

Overview

Workloads

Daemon Sets

Deployments

Jobs

Pods

Replica Sets

Replication Controllers

Stateful Sets

Allocated resources

CPU



Memory



Up Time


Since:

2017 Sep 14 14:56 IST

(1 hour 21 minutes)

Current Time: 4:17 PM

Storage Capacity



131



Should include outstanding alerts at the top of the page.

Cluster

Namespaces

Nodes

Persistent Volumes

Roles

Storage Classes

Namespace

default

Overview

Workloads

Daemon Sets

Deployments

Jobs

Pods

Replica Sets

Replication Controllers

Stateful Sets

Discovery and Load Balancing

Details

Name: pvc-f9fc9b42-af7b-11e7-9adc-080027b55b5e**Annotations:** openEBSProvisionerIdentity: minikube-dev pv.kubernetes.io/provisioned-by: openebs.io/provisioner-iscsi**Creation time:** 2017-10-12T18:34**Status:** Bound**Claim:** default/demo-vol1-claim**Reclaim policy:** Delete**Access modes:** ReadWriteOnce**Storage class:** openebs-percona**Capacity:** 5G**Reason:** -**Message:** -

Persistent volume source

ISCSI

Target portal: 10.0.0.16:3260**IQN:** iqn.2016-09.com.openebs.jiva:pvc-f9fc9b42-af7b-11e7-9adc-080027b55b5e**Lun:** 1

Cluster

Namespaces

Nodes

Persistent Volumes

Roles

Storage Classes

Namespace

default

Overview

Workloads

Daemon Sets

Deployments

Jobs



Pods

Replica Sets

Replication Controllers

Stateful Sets

Pods

| Name | Node | CPU (cores) | Memory (bytes) |
|--|----------------|---|--|
|  pvc-741c37cc-9920-11e7-8fdc-021c6f7db... | kubeminiion-01 |  0 |  320 Ki |
|  pvc-741c37cc-9920-11e7-8fdc-021c6f7db... | kubeminiion-01 |  0 |  764 Ki |
|  pvc-741c37cc-9920-11e7-8fdc-021c6f7db... | kubeminiion-02 |  0 |  760 Ki |

Storage:
avail/used

On each
pod.

Services

| Name | Labels | Cluster IP | Internal endpoints | External endpoints | Age |
|--|--|--------------|--|--------------------|------------|
|  pvc-741c37cc-9920... | openebs/controll... openebs/volume... vsm: pvc-741c37... | 10.106.44.21 | pvc-741c37cc-9920... pvc-741c37cc-9920... pvc-741c37cc-9920... pvc-741c37cc-9920... | - | 15 minutes |

Should include Events at the bottom of the page.

Cluster



- [Namespaces](#)
- [Nodes](#)
- [Persistent Volumes](#)
- [Roles](#)
- [Storage Classes](#)

Namespace

default ▾

Overview
Workloads

- [Daemon Sets](#)
- [Deployments](#)
- [Jobs](#)
- [Pods](#)
- [Replica Sets](#)
- [Replication Controllers](#)
- [Stateful Sets](#)
- [Discovery and Load Balancing](#)

CPU usage

Memory usage i

Details

Name: pvc-f9fc9b42-af7b-11e7-9adc-080027b55b5e-rep-1715657804-r0bhq

Namespace: default

Labels: openebs/replica: java-replica pod-template-hash: 1715657804 vsm: pvc-f9fc9b42-af7b-11e7-9adc-080027b55b5e

Annotations: Created by: ReplicaSet pvc-f9fc9b42-af7b-11e7-9adc-080027b55b5e-rep-1715657804

Creation time: 2017-10-12T18:34

Status: Running

Network

Node: minikube-dev

IP: 172.17.0.8

Containers

Cluster

- Namespaces
- Nodes
- Persistent Volumes
- Roles
- Storage Classes

Namespace

default

Overview

Workloads

- Daemon Sets
- Deployments
- Jobs
- Pods**
- Replica Sets
- Replication Controllers
- Stateful Sets

Discovery and Load Balancing

Containers

pvc-f9fc9b42-af7b-11e7-9adc-080027b55b5e-rep-con

Image: openebs/jiva:0.4.0

Environment variables: -

Commands: launch

Args: replica

--frontendIP

10.0.0.16

--size

5G

/openebs

Conditions

| Type | Status | Last heartbeat time | Last transition time | Reason | Message |
|--------------|--------|---------------------|----------------------|--------|---------|
| Initialized | True | - | 21 minutes | - | - |
| Ready | True | - | 19 minutes | - | - |
| PodScheduled | True | - | 21 minutes | - | - |