



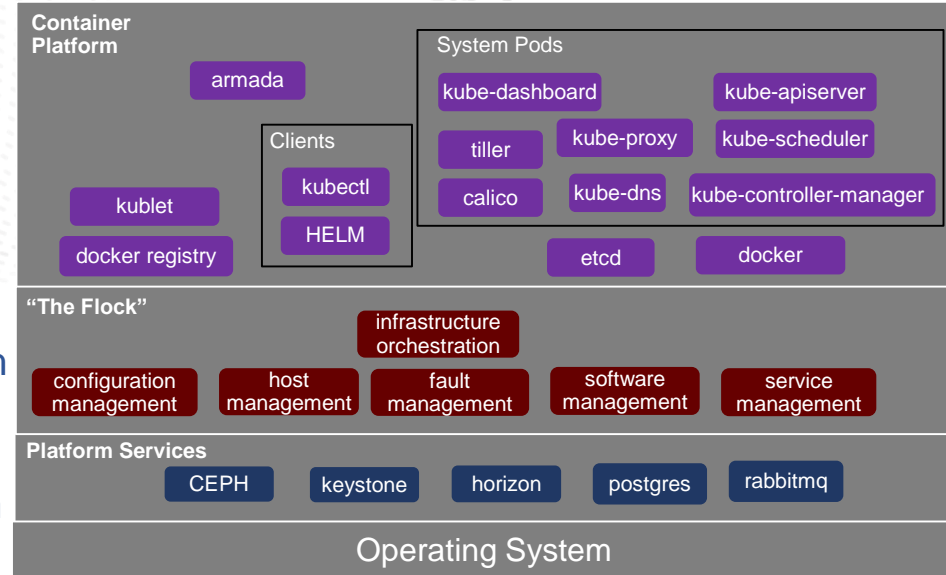
# Containerizing the Infrastructure Discussion

# Introduction

- StarlingX seed code provided a hardened openstack platform
- Evolution plan presented at the Vancouver summit is to move to a cloud native kubernetes based platform
- Along with providing an initial k8s hosting environment for applications, the infrastructure including openstack will also be containerized
- Infrastructure containerization would be done in phases – evolution vs revolution
- Containerization of openstack and dependencies will be the initial focus

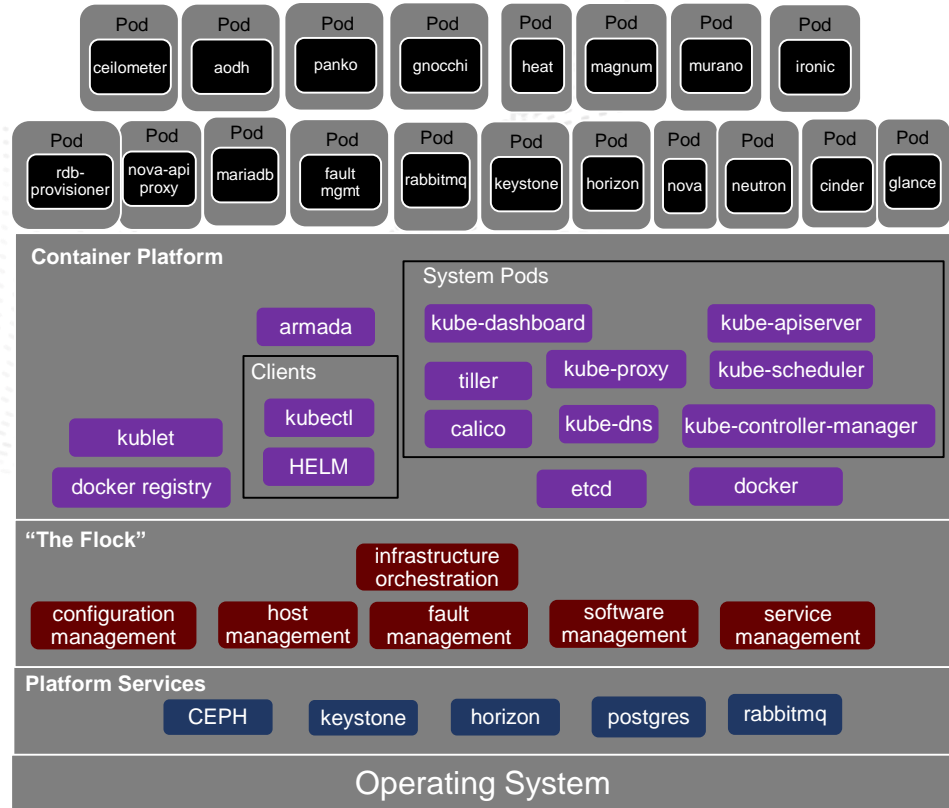
# Container Platform

- High availability k8s master configuration on two nodes
  - Run on existing StarlingX controller nodes
  - Deployed by StarlingX system configuration
- Calico CNI plugin
- Docker runtime
- CEPH as persistent storage backend
  - Leverage existing bare metal CEPH cluster, extending CEPH support to 1 and 2 nodes configurations
- Authentication/authorization of k8s API with keystone
- Local docker image registry, authentication with keystone
- HELM as the package manager
- Airship armada for orchestrating multiple Helm charts (ex. openstack)



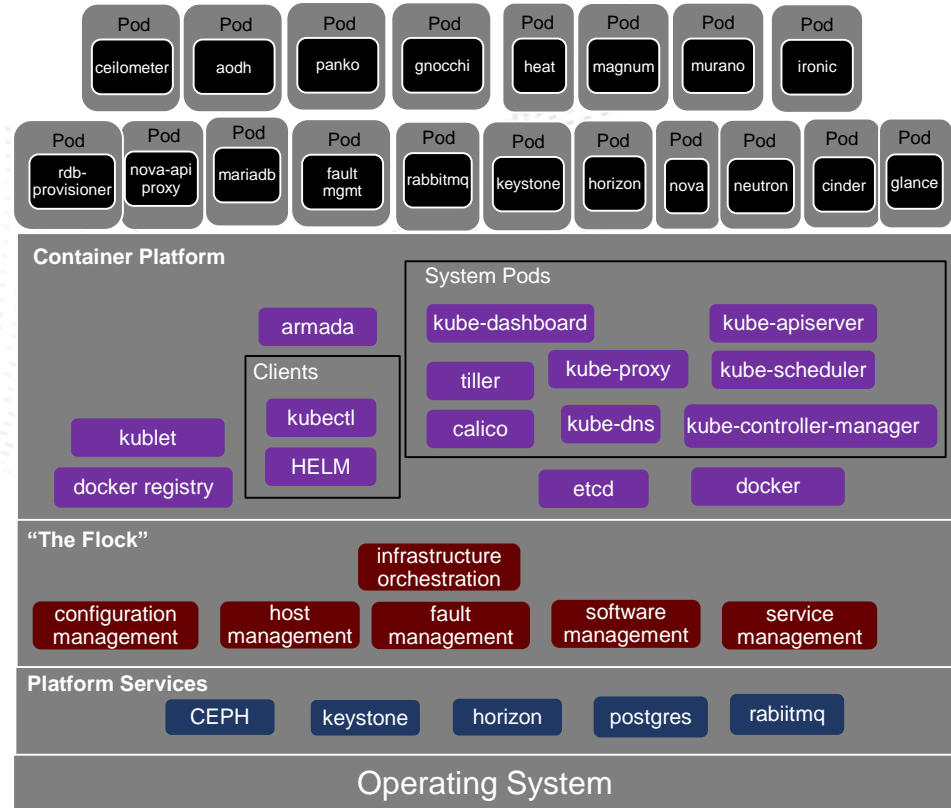
# Initial Infrastructure Containerizing

- The focus for this release is the containerizing of the openstack services
  - All current openstack services supported by StarlingX
    - Pike based
- Nova-api proxy from stx-nfv will also be containerized
- A rbd-provisioner pod will be introduced to enable pvc's from the ceph cluster
- An instance of the fault-management service will be containerized for openstack alarming
- Keystone container will run in the openstack namespace for the openstack services only



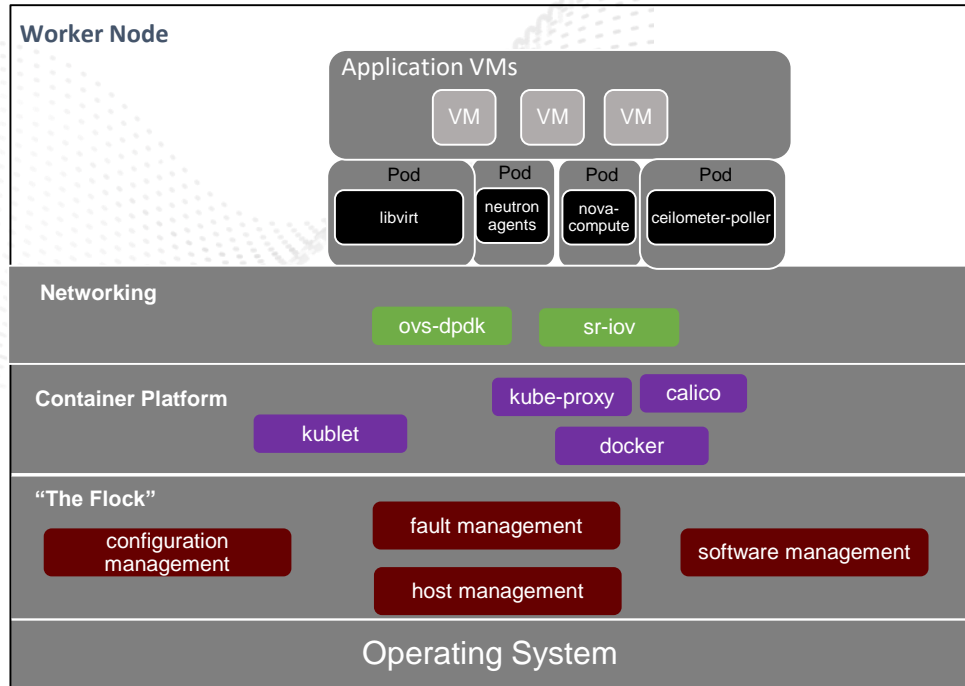
# Initial Infrastructure Containerizing

- An instance of keystone, rabbit and postgres will run on bare metal and be used for “the flock” services:
- An instance of horizon will run on bare metal for “the flock” services only, a containerized instance for openstack



# Initial Infrastructure Containerizing

- Current StarlingX compute nodes are re-purposed as k8s worker nodes
- Nova-compute and libvirt are containerized but are privileged containers with full access to the host
- Neutron agents are containerized but are privileged containers with full access to the host
- OVS-DPDK is not virtualized
  - No HELM chart currently available



# Deployment/Life Cycle Management

- Openstack helm will be used as a starting point
  - Extensions are required with the plan to work with the openstack helm team to upstream.
- Armada will be used to orchestrate the set of HELM charts that comprise openstack
- System configuration is extended to provide new API's to:
  - Upload application
    - Download docker images to local registry
    - Validate and install the helm charts
    - Install armada manifests templates
    - Invoke sysinv api to assign node labels based on provided mapping
  - Install application
    - For each chart in the manifest it will query the system overrides for that chart and write them to a file in a format suitable for armada.
    - Run "armada apply" with the "--values" option to specify our system overrides file(s).
  - Update application
    - Allow the end user to specify overrides to particular charts, and then trigger an update via the "armada apply" mechanism

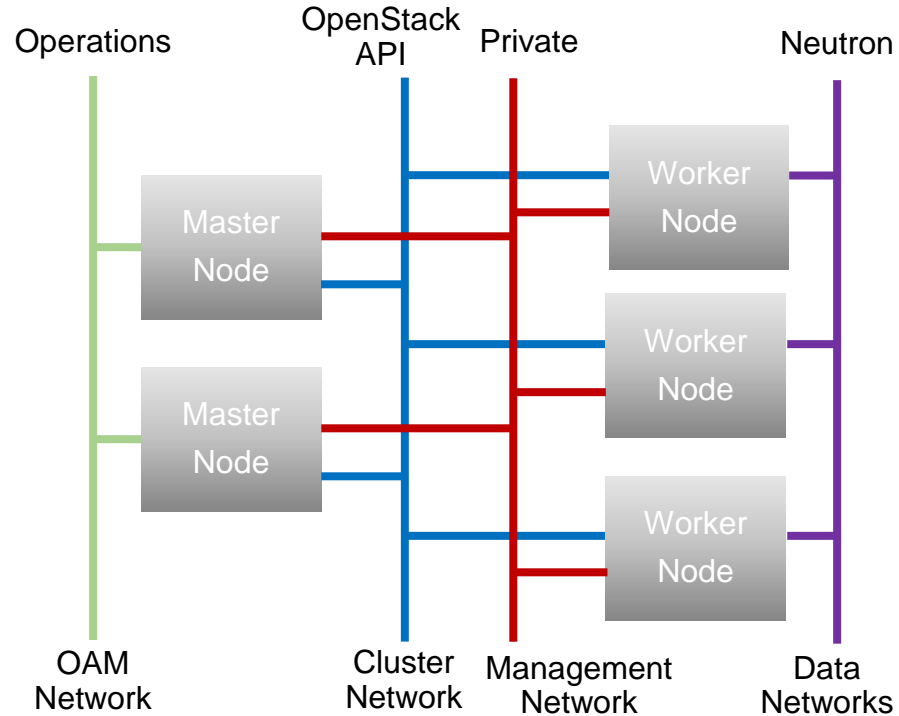
# Docker Images

- The build system will have to be enhanced to generate StarlingX based docker images for the containerized services
- A public repository will be required to host these docker images



# Containerized OpenStack Network Architecture

- Management network would be for platform services only and isolated from the cluster
- OpenStack services exposed on cluster network via an K8S ingress controller
  - Demarcation point for public APIs and TLS/SSL termination
- NFV-VIM APIs need to be accessible from both the OAM and Cluster networks
  - Demarcation point for public APIs and TLS/SSL termination via haproxy
- Platform services need access to several OpenStack APIs
  - VIM: neutron, nova
  - Sysinv: neutron, nova
- CEPH service must be accessible from the cluster network



# Software Updates

- Software updates for containerized services will drive change in strategy
- A software update will be the following:
  - A new docker image for a service(s) would be delivered to the system
  - A armada update would be invoked to update the applicable charts

# Future Release Candidates

## Infrastructure Services:

- Containerize “the flock” services
- Containerize OVS-DPDK
- Containerize CEPH

## Container Infrastructure:

- Accelerated container networking with sr-iov and dpdk
- Multi-tenancy support for containers
- Support for additional container runtimes including kata containers

