

Oracle® Communications

Cloud Native Environment (OCCNE)

Installation Guide



Release 1.6.0
F33920-03
October 2020

The Oracle logo, consisting of a solid red square with the word "ORACLE" in white, uppercase, sans-serif font centered within it.

ORACLE®

Copyright © 2019, 2020, Oracle and/or its affiliates.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish, or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, then the following notice is applicable:

U.S. GOVERNMENT END USERS: Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs) and Oracle computer documentation or other Oracle data delivered to or accessed by U.S. Government end users are "commercial computer software" or "commercial computer software documentation" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, reproduction, duplication, release, display, disclosure, modification, preparation of derivative works, and/or adaptation of i) Oracle programs (including any operating system, integrated software, any programs embedded, installed or activated on delivered hardware, and modifications of such programs), ii) Oracle computer documentation and/or iii) other Oracle data, is subject to the rights and limitations specified in the license contained in the applicable contract. The terms governing the U.S. Government's use of Oracle cloud services are defined by the applicable contract for such services. No other rights are granted to the U.S. Government.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Inside are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Epyc, and the AMD logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information about content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services unless otherwise set forth in an applicable agreement between you and Oracle. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services, except as set forth in an applicable agreement between you and Oracle.

Contents

1 Overview

Key terms	1-1
Key Acronyms and Abbreviations	1-2

2 Installation Procedures

Bare Metal Installation	2-1
OCCNE Installation Overview	2-1
Frame and Component Overview	2-1
Create OCCNE Instance	2-5
Installation Prerequisites	2-7
Configure Artifact Acquisition and Hosting	2-8
Initial Configuration - Prepare a Minimal Boot Strapping Environment	2-11
Installation of Oracle Linux 7.5 on Bootstrap Host	2-11
Configure the Installer Bootstrap Host BIOS	2-13
Configure Top of Rack 93180YC-EX Switches	2-17
Configure Addresses for RMS iLOs, OA, EBIPA	2-30
Configure Legacy BIOS on Remaining Hosts	2-37
Configure Enclosure Switches	2-43
Bastion Host Installation	2-48
Automated Installation	2-55
Virtualized CNE Installation	2-58
Prerequisites	2-58
Limitation/Expectations	2-59
Download qcow2 image	2-59
Upload an Image to an OpenStack Environment	2-60
Bootstrap Host Creation	2-60
Pre-deployment Configuration	2-62
Deploy the OCCNE Virtualized Cluster	2-74

3 Post Installation Activities

Post Install Verification	3-1
---------------------------	-----

A Artifact Acquisition and Hosting

Docker Image Registry Configuration	A-1
HTTP Repository Configuration	A-4
YUM Repository Configuration	A-5

B Reference Procedures

Inventory File Preparation	B-1
Installation Preflight Checklist	B-7
Inventory File Template	B-35
Install Additional Services/Network Functions	B-37
Change MySQL root user password	B-38
MySQL Repository Requirements	B-38
Oracle Linux 7.5 Download Instructions	B-39
vCNE VM Sizing	B-40
Environmental Variables	B-42

List of Figures

2-1	Frame Overview	2-2
2-2	Host Designations	2-3
2-3	Node Roles	2-4
2-4	Transient Roles	2-5
2-5	OCCNE Installation Overview	2-6
B-1	Rackmount ordering	B-8

List of Tables

1-1	Key Terms	1-1
1-2	Acronyms	1-2
2-1	OCCNE Artifacts	2-8
2-2	Terminology used in Procedure	2-48
2-3	Environmental Variables	2-55
3-1	Credentials	3-5
B-1	Base Groups	B-2
B-2	Data Tier Groups	B-4
B-3	Edit occne:vars	B-6
B-4	Edit mysqlndb_all_nodes:vars	B-7
B-5	Enclosure Switch Connections	B-8
B-6	ToR Switch Connections	B-10
B-7	Rackmount Server Connections	B-13
B-8	Complete Site Survey Subnet Table	B-14
B-9	Complete Site Survey Host IP Table	B-15
B-10	Complete VM IP Table	B-16
B-11	Complete OA and Switch IP Table	B-17
B-12	ToR and Enclosure Switches Variables Table (Switch Specific)	B-23
B-13	Complete Site Survey Repository Location Table	B-25
B-14	OCCNE Variables	B-35
B-15	Bootstrap VM	B-41
B-16	Kubernetes Master Node	B-41
B-17	Kubernetes Worker Node	B-41
B-18	Bastion Host VMs	B-42
B-19	Management nodes	B-42
B-20	Data Nodes	B-42
B-21	SQL Nodes	B-42
B-22	Environmental Variables	B-43

My Oracle Support

My Oracle Support (<https://support.oracle.com>) is your initial point of contact for all product support and training needs. A representative at Customer Access Support can assist you with My Oracle Support registration.

Call the Customer Access Support main number at 1-800-223-1711 (toll-free in the US), or call the Oracle Support hotline for your local country from the list at <http://www.oracle.com/us/support/contact/index.html>. When calling, make the selections in the sequence shown below on the Support telephone menu:

1. Select **2** for New Service Request.
2. Select **3** for Hardware, Networking and Solaris Operating System Support.
3. Select one of the following options:
 - For Technical issues such as creating a new Service Request (SR), select **1**.
 - For Non-technical issues such as registration or assistance with My Oracle Support, select **2**.

You are connected to a live agent who can assist you with My Oracle Support registration and opening a support ticket.

My Oracle Support is available 24 hours a day, 7 days a week, 365 days a year.

What's New in This Guide

New and Updated features in Release 1.6.0

Following procedures are newly added in this release:

1. MetalLB Configuration: Procedure to configure MetalLB Configuration is added under [Pre-deployment Configuration](#) section.
2. Configuration for using OCCNE LB (MetalLB): Procedure to configure occne_lb_flavor field is explained in [Pre-deployment Configuration](#) section.

1

Overview

Oracle Communications Cloud Native Environment (OCCNE) is an essence infrastructure code that provisions, configures, and manages reference cloud native environments.

This document details the procedure for installing the **Oracle Communications Cloud Native Environment (OCCNE)**. The intended audiences for this document are Oracle engineers who work with customers to install the **Cloud Native Environment (CNE)** on-site at customer facilities.

Key terms

The following table lists the terms used in this document.

Table 1-1 Key Terms

Term	Definition
Host	A computer running an instance of an operating system with an IP address. Hosts can be virtual or physical. The HP DL380 Gen10 Rack Mount Servers and BL460c Gen10 Blades are physical hosts. KVM based virtual machines are virtual hosts. Hosts are also referred to as nodes, machines, or computers.
Database Host	The Database (DB) Host is a physical machine that hosts guest virtual machines which in turn provides OCCNE's MySQL service and Database Management System (DBMS). The Database Hosts are comprised of two Rack Mount Servers (RMSs) below the <i>Top of Rack (TOR)</i> switches. For some customers, the DB host(s) can be VM such as in the vCNE environment.
Management Host	The Management Host is a physical machine in the frame that has a special configuration to support hardware installation and configuration of other components within a frame. For CNE, there is one machine with dedicated connectivity to out of band (OOB) interfaces on the Top of Rack switches. The OOB interfaces provide connectivity needed to initialize the ToR switches. In OCCNE, the Management Host role and Database Host roles are assigned to the same physical machine. When referring to a machine as a "Management Host", the context is with respect to its OOB connections which are unique to the Management Host hardware.
Bastion Host	The Bastion Host provides general orchestration support for the site. The Bastion Host runs as a virtual machine on a Database Host. Sometimes referred to as the Management VM. During the install process, the Bastion Host is used to host the automation environment and execute install automation. The install automation provisions and configures all other hosts, nodes, and switches within the frame. After the install process is completed, the Bastion Host continues to serve as the customer gateway to cluster operations and control.

Table 1-1 (Cont.) Key Terms

Term	Definition
Installer Bootstrap Host	As an early step in the site installation process, one of the hosts (which is eventually re-provisioned as a Database Server) is minimally provisioned to act as an Installer Bootstrap Host. The Installer Bootstrap Host has a very short lifetime as its job is to provision the first Database Server. Later in the install process, the server being used to host the Bootstrap server is re-provisioned as another Database Server. The Installer Bootstrap Host is also referred to simply as the Bootstrap Host.
Node	A logical computing node in the system. A node is usually a networking endpoint. May or may not be virtualized or containerized. Database nodes refer to hosts dedicated primarily to running Database services. Kubernetes nodes refer to hosts dedicated primarily to running Kubernetes.
Master Node	Some nodes in the system (three RMSs in the middle of the equipment rack) are dedicated to providing Container management. These nodes are responsible for managing all of the containerized services (which run on the worker nodes).
Worker Node	Some nodes in the system (the blade servers at the bottom of the equipment rack) are dedicated to hosting Containerized software and providing the 5G application services.
Container	An encapsulated software service. All 5G applications and OAM functions are delivered as containerized software. The purpose of the OCCNE is to host containerized software providing 5G Network Functions and services.
Cluster	A collection of hosts and nodes dedicated to providing either Database or Containerized services and applications. The Database service is comprised of the collection of Database nodes and is managed by MySQL. The Container cluster is comprised of the collection of Master and Worker Nodes and is managed by Kubernetes.
Virtualized CNE	A virtualized CNE is a cloud native environment that is deployed on VMs, rather than on bare metal servers.

Key Acronyms and Abbreviations

The following table lists the abbreviations, and acronyms specific to this document.

Table 1-2 Acronyms

Acronyms	Definition
5G NF	3GPP 5G Network Function
BIOS	Basic Input Output System
CLI	Command Line Interface
CNE	Cloud Native Environment
DB	Database
DBMS	Database Management System
DHCP(D)	Dynamic Host Configuration Protocol
DNS	Domain Name Server

Table 1-2 (Cont.) Acronyms

Acronyms	Definition
EBIPA	Enclosure Bay IP Addressing
FQDN	Fully Qualified Domain name
GUI	Graphical User Interface
HDD	Hard Disk Drive
HP	Hewlett Packard
HPE	Hewlett Packard Enterprise
HTTP	HyperText Transfer Protocol
iLO	HPE Integrated Lights-Out Management System
IP	Internet Protocol; may be used as shorthand to refer to an IP layer 3 address.
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
IRF	Intelligent Resilient Framework (IRF) is a proprietary software virtualization technology developed by H3C (3Com). Its core idea is to connect multiple network devices through physical IRF ports and perform necessary configurations, and then these devices are virtualized into a distributed device.
ISO	International Organization for Standardization; typically used as shorthand to refer to an ISO 9660 optical disk file system image
KVM	Keyboard, Video, Mouse
K8s	Shorthand alias for Kubernetes
MAC	Media Access Control address
MBE	Minimal Bootstrapping Environment
NFS	Network File System
NTP	Network Time Protocol
OA	HP BladeSystem Onboard Administrator
OAM	Operations, Administration, Maintenance
OCCNE	Oracle Communications Signaling, Network Function Cloud Native Environment
OS	Operating System
OSDC	Oracle Software Download Center
PKI	Public Key Infrastructure
POAP	PowerOn Auto Provisioning
PXE	Pre-Boot Execution Environment
RAID	Redundant Array of Independent Disks
RAM	Random Access Memory
RBSU	ROM Based Setup Utility
RMS	Rack Mount Server
RPM	Red Hat Package Manager
SAS	Serial Attached SCSI
SSD	Solid State Drive
TAR	Short for Tape Archive, and sometimes referred to as tarball, a file that has the TAR file extension is a file in the Consolidated Unix Archive format.

Table 1-2 (Cont.) Acronyms

Acronyms	Definition
TLA	Three Letter Acronym
TLD	Top Level Domain
ToR	Top of Rack - Colloquial term for the pair of Cisco 93180YC-EX switches
UEFI	Unified Extensible Firmware Interface
URL	Uniform Resource Locator
VM	Virtual Machine
vCNE	Virtualized CNE
VSP	Virtual Serial Port
YUM	Yellowdog Updator, Modified (a Linux Package Manager)

2

Installation Procedures

The installation procedures in this document explain the steps to provision and configure the Oracle Communications Cloud Native Environment (OCCNE). OCCNE offers the choice of deployment platform; the CNE can be deployed directly onto [dedicated hardware](#), (referred to as a bare metal CNE), or deployed onto [OpenStack-hosted VMs](#). (referred to as a virtualized CNE).

Regardless of which deployment platform is selected, OCCNE installation is highly automated. A collection of container-based utilities are used to automate the provisioning, installation, and configuration of OCCNE. These utilities are based on the following automation tools:

- PXE helps reliably automate provisioning the hosts with a minimal operating system.
- Terraform is used to create the virtual resources that the virtualized CNE is hosted on.
- Kubespray helps reliably install a base Kubernetes cluster, including all dependencies (like etcd), using the Ansible provisioning tool.
- Ansible is used to orchestrate the overall deployment.
- Helm is used to deploy and configure common services such as Prometheus, Grafana, ElasticSearch and Kibana.

Note:

Make sure that the shell is configured with Keepalive to avoid unexpected timeout.

Bare Metal Installation

This section describes the procedure to install OCCNE onto dedicated bare metal hardware.

OCCNE Installation Overview

Frame and Component Overview

The initial release of the OCCNE system provides support for on-prem deployment to a very specific target environment consisting of a frame holding switches and servers. This section describes the layout of the frame and the roles performed by the racked equipment.

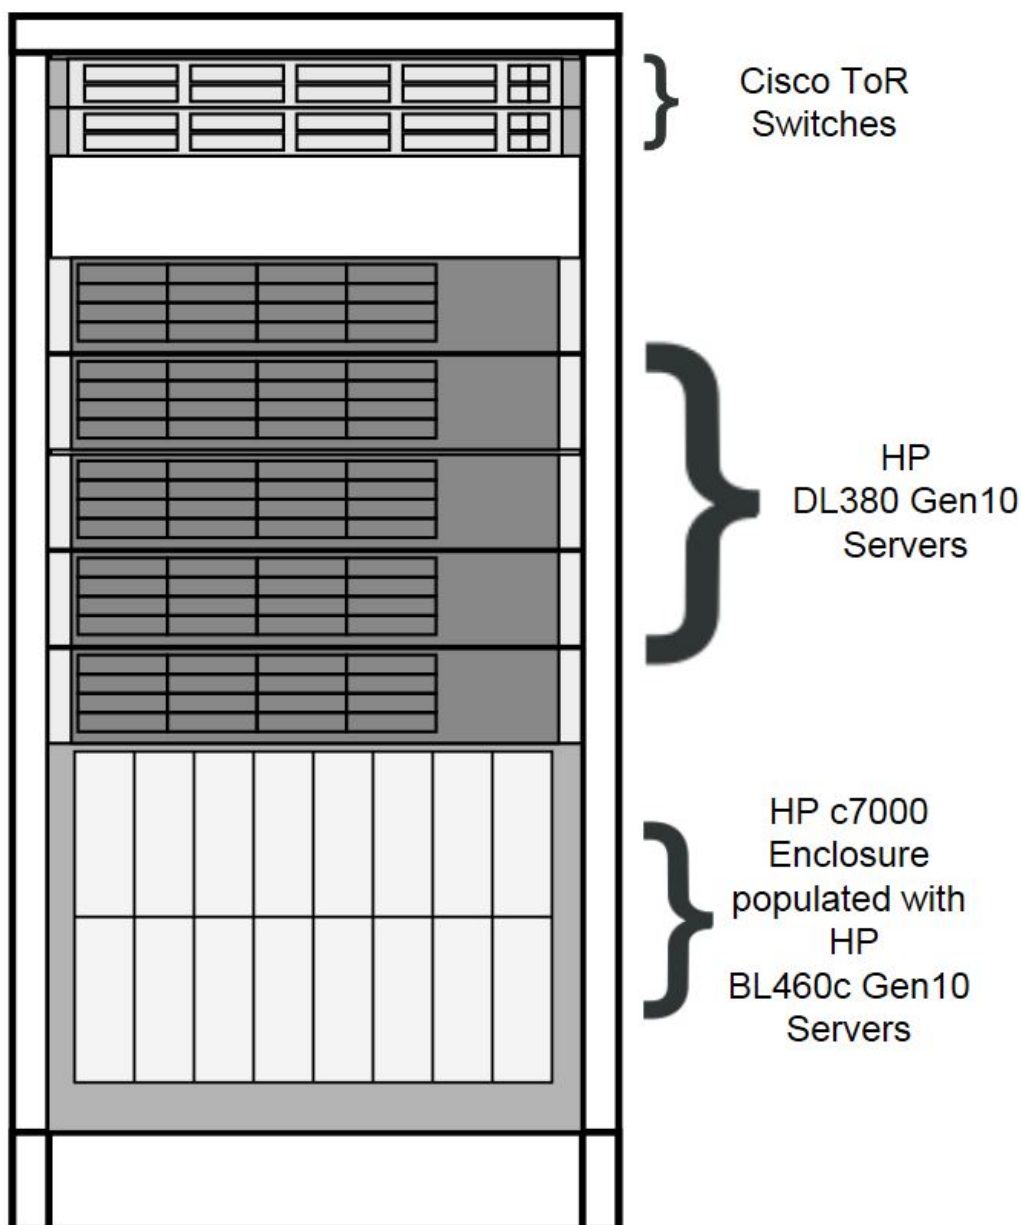
 **Note:**

In the installation process, some of the roles of servers change as the installation procedure proceeds.

Frame Overview

The physical frame is comprised of an HP c-Class enclosure and 5 DL380 rack mount servers, and 2 Top of Rack (ToR) Cisco switches. The frame components are added from the bottom up, thus designations found in the next section number from the bottom of the frame to the top of the frame.


Figure 2-1 Frame Overview



Host Designations

Each physical server has a specific role designation within the CNE solution.


Figure 2-2 Host Designations



Node Roles

Along with the primary role of each host, a secondary role may be assigned. The secondary role may be software related, or, in the case of the Bootstrap Host, hardware related, as there are unique OOB connections to the ToR switches.


Figure 2-3 Node Roles



Transient Roles

Transient role is unique in that it has OOB connections to the ToR switches, which includes the designation of Bootstrap Host. This role is only relevant during initial switch configuration and disaster recovery of the switch. RMS1 also has a transient role as the Installer Bootstrap Host, which is only relevant during initial install of the frame, and subsequent to getting an official install on RMS2, this host is re-paved to its Storage Host role.

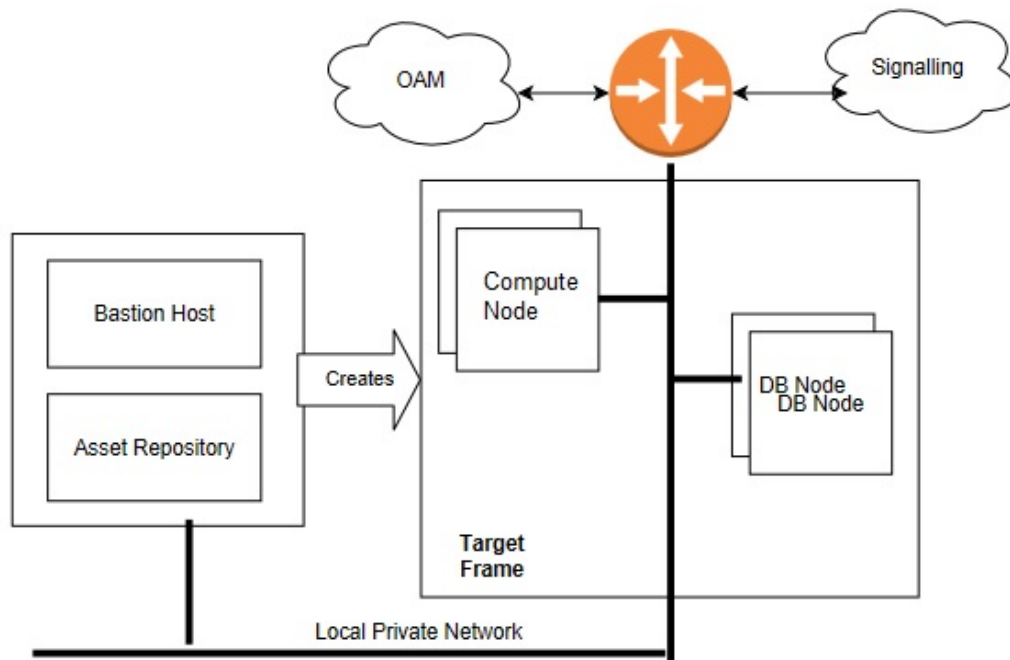
Figure 2-4 Transient Roles



Create OCCNE Instance

This section describes the procedures required to create the OCCNE instance at a customer site. The following diagram shows the installation context:

Figure 2-5 OCCNE Installation Overview



Following is the basic installation flow to understand the overall effort:

1. Check that the hardware is on-site and properly cabled and powered up.
2. Pre-assemble the basic ingredients needed to perform a successful install:
 - a. **Identify**
 - i. Download and stage software and other configuration files using provided manifests. Refer to [Artifact Acquisition and Hosting](#) for manifests information.
 - ii. Identify the layer 2 (MAC) and layer 3 (IP) addresses for the equipment in the target frame
 - iii. Identify the addresses of key external network services (e.g., NTP, DNS, etc.)
 - iv. Verify / Set all of the credentials for the target frame hardware to known settings
 - b. **Prepare**
 - i. Software Repositories: Load the various SW repositories (YUM, Helm, Docker, etc.) using the downloaded software and configuration
 - ii. Configuration Files: Populate the hosts inventory file with credentials and layer 2 and layer 3 network information, switch configuration files with assigned IP addresses, and yaml files with appropriate information.
3. Bootstrap the System:
 - a. Manually configure a Minimal Bootstrapping Environment (MBE); perform the minimal set of manual operations to enable networking and initial loading of a single Rack Mount Server - RMS1 - the transient Installer Bootstrap Host. In

this procedure, a minimal set of packages needed to configure switches, iLOs, PXE boot environment, and provision [RMS2](#) as an OCCNE Storage Host are installed.

- b. Using the newly constructed MBE, automatically create the first (complete) Management VM on RMS2. This freshly installed Storage Host will include a virtual machine for hosting the Bastion Host.
 - c. Using the newly constructed Bastion Host on [RMS2](#), automatically deploy and configure the OCCNE on the other servers in the frame.
4. Final Steps
 - a. Perform post installation checks
 - b. Perform recommended security hardening steps

Cluster Bootstrapping Overview

This install procedure is targeted to install the OCCNE onto a new hardware absent of any networking configurations to switches, or operating systems provisioned. Therefore, the initial step in the installation process is to provision RMS1 (see [Installation Procedures](#)) as a temporary Installer Bootstrap Host. The Bootstrap Host is configured with a minimal set of packages needed to configure switches, iLOs, PXE boot environment, and provision RMS2 as an OCCNE Storage Host. A virtual Bastion Host is also provisioned on RMS2. The Bastion Host is then used to provision (and in the case of the Bootstrap Host, re-provision) the remaining OCCNE hosts and install Kubernetes, Database services, and Common Services running within the Kubernetes cluster.

Installation Prerequisites

Complete the procedures outlined in this section before moving on to the [Install Procedures](#) section. OCCNE installation procedures require certain artifacts and information to be made available prior to executing installation procedures. Refer to [Configure Artifact Acquisition and Hosting](#) for the prerequisites.

Obtain Mate Site DB Replication Service Load Balancer IP

While installing MYSQL NDB on the second site the Mate Site DB Replication Service Load Balancer IP must be provided as the configuration parameter for the geo-replication process to start.

1. Login to Bastion Host of the first site.
2. Fetch DB Replication Service Load Balancer IP of Mate Site MYSQL NDB:

```
$ kubectl get svc --namespace=occne-infra | grep replication
```

Example:

```
$ kubectl get svc --namespace=occne-infra | grep replication
occne-db-replication-svc      LoadBalancer    10.233.3.117
10.75.182.88                  80:32496/TCP    2m8s
```

In the above example IPv4: 10.75.182.88 is the Mate Site DB Replication Service Load Balancer IP.

Configure Artifact Acquisition and Hosting

OCCNE requires artifacts from Oracle eDelivery and certain open-source projects. OCCNE deployment environments are not expected to have direct internet access. Thus, customer-provided intermediate repositories are necessary for the OCCNE installation process. These repositories will need OCCNE dependencies to be loaded into them. This section will address the artifacts list needed to be in these repositories.

Oracle eDelivery Artifact Acquisition

The OCCNE artifacts are posted on Oracle Software delivery Cloud (OSDC) and/or OHC.

Table 2-1 OCCNE Artifacts

Artifact	Description	File Type	Destination repository
occne-images-1.6.0.tgz	OCCNE Installers (Docker images) from OSDC (edelivery.oracle.com)	Tar GZ	Docker Registry
v980756-01.zip	Zip file of MySQL Cluster Manager 1.4.7+Cluster. Obtain p29060743_140_Linux-x86-64.zip from MOS (support.oracle.com) under patch 29060743 and rename it.	Zip of tar file	File repository
V987059-01.zip	Zip file of MySQL Cluster Manager 1.4.8+Cluster. Obtain from OSDC (edelivery.oracle.com)	Zip of tar file	File repository
V978737-01.iso	Oracle Linux Release 7 Update 5 for x86 (64 bit) from OSDC (edelivery.oracle.com)	ISO	File repository
Templates	Switch config files, metallb config file, snmp mib files, sample hosts.ini file vCNE deploy script and configuration files from OHC (docs.oracle.com)	Config files (.conf, .ini, .yaml, .mib, .sh, .txt)	Local media

Third Party Artifacts

OCCNE dependencies that come from open-source software must be available in repositories reachable by the OCCNE installation tools. For an accounting of third party artifacts needed for this installation, refer to the [Artifact Acquisition and Hosting](#).

Populate the MetalLB Configuration

Introduction

The metalLB configMap file (mb_configmap.yaml) contains the manifest for the metalLB configMap, this defines the BGP peers and address pools for metalLB. This file (mb_configmap.yaml) must be placed in the same directory (*/var/ocne/<cluster_name>*) as the hosts.ini file.

Following is the procedure to configure MetalLB pools and peers:

1. Add BGP peers and address groups: Referring to the data collected in the [Preflight Checklist](#), add BGP peers (ToRswitchA_Platform_IP, ToRswitchB_Platform_IP) and address groups for each address pool. Address-pools list the IP addresses that metalLB is allowed to allocate.
2. Edit the mb_configmap.yaml file with the site-specific values found in the [Preflight Checklist](#)

 **Note:**

The name "signaling" is prone to different spellings (UK vs US), therefore pay special attention to how this signaling pool is referenced.

```
configInline:
  peers:
    - peer-address: <ToRswitchA_Platform_IP>
      peer-asn: 64501
      my-asn: 64512
    - peer-address: <ToRswitchB_Platform_IP>
      peer-asn: 64501
      my-asn: 64512
  address-pools:
    - name: signaling
      protocol: bgp
      auto-assign: false
      addresses:
        - '<MetalLB_Signal_Subnet_IP_Range>'
    - name: oam
      protocol: bgp
      auto-assign: false
      addresses:
        - '<MetalLB_OAM_Subnet_IP_Range>'
```

Install Backup Bastion Host

Introduction

This procedure details the steps necessary to install the Backup Bastion Host on the Storage Host db-1/RMS1 and backing up the data from the active Bastion Host on db-2/RMS2 to the Backup Bastion Host.

Prerequisites

1. Bastion Host is already created on Storage Host db-2/RMS2.
2. Storage Host db-2/RMS2 and the Backup Bastion Host are defined in the Customer hosts.ini file as defined in procedure: [Inventory File Preparation](#).

 **Note:**

If the initial bootstrap host is RMS1 then the bastion host is created on the RMS2 and backup bastion host is created on the RMS1.

3. Host names and IP Address, network information assigned to Backup Management VM are captured in the [Installation PreFlight Checklist](#).
4. All the Network information should be configured in [Inventory File Preparation](#).

Expectations

1. Bastion Host VM on Storage Host db-1/RMS1 is created as a backup for Bastion Host VM on Storage Host db-2/RMS2.
2. All the required config files and data configured in the Backup Bastion Host on Storage Host db-1/RMS1 are copied from the active Bastion Host on Storage Host db-2/RMS2.

Procedure

All commands are executed from the active Bastion Host on db-2/RMS2.

Create the Backup Bastion Host on Storage Host db-1/RMS1

1. Login to the active Bastion Host (VM on RMS2) using the admusr/***** credentials.
2. Execute the deploy.sh script from the `/var/occne/` directory with the required parameters set.

```
$ export CENTRAL_REPO=<customer specific repo name>
$ export CENTRAL_REPO_IP=<customer_specific_repo_ipv4>
$ export OCCNE_CLUSTER=<cluster_name>
$ export OCCNE_BASTION=<bastion_full_name>
$ ./deploy.sh
```

Customer Example:

```
$ export CENTRAL_REPO=central-repo
$ export CENTRAL_REPO_IP=10.10.10.10
$ export OCCNE_CLUSTER=rainbow
$ export OCCNE_BASTION=bastion-1.rainbow.lab.us.oracle.com
$ ./deploy.sh
```

Note: The above example can be executed like the following:

```
CENTRAL_REPO=central-repo
CENTRAL_REPO_IP=10.10.10.10 OCCNE_CLUSTER=rainbow
OCCNE_BASTION=bastion-1.rainbow.lab.us.oracle.com ./deploy.sh
```

3. Verify installation of the Backup Bastion Host.

 **Note:**

The IP of the backup Bastion Host can be derived from the hosts.ini file under the group host_kernel_virtual for db-1 ansible_host IP.

```
$ ssh -i ~/.ssh/id_rsa admusr@<backup_bastion_ip_address>
```

Initial Configuration - Prepare a Minimal Boot Strapping Environment

In the first step of the installation, a minimal bootstrapping environment is established that is to support the automated installation of the CNE environment. The steps in this section provide the details necessary to establish this minimal bootstrap environment on the Installer Bootstrap Host using a Keyboard, Video, Mouse (KVM) connection.

Installation of Oracle Linux 7.5 on Bootstrap Host

This procedure outlines the installation steps for installing the OL7 onto the OCCNE Installer Bootstrap Host. This host is used to configure the networking throughout the system and install OL7 onto RMS1. It is re-paved as a Database Host in a later procedure.

Prerequisites

1. USB drive of sufficient size to hold the ISO (approximately 5Gb)
2. Oracle Linux 7.5 iso
3. YUM repository file
4. Keyboard, Video, Mouse (KVM)

Limitations and Expectations

1. The configuration of the Installer Bootstrap Host is meant to be quick and easy, without a lot of care on appropriate OS configuration. The Installer Bootstrap Host is re-paved with the appropriate OS configuration for cluster and DB operation at a later stage of installation. The Installer Bootstrap Host needs a Linux OS and some basic network to get the installation process started.
2. All steps in this procedure are performed using Keyboard, Video, Mouse (KVM).

References

1. Oracle Linux 7 Installation guide: https://docs.oracle.com/cd/E52668_01/E54695/html/index.html
2. [HPE Proliant DL380 Gen10 Server User Guide](#)

Bootstrap Install Procedure

1. Create Bootable USB Media:
 - a. Download the Oracle Linux 7.5.

On the installer's notebook, download the OL ISO from the customer's repository. Since Installer notebook may be Windows or Linux OS, and the Customer repository location may vary so the user executing this procedure determines the appropriate detail to execute this task.

- b. Push the OL ISO image onto the USB Flash Drive.
Since the installer's notebook may be Windows or Linux OS-based, the user executing this procedure determines the appropriate detail to execute this task. For a Linux based notebook, insert a USB Flash Drive of the appropriate size into a Laptop (or some other linux host where the iso can be copied to), and run the `dd` command to create a bootable USB drive with the Oracle Linux 7 iso.

```
$ dd -if=<path to ISO> -of=<USB device path> -bs=1m
```

2. Install OL7 on the Installer Bootstrap Host:
 - a. Connect a Keyboard, Video, and Mouse (KVM) into the Installer Bootstrap Host's monitor and USB ports.
 - b. Plug the USB flash drive containing the bootable iso into an available USB port on the Bootstrap host (usually in the front panel).
 - c. Reboot the host by momentarily pressing the power button on the host's front panel. The button will go yellow. If it holds at yellow, press the button again. The host should auto-boot to the USB flash drive.

 **Note:**

If the host was previously configured and the USB is not a bootable path in the boot order, it may not boot successfully.

- d. If the host does not boot to the USB, repeat step 3, and interrupt the boot process by pressing **F11** which brings up the Boot Menu. If the host has been recently booted with an OL, the **Boot Menu** will display **Oracle Linux** at the top of the list. Select **Generic USB Boot** as the first boot device and proceed.
- e. The host attempts to boot from the USB. The following menu is displayed on the screen. Select **Test this media & install Oracle Linux 7.x** and click ENTER. This begins the verification of the media and the boot process. After the verification reaches 100%, the **Welcome** screen is displayed. When prompted for the language to use, select the default setting: **English (United States)** and click **Continue** in the lower left corner.
- f. The **INSTALLATION SUMMARY** page, is displayed. The following settings are expected. If any of these are not set correctly then please select that menu item and make the appropriate changes.
 - i. **LANGUAGE SUPPORT:** English (United States)
 - ii. **KEYBOARD:** English (US)
 - iii. **INSTALLATION SOURCE:** Local Media
 - iv. **SOFTWARE SELECTION:** Minimal Install

INSTALLATION DESTINATION should display `No disks selected`. Select **INSTALLATION DESTINATION** to indicate the drive to install the OS on.

Select the first HDD drive (in this case that would be the first one listed or the 1.6 TB disk) and select **DONE** in the upper right corner.

If the server has already been installed a red banner at the bottom of the page may indicate there is an error condition. Selecting that banner causes a dialog to appear indicating there is not enough free space (which might mean an OS has already been installed). In the dialog it may show both 1.6 TB HDDs as claimed or just the one.

If only one HDD is displayed (or it could be both 1.6 TB drives selected, select the **Reclaim space** button. Another dialog appears. Select the **Delete all** button and the **Reclaim space** button again. Select **DONE** to return to the **INSTALLATION SUMMARY** screen.

If the disk selection dialog appears (after selecting the red banner at the bottom of the page), this implies a full installation of the RMS has already been performed (usually this is because the procedure had to be restarted after it was successfully completed). In this case select the **Modify Disk Selection**. This will return to the disk selection page. Select both HDDs and hit done. The red banner should now indicate the space must be reclaimed. The same steps to reclaim the space can be performed.

- g. Select **DONE**. This returns to the **INSTALLATION SUMMARY** page.
- h. At the **INSTALLATION SUMMARY** screen, select **Begin Installation**. The **CONFIGURATION** screen is displayed.
- i. At the **CONFIGURATION** screen, select **ROOT PASSWORD**. Enter a root password appropriate for this installation. It is good practice to use a customer provided secure password to minimize the host being compromised during installation.
- j. At the conclusion of the install, remove the USB and select **Reboot** to complete the install and boot to the OS on the host. At the end of the boot, the login prompt appears.

Configure the Installer Bootstrap Host BIOS

Introduction

These procedures define the steps necessary to set up the Legacy BIOS changes on the Bootstrap host using the KVM. Some of the procedures in this document require a reboot of the system and are indicated in the procedure.

Prerequisites

Procedure [OCCNE Installation of Oracle Linux 7.5 on Bootstrap Host](#) is complete.

Limitations and Expectations

1. Applies to HP Gen10 iLO 5 only.
2. The procedures listed here applies to the Bootstrap host only.

Steps to OCCNE Configure the Installer Bootstrap Host BIOS

1. Expose the System Configuration Utility: This step details how to expose the HP iLO 5 System Configuration Utility main page from the KVM. It does not provide instructions on how to connect the console as these may be different on each installation.

- a. After making the proper connections for the KVM on the back of the Bootstrap host to have access to the console, the user should reboot the host by momentarily pressing the power button on the front of the Bootstrap host.
 - b. Expose the **HP Proliant DL380 Gen10 System Utilities**.
Once the remote console has been exposed, the system must be reset to force it through the restart process. When the initial window is displayed, hit the F9 key repeatedly. Once the F9 is highlighted at the lower left corner of the remote console, it should eventually bring up the main System Utility.
 - c. The System Utilities screen is exposed in the remote console.
2. Change over from UEFI Booting Mode to Legacy BIOS Booting Mode: The System Utility must default the booting mode to UEFI or has been changed to UEFI, it will be necessary to switch the booting mode to Legacy.
 - a. Expose the System Configuration Utility by following Step 1.
 - b. Select System Configuration.
 - c. Select BIOS/Platform Configuration (RBSU).
 - d. Select Boot Options.
If the Boot Mode is set to UEFI Mode then this procedure should be used to change it to Legacy BIOS Mode.
Note: The server reset must go through an attempt to boot before the changes will actually apply.
 - e. The user is prompted to select the **Reboot Required** popup dialog. This will drop back into the boot process. The boot must go into the process of actually attempting to boot from the boot order. This should fail since the disks have not been installed at this point. The System Utility can be accessed again.
 - f. After the reboot and the user re-enters the System Utility, the Boot Options page should appear.
 - g. Select **F10: Save** if it's desired to save and stay in the utility or select the **F12: Save and Exit** if its desired to save and exit to complete the current boot process.
 3. Adding a New User Account: This step provides the steps required to add a new user account to the server iLO 5 interface.

 **Note:**

This user must match the pxe_install_lights_out_usrfields as provided in the hosts inventory files created using the template: [OCCNE Inventory File Preparation](#).

- a. Expose the System Utility by following Step 1.
- b. Select System Configuration.
- c. Select iLO 5 Configuration Utility.
- d. Select **User Management** → **Add User** .
- e. Select the appropriate permissions. For the root user set all permissions to **YES**. Enter **root** as **New User Name** and **Login Name** fields, and enter <password> in the **Password** field.

- f. Select **F10: Save** to save and stay in the utility or select the **F12: Save and Exit** to save and exit, to complete the current boot process.
4. Force PXE to boot from the first Embedded FlexibleLOM HPE Ethernet 10Gb 2-port Adapter. During host PXE, the DHCP DISCOVER requests from the hosts must be broadcast over the 10Gb port. This step provides the steps necessary to configure the broadcast to use the 10Gb ports before it attempts to use the 1Gb ports. Moving the 10Gb port up on the search order helps to speed up the response from the host servicing the DHCP DISCOVER. Enclosure blades have 2 10GE NICs which default to being configured for PXE booting. The RMS are re-configured to use the PCI NICs using this step.
 - a. Expose the System Utility by following Step 1.
 - b. Select System Configuration.
 - c. Select BIOS/Platform Configuration (RBSU).
 - d. Select Boot Options.

This menu defines the boot mode which should be set to Legacy BIOS Mode, the UEFI Optimized Boot which should be disabled, and the Boot Order Policy which should be set to Retry Boot Order Indefinitely (this means it will keep trying to boot without ever going to disk). In this screen select Legacy BIOS Boot Order. If not in Legacy BIOS Mode, please follow procedure 2.2 Change over from UEFI Booting Mode to Legacy BIOS Booting Mode to set the Configuration Utility to Legacy BIOS Mode.
 - e. Select Legacy BIOS Boot Order

This page defines the legacy BIOS boot order. This includes the list of devices from which the server will listen for the DHCP OFFER (includes the reserved IPv4) after the PXE DHCP DISCOVER message is broadcast out from the server.

In the default view, the 10Gb Embedded FlexibleLOM 1 Port 1 is at the bottom of the list. When the server begins the scan for the response, it scans down this list until it receives the response. Each NIC will take a finite amount of time before the server gives up on that NIC and attempts another in the list. Moving the 10Gb port up on this list should decrease the time that is required to finally process the DHCP OFFER.

To move an entry, select that entry, hold down the first mouse button and move the entry up in the list below the entry it must reside under.
 - f. Move the 10 Gb Embedded FlexibleLOM 1 Port 1 entry up above the 1Gb Embedded LOM 1 Port 1 entry.
 - g. Select **F10: Save** to save and stay in the utility or select the **F12: Save and Exit** to save and exit, to complete the current boot process.
5. Enabling Virtualization: This step provides the steps required to enable virtualization on a given Bare Metal Server. Virtualization can be configured using the default settings or via the Workload Profiles.
 - a. Verifying Default Settings
 - i. Expose the System Configuration Utility by following Step 1.
 - ii. Select System Configuration.
 - iii. Select BIOS/Platform Configuration (RBSU)
 - iv. Select Virtualization Options

This screen displays the settings for the Intel(R) Virtualization Technology (IntelVT), Intel(R) VT-d, and SR-IOV options (Enabled or Disabled). The default values for each option is Enabled.

- v. Select **F10: Save** to save and stay in the utility or select the **F12: Save** and **Exit** to save and exit, to complete the current boot process.
6. Disable RAID Configurations:
 - a. Expose the System Configuration Utility by following Step 1.
 - b. Select System Configuration.
 - c. Select Embedded RAID 1 : HPE Smart Array P408i-a SR Gen 10.
 - d. Select Array Configuration.
 - e. Select Manage Arrays.
 - f. Select Array A (or any designated Array Configuration if there are more than one).
 - g. Select Delete Array.
 - h. Select Submit Changes.
 - i. Select **F10: Save** to save and stay in the utility or select the **F12: Save** and **Exit** to save and exit, to complete the current boot process.
7. Enable the Primary Boot Device: This step provides the steps necessary to configure the primary bootable device for a given Gen10 Server. In this case the RMS would include two devices as Hard Drives (HDDs). Some configurations may also include two Solid State Drives (SSDs). The SSDs are not to be selected for this configuration. Only the primary bootable device is set in this procedure since RAID is being disabled. The secondary bootable device remains as Not Set.
 - a. Expose the System Configuration Utility by following Step 1.
 - b. Select System Configuration.
 - c. Select Embedded RAID 1 : HPE Smart Array P408i-a SR Gen 10.
 - d. Select Set Bootable Device(s) for Legacy Boot Mode. If the boot devices are not set then it will display Not Set for the primary and secondary devices.
 - e. Select **Select Bootable Physical Drive**.
 - f. Select Port 1| Box:3 Bay:1 Size:1.8 TB SAS HP EG00100JWJNR.
Note: This example includes two HDDs and two SSDs. The actual configuration may be different.
 - g. Select Set as Primary Bootable Device.
 - h. Select Back to Main Menu.
This will return to the HPE Smart Array P408i-a SR Gen10 menu. The secondary bootable device is left as Not Set.
 - i. Select **F10: Save** to save and stay in the utility or select the **F12: Save** and **Exit** to save and exit, to complete the current boot process.
8. Configure the iLO 5 Static IP Address: When configuring the Bootstrap host, the static IP address for the iLO 5 must be configured.

 **Note:**

This step requires a reboot after completion.

- a. Expose the System Configuration Utility by following Step 1.
- b. Select System Configuration.
- c. Select iLO 5 Configuration Utility.
- d. Select Network Options.
- e. Enter the IP Address, Subnet Mask, and Gateway IP Address fields provided in [Installation PreFlight Checklist](#).
- f. Select **F12: Save** and **Exit** to complete the current boot process. A reboot is required when setting the static IP for the iLO 5. A warning appears indicating that the user must wait 30 seconds for the iLO to reset and then a reboot is required. A prompt appears requesting a reboot. Select **Reboot**.
- g. Once the reboot is complete, the user can re-enter the System Utility and verify the settings if necessary.

Configure Top of Rack 93180YC-EX Switches

Introduction

This procedure provides the steps required to initialize and configure Cisco 93180YC-EX switches as per the topology defined in Physical Network Topology Design.

 **Note:**

All instructions in this procedure are executed from the Bootstrap Host.

Prerequisites

1. Procedure [OCCNE Installation of Oracle Linux 7.5 on Bootstrap Host](#) has been completed.
2. The switches are in factory default state.
3. The switches are connected as per [Installation PreFlight Checklist](#). Customer uplinks are not active before outside traffic is necessary.
4. DHCP, XINETD, and TFTP are already installed on the Bootstrap host but are not configured.
5. The Utility USB is available containing the necessary files as per: [Installation PreFlight checklist: Create Utility USB](#).

Limitations/Expectations

All steps are executed from a Keyboard, Video, Mouse (KVM) connection.

References

- <https://github.com/datacenter/nexus9000/blob/master/nx-os/poap/poap.py>

- <https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli/CLIConfigurationGuide/Licensing.html>

Configuration Procedure

Following is the procedure to configure Top of Rack 93180YC-EX Switches:

1. Login to the Bootstrap host as root.

 **Note:**

All instructions in this step are executed from the Bootstrap Host.

2. Insert and mount the Utility USB that contains the configuration and script files. Verify the files are listed in the USB using the `ls /media/usb` command.

 **Note:**

Instructions for mounting the USB can be found in: [Installation of Oracle Linux 7.5 on Bootstrap Server : Install Additional Packages](#). Only steps 2 and 3 need to be followed in that procedure.

3. Create bridge interface: Create bridge interface to connect both management ports and setup the management bridge to support switch initialization.

 **Note:**

<CNE_Management_IP_With_Prefix> is from [Installation PreFlight Checklist : Complete Site Survey Host IP Table](#). Row 1 CNE Management IP Addressess (VLAN 4) column.

<ToRSwitch_CNEManagementNet_VIP> is from [Installation PreFlight Checklist : Complete OA and Switch IP Table](#).

```
$ nmcli con add con-name mgmtBridge type bridge ifname mgmtBridge
$ nmcli con add type bridge-slave ifname eno2 master mgmtBridge
$ nmcli con add type bridge-slave ifname eno3 master mgmtBridge
$ nmcli con mod mgmtBridge ipv4.method manual ipv4.addresses
192.168.2.11/24
$ nmcli con up mgmtBridge

$ nmcli con add type team con-name team0 ifname team0 team.runner
lacp
$ nmcli con add type team-slave con-name team0-slave-1 ifname eno5
master team0
$ nmcli con add type team-slave con-name team0-slave-2 ifname eno6
master team0
$ nmcli con mod team0 ipv4.method manual ipv4.addresses
172.16.3.4/24
$ nmcli con add con-name team0.4 type vlan id 4 dev team0
$ nmcli con mod team0.4 ipv4.method manual
```

```
ipv4.addresses <CNE_Management_IP_Address_With_Prefix> ipv4.gateway  
<ToRswitch_CNEManagementNet_VIP>  
$ nmcli con up team0.4
```

4. Edit the `/etc/xinetd.d/tftp` file to enable TFTP service. Change the **disable** option to **no**, if it is set to **yes**.

```
$ vi /etc/xinetd.d/tftp  
# default: off  
# description: The tftp server serves files using the trivial file  
transfer \  
#          protocol.  The tftp protocol is often used to boot diskless  
\  
#          workstations, download configuration files to network-aware  
printers, \  
#          and to start the installation process for some operating  
systems.  
service tftp  
{  
    socket_type          = dgram  
    protocol             = udp  
    wait                 = yes  
    user                 = root  
    server               = /usr/sbin/in.tftpd  
    server_args          = -s /var/lib/tftpboot  
    disable              = no  
    per_source           = 11  
    cps                  = 100 2  
    flags                = IPv4  
}
```

5. Enable tftp on the Bootstrap host:

```
$ systemctl start tftp  
$ systemctl enable tftp
```

Verify tftp is active and enabled:

```
$ systemctl status tftp  
$ ps -elf | grep tftp
```

6. Copy the `dhcpd.conf` file from the Utility USB in [Installation PreFlight checklist](#) : [Create the dhcpd.conf File](#) to the `/etc/dhcp/` directory.

```
$ cp /media/usb/dhcpd.conf /etc/dhcp/
```

7. Restart and enable dhcpd service.

```
# /bin/systemctl restart dhcpd  
# /bin/systemctl enable dhcpd
```

Use the `systemctl status dhcpd` command to verify active and enabled.

```
# systemctl status dhcpd
```

- Copy the switch configuration and script files from the Utility USB to directory `/var/lib/tftpboot/`.

```
$ cp /media/usb/93180_switchA.cfg /var/lib/tftpboot/.  
$ cp /media/usb/93180_switchB.cfg /var/lib/tftpboot/.  
$ cp /media/usb/poap_nexus_script.py /var/lib/tftpboot/.
```

- Modify POAP script File. Change Username and password credentials used to login to the Bootstrap host.

```
# vi /var/lib/tftpboot/poap_nexus_script.py  
# Host name and user credentials  
options = {  
    "username": "<username>",  
    "password": "<password>",  
    "hostname": "192.168.2.11",  
    "transfer_protocol": "scp",  
    "mode": "serial_number",  
    "target_system_image": "nxos.9.2.3.bin",  
}
```

Note: The version `nxos.9.2.3.bin` is used by default. If different version is to be used, modify the `"target_system_image"` with new version.

- Modify POAP script file `md5sum` by executing the `md5Poap.sh` script from the Utility USB created from [Installation PreFlight checklist](#): Create the `md5Poap` Bash Script.

```
# cd /var/lib/tftpboot/  
# /bin/bash md5Poap.sh
```

- Create the files necessary to configure the ToR switches using the serial number from the switch. The serial number is located on a pullout card on the back of the switch in the left most power supply of the switch.



 **Note:**

The serial number is located on a pullout card on the back of the switch in the left most power supply of the switch. Be careful in interpreting the exact letters. If the switches are preconfigured then you can even verify the serial numbers using 'show license host-id' command.

12. Copy the `/var/lib/tftpboot/93180_switchA.cfg` into a file called `/var/lib/tftpboot/conf.<switchA serial number>` Modify the switch specific values in the `/var/lib/tftpboot/conf.<switchA serial number>` file, including all the values in the curly braces as following code block.

These values are contained at [Installation PreFlight checklist : ToR and Enclosure Switches Variables Table \(Switch Specific\)](#) and [Installation PreFlight Checklist : Complete OA and Switch IP Table](#). Modify these values with the following sed commands, or use an editor such as vi etc.

```
# sed -i 's/{switchname}/<switch_name>/' conf.<switchA serial
number>
# sed -i 's/{admin_password}/<admin_password>/' conf.<switchA
serial number>
# sed -i 's/{user_name}/<user_name>/' conf.<switchA serial number>
# sed -i 's/{user_password}/<user_password>/' conf.<switchA serial
number>
# sed -i 's/{ospf_md5_key}/<ospf_md5_key>/' conf.<switchA serial
number>
# sed -i 's/{OSPF_AREA_ID}/<ospf_area_id>/' conf.<switchA serial
number>
```

```
# sed -i 's/{NTPSERVER1}/<NTP_server_1>/' conf.<switchA serial
number>
# sed -i 's/{NTPSERVER2}/<NTP_server_2>/' conf.<switchA serial
number>
# sed -i 's/{NTPSERVER3}/<NTP_server_3>/' conf.<switchA serial
number>
# sed -i 's/{NTPSERVER4}/<NTP_server_4>/' conf.<switchA serial
number>
# sed -i 's/{NTPSERVER5}/<NTP_server_5>/' conf.<switchA serial
number>
```

Note: If less than 5 ntp servers available, delete the extra ntp server lines such as command:

```
# sed -i 's/{NTPSERVER5}/d' conf.<switchA serial number>
```

Note: different delimiter is used in next two commands due to '/' sign in the variables

```
# sed -i
's#{ALLOW_5G_XSI_LIST_WITH_PREFIX_LEN}#<MetalLB_Signal_Subnet_With_P
refix>#g' conf.<switchA serial number>
# sed -i
's#{CNE_Management_SwA_Address}#<ToRswitchA_CNEManagementNet_IP>#g'
conf.<switchA serial number>
# sed -i
's#{CNE_Management_SwB_Address}#<ToRswitchB_CNEManagementNet_IP>#g'
conf.<switchA serial number>
```

```
# sed -i 's#{CNE_Management_Prefix}#<CNEManagementNet_Prefix>#g'
conf.<switchA serial number>
# sed -i
's#{SQL_replication_SwA_Address}#<ToRswitchA_SQLreplicationNet_IP>#g
' conf.<switchA serial number>
# sed -i
's#{SQL_replication_SwB_Address}#<ToRswitchB_SQLreplicationNet_IP>#g
' conf.<switchA serial number>
# sed -i 's#{SQL_replication_Prefix}#<SQLreplicationNet_Prefix>#g'
conf.<switchA serial number>
# ipcalc -n <ToRswitchA_SQLreplicationNet_IP/
<SQLreplicationNet_Prefix> | awk -F=' ' '{print $2}'
# sed -i 's/{SQL_replication_Subnet}/<output from ipcalc command as
SQL_replication_Subnet>/' conf.<switchA serial number>

# sed -i 's/{CNE_Management_VIP}/
<ToRswitch_CNEManagementNet_VIP>/g' conf.<switchA serial number>
# sed -i 's/{SQL_replication_VIP}/
<ToRswitch_SQLreplicationNet_VIP>/g' conf.<switchA serial number>
# sed -i 's/{OAM_UPLINK_CUSTOMER_ADDRESS}/
<ToRswitchA_oam_uplink_customer_IP>/' conf.<switchA serial number>

# sed -i 's/{OAM_UPLINK_SwA_ADDRESS}/<ToRswitchA_oam_uplink_IP>/g'
conf.<switchA serial number>
# sed -i 's/{SIGNAL_UPLINK_SwA_ADDRESS}/
<ToRswitchA_signaling_uplink_IP>/g' conf.<switchA serial number>
# sed -i 's/{OAM_UPLINK_SwB_ADDRESS}/<ToRswitchB_oam_uplink_IP>/g'
conf.<switchA serial number>
# sed -i 's/{SIGNAL_UPLINK_SwB_ADDRESS}/
<ToRswitchB_signaling_uplink_IP>/g' conf.<switchA serial number>
# ipcalc -n <ToRswitchA_signaling_uplink_IP>/30 | awk -F=' '
'{print $2}'
# sed -i 's/{SIGNAL_UPLINK_SUBNET}/<output from ipcalc command as
signal_uplink_subnet>/' conf.<switchA serial number>

# ipcalc -n <ToRswitchA_SQLreplicationNet_IP> | awk -F=' '
'{print $2}'
# sed -i 's/{MySQL_Replication_SUBNET}/<output from the above
ipcalc command appended with prefix >/' conf.<switchA serial number>
```

Note: The version nxos.9.2.3.bin is used by default and hard-coded in the conf files. If different version is to be used, run the following command:

```
# sed -i 's/nxos.9.2.3.bin/<nxos_version>/' conf.<switchA serial
number>
```

Note: access-list Restrict_Access_ToR

The following line allow one access server to access the switch management and SQL vlan addresses while other accesses are denied. If no need, delete this line. If need more servers, add similar line.

```
# sed -i 's/{Allow_Access_Server}/<Allow_Access_Server>/'
conf.<switchA serial number>
```

13. Copy the `/var/lib/tftpboot/93180_switchB.cfg` into a file called `/var/lib/tftpboot/conf.<switchB serial number>`
 Modify the switch specific values in the `/var/lib/tftpboot/conf.<switchA serial number>` file, including: hostname, username/password, oam_uplink IP address, signaling_uplink IP address, access-list ALLOW_5G_XSI_LIST permit address, prefix-list ALLOW_5G_XSI.

These values are contained at [Installation PreFlight checklist : ToR and Enclosure Switches Variables Table](#) and [Installation PreFlight Checklist : Complete OA and Switch IP Table](#).

```
# sed -i 's/{switchname}/<switch_name>/' conf.<switchB serial
number>
# sed -i 's/{admin_password}/<admin_password>/' conf.<switchB
serial number>
# sed -i 's/{user_name}/<user_name>/' conf.<switchB serial number>
# sed -i 's/{user_password}/<user_password>/' conf.<switchB serial
number>
# sed -i 's/{ospf_md5_key}/<ospf_md5_key>/' conf.<switchB serial
number>
# sed -i 's/{OSPF_AREA_ID}/<ospf_area_id>/' conf.<switchB serial
number>
```

```
# sed -i 's/{NTPSERVER1}/<NTP_server_1>/' conf.<switchB serial
number>
# sed -i 's/{NTPSERVER2}/<NTP_server_2>/' conf.<switchB serial
number>
# sed -i 's/{NTPSERVER3}/<NTP_server_3>/' conf.<switchB serial
number>
# sed -i 's/{NTPSERVER4}/<NTP_server_4>/' conf.<switchB serial
number>
# sed -i 's/{NTPSERVER5}/<NTP_server_5>/' conf.<switchB serial
number>
```

Note: If less than 5 ntp servers available, delete the extra ntp server lines such as command:

```
# sed -i 's/{NTPSERVER5}/d' conf.<switchB serial number>
```

Note: different delimiter is used in next two commands due to '/' sign in in the variables

```
# sed -i
's#{ALLOW_5G_XSI_LIST_WITH_PREFIX_LEN}#<MetalLB_Signal_Subnet_With_P
refix>#g' conf.<switchB serial number>
# sed -i
's#{CNE_Management_SwA_Address}#<ToRswitchA_CNEManagementNet_IP>#g'
conf.<switchB serial number>
# sed -i
's#{CNE_Management_SwB_Address}#<ToRswitchB_CNEManagementNet_IP>#g'
conf.<switchB serial number>
# sed -i 's#{CNE_Management_Prefix}#<CNEManagementNet_Prefix>#g'
conf.<switchB serial number>
# sed -i
's#{SQL_replication_SwA_Address}#<ToRswitchA_SQLreplicationNet_IP>#g
' conf.<switchB serial number>
# sed -i
```

```

's#{SQL_replication_SwB_Address}#<ToRswitchB_SQLreplicationNet_IP>#g
' conf.<switchB serial number>
# sed -i 's#{SQL_replication_Prefix}#<SQLreplicationNet_Prefix>#g'
conf.<switchB serial number>
# ipcalc -n <ToRswitchB_SQLreplicationNet_IP/
<SQLreplicationNet_Prefix> | awk -F=' ' '{print $2}'
# sed -i 's/{SQL_replication_Subnet}/<output from ipcalc command as
SQL_replication_Subnet>/' conf.<switchB serial number>

# sed -i 's/{CNE_Management_VIP}/<ToRswitch_CNEManagementNet_VIP>/'
conf.<switchB serial number>
# sed -i 's/{SQL_replication_VIP}/
<ToRswitch_SQLreplicationNet_VIP>/' conf.<switchB serial number>
# sed -i 's/{OAM_UPLINK_CUSTOMER_ADDRESS}/
<ToRswitchB_oam_uplink_customer_IP>/' conf.<switchB serial number>

# sed -i 's/{OAM_UPLINK_SwA_ADDRESS}/<ToRswitchA_oam_uplink_IP>/g'
conf.<switchB serial number>
# sed -i 's/{SIGNAL_UPLINK_SwA_ADDRESS}/
<ToRswitchA_signaling_uplink_IP>/g' conf.<switchB serial number>
# sed -i 's/{OAM_UPLINK_SwB_ADDRESS}/<ToRswitchB_oam_uplink_IP>/g'
conf.<switchB serial number>
# sed -i 's/{SIGNAL_UPLINK_SwB_ADDRESS}/
<ToRswitchB_signaling_uplink_IP>/g' conf.<switchB serial number>
# ipcalc -n <ToRswitchB_signaling_uplink_IP>/30 | awk -F=' '
'{print $2}'
# sed -i 's/{SIGNAL_UPLINK_SUBNET}/<output from ipcalc command as
signal_uplink_subnet>/' conf.<switchB serial number>

```

Note: The version nxos.9.2.3.bin is used by default and hard-coded in the conf files. If different version is to be used, run the following command:

```
# sed -i 's/nxos.9.2.3.bin/<nxos_version>/' conf.<switchB serial
number>
```

Note: access-list Restrict_Access_ToR

The following line allow one access server to access the switch management and SQL vlan addresses while other accesses are denied. If no need, delete this line. If need more servers, add similar line.

```
# sed -i 's/{Allow_Access_Server}/<Allow_Access_Server>/'
conf.<switchB serial number>
```

14. Generate the md5 checksum for each conf file in `/var/lib/tftpboot` and copy that into a new file called **conf.<switchA/B serial number>.md5**.

```
$ md5sum conf.<switchA serial number> > conf.<switchA serial
number>.md5
$ md5sum conf.<switchB serial number> > conf.<switchB serial
number>.md5
```

15. Verify the `/var/lib/tftpboot` directory has the correct files. Make sure the file permissions are set as given below.

Note: The ToR switches are constantly attempting to find and execute the `poap_nexus_script.py` script which uses `tftp` to load and install the configuration files.

```
# ls -l /var/lib/tftpboot/
total 1305096
-rw-r--r--. 1 root root      7161 Mar 25 15:31 conf.<switchA
serial number>
-rw-r--r--. 1 root root       51 Mar 25 15:31 conf.<switchA
serial number>.md5
-rw-r--r--. 1 root root      7161 Mar 25 15:31 conf.<switchB
serial number>
-rw-r--r--. 1 root root       51 Mar 25 15:31 conf.<switchB
serial number>.md5
-rwxr-xr-x. 1 root root    75856 Mar 25 15:32 poap_nexus_script.py
```

16. Disable firewalld.

```
$ systemctl stop firewalld
$ systemctl disable firewalld
```

To verify:

```
$ systemctl status firewalld
```

Once this is complete, the ToR Switches will attempt to boot from the `tftpboot` files automatically. Eventually the verification steps can be executed below. It may take about 5 minutes for this to complete.

17. Un-mount the Utility USB and remove it: `umount /media/usb`

Verification

Following is the procedure to verify Top of Rack 93180YC-EX Switches

1. After the ToR switches configured, ping the switches from bootstrap server. The switches `mgmt0` interfaces are configured with the IP addresses which are in the `conf` files. **Note:** Wait till the device responds.

```
# ping 192.168.2.1
PING 192.168.2.1 (192.168.2.1) 56(84) bytes of data.
64 bytes from 192.168.2.1: icmp_seq=1 ttl=255 time=0.419 ms
64 bytes from 192.168.2.1: icmp_seq=2 ttl=255 time=0.496 ms
64 bytes from 192.168.2.1: icmp_seq=3 ttl=255 time=0.573 ms
64 bytes from 192.168.2.1: icmp_seq=4 ttl=255 time=0.535 ms
^C
--- 192.168.2.1 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3000ms
rtt min/avg/max/mdev = 0.419/0.505/0.573/0.063 ms
# ping 192.168.2.2
PING 192.168.2.2 (192.168.2.2) 56(84) bytes of data.
64 bytes from 192.168.2.2: icmp_seq=1 ttl=255 time=0.572 ms
64 bytes from 192.168.2.2: icmp_seq=2 ttl=255 time=0.582 ms
64 bytes from 192.168.2.2: icmp_seq=3 ttl=255 time=0.466 ms
64 bytes from 192.168.2.2: icmp_seq=4 ttl=255 time=0.554 ms
```

```
^C
--- 192.168.2.2 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3001ms
rtt min/avg/max/mdev = 0.466/0.543/0.582/0.051 ms
```

2. Attempt to ssh to the switches with the username/password provided in the conf files.

```
# ssh plat@192.168.2.1
The authenticity of host '192.168.2.1 (192.168.2.1)' can't be
established.
RSA key fingerprint is
SHA256:jEPSMHRNg9vejiLcEvw5qprjgt+4ua9jucUBhktH520.
RSA key fingerprint is
MD5:02:66:3a:c6:81:65:20:2c:6e:cb:08:35:06:c6:72:ac.
Are you sure you want to continue connecting (yes/no)? yes
Warning: Permanently added '192.168.2.1' (RSA) to the list of known
hosts.
User Access Verification
Password:
```

```
Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (C) 2002-2019, Cisco and/or its affiliates.
All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under their
own
licenses, such as open source. This software is provided "as is,"
and unless
otherwise stated, there is no warranty, express or implied,
including but not
limited to warranties of merchantability and fitness for a
particular purpose.
Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or
GNU General Public License (GPL) version 3.0 or the GNU
Lesser General Public License (LGPL) Version 2.1 or
Lesser General Public License (LGPL) Version 2.0.
A copy of each such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://opensource.org/licenses/gpl-3.0.html and
http://www.opensource.org/licenses/lgpl-2.1.php and
http://www.gnu.org/licenses/old-licenses/library.txt.
#
```

3. Verify the running-config has all expected configurations in the conf file using the **show running-config** command.

```
# show running-config
!Command: show running-config
!Running configuration last done at: Mon Apr 8 17:39:38 2019
!Time: Mon Apr 8 18:30:17 2019
version 9.2(3) Bios:version 07.64
hostname 12006-93108A
```

```
vdc 12006-93108A id 1
  limit-resource vlan minimum 16 maximum 4094
  limit-resource vrf minimum 2 maximum 4096
  limit-resource port-channel minimum 0 maximum 511
  limit-resource u4route-mem minimum 248 maximum 248
  limit-resource u6route-mem minimum 96 maximum 96
  limit-resource m4route-mem minimum 58 maximum 58
  limit-resource m6route-mem minimum 8 maximum 8
feature scp-server
feature sftp-server
cfs eth distribute
feature ospf
feature bgp
feature interface-vlan
feature lacp
feature vpc
feature bfd
feature vrrpv3
....
....
```

4. Verify license on the switches. In case some of the above features are missing, verify license on the switches and at least NXOS_ADVANTAGE level license is "In use". If license not installed or too low level, contact vendor for correct license key file, following Licensing document mentioned in reference section to install license key. Then run "write erase" and "reload" to set back to factory default. The switches will go to POAP configuration again.

```
# show license
```

Example output:

```
# show license
MDS20190215085542979.lic:
SERVER this_host ANY
VENDOR cisco
INCREMENT NXOS_ADVANTAGE_XF cisco 1.0 permanent uncounted \
  VENDOR_STRING=<LIC_SOURCE>MDS_SWIFT</LIC_SOURCE><SKU>NXOS-
AD-XF</SKU> \
  HOSTID=VDH=FDO22412J2F \
  NOTICE="<LicFileID>20190215085542979</
LicFileID><LicLineID>1</LicLineID> \
  <PAK></PAK>" SIGN=8CC8807E6918
```

```
# show license usage
```

Example output:

```
# show license usage
Feature                               Ins Lic  Status Expiry Date Comments
                                      Count
-----
...
NXOS_ADVANTAGE_M4                      No  -   Unused          -
NXOS_ADVANTAGE_XF                      Yes -   In use never     -
NXOS_ESSENTIALS_GF                     No  -   Unused          -
```

```
...
#
```

5. Verify the RMS1 can ping the CNE_Management VIP.

```
# ping <ToRSwitch_CNEManagementNet_VIP>
PING <ToRSwitch_CNEManagementNet_VIP>
(<ToRSwitch_CNEManagementNet_VIP>) 56(84) bytes of data.
64 bytes from <ToRSwitch_CNEManagementNet_VIP>: icmp_seq=2 ttl=255
time=1.15 ms
64 bytes from <ToRSwitch_CNEManagementNet_VIP>: icmp_seq=3 ttl=255
time=1.11 ms
64 bytes from <ToRSwitch_CNEManagementNet_VIP>: icmp_seq=4 ttl=255
time=1.23 ms
^C
--- 10.75.207.129 ping statistics ---
4 packets transmitted, 3 received, 25% packet loss, time 3019ms
rtt min/avg/max/mdev = 1.115/1.168/1.237/0.051 ms
```

6. Enable customer uplink.
7. Verify the RMS1 can be accessed from laptop. Use application such as putty etc to ssh to RMS1.

```
$ ssh root@<CNE_Management_IP_Address>
Using username "root".
root@<CNE_Management_IP_Address>'s password:<root password>
Last login: Mon May 6 10:02:01 2019 from 10.75.9.171
[root@RMS1 ~]#
```

SNMP Trap Configuration

The following procedure explains the steps to configure SNMP Trap:

1. SNMPv2c Configuration.

When SNMPv2c configuration is needed, ssh to the two switches, run the following commands:

These values <SNMP_Trapping_Receiver_Address> and <SNMP_Community_String> are from [Installation Preflight Checklist](#).

```
[root@RMS1 ~]# ssh <user_name>@<ToRswitchA_CNEManagementNet_IP>
# configure terminal
(config)# snmp-server host <SNMP_Trapping_Receiver_Address> traps
version 2c <SNMP_Community_String>
(config)# snmp-server host <SNMP_Trapping_Receiver_Address> use-vrf
default
(config)# snmp-server host <SNMP_Trapping_Receiver_Address> source-
interface Ethernet1/51
(config)# snmp-server enable traps
(config)# snmp-server community <SNMP_Community_String> group
network-admin
```


2. Restrict direct access to ToR switches. In order to restrict direct access to ToR switches, IP access list is created and applied on the uplink interfaces, the following commands are needed on ToR switches:

```
[root@RMS1 ~]# ssh <user_name>@<ToRswitchA_CNEManagementNet_IP>
# configure terminal
(config)#
ip access-list Restrict_Access_ToR
    permit ip {Allow_Access_Server}/32 any
    permit ip {NTPSERVER1}/32 {OAM_UPLINK_SwA_ADDRESS}/32
    permit ip {NTPSERVER2}/32 {OAM_UPLINK_SwA_ADDRESS}/32
    permit ip {NTPSERVER3}/32 {OAM_UPLINK_SwA_ADDRESS}/32
    permit ip {NTPSERVER4}/32 {OAM_UPLINK_SwA_ADDRESS}/32
    permit ip {NTPSERVER5}/32 {OAM_UPLINK_SwA_ADDRESS}/32
    deny ip any {CNE_Management_VIP}/32
    deny ip any {CNE_Management_SwA_Address}/32
    deny ip any {CNE_Management_SwB_Address}/32
    deny ip any {SQL_replication_VIP}/32
    deny ip any {SQL_replication_SwA_Address}/32
    deny ip any {SQL_replication_SwB_Address}/32
    deny ip any {OAM_UPLINK_SwA_ADDRESS}/32
    deny ip any {OAM_UPLINK_SwB_ADDRESS}/32
    deny ip any {SIGNAL_UPLINK_SwA_ADDRESS}/32
    deny ip any {SIGNAL_UPLINK_SwB_ADDRESS}/32
    permit ip any any

interface Ethernet1/51
    ip access-group Restrict_Access_ToR in

interface Ethernet1/52
    ip access-group Restrict_Access_ToR in
```

3. Traffic egress out of cluster, including snmptrap traffic to SNMP trap receiver, and traffic goes to signal server:

```
[root@RMS1 ~]# ssh <user_name>@<ToRswitchA_CNEManagementNet_IP>
# configure terminal
(config)#
feature nat
ip access-list host-snmptap
    10 permit udp 172.16.3.0/24 <snmp trap receiver>/32 eq snmptap log

ip access-list host-sigserver
    10 permit ip 172.16.3.0/24 <signal server>/32

ip nat pool sig-pool 10.75.207.211 10.75.207.222 prefix-length 27
ip nat inside source list host-sigserver pool sig-pool overload add-
route
ip nat inside source list host-snmptap interface Ethernet1/51
overload

interface Vlan3
    ip nat inside

interface Ethernet1/51
```

```
ip nat outside

interface Ethernet1/52
ip nat outside
```

Run the same commands on ToR switchB

Configure Addresses for RMS iLOs, OA, EBIPA

Introduction

This procedure is used to configure RMS iLO addresses and add a new user account for each RMS other than the Bootstrap Host. When the RMSs are shipped and out of box after hardware installation and powerup, the RMSs are in a factory default state with the iLO in DHCP mode waiting for DHCP service. DHCP is used to configure the ToR switches, OAs, Enclosure switches, and blade server iLOs, so DHCP can be used to configure RMS iLOs as well.

Prerequisites

Procedure [OCCNE Configure Top of Rack 93180YC-EX Switches](#) has been completed.

Limitations/Expectations

All steps are executed from the ssh session of the Bootstrap server.

References

[HPE BladeSystem Onboard Administrator User Guide](#)

Procedure

Following is the procedure to configure Addresses for RMS iLOs, OA, EBIPA:

1. Setup team0.2 interface:

```
$ nmcli con add con-name team0.2 type vlan id 2 dev team0
$ nmcli con mod team0.2 ipv4.method manual ipv4.addresses
192.168.20.11/24
$ nmcli con up team0.2
```

2. Subnet and conf file address.

The `/etc/dhcp/dhcpd.conf` file should already have been configured in [OCCNE Configure Top of Rack 93180YC-EX Switches](#) procedure.

[Configure Top of Rack 93180YC-EX Switches](#) and dhcp started/enabled on the bootstrap server. The second subnet 192.168.20.0 is used to assign addresses for OA and RMS iLOs. The "next-server 192.168.20.11" option is same as the server team0.2 IP address.

3. Display the dhcpd leases file at `/var/lib/dhcpd/dhcpd.leases`. The DHCPD lease file will display the DHCP addresses for all RMS iLOs, Enclosure OAs.

```
# cat /var/lib/dhcpd/dhcpd.leases
# The format of this file is documented in the dhcpd.leases(5)
```

```
manual page.
# This lease file was written by isc-dhcp-4.2.5
lease 192.168.20.101 {
    starts 4 2019/03/28 22:05:26;
    ends 4 2019/03/28 22:07:26;
    tstp 4 2019/03/28 22:07:26;
    cltt 4 2019/03/28 22:05:26;
    binding state free;
    hardware ethernet 48:df:37:7a:41:60;
}
lease 192.168.20.103 {
    starts 4 2019/03/28 22:05:28;
    ends 4 2019/03/28 22:07:28;
    tstp 4 2019/03/28 22:07:28;
    cltt 4 2019/03/28 22:05:28;
    binding state free;
    hardware ethernet 48:df:37:7a:2f:70;
}
lease 192.168.20.102 {
    starts 4 2019/03/28 22:05:16;
    ends 4 2019/03/28 23:03:29;
    tstp 4 2019/03/28 23:03:29;
    cltt 4 2019/03/28 22:05:16;
    binding state free;
    hardware ethernet 48:df:37:7a:40:40;
}
lease 192.168.20.106 {
    starts 5 2019/03/29 11:14:04;
    ends 5 2019/03/29 14:14:04;
    tstp 5 2019/03/29 14:14:04;
    cltt 5 2019/03/29 11:14:04;
    binding state free;
    hardware ethernet b8:83:03:47:5f:14;
    uid "\000\270\203\003G_\024\000\000\000";
}
lease 192.168.20.105 {
    starts 5 2019/03/29 12:56:23;
    ends 5 2019/03/29 15:56:23;
    tstp 5 2019/03/29 15:56:23;
    cltt 5 2019/03/29 12:56:23;
    binding state free;
    hardware ethernet b8:83:03:47:5e:54;
    uid "\000\270\203\003G^T\000\000\000";
}
lease 192.168.20.104 {
    starts 5 2019/03/29 13:08:21;
    ends 5 2019/03/29 16:08:21;
    tstp 5 2019/03/29 16:08:21;
    cltt 5 2019/03/29 13:08:21;
    binding state free;
    hardware ethernet b8:83:03:47:64:9c;
    uid "\000\270\203\003Gd\234\000\000\000";
}
lease 192.168.20.108 {
    starts 5 2019/03/29 09:57:02;
```

```
ends 5 2019/03/29 21:57:02;
tstp 5 2019/03/29 21:57:02;
cltt 5 2019/03/29 09:57:02;
binding state active;
next binding state free;
rewind binding state free;
hardware ethernet fc:15:b4:1a:ea:05;
uid "\001\374\025\264\032\352\005";
client-hostname "OA-FC15B41AEA05";
}
lease 192.168.20.107 {
starts 5 2019/03/29 12:02:50;
ends 6 2019/03/30 00:02:50;
tstp 6 2019/03/30 00:02:50;
cltt 5 2019/03/29 12:02:50;
binding state active;
next binding state free;
rewind binding state free;
hardware ethernet 9c:b6:54:80:d7:d7;
uid "\001\234\266T\200\327\327";
client-hostname "SA-9CB65480D7D7";
}
server-duid "\000\001\000\001$\#\364\344\270\203\003Gim";
lease 192.168.20.107 {
starts 5 2019/03/29 18:09:47;
ends 6 2019/03/30 06:09:47;
cltt 5 2019/03/29 18:09:47;
binding state active;
next binding state free;
rewind binding state free;
hardware ethernet 9c:b6:54:80:d7:d7;
uid "\001\234\266T\200\327\327";
client-hostname "SA-9CB65480D7D7";
}
lease 192.168.20.108 {
starts 5 2019/03/29 18:09:54;
ends 6 2019/03/30 06:09:54;
cltt 5 2019/03/29 18:09:54;
binding state active;
next binding state free;
rewind binding state free;
hardware ethernet fc:15:b4:1a:ea:05;
uid "\001\374\025\264\032\352\005";
client-hostname "OA-FC15B41AEA05";
}
lease 192.168.20.106 {
starts 5 2019/03/29 18:10:04;
ends 5 2019/03/29 21:10:04;
cltt 5 2019/03/29 18:10:04;
binding state active;
next binding state free;
rewind binding state free;
hardware ethernet b8:83:03:47:5f:14;
uid "\000\270\203\003G_\024\000\000\000";
client-hostname "ILO2M2909004B";
```

```

}
lease 192.168.20.104 {
  starts 5 2019/03/29 18:10:35;
  ends 5 2019/03/29 21:10:35;
  cltt 5 2019/03/29 18:10:35;
  binding state active;
  next binding state free;
  rewind binding state free;
  hardware ethernet b8:83:03:47:64:9c;
  uid "\000\270\203\003Gd\234\000\000\000";
  client-hostname "ILO2M2909004F";
}
lease 192.168.20.105 {
  starts 5 2019/03/29 18:10:40;
  ends 5 2019/03/29 21:10:40;
  cltt 5 2019/03/29 18:10:40;
  binding state active;
  next binding state free;
  rewind binding state free;
  hardware ethernet b8:83:03:47:5e:54;
  uid "\000\270\203\003G^T\000\000\000";
  client-hostname "ILO2M29090048";
}

```

4. Access RMS iLO from the DHCP address with default Administrator password. From the above `dhcpd.leases` file, find the IP address for the iLO name, the default username is Administrator, the password is on the label which can be pulled out from front of server.

 **Note:**

The DNS Name on the pull-out label. The DNS Name on the pull-out label should be used to match the physical machine with the iLO IP since the same default DNS Name from the pull-out label is displayed upon logging in to the iLO command line interface, as shown in the example below.

```

# ssh Administrator@192.168.20.104
Administrator@192.168.20.104's password:
User:Administrator logged-in to
ILO2M2909004F.labs.nc.tekelec.com(192.168.20.104 /
FE80::BA83:3FF:FE47:649C)
iLO Standard 1.37 at Oct 25 2018
Server Name:
Server Power: On

```

5. Create RMS iLO new user. Create new user with customized username and password.

```

</>hpiLO-> create /map1/accounts1 username=root password=TklcRoot
group=admin,config,oemHP_rc,oemHP_power,oemHP_vm

status=0

```

```
status_tag=COMMAND COMPLETED
Tue Apr  2 20:08:30 2019
User added successfully.
```

6. Disable the DHCP before able to setup static IP. Setup static failed before DHCP is disabled.

```
</>hpiLO-> set /map1/dhccpendpt1 EnabledState=NO
status=0
status_tag=COMMAND COMPLETED
Tue Apr  2 20:04:53 2019
Network settings change applied.
Settings change applied, iLO 5 will now be reset.
Logged Out: It may take several minutes before you can log back in.
CLI session stopped
packet_write_wait: Connection to 192.168.20.104 port 22: Broken pipe
```

7. Setup RMS iLO static IP address. After a while after previous step, can login back with the same address (which is static IP now) and new username/password. If don't want to use the same address, go to next step to change the IP address.

```
# ssh <new username>@192.168.20.104
<new username>@192.168.20.104's password: <new password>
User: logged-in to
ILO2M2909004F.labs.nc.tekelec.com(192.168.20.104 /
FE80::BA83:3FF:FE47:649C)
iLO Standard 1.37 at Oct 25 2018
Server Name:
Server Power: On
```

```
</>hpiLO-> set /map1/enetport1/lanendpt1/ipendpt1
IPv4Address=192.168.20.122 SubnetMask=255.255.255.0
```

```
status=0
status_tag=COMMAND COMPLETED
Tue Apr  2 20:22:23 2019
```

```
Network settings change applied.
Settings change applied, iLO 5 will now be reset.
Logged Out: It may take several minutes before you can log
back in.
```

```
CLI session stopped
```

```
packet_write_wait: Connection to 192.168.20.104 port 22:
Broken pipe
#
```

8. Set EBIPA addresses for InterConnect Bays (Enclosure Switches). From bootstrap server, login to OA, set EBIPA addressed for the two enclosure switches. The addresses have to be in the subnet with server team0.2 address in order for TFTP to work.

Set address for each enclosure switch, note the last number 1 or 2 is the interconnect bay number.

```
OA-FC15B41AEA05> set ebipa interconnect 192.168.20.133
255.255.255.0 1
Entering anything other than 'YES' will result in the command not
executing.
It may take each interconnect several minutes to acquire the new
settings.
Are you sure you want to change the IP address for the specified
interconnect bays? yes
Successfully set 255.255.255.0 as the netmask for interconnect bays.
Successfully set interconnect bay # 1 to IP address 192.168.20.133
For the IP addresses to be assigned EBIPA must be enabled.
```

```
OA-FC15B41AEA05> set ebipa interconnect 192.168.20.134
255.255.255.0 2
Entering anything other than 'YES' will result in the command not
executing.
It may take each interconnect several minutes to acquire the new
settings.
Are you sure you want to change the IP address for the specified
interconnect bays? yes
Successfully set 255.255.255.0 as the netmask for interconnect bays.
Successfully set interconnect bay # 2 to IP address 192.168.20.134
For the IP addresses to be assigned EBIPA must be enabled.
```

9. Set EBIPA addresses for Blade Servers. Set EBIPA addressed for all the blade servers. The addresses are in the same subnet with first server team0.2 address and enclosure switches.

```
OA-FC15B41AEA05> set ebipa server 192.168.20.141 255.255.255.0 1-16
Entering anything other than 'YES' will result in the command not
executing.
Changing the IP address for device (iLO) bays that are enabled
causes the iLOs in those bays to be reset.
Are you sure you want to change the IP address for the specified
device (iLO) bays? YES
Successfully set 255.255.255.0 as the netmask for device (iLO) bays.
Successfully set device (iLO) bay # 1 to IP address 192.168.20.141
Successfully set device (iLO) bay # 2 to IP address 192.168.20.142
Successfully set device (iLO) bay # 3 to IP address 192.168.20.143
Successfully set device (iLO) bay # 4 to IP address 192.168.20.144
Successfully set device (iLO) bay # 5 to IP address 192.168.20.145
Successfully set device (iLO) bay # 6 to IP address 192.168.20.146
Successfully set device (iLO) bay # 7 to IP address 192.168.20.147
Successfully set device (iLO) bay # 8 to IP address 192.168.20.148
Successfully set device (iLO) bay # 9 to IP address 192.168.20.149
Successfully set device (iLO) bay #10 to IP address 192.168.20.150
Successfully set device (iLO) bay #11 to IP address 192.168.20.151
Successfully set device (iLO) bay #12 to IP address 192.168.20.152
Successfully set device (iLO) bay #13 to IP address 192.168.20.153
Successfully set device (iLO) bay #14 to IP address 192.168.20.154
Successfully set device (iLO) bay #15 to IP address 192.168.20.155
Successfully set device (iLO) bay #16 to IP address 192.168.20.156
```

For the IP addresses to be assigned EBIPA must be enabled.
OA-FC15B41AEA05>

10. Add New User for OA.

Create new user, set access level as ADMINISTRATOR, and assign access to all blades, all enclosure switches and OAs. After that, the username and password can be used to access OAs.

```
OA-FC15B41AEA05> ADD USER <username>
New Password: *****
Confirm      : *****
User "<username>" created.
You may set user privileges with the 'SET USER ACCESS' and 'ASSIGN'
commands.
```

```
OA-FC15B41AEA05> set user access <username> ADMINISTRATOR
```

```
"<username>" has been given administrator level privileges.
```

```
OA-FC15B41AEA05> ASSIGN SERVER ALL <username>
```

```
<username> has been granted access to the valid requested bay(s)
```

```
OA-FC15B41AEA05> ASSIGN INTERCONNECT ALL <username>
```

```
<username> has been granted access to the valid requested bay(s)
```

```
OA-FC15B41AEA05> ASSIGN OA <username>
```

```
<username> has been granted access to the OA.
```

11. From OA, go to each blade with "connect server <bay number>", add New User for each blade.

```
OA-FC15B41AEA05> connect server 4
```

```
Connecting to bay 4 ...
```

```
User:OAtmp-root-5CBF2E61 logged-in to ILO2M290605KP.
```

```
(192.168.20.144 / FE80::AF1:EAFF:FE89:460)
```

```
iLO Standard Blade Edition 1.37 at Oct 25 2018
```

```
Server Name:
```

```
Server Power: On
```

```
</>hpiLO->
```

```
</>hpiLO-> create /map1/accounts1 username=root password=TKlcRoot
group=admin,config,oemHPE_rc,oemHPE_power,oemHPE_vm
```

```
status=2
```

```
status_tag=COMMAND PROCESSING FAILED
```

```
error_tag=COMMAND SYNTAX ERROR
```

```
Tue Apr 23 16:18:58 2019
```

```
User added successfully.
```


12. Change to static IP on OA. In order not reply on DHCP and make the OA address stable, change to static IP.

 **Note:**

After the following change, on the active OA (could be the bay1 OA or bay2 OA), the OA session will be stuck due to the address change, make another server session ready to ssh with the new IP address and new root user. The change on the standby OA will not stuck the OA session.

```
OA-FC15B41AEA05> SET IPCONFIG STATIC 1 192.168.20.131 255.255.255.0
Static IP settings successfully updated.
These setting changes will take effect immediately.
```

```
OA-FC15B41AEA05> SET IPCONFIG STATIC 2 192.168.20.132 255.255.255.0
Static IP settings successfully updated.
These setting changes will take effect immediately.
OA-FC15B41AEA05>
```

Configure Legacy BIOS on Remaining Hosts

These procedures define the steps necessary to configure additional Legacy BIOS for all hosts in OCCNE. This includes steps that cannot be performed from the HP iLO 5 CLI prompt such as RAID configuration, changing the boot mode, and setting the primary and secondary boot devices.

 **Note:**

The procedures in this document apply to the HP iLO console accessed via KVM. Each procedure is executed in the order listed.

Prerequisites

Procedure [OCCNE Configure Addresses for RMS iLOs, OA, EBIPA](#) is complete.

Limitations and Expectations

1. Applies to HP iLO 5 only.
2. Should the System Utility indicate (or defaults to) UEFI booting, then the user must go through the steps to reset booting back to the Legacy BIOS mode by following step: Change over from UEFI Booting Mode to Legacy BIOS Booting Mode.
3. The procedures listed here apply to both Gen10 DL380 RMSs and Gen10 BL460c Blades in a C7000 enclosure.
4. Access to the enclosure blades in these procedures is via the Bootstrap host using SSH on the KVM. This is possible because the prerequisites are complete. If the prerequisites are not completed before executing this procedure, the enclosure blades are only accessible via the KVM connected directly to the active OA. In this

case the mouse is not usable and screen manipulations are performed using the keyboard ESC and directional keys.

5. This procedure does NOT apply to the Bootstrap Host.

References

1. [HPE iLO 5 User Guide](#)
2. [UEFI System Utilities User Guide for HPE ProLiant Gen10 Servers and HPE Synergy](#)
3. [UEFI Workload-based Performance and Tuning Guide for HPE ProLiant Gen10 Servers and HPE Synergy](#)
4. [HPE BladeSystem Onboard Administrator User Guide](#)
5. [OCCNE Inventory File Preparation](#)

Procedure

Following is the procedure to configure the Legacy BIOS on Remaining Hosts:

1. Expose the System Configuration Utility on a RMS Host on the KVM. This procedure does not provide instructions on how to connect the KVM as this may be different on each installation.
 - a. Once the remote console has been exposed, the system must be reset by manually pressing the power button on the front of the RMS host to force it through the restart process. When the initial window is displayed, hit the F9 key repeatedly. Once the F9 is highlighted at the lower left corner of the remote console, it should eventually bring up the main System Utility.
 - b. The System Utilities screen is exposed in the remote console.
2. Expose the System Utility for an Enclosure Blade.
 - a. The blades are maintained via the OAs in the enclosure. Because each blade iLO has already been assigned an IP address from the prerequisites, the blades can each be reached using SSH from the Bootstrap host login shell on the KVM.
 - i. SSH to the blade using the iLO IP address and the root user and password. This brings up the HP iLO prompt.

```
$ ssh root@<blade_ilo_ip_address>
Using username "root".
Last login: Fri Apr 19 12:24:56 2019 from 10.39.204.17
[root@localhost ~]# ssh root@192.168.20.141
root@192.168.20.141's password:
User:root logged-in to ILO2M290605KM.(192.168.20.141 /
FE80::AF1:EAF:FE89:35E)
iLO Standard Blade Edition 1.37 at Oct 25 2018
Server Name:
Server Power: On

</>hpiLO->
```

- ii. Use VSP to connect to the blade remote console.

```
</>hpiLO->vsp
```

- iii. Power cycle the blade to bring up the System Utility for that blade.

 **Note:**

The System Utility is a text based version of that exposed on the RMS via the KVM. The user must use the directional (arrow) keys to manipulate between selections, **ENTER** key to select, and **ESC** to go back from the current selection.

- iv. Access the System Utility by hitting **ESC 9**.

- b. Enabling Virtualization

This procedure provides the steps required to enable virtualization on a given Bare Metal Server. Virtualization can be configured using the default settings or via the default Workload Profiles.

Verifying Default Settings

- i. Expose the **System Utility** by following step 1 or 2 depending on the hardware being configured.
 - ii. Select **System Configuration**
 - iii. Select **BIOS/Platform Configuration (RBSU)**
 - iv. Select **Virtualization Options**
This view displays the settings for the **Intel(R) Virtualization Technology (IntelVT)**, **Intel(R) VT-d**, and **SR-IOV** options (Enabled or Disabled). The default values for each option is **Enabled**.
 - v. Select **F10** if it is desired to save and stay in the utility or select the **F12** if it is desired to save and exit to continue the current boot process.
3. Change over from UEFI Booting Mode to Legacy BIOS Booting Mode:
 - a. Expose the **System Utility** by following step 1 or 2 depending on the hardware being configured.
 - b. Select **System Configuration**
 - c. Select **BIOS/Platform Configuration (RBSU)**
 - d. Select **Boot Options**.
This menu defines the boot mode.

If the **Boot Mode** is set to **UEFI Mode** then continue this procedure. Otherwise there is no need to make any of the changes below.
 - e. Select **Boot Mode**
This generates a warning indicating the following:

Boot Mode changes require a system reboot in order to take effect. Changing the Boot Mode can impact the ability of the server to boot the installed operating system. An operating system is installed in the same mode as the platform during the installation. If the Boot Mode does not match the operating

system installation, the system cannot boot. The following features require that the server be configured for UEFI Mode: Secure Boot, IPv6 PXE Boot, Boot > 2.2 TB Disks in AHCI SATA Mode, and Smart Array SW RAID.

Hit the ENTER key and two selections appear: UEFI Mode(highlighted) and Legacy BIOS Mode

- f. Use the down arrow key to select **Legacy BIOS Mode** and hit the ENTER. The screen indicates: A reboot is required for the Boot Mode changes.
- g. Hit **F12**. This displays the following: Changes are pending. Do you want to save changes? Press 'Y' to save and exit, 'N' to discard and stay, or 'ESC' to cancel.
- h. Hit the **y** key and an additional warning appears indicating: System configuration changed. A system reboot is required. Press ENTER to reboot the system.
- i. i. Hit **ENTER** to force a reboot.

 **Note:**

The boot must go into the process of actually trying to boot from the boot devices using the boot order (not just go back through initialization and access the System Utility again). The boot should fail and the System Utility can be accessed again to continue any further changes needed.

- ii. After the reboot, hit the ESC 9key sequence to re-enter the System Utility. Selecting System Configuration->BIOS/Platform Configuration (RBSU)->Boot Options. Verify the Boot Mode is set to Legacy Boot Mode UEFI Optimized Boot is set to Disabled
 - j. Select **F10** if it is desired to save and stay in the utility or select the **F12** if it is desired to save and exit to complete the current boot process.
4. Force PXE to boot from the first Embedded FlexibleLOM HPE Ethernet 10Gb 2-port Adapter.
 - a. Expose the **System Utility** by following step 1 or 2 depending on the hardware being configured.
 - b. Select **System Configuration**.
 - c. Select **BIOS/Platform Configuration (RBSU)** .
 - d. Select **Boot Options**. This menu defines the boot mode.
 - e. Confirm the following settings: Boot Mode Legacy BIOS Mode UEFI Optimized Boot, and Boot Order Policy Retry Boot Order Indefinitely(this means it keeps trying to boot without ever going to disk). If not in Legacy BIOS Mode, follow procedure Change over from UEFI Booting Mode to Legacy BIOS Booting Mode.
 - f. Select **Legacy BIOS Boot Order** In the default view, the 10Gb Embedded FlexibleLOM 1 Port 1 is at the bottom of the list.

- g. Move the 10 Gb Embedded FlexibleLOM 1 Port 1 entry up above the 1Gb Embedded LOM 1 Port 1 entry. To move an entry press the '+' key to move an entry higher in the boot list and the '-' key to move an entry lower in the boot list. Use the arrow keys to navigate through the Boot Order list.
- h. Select **F10** if it is desired to save and stay in the utility or select the **F12** if it is desired to save and exit to continue the current boot process.

5. Enabling Virtualization:

This step provides the steps required to enable virtualization on a given Bare Metal Server. Virtualization can be configured using the default settings or via the Workload Profiles.

Verifying Default Settings

- a. Expose the **System Utility** by following step 1 or 2 depending on the hardware being configured.
 - b. Select **System Configuration**
 - c. Select **BIOS/Platform Configuration (RBSU)**
 - d. Select **Virtualization Options**
This view displays the settings for the **Intel(R) Virtualization Technology (IntelVT)**, **Intel(R) VT-d**, and **SR-IOV** options (Enabled or Disabled). The default values for each option is **Enabled**.
 - e. Select **F10** if it is desired to save and stay in the utility or select the **F12** if it is desired to save and exit to continue the current boot process.
6. Disable RAID Configurations:
- OCCNE does not currently support any RAID configuration. Follow this step to disable RAID settings if the default settings of the System Utility include any RAID configuration(s).

Note: There may be more than one RAID Array set up. This procedure should be repeated for any RAID configuration.

- a. Expose the **System Utility** by following step 1 or 2 depending on the hardware being configured.
- b. Select **System Configuration**.
- c. Select **Embedded RAID 1: HPE Smart Array P408i-a SR Gen 10**.
- d. Select **Array Configuration**.
- e. Select **Manage Arrays**.
- f. Select **Array A (or any designated Array Configuration if there are more than one)**.
- g. Select **Delete Array**. A warning is displayed indicating the following:

Deletes an Array. All the data on the logical drives that are part of deleted array will be lost. Also if the deleted array is the only one on the controller, the controller settings will be erased and its default configuration is restored.

- h. Hit **ENTER**, the changes are submitted and Delete Array Successful is displayed.
- i. Hit **ENTER** to go back to the main menu for the HPE Smart Array.

- j. Select **F10** if it is desired to save and stay in the utility or select the **F12** if it is desired to save and exit to continue the current boot process.
7. Enable the Primary and Secondary Boot Devices:
This steps provide necessary to configure the primary and secondary bootable devices for a Gen10 Server.

 **Note:**

There can be multiple configurations of hardware drives on the server that include both Hard Drives (HDD) and Solid State Hard Drives (SSD). SSDs are indicated by SATA-SSD ATA in the drive description. The commands below include two HDDs and two SSDs. The SSDs are not to be selected for this configuration. The actual selections may be different based on the hardware being updated.

- a. Expose the **System Utility** by following step 1 or 2 depending on the hardware being configured.
- b. Select **System Configuration**.
- c. Select **Embedded RAID 1 : HPE Smart Array P408i-a SR Gen 10**.
- d. Select **Set Bootable Device(s) for Legacy Boot Mode**.
If the boot devices are not set then **Not Set** is displayed for the primary and secondary devices.
- e. Examine the list of available hardware drives. If one or more HDDs are available, continue with this step.

 **Note:**

A single drive can be set as both the primary and secondary boot device but that is not part of this configuration.

- f. Select **Bootable Physical Drive**
- g. Select **Port 1| Box:3 Bay:1 Size:1.8 TB SAS HP EG00100JWJNR**. **Note:**
This example includes two HDDs and two SSDs. The actual configuration may be different.
- h. Select **Set as Primary Bootable Device**.
- i. Hit **ENTER**.

 **Note:**

There is no need to set the secondary boot device. Leave it as **Not Set**.

- j. Hit the **ESC** key to back out to the **System Utilities** menu.
- k. Select **F10** if it is desired to save and stay in the utility or select the **F12** if it is desired to save and exit to continue the current boot process.

Configure Enclosure Switches

Introduction

This procedure is used to configure the 6127XLG enclosure switches.

Prerequisites

- Procedure [Configure Top of Rack 93180YC-EX Switches](#) has been completed.
- Procedure [Configure Addresses for RMS iLOs, OA, EBIPA](#) has been completed.
- The Utility USB is available containing the necessary files as per: [Installation PreFlight checklist: Create Utility USB](#).

Limitations/Expectations

All steps are executed from a Keyboard, Video, Mouse (KVM) connection.

References

1. <https://support.hpe.com/hpsc/doc/public/display?docId=c04763537>

Procedure

Following is the procedure to configure enclosure switches:

1. Copy the 6127XLG configuration file from the Utility USB (See [Installation PreFlight checklist: Create the OA 6127XLG Switch Configuration File](#)) to the `/var/lib/tftpboot` directory on the Installer Bootstrap Host and verify it exists and the permissions.

```
$ cp /media/usb/6127xlg_irf.cfg /var/lib/tftpboot/6127xlg_irf.cfg
```

```
$ ls -l /var/lib/tftpboot/
```

```
total 1305096
```

```
-rw-r--r--. 1 root root      311 Mar 25 08:41 6127xlg_irf.cfg
```

2. Modify the switch specific values in the `/var/lib/tftpboot/6127xlg_irf.cfg` file. These values are contained at [Installation PreFlight checklist : Create the OA 6127XLG Switch Configuration File](#) from column `Enclosure_Switch`.

```
$ cd /var/lib/tftpboot
```

```
$ sed -i 's/{switchname}/<switch_name>/' 6127xlg_irf.cfg
```

```
$ sed -i 's/{admin_password}/<admin_password>/' 6127xlg_irf.cfg
```

```
$ sed -i 's/{user_name}/<user_name>/' 6127xlg_irf.cfg
```

```
$ sed -i 's/{user_password}/<user_password>/' 6127xlg_irf.cfg
```

3. Access the InterConnect Bay1 6127XLG switch to configure the IRF (Intelligent Resilient Framework).

 **Note:**

On a new switch the user is presented with the following when connecting to the console and must type CTRL_C or CTRL_D to break out of the loop.

When trying to save the config, the following prompt is received: [HPE] [HPE] save The current configuration will be written to the device. Are you sure? [Y/N]: Before pressing ENTER you must choose 'YES' or 'NO'[Y/N]:y Please input the file name(*.cfg)[flash:/startup.cfg] (To leave the existing filename unchanged, press the enter key): User can leave this default startup.cfg unchanged, or change to another name. The cfg file will be used for next reboot.

```
$ ssh <oa username>@<oa address>
```

If it shows standby, ssh to the other OA address.

```
OA-FC15B41AEA05> connect interconnect 1
```

```
....
```

```
<HPE>system-view
```

System View: return to User View with Ctrl+Z.

(Note: Run the following commands:)

```
irf member 1 priority 32

interface range Ten-GigabitEthernet 1/0/17 to Ten-GigabitEthernet
1/0/20

shutdown

quit

irf-port 1/1

port group interface Ten-GigabitEthernet1/0/17

port group interface Ten-GigabitEthernet1/0/18

port group interface Ten-GigabitEthernet1/0/19

port group interface Ten-GigabitEthernet1/0/20

quit
```



```
interface range Ten-GigabitEthernet 1/0/17 to Ten-GigabitEthernet
1/0/20

    undo shutdown

    quit

save

irf-port-configuration active
```

4. Access the InterConnect Bay2 6127XLG switch to re-number to IRF 2.

```
OA-FC15B41AEA05> connect interconnect 2

....

<HPE>system-view

System View: return to User View with Ctrl+Z.

[HPE] irf member 1 renumber 2

Renumbering the member ID may result in configuration change or
loss. Continue?[Y/N]Y

[HPE]save

The current configuration will be written to the device. Are you
sure? [Y/N]:Y

Please input the file name(*.cfg)[flash:/startup.cfg]

(To leave the existing filename unchanged, press the enter key):

Validating file. Please wait...

Saved the current configuration to mainboard device successfully.

[HPE]quit

<HPE>reboot

Start to check configuration with next startup configuration file,
please wait.....DONE!

This command will reboot the device. Continue? [Y/N]:Y

Now rebooting, please wait...

System is starting...
```

5. Configure the IRF on Bay2 6127XLG switch

After rebooting, the interfaces will begin with number 2 such as Ten-GigabitEthernet2/0/17, Ten-GigabitEthernet2/1/5. Run the following commands:

```
system-view

interface range Ten-GigabitEthernet 2/0/17 to Ten-GigabitEthernet
2/0/20

shutdown

quit

irf-port 2/2

port group interface Ten-GigabitEthernet2/0/17

port group interface Ten-GigabitEthernet2/0/18

port group interface Ten-GigabitEthernet2/0/19

port group interface Ten-GigabitEthernet2/0/20

quit

interface range Ten-GigabitEthernet 2/0/17 to Ten-GigabitEthernet
2/0/20

undo shutdown

quit

save

irf-port-configuration active
```

6. Run "reboot" command on both switches.

```
<HPE>reboot
Start to check configuration with next startup configuration file,
please wait.....DONE!
This command will reboot the device. Continue? [Y/N]:Y
Now rebooting, please wait...

System is starting...
```

7. Verify the IRF for the 6127XLG switches. When reboot is finished, verify IRF is working with both member and ports from previous two switches, which form IRF to act as one switch now.

```
<HPE>system-view

System View: return to User View with Ctrl+Z.
```

```
[HPE]display irf configuration

MemberID NewID   IRF-Port1                IRF-Port2
-----
1         1         Ten-GigabitEthernet1/0/17  disable
                Ten-GigabitEthernet1/0/18
                Ten-GigabitEthernet1/0/19
                Ten-GigabitEthernet1/0/20

2         2         disable                  Ten-
GigabitEthernet2/0/17

                Ten-
GigabitEthernet2/0/18

                Ten-
GigabitEthernet2/0/19

                Ten-
GigabitEthernet2/0/20

[HPE]
```

8. Configure the IRF switch with predefined configuration file.

```
<HPE>tftp 192.168.20.11 get 6127xlg_irf.cfg startup.cfg
startup.cfg already exists. Overwrite it? [Y/N]:Y
Press CTRL+C to abort.

% Total % Received % Xferd Average Speed Time Time Time Current
Dload Upload Total Spent Left Speed
100 9116 100 9116 0 0 167k 0 --:--:-- --:--:-- --:--:-- 178k

<HPE>system-view
System View: return to User View with Ctrl+Z.

[HPE]configuration replace file flash:/startup.cfg

Current configuration will be lost, save current configuration? [Y/
N]:N

Now replacing the current configuration. Please wait ...

Succeeded in replacing current configuration with the file flash:/
```

```

startup.cfg.

[<switch_name>]save flash:/startup.cfg

The current configuration will be saved to flash:/startup.cfg.
Continue? [Y/N]:Y

flash:/startup.cfg exists, overwrite? [Y/N]:Y

Now saving current configuration to the device.

Saving configuration flash:/startup.cfg.Please wait...

Configuration is saved to device successfully.

[<switch_name>]

```

Bastion Host Installation

This section outlines the use of the Installer Bootstrap Host to provision db-2/RMS2 with an operating system and configure it to fulfill the role of Database Host. After the Bastion Host is created, it is used to complete the installation of OCCNE.

Provision Second Database Host (RMS2) from Installer Bootstrap Host (RMS1)

Table 2-2 Terminology used in Procedure

Name	Description
bastion_full_name	This is the full name of the Bastion Host as defined in the hosts.ini file. Example: bastion-2.rainbow.us.labs.oracle.com
bastion_kvm_host_full_name	This is the full name of the KVM server (usually RMS2/db-2) that hosts the Bastion Host VM. Example: db-2.rainbow.us.labs.oracle.com
bastion_kvm_host_ip_address	This is the IPv4 ansible_host IP address of the server (usually RMS2/db-2) that hosts the Bastion Host VM. Example: 172.16.3.5
bastion_short_name	This is the name of the Bastion Host derived from the bastion_full_name up to the first ".". Example: bastion-2
bastion_external_ip_address	This is the external address for the Bastion Host Example : 10.75.148.5 for bastion-2
bastion_ip_address	This is the internal IPv4 "ansible_host" address of the Bastion Host as defined within the hosts.ini file. Example: 172.16.3.100 for bastion-2
cluster_full_name	This is the name of the cluster as defined in the hosts.ini file field: occne_cluster_name. Example: rainbow.us.labs.oracle.com
cluster_short_name	This is the short name of the cluster derived from the cluster_full_name up to the the first ".". Example: rainbow

 **Note:**

The Bootstrap Host must be setup to use **root/**
<customer_specific_root_password> as the credentials to access it.
Setting that user/password is part of the instructions at: [Installation of Oracle Linux 7.x on Bootstrap Host](#).

Procedure

Following is the procedure to install Bastion:

1. Copy the necessary files from the Utility USB to support the OS Install:
This procedure is used to provide the steps for copying all supporting files from the Utility USB to the appropriate directories so that the Provision Container successfully installs OL7 onto RMS2.
 - a. Login to the Bootstrap Host using the root credentials configured during OS installation of the Bootstrap Host.
 - b. Create the directories needed on the Installer Bootstrap Host.

```
$ mkdir -p /var/ocne/cluster/<cluster_short_name>/yum.repos.d
```

- c. Mount the Utility USB.

 **Note:**

Follow the instructions for mounting a USB in Linux are at:
[Installation of Oracle Linux 7.x on Bootstrap Host](#).

- d. Copy the hosts.ini file (created using procedure: [OCCNE Inventory File Preparation](#)) into the /var/ocne/cluster/<cluster_short_name>/ directory. This hosts.ini file defines the Bastion KVM Host to the Provision Container running the provision image downloaded from the repo.

```
$ cp /<path_to_usb>/hosts.ini /var/ocne/cluster/  
<cluster_short_name>/hosts.ini
```

Example:

```
$ cp /media/usb/hosts.ini /var/ocne/cluster/rainbow/hosts.ini
```

- e. Edit the /var/ocne/cluster/<cluster_short_name>/hosts.ini file to include the ToR host_net (vlan3) VIP for NTP clock synchronization. Use the ToR VIP address (ToRswitch_Platform_VIP) as defined in procedure: [Installation PreFlight Checklist : Complete OA and Switch IP SwitchTable](#) as the NTP source. Update the ntp_server field with the VIP address. Update the ocne_repo_host_address to point to this Bootstrap Host internal address. This is being used for PXE booting and accessing the NFS share on the Installer Bootstrap Host (db-1/RMS1).

Example (from hosts.sample.ini):

```
ntp_server='172.16.3.1'
occne_repo_host_address='172.16.3.4'
```

- f. Follow procedure [Populate the MetalLB Configuration File](#) to copy the MetalLB configuration file into the `/var/occne/cluster/<cluster_short_name>/` directory and configure it.

```
$ cp /<path_to_usb>/mb_configmap.yaml /var/occne/cluster/
<cluster_short_name>/mb_configmap.yaml
```

Example:

```
$ cp /media/usb/mb_configmap.yaml /var/occne/cluster/rainbow/
mb_configmap.yaml
```

- g. Copy the customer specific `.repo` file from the Utility USB to the Installer Bootstrap Host.
This is the `.repo` file created by the customer that provides access to the onsite (within their network) yum repositories needed to complete the full deployment of OCCNE onto the Installer Bootstrap Host. It needs to be in two places, one for the local system, and one for the target systems.

```
$ cp /<path_to_usb>/<customer_specific_repo>.repo /var/occne/
cluster/<cluster_short_name>/yum.repos.d/.
$ cp -r /var/occne/cluster/<cluster_short_name>/yum.repos.d /var/
occne/.
$ echo "reposdir=/var/occne/yum.repos.d" >> /etc/yum.conf
```

Example:

```
$ cp /media/usb/ol7-mirror.repo /var/occne/cluster/rainbow/
yum.repos.d/
$ cp -r /var/occne/cluster/rainbow/yum.repos.d /var/occne/
$ echo "reposdir=/var/occne/yum.repos.d" >> /etc/yum.conf
```

2. Set up the `/etc/hosts` file for the Central Repo and Verify Access:

- a. Add an entry to the `/etc/hosts` file on the Installer Bootstrap Host to provide a name mapping for the Customer Central Repository.

```
$ vi /etc/hosts
```

Example:

```
10.75.200.217 rainbow-reg
```

To Verify:

```
$ ping <central_repo_name>
```

Example:

```
# ping rainbow-reg
PING reg-1 (10.75.200.217) 56(84) bytes of data.
64 bytes from reg-1 (10.75.200.217): icmp_seq=1 ttl=61
time=0.248 ms
64 bytes from reg-1 (10.75.200.217): icmp_seq=2 ttl=61
time=0.221 ms
64 bytes from reg-1 (10.75.200.217): icmp_seq=3 ttl=61
time=0.239 ms
```

- b. Verify the repo access execute the following command:

```
$ yum repolist
```

Example:

```
$ yum repolist
Loaded plugins: ulninfo
repo id          repo
name
      status
!UEKR5/x86_64    Unbreakable Enterprise Kernel Release 5
for Oracle Linux 7 (x86_64)      80
!addons/x86_64   Oracle Linux 7 Addons
(x86_64)                                               91
!developer/x86_64 Packages for creating test and
development environments for Oracle Linux 7      226
!developer_EPEL/x86_64 Packages for creating test and
development environments for Oracle Linux 7 13,246
!ksplce/x86_64   Ksplice for Oracle Linux 7
(x86_64)                                               393
!latest/x86_64   Oracle Linux 7 Latest
(x86_64)                                               5,401
repolist: 19,437
```

3. Copy the OL7 ISO to the Installer Bootstrap Host:
The iso file must be accessible from a Customer Site Specific repository. This file should be accessible because the ToR switch configurations were completed in procedure: [Configure Top of Rack 93180YC-EX Switches](#)

Copy the OL7 ISO file to the /var/occne directory. The example below uses OracleLinux-7.5-x86_64-disc1.iso. If this file was copied to the Utility USB, it can be copied from there into the same directory on the Bootstrap Host.

 **Note:**

If the user copies this ISO from their laptop then they must use an application like WinSCP pointing to the Management Interface IP.

```
$ scp <usr>@<site_specific_address>:/<path_to_iso>/OracleLinux-7.5-
x86_64-disc1.iso /var/occne/OracleLinux-7.5-x86_64-disc1.iso
```

4. Install Packages onto the Installer Bootstrap Host. Use YUM to install necessary packages onto the installer Bootstrap Host.

```
$ yum install docker-engine nfs-utils ansible
```

5. Set up access to the Docker Registry on the Installer Bootstrap Host:
- a. Copy the docker registry certificate to two places on the Bootstrap Host.

 **Note:**

How to obtain the docker registry certificate <source> is not necessarily covered in the procedure. The user can use the instructions at reference 1 to understand this more thoroughly.

```
$ mkdir -p /var/ocne/certificates
$ cp <source>.cert /var/ocne/certificates/
<ocne_private_registry>:<ocne_private_registry_port>.cert
$ mkdir -p /etc/docker/certs.d/
<ocne_private_registry>:<ocne_private_registry_port>
$ cp <source>.cert /etc/docker/certs.d/
<ocne_private_registry>:<ocne_private_registry_port>/ca.cert
```

Example:

```
$ mkdir -p /var/ocne/certificates
$ cp <source>.cert /var/ocne/certificates/rainbow_reg:5000.cert
$ mkdir -p /etc/docker/certs.d/rainbow_reg:5000
$ cp <source>.cert /etc/docker/certs.d/rainbow_reg:5000/ca.cert
```

b. Start the docker daemon.

```
$ systemctl daemon-reload
$ systemctl restart docker
$ systemctl enable docker
```

Verify docker is running:

```
$ ps -elf | grep docker
$ systemctl status docker
```

6. Setup NFS on the Installer Bootstrap Host:

Run the following commands using sudo (nfs-utils has already been installed in previous steps).

 **Note:**

The IP address used in the echo command is the Platform VLAN IP Address (VLAN 3) of the Bootstrap Host (RMS 1) as given in: [Installation PreFlight Checklist : Site Survey Host Table](#).

```
$ echo '/var/ocne 172.16.3.4/24(ro,no_root_squash)' >> /etc/exports
$ systemctl start nfs-server
$ systemctl enable nfs-server
```

Verify nfs is running:

```
$ ps -elf | grep nfs
$ systemctl status nfs-server
```


7. Set up the Boot Loader on the Installer Bootstrap Host:

```
$ mkdir -p /var/occne/pxelinux
$ mount -t iso9660 -o loop /var/occne/OracleLinux-7.5-x86_64-
disc1.iso /mnt
$ cp /mnt/isolinux/initrd.img /var/occne/pxelinux
$ cp /mnt/isolinux/vmlinuz /var/occne/pxelinux
```

8. Verify and Set the PXE Configuration File Permissions on the Installer Bootstrap Host:

 **Note:**

Each file configured in the step above must be open for read and write permissions.

```
$ chmod -R 777 /var/occne/pxelinux
```

9. Disable DHCP and TFTP on the Installer Bootstrap Host. The TFTP and DHCP services running on the Installer Bootstrap Host may still be running. These services must be disabled.

```
$ systemctl stop dhcpd
$ systemctl disable dhcpd
$ systemctl stop tftp
$ systemctl disable tftp
```

10. Disable SELINUX: Set SELINUX to permissive mode. In order to successfully set the SELINUX mode, a reboot of the system is required. The getenforce command is used to determine the status of SELINUX.

```
$ getenforce
active
```

If the output of this command displays "active", change it to "permissive" by editing the `/etc/selinux/config` file.

```
$ vi /etc/selinux/config
```

Change the SELINUX variable to passive: `SELINUX=permissive`
save the file

Reboot the system: `reboot`

11. Generate the SSH private and public keys on Bootstrap Host.

This command generates a private and public key for the cluster. These keys are passed to the Bastion Host and used to communicate to other nodes from that Bastion Host. The public key is passed to each node on OS install. Do not supply a passphrase when it asks for one. Click Enter.

 **Note:**

The private key (`occne_id_rsa`) must be copied to a server that going to access the Bastion Host because the Bootstrap Host is repaved. This key is used later in the procedure to access the Bastion Host after it has been created.

Execute the following commands on the Bootstrap Host:

```
$ mkdir -m 0700 /var/occne/cluster/<cluster_short_name>/.ssh
$ ssh-keygen -b 4096 -t rsa -C "occne installer key" -f "/var/occne/cluster/<cluster_short_name>/.ssh/occne_id_rsa" -q -N ""
```

12. Execute the OS Install and Bastion VM Creation on Bastion KVM Host (RMS2) from the Installer Bootstrap Host:

- a.** Run the docker commands below to perform the OS install and Bastion Host VM creation on the Bastion KVM Host (RMS2):

```
$ docker run --rm --network host --cap-add=NET_ADMIN -v /var/occne/cluster/<cluster_short_name>:/host -v /var/occne:/var/occne:rw -e "OCCNEARGS=--limit=<bastion_full_name>,<bastion_kvm_host_full_name>,localhost" <image_name>:<image_tag>
```

Example:

```
$ docker run -it --rm --network host --cap-add=NET_ADMIN -v /var/occne/cluster/rainbow:/host -v /var/occne:/var/occne:rw -e "OCCNEARGS=--limit=bastion-2.rainbow.lab.us.oracle.com,db-2.rainbow.lab.us.oracle.com,localhost" winterfell:5000/occne/provision:1.3.0
```

- b.** Verify that Bastion Host VM is installed by logging into RMS2/db-2 and issuing the following command. The `<ansible_host>` field (which is an IPv4 address) is derived from the `hosts.ini` file db-2 entry for `host_hp_gen_x` groups.

 **Note:**

This command is optional. Had a failure actually occurred, the docker run command would have experienced failures.

```
ssh -i /var/occne/cluster/<cluster_short_name>/.ssh/occne_id_rsa admusr@<oam_host>
```

```
$ sudo virsh list
```

Example:

```
ssh -i /var/occne/cluster/rainbow/.ssh/occne_id_rsa admusr@10.75.148.6
```

```
$ sudo virsh list
```

```

Id      Name
-----
11      bastion-2.rainbow.lab.us.oracle.com running

```

- c. Login to Bastion Host from the Bootstrap Host as admusr using the generated key from the `/var/ocne/cluster/<cluster_short_name>` directory to confirm the VM is set up correctly. The `<oam_host>` field (which is an IPv4 address) is derived from the `hosts.ini` file `bastion-2` entry for the `host_kernel_virtual` group.

 **Note:**

This command is optional. Had a failure actually occurred, the `docker run` command would have experienced failures.

```
ssh -i /var/ocne/cluster/<cluster_short_name>/.ssh/
ocne_id_rsa admusr@<oam_host>
```

Example:

```
ssh -i /var/ocne/cluster/rainbow/.ssh/ocne_id_rsa
admusr@10.75.148.5
```

Automated Installation

This section details the steps required to execute the automated configuration of the Bastion Host VM. This consists of two main section:

1. Setting up and executing the `deploy.sh` script on the Bootstrap Host.
2. Accessing the Bastion Host and executing the final commands to execute the `pipeline.sh` script to complete the Bastion Host configuration and deploy the OCCNE cluster.

Following is the procedure to perform automated installation of the Bastion Host VM:

1. Setting up for and executing the `deploy.sh` script on the Bootstrap Host:

The `deploy.sh` script performs the initial pre-configuration of the Bastion host. This includes installing `ansible`, executing the `ansible` playbook `configBastionHost.yaml` to setup the initial files and staging directories on the Bastion Host and executing the `pipeline` to setup the artifacts directory. The script is executed on the Bootstrap Host using a set of environment variables that can be initialized on the command line along with the `deploy.sh` script. These variables include the following:

Table 2-3 Environmental Variables

Name	Comment	Example usage
CENTRAL_REPO	Defines the customer specific repository host name. Note: This would be used in conjunction with <code>CENTRAL_REPO_IP</code> and <code>CENTRAL_REPO_DOCKER_PORT</code> .	<code>CENTRAL_REPO=customer_repo \</code> <code>CENTRAL_REPO_IP=<IP_Address> \</code> <code>CENTRAL_REPO_DOCKER_PORT=5000</code>

Table 2-3 (Cont.) Environmental Variables

Name	Comment	Example usage
CENTRAL_REPO_IP	Defines the customer specific repository IPv4 address.	See above.
CENTRAL_REPO_DOCKER_PORT	Defines the customer specific repository docker port number. Defaults to 5000.	See above.
OCCNE_CLUSTER	Defines the cluster short name.	OCCNE_CLUSTER=rainbow
OCCNE_BASTION	Bastion Host full name	OCCNE_BASTION=bastion-2.rainbow.us.labs.oracle.com
OCCNE_VERSION	The version tag of the image releases	OCCNE_VERSION=1.6.0

2. Copy necessary files from Utility USB to the Bootstrap Host staging directory:
 - a. The MySQL .zip file (ex: V980756-01.zip) must be copied to the staging directory `/var/occne` directory. This file should be provided from the Utility USB. This file is used for installing the ndb MySQL nodes.

```
$ cp /<usb_dev>/db/*.zip /var/occne/*.zip
```

- b. Copy the `configBastionHost.yaml` file from the Customer Utility USB to the staging directory `/var/occne/`.

```
$ cp /<usb_dev>/configBastionHost.yaml /var/occne/.
```

- c. Copy the `deploy.sh` script from the Customer Utility USB to the staging directory at `/var/occne/` and set the file to executable mode.

```
$ cp /<usb_dev>/deploy.sh /var/occne/.
$ chmod +x /var/occne/deploy.sh
```

3. Execute the `deploy.sh` script from the `/var/occne/` directory with the required parameters set.

```
$ export CENTRAL_REPO=<customer_specific_repo_name>
$ export CENTRAL_REPO_IP=<customer_specific_repo_ipv4>
$ export OCCNE_CLUSTER=<cluster_short_name>
$ export OCCNE_BASTION=<bastion_full_name>
$ ./deploy.sh
```

Customer Example:

```
$ export CENTRAL_REPO=central-repo
$ export CENTRAL_REPO_IP=10.10.10.10
$ export OCCNE_CLUSTER=rainbow
$ export OCCNE_BASTION=bastion-2.rainbow.us.labs.oracle.com
$ ./deploy.sh
```

The command above can be executed like the following:

```
$ CENTRAL_REPO=central-repo
CENTRAL_REPO_IP=10.10.10.10 OCCNE_CLUSTER=rainbow
OCCNE_BASTION=bastion-2.rainbow.us.labs.oracle.com ./deploy.sh
```

4. Execute the final deploy on Bastion Host:
The following commands are executed from the Bastion Host to complete the Bastion Host configuration and deploy OCCNE on the Bare Metal system.

 **Note:**

The Bootstrap Host cannot be used to access the Bastion Host as it will be re-paved from execution of this command.

- a. Login to the Bastion Host as *admusr*. The private key that was saved earlier should be used to access the Bastion Host from a server other than the Bootstrap Host using the `ssh` command. This private key should be copied to the `/home/<user>/.ssh` directory on that server as `id_rsa` using `scp` (or `winSCP` from a desktop PC). The permissions of the key must be set to `0600` using the command: `chmod 0600 ~/.ssh/id_rsa`

```
$ ssh -i ~/.ssh/id_rsa admusr@<bastion_external_ip_address>
```

- b. Execute the following command to complete the deployment of OCCNE from the Bastion Host (excluding re-install on the Bastion Host and its KVM host, which are already setup). This action will re-pave the Bootstrap Host RMS.

```
$ export OCCNE_PROV_DEPLOY_ARGS='--limit=!<bastion_full_name>,!<bastion_kvm_host_full_name>'
$ export OCCNE_DB_ARGS='--extra-vars=mate_site_db_replication_ip=<remote_site_db_replication_service_ip>,occne_mysqlndb_cluster_id=<mysqlndb_cluster_identifier>'
$ /var/occne/cluster/<cluster_short_name>/artifacts/pipeline.sh
```

Customer Example:

```
$ export OCCNE_PROV_DEPLOY_ARGS='--limit=!bastion-2.rainbow.lab.us.oracle.com,!db-2.rainbow.lab.us.oracle.com'
$ export OCCNE_DB_ARGS='--extra-vars=mate_site_db_replication_ip=10.75.182.88,occne_mysqlndb_cluster_id=2'
$ /var/occne/cluster/rainbow/artifacts/pipeline.sh
```

To save the output from the `pipeline.sh` script the command can be written as:

```
$ /var/occne/cluster/rainbow/artifacts/pipeline.sh | tee pipeline$(date +%F_%H%M%S).log
```

 **Note:**

While Installing on the First Site ignore `OCCNE_DB_ARGS` configuration parameter which Provides Mate Site DB Replication Service Load Balancer IP. Provide the Mate Site DB Replication Service Load Balancer IP and MySQL Cluster identifier while Installing MySQL NDB Cluster on second site.

5. Change MySQL root user password. Refer to [Change MySQL root user password](#)
6. Update the Bastion KVM Host repo file:
Since db-2 was not part of the final OS install and cluster deploy, it's `/var/occne/yum.repos.d/*.repo` file is not pointing to the Bastion Host as its YUM repo. That file on RMS2/db-2 must be updated so that it now points to the Bastion Host as the repo. After the Bastion Host was created, the `.repo` file that was copied onto the Bastion Host has the correct settings. That file can just be copied back to RMS2/db-2.
 - a. From the bastion Host, login to the Bastion Host KVM Host, using the `occne` private key and the internal host IP address for that node (extracted from the `hosts.ini` file):

```
$ ssh -i /var/occne/cluster/<cluster_short_name>/.ssh/  
occne_id_rsa admusr@<bastion_ip_address>
```

- b. Remove the existing `.repo` files:

```
$ sudo rm /var/occne/yum.repos.d/*.repo
```
- c. Copy the Bastion specific `.repo` file from the Bastion Host to the Bastion KVM Host. Execute this command from the Bastion KVM Host:

```
scp <bastion_short_name>:/var/occne/yum.repos.d/*.repo /var/  
occne/yum.repos.d/
```

Example:

```
scp bastion-2:/var/occne/yum.repos.d/*.repo /var/occne/  
yum.repos.d/
```

Virtualized CNE Installation

This procedure details the steps necessary to configure and install an OCCNE cluster in an OpenStack Environment.

Prerequisites

1. The user must have access to an existing OpenStack Environment including the OpenStack Desktop.
2. The OpenStack Environment is configured with appropriate resource flavors and network resources for resource allocation to the VMs created using this procedure.
3. MetalLB configMap file `mb_configmap.yaml` for deploying MetalLB as LBaaS (**Optional**: Required only when LBaaS is selected as MetalLB).
4. Octavia Load Balancing plugin must be installed on the OpenStack Environment.
5. **All necessary images, binaries, and files have been downloaded from Oracle OSDC prior to executing this procedure and these resources are available for use in this procedure (refer [Artifact Acquisition and Hosting](#) for instructions on how to generate the lists of images and binaries).**
6. Users must have a public SSH key that can be configured for logging into the Bootstrap Host. This key should be placed into the customer OpenStack Environment prior to running this procedure using the following:

Use the **Import Key** tab on the **Launch Instance** → **Key Pair** dialog or via the **Compute** → **Access and Security** screen.

Limitation/Expectations

1. It is expected that the user is familiar with the use of OpenStack as a virtualized provider and the OpenStack client.

Note:

All OpenStack client commands listed in this procedure are executed from the Bootstrap Host after it has been instantiated.

2. The customer has made available a central repository for all images, binaries, helm charts, etc, prior to executing this procedure.
3. All necessary networking (whether using fixed IPs or floating IPs) has been defined for use on the Openstack provider.

Note:

The OpenStack commands in the procedures below are from a specific version of the OpenStack Desktop. The desktop used at the customer site may be slightly different depending on the version installed. The operations should be compatible.

Download qcow2 image

This section describes procedure on how to download qcow2 image from OSDC.

1. Go to <https://edelivery.oracle.com> and login using appropriate credentials.
2. Enter “**Oracle Linux Virtual Machine Image for Openstack**” in search field and click on **Search**.
3. Select *DLP: Oracle Linux 7 Virtual Machine Image for Openstack 7.0.0.0.0* and click on **Add to Cart**.
4. Click on **Checkout**.
5. Click on **Continue**.
6. Accept the license agreement. Click on **Continue**.
7. Click on link for : **V979992-01.zip** Oracle Linux 7.5 Virtual Machine Image for OpenStack, 458.4 MB
8. Click on **Download**.

 **Note:**

- a. Optionally use the link for 7.6 instead: **V981347-01.zip** Oracle Linux 7.6 Virtual Machine Image for OpenStack, 487.9 MB
- b. These zip archives files contains the following respective files, *OracleLinux-7.5-x86_64.qcow2* and *OracleLinux-7.6-x86_64.qcow2*.

Upload an Image to an OpenStack Environment

This is the process of uploading the qcow2 image to an openstack environment. The image is provided at OSDC.

 **Note:**

This procedure is executed from the OpenStack Desktop.

1. Login to the Openstack Desktop using the specific credentials.
2. Select **Compute** → **Images**.
3. Select the **+Create Image** button. This brings up a new dialog.
4. Enter a name for the image. Use the same name of the image which was used to create and download the image (ex: occne_bootstrap-1.6.0).
5. Under **Image Source** → **select: File**.
6. This will enable a **File*** → **Browse** button. Select this button to bring up a Windows Explorer dialog.
7. From the Windows dialog, select the image that was downloaded from the <Release Artifacts> above. This will insert the image name into the OpenStack Create Image dialog and set the Format.

 **Note:**

This should be set to QCOW2 - QEMU Emulator. If not set, use the pulldown menu to select that format.

8. Select the **Create Image** button at the bottom right of the dialog. This will start the image upload process. It will take a while so be patient. You will not be able to actually see the image being uploaded even if you log into another OpenStack instance.
9. When the process is complete, the image should be listed in the **Compute** → **Images** screen.

Bootstrap Host Creation

The Bootstrap Host is provisioned to drive the creation of the virtualized cluster using Terraform, the OpenStack Client, and Ansible Playbook(s). A qcow2 image

was provided as part of the OSDC download and should be available on the users OpenStack Environment as per the previous section of this document.

 **Note:**

The examples below are for reference only. While the steps are correct the actual values used will be different. The following steps are to be performed manually on the customer specific Openstack Environment Desktop.

1. Login to the OpenStack Environment using your OpenStack credentials, the appropriate domain and project name.
2. Select **Compute** → **Instances**
3. Select the **Launch Instances** tab on the upper right. A dialog will appear to configure a VM instance.
4. Enter an Instance Name (for example: occne-<name>). Leave the Availability Zone and Count set as is.
5. Select **Source** on the left hand side of the dialog. A new dialog appears (**Note:** there might be a long list of available images to choose from).
6. Make sure the **Select Boot Source** pulldown is set to **Image**.
7. Enter occne-bootstrap in the **Available** filter. This will display the occne-bootstrap-<x.y.z> image uploaded in the previous sections of this procedure.
8. Select the **OCCNE Bootstrap Host** image by selecting the "+" on the right side of the image listing. This adds the image as the source for this VM.
9. Select **Flavor** on the left hand side.
10. Enter a string (not case sensitive) which best describes the flavor being used for this customer specific OpenStack Environment in the **Available** search filter. This shortens the list of possible choices.
11. Select the appropriate customer specific flavor (for example: OCCNE-Bootstrap-host) by selecting the "+" on the right side of the flavor listings. This adds the resources to the Launch Instance dialog.

 **Note:**

The BSH image requires a flavor that includes a disk size of 40GB or higher. The RAM size should be 8GB or higher although that is not a restriction.

12. Select **Networks** on the left hand side.
13. Enter the appropriate network name as defined by the customer with the OpenStack Environment (example: ext-net) in the Available search filter. This shortens the list of possible choices.
14. Select the appropriate network by selecting the "+" on the right side of the flavor listings. This adds the external network interface to the Launch Instance dialog.
15. Select **Key Pair**. This dialog assumes you have already uploaded a public key to OpenStack (see Prerequisites above). Select the appropriate key by selecting

the "↑" on the right side of the key pair listings. This adds the public key to the `authorized_keys` file on the Bootstrap Host.

16. Select **Configuration**. This screen allows the user to add configuration data which is used by cloud-init to set on the VM, the initial cloud-user and hostname/FQDN additions to the `/etc/hosts` file. Use the following configuration. This must be copied into the Customization Script text box. Make sure the fields marked as `<instance_name_from_details_screen>` are updated with the instance name you used in step 5 above. Leave the other fields on this dialog set to their default setting.

```
#cloud-config
hostname: <instance_name_from_details_screen>
fqdn: <instance_name_from_details_screen>
system_info:
  default_user:
    name: cloud-user
    lock_passwd: false
  write_files:
    - content: |
        127.0.0.1 localhost localhost4 localhost4.localdomain4
<instance_name_from_details_screen>
        ::1 localhost localhost6 localhost6.localdomain6
<instance_name_from_details_screen>
      path: /etc/hosts
      owner: root:root
      permissions: '0644'
```

17. Select **Launch Instance** at the lower right side of the initial dialog. This will launch the creation of the VM. This can be observed back at the **Compute → Instances** screen.

Pre-deployment Configuration

The commands in this procedure are executed from the Bootstrap Host. All terraform commands are executed from the `/var/terraform` directory.

Obtain the TLS Certificate for OpenStack

Depending on the Customer's environment it is very likely that the customer's OpenStack uses a TLS certificate for access to the Openstack controller API. Without this certificate, OpenStack commands will not work. Customer's must obtain this certificate before using OpenStack client commands.

Note:

This step is required only if the customer Openstack environment requires a TLS certificate to access the controller from the cluster nodes and the bootstrap host.

1. Contact the OpenStack admin to provide the required TLS certificate to access the client commands (example: in an Oracle OpenStack system installed with kolla, the certificate is located at `/etc/kolla/certificates/openstack-cacert.pem`)

2. Copy the certificate to the Bootstrap Host to directory `/tmp/cacertificates/` (should already exist by using the Bootstrap Host image provided). Ensure that the certificate file name is `'openstack-cacert.pem'`, when it is copied to directory `'/tmp/cacertificates/'`. (If the certificate file name is different, rename it to `openstack-cacert.pem` before copying to directory `/tmp/cacertificates/`)
3. Set the environment variable `OS_CACERT` to `/var/occne/cacertificates/openstack-cacert.pem` using the command: `export OS_CACERT=/var/occne/cacertificates/openstack-cacert.pem`.

Get the Openstack RC (API v3) File

This file exports a number of environment variables on the Bootstrap Host for the given user which directs the OpenStack Client commands towards the particular OpenStack Environment. It must be copied to the users home directory on the Bootstrap Host so that the OpenStack Client commands can be executed.

Note:

These instructions may slightly vary for OpenStack Desktops.

1. From the OpenStack Desktop: go to **Project - > API Access**.
2. On the right hand side, **Download OpenStack RC File** pulldown menu and select **Openstack RC File (Identity API v3)**.
This will download a `openrc.sh` file prefixed with the OpenStack project name (ex: `OCCNE-openrc.sh`) to your PC.
3. SCP this file (ie. winSCP) to the Bootstrap Host in the `/home/admusr` directory as `.<project_name>-openrc.sh`

Note:

In order for SCP/winSCP to work properly, the key mentioned in the Prerequisites above must be used to access the Bootstrap Host. It may also be necessary to add the appropriate Security Group Rules to support SSH (**Rule:** SSH, **Remote:** CIDR **CIDR:** 0.0.0.0/0) under the **Network** → **Security Groups** page in the OpenStack Environment. Contact the OpenStack Administrator to get the proper rules added if necessary.

4. Execute the following command: `source .<project_name>-openrc.sh`
5. Execute the following command to verify the OpenStack Client is working:
`openstack image list`

Create SSH Key on Bootstrap Host

Create the keys that will be used to access the other VMs. This command generates the private and public keys that are passed to the Bastion Host and used to communicate with other node from that Bastion Host. Do not supply a passphrase

when it asks for one. Just hit enter. Also the private key should be copied to a place for safe keeping should the Bootstrap Host be destroyed.

```
$ ssh-keygen -m PEM -t rsa -b 2048
```

Add Files to /tmp Directory

These files must be copied to the directories listed using scp or some other means (ie. winSCP).

1. There are three directories on the Bootstrap Host. These three directories are as follows:
 - a. /tmp/yum.repos.d
 - b. /tmp/db
 - c. /tmp/certificates
2. Within these three directories the user must supply the following mandatory files:
 - a. Add a customer specific central repo .repo file to the */tmp/yum.repos.d* directory which allows for access to the customer specific repo (ex: winterfell-mirror.repo).
 - b. Add a mysql .zip file. (ex: V980756-01.zip) to the */tmp/db* directory. This file is used for installing the ndb MySQL cluster and is downloaded from OSDC
 - c. Add a docker registry certificate to the */tmp/certificates* directory for the central docker registry. This file must be copied using the following format:<central_repo_hostname>:<central_repo_port>.crt(ex: winterfell:5000.crt).
3. Make sure the permissions of each file is at least readable (using 0644) using the `sudo chmod 644 <filename>` command.

MetalLB Configuration

MetalLB can be used to replace Octavia on the vCNE deployment in Openstack. If the user creates and configures the **mb_configmap.yaml** file as indicated below, sets the variable **OS_LBAAS_ENABLED=false** on the command line when executing the `deploy.sh` command, and configures the fields in the `cluster.tfvars` file as indicated in section [Configuration for using OCCNE LB](#), the deploy will automatically disable Octavia, enable/install MetalLB, and create the necessary Load Balancing VMs (LBVMs) to support service load balancing for OCCNE (and the NFs as they are configured). A LBVM will be created per Peer Address Pool name, as defined in the **mb_configmap.yaml** file, which processes the load balancing functionality. These VMs are named according to the following format in Openstack: **<cluster_name>-<peer_address_pool_name>-lbvm** and are visible from the Openstack desktop.

Copy the MetalLB configMap file: `mb_configmap.yaml` to directory `/var/terraform` on the Bootstrap Host

The contents of the `mb_configmap.yaml` file includes separate peer pool names (see the `- name` field in the example file below). The peer pool names are user specific, defined in any order, there is no limitation on the number of peer pools, and they should all be lowercase chars.

The `addresses` field can be configured in one of the following three ways:

1. As a single cidr in the form of: 'xxx.xxx.xxx.xxx/xx' which defines a range of IPs in a given subnet.
2. As a list of specific IP addresses (not in any specific order) in the form of: 'xxx.xxx.xxx.xxx/32'.
3. As a specific range of IP addresses in the form of: 'xxx.xxx.xxx.xxx - xxx.xxx.xxx.xxx'

 **Note:**

The first peer address pool name (oam) in the example below defines the LoadBalancer peer address pool for the OCCNE OAM common services. Currently there are 6 total if the DB is included in the deployment. The others are just for example.

If a cidr is used, MetalLB will attempt to use the IPs in a sequential manner. It will start with the first indicated IP and apply them sequentially. For example: 10.75.235.0/25 will cause MetalLB to assign the IPs starting at .0 - .5. The user must guarantee these addresses are available. If the field `avoid-buggy-ips: true` is added to the peer address-pool, then MetalLB will skip addresses .0 and .255.

Example:

```
configInline:
  peers:
    - peer-address: 172.17.0.2
      peer-asn: 65000
      my-asn: 64512
  address-pools:
    - name: oam
      protocol: bgp
      addresses:
        - '10.75.235.20/32'
        - '10.75.235.21/32'
        - '10.75.235.25/32'
        - '10.75.235.44/32'
        - '10.75.235.23/32'
        - '10.75.235.27/32'
    - name: signal
      protocol: bgp
      addresses:
        - '10.75.236.25 - 10.75.236.31'
    - name: signal1
      protocol: bgp
      addresses:
        - '10.75.236.0/28'
```

⚠ WARNING:

It is advised to place the address-pools IPs on a different subnet from the floatingip_network subnet (see **Configuring Floating IPs** below). MetalLB has no direct interface with Openstack. If the same network/subnet is used in both cases it is possible that the allocation of the IPs can collide with one another (that is, an IP listed in the mb_configmap.yaml file could be inadvertently use as a worker node IP since they are DHCP assigned by the terraform apply).

Updating cluster.tfvars File

The **cluster.tfvars** file should exist on the Bootstrap Host after it has been deployed in `/var/terraform/occne_example/cluster.tfvars`. The user must create a directory called `occne-<user>-<name>` with a copy of the example tfvars file included in the `/var/terraform` directory.

The fields in the **cluster.tfvars** file must be configured to adapt to the current customer Openstack Environment. The steps below detail how to collect and set the fields that must be changed. These instructions may vary depending on whether the user is deploying the cluster using floating IPs or using fixed IP. Differences are noted below when needed.

The **cluster.tfvars** file includes a parameter named: **use_floating_ip**. This field defaults to true and should remain as set if the user is configuring their system to use floating IPs allocated from existing DHCP enabled networks defined within their Openstack Provider. If the user wishes to use fixed IPs, this field must be set to false.

Note: All fields in the **cluster.tfvars** file are unique and should not be duplicated (even using comments - #) as this can cause possible parsing errors.

Common Configuration

This procedure applies to both floating IPs or Fixed IPs.

1. From the `/var/terraform` directory, copy directory `occne_example` and its contents **cluster.tfvars**, using the command below to create a new directory. Change the name of the new directory to include your name or user name to distinguish it from the `occne_example` directory.

```
$ cp -R occne_example occne_<user>
```

Example:

```
$ cp -R occne-abc-xyz
```

2. Use the following commands to retrieve the information necessary to configure the `cluster.tfvars`
 - a. The different flavor settings should be set according to the recommendations from [vCNE VM Sizing](#) in this document. An Admin user of the customer specific OpenStack Environment must add the flavors and provide the name of those flavors for configuration into the **cluster.tfvars** file. The name of each specific flavor that is used must be added as the value field of the key/value fields in the **cluster.tfvars** file.

- b. Once flavors have been added to the OpenStack Environment, the flavor name can be retrieved via the following OpenStack Client command from the Bootstrap Host shell or the Openstack jumpbox:

```
$ openstack flavor list | grep <flavor_name>
```

Example:

```
$ openstack flavor list | grep OCCNE
| 740c9472-db2c-4015-be5d-04eea1e9f647 | OCCNE-Bastion-Host
| 3840 | 100 | 0 | 2 | True |
| 337263ef-f69e-43b5-9c6f-ddfc7bb8237e | OCCNE-Bootstrap-host
| 8192 | 40 | 0 | 2 | True |
| 2de71982-3a27-45bf-9232-27a3e03c685f | OCCNE-DB-Data-Large
| 65536 | 60 | 0 | 8 | True |
| 27bdc547-a180-45a3-a349-7c574638ba27 | OCCNE-DB-Data-Medium
| 32768 | 40 | 0 | 4 | True |
| f1f1fce6-fc7a-4e33-9711-ea47d476a843 | OCCNE-DB-Data-Small
| 16384 | 20 | 0 | 2 | True |
| 4fbe0ac8-4e63-439b-ac6e-c54f40d6e116 | OCCNE-DB-Mgmt-Large
| 15360 | 60 | 0 | 4 | True |
| 252d96c6-956c-4e99-bae1-clff8c09cc8e | OCCNE-DB-Mgmt-Medium
| 7680 | 40 | 0 | 2 | True |
| 1c42a100-9923-4a82-b56c-579ac5060ff9 | OCCNE-DB-Mgmt-Small
| 3840 | 20 | 0 | 1 | True |
| 0ac5719e-0aa6-480e-9835-7ea091e53949 | OCCNE-DB-Sql-Large
| 32768 | 60 | 0 | 8 | True |
| adbe5fa8-09e5-41c1-a6db-91a3650bea95 | OCCNE-DB-Sql-Medium
| 16384 | 40 | 0 | 4 | True |
| 7c40962d-5992-4a8a-9668-f6d0dc0107c4 | OCCNE-DB-Sql-Small
| 8192 | 20 | 0 | 2 | True |
| b219a120-3031-4558-8214-00f5184fc4b4 | OCCNE-K8s-Master-
Large| 15360 | 80 | 0 | 4 | True |
| 0d37d92c-8b06-4052-a1cc-94368d63f898 | OCCNE-K8s-Master-Medium
| 7680 | 40 | 0 | 2 | True |
| 99883b4d-8532-4c72-acd8-00955a772b9d | OCCNE-K8s-Master-Small
| 3840 | 20 | 0 | 1 | True |
| facc50fa-e784-4259-aac4-731ce0cbf3ce | OCCNE-K8s-Worker-Large
| 32768 | 60 | 0 | 8 | True |
| 7307d7bf-bc78-43fd-a5a7-b6e59eeace81 | OCCNE-K8s-Worker-Medium
| 16384 | 40 | 0 | 4 | True |
| 17d58f3d-3f62-4acd-bb2b-b8b29fbb76a4 | OCCNE-K8s-Worker-Small
| 8192 | 20 | 0 | 2 | True |
| f407e26d-c677-47b3-b5ed-22951ad77e49 | OCCNE-K8s-Worker-XLarge
| 65536 | 80 | 0 | 16 | True |
```

- c. Use the following command to retrieve the **external_net** UUID.

```
$ openstack network list
```

```
+-----+
+-----+
| ID                                     | Name
| Subnets                               |
+-----+
+-----+
```

```

| 1d25d5ea-77ca-4f56-b364-f53b09292e7b | ext-net2
| f5c5ee71-8688-466d-a79f-4306e2bf3f6a |
| 668bc488-5307-49ad-9332-24fb0767bb39 | test-network
| 9432b2d5-99c0-43ee-8f8c-4709f38b68d9 |
| 903155c7-c3ff-4283-bc2b-f34e8b6e76b0 | occne-ebadger-
tc1 | ecbadd3e-e239-4830-b8c1-5ff94fa64c3a |
| 90c160aa-2ef7-47d3-a212-e1790d56c971 | ext-net-ipv6
| 4c0b844f-1557-4454-b561-88fa31f657f3 |
| c4a7569b-5448-4add-8c4e-006bbdd984ef | cluster1
| 4cf62be3-05e9-4a5b-b2a9-6aceee3c860f |
| e4351e3e-81e3-4a83-bdc1-dde1296690e3 | ext-net
| c0e0c185-ed65-4a53-a7a3-418277fb9a20 |
| fc36d63f-b30b-4c7f-979f-9b52b614bbd7 | occne-mkingre
| 7631612f-5d22-49be-975c-6e0a9329339b |
+-----+
+-----+

```

- Once the necessary data has been collected, navigate to the `occne_<user>` directory and edit the contents of the `cluster.tfvars` file in the newly created directory:

```
$ vi occne_<user>/cluster.tfvars
```

- The following miscellaneous fields should remain as set in the example `cluster.tfvars` file.
 - `public_key_path`
 - `image`
 - `ssh_user`
- The fields which define the number of each node type have defaulted values. For a standard deployment the defaults should be used. The user can update these values if their deployment requires additional nodes. Note these fields are integer values and do not require double quotes.
 - `number_of_bastionsnumber_of_k8s_nodes`
 - `number_of_k8s_masters_no_floating_ip`
WARNING: The number of master nodes must be set to an odd number. The recommended value for `number_of_k8s_masters_no_floating_ip` is 3.
 - `number_of_db_tier_management_nodes`
 - `number_of_db_tier_data_nodes`
 - `number_of_db_tier_sql_nodes`
- The corresponding flavors for each node type must be set to a unique flavor name. Flavors are provided from the Openstack Provider administrator.
 - `flavor_bastion`
 - `flavor_k8s_master`
 - `flavor_k8s_node`
 - `flavor_db_tier_management_node`
 - `flavor_db_tier_data_node`
 - `flavor_db_tier_sql_node`

7. The **cluster_name** and **network_name** should be set as the same value.
Note: This field is used to create each of the nodes contained within the cluster. Kubespray does not allow upper case alphanumeric characters to be used in the node hostname. Do not use any uppercase characters when defining this cluster-name field.

```
# Kubernetes short
cluster_namecluster_name = "<cluster-name>"
# networking
network_name = "<cluster-name>"
```

8. The **subnet_cidr** defines the tenant side network ip address range. This field should remain set to the default value.
9. The field **bastion_allowed_remote_ips** defines the configuration for the bastion networking security group. This field should remain set to the default value.
10. For setting the **ntp_server** value in the **cluster.tfvars** file, use the IP Address of your cloud URL. One way of obtaining this is using the ping command on your Bootstrap Host. (For example: ping thundercloud.us.oracle.com)

```
$ ping <openstack_provider_url>
$ ping thundercloud.us.oracle.com
PING srv-10-75-171-2.us.oracle.com (10.75.171.2) 56(84) bytes of
data.
64 bytes from pc1011601.labs.nc.tekelec.com (10.75.171.2):
icmp_seq=1 ttl=63 time=0.283 ms
```

11. If your deployment requires a specific **Availability Zone** other than the default availability zone called nova, make the following changes in the **cluster.tfvars** file. This value can be added just after the cluster_name field. If you wish to use nova then do not add these changes.

```
az_list = [<your_availability_zone_name>"]
```

```
Example:
# Authorizing_Zone
az_list = ["foobar"]
```

12. If your deployment requires **additional** external **dns_nameservers** to be configured, make the following changes in the **cluster.tfvars** file. There can be multiple ipv4 addresses in this list. This value can be added at the end of the file after **wait_for_floatingip**. This is an optional value and the default is empty list []. The sample configuration format is as below:

```
dns_nameservers = [<ipv4_address>,...]
Example:
dns_nameservers = ["10.75.xxx.xxx",...]
```

13. If your deployment requires the delivery of metadata and user data through a config drive to each instance, add the following changes in the **cluster.tfvars** file. This is an optional value. The default is **ocne_config_drive = "false"** which indicates using a metadata server instead of a config drive for OpenStack.

 **Note:**

Make sure the OpenStack administrator did not set **force_config_drive=true** in the `/etc/nova/nova.conf` file, otherwise it will use config drive in either case. The sample configuration format is as below:

```
occne_config_drive = "<true/false>"
Example:
occne_config_drive = "true"
```

Configuration for using OCCNE LB (MetalLB)

These settings are optional and only used in conjunction with the **mb_configmap.yaml** file as explained in section [MetalLB Configuration](#). This is common to both Floating IP and Fixed IP deployments and only used if the customer is using MetalLB BGP instead of Octavia as a load balancing service.

Update the **occne_lb_flavor** field. This should be set as per the Oracle sizing charts: [vCNE VM Sizing](#). The flavor must be provided by the Openstack administrator.

```
occne_lb_flavor="<lbvm_flavor_name>"
Example:
occne_lb_flavor="OCCNE-lbvm-host"
```

Configuring Floating IPs

This procedure defines configuration that apply to floating IPs only.

1. Use the following command to retrieve the `floatingip_pool` name. This would be one of the existing provider networks.

```
$ openstack network list
+-----+-----+
+-----+-----+
| ID                                     | Name |
| Subnets                               |      |
+-----+-----+
+-----+-----+
| 1d25d5ea-77ca-4f56-b364-f53b09292e7b | ext- |
| net2                                   |     |
| f5c5ee71-8688-466d-a79f-4306e2bf3f6a |     |
| 90c160aa-2ef7-47d3-a212-e1790d56c971 | ext- |
| net-ipv6                               |     |
| 4c0b844f-1557-4454-b561-88fa31f657f3 |     |
| de033f1f-1e32-4b1f-b69a-dbf1a7a129f5 | ext- |
| net3                                   |     |
| c4c5375c-4a28-4775-8e42-b2b806514e87 |     |
| e4351e3e-81e3-4a83-bdc1-dde1296690e3 | ext- |
| net                                     |     |
| c0e0c185-ed65-4a53-a7a3-418277fb9a20 |     |
+-----+-----+
+-----+-----+
```

2. Assign the `floatingip_pool` field as follows:

```
floatip_pool = "<floating_ip_pool_name>"
Example:
floatingip_pool = ext-net2
```

3. **wait_for_floatingip** provides the ability for the Terraform deployment to poll the instance until the floating IP has been associated. It is set to true by default in the vCNE deployment. There is no need to change this.

 **Note:**

This field is a string and requires double quotes around it.

Configuring Fixed IPs

This procedure defines the configuration that applies to fixed IPs only.

Note: The fields below also includes the openstack command used to retrieve the necessary information. It is possible for a network to be defined that does not include a Name value assigned to it. It is important that the network and subnet include a Name value assigned to it.

 **Note:**

These fields can be left undefined if using the Floating IP configuration.

1. Determine and configure the network name (**fixedip_network_name**) used to provide the IPs configured in step 2 below.

```
$ openstack network list
+-----+-----+
+-----+-----+
| ID                | Name |
| Subnets          |      |
+-----+-----+
+-----+-----+
| 1d25d5ea-77ca-4f56-b364-f53b09292e7b | ext- |
net2      | f5c5ee71-8688-466d-a79f-4306e2bf3f6a |   |
| 90c160aa-2ef7-47d3-a212-e1790d56c971 | ext- |
net-ipv6  | 4c0b844f-1557-4454-b561-88fa31f657f3 |   |
| de033f1f-1e32-4b1f-b69a-dbf1a7a129f5 | ext- |
net3      | c4c5375c-4a28-4775-8e42-b2b806514e87 |   |
| e4351e3e-81e3-4a83-bdc1-dde1296690e3 | ext- |
net       | c0e0c185-ed65-4a53-a7a3-418277fb9a20 |   |
+-----+-----+
+-----+-----+
```

```
fixedip_network_name = "<fixedip_network_name>"
```

```
Example:
fixed_ip_network_name = "ext-net3"
```

2. Determine the set of fixed IPs that are to be used for the Bastion Host, K8s Nodes and DB Tier SQL Nodes.

 **Note:**

Each IP must be in double quotes and separated by commas. The number of IPs MUST match the number of the specific node type.

- a. Use the following command to retrieve the list of networks defined on the Openstack Provider.

 **Note:**

In this example, network ext-net3 is the fixed IP network selected.

```
$ openstack network list
+-----+-----+
+-----+-----+
| ID                                     | Name |
| Subnets                               |      |
+-----+-----+
+-----+-----+
| 1d25d5ea-77ca-4f56-b364-f53b09292e7b | ext-net2 | |
| f5c5ee71-8688-466d-a79f-4306e2bf3f6a |      |
| 90c160aa-2ef7-47d3-a212-e1790d56c971 | ext-net- |
| ipv6 | 4c0b844f-1557-4454-b561-88fa31f657f3 |      |
| de033f1f-1e32-4b1f-b69a-dbf1a7a129f5 | ext-net3 |
| c4c5375c-4a28-4775-8e42-b2b806514e87 |      |
| e4351e3e-81e3-4a83-bdc1-dde1296690e3 | ext-net  |
| c0e0c185-ed65-4a53-a7a3-418277fb9a20 |      |
+-----+-----+
+-----+-----+
```

- b. The user must determine the list of available IPs to use in the fixed IP lists given below. Contact your administrator for available IPs. The following command can be used by a user to retrieve the range of ports that are available in the fixed IP network. Use the subnet name from the fixed IP network to display the range.

To retrieve the number of ips in a subnet, use the following command:

```
$ openstack subnet list --network ext-net3
+-----+-----+
+-----+-----+
| ID                                     | Name |
| Network                               | Subnet |
+-----+-----+
+-----+-----+
| c4c5375c-4a28-4775-8e42-b2b806514e87 | ext-net3-subnet |
+-----+-----+
```

```
| de033f1f-1e32-4b1f-b69a-dbf1a7a129f5 | 10.75.235.0/25 |
+-----+-----+
+-----+-----+
```

- c. Select the port IPs that are to be used for the Bastion Host depending on the number of Bastion Hosts as defined in the **cluster.tfvars** file by field `number_of_bastions`.
- d. Select the port IPs that are to be used for the K8s Nodes depending on the number of Worker Nodes as defined in the **cluster.tfvars** file by field `number_of_k8s_nodes`.
- e. Select the port IPs that are to be used for the DB Tier SQL Nodes depending on the number of DB SQL Nodes as defined in the **cluster.tfvars** file by field `number_of_db_tier_sql_nodes`.
- f. Assign the IPs to the field `bastions_fixed_ip_list` in the following format (example assumes `number_of_bastions = 1`):

```
bastions_fixed_ip_list =
["<fixed_ip_0>...<fixed_ip_(number_of_bastions-1)>"]
Example:
bastions_fixed_ip_list = ["10.75.10.20"]
```

- g. Assign the IPs to the field `k8s_nodes_fixed_ip_list` in the following format (example assumes `number_of_k8s_nodes = 4`):

```
k8s_nodes_fixed_ip_list =
["<fixed_ip_0>...<fixed_ip_(number_of_k8s_nodes-1)>"]
Example:
k8s_nodes_fixed_ip_list =
["10.75.10.21", "10.75.10.22", "10.75.10.23", "10.75.10.24"]
```

- h. Assign the IPs to the field `db_tier_sql_nodes_fixed_ip_list` in the following format (example assumes the `number_of_db_tier_sql_nodes = 2`):

```
db_tier_sql_nodes_fixed_ip_list =
["<fixed_ip_0>...<fixed_ip_(number_of_db_tier_sql_nodes-1)>"]
Example:
db_tier_sql_nodes_fixed_ip_list = ["10.75.10.25", "10.75.10.26"]
```

Obtain Mate Site DB Replication Service Load Balancer IP

While installing MYSQL NDB on the second site we need to provide the Mate Site DB Replication Service Load Balancer IP as the configuration parameter for the geo-replication process to start.

Note:

If this is a single deploy and or a mated site with this being the first site deployed, this step can be skipped.

1. In order to obtain the Mate Site DB Replication Service Load Balancer IP, login to the Bastion Host of first site and execute the following command to retrieve the DB Replication Service Load Balancer IP.
2. Fetch DB Replication Service Load Balancer IP of Mate Site MYSQL NDB.

```
[cloud-user@bastion ~]$ kubectl get svc --namespace=occne-infra |
grep replication
occne-db-replication-svc      LoadBalancer    10.233.3.117
10.75.182.88      80:32496/TCP    2m8s
```

In the above example, "10.75.182.88" is the Mate Site DB Replication Service Load Balancer IP.

Deploy the OCCNE Virtualized Cluster

This section describes the procedure to deploy the VMs in the OpenStack Environment, configure the Bastion Host, and deploy and configure the Kubernetes clusters.

Deploying Using Fixed IPs

If the field `use_floating_ip` in the `cluster.tfvars` file is set to `false` (for using fixed IPs), the `deploy.sh` script performs a validation of the settings for fixed IP solution in the `tfvars` file. If the script detects any incorrect configurations within the fixed IP specific fields, it will produce an error explaining why and stop execution before the VMs are instantiated on the Openstack Provider using terraform. If `use_floating_ip = true` the fixed IP validation is not executed.

The following errors are currently detected by the verification:

- If the `use_floating_ip` setting is true or false. Anything else and that verification will fail.
- Invalid fixed IP network name (network does not exist on the Openstack Provider).
- Number of nodes of a given type does not match with the number of IP addresses in the IP list for that node type (for example, if `number_of_bastion = 1` and the fixed IP list for bastions includes 2 IPs).
- Invalid IP address (for example, 10.75.10.999)
- IP address used is currently in use.

Example verification output (assumes the IPs assigned in the `cluster.tfvars` file are already in use):

```
Example (failed case):
..Verifying use_floating_ip setting
..Verifying Fixed IP network info
....Verifying network name: ext-net3
..Verifying Fixed IP lists
....Scanning ip list: bastions_fixed_ip_list
.....ip: 10.75.235.4
          ERROR: IPv4: 10.75.235.4 from Network: ext-net3 is already in
use.
....Scanning ip list: k8s_nodes_fixed_ip_list
.....ip: 10.75.235.23
```

```

        ERROR: IPv4: 10.75.235.23 from Network: ext-net3 is already in
use.
.....ip: 10.75.235.21
        ERROR: IPv4: 10.75.235.21 from Network: ext-net3 is already in
use.
.....ip: 10.75.235.19
        ERROR: IPv4: 10.75.235.19 from Network: ext-net3 is already in
use.
.....ip: 10.75.235.27
        ERROR: IPv4: 10.75.235.27 from Network: ext-net3 is already in
use.
....Scanning ip list: db_tier_sql_nodes_fixed_ip_list
...occn1-John-Doe/cluster.tfvars verification completed.
    5 errors detected. Please address before continuing.

```

Example (pass case):

```

Verifying tfvars file...
..Verifying use_floating_ip setting
..Verifying Fixed IP network info
....Verifying network name: ext-net3
..Verifying Fixed IP lists
....Scanning ip list: bastions_fixed_ip_list
.....ip: 10.75.235.23
....Scanning ip list: k8s_nodes_fixed_ip_list
.....ip: 10.75.235.21
.....ip: 10.75.235.15
.....ip: 10.75.235.14
.....ip: 10.75.235.10
....Scanning ip list: db_tier_sql_nodes_fixed_ip_list
.....ip: 10.75.235.4
.....ip: 10.75.235.31
...occn1-John-Doe/cluster.tfvars verification successful.

```

Deploy Command Execution

[Environmental Variables](#) describes the list of possible environment variables that can be combined with the `deploy.sh` command.

1. If a Fixed IP deployment is being performed, execute the following commands from the `/var/terraform` directory prior to executing the `deploy.sh` script as given in the next step:

```

$ sed -i '/wait_for_connection/a \                                ssh -t -t -i
~/ssh/id_rsa ${OCCNE_USER}@${OCCNE_BASTION} \'sudo service docker
restart\' deploy.sh
$ sed -i '/ - \[ systemctl, restart, network \]/a
cloud_final_modules:\n - \[ scripts-user, always \]' modules/
compute/main.tf

```

2. Execute the following command from the `/var/terraform` directory on the Bootstrap Host. This command may take a while to run (can be up to 2 hours depending on the machines its run on):

```

First Site:
$ OCCNE_TFVARS_DIR=occn1_<user>

```

```
CENTRAL_REPO=<central_repo_hostname>  
CENTRAL_REPO_IP=<central_repo_ipv4_address> ./deploy.sh
```

Replica Site:

```
$ OCCNE_TFVARS_DIR=occne_<user>  
CENTRAL_REPO=<central_repo_hostname>  
CENTRAL_REPO_IP=<central_repo_ipv4_address>  
OCCNE_DB_REPLICATION_MATE_IP=<db_replication_service_load_balancer_i  
p> OCCNE_DB_REPLICATION_CLUSTER_ID=<mysqlndb_cluster_identifi  
er> ./  
deploy.sh
```

Examples:

```
$ OCCNE_TFVARS_DIR=occne_john_doe CENTRAL_REPO=winterfell  
CENTRAL_REPO_IP=10.75.216.10 ./deploy.sh  
$ OCCNE_TFVARS_DIR=occne_john_doe  
CENTRAL_REPO=winterfell CENTRAL_REPO_IP=10.75.216.10  
OCCNE_DB_REPLICATION_MATE_IP=10.75.182.88  
OCCNE_DB_REPLICATION_CLUSTER_ID=2 ./deploy.sh
```

 **Note:**

While Installing on the First Site ignore OCCNE_DB_REPLICATION configuration parameters which Provides Mate Site DB Replication Service Load Balancer IP and ID. Provide the Mate Site DB Replication Service Load Balancer IP and MySQL Cluster identifier while Installing MYSQL NDB on second site.

Change MySQL root user password

Refer to [Change MySQL root user password](#).

3

Post Installation Activities

This chapter describes the verification and security hardening procedures post installation of OCCNE.

Post Install Verification

Introduction

This document verifies installation of CNE Common services on all nodes hosting the cluster. There are different UI end points installed with common services like Kibana, Grafana, Prometheus Server, Alert Manager; below are the steps to launch different UI endpoints and verify the services are installed and working properly.

Prerequisites

1. Common services have been installed on all nodes hosting the cluster.
2. Gather list of cluster names and version tags for docker images that were used during install.
3. All cluster nodes and services pods should be up and running.
4. Commands are required to be run on Management server.
5. Any Modern browser (HTML5 compliant) with network connectivity to CNE.

Verify Kibana is Running and Accessible

1. Run the commands to get the load-balancer IP address and port number for Kibana Web Interface:

```
# LoadBalancer ip address of the kibana service is retrieved with
below command
$ export KIBANA_LOADBALANCER_IP=$(kubectl get
services occne-kibana --namespace occne-infra -o
jsonpath="{.status.loadBalancer.ingress[*].ip}")

# LoadBalancer port number of the kibana service is retrieved with
below command
$ export KIBANA_LOADBALANCER_PORT=$(kubectl get services occne-
kibana --namespace occne-infra -o jsonpath="{.spec.ports[*].port}")

# Complete url for accessing kibana in external browser
$ echo http://$KIBANA_LOADBALANCER_IP:$KIBANA_LOADBALANCER_PORT
http://10.75.182.51:80
```

2. Launch the browser and navigate to Kibana Web Interface.
The browser url to access kibana in external browser
is: **http://\$KIBANA_LOADBALANCER_IP:\$KIBANA_LOADBALANCER_PORT**

(for example: `http://10.75.182.51:80` in the example above) received in the output of the above commands.

Using Kibana verify Log and Tracer data is stored in Elasticsearch

1. Navigate to "**Management**" Tab in Kibana.
2. Click on "**Index Patterns**". You should be able to see the two patterns as below which confirms Log and Tracer data been stored in Elastic-Search successfully.
 - a. jaeger-*
 - b. logstash-*
3. Type `logstash*` in the index pattern field and wait for few seconds.
4. Verify the "**Success**" message and index pattern "`logstash-YYYY.MM.DD`" appeared as highlighted in the bottom red box. Click on "**Next step**".
5. Select "I don't want to use the Time Filter" and click on "**Create index pattern**".
6. Ensure the Web page having the indices appear in the main viewer frame.
7. Click on "**Discover**" Tab and you should be able to view raw Log records.
8. Repeat steps 3-6 using "`jaeger*`" instead of "`logstash*`" to ensure the data is stored in elastic search.

Verify Elasticsearch cluster health

1. Navigate to "**Dev Tools**" in Kibana
2. Enter the command "`GET _cluster/health`" and press on the green arrow mark. You should see the status as "green" on the right side of the screen.

Verify Prometheus Alert manager is accessible

1. Run the following commands to get the load-balancer IP address and port number for Prometheus Alert Manager Web Interface.

```
# LoadBalancer ip address of the alertmanager service is retrieved
with below command
$ export ALERTMANAGER_LOADBALANCER_IP=$(kubectl get services
occne-prometheus-alertmanager --namespace occne-infra -o
jsonpath="{.status.loadBalancer.ingress[*].ip}")

# LoadBalancer port number of the alertmanager service is retrieved
with below command
$ export ALERTMANAGER_LOADBALANCER_PORT=$(kubectl get services
occne-prometheus-alertmanager --namespace occne-infra -o
jsonpath="{.spec.ports[*].port}")

# Complete url for accessing alertmanager in external browser
$ echo
http://$ALERTMANAGER_LOADBALANCER_IP:$ALERTMANAGER_LOADBALANCER_PORT
http://10.75.182.53:80
```

2. Launch the Browser and navigate to `http://$ALERTMANAGER_LOADBALANCER_IP:$ALERTMANAGER_LOADBALA`

NCER_PORT (e.g.: http://10.75.182.53:80 in the example above) received in the output of the above commands. Ensure the AlertManager GUI is accessible.

Verify metrics are scraped and stored in prometheus server

1. Run the following commands to get the load-balancer IP address and port number for Prometheus Server Web Interface.

```
# LoadBalancer ip address of the prometheus service is retrieved
with below command
$ export PROMETHEUS_LOADBALANCER_PORT=$(kubectl get
services occne-prometheus-server --namespace occne-infra -o
jsonpath="{.spec.ports[*].port}")

# LoadBalancer port number of the prometheus service is retrieved
with below command
$ export PROMETHEUS_LOADBALANCER_IP=$(kubectl get services
occne-prometheus-server --namespace occne-infra -o
jsonpath="{.status.loadBalancer.ingress[*].ip}")

# Complete url for accessing prometheus in external browser
$ echo
http://$PROMETHEUS_LOADBALANCER_IP:$PROMETHEUS_LOADBALANCER_PORT
http://10.75.182.54:80
```

2. Launch the Browser and navigate to http://\$PROMETHEUS_LOADBALANCER_IP:\$PROMETHEUS_LOADBALANCER_PORT (e.g.: http://10.75.182.54:80 in the example above) received in the output of the above commands. Ensure the Prometheus server GUI is accessible.
3. Select **"UP"** option from **"insert metric at cursor"** drop down and click on **"Execute"** button.
4. Here the entries present under the Element section are scrape endpoints and under the value section its corresponding status(1 for up 0 for down). Ensure all the scrape endpoints have value as 1 (means up and running).

Verify Alerts are configured

1. Navigate to alerts tab of Prometheus server GUI or navigate using URL: http://\$PROMETHEUS_LOADBALANCER_IP:\$PROMETHEUS_LOADBALANCER_PORT/alerts. For <PROMETHEUS_LOADBALANCER_IP> and <PROMETHEUS_LOADBALANCER_PORT>, refer to above section.
2. If below alerts are seen in "Alerts" tab of prometheus GUI, then Alerts are configured properly.

Verify grafana is accessible and change the default password for admin user

1. Run below commands to get the load-balancer IP address and port number for Grafana Web Interface.

```
# LoadBalancer ip address of the grafana service is retrieved with
below command
$ export GRAFANA_LOADBALANCER_IP=$(kubectl get
services occne-grafana --namespace occne-infra -o
jsonpath="{.status.loadBalancer.ingress[*].ip}")
```

```
# LoadBalancer port number of the grafana service is retrieved with
below command
$ export GRAFANA_LOADBALANCER_PORT=$(kubectl get services occne-
grafana --namespace occne-infra -o jsonpath="{.spec.ports[*].port}")

# Complete url for accessing grafana in external browser
$ echo http://$GRAFANA_LOADBALANCER_IP:$GRAFANA_LOADBALANCER_PORT
http://10.75.182.55:80
```

2. Launch the Browser and navigate to `http://$GRAFANA_LOADBALANCER_IP:$GRAFANA_LOADBALANCER_PORT` (e.g.: `http://10.75.182.55:80` in the example above) received in the output of the above commands. Ensure the Prometheus server GUI is accessible. The default username and password is admin/admin for the 1st time access.
3. At first connection to the Grafana dashboard, a 'Change Password' screen will appear. Change the password to the customer provided credentials.
Note: Grafana data is not persisted, so if Grafana services restarted for some reason change password screen will appear again.
4. Grafana dashboards are accessed after the changing the default password in the above step.
5. Click on **"New dashboard"** as marked red below.
6. Click on **"Add Query"**
7. From **"Queries to"** drop down select **"Prometheus"** as data source. Presence of "Prometheus" entry in the "Queries to" drop down ensures Grafana is connected to Prometheus time series database.
8. In the Query Section marked in Red below put "sum by(__name__)
({kubernetes_namespace="occne-infra"})" and then click anywhere outside of the textbox and wait for few seconds. Ensure the dashboard appearing in the top section of the page. Example for using the metrics list, write a promQL query: `sum($metricnamefromlist)sum by(kubernetes_pod_name)($metricnamefromlist{kubernetes_namespace="occne-infra"})`. For more details about promQL, follow the [Prometheus Query Examples](#).

Post-Installation Security Hardening

Introduction

After installation, the OC-CNE system security stance should be audited prior to placing the system into service. This primarily consists of changing credentials and sequestering SSH keys to trusted servers. The following table lists all the credentials that need to be checked / changed / retained:



Note:

Refer to this section if you are performing bare metal installation.

Table 3-1 Credentials

Credential Name	Type	Associated Resource	Initial Setting	Credential Rotation
TOR Switch	username / password	Cisco Top or Rack Switch	username/ password from PreFlight Checklist	Reset post-install
Enclosure Switch	username / password	HP Enclosure Switch	username/ password from PreFlight Checklist	Reset post-install
OA Admin	username / password	HP On-board Administrator Console	username/ password from PreFlight Checklist	Reset post-install
ILO Admin	username / password	HP Integrated Lights Out Manger	username/ password from PreFlight Checklist	Reset post-install
Server Super User (root)	username / password	Server Super User	Set to well-known Oracle default during server installation	Reset post-install
Server Admin User (admusr)	username / password	Server Admin User	Set to well-known Oracle default during server installation	Reset post-install
Server Admin User SSH	SSH Key Pair	Server Admin User	Key Pair generated at install time	Can rotate keys at any time; key distribution manual procedure
MySQL Admin	username / password	MySQL Database	Set by customer during initial install	Reset post-install

If factory or Oracle defaults were used for any of these credentials, they should be changed prior to placing the system into operation. The customer should then store these credentials in a safe a secure way off site. It is recommended that the customer may plan a regular schedule for updating (rotating) these credentials.

Prerequisites

This procedure is performed after the site has been deployed and prior to placing the site into service.

Limitations and Expectations

The focus of this procedure is to secure the various credentials used or created during the install procedure. There are additional security audits that the CNE operator should perform such as scanning repositories for vulnerabilities, monitoring the system for anomalies, regularly checking security logs. These are outside the scope of this post-installation procedure.

References

1. Nexus commands to configure Top of Rack switch username and password:<https://www.cisco.com/c/en/us/td/docs/switches/datacenter/>

[nexus9000/sw/6-x/security/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_chapter_01001.html](https://www.cisco.com/c/en/us/td/docs/switches/9000/sw/6-x/security/configuration/guide/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide/b_Cisco_Nexus_9000_Series_NX-OS_Security_Configuration_Guide_chapter_01001.html)

2. HP commands to configure Enclosure switch username and password: <https://support.hpe.com/hpsc/doc/public/display?docId=c04763521>
3. HP OA commands to configure OA username and password: https://support.hpe.com/hpsc/doc/public/display?docId=emr_na-a00040582en_us&docLocale=en_US#N101C8
4. HP iLO commands to configure iLO username and password: <https://www.golinuxhub.com/2018/02/hp-ilo4--cli-guide-cheatsheet-example.html>
5. See ToR switch procedure for initial username/password configuration: [Configure Top of Rack 93180YC-EX Switches](#)
6. See procedure to configure initial iLO/OA username/password: [Configure Addresses for RMS iLOs, OA, EBIPA](#)
7. See Enclosure switch procedure for initial username/password: [Configure Enclosure Switches](#)

Procedure

1. Reset Credentials on the TOR Switch:

- a. From bastion host, login to the switch with username and password from the procedure:

```
[bastion host]# ssh <username>@<switch IP address>  
User Access Verification  
Password: <password>
```

```
Cisco Nexus Operating System (NX-OS) Software  
TAC support: http://www.cisco.com/tac  
...  
...  
<switch name>#
```

- b. Change the password for current username:

```
#  
# configure  
Enter configuration commands, one per line. End with CNTL/Z.  
(config)# username <username> password <newpassword>  
(config)#exit  
#
```

- c. Create new username:

```
#  
# configure  
Enter configuration commands, one per line. End with CNTL/Z.  
(config)# username <newusername> password <newpassword> role  
[network-operator|network-admin|vdc-admin|vdc-operator]  
(config)#exit  
#
```

- d. Exit from the switch and login with the new username and password to verify the new change works:

```
# exit
Connection to <switch IP address> closed.
[bastion host]#

[some server]# ssh <newusername>@<switch IP address>
User Access Verification
Password: <newpassword>

Cisco Nexus Operating System (NX-OS) Software
TAC support: http://www.cisco.com/tac
...
...
<switch name>#
```

- e. Delete the previous old username if it is not needed:

```
#
# configure
Enter configuration commands, one per line. End with CNTL/Z.
(config)# no username <username>
(config)#exit
#
```

- f. Change the enable secret when needed:

```
#
(config)# enable secret <newenablepassword>
(config)# exit
#
```

- g. Save the above configuration:

```
# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete.
#
```

2. Reset Credentials on the Enclosure Switch:

- a. From bastion host, login to the switch with username and password from the procedure:

```
[bastion host]# ssh <username>@<switch IP address>
<username>@<switch IP address>'s password: <password>

*****
*****
* Copyright (c) 2010-2017 Hewlett Packard Enterprise Development
LP          *
* Without the owner's prior written
consent,                                         *
```

```
* no decompiling or reverse-engineering shall be
allowed.
*
*****
*****
```

```
<switchname>
<switchname>sys
System View: return to User View with Ctrl+Z.
[switchname]
```

b. Change the password for current username:

```
[switchname]local-user <username> class <current class>
[switchname-luser-manage-<username>]password simple <newpassword>
[switchname-luser-manage-<username>]quit
[switchname]
```

c. Create new username:

```
[switchname]local-user <newusername> class [manage|network]
New local user added.
[switchname-luser-manage-<newusername>]password simple
<newpassword>
[switchname-luser-manage-<newusername>]quit
[switchname]
```

d. Exit from the switch and login with the new username and password to verify the new change works:

```
<switchname>quit
Connection to <switch IP address> closed.
[bastion host]#

[bastion host]# ssh <newusername>@<switch IP address>
<newusername>@<switch IP address>'s password: <newpassword>

*****
*****
* Copyright (c) 2010-2017 Hewlett Packard Enterprise Development
LP
*
* Without the owner's prior written
consent,
* no decompiling or reverse-engineering shall be
allowed.
*
*****
*****

<switchname>
<switchname>sys
System View: return to User View with Ctrl+Z.
[switchname]
```


- e. Delete the previous old username if it is not needed:

```
[switchname]undo local-user <username> class <current class>
```

- f. Save the above configuration:

```
[switchname]save
The current configuration will be written to the device. Are you
sure? [Y/N]:y
Please input the file name(*.cfg)[flash:/<filename>]
(To leave the existing filename unchanged, press the enter key):
flash:/<filename> exists, overwrite? [Y/N]:y
Validating file. Please wait...
Saved the current configuration to mainboard device successfully.
Slot 1:
Save next configuration file successfully.
[switchname]
```

3. Reset Credentials for the OA Admin Console:

- a. From bastion host, login to the OA with username and password from the procedure: (Note: If Standby OA, exit and login with the other OA address)

```
[bastion host]# ssh <username>@<OA address>
```

```
-----
-----
WARNING: This is a private system. Do not attempt to login
unless you are an
authorized user. Any authorized or unauthorized access and use
may be moni-
tored and can result in criminal or civil prosecution under
applicable law.
-----
-----
```

```
Firmware Version: 4.85
Built: 04/06/2018 @ 06:14
OA Bay Number: 1
OA Role: Active
<username>@<OA address>'s password:<password>
```

```
HPE BladeSystem Onboard Administrator
(C) Copyright 2006-2018 Hewlett Packard Enterprise Development LP
```

```
Type 'HELP' to display a list of valid commands.
Type 'HELP <command>' to display detailed information about a
specific command.
Type 'HELP HELP' to display more detailed information about the
help system.
```

```
OA-A45D36FD5FB1>
```

b. Change the password for current username:

```
OA-A45D36FD5FB1> set password <newpassword>

Changed password for the "<username>" user account.

OA-A45D36FD5FB1>
```

c. Add new user:

```
OA-A45D36FD5FB1> add user <newusername>

New Password: <newpassword>
Confirm      : <newpassword>
User "<newusername>" created.
You may set user privileges with the 'SET USER ACCESS' and
'ASSIGN' commands.

OA-A45D36FD5FB1> set user access <newusername> [ADMINISTRATOR|
OPERATOR|USER]

"<newusername>" has been given [administrator|operator|user]
level privileges.
```

d. Assign full access to the enclosure for the user:

```
OA-A45D36FD5FB1> assign server all <newusername>

<newusername> has been granted access to the valid requested
bay(s)

OA-A45D36FD5FB1> assign interconnect all <newusername>

<newusername> has been granted access to the valid requested
bay(s)

OA-A45D36FD5FB1> assign oa <newusername>

<newusername> has been granted access to the OA.
```

e. Exit from the OA and login with the new username and password to verify the new change works:

```
OA-A45D36FD5FB1> exit

Connection to <OA address> closed.
[bastion host]# ssh <newusername>@<OA address>
```

```
-----
-----
WARNING: This is a private system. Do not attempt to login
unless you are an
authorized user. Any authorized or unauthorized access and use
may be moni-
```

tored and can result in criminal or civil prosecution under applicable law.

```
-----  
-----  
Firmware Version: 4.85  
Built: 04/06/2018 @ 06:14  
OA Bay Number: 1  
OA Role: Active  
<newusername>@<OA address>'s password:<newpassword>
```

```
HPE BladeSystem Onboard Administrator  
(C) Copyright 2006-2018 Hewlett Packard Enterprise Development LP
```

```
Type 'HELP' to display a list of valid commands.  
Type 'HELP <command>' to display detailed information about a  
specific command.  
Type 'HELP HELP' to display more detailed information about the  
help system.
```

```
OA-A45D36FD5FB1>
```

f. Delete previous user if not needed:

```
OA-A45D36FD5FB1> remove user <username>
```

```
Entering anything other than 'YES' will result in the command  
not executing.
```

```
Are you sure you want to remove testuser1? yes
```

```
User "<username>" removed.
```

4. Reset Credentials for the iLO Admin Console:

a. From bastion host, login to the iLO with username and password from the procedure:

```
[root@winterfell ~]# ssh <username>@<iLO address>  
<username>@<iLO address>'s password: <password>  
User:<username> logged-in to ...(<iLO address> / <ipv6 address>)
```

```
iLO Advanced 2.61 at Jul 27 2018  
Server Name: <server name>  
Server Power: On
```

```
</>hpiLO->
```

b. Change current password:

```
</>hpiLO-> set /map1/accounts1/<username> password=<newpassword>

status=0
status_tag=COMMAND COMPLETED
Tue Aug 20 13:27:08 2019

</>hpiLO->
```

c. Create new user:

```
</>hpiLO-> create /map1/accounts1
username=<newusername> password=<newpassword>
group=admin,config,oemHP_rc,oemHP_power,oemHP_vm
status=0
status_tag=COMMAND COMPLETED
Tue Aug 20 13:47:56 2019

User added successfully.
```

d. Exit from the iLO and login with the new username and password to verify the new change works:

```
</>hpiLO-> exit

status=0
status_tag=COMMAND COMPLETED
Tue Aug 20 13:30:52 2019

CLI session stopped
Received disconnect from <iLO address> port 22:11: Client
Disconnect
Disconnected from <iLO address> port 22

[bastion host]# ssh <newusername>@<iLO address>
<newusername>@<iLO address>'s password: <newpassword>
User:<newusername> logged-in to ...(<iLO address> / <ipv6
address>)

iLO Advanced 2.61 at Jul 27 2018
Server Name: <server name>
Server Power: On

</>hpiLO->
```

e. Delete the previous username if not needed:

```
</>hpiLO-> delete /map1/accounts1/<username>

status=0
```

```
status_tag=COMMAND COMPLETED
Tue Aug 20 13:59:04 2019
```

```
User deleted successfully.
```

5. Reset Credentials for the root account on each server:
Login to each server in the cluster (`ssh admusr@cluster_host`) and perform the following command:

```
sudo passwd root
```

6. Reset (or Delete) Credentials for the `admusr` account on each server:
Login to each and every server in the cluster (`ssh admusr@cluster_host`) and perform the following command:

```
sudo passwd -l admusr
```

7. Regenerate / Redistribute SSH Keys Credentials for the `admusr` account:
Log into the Bastion Host VM and generate a new cluster-wide keypair by perform the following:

```
ssh-keygen -b 4096 -t rsa -C "New SSH Key" -f .ssh/new_occne_id_rsa
-q -N ""
```

Now, for each server in the cluster, perform these actions:

```
# for each cluster_host in the cluster; do
# copy the public key to the node
scp .ssh/new_occne_id_rsa.pub admusr@cluster_host:.ssh/

# install the key
ssh admusr@cluster_host "cat .ssh/new_occne_id_rsa.pub >> .ssh/
authorized_keys"
# done
```

At this point, the new key should be usable. Switch from using the old key to the new key, and confirm that each and every cluster host is still reachable. On the Bastion Host VM, perform these actions:

```
# remove the old keys from the agent (assuming you are using an
agent)
ssh-add -D
# add the new key to the agent
ssh-add .ssh/new_occne_is_rsa

# for each cluster_host in the cluster; do
# confirm access to the cluster host(s) and remove the old key
ssh admusr@cluster_host "sed -i '/ occne installer key$/d' .ssh/
authorized_keys"
# done
```

The new private key (`new_occne_id_rsa`) should also be copied to any secondary Bastion Host VM, and possibly copied off site and securely saved.

A

Artifact Acquisition and Hosting

Introduction

The OCCNE deployment containers require access to a number of resources that are usually downloaded from the internet. For cases where the target system is isolated from the internet, locally available repositories may be used. These repositories require provisioning with the proper files and versions, and some of the cluster configuration needs to be updated to allow the installation containers to locate these local repositories.

- [YUM Repository Configuration](#) is needed to hold a mirror of a number of OL7 repositories, as well as the version of docker-ce that is required by OCCNE's Kubernetes deployment
- [HTTP Repository Configuration](#) is needed to hold Kubernetes binaries and Helm charts
- [Docker Image Registry Configuration](#) is needed to hold the proper Docker images to support the containers that run Kubernetes and the common services that Kubernetes will manage
- A copy of the Oracle Linux ISO. See [Oracle Linux 7.5 Download Instructions](#) for OS installation.
- A copy of the [MySQL NDB archive](#) for database nodes.

Docker Image Registry Configuration

Introduction

To perform an installation without the system needing access to the internet, a local Docker registry must be created, and provisioned with the necessary docker images. These docker images are used to populate the Kubernetes pods once Kubernetes is installed, as well as providing the services installed during Common Services installation.

Prerequisites

Docker images for OCCNE release must be pulled to the executing system.

1. Docker images for the OCCNE release should be pulled to the executing system.
2. Docker is installed and docker commands can be run.
3. Make sure docker registry is running:

```
$ docker ps
```

4. If not then creating a local docker registry accessible by the target of the installation:

```
$ docker run -d -p <port>:<port> --restart=always --name  
<registryname> registry:2
```

(For more directions refer: <https://docs.docker.com/registry/deploying/>)

References

<https://docs.docker.com/registry/deploying/>

<https://docs.docker.com/registry/configuration/>

Provision the registry with the necessary images

On the repo server that can reach the internet AND reach the registry, populate the registry with the following images.

Run the following commands on repo server to generate bastion, k8s install, and configure dependencies.

First retrieve the docker registry image which will be used by the bastion-host to serve up docker images to the rest of the cluster:

```
$ docker pull registry:2  
$ docker tag registry:2 <registryaddress>:<port>/registry:2  
$ docker push <registryaddress>:<port>/registry:2
```

Then retrieve the lists of required docker images from each container:

```
$ docker run --rm -it -v /var/occne/<cluster>/:/host occne/  
<configure_install_image_name>:<1.6.x_tag> /getdeps/getdeps  
$ docker run --rm -it -v /var/occne/<cluster>/:/host occne/  
<k8s_install_image_name>:<1.6.x_tag> /getdeps/getdeps
```

Example:

```
$ docker run --rm -it -v /var/occne/rainbow:/:/host occne/  
configure:1.6.0 /getdeps/getdeps  
$ docker run --rm -it -v /var/occne/rainbow:/:/host occne/  
k8s_install:1.6.0 /getdeps/getdeps
```

Once the above command is successfully executed, go to `/var/occne/<cluster>/artifacts` directory and verify that there are `retrieve_docker.sh` script and `k8s_docker_images.txt` file in the directory and execute:

```
$ sh /var/occne/<cluster>/artifacts/retrieve_docker.sh  
docker.io <registryaddress>:<port> < /var/occne/<cluster>/artifacts/  
k8s_docker_images.txt
```


Once the above command is successfully executed, go to the `/var/ocne/<cluster>/artifacts` directory and verify that there are `retrieve_docker.sh` script and `config_docker_images.txt` file in the directory and execute:

```
$ sh /var/ocne/<cluster>/artifacts/retrieve_docker.sh
docker.io <registryaddress>:<port> < /var/ocne/<cluster>/artifacts/
config_docker_images.txt
```

Verify the list of repositories in the docker registry

Access endpoint `<registryaddress>:<port>/v2/_catalog` using a browser or from any linux server with `curl` command available and can query the repo server address, using `curl` command:

```
$ curl http://<registryaddress>:5000/v2/_catalog
```

Sample:

```
$ { "repositories":["coredns/coredns", "docker.elastic.co/
elasticsearch/elasticsearch-oss", "docker.elastic.co/kibana/kibana-
oss", "gcr.io/google-containers/fluentd-elasticsearch", "gcr.io/
google-containers/kube-apiserver", "gcr.io/google-containers/kube-
controller-manager", "gcr.io/google-containers/kube-proxy", "gcr.io/
google-containers/kube-scheduler", "gcr.io/google-containers/
pause", "gcr.io/google_containers/cluster-proportional-autoscaler-
amd64", "gcr.io/google_containers/metrics-server-amd64", "gcr.io/
google_containers/pause-amd64", "gcr.io/kubernetes-helm/tiller", "grafana/
grafana", "jaegertracing/jaeger-agent", "jaegertracing/jaeger-
collector", "jaegertracing/jaeger-query", "jimmidyson/configmap-
reload", "justwatch/elasticsearch_exporter", "k8s.gcr.io/addon-
resizer", "lachlanevenson/k8s-helm", "metallb/controller", "metallb/
speaker", "nginx", "prom/alertmanager", "prom/prometheus", "prom/
pushgateway", "quay.io/calico/cni", "quay.io/calico/ctl", "quay.io/calico/
kube-controllers", "quay.io/calico/node", "quay.io/coreos/etcd", "quay.io/
coreos/kube-state-metrics", "quay.io/external_storage/local-volume-
provisioner", "quay.io/jetstack/cert-manager-controller", "quay.io/pires/
docker-elasticsearch-curator", "quay.io/prometheus/node-exporter" ] }
```

Set hosts.ini variables

The [hosts.ini inventory file](#) for the cluster needs to have a few variables set in the `[ocne:vars]` section to direct the installation logic to the registry, these variables need to be set to the your docker registry configuration:

```
...
[ocne:vars]
...
ocne_private_registry=winterfell
ocne_private_registry_address='10.75.216.114'
ocne_private_registry_port=5002
ocne_helm_images_repo='winterfell:5002'
...
```

HTTP Repository Configuration

Introduction

To perform an installation without the system needing access to the internet, a local HTTP repository must be created and provisioned with the necessary files. These files are used to provide the binaries for [Kubernetes installation](#), as well as the Helm charts used during [Common Services installation](#).

Prerequisites

1. [Docker Image Registry Configuration](#) procedure should be completed before starting this procedure.
2. Docker is setup and docker commands can be run by the target system.
3. HTTP server that is reachable by the target system. Example: Running Nginx in docker container.

```
$ docker run --name mynginx1 -p <port>:<port> -d nginx
```

More information can be found out on configuring and installing Nginx using docker here: <https://docs.nginx.com/nginx/admin-guide/installing-nginx/installing-nginx-docker/>

OR

Use the html directory of Apache http server created during setting up yum mirror to perform the tasks listed below.

 **Note:**

Create new directories for kubernetes binaries and helm charts in html folder.

Procedure Steps

1. Retrieve Kubernetes Binaries:
The [Kubernetes installer](#) requires access to an HTTP server (possibly use an nginx container run via docker) from which it can download the proper version of a set of binary files. To provision an internal HTTP repository one will need to obtain these files from the internet, and place them at a known location on the internal HTTP server.

The Kubernetes Installer requires access to a HTTP server.

The following command will retrieve the proper binaries and place them in a directory named binaries under the command-line specified directory. This 'binaries' directory needs to then be placed on the HTTP server where it can be served up, with the URL identified in the clusters hosts.ini inventory file.

```
$ sh /var/ocne/<cluster>/artifacts/k8s_retrieve_bin.sh /var/www/html
```

Example:

```
$ sh /var/occne/rainbow/artifacts/k8s_retrieve_bin.sh /var/www/html
```

2. Retrieve Helm binaries and charts:

The [Configuration installer](#) requires access to an HTTP server from which it can download the proper version of a set of Helm charts for the common services. To provision an internal HTTP repository one will need to obtain these charts from the internet, and place them at a known location on the internal HTTP server using the following command.

```
$ sh /var/occne/<cluster>/artifacts/retrieve_helm.sh /var/www/html  
<helm path> < /var/occne/<cluster>/artifacts/config_helm_charts.txt
```

Example:

```
$ sh /var/occne/rainbow/artifacts/retrieve_helm.sh /var/www/html ./  
< /var/occne/rainbow/artifacts/config_helm_charts.txt
```

YUM Repository Configuration

Introduction

To perform an installation without the system needing access to the internet, a local YUM mirror must be made of a number of the OL7 repositories for use by the [OS installation](#).

A repository file will need to be created to reference this local YUM repository, and placed on the necessary machines (those which run the OCCNE installation Docker instances).

Prerequisites

1. Local YUM mirror repository for the OL7 'latest', 'epel', and 'addons' repositories. Directions here: <https://www.oracle.com/technical-resources/articles/it-infrastructure/unbreakable-linux-network.html>
2. Subscribe to following channels while creating the yum mirror from uln:

```
[ol7_x86_64_UEKR5]  
[ol7_x86_64_latest]  
[ol7_x86_64_addons]  
[ol7_x86_64_developer_EPEL]  
[ol7_x86_64_developer]
```

References

Oracle YUM mirroring directions:

<https://www.oracle.com/technetwork/articles/servers-storage-admin/yum-repo-setup-1659167.html>

Procedure Steps

1. Create OL7 repository mirror repo:

Below is an example of a repository file providing the details on a mirror with the necessary repositories. This repository file would be placed on the OCCNE Bootstrap machine that will setup the OCCNE Bastion Host. (directions on the locations in the installation procedure)

The OCCNE logic expects the exact repo names as described below: **'UEKR5'**, **'latest'**, **'addons'**, **'developer'**, and **'developer_EPEL'**

```
[UEKR5]
name=Unbreakable Enterprise Kernel Release 5 for Oracle Linux 7
(x86_64)
baseurl=http://10.75.155.195/yum/OracleLinux/OL7/UEKR5/$basearch/
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY
enabled=1
proxy=_none_

[latest]
name=Oracle Linux 7 Latest (x86_64)
baseurl=http://10.75.155.195/yum/OracleLinux/OL7/latest/$basearch/
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY
enabled=1
proxy=_none_

[addons]
name=Oracle Linux 7 Addons (x86_64)
baseurl=http://10.75.155.195/yum/OracleLinux/OL7/addons/$basearch/
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY
enabled=1
proxy=_none_

[developer]
name=Packages for creating test and development environments for
Oracle Linux 7 (x86_64)
baseurl=http://10.75.155.195/yum/OracleLinux/OL7/
developer/$basearch/
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY
enabled=1
proxy=_none_

[developer_EPEL]
name=EPEL Packages for creating
test and development environments for Oracle Linux 7 (x86_64)
baseurl=http://10.75.155.195/yum/OracleLinux/OL7/developer/
EPEL/$basearch/
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY
```

```
enabled=1  
proxy=_none_
```

B

Reference Procedures

This appendix lists the procedures which are referred in various installation procedures.

Inventory File Preparation

Introduction

OCCNE Installation automation uses information within an OCCNE Inventory file to provision servers and virtual machines, install cloud native components, as well as configure all of the components within the cluster such that they constitute a cluster conformant to the OCCNE platform specifications. To assist with the creation of the OCCNE Inventory, a boilerplate OCCNE Inventory is provided. The boilerplate inventory file requires the input of site-specific information.

This document outlines the procedure for taking the OCCNE Inventory boilerplate and creating a site specific OCCNE Inventory file usable by the OCCNE Install Procedures.

Inventory File Overview

The inventory file is an Initialization (INI) formatted file. The basic elements of an inventory file are hosts, properties, and groups.

1. A host is defined as a Fully Qualified Domain Name (FQDN). Properties are defined as key=value pairs.
2. A property applies to a specific host when it appears on the same line as the host.
3. Square brackets define group names. For example `host_hp_gen_10` defines the group of physical HP Gen10 machines. There is no explicit "end of group" delimiter, rather group definitions end at the next group declaration or the end of the file. Groups can not be nested.
4. A property applies to an entire group when it is defined under a group heading not on the same line as a host.
5. Groups of groups are formed using the `children` keyword. For example, the `occne:children` creates an `occne` group comprised of several other groups.
6. Inline comments are not allowed.

Table B-1 Base Groups

Group Name	Description/Comments
host_hp_gen_10	<p>The list of all physical hosts in the OCCNE cluster. Each host in this group must also have several properties defined (outlined below).</p> <ul style="list-style-type: none"> • <code>ansible_host</code>: The IP address for the host's teamed primary interface. The occne containers use this IP to communicate with the host • <code>ilo</code>: The IP address of the host's iLO interface. This IP is manually configured as part of the Configure Addresses for RMS iLOs, OA, EBIPA process. • <code>mac</code>: The MAC address of the host's network bootable interface. This is typically <code>eno5</code> for Gen10 RMS hardware and <code>eno1</code> for Gen10 bladed hardware. MAC addresses must use all lowercase alphanumeric values with a dash as the separator. To get the mac address, login to the above ilo address with ssh, the username and password are the <code>pxe_install_lights_out_usr</code> and <code>pxe_install_lights_out_passwd</code>, which are created in the Configure Addresses for RMS iLOs, OA, EBIPA process. After login, run command "<code>show /system1/network1/Integrated_NICs</code>", <code>Port1NIC_MACAddress</code> is for <code>eno1</code>, <code>Port5NIC_MACAddress</code> is for <code>eno5</code>. <p>The default configuration of a node in this group is for a Gen 10 RMS with modules providing boot interfaces at Linux interface identifiers '<code>eno5</code>' and '<code>eno6</code>'. For Gen 10 blades the boot interfaces are usually '<code>eno1</code>' and '<code>eno2</code>' and should be specified by adding the following properties (outlined below).</p> <ul style="list-style-type: none"> • <code>pxe_config_ks_nic</code>: The bootable interface to initiate the installation process (for Gen10 blades ='<code>eno1</code>') • <code>pxe_config_nic_list</code>: The set of interfaces to team together (for Gen10 blades ='<code>eno1,eno2</code>')

Table B-1 (Cont.) Base Groups

Group Name	Description/Comments
host_kernel_virtual	<p>The list of all virtual hosts in the OCCNE cluster. Each host in this group must have the same properties defined as above with the exception of the ilo</p> <ul style="list-style-type: none"> • <code>ansible_host</code>:The IP address for the host's primary interface. The occne containers use this interface to communicate with the host. • <code>kvm_host</code>: The physical host name.fqdn running KVM that hosts this VM host (ex. for guest db-10.icemark.lab.us.oracle.com the <code>kvm_host</code> is db-1.icemark.lab.us.oracle.com). Bastion-1 should be on db-1, bastion-2 should be on db-2. • <code>mac</code>: Always begin with 52-54-00 with the last 3 hex values being unique within the hosts.ini file (ex: <code>mac=52-54-00-c1-8e-38</code>) • <code>signal_host</code>: The Signaling network IPv4 address assigned to the MySQL NDB SQL Nodevirtual machines. The ILO network IPv4 address assigned to the Bastion host virtual machines. This IP is manually assigned and should be on the same network as the <code>host_hp_gen_10/9/8</code> node iLo address except for the final octet. For example: If the <code>kvm_host=db-2</code> and the iLo field in <code>host_hp_gen_10/9/8</code> for db-2 is set to 192.168.20.101, this value can be set to 192.168.20.201. (be sure .201 is not being used within that network by executing a ping on that address from the host node, db-1 in this example, the ping should fail). • <code>ilo_host</code>: The ILO network IPv4 address assigned to the Bastion host virtual machines. This IP is manually assigned and should be on the same network as the <code>host_hp_gen_10/9/8</code> node iLo address except for the final octet. For example: If the <code>kvm_host=db-2</code> and the iLo field in <code>host_hp_gen_10/9/8</code> for db-2 is set to 192.168.20.101, this value can be set to 192.168.20.201. (be sure .201 is not being used within that network by executing a ping on that address from the host node, db-1 in this example, the ping should fail). • <code>oam_host</code>: The OAM network IPv4 address assigned to the Bastionhost virtual machines.
kvm_hosts:children	The list of all physical hosts which will be hosting the virtual hosts. This should be the set <code>data_store</code> and <code>kube-master</code> . Do not modify.
occne:children	Do not modify the children of the occne
occne:vars	This is a list of variables representing configurable site-specific data. While some variables are optional, the ones listed in the boilerplate should be defined with valid values. If a given site does not have applicable data to fill in for a variable, the OCCNE installation or engineering team should be consulted. Individual variable values are explained in subsequent sections.
data_store	The list of Storage Hosts
kube-master	The list of Master Node hosts where kubernetes master components run.
etcd	The list of hosts that compose the etcd server. Should always be an odd number. This set is the same list of nodes as the <code>kube-master</code>

Table B-1 (Cont.) Base Groups

Group Name	Description/Comments
kube-node	The list of Worker Nodes. Worker Nodes are where kubernetes pods run and should be comprised of the bladed hosts.
k8s-cluster:children	Do not modify the children of k8s-cluster
bastion_hosts	The list of Bastion Hosts name.fqdn (ex: bastion-1.icemark.lab.us.oracle.com)

Data Tier Groups

The MySQL service is comprised of several nodes running on virtual machines on RMS hosts. This collection of hosts is referred to as the MySQL Cluster. Each host in the MySQL Cluster requires a NodeID parameter. Each host in the MySQL cluster is required to have a NodeID value that is unique per site across the MySQL cluster. Additional parameter range limitations are outlined below.

Table B-2 Data Tier Groups

Group Name	Description/Comments
mysqlndb_mgm_nodes	The list of MySQL Management nodes. In OCCNE, this group consists of three virtual machines distributed equally among the kube-masternodes. These nodes must have a NodeId parameter defined: <ul style="list-style-type: none"> • NodeId: Parameter must be unique across the MySQL Cluster and have a value between 49 and 255.
mysqlndb_data_nodes_ng0	The list of MySQL Data nodes, In OCCNE, this group consists of two virtual machine distributed equally among the Storage Hosts. Each VM in this group should belong to the different Storage Hosts. Requires a NodeId parameter. <ul style="list-style-type: none"> • NodeId: Parameter must be unique across the MySQL Cluster and have a value between 1 and 48. For example: NodeId should be assigned with value 1 and 2 [mysqlndb_data_nodes_ng0] db-6.foo.lab.us.oracle.com NodeId=1 db-7.foo.lab.us.oracle.com NodeId=2
mysqlndb_data_nodes_ng1	The list of MySQL Data nodes, In OCCNE, this group consists of two virtual machine distributed equally among the Storage Hosts. Each VM in this group should belong to the different Storage Hosts. Requires a NodeId parameter. <ul style="list-style-type: none"> • NodeId: Parameter must be unique across the MySQL Cluster and have a value between 1 and 48. For example: NodeId should be assigned with value 3 and 4 [mysqlndb_data_nodes_ng1] db-8.foo.lab.us.oracle.com NodeId=3 db-9.foo.lab.us.oracle.com NodeId=4
mysqlndb_data_nodes	The list of MySQL Data node groups. In OCCNE, this group consists of 2 groups, each groups consists of two virtual machines distributed equally among the Storage Hosts.
mysqlndb_sql_nodes	List of MySQL nodes. In OCCNE, this group consists of two virtual machines distributed equally among the Storage Hosts. Requires a NodeId parameters. <ul style="list-style-type: none"> • NodeId: Parameter must be unique across the MySQL Cluster and have a value between 49 and 255.

Table B-2 (Cont.) Data Tier Groups

Group Name	Description/Comments
mysqlndb_all_nodes:children	Do not modify the children of the mysqlndb_all_nodes group.
mysqlndb_all_nodes:vars	This is a list of variables representing configurable site-specific data. While some variables are optional, the ones listed in the boilerplate should be defined with valid values. If a given site does not have applicable data to fill in for a variable, the OCCNE installation or engineering team should be consulted. Individual variable values are explained in subsequent sections.

Prerequisites

1. Prior to initiating the procedure steps, the Inventory Boilerplate should be copied to a system where it can be edited and saved for future use. Eventually the hosts.ini file needs to be transferred to OCCNE servers.

Procedure

OCCNE Cluster Name

In order to provide each OCCNE host with a unique FQDN, the first step in composing the OCCNE Inventory is to create an OCCNE Cluster domain suffix. The OCCNE Cluster domain suffix starts with a Top-level Domain (TLD). The structure of a TLD is maintained by various government and commercial authorities. Additional domain name levels help identify the cluster and are added to help convey additional meaning. OCCNE suggests adding at least one "ad hoc" identifier and at least one "geographic" and "organizational" identifier.

Geographic and organizational identifiers may be multiple levels deep.

An example OCCNE Cluster Name using the following identifiers is below:

- Ad hoc Identifier: atlantic
- Organizational Identifier: lab1
- Organizational Identifier: research
- Geographical Identifier (State of North Carolina): nc
- Geographical Identifier (Country of United States): us
- TLD: oracle.com

Example: OCCNE Cluster name: atlantic.lab1.research.nc.us.oracle.com

Create host_hp_gen_10 and host_kernel_virtual group lists

Using the OCCNE Cluster domain suffix created above, fill out the inventory boilerplate with the list of hosts in the host_hp_gen_10 and host_kernel_virtual groups. The recommended host name prefix for nodes in the host_hp_gen_10 groups is "k8s-x" where x is a number 1 to N. Kubernetes "master" and "worker" nodes should not be differentiated using the host name. The recommended host name prefix for nodes in the host_kernel_virtual group is "db-x" where x is a number 1 to N. MySQL Cluster nodes should not be differentiated using host names.

Edit occne:vars

Edit the values in the occne:vars group to reflect site specific data. Values in the occne:vars group are defined below:

Table B-3 Edit occne:vars

Var Name	Description/Comment
occne_cluster_name	Set to the OCCNE Cluster Name generated in step 2.1 above.
subnet_ipv4	Set to the subnet of the network used to assign IPs for OCCNE hosts
subnet_cidr	Appears this is not used so does not need to be included. If it does need to be included, set to the cidr notation for the subnet. For example /24
netmask	Set appropriately for the network used to assign IPs for OCCNE hosts.
broadcast_address	Set appropriately for the network used to assign IPs for OCCNE hosts.
default_route	Set to the IP of the TOR switch.
name_server	'none'
ntp_server	Set to the IP of the TOR switch.
occne_repo_host	Set to the hostname of the bootstrap host initially. This defaults to "bootstrap". It can remain as that value or the user can change it to their own specifications but they must adhere to hostname conventions.
occne_repo_host_address	Set to the internal (ansible_host) IPv4 address of the occne_repo_host.
pxe_install_lights_out_usr	Set to the user name configured for iLO admins on each host in the OCCNE Frame.
pxe_install_lights_out_pas swd	Set to the password configured for iLO admins on each host in the OCCNE Frame.
ilo_vlan_id	Set to the VLAN ID of the ILO network For Ex: 2
ilo_subnet_ipv4	Set to the subnet of the ILO network used to assign IPs for Storage hosts
ilo_subnet_cidr	Set to the cidr notation for the subnet. For example 24
ilo_netmask	Set appropriately for the network used to assign ILO IPs for Storage hosts.
ilo_broadcast_address	Set appropriately for the network used to assign ILO IPs for OCCNE hosts.
ilo_default_route	Set to the ILO VIP of the TOR switch.
mgmt_vlan_id	Set to the VLAN ID of the Management network For example: 4
mgmt_subnet_ipv4	Set to the subnet of the Management network used to assign IPs for Storage hosts
mgmt_subnet_cidr	Set to the cidr notation for the Management subnet. For example: 29
mgmt_netmask	Set appropriately for the network used to assign Management IPs for Storage hosts.
mgmt_broadcast_address	Set appropriately for the network used to assign Management IPs for Storage hosts.
mgmt_default_route	Set to the Management VIP of the TOR switch.

Table B-3 (Cont.) Edit occne:vars

Var Name	Description/Comment
signal_vlan_id	Set to the VLAN ID of the Signaling network. For example: 5
signal_subnet_ipv4	Set to the subnet of the Signaling network used to assign IPs for Storage hosts
signal_subnet_cidr	Set to the cidr notation for the Signaling subnet. For example: 29
signal_netmask	Set appropriately for the network used to assign Signaling IPs for Storage hosts and MySQL SQL Node VM's.
signal_broadcast_address	Set appropriately for the network used to assign Signaling IPs for Storage hosts and MySQL SQL Node VM's.
signal_default_route	Set to the Signaling VIP of the TOR switch.
occne_snmp_notifier_destination	Set to the address of SNMP trap receiver. For example: "127.0.0.1:162"

Edit mysqlndb_all_nodes:vars

Table B-4 Edit mysqlndb_all_nodes:vars

Num	Var Name	Description/Comment
1	occne_mysqlndb_NoOfReplicas	Number of Replicas with in the MySQL NDB Cluster. For example: 2
2	occne_mysqlndb_DataMemory	Size of Data Memory(RAM) assigned to each MySQL Data Nodes. For example: 12G

OCCNE Inventory Boilerplate

The hosts_sample.ini file is obtained via MOS. It is delivered in the occne-config-<release_number>.tgz file.

Installation Preflight Checklist

Introduction

This procedure identifies the pre-conditions necessary to begin installation of a CNE frame. This procedure is to be referenced by field install personnel to ensure the frame is properly assembled and the inventory of needed artifacts are present before installation activities are attempted.

Prerequisites


The primary function of this procedure is to identify the prerequisites necessary for installation to begin.

Confirm Hardware Installation

Confirm hardware components are installed in the frame and connected as per the tables below

Rackmount ordering (frame not to scale)

Figure B-1 Rackmount ordering



Enclosure, ToR, and RMS Connections

OCCNE frame installation is expected to be complete prior to executing any software installation. This section provides reference to prove the frame installation is completed as expected by software installation tools.

Enclosure Switch Connections

The HP 6127XLG switch (<https://www.hpe.com/us/en/product-catalog/servers/server-interconnects/pip.hpe-6127xlg-blade-switch.8699023.html>) will have 4x10GE fiber (or DAC) connections between it and ToR respective switches' SFP+ ports.

Table B-5 Enclosure Switch Connections

Switch Port Name/ID (From)	Destination (To)	Cable Type	Module Required
Internal 1	Blade 1, NIC (1 for IObay1, 2 for IObay2)	Internal	None
Internal 2	Blade 2, NIC (1 for IObay1, 2 for IObay2)	Internal	None

Table B-5 (Cont.) Enclosure Switch Connections

Switch Port Name/ID (From)	Destination (To)	Cable Type	Module Required
Internal 3	Blade 3, NIC (1 for IObay1, 2 for IObay2)	Internal	None
Internal 4	Blade 4, NIC (1 for IObay1, 2 for IObay2)	Internal	None
Internal 5	Blade 5, NIC (1 for IObay1, 2 for IObay2)	Internal	None
Internal 6	Blade 6, NIC (1 for IObay1, 2 for IObay2)	Internal	None
Internal 7	Blade 7, NIC (1 for IObay1, 2 for IObay2)	Internal	None
Internal 8	Blade 8, NIC (1 for IObay1, 2 for IObay2)	Internal	None
Internal 9	Blade 9, NIC (1 for IObay1, 2 for IObay2)	Internal	None
Internal 10	Blade 10, NIC (1 for IObay1, 2 for IObay2)	Internal	None
Internal 11	Blade 11, NIC (1 for IObay1, 2 for IObay2)	Internal	None
Internal 12	Blade 12, NIC (1 for IObay1, 2 for IObay2)	Internal	None
Internal 13	Blade 13, NIC (1 for IObay1, 2 for IObay2)	Internal	None
Internal 14	Blade 14, NIC (1 for IObay1, 2 for IObay2)	Internal	None
Internal 15	Blade 15, NIC (1 for IObay1, 2 for IObay2)	Internal	None
Internal 16	Blade 16, NIC (1 for IObay1, 2 for IObay2)	Internal	None
External 1	Uplink 1 to ToR Switch (A for IObay1, B for IObay2)	Fiber (multi-mode)	10GE Fiber
External 2	Uplink 2 to ToR Switch (A for IObay1, B for IObay2)	Fiber (multi-mode)	10GE Fiber
External 3	Uplink 3 to ToR Switch (A for IObay1, B for IObay2)	Fiber (multi-mode)	10GE Fiber
External 4	Uplink 4 to ToR Switch (A for IObay1, B for IObay2)	Fiber (multi-mode)	10GE Fiber
External 5	Not Used	None	None
External 6	Not Used	None	None
External 7	Not Used	None	None
External 8	Not Used	None	None
Internal 17	Crosslink to IObay (2 for IObay1, 1 for IObay2)	Internal	None
Internal 18	Crosslink to IObay (2 for IObay1, 1 for IObay2)	Internal	None
Management	OA	Internal	None

ToR Switch Connections

This section contains the point to point connections for the switches. The switches in the solution will follow the naming scheme of "Switch<series number>", i.e. Switch1, Switch2, etc; where Switch1 is the first switch in the solution, and switch2 is the second. These two form a redundant pair. The switch datasheet is linked here: <https://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78-736651.html>.

The first switch in the solution will serve to connect each server's first NIC in their respective NIC pairs to the network. The next switch in the solution will serve to connect each server's redundant (2nd) NIC in their respective NIC pairs to the network.

Table B-6 ToR Switch Connections

Switch Port Name/ID (From)	From Switch 1 to Destination	From Switch 2 to Destination	Cable Type	Module Required
1	RMS 1, FLOM NIC 1	RMS 1, FLOM NIC 2	Cisco 10GE DAC	Integrated in DAC
2	RMS 1, iLO	RMS 2, iLO	CAT 5e or 6A	1GE Cu SFP
3	RMS 2, FLOM NIC 1	RMS 2, FLOM NIC 2	Cisco 10GE DAC	Integrated in DAC
4	RMS 3, FLOM NIC 1	RMS 3, FLOM NIC 2	Cisco 10GE DAC	Integrated in DAC
5	RMS 3, iLO	RMS 4, iLO	CAT 5e or 6A	1GE Cu SFP
6	RMS 4, FLOM NIC 1	RMS 4, FLOM NIC 2	Cisco 10GE DAC	Integrated in DAC
7	RMS 5, FLOM NIC 1	RMS 5, FLOM NIC 2	Cisco 10GE DAC	Integrated in DAC
8	RMS 5, iLO	RMS 6, iLO	CAT 5e or 6A	1GE Cu SFP
9	RMS 6, FLOM NIC 1	RMS 6, FLOM NIC 2	Cisco 10GE DAC	Integrated in DAC
10	RMS 7, FLOM NIC 1	RMS 7, FLOM NIC 2	Cisco 10GE DAC	Integrated in DAC
11	RMS 7, iLO	RMS 8, iLO	CAT 5e or 6A	1GE Cu SFP
12	RMS 8, FLOM NIC 1	RMS 8, FLOM NIC 2	Cisco 10GE DAC	Integrated in DAC
13	RMS 9, FLOM NIC 1	RMS 9, FLOM NIC 2	Cisco 10GE DAC	Integrated in DAC
14	RMS 9, iLO	RMS 10, iLO	CAT 5e or 6A	1GE Cu SFP
15	RMS 10, FLOM NIC 1	RMS 10, FLOM NIC 2	Cisco 10GE DAC	Integrated in DAC
16	RMS 11, FLOM NIC 1	RMS 11, FLOM NIC 2	Cisco 10GE DAC	Integrated in DAC
17	RMS 11, iLO	RMS 12, iLO	CAT 5e or 6A	1GE Cu SFP
18	RMS 12, FLOM NIC 1	RMS 12, FLOM NIC 2	Cisco 10GE DAC	Integrated in DAC
19	Enclosure 6, OA 1, Mngt	Enclosure 6, OA 2, Mngt	CAT 5e or 6A	1GE Cu SFP

Table B-6 (Cont.) ToR Switch Connections

Switch Port Name/ID (From)	From Switch 1 to Destination	From Switch 2 to Destination	Cable Type	Module Required
20	Enclosure 6, IOBay 1, Port 17	Enclosure 6, IOBay 2, Port 17	Cisco 10GE DAC	Integrated in DAC
21	Enclosure 6, IOBay 1, Port 18	Enclosure 6, IOBay 2, Port 18	Cisco 10GE DAC	Integrated in DAC
22	Enclosure 6, IOBay 1, Port 19	Enclosure 6, IOBay 2, Port 19	Cisco 10GE DAC	Integrated in DAC
23	Enclosure 6, IOBay 1, Port 20	Enclosure 6, IOBay 2, Port 20	Cisco 10GE DAC	Integrated in DAC
24	Enclosure 5, OA 1, Mngt	Enclosure 5, OA 2, Mngt	CAT 5e or 6A	1GE Cu SFP
25	Enclosure 5, IOBay 1, Port 17	Enclosure 5, IOBay 2, Port 17	Cisco 10GE DAC	Integrated in DAC
26	Enclosure 5, IOBay 1, Port 18	Enclosure 5, IOBay 2, Port 18	Cisco 10GE DAC	Integrated in DAC
27	Enclosure 5, IOBay 1, Port 19	Enclosure 5, IOBay 2, Port 19	Cisco 10GE DAC	Integrated in DAC
28	Enclosure 5, IOBay 1, Port 20	Enclosure 5, IOBay 2, Port 20	Cisco 10GE DAC	Integrated in DAC
29	Enclosure 4, OA 1, Mngt	Enclosure 4, OA 2, Mngt	CAT 5e or 6A	1GE Cu SFP
30	Enclosure 4, IOBay 1, Port 17	Enclosure 4, IOBay 2, Port 17	Cisco 10GE DAC	Integrated in DAC
31	Enclosure 4, IOBay 1, Port 18	Enclosure 4, IOBay 2, Port 18	Cisco 10GE DAC	Integrated in DAC
32	Enclosure 4, IOBay 1, Port 19	Enclosure 4, IOBay 2, Port 19	Cisco 10GE DAC	Integrated in DAC
33	Enclosure 4, IOBay 1, Port 20	Enclosure 4, IOBay 2, Port 20	Cisco 10GE DAC	Integrated in DAC
34	Enclosure 3, OA 1, Mngt	Enclosure 3, OA 2, Mngt	CAT 5e or 6A	1GE Cu SFP
35	Enclosure 3, IOBay 1, Port 17	Enclosure 3, IOBay 2, Port 17	Cisco 10GE DAC	Integrated in DAC
36	Enclosure 3, IOBay 1, Port 18	Enclosure 3, IOBay 2, Port 18	Cisco 10GE DAC	Integrated in DAC
37	Enclosure 3, IOBay 1, Port 19	Enclosure 3, IOBay 2, Port 19	Cisco 10GE DAC	Integrated in DAC
38	Enclosure 3, IOBay 1, Port 20	Enclosure 3, IOBay 2, Port 20	Cisco 10GE DAC	Integrated in DAC
39	Enclosure 2, OA 1, Mngt	Enclosure 2, OA 2, Mngt	CAT 5e or 6A	1GE Cu SFP
40	Enclosure 2, IOBay 1, Port 17	Enclosure 2, IOBay 2, Port 17	Cisco 10GE DAC	Integrated in DAC
41	Enclosure 2, IOBay 1, Port 18	Enclosure 2, IOBay 2, Port 18	Cisco 10GE DAC	Integrated in DAC

Table B-6 (Cont.) ToR Switch Connections

Switch Port Name/ID (From)	From Switch 1 to Destination	From Switch 2 to Destination	Cable Type	Module Required
42	Enclosure 2, IOBay 1, Port 19	Enclosure 2, IOBay 2, Port 19	Cisco 10GE DAC	Integrated in DAC
43	Enclosure 2, IOBay 1, Port 20	Enclosure 2, IOBay 2, Port 20	Cisco 10GE DAC	Integrated in DAC
44	Enclosure 1, OA 1, Mngt	Enclosure 1, OA 2, Mngt	CAT 5e or 6A	1GE Cu SFP
45	Enclosure 1, IOBay 1, Port 17	Enclosure 1, IOBay 2, Port 17	Cisco 10GE DAC	Integrated in DAC
46	Enclosure 1, IOBay 1, Port 18	Enclosure 1, IOBay 2, Port 18	Cisco 10GE DAC	Integrated in DAC
47	Enclosure 1, IOBay 1, Port 19	Enclosure 1, IOBay 2, Port 19	Cisco 10GE DAC	Integrated in DAC
48	Enclosure 1, IOBay 1, Port 20	Enclosure 1, IOBay 2, Port 20	Cisco 10GE DAC	Integrated in DAC
49	Mate Switch, Port 49	Mate Switch, Port 49	Cisco 40GE DAC	Integrated in DAC
50	Mate Switch, Port 50	Mate Switch, Port 50	Cisco 40GE DAC	Integrated in DAC
51	OAM Uplink to Customer	OAM Uplink to Customer	40GE (MM or SM) Fiber	40GE QSFP
52	Signaling Uplink to Customer	Signaling Uplink to Customer	40GE (MM or SM) Fiber	40GE QSFP
53	Unused	Unused		
54	Unused	Unused		
Management (Ethernet)	RMS 1, NIC 2 (1GE)	RMS 1, NIC 3 (1GE)	CAT5e or CAT 6A	None (RJ45 port)
Management (Serial)	Unused	Unused	None	None

Rackmount Server Connections

Server quickspecs can be found here: <https://h20195.www2.hp.com/v2/getdocument.aspx?docname=a00008180enw>

The HP DL380 Gen10 RMS will be configured with an iLO, a 4x1GE LOM, and a 2x10GE SFP+ FLOM.

- iLO. The integrated Lights Out management interface (iLO) contains an ethernet out of band management interface for the server. This connection is 1GE RJ45.
- 4x1GE LOM. For most servers in the solution, their 4x1GE LOM ports will be unused. The exception is the first server in the first frame. This server will serve as the management server for the ToR switches. In this case, the server will use

2 of the LOM ports to connect to ToR switches' respective out of band ethernet management ports. These connections will be 1GE RJ45 (CAT 5e or CAT 6).

- 2x10GE FLOM. Every server will be equipped with a 2x10GE Flex LOM card (or FLOM). These will be for in-band, or application and solution management traffic. These connections are 10GE fiber (or DAC) and will terminate to the ToR switches' respective SFP+ ports.

All RMS in the frame will only use the 10GE FLOM connections, except for the "management server", the first server in the frame, which will have some special connections as listed below.

Table B-7 Rackmount Server Connections

Server Interface	Destination	Cable Type	Module Required	Notes
Base NIC1 (1GE)	Unused	None	None	
Base NIC2 (1GE)	Switch1A Ethernet Mngt	CAT5e or 6a	None	Switch Initialization
Base NIC3 (1GE)	Switch1B Ethernet Mngt	CAT5e or 6a	None	Switch Initialization
Base NIC4 (1GE)	Unused	None	None	
FLOM NIC1	Switch1A Port 1	Cisco 10GE DAC	Integrated in DAC	OAM, Signaling, Cluster
FLOM NIC2	Switch1B Port 1	Cisco 10GE DAC	Integrated in DAC	OAM, Signaling, Cluster
USB Port1	USB Flash Drive	None	None	Bootstrap Host Initialization Only (temporary)
USB Port2	Keyboard	USB	None	Bootstrap Host Initialization Only (temporary)
USB Port3	Mouse	USB	None	Bootstrap Host Initialization Only (temporary)
Monitor Port	Video Monitor	DB15	None	Bootstrap Host Initialization Only (temporary)

OCCNE Artifacts

Ensure artifacts listed in the [Artifact Acquisition and Hosting](#) are available in repositories accessible from the OCCNE Frame.

Keyboard, Video, Mouse (KVM) Availability

The beginning stage of installation requires a local KVM for installing the bootstrap environment.

Procedure

Complete Site Survey Subnet Table

Table values that are prefilled are fixed in the topology and do not need to be changed. Blank values indicate that customer engagement is needed to determine the appropriate value.

Table B-8 Complete Site Survey Subnet Table

SI No.	Network Description	Subnet Allocation	Bitmask	VLAN ID	Gateway Address
1	iLO/OA Network	192.168.20.0	24	2	N/A
2	Platform Network	172.16.3.0	24	3	172.16.3.1
3	Switch Configuration Network	192.168.2.0	24	N/A	N/A
4	Management Network - Bastion Hosts		28	4	
5	Signaling Network - MySQL Replication		29	5	
6	OAM Pool - metallB pool for common services			N/A	N/A (BGP redistribution)
7	Signaling Pool - metallB pool for 5G NFs			N/A	N/A (BGP redistribution)
8	Other metallB pools (Optional)			N/A	N/A (BGP redistribution)
9	Other metallB pools (Optional)			N/A	N/A (BGP redistribution)
10	Other metallB pools (Optional)			N/A	N/A (BGP redistribution)
11	ToR Switch A OAM Uplink Subnet		30	N/A	
12	ToR Switch B OAM Uplink Subnet		30	N/A	
13	ToR Switch A Signaling Uplink Subnet		30	N/A	
14	ToR Switch B Signaling Uplink Subnet		30	N/A	
15	ToR Switch A/B Crosslink Subnet (OSPF link)	172.16.100.0	30	100	

Complete Site Survey Host IP Table

Table values that are prefilled are fixed in the topology and do not need to be changed. Blank values indicate that customer engagement is needed to determine the appropriate value.

Table B-9 Complete Site Survey Host IP Table

SI No.	Component/ Resource	Platform VLAN IP Addresses (VLAN 3)	iLO VLAN IP Addresses (VLAN 2)	CNE Management IP Addresses (VLAN 4)	Device iLO IP Addresses	MAC of Primary NIC	Notes
1	RMS 1 Host IP	172.16.3.4	192.168.20.11		192.168.20.121	Eno5:	
2	RMS 2 Host IP	172.16.3.5	192.168.20.12		192.168.20.122	Eno5:	
3	RMS 3 Host IP	172.16.3.6	N/A	N/A	192.168.20.123	Eno5:	
4	RMS 4 Host IP	172.16.3.7	N/A	N/A	192.168.20.124	Eno5:	
5	RMS 5 Host IP	172.16.3.8	N/A	N/A	192.168.20.125	Eno5:	
6	Enclosure 1 Bay 1 Host IP	172.16.3.11	N/A	N/A	192.168.20.141	Eno1:	
7	Enclosure 1 Bay 2 Host IP	172.16.3.12	N/A	N/A	192.168.20.142	Eno1:	
8	Enclosure 1 Bay 3 Host IP	172.16.3.13	N/A	N/A	192.168.20.143	Eno1:	
9	Enclosure 1 Bay 4 Host IP	172.16.3.14	N/A	N/A	192.168.20.144	Eno1:	
10	Enclosure 1 Bay 5 Host IP	172.16.3.15	N/A	N/A	192.168.20.145	Eno1:	
11	Enclosure 1 Bay 6 Host IP	172.16.3.16	N/A	N/A	192.168.20.146	Eno1:	
12	Enclosure 1 Bay 7 Host IP	172.16.3.17	N/A	N/A	192.168.20.147	Eno1:	
13	Enclosure 1 Bay 8 Host IP	172.16.3.18	N/A	N/A	192.168.20.148	Eno1:	
14	Enclosure 1 Bay 9 Host IP	172.16.3.19	N/A	N/A	192.168.20.149	Eno1:	

Table B-9 (Cont.) Complete Site Survey Host IP Table

SI No.	Component/ Resource	Platform VLAN IP Addresses (VLAN 3)	iLO VLAN IP Addresses (VLAN 2)	CNE Management IP Addresses (VLAN 4)	Device iLO IP Addresses	MAC of Primary NIC	Notes
15	Enclosure 1 Bay 10 Host IP	172.16. 3.20	N/A	N/A	192.16 8.20.15 0	Eno1:	
16	Enclosure 1 Bay 11 Host IP	172.16. 3.21	N/A	N/A	192.16 8.20.15 1	Eno1:	
17	Enclosure 1 Bay 12 Host IP	172.16. 3.22	N/A	N/A	192.16 8.20.15 2	Eno1:	
18	Enclosure 1 Bay 13 Host IP	172.16. 3.23	N/A	N/A	192.16 8.20.15 3	Eno1:	
19	Enclosure 1 Bay 14 Host IP	172.16. 3.24	N/A	N/A	192.16 8.20.15 4	Eno1:	
20	Enclosure 1 Bay 15 Host IP	172.16. 3.25	N/A	N/A	192.16 8.20.15 5	Eno1:	
21	Enclosure 1 Bay 16 Host IP	172.16. 3.26	N/A	N/A	192.16 8.20.15 6	Eno1:	

Complete VM IP Table

Table values that are prefilled are fixed in the topology and do not need to be changed. Blank values indicate that customer engagement is needed to determine the appropriate value.

Table B-10 Complete VM IP Table

SI No.	Component/ Resource	Platform VLAN IP Addresses (VLAN 3)	iLO VLAN IP Addresses (VLAN 2)	CNE Management IP Addresses (VLAN 4)	SQL Replication IP Address(VLAN 5)	Notes
1	Bastion Host 1	172.16.3 .100	192.168. 20.100		N/A	
2	Bastion Host 2	172.16.3 .101	192.168. 20.101		N/A	
3	MySQL SQL Node 1	172.16.3 .102	N/A	N/A		
4	MySQL SQL Node 2	172.16.3 .103	N/A	N/A		

Complete OA and Switch IP Table

Table values that are prefilled are fixed in the topology and do not need to be changed. Blank values indicate that customer engagement is needed to determine the appropriate value.

Table B-11 Complete OA and Switch IP Table

SI No.	Procedure Reference Variable Name	Description	IP Address	VLAN ID	Notes
1	N/A	Enclosure 1 IObay1	192.168.20.133	N/A	
2	N/A	Enclosure 1 IObay2	192.168.20.134	N/A	
3	N/A	Enclosure 1 OA1	192.168.20.131	N/A	
4	N/A	Enclosure 1 OA2	192.168.20.132	N/A	
5	ToRswitchA_Platform_IP	Host Platform Network	172.16.3.2	3	
6	ToRswitchB_Platform_IP	Host Platform Network	172.16.3.3	3	
7	ToRswitch_Platform_VIP	Host Platform Network Default Gateway	172.16.3.1	3	This address is also used as the source NTP address for all servers.
8	ToRswitchA_CNE ManagementNet_IP	Bastion Host Network		4	Address needs to be without prefix length, such as 10.25.100.2
9	ToRswitchB_CNE ManagementNet_IP	Bastion Host Network		4	Address needs to be without prefix length, such as 10.25.100.3
10	ToRswitch_CNE ManagementNet_VIP	Bastion Host Network Default Gateway		4	No prefix length, address only for VIP

Table B-11 (Cont.) Complete OA and Switch IP Table

SI No.	Procedure Reference Variable Name	Description	IP Address	VLAN ID	Notes
11	CNEManagementNet_Prefix	Bastion Host Network Prefix Length		4	number only such as 29
12	ToRswitchA_SQLreplicationNet_IP	SQL Replication Network		5	Address needs to be with prefix length, such as 10.25.200.2
13	ToRswitchB_SQLreplicationNet_IP	SQL Replication Network		5	Address needs to be with prefix length, such as 10.25.200.3
14	ToRswitch_SQLreplicationNet_VIP	SQL Replication Network Default Gateway		5	No prefix length, address only for VIP
15	SQLreplicationNet_Prefix	SQL Replication Network Prefix Length		5	number only such as 28
16	ToRswitchA_oam_uplink_customer_IP	ToR Switch A OAM uplink route path to customer network		N/A	No prefix length in address, static to be /30
17	ToRswitchA_oam_uplink_IP	ToR Switch A OAM uplink IP		N/A	No prefix length in address, static to be /30

Table B-11 (Cont.) Complete OA and Switch IP Table

SI No.	Procedure Reference Variable Name	Description	IP Address	VLAN ID	Notes
18	ToRswitchB_oam_uplink_customer_IP	ToR Switch B OAM uplink route path to customer network		N/A	No prefix length in address, static to be /30
19	ToRswitchB_oam_uplink_IP	ToR Switch B OAM uplink IP		N/A	No prefix length in address, static to be /30
20	ToRswitchA_signaling_uplink_customer_IP	ToR Switch A Signaling uplink route path to customer network		N/A	No prefix length in address, static to be /30
21	ToRswitchA_signaling_uplink_IP	ToR Switch A Signaling uplink IP		N/A	No prefix length in address, static to be /30
22	ToRswitchB_signaling_uplink_customer_IP	ToR Switch B Signaling uplink route path to customer network		N/A	No prefix length in address, static to be /30
23	ToRswitchB_signaling_uplink_IP	ToR Switch B Signaling uplink IP		N/A	No prefix length in address, static to be /30
24	ToRswitchA_mgmt_IP	ToR Switch A Out of Band Management IP	192.168.2.1	N/A	

Table B-11 (Cont.) Complete OA and Switch IP Table

SI No.	Procedure Reference Variable Name	Description	IP Address	VLAN ID	Notes
25	ToRswitchB_mngt_IP	ToR Switch A Out of Band Management IP	192.168.2.2	N/A	
26	MetallB_Signal_Subnet_With_Prefix	ToR Switch route provisioning for metallB		N/A	From Section 2.1
27	MetallB_Signal_Subnet_IP_Range	Used for mb_configmap.yaml signaling address pool			host address range from the above row subnet, exclude network and broadcast address, such as 1.1.1.1-1.1.1.14 for 1.1.1.0/28 subnet
28	MetallB_OAM_Subnet_With_Prefix	ToR Switch route provisioning for metallB		N/A	From Section 2.1

Table B-11 (Cont.) Complete OA and Switch IP Table

SI No.	Procedure Reference Variable Name	Description	IP Address	VLAN ID	Notes
29	MetaLB_OAM_Subnet_IP_Range	Used for mb_configmap.yml OAM address pool			host address range from the above row subnet, exclude network and broadcast address, such as 1.1.1.1-1.1.1.14 for 1.1.1.0/28 subnet

Table B-11 (Cont.) Complete OA and Switch IP Table

SI No.	Procedure Reference Variable Name	Description	IP Address	VLAN ID	Notes
30	Allow_Access_Server	IP address of external management server to access ToR switches			access-list Restrict_Access_ToR denied all direct external access to ToR switch vlan interfaces, in case of troubleshooting or management need to access direct access from outside, allow specific server to access. If no need, delete this line from switch configuration file. If need more than one, add similar line.
31	SNMP_Trap_Receiver_Address	IP address of the SNMP trap receiver			

Table B-11 (Cont.) Complete OA and Switch IP Table

SI No.	Procedure Reference Variable Name	Description	IP Address	VLAN ID	Notes
32	SNMP_Community_String	SNMP v2c community string			To be easy, same for snmpget and snmp traps

ToR and Enclosure Switches Variables Table (Switch Specific)

Table values that are prefilled are fixed in the topology and do not need to be changed. Blank values indicate that customer engagement is needed to determine the appropriate value.

Table B-12 ToR and Enclosure Switches Variables Table (Switch Specific)

	Key/Vairable Name	ToR_SwitchA Value	ToR_SwitchB Value	Enclosure_Switch1 Value	Enclosure_Switch2 Value	Notes
1	switch_name				N/A (This switch will assume the name of Enclosure_Switch1 after IRF is applied in configuration procedures)	Customer defined switch name for each switch.

Table B-12 (Cont.) ToR and Enclosure Switches Variables Table (Switch Specific)

	Key/Vairable Name	ToR_S witchA Value	ToR_S witchB Value	Enclosure_S witch1 Value	Enclosure_Switch2 Value	Notes
2	admin_password					Password for admin user. Strong password requirement: Length should be at least 8 characters Contain characters from at least three of the following classes: lower case letters, upper case letters, digits and special characters. No '?' as special character due to not working on switches. No '/' as special character due to the procedures.
3	user_name					Customer defined user.
4	user_password					Password for <user_name> Strong password requirement: Length should be at least 8 characters. Contain characters from at least three of the following classes: lower case letters, upper case letters, digits and special characters. No '?' as special character due to not working on switches. No '/' as special character due to the procedures.

Table B-12 (Cont.) ToR and Enclosure Switches Variables Table (Switch Specific)

	Key/Vairable Name	ToR_S witchA Value	ToR_S witchB Value	Enclos ure_S witch1 Value	Enclosure_Swit ch2 Value	Notes
5	ospf_md5_key			N/A	N/A	The key has to be same on all ospf interfaces on ToR switches and connected customer switches
6	ospf_area_id			N/A	N/A	The number as OSPF area id.
7	nxos_version			N/A	N/A	The version nxos.9.2.3.bin is used by default and hard-coded in the configuration template files. If the installed ToR switches use a different version, record the version here. The installation procedures will reference this variable and value to update a configuration template file.

Complete Site Survey Repository Location Table

Table B-13 Complete Site Survey Repository Location Table

Repository	Location Override Value
Yum Repository	
Docker Registry	
Binary Location (mysql)	
Helm Repository	

Set up the Host Inventory File (hosts.ini)

Execute the [Inventory File Preparation](#) Procedure to populate the inventory file.

Assemble 2 USB Flash Drives

Given that the bootstrap environment isn't connected to the network until the ToR switches are configured, it is necessary to provide the bootstrap environment with certain software via USB flash drives to begin the install process.

One flash drive will be used to install an OS on the Installer Bootstrap Host. The setup of this USB will be handled in a different procedure. This flash drive should have approximately 6GB capacity.

Another flash drive will be used to transfer necessary configuration files to the Installer Bootstrap Host once it has been setup with an OS. This flash drive should have approximately 6GB capacity.

Create the Utility USB

This Utility USB flash drive is used to transfer configuration and script files to the Bootstrap Host during initial installation. This USB must include enough space to accommodate all the necessary files listed below (approximately 6Gb).

Note:

- The instructions listed here are for a linux host. Instructions to do this on a PC can be obtained from the Web if needed. The mount instructions are for a Linux machine.
- When creating these files on a USB from Windows (using notepad or some other Windows editor), the files may contain control characters that are not recognized when using in a Linux environment. Usually this includes a **^M** at the end of each line. These control characters can be removed by using the dos2unix command in Linux with the file: dos2unix <filename>.
- When copying the files to this USB, make sure the USB is formatted as FAT32.

Miscellaneous Files

This procedure details any miscellaneous files that need to be copied to the Utility USB.

1. Copy the hosts.ini file from step 2.7 onto the Utility USB.
2. Copy the ol7-mirror.repo file from the customer's OL YUM mirror instance onto the Utility USB. Reference procedure: [YUM Repository Configuration](#)
3. Copy the docker-ce-stable.repo file from procedure: [YUM Repository Configuration](#) onto the Utility USB.
4. Copy the following switch configuration template files from OHC to the Utility USB:
 - a. 93180_switchA.cfg
 - b. 93180_switchB.cfg
 - c. 6127xlg_irf.cfg
 - d. ifcfg-vlan
 - e. ifcfg-bridge
5. Copy VM kickstart template file bastion_host.ks from OHC onto the Utility USB.

Copy and Edit the poap.py Script

This procedure is used to create the dhcpd.conf file that will be needed in procedure: [Configure Top of Rack 93180YC-EX Switches](#).

1. Mount the Utility USB.

 **Note:**

Instructions for mounting a USB in linux are at: [Installation of Oracle Linux 7.5 on Bootstrap Server : Install Additional Packages](#). Only follow steps 1-3 to mount the USB.

2. cd to the mounted USB directory.
3. Download the poap.py straight to the usb. The file can be obtained using the following command:

```
wget https://raw.githubusercontent.com/datacenter/nexus9000/master/nx-os/poap/poap.py
on any linux server or laptop
```

4. Rename the poap.py script to poap_nexus_script.py.

```
mv poap.py poap_nexus_script.py
```

5. The switches' firmware version is handled before the installation procedure, no need to handle it from here. Comment out the lines to handle the firmware at lines 1931-1944.

```
vi poap_nexus_script.py

# copy_system()

# if single_image is False:

#     copy_kickstart()

# signal.signal(signal.SIGTERM, sig_handler_no_exit)

# # install images

# if single_image is False:

#     install_images()

# else:

#     install_images_7_x()

# # Cleanup midway images if any

# cleanup_temp_images()
```

Create the dhcpd.conf File

This procedure is used to create the dhcpd.conf file that will be needed in procedure: [Configure Top of Rack 93180YC-EX Switches](#).

1. Edit file: dhcpd.conf.

2. Copy the following contents to that file and save it on the USB.

```
# DHCP Server Configuration file.

# see /usr/share/doc/dhcp*/dhcpd.conf.example

# see dhcpd.conf(5) man page

#

subnet 192.168.2.0 netmask 255.255.255.0 {

    range 192.168.2.101 192.168.2.102;

    default-lease-time 10800;

    max-lease-time 43200;

    allow unknown-clients;

    filename "poap_nexus_script.py";

    option domain-name-servers 192.168.2.11;

    option broadcast-address 192.168.2.255;

    option tftp-server-name "192.168.2.11";

    option routers 192.168.2.11;

    next-server 192.168.2.11;

}

subnet 192.168.20.0 netmask 255.255.255.0 {

    range 192.168.20.101 192.168.20.120;

    default-lease-time 10800;

    max-lease-time 43200;

    allow unknown-clients;

    option domain-name-servers 192.168.20.11;

    option broadcast-address 192.168.20.255;

    option tftp-server-name "192.168.20.11";

    option routers 192.168.20.11;

    next-server 192.168.20.11;
```

}

Create the md5Poap Bash Script

This procedure is used to copy the sed command to a script and copy this to the USB.

This script is needed in procedure: [Configure Top of Rack 93180YC-EX Switches](#).

1. Edit file: md5Poap.sh
2. Copy the following contents to that file and save it on the USB.

```
#!/bin/bash

f=poap_nexus_script.py ; cat $f | sed '/^#md5sum/d' > $f.md5 ;
sed -i "s/^#md5sum=.*#md5sum=\"\$(md5sum $f.md5 | sed 's/ .*//')
\"/" $f
```

Create the Bastion Host Kickstart File

This procedure is used to create the Bastion Host kickstart file. This file can be copied as is written.

The file is used in procedure: [Installation of the Bastion Host](#).

Copy the following contents to the Utility USB as bastion_host.ks.

 **Note:**

This file includes some variables that must be updated when used in procedure: [Installation of the Bastion Host](#).

 **Note:**

The steps to update those variables are contained in that procedure.

```
#version=DEVEL

# System authorization information

auth --enablesshadow --passalgo=sha512

repo --name="Server-HighAvailability" --baseurl=file:///run/install/
repo/addons/HighAvailability

repo --name="Server-ResilientStorage" --baseurl=file:///run/install/
repo/addons/ResilientStorage

# Use CDROM installation media

cdrom
```

```

# Use text mode install

text

# Run the Setup Agent on first boot

firstboot --enable

ignoredisk --only-use=sda

# Keyboard layouts

keyboard --vckeymap=us --xlayouts=''

# System language

lang en_US.UTF-8

# Network information

network --bootproto=static --device=ens3 --ip=BASTION_VLAN3_IP --
nameserver=NAMESEVERIPS --netmask=255.255.255.0 --ipv6=auto --activate

network --bootproto=static --device=ens4 --ip=BASTION_VLAN2_IP --
netmask=255.255.255.0 --ipv6=auto --activate

network --bootproto=static --device=ens5 --gateway=GATEWAYIP --
ip=BASTION_VLAN4_IP --netmask=BASTION_VLAN4_MASK --ipv6=auto --activate

network --hostname=NODEHOSTNAME

# Root password

rootpw --
iscrypted $6$etqyspJhPUG440VO$0FqnB.agxmnDqb.Bh0sSLhq7..t37RwUZr7S1VmIBv
MmWVoUjb2DJJ2f4VlrW9RdfVi.IDXxd2/Eeo41FCCJ01

# System services

services --enabled="chronyd"

# Do not configure the X Window System

skipx

# System timezone

timezone Etc/GMT --isUtc --ntpservers=NTPSERVERIPS

user --groups=wheel --name=admusr --
password=$6$etqyspJhPUG440VO$0FqnB.agxmnDqb.Bh0sSLhq7..t37RwUZr7S1VmIBvM

```

```
mWVoUjb2DJJ2f4VlrW9RdfVi.IDXxd2/Eeo41FCCJ01 --iscrypted --gecos="admusr"

# System bootloader configuration

bootloader --append=" crashkernel=auto" --location=mbr --boot-drive=sda

#autopart --type=lvm

# Partition clearing information

clearpart --all --initlabel --drives=sda

# Disk partitioning information

part /boot --fstype="xfs" --ondisk=sda --size=1024

part pv.11 --size 1 --grow --ondisk=sda

volgroup ol pv.11

logvol / --fstype="xfs" --size=20480 --name=root --vgname=ol

logvol /var --fstype="xfs" --size=1 --grow --name=var --vgname=ol

%packages

@^minimal

@compat-libraries

@base

@core

@debugging

@development

chrony

kexec-tools

%end

%addon com_redhat_kdump --enable --reserve-mb='auto'

%end
```

```

%anaconda

pwpolicy root --minlen=6 --minquality=1 --notstrict --nochanges --
notempty

pwpolicy user --minlen=6 --minquality=1 --notstrict --nochanges --
emptyok

pwpolicy luks --minlen=6 --minquality=1 --notstrict --nochanges --
notempty

%end

%post --log=/root/ocne-ks.log

echo "===== Running Post Configuration
===== "

# Set shell editor to vi

echo set -o vi >> /etc/profile.d/sh.local

# selinux set to permissive

setenforce permissive

sed -i 's/SELINUX=enforcing/SELINUX=permissive/g' /etc/selinux/config

# Set sudo to nopassword

sed --in-place 's/^#\s*(%wheel\s\+ALL=(ALL)\s\+NOPASSWD:\s\
+ALL\)\1/' /etc/sudoers

echo "proxy=HTTP_PROXY" >> /etc/yum.conf

# Configure keys for admusr

mkdir -m0700 /home/admusr/.ssh/

chown admusr:admusr /home/admusr/.ssh

```

```
cat <<EOF >/home/admusr/.ssh/authorized_keys

PUBLIC_KEY

EOF

echo "Configuring SSH..."

cp /etc/ssh/sshd_config /etc/ssh/sshd_config.orig && \
sed -i 's/#Protocol 2/Protocol 2/' /etc/ssh/sshd_config && \
sed -i 's/#LogLevel.*LogLevel INFO/' /etc/ssh/sshd_config && \
sed -i 's/X11Forwarding yes/X11Forwarding no/' /etc/ssh/sshd_config && \
sed -i 's/#MaxAuthTries.*MaxAuthTries 4/' /etc/ssh/sshd_config && \
sed -i 's/#IgnoreRhosts.*IgnoreRhosts yes/' /etc/ssh/sshd_config

if [ `grep HostBasedAuthentication /etc/ssh/sshd_config | wc -l` -lt
1 ]; then

    echo 'HostBasedAuthentication no' >> /etc/ssh/sshd_config

fi

sed -i 's/#PermitRootLogin.*PermitRootLogin no/' /etc/ssh/sshd_config
&& \

sed -i 's/PermitRootLogin.*PermitRootLogin no/' /etc/ssh/sshd_config
&& \

sed -i 's/#PermitEmptyPasswords.*PermitEmptyPasswords no/' /etc/ssh/
sshd_config && \

sed -i 's/#PermitUserEnvironment.*PermitUserEnvironment no/' /etc/ssh/
sshd_config && \

sed -i 's/PermitUserEnvironment.*PermitUserEnvironment no/' /etc/ssh/
sshd_config
```

```

if [ `grep -i 'Ciphers aes128-ctr,aes192-ctr,aes256-ctr' /etc/ssh/
sshd_config | wc -l` -lt 1 ]; then

    echo 'Ciphers aes128-ctr,aes192-ctr,aes256-ctr' >> /etc/ssh/
sshd_config

    if [ $? -ne 0 ]; then

        echo " ERROR: echo 1 failed"

    fi

fi

if [ `grep '^MACs' /etc/ssh/sshd_config | wc -l` -lt 1 ]; then

    echo 'MACs hmac-sha2-512-etm@openssh.com,hmac-sha2-256-
etm@openssh.com,umac-128-etm@openssh.com,hmac-sha2-512,hmac-
sha2-256,umac-128@openssh.com' >> /etc/ssh/sshd_config

    if [ $? -ne 0 ]; then

        echo " ERROR: echo 2 failed"

    fi

fi

sed -i 's/#ClientAliveInterval.*/ClientAliveInterval 300/' /etc/ssh/
sshd_config

sed -i 's/#ClientAliveCountMax.*/ClientAliveCountMax 0/' /etc/ssh/
sshd_config

sed -i 's/#Banner.*/Banner \/etc\/issue.net/' /etc/ssh/sshd_config

egrep -q "^(\\s*)LoginGraceTime\\s+\\S+(\\s*#.*)?\\s*$" /etc/ssh/sshd_config
&& sed -ri "s/^(\\s*)LoginGraceTime\\s+\\S+(\\s*#.*)?\\s*$\\/\\1LoginGraceTime
60\\2/" /etc/ssh/sshd_config || echo "LoginGraceTime 60" >> /etc/ssh/
sshd_config

echo 'This site is for the exclusive use of Oracle and its authorized
customers and partners. Use of this site by customers and partners is
subject to the Terms of Use and Privacy Policy for this site, as well
as your contract with Oracle. Use of this site by Oracle employees is
subject to company policies, including the Code of Conduct.
Unauthorized access or breach of these terms may result in termination

```

```
of your authorization to use this site and/or civil and criminal
penalties.' > /etc/issue
```

```
echo 'This site is for the exclusive use of Oracle and its authorized
customers and partners. Use of this site by customers and partners is
subject to the Terms of Use and Privacy Policy for this site, as well
as your contract with Oracle. Use of this site by Oracle employees is
subject to company policies, including the Code of Conduct.
Unauthorized access or breach of these terms may result in termination
of your authorization to use this site and/or civil and criminal
penalties.' > /etc/issue.net
```

```
%end
```

```
reboot
```

Inventory File Template

The `host.ini` file contains the inventory used by the various OCCNE deployment containers that will instantiate the OCCNE cluster.

Template example

The inventory is composed of multiple groups (indicated by bracketed strings):

- `local`: OCCNE ansible use. Do not modify.
- `occne`: list of servers in the OCCNE cluster that will be installed by the `os_install` container.
- `k8s-cluster`: list of servers in the kubernetes cluster.
- `kube-master`: list of servers that will be provisioned as kubernetes master nodes by the `k8s_install` container.
- `kube-node`: list of servers that will be provisioned as kubernetes worker nodes by the `k8s_install` container.
- `etcd`: list of servers that will be provisioned as part of kubernetes etcd cluster by the `k8s_install` container.
- `data_store`: list of servers that will be host the VMs of the MySQL database cluster, `os_install` container will install `kvm` on them.
- `occne:vars`: list of `occne` environment variables. Values for variables are required. See below for description.

OCCNE Variables

Table B-14 OCCNE Variables

Var Name	Description/Comment
<code>occne_cluster_name</code>	Set to the OCCNE Cluster Name.

Table B-14 (Cont.) OCCNE Variables

Var Name	Description/Comment
subnet_ipv4	Set to the subnet of the network used to assign IPs for OCCNE hosts
subnet_cidr	Appears this is not used so does not need to be included. If it does need to be included, set to the cidr notation for the subnet. For example /24
netmask	Set appropriately for the network used to assign IPs for OCCNE hosts.
broadcast_address	Set appropriately for the network used to assign IPs for OCCNE hosts.
default_route	Set to the IP of the TOR switch.
name_server	'none'
ntp_server	Set to the IP of the TOR switch.
occne_repo_host	Set to the hostname of the bootstrap host initially. This defaults to "bootstrap". It can remain as that value or the user can change it to their own specifications but they must adhere to hostname conventions.
occne_repo_host_address	Set to the internal (ansible_host) IPv4 address of the occne_repo_host.
pxe_install_lights_out_usr	Set to the user name configured for iLO admins on each host in the OCCNE Frame.
pxe_install_lights_out_passwd	Set to the password configured for iLO admins on each host in the OCCNE Frame.
ilo_vlan_id	Set to the VLAN ID of the ILO network. For example: 2
ilo_subnet_ipv4	Set to the subnet of the ILO network used to assign IPs for Storage hosts
ilo_subnet_cidr	Set to the cidr notation for the subnet. For example: 24
ilo_netmask	Set appropriately for the network used to assign ILO IPs for Storage hosts.
ilo_broadcast_address	Set appropriately for the network used to assign ILO IPs for OCCNE hosts.
ilo_default_route	Set to the ILO VIP of the TOR switch.
mgmt_vlan_id	Set to the VLAN ID of the Management network. For example: 4
mgmt_subnet_ipv4	Set to the subnet of the Management network used to assign IPs for Storage hosts.
mgmt_subnet_cidr	Set to the cidr notation for the Management subnet. For example: 29
mgmt_netmask	Set appropriately for the network used to assign Management IPs for Storage hosts.
mgmt_broadcast_address	Set appropriately for the network used to assign Management IPs for Storage hosts.
mgmt_default_route	Set to the Management VIP of the TOR switch.
signal_vlan_id	Set to the VLAN ID of the Signaling network. For example: 5
signal_subnet_ipv4	Set to the subnet of the Signaling network used to assign IPs for Storage hosts
signal_subnet_cidr	Set to the cidr notation for the Signaling subnet. For example: 29
signal_netmask	Set appropriately for the network used to assign Signaling IPs for Storage hosts and MySQL SQL Node VM's.

Table B-14 (Cont.) OCCNE Variables

Var Name	Description/Comment
signal_broadcast_address	Set appropriately for the network used to assign Signaling IPs for Storage hosts and MySQL SQL Node VM's.
signal_default_route	Set to the Signaling VIP of the TOR switch.
occne_snmp_notifier_destination	Set to the address of SNMP trap receiver. For example: "127.0.0.1:162"

Install Additional Services/Network Functions

This assumes the service has docker images located on a docker registry that is reachable by the cluster's bastion, and associated helm charts located at a URL also accessible by the bastion.

Run the following commands from the cluster bastion:

1. Copy docker images needed for the service into the bastion-host docker registry:

- a. Create a file `docker_images.txt` listing the required docker images and tags

```
dockerRepo:5000/<serviceNameImage>
```

Example:

```
dockerRepo:5000/serviceNameImage2:1.2.2
```

- b. Load these images into the cluster-local docker registry by running:

```
$ /var/occne/cluster/<cluster>/artifacts/retrieve_docker.sh <<  
docker_images.txt
```

2. Copy helm charts needed for the service into the bastion-host helm chart repository:

- a. Create a file `helm_charts.txt` listing the required helm charts and versions:

```
helmRepoName/chart_name <version>
```

- b. Add the source helm repository to the cluster-local helm configuration (this need only be done once for each repo):

```
$ helm repo add helmRepoName https://someUrl.oracle.com
```

- c. Load the chart(s) into the cluster-local helm chart repository by running:

```
$ /var/occne/cluster/<cluster>/artifacts/retrieve_helm.sh <<  
helm_charts.txt
```

3. Install the service:

Create a `values.yaml` file on the Bastion Host that contains the values needed by the Helm chart

To install the service run:

```
$ helm install --name <release-name> --namespace <service-namespace> -f values.yaml <chart_name>
```

Change MySQL root user password

Following is the procedure to change MySQL root user password.

As part of the installation of the MySQL Cluster, db_install container generates the random password and marked as expired in the MySQL SQL nodes. This password is stored in "/var/occnedb/mysql_d_expired.log" file. Login to the each of the MySQL nodes and change the MySQL root user password.

1. Login to MySQL Node VM.
2. Login to mysql client as a root user:

```
$ sudo su  
$ mysql -h 127.0.0.1 -uroot -p
```

3. Enter expired random password stored in "/var/occnedb/mysql_d_expired.log" file.

```
$ mysql -h 127.0.0.1 -uroot -p  
$ Enter password:
```

4. Execute the following commands to change the root password:

```
$ mysql> ALTER USER 'root'@'localhost' IDENTIFIED BY  
'<NEW_PASSWORD>';  
$ mysql> FLUSH PRIVILEGES;
```

Note:

Here 'NEW_PASSWORD' is the password of the mysql root user.

MySQL Repository Requirements

MySQL Cluster Manager is a distributed client/server application consisting of two main components. The MySQL Cluster Manager agent is a set of one or more agent processes that manage NDB Cluster nodes, and the MySQL Cluster Manager client provides a command-line interface to the agent's management functions.

In OCCNE MySQL Cluster Manager 1.4.7 binary distributions that include MySQL NDB Cluster will be used for installing MySQL Cluster Manager 1.4.7 and MySQL NDB Cluster 7.6.8. The complete MySQL NDB Cluster 7.6.8 binary distribution is included in this below software.

MySQL Cluster Binaries


Below binary is used for installation of MySQL Cluster Manager along with the MySQL NDB Cluster, This binary distributions includes MySQL Cluster Manager 1.4.7 and

MySQL NDB Cluster 7.6.8 in it. This software will be downloaded from the Oracle Software Delivery Cloud (OSDC) i.e. <https://edelivery.oracle.com>.

Download MySQL Cluster Manager

1. Login/Access <https://edelivery.oracle.com> Oracle Software Delivery Cloud (OSDC) page, to download MySQL Cluster Manager 1.4.7+Cluster TAR for Oracle Linux / RHEL 7 x86
2. Enter "MySQL Cluster Carrier Grade Edition" and click on **Search**.
3. "DLP:MySQL Cluster Carrier Grade Edition 0.0.0.0.0 (MySQL Cluster Carrier Grade Edition)" is listed, click on **Add to Cart**.
4. Click on **Checkout**, following page will be displayed, deselect "Selected Software".
5. Select MySQL Cluster Manager 1.4.7 and Select Platform as "Linux x86-64", Click on **Continue**.
6. Accept the licence agreement and click on **Continue**.
7. Select "MySQL Cluster Manager 1.4.7+Cluster TAR for Oracle Linux / RHEL 7 x86 (64bit)" as shown below and Click on **Download**.

This will install download manager and then provide the path where to download, download manager will download the MySQL Cluster Manager 1.4.7+Cluster software.



8. View the download progress in Download manager, once download is completed, MySQL Cluster Manager software (V980756-01.zip) is downloaded. Once download is completed, V980756-01.zip file is used to install MySQL Cluster Manager 1.4.7(MCM) and MySQL NDB Cluster 7.6.8.

Oracle Linux 7.5 Download Instructions

The procedure to download Oracle Linux 7.5 is explained below:

1. Login to <https://edelivery.oracle.com> Oracle Software Delivery Cloud (OSDC) page, to download Oracle Linux 7.5 ISO.
2. Enter "Oracle Linux 7.5" and click on **Search**

"DLP: Oracle Linux 7.5.0.0.0 (Oracle Linux) " will be listed as shown below, click on **Add to Cart**.

3. Click on **Checkout**, following page will be displayed.
4. Select "Oracle Linux 7.5.0.0.0" and Select Platform as "x86-64 bit", Click on "**Continue**".
5. Accept the licence agreement and click on "Continue".
6. Select "Oracle Linux Release 7 Update 5 for x86 (64 bit), 4.1 GB" as shown below and Click on "**Download**". V975367-01.iso Oracle Linux Release 7 Update 5 for x86 (64 bit), 4.1 GB

This will install download manager and then provide the path where to download, download manager will download the Oracle Linux 7.5 software.

Oracle Linux 7.5.0.0.0	
Oracle Linux 7.5.0.0.0 for x86 64 bit	
<input type="checkbox"/> V975332-01.zip	Readme for Driver Update Disk, 1.1 KB
<input type="checkbox"/> V975333-01.iso	Driver Update Disk for Oracle Linux 7 x86_64, 3.8 MB
<input type="checkbox"/> V975334-01.zip	Oracle Container Services for use with Kubernetes 1.1.9.1, 633.2 MB
<input type="checkbox"/> V975335-01.iso	Oracle VirtIO Drivers Version for Microsoft Windows 1.1.2, 59.2 MB
<input type="checkbox"/> V975336-01.zip	Oracle Container Services for use with Kubernetes 1.1.8, 713.1 MB
<input type="checkbox"/> V975363-01.iso	Oracle Linux Release 7 Update 5 Boot ISO image for x86 (64 bit), 540.0 MB
<input type="checkbox"/> V975364-01.iso	Oracle Linux Release 7 Update 5 UEK Boot ISO image for x86 (64 bit), 566.0 MB
<input type="checkbox"/> V975365-01.iso	Oracle Linux Release 7 Update 5 source DVD 1, 3.4 GB
<input type="checkbox"/> V975366-01.iso	Oracle Linux Release 7 Update 5 source DVD 2, 4.1 GB
<input checked="" type="checkbox"/> V975367-01.iso	Oracle Linux Release 7 Update 5 for x86 (64 bit), 4.1 GB

Total 10 distinct files Total Size 14.1 GB

7. View the download progress in Download manager, once download is completed, Oracle Linux 7.5(V975367-01.iso) is downloaded. Once the download is complete the Oracle Linux 7.5(V975367-01.iso) is used for installing host servers and VM creation.

vCNE VM Sizing

For a virtualized CNE, the VM can be sized to host each node in the CNE so that the resources used by each node closely matches the expected workload. This page gives recommendations on VM sizes for each node type. Note that these are sizing *guidelines*. Customers do not have to use these exact sizes, although creating VMs that are smaller than the minimum recommended sizes may result in a CNE that performs poorly.

Bootstrap VM

The Bootstrap host has fixed sizing that it needs to bootstrap the installation of OC-CNE.

Table B-15 Bootstrap VM

VM name	vCPUs	RAM	DISK	Comments
Bootstrap host	2	8 GB	40 GB	

Kubernetes VMs**Master nodes**

GCE and AWS have established consistent sizing guidelines for master node VMs. OC-CNE follows these generally accepted guidelines. Use the "medium" size unless you have more than 10 worker nodes.

Table B-16 Kubernetes Master Node

VM name	vCPUs	RAM	DISK	Comments
K8s Master - small	1	3.75 GB	20 GB	For K8s clusters with 1-5 worker nodes
K8s Master - medium	2	7.5 GB	20 GB	For K8s clusters with 6-10 worker nodes
K8s Master - large	4	15 GB	20GB	For K8s clusters with 11-100 worker nodes

Worker nodes

Use the "medium" worker node size for deployments with 2 or less 5G NFs. Else choose the "large" size. Lab deployments may use the "small" size if they have resource limitations.

Table B-17 Kubernetes Worker Node

VM Name	vCPUs	RAM	DISK	Comments
K8s worker - small	4	15 GB	20 GB	
K8s worker - medium	8	30 GB	20 GB	
K8s worker - large	16	60 GB	20 GB	
K8s worker - extra large	32	120 GB	20 GB	

 **Note:**

Above mentioned values are suggested for worker node size. Actual size is determined only after testing the environment.

Bastion host VMs

Bastion hosts are expected to have light, occasional, workloads, with few persistent processes.

Table B-18 Bastion Host VMs

VM name	vCPUs	RAM	DISK	Comments
Bastion host	1	8 Gb	100 GB	

DB Tier VMs

For bare metal, VMs are created for the DB Tier (via libvirt). For vCNE, use the VM sizes shown below:

Management nodes

Table B-19 Management nodes

VM Name	vCPUs	RAM	DISK	Comments
DB mgmt	1	3.75 GB	20 GB	

Data nodes

NDB data nodes store data tables in memory. They should have more memory than the average VM.

Table B-20 Data Nodes

VM Name	vCPUs	RAM	DISK	Comments
DB data	8	60 GB	100 GB	

SQL nodes

SQL nodes need enough CPU to parse large SQL queries. They only store data in memory while an SQL transaction is in progress.

Table B-21 SQL Nodes

VM Name	vCPUs	RAM	DISK	Comments
DB SQL	8	30 GB	50 GB	



Note:

In case of "high CPU" profile, VM uses a profile with less than 1GB per vCPU.

Environmental Variables

The following table describes the list of possible environment variables that can be combined with the *deploy.sh* command to further define the execution of the deployment.

 **Note:**

Those marked **Y** under *Required (Y/N)* column are necessary but the defaults can be used if they meet the user's requirements for deployment.

Table B-22 Environmental Variables

Environment Variable	Definition	Default Value	Required (Y/N)
OCCNE_PIPELINE_ARGS	Additional parameters to the installation process		N
CENTRAL_REPO_DOCKER_PORT	Central Repository Docker Port	5000	Y
CENTRAL_REPO	Central Repository Hostname	winterfell	Y
CENTRAL_REPO_IP	Central Repository IPv4 Address	10.75.216.10	Y
OCCNE_DB_REPLICATION_CLUSTER_ID	DB Replication cluster ID. Set this for the second site replica installation.		N
OCCNE_DB_REPLICATION_MATE_IP	DB Replication mate IP. Set this for the second site replica installation.		N
OCCNE_PREFIX	Development time prefix for OCCNE image names		N
OS_LBAAS_USE_OCTAVIA	Flag to use Octavia load-balancer in OpenStack ("true"/"false")	true	N
OCCNE_VERIFY_TFVARS	Instructs the deploy.sh script to execute verification of the cluster.tfvars file <ul style="list-style-type: none"> 0 = Skip tfvars verification 1 = Run tfvars verification 	1	N
OS_AUTH_URL	OpenStack authorization URL (Should be set by the OpenStack RC file)	(Set by .rc file)	Y
OS_CINDER_AZ	OpenStack Cinder storage volume availability zone (Should be set on bootstrap host if OpenStack Cinder availability zone needs to be different from default zone 'nova')		N
OS_USER_DOMAIN_NAME	OpenStack domain name (Should be set by the OpenStack RC file)	(Set by .rc file)	Y
OS_LBAAS_FLOATING_NETWORK_ID	OpenStack floating network ID for allocating floating IP addresses (Should be automatically set by "external_net" content from tfvars file) (Currently applicable only to Octavia)	(Automatically derived)	N

Table B-22 (Cont.) Environmental Variables

Environment Variable	Definition	Default Value	Required (Y/N)
OS_LBAAS_ENABLED	OpenStack Load-Balancing as a service ("true"/"false") (Currently applicable only to Octavia)	true	N
OS_LBAAS_METHOD	OpenStack load-balancing method (Currently applicable only to Octavia)	ROUND_ROBIN	N
OS_PASSWORD	OpenStack password for the account for deployment (Should be set by the OpenStack RC file)	(Set by .rc file)	Y
OS_PROJECT_ID	OpenStack project ID (Should be set by the OpenStack RC file)	(Set by .rc file)	Y
OS_REGION_NAME	OpenStack region name (Should be set by the OpenStack RC file)	(Set by .rc file)	Y
OS_LBAAS_SUBNET_ID	OpenStack subnet ID for load-balancer (Should be automatically set by using cluster private subnet) (Currently applicable only to Octavia)	(Automatically derived)	N
OS_USERNAME	OpenStack user name account for deployment (Should be set by the OpenStack RC file)	(Set by .rc file)	Y
OCCNE_TFVARS_DIR	Provides the path to the clusters.tfvars file in reference to the current directory.	NA	Y
OCCNE_VERSION	Used to define the version of the container images used during deployment	Defaults to current release	Y
OCCNE_USER	User account for accessing OpenStack VM instances	cloud-user	Y