
Airship Integration Documentation

Release 0.1.0

Airship Authors

Apr 26, 2019

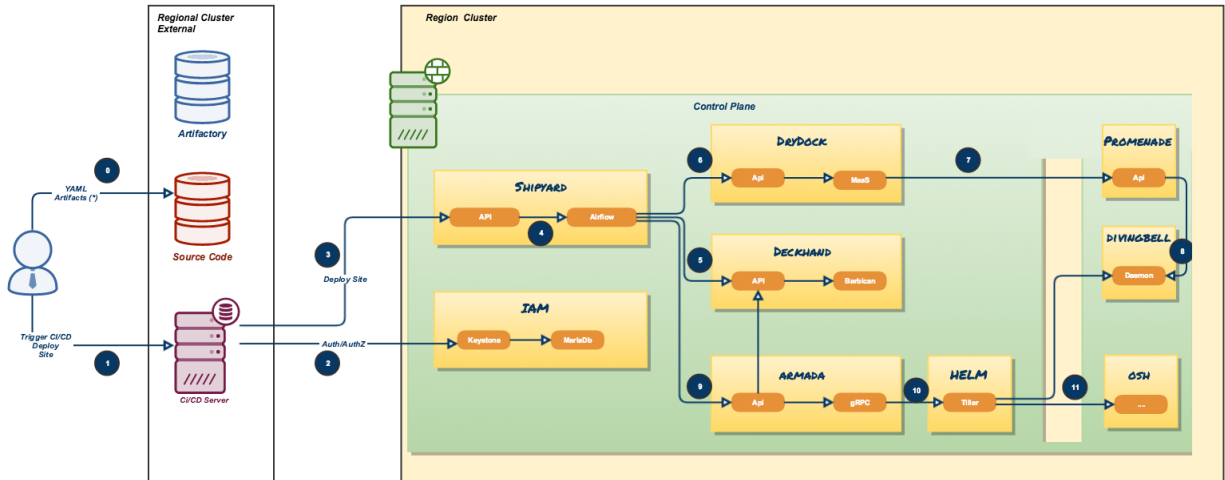
Contents

1	Component Projects	3
1.1	Pegleg	3
1.2	Shipyard	3
1.3	Drydock	4
1.4	Deckhand	4
1.5	Armada	4
1.6	Kubernetes	5
1.7	Promenade	5
1.8	Helm	5
1.9	OpenStack-Helm	5
1.10	Divingbell	5
1.11	Berth	6
2	Process Flows	7
2.1	Site Authoring and Deployment Guide	8
2.2	Configuration Update Guide	19
2.3	Troubleshooting Guide	22
2.4	Airship Seaworthy	24
2.5	Airskiff	25
2.6	Airsloop	32

Airship is a collection of components that coordinate to form means of configuring and deploying and maintaining a [Kubernetes](#) environment using a declarative set of [yaml](#) documents.

More specifically, the current focus of this project is the implementation of OpenStack on Kubernetes (OOK).

ARCHITECTURE



1.1 Pegleg

[Pegleg](#) is a document aggregator that provides early linting and validations via [Deckhand](#), a document management micro-service within Airship.

1.2 Shipyard

[Shipyard](#) is the directed acyclic graph controller for Kubernetes and OpenStack control plane life cycle management. Shipyard provides the entrypoint for the following aspects of the control plane:

1.2.1 Designs and Secrets

Site designs, including the configuration of bare metal host nodes, network design, operating systems, Kubernetes nodes, Armada manifests, Helm charts, and any other descriptors that define the build out of a group of servers enter the Airship via Shipyard. Secrets, such as passwords and certificates, use the same mechanism. The designs and secrets are stored in Airship's [Deckhand](#), providing for version history and secure storage among other document-based conveniences.

1.2.2 Actions

Interaction with the site's control plane is done via invocation of actions in Shipyard. Each action is backed by a workflow implemented as a directed acyclic graph (DAG) that runs using Apache Airflow. Shipyard provides a mechanism to monitor and control the execution of the workflow.

1.3 Drydock

Drydock is a provisioning orchestrator for baremetal servers that translates a YAML-based declarative site topology into a physical undercloud that can be used for building out an enterprise Kubernetes cluster. It uses plugins to leverage existing provisioning systems to build the servers allowing integration with the provisioning system that best fits the goals and environment of a site.

1.3.1 Capabilities

- Initial IPMI configuration for PXE booting new servers.
- Support for Canonical MAAS provisioning.
- Configuration of complex network topologies including bonding, tagged VLANs and static routes
- Support for running behind a corporate proxy
- Extensible boot action system for placing files and SystemD units on nodes for post-deployment execution
- Supports Keystone-based authentication and authorization

1.4 Deckhand

Deckhand is a document-based configuration storage service built with auditability and validation in mind.

1.4.1 Core Responsibilities

- layering - helps reduce duplication in configuration by applying the notion of inheritance to documents
- substitution - provides separation between secret data and other configuration data for security purposes and reduces data duplication by allowing common data to be defined once and substituted elsewhere dynamically
- revision history - maintains well-defined collections of documents within immutable revisions that are meant to operate together, while providing the ability to rollback to previous revisions
- validation - allows services to implement and register different kinds of validations and report errors
- secret management - leverages existing OpenStack APIs – namely **Barbican** – to reliably and securely store sensitive data

1.5 Armada

Armada is a tool for managing multiple Helm charts with dependencies by centralizing all configurations in a single Armada YAML and providing life-cycle hooks for all Helm releases.

1.5.1 Core Responsibilities

- Multiple Chart Deployments and Upgrades driven by Armada Manifests
- Manage multiple chart dependencies using Chart Groups
- Enhancing base Helm functionality

- Supports Keystone-based authentication and authorization

1.6 Kubernetes

[Kubernetes](#) is an open source system for managing containerized applications across multiple hosts, providing basic mechanisms for deployment, maintenance, and scaling of applications.

1.7 Promenade

[Promenade](#) is a tool for bootstrapping a resilient, self-hosted Kubernetes cluster and managing its life-cycle.

Bootstrapping begins by provisioning a single-node cluster with a complete, configurable Airship infrastructure. After hosts are added to the cluster, the original bootstrapping node can be re-provisioned to avoid subtle differences that could result in future issues.

Promenade provides cluster resiliency against both node failures and full cluster restarts. It does so by leveraging [Helm](#) charts to manage core Kubernetes assets directly on each host, to ensure their availability.

1.8 Helm

[Helm](#) is a package manager for Kubernetes. It helps you define, install, and upgrade even the most complex Kubernetes applications using Helm charts.

A chart is a collection of files that describe a related set of Kubernetes resources. Helm wraps up each chart's deployment into a concrete release, a tidy little box that is a collection of all the Kubernetes resources that compose that service, and so you can interact with a collection of Kubernetes resources that compose a release as a single unit, either to install, upgrade, or remove.

At its core, the value that Helm brings to the table – at least for us – is allowing us to templatize our experience with Kubernetes resources, providing a standard interface for operators or high-level software orchestrators to control the installation and life cycle of Kubernetes applications.

1.9 OpenStack-Helm

The [OpenStack-Helm](#) project provides a framework to enable the deployment, maintenance, and upgrading of loosely coupled OpenStack services and their dependencies individually or as part of complex environments.

OpenStack-Helm is essentially a marriage of Kubernetes, Helm, and OpenStack, and seeks to create Helm charts for each OpenStack service. These Helm charts provide complete life cycle management for these OpenStack services.

Users of OpenStack-Helm either deploy all or individual OpenStack components along with their required dependencies. It heavily borrows concepts from Stackanetes and complex Helm application deployments. Ideally, at the end of the day, this project is meant to be a collaborative project that brings OpenStack applications into a cloud-native model.

1.10 Divingbell

[Divingbell](#) is a lightweight solution for:

1. Bare metal configuration management for a few very targeted use cases
2. Bare metal package manager orchestration

1.10.1 What problems does it solve?

The needs identified for Divingbell were:

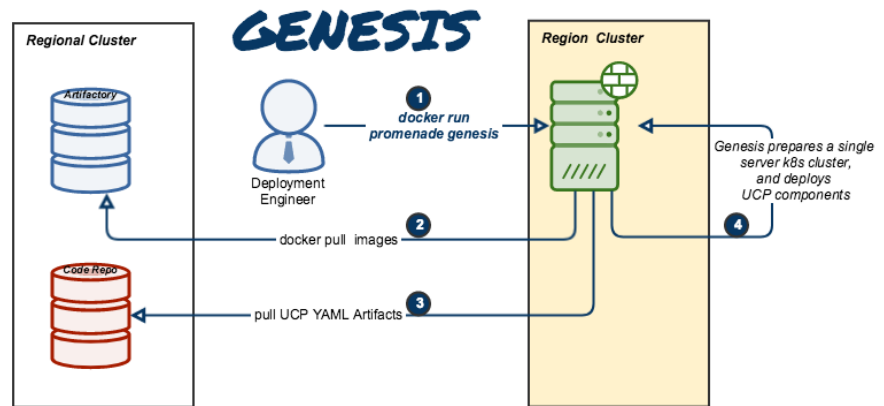
1. To plug gaps in day 1 tools (e.g., Drydock) for node configuration
2. To provide a day 2 solution for managing these configurations going forward
3. [Future] To provide a day 2 solution for system level host patching

1.11 Berth

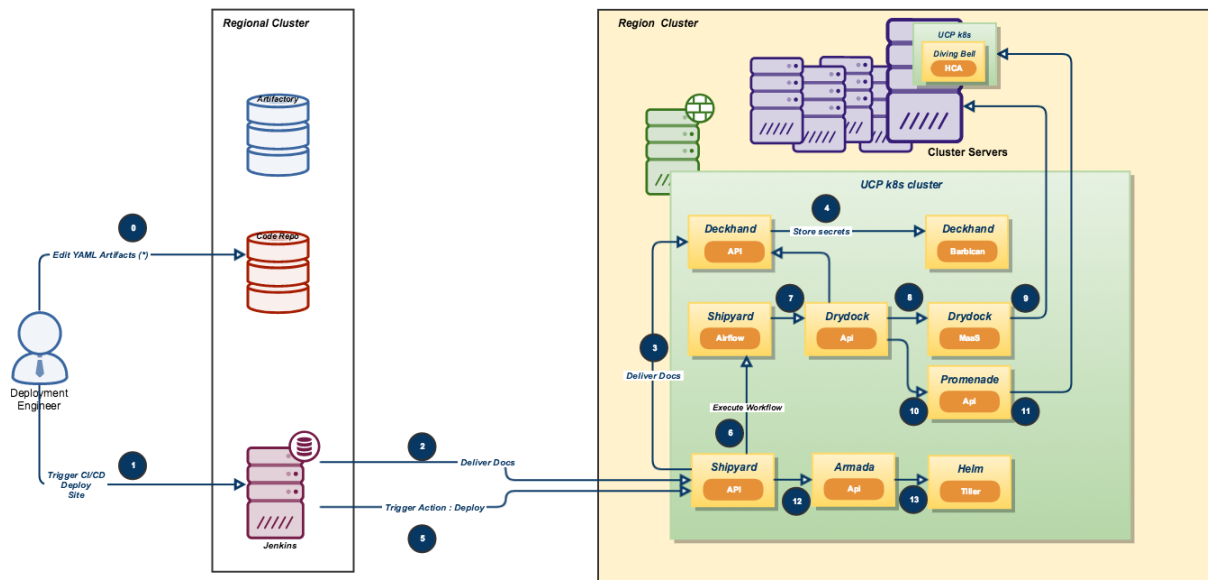
Berth is a deliberately minimalist VM runner for Kubernetes.

CHAPTER 2

Process Flows



ACTION : DEPLOY OR UPDATE SITE



2.1 Site Authoring and Deployment Guide

The document contains the instructions for standing up a greenfield Airship site. This can be broken down into two high-level pieces:

1. **Site authoring guide(s):** Describes how to craft site manifests and configs required to perform a deployment. The primary site authoring guide is for deploying Airship sites, where OpenStack is the target platform deployed on top of Airship.
2. **Deployment guide(s):** Describes how to apply site manifests for a given site.

This document is an “all in one” site authoring guide + deployment guide for a standard Airship deployment. For the most part, the site authoring guidance lives within `airship-seaworthy` reference site in the form of YAML comments.

2.1.1 Terminology

Cloud: A platform that provides a standard set of interfaces for IaaS consumers.

OSH: (OpenStack Helm) is a collection of Helm charts used to deploy OpenStack on kubernetes.

Undercloud/Overcloud: Terms used to distinguish which cloud is deployed on top of the other. In Airship sites, OpenStack (overcloud) is deployed on top of Kubernetes (undercloud).

Airship: A specific implementation of OpenStack Helm charts onto kubernetes, the deployment of which is the primary focus of this document.

Control Plane: From the point of view of the cloud service provider, the control plane refers to the set of resources (hardware, network, storage, etc) sourced to run cloud services.

Data Plane: From the point of view of the cloud service provider, the data plane is the set of resources (hardware, network, storage, etc.) sourced to run consumer workloads. When used in this document, “data plane” refers to the data plane of the overcloud (OSH).

Host Profile: A host profile is a standard way of configuring a bare metal host. Encompasses items such as the number of bonds, bond slaves, physical storage mapping and partitioning, and kernel parameters.

Component Overview



Node Overview

This document refers to several types of nodes, which vary in their purpose, and to some degree in their orchestration / setup:

- **Build node:** This refers to the environment where configuration documents are built for your environment (e.g., your laptop)
- **Genesis node:** The “genesis” or “seed node” refers to a node used to get a new deployment off the ground, and is the first node built in a new deployment environment.
- **Control / Controller nodes:** The nodes that make up the control plane. (Note that the Genesis node will be one of the controller nodes.)
- **Compute nodes / Worker Nodes:** The nodes that make up the data plane

2.1.2 Support

Bugs may be viewed and reported at the following locations, depending on the component:

- OpenStack Helm: [OpenStack Storyboard group](#)
- Airship: Bugs may be filed using OpenStack Storyboard for specific projects in [Airship group](#):
 - [Airship Armada](#)
 - [Airship Berth](#)
 - [Airship Deckhand](#)
 - [Airship Divingbell](#)
 - [Airship Drydock](#)
 - [Airship MaaS](#)
 - [Airship Pegleg](#)
 - [Airship Promenade](#)
 - [Airship Shipyard](#)
 - [Airship in a Bottle](#)
 - [Airship Treasuremap](#)

2.1.3 Hardware Prep

Disk

1. For servers that are in the control plane (including Genesis):
 - Two-disk RAID-1: Operating System
 - Two disks JBOD: Ceph Journal/Meta for control plane
 - Remaining disks JBOD: Ceph OSD for control plane
2. For servers that are in the tenant data plane (compute nodes):
 - Two-disk RAID-1: Operating System
 - Two disks JBOD: Ceph Journal/Meta for tenant-ceph
 - Two disks JBOD: Ceph OSD for tenant-ceph
 - Remaining disks need to be configured according to the host profile target for each given server (e.g. RAID-10 for OpenStack Ephemeral).

BIOS and IPMI

1. Virtualization enabled in BIOS
2. IPMI enabled in server BIOS (e.g., IPMI over LAN option enabled)
3. IPMI IPs assigned, and routed to the environment you will deploy into Note: Firmware bugs related to IPMI are common. Ensure you are running the latest firmware version for your hardware. Otherwise, it is recommended to perform an iLo/iDrac reset, as IPMI bugs with long-running firmware are not uncommon.
4. Set PXE as first boot device and ensure the correct NIC is selected for PXE

Network

1. You have a network you can successfully PXE boot with your network topology and bonding settings (dedicated PXE interface on untagged/native VLAN in this example)
2. You have (VLAN) segmented, routed networks accessible by all nodes for:
 1. Management network(s) (k8s control channel)
 2. Calico network(s)
 3. Storage network(s)
 4. Overlay network(s)
 5. Public network(s)

2.1.4 HW Sizing and minimum requirements

Node	disk	memory	cpu
Build	10 GB	4 GB	1
Genesis	100 GB	16 GB	8
Control	10 TB	128 GB	24
Compute	N/A*	N/A*	N/A*

- Workload driven (determined by host profile)

2.1.5 Establishing build node environment

1. On the machine you wish to use to generate deployment files, install required tooling:

```
sudo apt -y install docker.io git
```

2. Clone and link the required git repos as follows:

```
git clone https://git.openstack.org/openstack/airship-pegleg
git clone https://git.openstack.org/openstack/airship-treasuremap
```

2.1.6 Building Site documents

This section goes over how to put together site documents according to your specific environment, and generate the initial Promenade bundle needed to start the site deployment.

Preparing deployment documents

In its current form, pegleg provides an organized structure for YAML elements, in order to separate common site elements (i.e., `global` folder) from unique site elements (i.e., `site` folder).

To gain a full understanding of the pegleg structure, it is highly recommended to read pegleg documentation on this [here](#).

The `airship-seaworthy` site may be used as reference site. It is the principal pipeline for integration and continuous deployment testing of Airship.

Change directory to the `airship-treasuremap/site` folder and copy the `airship-seaworthy` site as follows:

```
NEW_SITE=mySite # replace with the name of your site
cd airship-treasuremap/site
cp -r airship-seaworthy $NEW_SITE
```

Remove `airship-seaworthy` specific certificates.

```
rm -f airship-treasuremap/site/${NEW_SITE}/secrets/certificates/certificates.yaml
```

You will then need to manually make changes to these files. These site manifests are heavily commented to explain parameters, and importantly identify all of the parameters that need to change when authoring a new site.

These areas which must be updated for a new site are flagged with the label `NEWSITE-CHANGEME` in YAML commentary. Search for all instances of `NEWSITE-CHANGEME` in your new site definition, and follow the instructions that accompany the tag in order to make all needed changes to author your new Airship site.

Because some files depend on (or will repeat) information from others, the order in which you should build your site files is as follows:

1. `site/$NEW_SITE/networks/physical/networks.yaml`
2. `site/$NEW_SITE/baremetal/nodes.yaml`
3. `site/$NEW_SITE/networks/common-addresses.yaml`
4. `site/$NEW_SITE/pki/pki-catalog.yaml`
5. All other site files

Register DNS names

Register the following list of DNS names:

```
cloudformation.DOMAIN
compute.DOMAIN
dashboard.DOMAIN
grafana.DOMAIN
iam.DOMAIN
identity.DOMAIN
image.DOMAIN
kibana.DOMAIN
nagios.DOMAIN
network.DOMAIN
nova-novncproxy.DOMAIN
object-store.DOMAIN
orchestration.DOMAIN
placement.DOMAIN
shipyard.DOMAIN
volume.DOMAIN
```

Here `DOMAIN` is a name of ingress domain, you can find it in the `data.dns.ingress_domain` section of `site/${NEW_SITE}/secrets/certificates/ingress.yaml` configuration file.

Run the following command to get up to date list of required DNS names:

```
grep -E 'host: .+DOMAIN' site/${NEW_SITE}/software/config/endpoints.yaml | \
  sort -u | awk '{print $2}'
```


Update Secrets

Replace passphrases under `site/${NEW_SITE}/secrets/passphrases/` with random generated ones:

- Passphrases generation `openssl rand -hex 10`
- UUID generation `uuidgen` (e.g. for Ceph filesystem ID)
- Update `secrets/passphrases/ipmi_admin_password.yaml` with IPMI password
- Update `secrets/passphrases/ubuntu_crypt_password.yaml` with password hash:

```
python3 -c "from crypt import *; print(crypt('<YOUR_PASSWORD>', METHOD_SHA512))"
```

Configure certificates in `site/${NEW_SITE}/secrets/certificates/ingress.yaml`, they need to be issued for the domains configured in Register DNS names section.

Caution: It is required to configure valid certificates, self-signed certificates are not supported.

Control Plane & Tenant Ceph Cluster Notes

Configuration variables for ceph control plane are located in:

- `site/${NEW_SITE}/software/charts/ucp/ceph/ceph-osd.yaml`
- `site/${NEW_SITE}/software/charts/ucp/ceph/ceph-client.yaml`

Configuration variables for tenant ceph are located in:

- `site/${NEW_SITE}/software/charts/osh/openstack-tenant-ceph/ceph-osd.yaml`
- `site/${NEW_SITE}/software/charts/osh/openstack-tenant-ceph/ceph-client.yaml`

Setting highlights:

- `data/values/conf/storage/osd[*]/data/location`: The block device that will be formatted by the Ceph chart and used as a Ceph OSD disk
- `data/values/conf/storage/osd[*]/journal/location`: The block device backing the ceph journal used by this OSD. Refer to the journal paradigm below.
- `data/values/conf/pool/target/osd`: Number of OSD disks on each node

Assumptions:

1. Ceph OSD disks are not configured for any type of RAID, they are configured as JBOD when connected through a RAID controller. If RAID controller does not support JBOD, put each disk in its own RAID-0 and enable RAID cache and write-back cache if the RAID controller supports it.
2. Ceph disk mapping, disk layout, journal and OSD setup is the same across Ceph nodes, with only their role differing. Out of the 4 control plane nodes, we expect to have 3 actively participating in the Ceph quorum, and the remaining 1 node designated as a standby Ceph node which uses a different control plane profile (`cp_*-secondary`) than the other three (`cp_*-primary`).
3. If doing a fresh install, disk are unlabeled or not labeled from a previous Ceph install, so that Ceph chart will not fail disk initialization.

It's highly recommended to use SSD devices for Ceph Journal partitions.

If you have an operating system available on the target hardware, you can determine HDD and SSD devices with:

```
lsblk -d -o name,rota
```

where a `rota` (rotational) value of 1 indicates a spinning HDD, and where a value of 0 indicates non-spinning disk (i.e. SSD). (Note - Some SSDs still report a value of 1, so it is best to go by your server specifications).

For OSDs, pass in the whole block device (e.g., `/dev/sdd`), and the Ceph chart will take care of disk partitioning, formatting, mounting, etc.

For Ceph Journals, you can pass in a specific partition (e.g., `/dev/sdb1`), note that it's not required to pre-create these partitions, Ceph chart will create journal partitions automatically if they don't exist. By default the size of every journal partition is 10G, make sure there is enough space available to allocate all journal partitions.

Consider the following example where:

- `/dev/sda` is an operating system RAID-1 device (SSDs for OS root)
- `/dev/sd[bc]` are SSDs for ceph journals
- `/dev/sd[efgh]` are HDDs for OSDs

The data section of this file would look like:

```
data:
  values:
    conf:
      storage:
        osd:
          - data:
              type: block-logical
              location: /dev/sde
            journal:
              type: block-logical
              location: /dev/sdb1
          - data:
              type: block-logical
              location: /dev/sdf
            journal:
              type: block-logical
              location: /dev/sdb2
          - data:
              type: block-logical
              location: /dev/sdg
            journal:
              type: block-logical
              location: /dev/sdc1
          - data:
              type: block-logical
              location: /dev/sdh
            journal:
              type: block-logical
              location: /dev/sdc2
```

Manifest linting and combining layers

After constituent YAML configurations are finalized, use `Pegleg` to lint your manifests, and resolve any issues that result from linting before proceeding:

```
sudo airship-pegleg/tools/pegleg.sh repo \
-r airship-treasuremap lint
```

Note: P001 and P003 linting errors are expected for missing certificates, as they are not generated until the next section. You may suppress these warnings by appending `-x P001 -x P003` to the lint command.

Next, use `pegleg` to perform the merge that will yield the combined global + site type + site YAML:

```
sudo sh airship-pegleg/tools/pegleg.sh site \
-r airship-treasuremap \
collect $NEW_SITE
```

Perform a visual inspection of the output. If any errors are discovered, you may fix your manifests and re-run the `lint` and `collect` commands.

After you have an error-free output, save the resulting YAML as follows:

```
sudo airship-pegleg/tools/pegleg.sh site \
-r airship-treasuremap \
collect $NEW_SITE -s ${NEW_SITE}_collected
```

It is this output which will be used in subsequent steps.

Lastly, you should also perform a `render` on the documents. The resulting render from `Pegleg` will not be used as input in subsequent steps, but is useful for understanding what the document will look like once `Deckhand` has performed all substitutions, replacements, etc. This is also useful for troubleshooting, and addressing any `Deckhand` errors prior to submitting via `Shipyard`:

```
sudo airship-pegleg/tools/pegleg.sh site \
-r airship-treasuremap \
render $NEW_SITE
```

Inspect the rendered document for any errors. If there are errors, address them in your manifests and re-run this section of the document.

Building the Promenade bundle

Clone the Promenade repo, if not already cloned:

```
git clone https://opendev.org/airship/promenade
```

Refer to the `data/charts/ucp/promenade/reference` field in `airship-treasuremap/global/software/config/versions.yaml`. If this is a pinned reference (i.e., any reference that's not `master`), then you should checkout the same version of the Promenade repository. For example, if the Promenade reference was `86c3c11...` in the versions file, checkout the same version of the Promenade repo which was cloned previously:

```
(cd promenade && git checkout 86c3c11)
```

Likewise, before running the `simple-deployment.sh` script, you should refer to the `data/images/ucp/promenade/promenade` field in `~/airship-treasuremap/global/software/config/versions.yaml`. If there is a pinned reference (i.e., any image reference that's not `latest`), then this reference should be used to set the `IMAGE_PROMENADE` environment variable. For example, if the Promenade image was pinned to `quay.io/airshipit/promenade:d30397f...` in the versions file, then export the previously mentioned environment variable like so:

```
export IMAGE_PROMENADE=quay.io/airshipit/promenade:d30397f...
```

Now, create an output directory for Promenade bundles and run the `simple-deployment.sh` script:

```
mkdir ${NEW_SITE}_bundle
sudo -E promenade/tools/simple-deployment.sh ${NEW_SITE}_collected ${NEW_SITE}_bundle
```

Estimated runtime: About **1 minute**

After the bundle has been successfully created, copy the generated certificates into the security folder. Ex:

```
mkdir -p airship-treasuremap/site/${NEW_SITE}/secrets/certificates
sudo cp ${NEW_SITE}_bundle/certificates.yaml \
  airship-treasuremap/site/${NEW_SITE}/secrets/certificates/certificates.yaml
```

Regenerate collected YAML files to include copied certificates:

```
sudo airship-pegleg/tools/pegleg.sh site \
  -r airship-treasuremap \
  collect $NEW_SITE -s ${NEW_SITE}_collected
```

2.1.7 Genesis node

Initial setup

Before starting, ensure that the BIOS and IPMI settings match those stated previously in this document. Also ensure that the hardware RAID is setup for this node per the control plane disk configuration stated previously in this document.

Then, start with a manual install of Ubuntu 16.04 on the node you wish to use to seed the rest of your environment standard [Ubuntu ISO](#). Ensure to select the following:

- UTC timezone
- Hostname that matches the Genesis hostname given in `/data/genesis/hostname` in `airship-treasuremap/site/${NEW_SITE}/networks/common-addresses.yaml`.
- At the Partition Disks screen, select Manual so that you can setup the same disk partitioning scheme used on the other control plane nodes that will be deployed by MaaS. Select the first logical device that corresponds to one of the RAID-1 arrays already setup in the hardware controller. On this device, setup partitions matching those defined for the `bootdisk` in your control plane host profile found in `airship-treasuremap/site/${NEW_SITE}/profiles/host`. (e.g., 30G for `/`, 1G for `/boot`, 100G for `/var/log`, and all remaining storage for `/var`). Note that the volume size syntax looking like `>300g` in Drydock means that all remaining disk space is allocated to this volume, and that volume needs to be at least 300G in size.
- Ensure that OpenSSH and Docker (Docker is needed because of `miniMirror`) are included as installed packages
- When you get to the prompt, “How do you want to manage upgrades on this system?”, choose “No automatic updates” so that packages are only updated at the time of our choosing (e.g. maintenance windows).
- Ensure the grub bootloader is also installed to the same logical device as in the previous step (this should be default behavior).

After installation, ensure the host has outbound internet access and can resolve public DNS entries (e.g., `nslookup google.com`, `curl https://www.google.com`).

Ensure that the deployed Genesis hostname matches the hostname in `data/genesis/hostname` in `airship-treasuremap/site/${NEW_SITE}/networks/common-addresses.yaml`. If it does not match, then either change the hostname of the node to match the configuration documents, or re-generate the configuration with the correct hostname. In order to change the hostname of the deployed node, you may run the following:

```
sudo hostname $NEW_HOSTNAME
sudo sh -c "echo $NEW_HOSTNAME > /etc/hostname"
sudo vi /etc/hosts # Anywhere the old hostname appears in the file, replace
                  # with the new hostname
```

Or to regenerate manifests, re-run the previous two sections with the after updating the genesis hostname in the site definition.

Installing matching kernel version

Install the same kernel version on the Genesis host that MaaS will use to deploy new baremetal nodes.

In order to do this, first you must determine the kernel version that will be deployed to those nodes. Start by looking at the host profile definition used to deploy other control plane nodes by searching for `control-plane: enabled`. Most likely this will be a file under `global/profiles/host`. In this file, find the kernel info - e.g.:

```
platform:
  image: 'xenial'
  kernel: 'hwe-16.04'
  kernel_params:
    kernel_package: 'linux-image-4.15.0-34-generic'
```

In this example, the kernel version is `4.15.0-34-generic`. Define any proxy environment variables needed for your environment to reach public ubuntu package repos, and install the matching kernel on the Genesis host (make sure to run on Genesis host, not on the build host):

```
sudo apt -y install linux-modules-4.15.0-34-generic \
  linux-modules-extra-4.15.0-34-generic \
  linux-image-4.15.0-34-generic linux-headers-4.15.0-34-generic \
  linux-headers-4.15.0-34
```

Check the installed packages on the genesis host with `dpkg --get-selections`. If there are any later kernel versions installed, remove them with `sudo apt remove`, so that the newly install kernel is the latest available.

Install ntpdate/ntp

Install and run `ntpdate`, to ensure a reasonably sane time on genesis host before proceeding:

```
sudo apt -y install ntpdate
sudo ntpdate ntp.ubuntu.com
```

If your network policy does not allow time sync with external time sources, specify a local NTP server instead of using `ntp.ubuntu.com`.

Then, install the NTP client:

```
sudo apt -y install ntp
```

Add the list of NTP servers specified in `data/ntp/servers_joined` in file `airship-treasuremap/site/${NEW_SITE}/networks/common-address.yaml` to `/etc/ntp.conf` as follows:

```
pool NTP_SERVER1 iburst
pool NTP_SERVER2 iburst
(repeat for each NTP server with correct NTP IP or FQDN)
```

Then, restart the NTP service:

```
sudo service ntp restart
```

If you cannot get good time to your selected time servers, consider using alternate time sources for your deployment.

Disable the apparmor profile for ntpd:

```
sudo ln -s /etc/apparmor.d/usr.sbin.ntpd /etc/apparmor.d/disable/
sudo apparmor_parser -R /etc/apparmor.d/usr.sbin.ntpd
```

This prevents an issue with the MaaS containers, which otherwise get permission denied errors from apparmor when the MaaS container tries to leverage libc6 for /bin/sh when MaaS container ntpd is forcefully disabled.

Promenade bootstrap

Copy the `${NEW_SITE}_bundle` and `${NEW_SITE}_collected` directories from the build node to the genesis node, into the home directory of the user there (e.g., `/home/ubuntu`). Then, run the following script as sudo on the genesis node:

```
cd ${NEW_SITE}_bundle
sudo ./genesis.sh
```

Estimated runtime: **40m**

Following completion, run the `validate-genesis.sh` script to ensure correct provisioning of the genesis node:

```
cd ${NEW_SITE}_bundle
sudo ./validate-genesis.sh
```

Estimated runtime: **2m**

2.1.8 Deploy Site with Shipyard

Start by cloning the shipyard repository to the Genesis node:

```
git clone https://opendev.org/airship/shipyard
```

Refer to the `data/charts/ucp/shipyard/reference` field in `airship-treasuremap/global/software/config/versions.yaml`. If this is a pinned reference (i.e., any reference that's not `master`), then you should checkout the same version of the Shipyard repository. For example, if the Shipyard reference was `7046ad3...` in the versions file, checkout the same version of the Shipyard repo which was cloned previously:

```
(cd shipyard && git checkout 7046ad3)
```

Likewise, before running the `deckhand_load_yaml.sh` script, you should refer to the `data/images/ucp/shipyard/shipyard` field in `airship-treasuremap/global/software/config/versions.yaml`. If there is a pinned reference (i.e., any image reference that's not `latest`), then this reference should be used to set the `SHIPYARD_IMAGE` environment variable. For example, if the Shipyard image was pinned to `quay.io/airshipit/shipyard@sha256:dfc25e1...` in the versions file, then export the previously mentioned environment variable:

```
export SHIPYARD_IMAGE=quay.io/airshipit/shipyard@sha256:dfc25e1...
```

Export valid login credentials for one of the Airship Keystone users defined for the site. Currently there is no authorization checks in place, so the credentials for any of the site-defined users will work. For example, we can use the `shipyard` user, with the password that was defined in `airship-treasuremap/site/${NEW_SITE}/secrets/passphrases/ucp_shipyard_keystone_password.yaml`. Ex:

```
export OS_USERNAME=shipyard
export OS_PASSWORD=46a75e4...
```

(Note: Default auth variables are defined [here](#), and should otherwise be correct, barring any customizations of these site parameters).

Next, run the `deckhand_load_yaml.sh` script providing an absolute path to a directory that contains collected manifests:

```
sudo -E shipyard/tools/deckhand_load_yaml.sh ${NEW_SITE} $(pwd)/${NEW_SITE}_collected
```

Estimated runtime: **3m**

Now deploy the site with `shipyard`:

```
cd shipyard/tools/
sudo -E ./deploy_site.sh
```

Estimated runtime: **1h30m**

The message `Site Successfully Deployed` is the expected output at the end of a successful deployment. In this example, this means that Airship and OSH should be fully deployed.

2.1.9 Disable password-based login on Genesis

Before proceeding, verify that your SSH access to the Genesis node is working with your SSH key (i.e., not using password-based authentication).

Then, disable password-based SSH authentication on Genesis in `/etc/ssh/sshd_config` by uncommenting the `PasswordAuthentication` and setting its value to `no`. Ex:

```
PasswordAuthentication no
```

Then, restart the `ssh` service:

```
sudo systemctl restart ssh
```

2.2 Configuration Update Guide

The guide contains the instructions for updating the configuration of a deployed Airship environment. Please refer to [Site Authoring and Deployment Guide](#) if you do not have an Airship environment already deployed.

Update of an Airship environment consists of the following stages:

1. **Prepare the configuration:** before deploying any changes, a user should prepare and validate the manifests on a build node using [Airship Pegleg](#).
2. **Deploy the changes:** during this stage, a user uploads the configuration to the Airship environment and starts the deployment using [Airship Shipyard](#).

Note: This guide assumes you have [Airship Pegleg](#) and [Airship Shipyard](#) tools installed and configured; please refer to [Site Authoring and Deployment Guide](#) for the details.

2.2.1 Configuring Airship CLI

Clone the Airship Treasuremap repository and switch to correct version.

```
git clone https://opendev.org/airship/treasuremap
cd treasuremap/
# List available tags.
git tag --list
# Switch to the version your site is using.
git checkout {your-tag}
# Go back to a previous directory.
cd ..
```

Configure environment variables with the name of your site, and specify a path to the directory where site configuration is stored; for this example, we use [Airship Seaworthy](#) site:

```
export SITE=airship-seaworthy
export SITE_PATH=treasuremap/site/airship-seaworthy
```

2.2.2 Updating the manifests

Changing the configuration consists of the following steps:

1. Change site manifests.
2. Lint the manifests.
3. Collect the manifests.
4. Copy the manifests to the Airship environment.

Linting and collecting the manifests is done using [Airship Pegleg](#).

For this example, we are going to update a debug level for keystone logs in a site layer.

Note: It is also possible to update the configuration in a global layer; for more details on Airship layering mechanism see [Pegleg Definition Artifact Layout](#) documentation.

Create an override file `${SITE_PATH}/software/charts/osh/openstack-keystone/keystone.yaml` with the following content:

```
---
schema: armada/Chart/v1
metadata:
  schema: metadata/Document/v1
  name: keystone
  replacement: true
  layeringDefinition:
    abstract: false
    layer: site
```

(continues on next page)

(continued from previous page)

```

parentSelector:
  name: keystone-global
actions:
  - method: merge
    path: .
storagePolicy: cleartext
data:
  values:
    conf:
      logging:
        logger_keystone:
          level: DEBUG
...

```

Check that the configuration is valid:

```

sudo ./treasuremap/tools/airship pegleg site -r treasuremap/ \
  lint ${SITE}

```

Collect the configuration:

```

sudo ./treasuremap/tools/airship pegleg site \
  -r treasuremap/ collect ${SITE} -s ${SITE}_collected

```

Copy the configuration to a node that has the access to the site's Shipyard API, if the current node does not; this node can be one of your controllers:

```

scp -r ${SITE}_collected {genesis-ip}:/home/{user-name}/${SITE}_collected

```

2.2.3 Deploying the changes

After you copied the manifests, there are just a few steps needed to start the deployment:

1. Upload the changes to [Airship Deckhand](#).
2. Start the deployment using [Airship Shipyard](#).

Install Airship CLI as described in [Configuring Airship CLI](#) section.

Set the name of your site:

```

:: export SITE=airship-seaworthy

```

Configure credentials for accessing Shipyard; the password is stored in `ucp_shipyard_keystone_password` secret, you can find it in `site/airship-seaworthy/secrets/passphrases/ucp_shipyard_keystone_password.yaml` configuration file of your site.

```

export OS_USERNAME=shipyard
export OS_PASSWORD={shipyard_password}

```

Upload the changes to [Airship Deckhand](#):

```

# Upload the configuration.
sudo -E ./treasuremap/tools/airship shipyard \
  create configdocs ${SITE} --replace --directory=${SITE}_collected

```

(continues on next page)

(continued from previous page)

```
# Commit the configuration.
sudo -E ./treasuremap/tools/airship shipyard commit configdocs
```

Run the deployment:

```
sudo -E ./treasuremap/tools/airship shipyard create action update_site
```

You can also run `update_software` instead of `update_site` which skips hardware configuration and only applies the changes to services that are running on top of Kubernetes.

Now you can track the deployment progress using the following commands:

```
# Get all actions that were executed on your environment.
sudo -E ./treasuremap/tools/airship shipyard get actions

# Show all the steps within the action.
sudo -E ./treasuremap/tools/airship shipyard describe action/{action_id}
```

All steps will have status `success` when the update finishes.

2.3 Troubleshooting Guide

This guide provides information on troubleshooting of an Airship environment. Debugging of any software component starts with gathering more information about the failure, so the intention of the document is not to describe specific issues that one can encounter, but to provide a generic set of instructions that a user can follow to find the root cause of the problem.

For additional support you can contact the Airship team via [IRC or mailing list](#), use [Airship bug tracker](#) to search and create issues.

2.3.1 Configuring Airship CLI

Many commands from this guide use Airship CLI, this section describes how to get it configured on your environment.

```
git clone https://opendev.org/airship/treasuremap
cd treasuremap/
# List available tags.
git tag --list
# Switch to the version your site is using.
git checkout {your-tag}
# Go back to a previous directory.
cd ..
# Run it without arguments to get a help message.
sudo ./treasuremap/tools/airship
```

2.3.2 Manifests Preparation

When you do any configuration changes to the manifests, there are a few commands that you can use to validate the changes without uploading them to the Airship environment.

Run `lint` command for your site; it helps to catch the errors related to documents duplication, broken references, etc.

Example:

```
sudo ./treasuremap/tools/airship pegleg site -r airship-treasuremap/ \
  lint {site-name}
```

If you create configuration overrides or do changes to substitutions, it is recommended to run `render` command this command merges the layers and renders all substitutions. This allows finding what parameters are passed to Helm as overrides for Charts' defaults.

Example:

```
# Saves the result into rendered.txt file.
sudo ./treasuremap/tools/airship pegleg site -r treasuremap/ \
  render -o rendered.txt ${SITE}
```

2.3.3 Deployment Failure

During the deployment, it is important to identify a specific step where it fails, there are two major deployment steps:

1. **Drydock build:** deploys Operating System.
2. **Armada build:** deploys Helm Charts.

After *Configuring Airship CLI*, setup credentials for accessing Shipyard; the password is stored in `ucp_shipyard_keystone_password` secret, you can find it in `site/airship-seaworthy/secrets/passphrases/ucp_shipyard_keystone_password.yaml` configuration file of your site.

```
export OS_USERNAME=shipyard
export OS_PASSWORD={shipyard_password}
```

Now you can use the following commands to access Shipyard:

```
# Get all actions that were executed on you environment.
sudo ./treasuremap/tools/airship shipyard get actions
# Show all the steps within the action.
sudo ./treasuremap/tools/airship shipyard describe action/{action_id}
# Get a bit more details on the step.
sudo ./treasuremap/tools/airship shipyard describe step/{action_id}/armada_build
# Print the logs from the step.
sudo ./treasuremap/tools/airship shipyard logs step/{action_id}/armada_build
```

After the failed step is determined, you can access the logs of a specific service (e.g., `drydock-api/maas` or `armada-api`) to get more information on the failure, note that there may be multiple pods of a single service running, you need to check all of them to find where the most recent logs are available.

Example of accessing Armada API logs:

```
# Get all pods running on the cluster and find a name of the pod you are
# interested in.
kubectl get pods -o wide --all-namespaces

# See the logs of specific pod.
kubectl logs -n ucp -f --tail 200 armada-api-d5f757d5-6z6nv
```

In some cases you want to restart your pod, there is no dedicated command for that in Kubernetes. However, you can delete the pod, it will be restarted by Kubernetes to satisfy replication factor.

```
# Restart Armada API service.
kubectl delete pod -n ucp armada-api-d5f757d5-6z6nv
```

2.3.4 Ceph

Many stateful services in Airship rely on Ceph to function correctly. For more information on Ceph debugging follow an official [Ceph debugging guide](#).

Although Ceph tolerates failures of multiple OSDs, it is important to make sure that your Ceph cluster is healthy.

Example:

```
# Get a name of Ceph Monitor pod.
CEPH_MON=$(sudo kubectl get pods --all-namespaces -o=name | \
  grep ceph-mon | sed -n 1p | sed 's|pod|'|)
# Get the status of the Ceph cluster.
sudo kubectl exec -it -n ceph ${CEPH_MON} -- ceph -s
```

Cluster is in a healthy state when health parameter is set to HEALTH_OK.

When the cluster is unhealthy, and some Placement Groups are reported to be in degraded or down states, determine the problem by inspecting the logs of Ceph OSD that is down using `kubectl`.

```
# Get a name of Ceph Monitor pod.
CEPH_MON=$(sudo kubectl get pods --all-namespaces -o=name | \
  grep ceph-mon | sed -n 1p | sed 's|pod|'|)
# List a hierarchy of OSDs in the cluster to see what OSDs are down.
sudo kubectl exec -it -n ceph ${CEPH_MON} -- ceph osd tree
```

There are a few other commands that may be useful during the debugging:

```
# Get a name of Ceph Monitor pod.
CEPH_MON=$(sudo kubectl get pods --all-namespaces -o=name | \
  grep ceph-mon | sed -n 1p | sed 's|pod|'|)

# Get a detailed information on the status of every Placement Group.
sudo kubectl exec -it -n ceph ${CEPH_MON} -- ceph pg dump

# List allocated block devices.
sudo kubectl exec -it -n ceph ${CEPH_MON} -- rbd ls
# See what client uses the device.
sudo kubectl exec -it -n ceph ${CEPH_MON} -- rbd status \
  kubernetes-dynamic-pvc-e71e65a9-3b99-11e9-bf31-e65b6238af01

# List all Ceph block devices mounted on a specific host.
mount | grep rbd
```

2.4 Airship Seaworthy

Airship Seaworthy is a multi-node site deployment reference and continuous integration pipeline.

The site manifests are available at [site/airship-seaworthy](https://github.com/airshipit/airship-seaworthy).

2.4.1 Pipeline

Airship Seaworthy pipeline automates deployment flow documented in [Site Authoring and Deployment Guide](#).

The pipeline is implemented as Jenkins Pipeline (Groovy), see code for the pipeline at [Jenkinsfile](#).

2.4.2 Versions

The manifest overrides ([versions.yaml](#)) are setup to deploy OpenStack Ocata.

The versions are kept up to date via [updater.py](#), a utility that updates [versions.yaml](#) latest charts and (selected) images.

The pipeline attempts to uplift and deploy latest versions on daily bases.

2.4.3 Hardware

While HW configuration is flexible, Airship Seaworthy reference manifests reflect full HA deployment, similar to what might be expected in production.

Reducing number of control/compute nodes will require site overrides to align parts of the system such as Ceph replication, etcd, etc.

Airship Seaworthy site has 6 DELL R720xd bare-metal servers: 3 control, and 3 compute nodes. See host profiles for the servers [here](#).

Control (masters)

- cab23-r720-11
- cab23-r720-12
- cab23-r720-13

Compute (workers)

- cab23-r720-14
- cab23-r720-17
- cab23-r720-19

2.4.4 Network

Physical (underlay) networks are described in Drydock site configuration [here](#). It defines OOB (iLO/IPMI), untagged PXE, and multiple tagged general use networks.

Calico overlay for k8s POD networking uses IPIP mesh.

BGP peering is supported but not enabled in this setup, see [Calico chart](#).

2.5 Airskiff

- Skiff (n): a shallow, flat-bottomed, open boat
- Airskiff (n): a learning development, and gating environment for Airship

2.5.1 What is Airskiff

Airskiff is an easy way to get started with the software delivery components of Airship:

- Armada
- Deckhand
- Pegleg

- [Shipyards](#)

Airskiff is packaged with a set of deployment scripts modeled after the [OpenStack-Helm project](#) for seamless developer setup.

These scripts:

- Download, build, and containerize the Airship components above from source.
- Deploy a Kubernetes cluster using Minikube.
- Deploy Armada, Deckhand, and Shipyards using the latest [Armada image](#).
- Deploy OpenStack using the Airskiff site and charts from the [OpenStack-Helm project](#).

Warning: Airskiff is not safe for production use. These scripts are only intended to deploy a minimal development environment.

2.5.2 Common Deployment Requirements

This section covers actions that may be required for some deployment scenarios.

Passwordless sudo

Airskiff relies on scripts that utilize the `sudo` command. Throughout this guide the assumption is that the user is: `ubuntu`. It is advised to add the following lines to `/etc/sudoers`:

```
root    ALL=(ALL) NOPASSWD: ALL
ubuntu  ALL=(ALL) NOPASSWD: ALL
```

Proxy Configuration

Note: This section assumes you have properly defined the standard `http_proxy`, `https_proxy`, and `no_proxy` environment variables and have followed the [Docker proxy guide](#) to create a systemd drop-in unit.

In order to deploy Airskiff behind proxy servers, define the following environment variables:

```
export USE_PROXY=true
export PROXY=${http_proxy}
export no_proxy=${no_proxy},10.0.2.15,.svc.cluster.local
export NO_PROXY=${NO_PROXY},10.0.2.15,.svc.cluster.local
```

Note: The `.svc.cluster.local` address is required to allow the OpenStack client to communicate without being routed through proxy servers. The IP address `10.0.2.15` is the advertised IP address of the minikube Kubernetes cluster. Replace the addresses if your configuration does not match the one defined above.

2.5.3 Deploy Airskiff

Deploy Airskiff using the deployment scripts contained in the `tools/deployment/airskiff` directory of the `airship-treasuremap` repository.

Note: Scripts should be run from the root of `airship-treasuremap` repository.

Install required packages

```
#!/bin/bash
set -xe

# Docker CE Repository
curl -fsSL https://download.docker.com/linux/ubuntu/gpg | sudo apt-key add -
sudo apt-add-repository \
    "deb [arch=amd64] https://download.docker.com/linux/ubuntu \
    $(lsb_release -cs) \
    stable"

# Purge Docker and install Docker CE
sudo systemctl unmask docker.service
sudo apt-get remove --no-install-recommends -y docker docker-engine docker.io

# TODO(drewwalters96): Update to Docker 18.09 when supported by Minikube.
sudo apt-get update
sudo apt-get install --allow-downgrades --no-install-recommends -y \
    apparmor \
    ca-certificates \
    docker-ce=18.06.1~ce~3~0~ubuntu \
    git \
    make \
    jq \
    nmap \
    curl \
    python-pip \
    uuid-runtime \
    apt-transport-https \
    ca-certificates \
    software-properties-common

# Enable apparmor
sudo systemctl enable apparmor
sudo systemctl start apparmor

# Add $USER to docker group
# NOTE: This requires re-authentication. Restart your shell.
sudo adduser "$(whoami)" docker
```

Alternatively, this step can be performed by running the script directly:

```
./tools/deployment/airskiff/developer/000-install-packages.sh
```

Restart your shell session

At this point, restart your shell session to complete adding \$USER to the docker group.

Clone Dependencies

```
#!/bin/bash
set -xe

CURRENT_DIR="$(pwd)"
: "${INSTALL_PATH:= "../}"
: "${OSH_INFRA_COMMIT:= "6d0a9c21b0455e9d674d525c5e0cd0d5d53f0f85"}"

cd ${INSTALL_PATH}

# Clone Airship projects
git clone https://opendev.org/airship/armada.git
git clone https://opendev.org/airship/deckhand.git
git clone https://opendev.org/airship/pegleg.git
git clone https://opendev.org/airship/shipyard.git

# Clone dependencies
git clone https://opendev.org/openstack/openstack-helm-infra.git

cd openstack-helm-infra
git checkout "${OSH_INFRA_COMMIT}"

cd "${CURRENT_DIR}"
```

Alternatively, this step can be performed by running the script directly:

```
./tools/deployment/airskiff/developer/005-clone-dependencies.sh
```

Deploy Kubernetes with Minikube

```
#!/bin/bash

CURRENT_DIR="$(pwd)"
: "${OSH_INFRA_PATH:= "../openstack-helm-infra"}"

# Configure proxy settings if $PROXY is set
if [ -n "${PROXY}" ]; then
    . tools/deployment/airskiff/common/setup-proxy.sh
fi

# Deploy K8s with Minikube
cd "${OSH_INFRA_PATH}"
bash -c "./tools/deployment/common/005-deploy-k8s.sh"

kubectl label nodes --all --overwrite ucp-control-plane=enabled

cd "${CURRENT_DIR}"
```

Alternatively, this step can be performed by running the script directly:


```
./tools/deployment/airskiff/developer/010-deploy-k8s.sh
```

Setup OpenStack Client

```
#!/bin/bash
set -xe

# Install OpenStack client and create OpenStack client configuration file.
sudo -H -E pip install "cmd2<=0.8.7"
sudo -H -E pip install python-openstackclient python-heatclient

sudo -H mkdir -p /etc/openstack
sudo -H chown -R "$(id -un)": /etc/openstack
tee /etc/openstack/clouds.yaml << EOF
clouds:
  airship:
    region_name: RegionOne
    identity_api_version: 3
    auth:
      username: 'admin'
      password: 'password'
      project_name: 'admin'
      project_domain_name: 'default'
      user_domain_name: 'default'
      auth_url: 'http://keystone.ucp.svc.cluster.local/v3'
  openstack:
    region_name: RegionOne
    identity_api_version: 3
    auth:
      username: 'admin'
      password: 'password123'
      project_name: 'admin'
      project_domain_name: 'default'
      user_domain_name: 'default'
      auth_url: 'http://keystone.openstack.svc.cluster.local/v3'
EOF
```

Alternatively, this step can be performed by running the script directly:

```
./tools/deployment/airskiff/developer/020-setup-client.sh
```

Deploy Airship components using Armada

```
#!/bin/bash
set -xe

: "${INSTALL_PATH:="$(pwd)/../"}"

# Download latest Armada image and deploy Airship components
docker run --rm --net host -p 8000:8000 --name armada \
  -v ~/.kube/config:armada/.kube/config \
  -v "$(pwd)/tools/deployment/airskiff/manifests:/manifests \
  -v "${INSTALL_PATH}"/airship-components \
```

(continues on next page)

(continued from previous page)

```
quay.io/airshipit/armada:latest \  
apply /manifests/airship.yaml
```

Alternatively, this step can be performed by running the script directly:

```
./tools/deployment/airskiff/developer/030-armada-bootstrap.sh
```

Deploy OpenStack using Airship

```
#!/bin/bash  
set -xe  
  
# Deploy OpenStack using Airship  
CURRENT_DIR="$(pwd)"  
: "${PL_PATH:=../pegleg}"  
: "${SY_PATH:=../shipyard}"  
  
# NOTE: Image to use for all Pegleg operations  
: "${PL_IMAGE:=quay.io/airshipit/pegleg:latest}"  
  
# Lint deployment documents  
: "${PEGLEG:=${PL_PATH}/tools/pegleg.sh}"  
: "${PL_SITE:=airskiff}"  
  
# NOTE(drewwalters96): Disable Pegleg linting errors P001 and P009; a  
# a cleartext storage policy is acceptable for non-production use cases  
# and maintain consistency with other treasuremap sites.  
IMAGE=${PL_IMAGE} TERM_OPTS=" " ${PEGLEG} site -r . lint "${PL_SITE}" -x P001 -x P009  
  
# Collect deployment documents  
: "${PL_OUTPUT:=peggles}"  
mkdir -p ${PL_OUTPUT}  
  
IMAGE=${PL_IMAGE} TERM_OPTS="-l info" ${PEGLEG} site -r . collect ${PL_SITE} -s ${PL_  
↪OUTPUT}  
cp -rp "${CURRENT_DIR}/${PL_OUTPUT}" ${SY_PATH}  
  
# Deploy Airskiff site  
cd ${SY_PATH}  
: "${SHIPYARD:=../tools/shipyard.sh}"  
  
${SHIPYARD} create configdocs airskiff-design \  
    --replace \  
    --directory=/target/${PL_OUTPUT}  
  
${SHIPYARD} commit configdocs  
${SHIPYARD} create action update_software --allow-intermediate-commits  
cd "${CURRENT_DIR}"
```

Alternatively, this step can be performed by running the script directly:

```
./tools/deployment/airskiff/developer/100-deploy-osh.sh
```

2.5.4 Use Airskiff

The Airskiff deployment scripts install and configure the OpenStack client for usage on your host machine.

Airship Examples

To use Airship services, set the `OS_CLOUD` environment variable to `airship`.

```
export OS_CLOUD=airship
```

List the Airship service endpoints:

```
openstack endpoint list
```

Note: `${SHIPYARD}` is the path to a cloned [Shipyards](#) repository.

Run Helm tests for all deployed releases:

```
${SHIPYARD}/tools/shipyards.sh create action test_site
```

List all [Shipyards](#) actions:

```
${SHIPYARD}/tools/shipyards.sh get actions
```

For more information about Airship operations, see the [Shipyards actions](#) documentation.

OpenStack Examples

To use OpenStack services, set the `OS_CLOUD` environment variable to `openstack`:

```
export OS_CLOUD=openstack
```

List the OpenStack service endpoints:

```
openstack endpoint list
```

List Glance images:

```
openstack image list
```

Issue a new Keystone token:

```
openstack token issue
```

Note: Airskiff deploys identity, network, cloudformation, placement, compute, orchestration, and image services. You can deploy more services by adding chart groups to `site/airskiff/software/manifests/full-site.yaml`. For more information, refer to the [site authoring and deployment guide](#).

2.5.5 Develop with Airskiff

Once you have successfully deployed a running cluster, changes to Airship and OpenStack components can be deployed using [Shipyards actions](#) or the Airskiff deployment scripts.

This example demonstrates deploying [Armada](#) changes using the Airskiff deployment scripts.

Note: `${ARMADA}` is the path to your cloned Armada repository that contains the changes you wish to deploy. `${TREASUREMAP}` is the path to your cloned Treasuremap repository.

Build Armada:

```
cd ${ARMADA}
make images
```

Update Airship components:

```
cd ${TREASUREMAP}
./tools/deployment/developer/airskiff/030-armada-bootstrap.sh
```

2.5.6 Troubleshooting

This section is intended to help you through the initial troubleshooting process. If issues persist after following this guide, please join us on [IRC: #airshipit](#) (freenode)

Missing value `auth-url` required for `auth` plugin password

If this error message appears when using the OpenStack client, verify your client is configured for authentication:

```
# For Airship services
export OS_CLOUD=airship

# For OpenStack services
export OS_CLOUD=openstack
```

2.6 Airsloop

Airsloop is a two node site deployment reference and continuous integration pipeline. The goal of this site is to be used as a reference for mini Airship deployments with one controller and one or more compute nodes.

Airsloop is meant to get an operator up and running quickly with an Airship lab environment, and has some resiliency and security features tuned down. Please see the [Seaworthy](#) manifests for a production-oriented example.

The site manifests are available at [site/airsloop](#).

2.6.1 Pipeline

Airsloop pipeline automates deployment flow documented in [Site Authoring and Deployment Guide](#).

The pipeline is implemented as Jenkins Pipeline (Groovy), see code for the pipeline at [Jenkinsfile](#).

2.6.2 Versions

The manifest software version overrides (`versions.yaml`) are setup to deploy OpenStack Ocata.

The versions are kept up to date via `updater.py`, a utility that updates `versions.yaml` latest charts and (selected) images.

The pipeline attempts to uplift and deploy latest versions on daily basis.

2.6.3 Hardware

While HW configuration is flexible, Airsloop reference manifests reflect a single controller and a single compute node. The aim of this is to create a minimalistic lab/demo reference environment.

Increasing the number of compute nodes will require site overrides to align parts of the system such as Ceph osds, etcd, etc.

This site is stripped down from all the complicated hardware and configuration requirements that Airship Seaworthy site is using and that leads to simplified deployments from disk, networking and hardware perspective.

Airsloop site has 2 bare-metal servers: 1 controller, and 1 compute node. See host profiles for the servers [here](#).

Control (masters)

- `airsloop-control-1`

Compute (workers)

- `airsloop-compute-1`

For simplification of the configuration, we recommend using hostnames provided below.

For a two node deployment the nodes and their hostnames are:

Node	Hostnames
controller	<code>airsloop-controller-1</code>
compute	<code>airsloop-compute-1</code>

2.6.4 Network

Physical (underlay) networks are described in Drydock site configuration [here](#). It defines OOB (iLO/IPMI), untagged PXE, and multiple tagged general use networks. Also no bond interfaces are used in Airsloop deployment.

In this deployment the networking is simplified compared to Airship Seaworthy site. There are only two NICs required (excluding oob), one for PXE and another one for the rest of the networks separated using VLAN segmentation. The recommendation is to use the highest bandwidth device available to carry the data networks.

Below is an example of network configuration:

NICs	VLANs	Names	CIDRs
oob	N/A	oob	10.22.104.0/24
pxe	N/A	pxe	10.22.70.0/24
data	71	oam	10.22.71.0/24
	72	calico	10.22.72.0/24
	73	storage	10.22.73.0/24
	74	overlay	10.22.74.0/24

Calico overlay for k8s POD networking uses IPIP mesh.

2.6.5 Storage

Because Airsloop is a minimalistic deployment the required number of disks is just one per node. That disk is not only used by the OS but also by Ceph Journals and OSDs. The way that this is achieved is by using directories and not extra disks for Ceph storage. Ceph OSD configuration can be changed in a [Ceph chart override](#).

2.6.6 Host Profiles

Host profiles in Airship are tightly coupled with the hardware profiles. That means every disk or interface which is described in host profiles should have a corresponding reference to the hardware profile which is being used.

Airship always identifies every NIC or disk by its PCI or SCSI address and that means that the interfaces and the disks that are defined in host and hardware profiles should have the correct PCI and SCSI addresses objectively.

Let's give an example by following the host profile of Airsloop site.

In this [Host Profile](#) is defined that the slave interface that will be used for the pxe boot will be the pxe_nic01. That means a corresponding entry should exist in this [Hardware Profile](#) which it does. So when drydock and maas try to deploy the node it will identify the interface by the PCI address that is written in the Hardware profile.

A simple way to find out which PCI or SCSI address corresponds to which NIC or Disk is to use the lshw command. More information about that command can be found [Here](#).

2.6.7 Extend Cluster

This section describes what changes need to be made to the existing manifests of Airsloop for the addition of an extra compute node to the cluster.

First and foremost the user should go to the [nodes.yaml](#) file and add an extra section for the new compute node.

The next step is to add a similar section as the existing airship-compute-1 section to the [pki-catalog.yaml](#). This is essential for the correct generation of certificates and the correct communication between the nodes in the cluster.

Also every time the user adds an extra compute node to the cluster then the number of OSDs that are managed by this manifest [Ceph-client](#) should be increased by one.

Last step is to regenerate the certificates which correspond to this [certificates.yaml](#) file so the changes in the pki-catalog.yaml file takes place. This can be done through the promenade CLI.

2.6.8 Getting Started

TODO: Specify which node(s) the command(s) in this section are run on. Also if there is an assumption that we have a node with Ubuntu already provisioned, that assumption or steps should be specified along with any Ubuntu version requirements/assumptions and other pre-requisite steps (e.g., installing NTP)

Below are the steps that a user should follow to deploy the Airsloop site:

TODO: Add the prerequisite steps that the user needs to do before starting executing the below steps such as: installing git, installing docker, clone several repos etc.

1. Collect manifests

```
tools/airship pegleg site -r /target collect airship -s collect
```

2. Generate certs

```
tools/airship promenade generate-certs -o /target/certs /target/collect/*.yaml
```

3. Generate genesis.sh script

```
tools/airship promenade build-all -o /target/bundle /target/collect/*.yaml /target/certs/*.yaml
```

4. Execute the genesis.sh script

- `cd /target/bundle`
- `./genesis.sh`

If the genesis.sh script completed successfully

5. Deploy site through shipyard

- `tools/airship shipyard create configdocs design --directory=/target/collect`
- `tools/airship shipyard commit configdocs`
- `tools/airship shipyard create action deploy_site`

6. Check the actions that are already created

```
tools/shipyard get actions
```