LF NETWORKING

LFN Developer & Testing Forum

High Availability Deployment of ONAP Components

Amar Kapadia, Aarna Networks

https://www.linkedin.com/in/amarkapadia/

Agenda

- Subset of ONAP projects used
- What is HA on Kubernetes
- Creating a HA Kubernetes cluster on bare-metal
- Creating a HA Kubernetes cluster on AWS
- Changes in Helm charts for deploying on HA setup
- Storageclass and Shared PV in Kubernetes
- Deploying a subset of ONAP (eg., CDS and Camunda) on HA setup
- Demo video of a HA setup
- Future plans and challenges

Subset of ONAP Projects Used

For this setup and demo we used the below ONAP projects.

- CDS
- Camunda
- Mariadb

High Availability on Kubernetes

- Multiple master nodes
- Load Balancer using HA Proxy and KeepaliveD for ACTIVE/PASSIVE setup
- Stacked ETCD cluster across nodes for kubernetes data
- Multiple Worker nodes to create multiple replicas of application

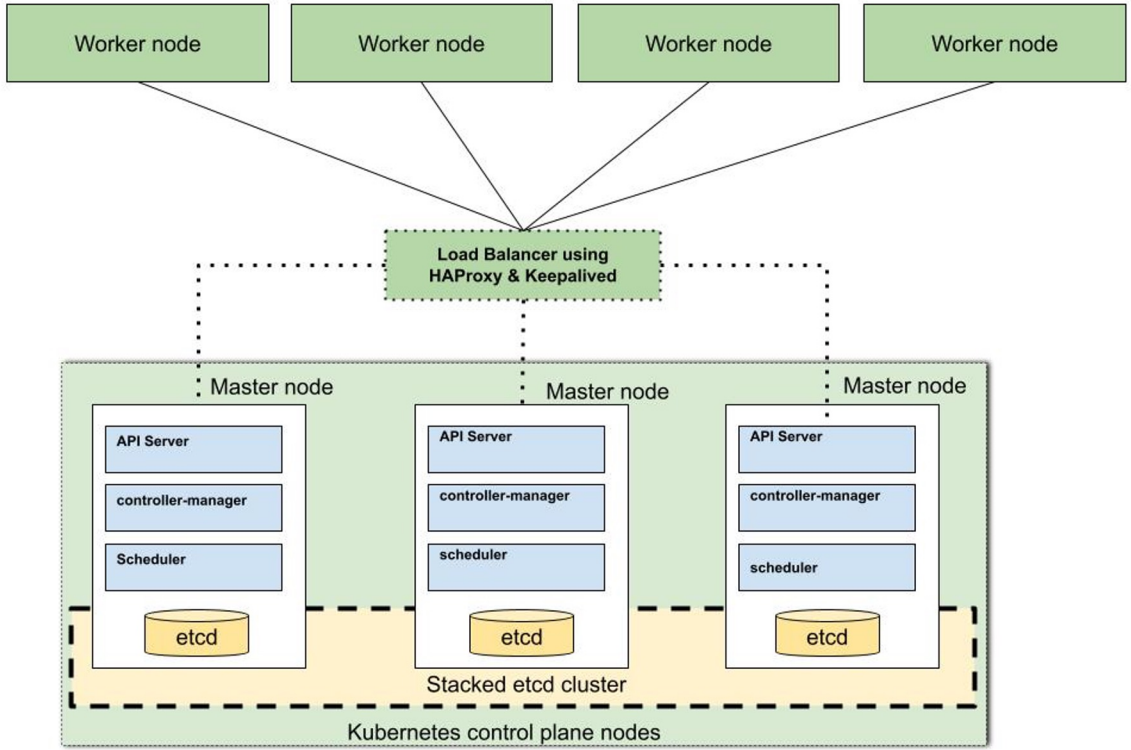

HA Proxy With Keepalive Daemon

- Keepalived uses **Virtual Router Redundancy Protocol (VRRP) as an election protocol**
- **It determine which master or proxy node holds the virtual IP**
- Keepalived virtual IP manager implements a set of checkers to dynamically and adaptively maintain and manage a load balanced server pool according to its health
- For our setup we configured the keepalived and HAPROXY on each master node
- In case of a master node failure the other node takes over as API server

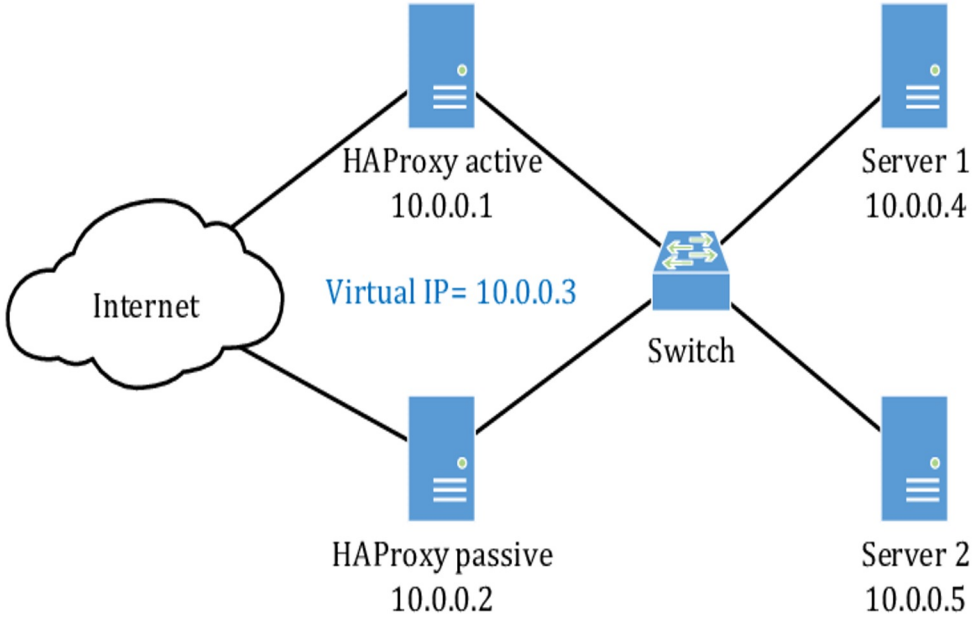

HA Cluster on AWS

- KOPS uses operator to deploy
- We can build a multi zone cluster on AWS
- Information about the cluster config is stored on S3 bucket

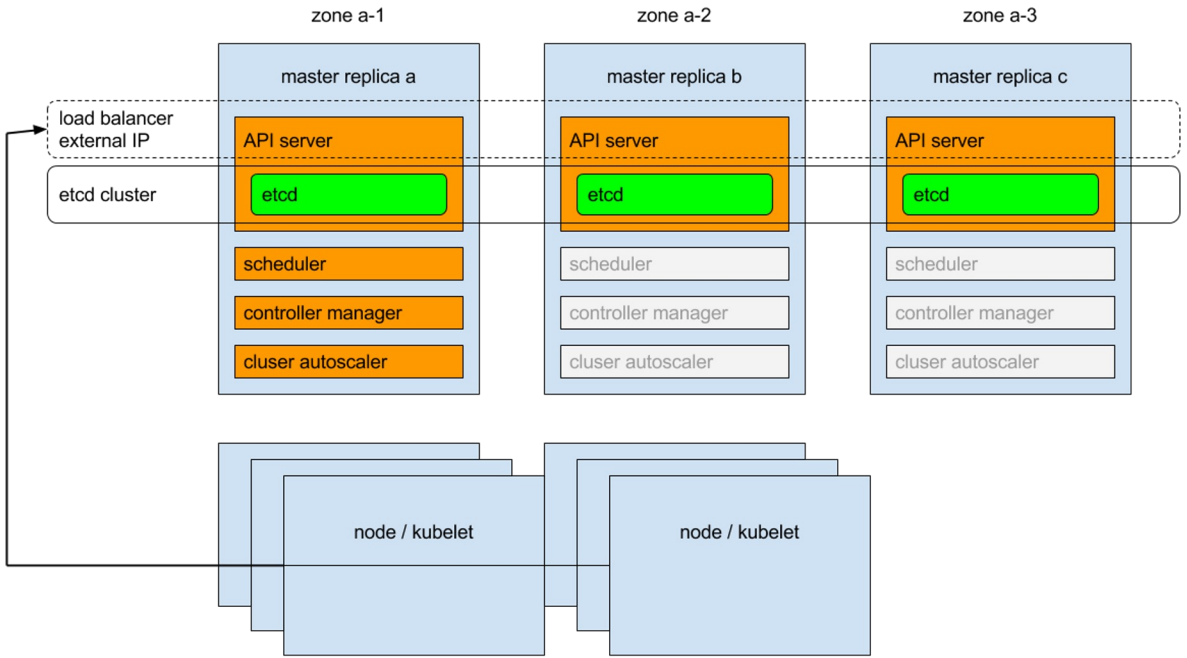

HA Cluster on Bare Metal Servers

- Require at least 3 VMs/servers
- Deploy kubernetes using kubeadm
- Deploy and configure HAPROXY + KEEPALIVED on three servers
- Requires VIRTUAL IP for failover
- Currently done using a script
  - It creates 7 VM’s (3 master and 4 worker nodes) using virsh+qemu
  - Installs all the tools in the above VM’s
  - Configures HAPROXY+KEEPALIVED
  - Deploys and configures KUBEADM
- Testing is done by abruptly shutting down one of the VM

Stateful Application (PV/Storageclass)

- Some application require to store data (e.g. DB, CAMUNDA, CDS etc.)
- Data is stored in Persistent Volumes
- Currently we provide mechanism for NFS Storage i.e. PV is created on the node on NFS server
- Since NFS is not in HA mode, in case of NFS server failure, all data is lost
- Having multiple replicas helps, however consistency of data has to be maintained
    - Replication maintained by app (eg. mariadb has synchronous replication and can recover in case of failure)
    - Replication not done by app. Option to used shared redundant storage (CEPH), each pod can mount same PV (**ReadWriteMany option**) (CDS, Camunda frameworks to support this)
    - Application level HA not solved at this point, we depend on the application e.g. CDS, Camunda – we can use community help/collaboration on this

CEPH and Distributed Storage

- CEPH deploys on multiple servers
- Ability to retrieve data in case of server failure due to distributed storage
- Supports Kubernetes deployment

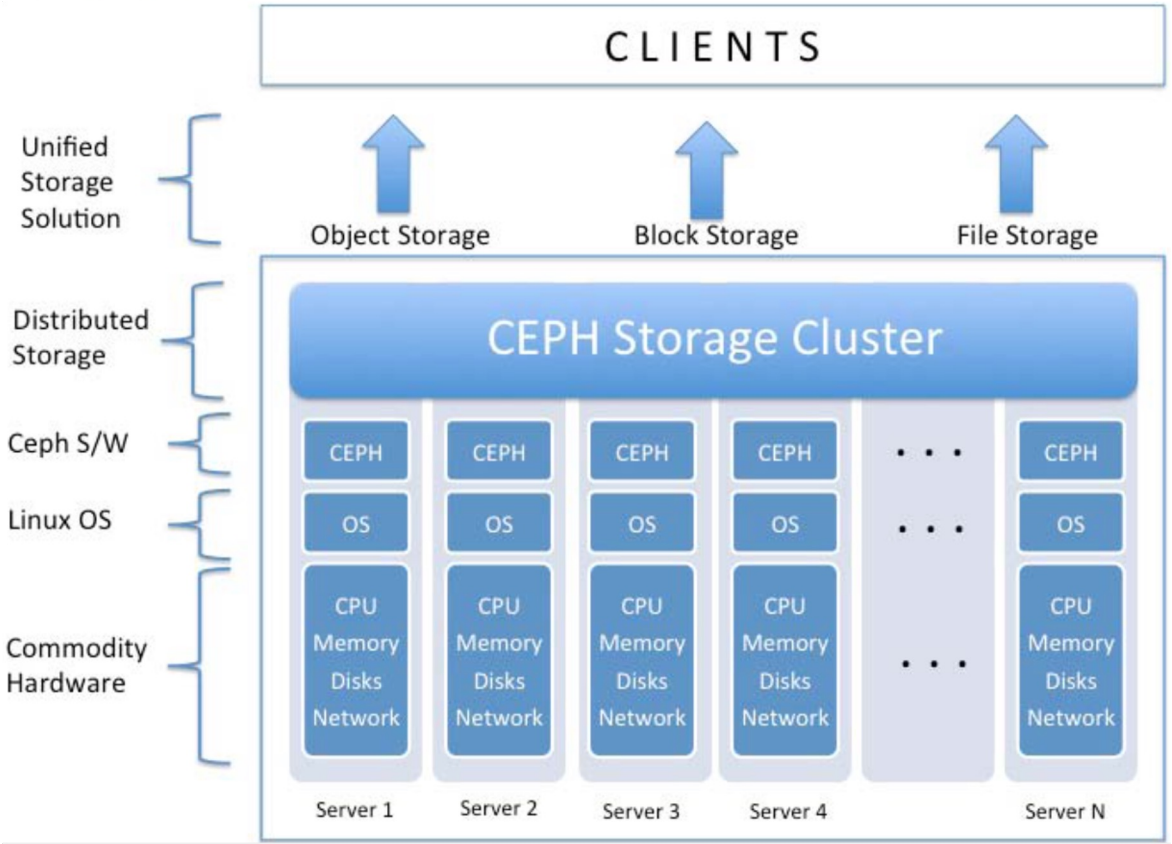

CEPH deployment on K8s using ROOK

- We have single script to deploy ROOK/CEPH on K8s cluster
- It adds disk (for cloud VMs), configures ROOK, creates a StorageClass on K8s
- PVs from the SC can then be used by pods

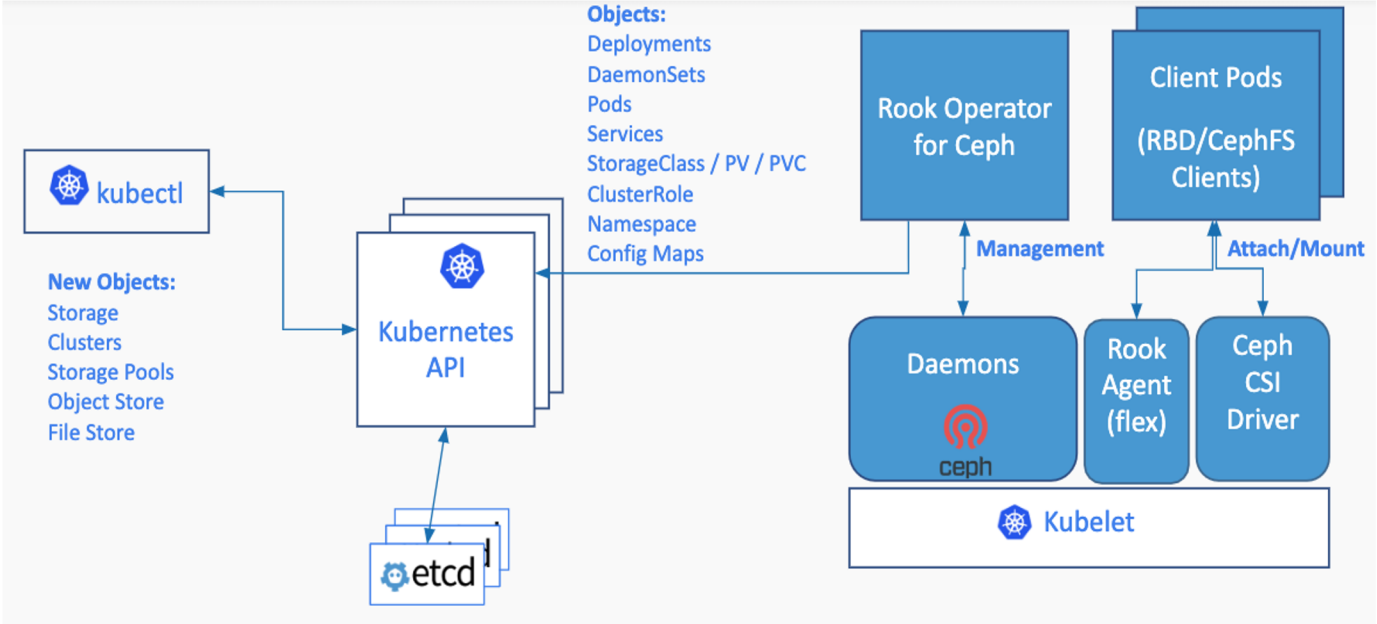

ONAP (CDS + Camunda) Changes

- Cluster will require twice the resources (compared to no HA deployment)
- Additional storage partition on each NODE (to store CEPH DATA)
- Make sure all pods have at least 2 replicas
- We should use CEPH PV instead of NFS (change in Helm charts)
- All deployment to have anti affinity rules (this will distribute the pods across nodes)
- CAMUNDA/CDS to use shared PV from CEPHFS

Testing HA on ONAP Components

- Deploy a subset of ONAP (e.g., CDS + Camunda) on a HA cluster with multiple nodes
- A script to shutdown/destroy one of the NODE. While the node is being brought down, following to be done in parallel:
  a. CDS Workflow test: Run a script that connects with CDS to run CBA
  b. This script triggers CDS CBA which in turn will trigger a Camunda Workflow
  c. In all we access CDS, Camunda & mariadb pods
  d. Above steps are done in a LOOP for multiple iterations
- Ensure that all services are available and none of the above scripts is failing while the nodes were abruptly brought down
- All the above scripts use REST APIs to access the K8 services

Demo Video

https://drive.google.com/file/d/1QP_IMIwOQSjAGNBx3Lwia0qv6Xa4Apuc/view?usp=sharing

Future Work

- Add and test more components of ONAP e.g. A&AI
- Support more StorageClass
- Test with Multiple Availability Zones (HA across DC’s)
- Support newer features of K8 on newer versions (1.25 onwards)

Thank You!