

Oracle® Communications

Oracle Communications Signaling, Cloud Native Environment (OC-CNE) Installation Guide



Release 1.4
F24135-01
March 2020

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

ORACLE®

F24135-01

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 What's New in This Guide

2 Introduction

Glossary	2-1
Key terms	2-1
Key Acronyms and Abbreviations	2-2
How to use this document	2-4
Documentation Admonishments	2-5
Locate Product Documentation on the Oracle Help Center Site	2-5
Customer Training	2-5
My Oracle Support	2-6
Emergency Response	2-6

3 Installation Procedures

Bare Metal Installation	3-1
OCCNE Installation Overview	3-1
Frame and Component Overview	3-1
Create OCCNE Instance	3-5
Installation Prerequisites	3-7
Configure Artifact Acquisition and Hosting	3-8
Initial Configuration - Prepare a Minimal Boot Strapping Environment	3-10
Installation of Oracle Linux 7.5 on Bootstrap Host	3-10
Configure the Installer Bootstrap Host BIOS	3-19
Configure Top of Rack 93180YC-EX Switches	3-25
Configure Addresses for RMS iLOs, OA, EBIPA	3-45
Configure Legacy BIOS on Remaining Hosts	3-56
Configure Enclosure Switches	3-64
Bastion Host Installation	3-73
Automated Installation	3-83
Virtualized CNE Installation	3-89
Prerequisites	3-89

Limitation/Expectations	3-89
Upload an Image to an OpenStack Environment	3-89
Bootstrap Host Creation	3-90
Pre-deployment Configuration	3-92
Deploy the OCCNE Virtualized Cluster	3-104

4 Post Installation Activities

Post Install Verification	4-1
Post-Installation Security Hardening	4-8

A Artifact Acquisition and Hosting

YUM Repository Configuration	A-1
HTTP Repository Configuration	A-4
Docker Image Registry Configuration	A-6

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-39
Oracle Linux 7.5 Download Instructions	B-40

List of Figures

2-1	Example of a Procedure Steps Used in This Document	2-4
3-1	Frame Overview	3-2
3-2	Host Designations	3-3
3-3	Node Roles	3-4
3-4	Transient Roles	3-5
3-5	OCCNE Installation Overview	3-6
B-1	Rackmount ordering	B-8

List of Tables

2-1	Key Terms	2-1
2-2	Key Acronyms and Abbreviations	2-2
2-3	Admonishments	2-5
3-1	Oracle eDelivery Artifact Acquisition	3-8
3-2	Procedure to configure MetalLB pools and peers	3-9
3-3	Bootstrap Install Procedure	3-12
3-4	Procedure to configure the Installer Bootstrap Host BIOS	3-20
3-5	Procedure to configure Top of Rack 93180YC-EX Switches	3-26
3-6	Procedure to verify Top of Rack 93180YC-EX Switches	3-38
3-7	Procedure to configure SNMP Trap	3-43
3-8	Procedure to configure Addresses for RMS iLOs, OA, EBIPA	3-46
3-9	Procedure to configure the Legacy BIOS on Remaining Hosts	3-57
3-10	Procedure to configure enclosure switches	3-65
3-11	Terminology used in Procedure	3-73
3-12	Bastion Installation	3-75
3-13	Automated Installation	3-84
3-14	Upload an Image to an OpenStack Environment	3-90
3-15	Bootstrap Host Creation	3-91
3-16	Pre-deployment Configuration	3-93
3-17	Deploy the OCCNE Virtualized Cluster	3-105
4-1	OCCNE Post Install Verification	4-2
4-2	Credentials	4-8
4-3	Post-Installation Security Hardening	4-11
A-1	Procedure to configure OCCNE YUM Repository	A-3
A-2	Steps to configure OCCNE HTTP Repository	A-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	Install Additional Services/Network Functions	B-37
B-16	Change MySQL root user password	B-38

1

What's New in This Guide

New and Updated features in Release 1.4

OC-CNE has been updated with the following enhancements in the version 1.4.0:

- **Continuous Deployment**
OC-CNE 1.4 includes a Continuous Deployment (CD) pipeline to automate OC-CNE lifecycle operations.
- **Automated Upgrade**
OC-CNE 1.4 offers a fully automated upgrade from OC-CNE 1.3. Host/guest OS, Kubernetes, and common services are all upgraded via a single CD pipeline command.
- **Automated Replication Configuration for DB Tier**
Prior to OC-CNE 1.4, configuring cross-site DB replication was a manual process. OCCNE 1.4 automates the initial configuration of cross-site replication.

2

Introduction

This document details the procedure for installing an **Oracle Communications Signaling, Network Function Cloud Native Environment**, referred to in these installation procedures simply as *OCCNE*. The intended audiences for this document are Oracle engineers who work with customers to install a **Cloud Native Environment (CNE)** on-site at customer facilities.

Glossary

Key terms

This table below lists terms used in this document.

Table 2-1 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 provide 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, these will be HP Gen10 servers.
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 1.0, 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 2-1 (Cont.) Key Terms

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

This table below lists abbreviations, and acronyms specific to this document.

Table 2-2 Key Acronyms and Abbreviations

Acronym/Abbreviation/ Term	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
EBIPA	Enclosure Bay IP Addressing
FQDN	Fully Qualified Domain name

Table 2-2 (Cont.) Key Acronyms and Abbreviations

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.
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

Table 2-2 (Cont.) Key Acronyms and Abbreviations

URL	Uniform Resource Locator
VM	Virtual Machine
vCNE	Virtualized CNE
VSP	Virtual Serial Port
YUM	Yellowdog Updator, Modified (a Linux Package Manager)

How to use this document

Although this document is primarily to be used as an initial installation guide, its secondary purpose is to be used as a reference for Disaster Recovery procedures.

When executing this document for either purpose, there are a few points which help to ensure that the user understands the author’s intent. These points are as follows:

1. Before beginning a procedure, completely read the instructional text (it will appear immediately after the Section heading for each procedure) and all associated procedural WARNINGS or NOTES.
2. Before execution of a STEP within a procedure, completely read the left and right columns including any STEP specific WARNINGS or NOTES.

If a procedural STEP fails to execute successfully, STOP and contact Oracle’s Customer Service for assistance before attempting to continue. [My Oracle Support](#) for information on contacting Oracle Customer Support.

Figure 2-1 Example of a Procedure Steps Used in This Document

Each step has a checkbox the user should check to keep track of the progress of the procedure.

The Title column describes the operations to perform during that step.




Each command the user enters, and any response output, is formatted in 10-point Courier font.

	Title	Directive/Result Step
1. <input type="checkbox"/>	Change directory	Change to the backout directory. <pre>\$ cd /var/TKLC/backout</pre>
2. <input type="checkbox"/>	ServerX: Connect to the console of the server	Establish a connection to the server using cu on the terminal server/console. <pre>\$ cu -l /dev/ttyS7</pre>
3. <input type="checkbox"/>	Verify Network Element data	View the Network Elements configuration data; verify the data; save and print report. 3. Select Configuration > Network Elements to view Network Elements Configuration screen.

Documentation Admonishments

Admonishments are icons and text throughout this manual that alert the reader to assure personal safety, to minimize possible service interruptions, and to warn of the potential for equipment damage.

Table 2-3 Admonishments

Icon	Description
 DANGER	Danger: (This icon and text indicate the possibility of personal injury.)
 WARNING	Warning: (This icon and text indicate the possibility of equipment damage.)
 CAUTION	Caution: (This icon and text indicate the possibility of service interruption.)

Locate Product Documentation on the Oracle Help Center Site

Oracle Communications customer documentation is available on the web at the Oracle Help Center site, <http://docs.oracle.com>. You do not have to register to access these documents. Viewing these files requires Adobe Acrobat Reader, which can be downloaded at <http://www.adobe.com>.

1. Access the Oracle Help Center site at <http://docs.oracle.com>.
2. Click **Industries**.
3. Under the Oracle Communications subheading, click **Oracle Communications documentation** link.
The Communications Documentation page displays.
4. Click on your product and then the release number.
A list of the documentation set for the selected product and release displays.
5. To download a file to your location, right-click the **PDF** link, select **Save target as** (or similar command based on your browser), and save to a local folder.

Customer Training

Oracle University offers training for service providers and enterprises. Visit our web site to view, and register for, Oracle Communications training at <http://education.oracle.com/communication>.

To obtain contact phone numbers for countries or regions, visit the Oracle University Education web site at www.oracle.com/education/contacts.

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.

Emergency Response

In the event of a critical service situation, emergency response is offered by the Customer Access Support (CAS) main number at 1-800-223-1711 (toll-free in the US), or by calling the Oracle Support hotline for your local country from the list at <http://www.oracle.com/us/support/contact/index.html>. The emergency response provides immediate coverage, automatic escalation, and other features to ensure that the critical situation is resolved as rapidly as possible.

A critical situation is defined as a problem with the installed equipment that severely affects service, traffic, or maintenance capabilities, and requires immediate corrective action. Critical situations affect service and/or system operation resulting in one or several of these situations:

- A total system failure that results in loss of all transaction processing capability
- Significant reduction in system capacity or traffic handling capability
- Loss of the system's ability to perform automatic system reconfiguration
- Inability to restart a processor or the system
- Corruption of system databases that requires service affecting corrective actions
- Loss of access for maintenance or recovery operations
- Loss of the system ability to provide any required critical or major trouble notification

Any other problem severely affecting service, capacity/traffic, billing, and maintenance capabilities may be defined as critical by prior discussion and agreement with Oracle.

3

Installation Procedures

The installation procedures in this document provision and configure an Oracle Communications Signaling, Network Function 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:

In case any procedure requires Linux Shell access, make sure that the shell is 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 describes 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 HP c-Class enclosure (BL460c blade servers), 5 DL380 rack mount servers, and 2 Top of Rack (ToR) Cisco switches. In case any procedure requires Linux Shell access, make sure that the shell is with Keepalive to avoid unexpected timeout.


Figure 3-1 Frame Overview



Host Designations

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


Figure 3-2 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 3-3 Node Roles



Transient Roles

Transient role is unique in that it has OOB connections to the ToR switches, which brings 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 3-4 Transient Roles



Create OCCNE Instance

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

Figure 3-5 OCCNE Installation Overview



The following is an overview or basic install flow for reference to understand the overall effort contained within these procedures:

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 at installing 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, 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 and execute the following command to retrieve DB Replication Service Load Balancer IP
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 following artifacts require download from eDelivery and/or OHC.

Table 3-1 Oracle eDelivery Artifact Acquisition

Artifact	Description	File Type	Destination Repository
occne-images-1.3.x.tgz	OCCNE Installers (Docker images)	Tar GZ	Docker Registry
v980756-01.zip	Zip file of MySQL Cluster Manager 1.4.7+Cluster	Zip of tar file	File repository MySQL Repository Requirements
v975367-01.iso	OL7 ISO	ISO	File repository Oracle Linux 7.5 Download Instructions
Install Docs	These Install Procedures from OHC	PDFs	N/A
Templates	Switch config files, hosts.ini file templates from OHC	Config files (.conf, .ini)	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/occne/<cluster_name>*) as the hosts.ini file.

Table 3-2 Procedure to configure MetalLB pools and peers

Step #	Procedure	Description
1. <input type="checkbox"/>	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. <input type="checkbox"/>	Edit the mb_configmap.yaml file	<p>Edit the mb_configmap.yaml file with the site-specific values found in the Preflight Checklist</p> <p>Note: The name "signaling" is prone to different spellings (UK vs US), therefore pay special attention to how this signaling pool is referenced.</p> <pre> 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>' </pre>

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](#).
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 in 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

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 RMS2. 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.x 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

Table 3-3 Bootstrap Install Procedure

Step #	Procedure	Description
1. <input type="checkbox"/>	Create Bootable USB Media	<ol style="list-style-type: none">1. Download the Oracle Linux Download the Oracle Linux ISO from OHC onto a user accessible location (eg. Installer's notebook). The exact details on how to perform this step is specific to the users equipment).2. 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. <pre>\$ dd -if=<path to ISO> -of=<USB device path> -bs=1m</pre> Example (assuming the USB is on /dev/sdf and the iso file is at /var/ocne) <pre>\$ dd -if=/var/ocne/OracleLinux-7.5-x86_64-disc1.iso -of=/dev/sdf -bs=1m</pre>

Table 3-3 (Cont.) Bootstrap Install Procedure

Step #	Procedure	Description
2. <input type="checkbox"/>	Install OL7 on the Installer Bootstrap Host.	<ol style="list-style-type: none"> <li data-bbox="862 352 1466 436">1. Connect a Keyboard, Video, and Mouse (KVM) into the Installer Bootstrap Host's monitor and USB ports. <li data-bbox="862 457 1466 541">2. Plug the USB flash drive containing the bootable iso into an available USB port on the Bootstrap host (usually in the front panel). <li data-bbox="862 562 1466 772">3. 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. <li data-bbox="862 793 1466 982">4. 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. <li data-bbox="862 1003 1466 1276">5. The host attempts to boot from the USB. 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. <li data-bbox="862 1297 1466 1927">6. 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. <ol style="list-style-type: none"> <li data-bbox="911 1430 1466 1486">a. LANGUAGE SUPPORT: English (United States) <li data-bbox="911 1507 1466 1535">b. KEYBOARD: English (US) <li data-bbox="911 1556 1466 1583">c. INSTALLATION SOURCE: Local Media <li data-bbox="911 1604 1466 1631">d. SOFTWARE SELECTION: Minimal Install <p>INSTALLATION DESTINATION should display No disks selected. Select INSTALLATION DESTINATION to indicate the drive to install the OS on.</p> <p>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.</p> <p>If the server has already been installed a red banner at the bottom of the page may indicate</p>

Table 3-3 (Cont.) Bootstrap Install Procedure

Step #	Procedure	Description
		<p>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.</p> <p>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.</p> <p>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.</p> <p>7. Select DONE. This returns to the INSTALLATION SUMMARY page.</p> <p>8. At the INSTALLATION SUMMARY screen, select Begin Installation. The CONFIGURATION screen is displayed.</p> <p>9. At the CONFIGURATION screen, select ROOT PASSWORD.</p> <p>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.</p> <p>10. 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.</p>

Table 3-3 (Cont.) Bootstrap Install Procedure

Step #	Procedure	Description
3. <input type="checkbox"/>	Install Additional Packages.	<p>Additional packages are needed to complete the installation and move on to the next step in the overall procedure. These additional packages are available within the OL install media on the USB. To install these packages, a YUM repo file is configured to use the install media. The additional packages to install are:</p> <ul style="list-style-type: none"> • dnsmasq • dhcp • xinetd • tftp-server • dos2unix • nfs-utils <ol style="list-style-type: none"> 1. Login with the root user and password configured above. 2. Create the mount directory: \$ mkdir /media/usb 3. Insert the USB into an available USB port (usually the front USB port) of the Installer Bootstrap Host. 4. Find and mount the USB partition. Typically the USB device is enumerated as /dev/sda but that is not always the case. Use the lsblk command to find the USB device. An example lsblk output is below. The capacity of the USB drive is expected to be approximately 30GiB, therefore the USB drive is enumerated as device /dev/sda in the example below: <pre>\$ lsblk sdd 8:48 0 894.3G 0 disk sde 8:64 0 1.7T 0 disk sdc 8:32 0 894.3G 0 disk sdc2 8:34 0 1G 0 part / boot sdc3 8:35 0 893.1G 0 part ol-swap 252:1 0 4G 0 lvm [SWAP] ol-home 252:2 0 839.1G 0 lvm /home ol-root 252:0 0 50G 0 lvm / sdc1 8:33 0 200M 0 part / boot/efi sda 8:0 1 29.3G 0 disk sda2 8:2 1 8.5M 0 part sda1 8:1 1 4.3G 0 part</pre> <p>The dmesg command also provides information about how the operating system enumerates devices. In the example below, the dmesg output indicates the USB drive is enumerated as device /dev/sda.</p>

Table 3-3 (Cont.) Bootstrap Install Procedure

Step #	Procedure	Description
		<p>Note: The output is shortened here for display purposes.</p> <pre> \$ dmesg ... [8850.211757] usb-storage 2-6:1.0: USB Mass Storage device detected [8850.212078] scsi host1: usb-storage 2-6:1.0 [8851.231690] scsi 1:0:0:0: Direct- Access SanDisk Cruzer Glide 1.00 PQ: 0 ANSI: 6 [8851.232524] sd 1:0:0:0: Attached scsi generic sg0 type 0 [8851.232978] sd 1:0:0:0: [sda] 61341696 512-byte logical blocks: (31.4 GB/29.3 GiB) [8851.234598] sd 1:0:0:0: [sda] Write Protect is off [8851.234600] sd 1:0:0:0: [sda] Mode Sense: 43 00 00 00 [8851.234862] sd 1:0:0:0: [sda] Write cache: disabled, read cache: enabled, doesn't support DPO or FUA [8851.255300] sda: sda1 sda2 ... The USB device should contain at least two partitions. One is the boot partition and the other is the install media. The install media is the larger of the two partitions. To find information about the partitions use the fdisk command to list the filesystems on the USB device. Use the device name discovered via the steps outlined above. In the examples above, the USB device is /dev/sda. \$ fdisk -l /dev/sda Disk /dev/sda: 31.4 GB, 31406948352 bytes, 61341696 sectors Units = sectors of 1 * 512 = 512 bytes Sector size (logical/physical): 512 bytes / 512 bytes I/O size (minimum/optimal): 512 bytes / 512 bytes Disk label type: dos Disk identifier: 0x137202cf Device Boot Start </pre>

Table 3-3 (Cont.) Bootstrap Install Procedure

Step #	Procedure	Description
		<pre> End Blocks Id System /dev/sda1 * 0 8929279 4464640 0 Empty /dev/sda2 3076 20503 8714 ef EFI (FAT-12/16/32) </pre> <p>In the example output above, the <code>/dev/sda2</code> partition is the EFI boot partition. Therefore the install media files are on <code>/dev/sda1</code>. Use the <code>mount</code> command to mount the install media file system. The same command without any options is used to verify the device is mounted to <code>/media/usb</code>.</p> <pre> \$ mount /dev/sda1 /media/usb \$ mount ... /dev/sda1 on /media/usb type iso9660 (ro,relatime,nojoliet,check=s,map=n,blocksize=2048) </pre> <p>5. Create a <code>yum</code> config file to install packages from local install media. Create a repo file <code>/etc/yum.repos.d/Media.repo</code> with the following information:</p> <pre> [ol7_base_media] name=Oracle Linux 7 Base Media baseurl=file:///media/usb gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-oracle gpgcheck=1 enabled=1 </pre> <p>6. Disable the default public yum repo. This is done by renaming the current <code>.repo</code> file to end with something other than <code>.repo</code>. Adding <code>.disabled</code> to the end of the file name is standard.</p> <p>Note: This can be left in this state as the Installer Bootstrap Host is re-paved in a later procedure.</p> <pre> \$ mv /etc/yum.repos.d/public-yum-ol7.repo /etc/yum.repos.d/public-yum-ol7.repo.disabled </pre> <p>7. Use the <code>yum repolist</code> command to check the repository configuration.</p>

Table 3-3 (Cont.) Bootstrap Install Procedure

Step #	Procedure	Description
		<p>The output of <code>yum repolist</code> should look like the example below. Verify there no errors regarding unreachable yum repos.</p> <pre>\$ yum repolist Loaded plugins: langpacks, ulninfo repo id repo name status ol7_base_media Oracle Linux 7 Base Media 5,134 repolist: 5,134</pre> <p>8. Use yum to install the additional packages from the USB repo.</p> <pre>\$ yum install dnsmasq \$ yum install dhcp \$ yum install xinetd \$ yum install tftp-server \$ yum install dos2unix \$ yum install nfs-utils</pre> <p>9. Verify installation of dhcp, xinetd, and tftp-server. Note: Currently dnsmasq is not being used. The verification of tftp makes sure the tftp file is included in the <code>/etc/xinetd.d</code> directory. Installation/Verification does not include actually starting any of the services. Service configuration/starting is performed in a later procedure.</p> <p>Verify dhcp is installed: -----</p> <pre>\$ cd /etc/dhcp \$ ls dhclient.d dhclient-exit-hooks.d dhcpd6.conf dhcpd.conf scripts</pre> <p>Verify xinetd is installed: -----</p> <pre>\$ cd /etc/xinetd.d \$ ls chargen-dgram chargen-stream daytime- dgram daytime-stream discard-dgram discard-stream echo-dgram echo-stream tcpmux-server</pre>

Table 3-3 (Cont.) Bootstrap Install Procedure

Step #	Procedure	Description
		<pre>time-dgram time-stream Verify tftp is installed: ----- \$ cd /etc/xinetd.d \$ ls chargen-dgram chargen-stream daytime- dgram daytime-stream discard-dgram discard-stream echo-dgram echo-stream tcpmux-server tftp time-dgram time-stream</pre> <p>10. Unmount the USB and remove the USB from the host. The mount command can be used to verify the USB is no longer mounted to /media/usb.</p> <pre>\$ umount /media/usb \$ mount Verify that /dev/sda1 is no longer shown as mounted to /media/usb.</pre>

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

Table 3-4 Procedure to configure the Installer Bootstrap Host BIOS

Step #	Procedure	Description
1. <input type="checkbox"/>	Expose the System Configuration Utility	<p>This procedure 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.</p> <ol style="list-style-type: none"> 1. 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. 2. 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. 3. The System Utilities screen is exposed in the remote console.
2. <input type="checkbox"/>	Change over from UEFI Booting Mode to Legacy BIOS Booting Mode	<p>Should the System Utility default the booting mode to UEFI or has been changed to UEFI, it will be necessary to switch the booting mode to Legacy.</p> <ol style="list-style-type: none"> 1. Expose the System Configuration Utility by following Step 1. 2. Select System Configuration. 3. Select BIOS/Platform Configuration (RBSU). 4. 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. 5. 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. 6. After the reboot and the user re-enters the System Utility, the Boot Options page should appear. 7. 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.

Table 3-4 (Cont.) Procedure to configure the Installer Bootstrap Host BIOS

Step #	Procedure	Description
3. <input type="checkbox"/>	Adding a New User Account	<p>This procedure provides the steps required to add a new user account to the server iLO 5 interface.</p> <p>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.</p> <ol style="list-style-type: none">1. Expose the System Utility by following Step 1.2. Select System Configuration.3. Select iLO 5 Configuration Utility.4. Select User Management, and then Add User.5. 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.6. 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.

Table 3-4 (Cont.) Procedure to configure the Installer Bootstrap Host BIOS

Step #	Procedure	Description
4. <input type="checkbox"/>	Force PXE to boot from the first Embedded FlexibleLOM HPE Ethernet 10Gb 2-port Adapter	<p>During host PXE, the DHCP DISCOVER requests from the hosts must be broadcast over the 10Gb port. This procedure 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 procedure.</p> <ol style="list-style-type: none"> 1. Expose the System Utility by following Step 1. 2. Select System Configuration. 3. Select BIOS/Platform Configuration (RBSU). 4. 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. 5. 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. 6. Move the 10 Gb Embedded FlexibleLOM 1 Port 1 entry up above the 1Gb Embedded LOM 1 Port 1 entry. 7. 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.

Table 3-4 (Cont.) Procedure to configure the Installer Bootstrap Host BIOS

Step #	Procedure	Description
5. <input type="checkbox"/>	Enabling Virtualization	<p>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 Workload Profiles.</p> <ol style="list-style-type: none"> 1. Verifying Default Settings <ol style="list-style-type: none"> a. Expose the System Configuration Utility by following Step 1. b. Select System Configuration. c. Select BIOS/Platform Configuration (RBSU) d. 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. e. 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. <input type="checkbox"/>	Disable RAID Configurations	<ol style="list-style-type: none"> 1. Expose the System Configuration Utility by following Step 1. 2. Select System Configuration. 3. Select Embedded RAID 1 : HPE Smart Array P408i-a SR Gen 10. 4. Select Array Configuration. 5. Select Manage Arrays. 6. Select Array A (or any designated Array Configuration if there are more than one). 7. Select Delete Array. 8. Select Submit Changes. 9. 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.

Table 3-4 (Cont.) Procedure to configure the Installer Bootstrap Host BIOS

Step #	Procedure	Description
7. <input type="checkbox"/>	Enable the Primary Boot Device	<p>This procedure 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.</p> <ol style="list-style-type: none"> 1. Expose the System Configuration Utility by following Step 1. 2. Select System Configuration. 3. Select Embedded RAID 1 : HPE Smart Array P408i-a SR Gen 10. 4. 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. 5. Select Select Bootable Physical Drive. 6. 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. 7. Select Set as Primary Bootable Device. 8. 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. 9. 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.

Table 3-4 (Cont.) Procedure to configure the Installer Bootstrap Host BIOS

Step #	Procedure	Description
8. <input type="checkbox"/>	Configure the iLO 5 Static IP Address	<p>When configuring the Bootstrap host, the static IP address for the iLO 5 must be configured.</p> <p>Note: This procedure requires a reboot after completion.</p> <ol style="list-style-type: none"> 1. Expose the System Configuration Utility by following Step 1. 2. Select System Configuration. 3. Select iLO 5 Configuration Utility. 4. Select Network Options. 5. Enter the IP Address, Subnet Mask, and Gateway IP Address fields provided in Installation PreFlight Checklist. 6. 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. 7. 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.

- 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>

Procedure Configuration Procedures

Table 3-5 Procedure to configure Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
1. <input type="checkbox"/>	Login to the Bootstrap host as root.	Using the KVM, login to the Bootstrap host as root. Note: All instructions in this procedure are executed from the Bootstrap Host.
2. <input type="checkbox"/>	Insert and mount the Utility USB	Insert and mount the Utility USB that contains the configuration and script files. Verify the files are listed in the USB using the <code>ls /media/usb</code> 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.

Table 3-5 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
3. <input type="checkbox"/>	Create bridge interface	<p>Create bridge interface to connect both management ports and setup the management bridge to support switch initialization.</p> <p>Note: <CNE_Management_IP_With_Prefix> is from Installation PreFlight Checklist : Complete Site Survey Host IP Table. Row 1 CNE Management IP Addresses (VLAN 4) column.</p> <p><ToRswitch_CNEManagementNet_VIP> is from Installation PreFlight Checklist : Complete OA and Switch IP Table.</p> <pre> \$ 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 </pre>

Table 3-5 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
4. <input type="checkbox"/>	Edit the <code>/etc/xinetd.d/tftp</code> file	<p>Edit the <code>/etc/xinetd.d/tftp</code> file to enable TFTP service. Change the disable option to no, if it is set to yes.</p> <pre>\$ 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 }</pre>
5. <input type="checkbox"/>	Enable tftp on the Bootstrap host.	<pre>\$ systemctl start tftp \$ systemctl enable tftp</pre> <p>Verify tftp is active and enabled:</p> <pre>\$ systemctl status tftp \$ ps -elf grep tftp</pre>
6. <input type="checkbox"/>	Copy the <code>dhcpd.conf</code> file	<p>Copy the <code>dhcpd.conf</code> file from the Utility USB in Installation PreFlight checklist : Create the dhcpd.conf File to the <code>/etc/dhcp/</code> directory.</p> <pre>\$ cp /media/usb/dhcpd.conf /etc/dhcp/</pre>

Table 3-5 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
7. <input type="checkbox"/>	Restart and enable dhcpd service.	<pre># /bin/systemctl restart dhcpd.service # /bin/systemctl enable dhcpd.service</pre> <p>Use the systemctl status dhcpd command to verify active and enabled.</p> <pre># systemctl status dhcpd</pre>
8. <input type="checkbox"/>	Copy the switch configuration and script files	<p>Copy the switch configuration and script files from the Utility USB to directory <i>/var/lib/tftpboot/</i>.</p> <pre>\$ 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/.</pre>
9. <input type="checkbox"/>	Modify POAP script File.	<p>Modify POAP script File. Change Username and password credentials used to login to the Bootstrap host.</p> <pre># 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", }</pre> <p>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.</p>
10. <input type="checkbox"/>	Modify POAP script file	<p>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.</p> <pre># cd /var/lib/tftpboot/ # /bin/bash md5Poap.sh</pre>

Table 3-5 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches


Step #	Procedure	Description
11. <input type="checkbox"/>	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.

Table 3-5 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
12. <input type="checkbox"/>	Copy the /var/lib/tftpboot/93180_switchA.cfg into a file called /var/lib/tftpboot/conf.<switchA serial number>	<p>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.</p> <pre># 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}#<Met alLB_Signal_Subnet_With_Prefix>#g'</pre>

Table 3-5 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
		<pre> 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}#<CNEManagementNe t_Prefix>#g' conf.<switchA serial number> # sed -i 's#{SQL_replication_SwA_Address}#<ToRswitch A_SQLreplicationNet_IP>#g' conf.<switchA serial number> # sed -i 's#{SQL_replication_SwB_Address}#<ToRswitch B_SQLreplicationNet_IP>#g' conf.<switchA serial number> # sed -i 's#{SQL_replication_Prefix}#<SQLreplication Net_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> </pre>

Table 3-5 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
		<pre># 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></pre>

Table 3-5 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
13. <input type="checkbox"/>	Copy the /var/lib/tftpboot/93180_switchB.cfg into a file called /var/lib/tftpboot/conf.<switchB serial number>	<p>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.</p> <p>These values are contained at Installation PreFlight checklist : ToR and Enclosure Switches Variables Table and Installation PreFlight Checklist : Complete OA and Switch IP Table.</p> <pre># 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></pre> <p>Note: different delimiter is used in next two commands due to '/' sign in in the variables</p> <pre># sed -i 's#{ALLOW_5G_XSI_LIST_WITH_PREFIX_LEN}#<Met</pre>

Table 3-5 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
		<pre> alLB_Signal_Subnet_With_Prefix>#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}#<CNEManagementNe t_Prefix>#g' conf.<switchB serial number> # sed -i 's#{SQL_replication_SwA_Address}#<ToRswitch A_SQLreplicationNet_IP>#g' conf.<switchB serial number> # sed -i 's#{SQL_replication_SwB_Address}#<ToRswitch B_SQLreplicationNet_IP>#g' conf.<switchB serial number> # sed -i 's#{SQL_replication_Prefix}#<SQLreplication Net_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' </pre>

Table 3-5 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
		<pre> 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> </pre>
14. <input type="checkbox"/>	Generate the md5 checksum	<p>Generate the md5 checksum for each conf file in <code>/var/lib/tftpboot</code> and copy that into a new file called conf.<switchA/B serial number>.md5.</p> <pre> \$ md5sum conf.<switchA serial number> > conf.<switchA serial number>.md5 \$ md5sum conf.<switchB serial number> > conf.<switchB serial number>.md5 </pre>

Table 3-5 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
15. <input type="checkbox"/>	Verify the <code>/var/lib/tftpboot</code> directory has the correct files.	<p>Make sure the file permissions are set as given below.</p> <p>Note: The ToR switches are constantly attempting to find and execute the <code>poap_nexus_script.py</code> script which uses <code>tftp</code> to load and install the configuration files.</p> <pre># 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</pre>
16. <input type="checkbox"/>	Disable firewalld.	<pre>\$ systemctl stop firewalld \$ systemctl disable firewalld</pre> <p>To verify:</p> <pre>\$ systemctl status firewalld</pre> <p>Once this is complete, the ToR Switches will attempt to boot from the <code>tftpboot</code> files automatically. Eventually the verification steps can be executed below. It may take about 5 minutes for this to complete.</p>
17. <input type="checkbox"/>	Un-mount the Utility USB	Un-mount the Utility USB and remove it: <code>umount /media/usb</code>

Verification

Table 3-6 Procedure to verify Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
1. <input type="checkbox"/>	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.	<p>Note: Wait till the device responds.</p> <pre># 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</pre>

Table 3-6 (Cont.) Procedure to verify Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
2. <input type="checkbox"/>	Attempt to ssh to the switches with the username/password provided in the conf files.	<pre># 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/</pre>

Table 3-6 (Cont.) Procedure to verify Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
		<pre>lgpl-2.1.php and http://www.gnu.org/licenses/old-licenses/ library.txt. #</pre>
<p>3. <input type="checkbox"/></p>	<p>Verify the running-config has all expected configurations in the conf file using the show running-config command.</p>	<pre># 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</pre>

Table 3-6 (Cont.) Procedure to verify Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
4. <input type="checkbox"/>	Verify license on the switches	<p>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.</p> <pre># show license</pre> <p>Example output:</p> <pre># 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</pre> <pre># show license usage</pre> <p>Example output:</p> <pre># 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 - ... #</pre>

Table 3-6 (Cont.) Procedure to verify Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
5. <input type="checkbox"/>	Verify the RMS1 can ping the CNE_Management VIP	<pre># 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</pre>
6. <input type="checkbox"/>	Enable customer uplink	Connect or enable customer uplink.
7. <input type="checkbox"/>	Verify the RMS1 can be accessed from laptop. Use application such as putty etc to ssh to RMS1.	<pre>\$ 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 ~]#</pre>

SNMP Trap Configuration

Table 3-7 Procedure to configure SNMP Trap

Step #	Procedure	Description
1. <input type="checkbox"/>	SNMPv2c Configuration	<p>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 #unique_45</p> <pre>[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</pre>

Table 3-7 (Cont.) Procedure to configure SNMP Trap

Step #	Procedure	Description
2. <input type="checkbox"/>	Restrict direct access to ToR switches	<p>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:</p> <pre> [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 </pre>

Table 3-7 (Cont.) Procedure to configure SNMP Trap

Step #	Procedure	Description
3. <input type="checkbox"/>	Traffic egress	<p>Traffic egress out of cluster, including snmptrap traffic to SNMP trap receiver, and traffic goes to signal server:</p> <pre>[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</pre>

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](#)

Steps to configure Addresses for RMS iLOs, OA, EBIPA

Table 3-8 Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
1. <input type="checkbox"/>	Setup team0.2 interface	<pre>\$ 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</pre>
2. <input type="checkbox"/>	Subnet and conf file address	<p>The /etc/dhcp/dhcpd.conf file should already have been configured in 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.</p>

Table 3-8 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
3. <input type="checkbox"/>	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.	<pre># 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 {</pre>

Table 3-8 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
		<pre> 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; </pre>

Table 3-8 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
		<pre> 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"; </pre>

Table 3-8 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
4. <input type="checkbox"/>	Access RMS iLO from the DHCP address with default Administrator password. From the above <code>dhcpd.leases</code> 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.	<p>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.</p> <pre># ssh Administrator@192.168.20.104 Administrator@192.168.20.104's password: User:Administrator logged-in to ILO2M2909004F.labs.nc.tekelec.com(192.168.2 0.104 / FE80::BA83:3FF:FE47:649C) iLO Standard 1.37 at Oct 25 2018 Server Name: Server Power: On</pre>
5. <input type="checkbox"/>	Create RMS iLO new user. Create new user with customized username and password.	<pre></>hpiLO-> create /map1/accounts1 username=root password=TklcRoot group=admin,config,oemHP_rc,oemHP_power,oem HP_vm status=0 status_tag=COMMAND COMPLETED Tue Apr 2 20:08:30 2019 User added successfully.</pre>
6. <input type="checkbox"/>	Disable the DHCP before able to setup static IP. Setup static failed before DHCP is disabled.	<pre></>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</pre>

Table 3-8 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
7. <input type="checkbox"/>	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.	<pre># 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(192.168.2 0.104 / 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 #</pre>

Table 3-8 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
8. <input type="checkbox"/>	Set EBIPA addresses for InterConnect Bays (Enclosure Switches).	<p>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.</p> <p>Set address for each enclosure switch, note the last number 1 or 2 is the interconnect bay number.</p> <pre> 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. </pre>

Table 3-8 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
9. <input type="checkbox"/>	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.	<pre> 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> </pre>

Table 3-8 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
10. <input type="checkbox"/>	Add New User for OA.	<p>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.</p> <pre> 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. </pre>

Table 3-8 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
11. <input type="checkbox"/>	From OA, go to each blade with "connect server <bay number>", add New User for each blade.	<pre> 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,o emHPE_vm status=2 status_tag=COMMAND PROCESSING FAILED error_tag=COMMAND SYNTAX ERROR Tue Apr 23 16:18:58 2019 User added successfully. </pre>
12. <input type="checkbox"/>	Change to static IP on OA. In order not reply on DHCP and make the OA address stable, change to static IP.	<p>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.</p> <pre> 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> </pre>

Configure Legacy BIOS on Remaining Hosts

These procedures define the steps necessary to configure additional Legacy BIOS for all hosts in OCCNE 1.2. 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 in [Table 3-9](#).
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 1.15](#)
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](#)

Steps to configure the Legacy BIOS on Remaining Hosts

Table 3-9 Procedure to configure the Legacy BIOS on Remaining Hosts

Step #	Procedure	Description
1. <input type="checkbox"/>	Expose the System Configuration Utility on a RMS Host	<p>Expose the System Utility screen to the user for 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.</p> <ol style="list-style-type: none">1. 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.2. The System Utilities screen is exposed in the remote console.

Table 3-9 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

Step #	Procedure	Description
2. <input type="checkbox"/>	Expose the System Utility for an Enclosure Blade	<p>1. 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.</p> <p>a. SSH to the blade using the iLO IP address and the root user and password. This brings up the HP iLO prompt.</p> <pre>\$ 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:EAFF:FE89:35E) iLO Standard Blade Edition 1.37 at Oct 25 2018 Server Name: Server Power: On </>hpiLO-></pre> <p>b. Use VSP to connect to the blade remote console.</p> <pre></>hpiLO->vsp</pre> <p>c. Power cycle the blade to bring up the System Utility for that blade.</p> <p>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.</p> <p>d. Access the System Utility by hitting ESC 9.</p> <p>2. 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</p>

Table 3-9 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

Step #	Procedure	Description
		<ul style="list-style-type: none"><li data-bbox="906 352 1448 436">a. Expose the System Utility by following step 1 or 2 depending on the hardware being configured.<li data-bbox="906 457 1279 485">b. Select System Configuration<li data-bbox="906 506 1448 533">c. Select BIOS/Platform Configuration (RBSU)<li data-bbox="906 554 1448 716">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.<li data-bbox="906 737 1448 842">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.

Table 3-9 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

Step #	Procedure	Description
3. <input type="checkbox"/>	Change over from UEFI Booting Mode to Legacy BIOS Booting Mode	<ol style="list-style-type: none"> 1. Expose the System Utility by following step 1 or 2 depending on the hardware being configured. 2. Select System Configuration 3. Select BIOS/Platform Configuration (RBSU) 4. 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. 5. Select Boot Mode This generates a warning indicating the following: <div style="border: 1px solid black; padding: 5px; margin: 5px 0;"> <p>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.</p> </div> <p>Hit the ENTER key and two selections appear: UEFI Mode(highlighted) and Legacy BIOS Mode</p> 6. 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. 7. 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. 8. Hit the y key and an additional warning appears indicating: System configuration changed. A system reboot is required. Press ENTER to reboot the system. 9. a. Hit ENTER to force a reboot. Note: The boot must go into the process of actually trying to boot from the boot devices

Table 3-9 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

Step #	Procedure	Description
		<p>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.</p> <p>b. 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</p> <p>10. 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.</p>
4. <input type="checkbox"/>	Force PXE to boot from the first Embedded FlexibleLOM HPE Ethernet 10Gb 2-port Adapter	<p>1. Expose the System Utility by following step 1 or 2 depending on the hardware being configured.</p> <p>2. Select System Configuration.</p> <p>3. Select BIOS/Platform Configuration (RBSU) .</p> <p>4. Select Boot Options. This menu defines the boot mode.</p> <p>5. 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 2.1 Change over from UEFI Booting Mode to Legacy BIOS Booting Mode.</p> <p>6. Select Legacy BIOS Boot Order In the default view, the 10Gb Embedded FlexibleLOM 1 Port 1 is at the bottom of the list.</p> <p>7. 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.</p> <p>8. Select F10 if it is desired to save and stay in the utility or select the F12 it is desired to save and exit to continue the current boot process.</p>

Table 3-9 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

Step #	Procedure	Description
5. <input type="checkbox"/>	Enabling Virtualization	<p>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 Workload Profiles.</p> <p>Verifying Default Settings</p> <ol style="list-style-type: none"> 1. Expose the System Utility by following step 1 or 2 depending on the hardware being configured. 2. Select System Configuration 3. Select BIOS/Platform Configuration (RBSU) 4. 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. 5. 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.

Table 3-9 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

Step #	Procedure	Description
6. <input type="checkbox"/>	Disable RAID Configurations	<p>OCCNE does not currently support any RAID configuration. Follow this procedure to disable RAID settings if the default settings of the System Utility include any RAID configuration(s).</p> <p>Note: There may be more than one RAID Array set up. This procedure should be repeated for any RAID configuration.</p> <ol style="list-style-type: none"> 1. Expose the System Utility by following step 1 or 2 depending on the hardware being configured. 2. Select System Configuration. 3. Select Embedded RAID 1 : HPE Smart Array P408i-a SR Gen 10. 4. Select Array Configuration. 5. Select Manage Arrays. 6. Select Array A (or any designated Array Configuration if there are more than one). 7. Select Delete Array. A warning is displayed indicating the following: <p>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.</p> 8. Hit ENTER, the changes are submitted and Delete Array Successful is displayed. 9. Hit ENTER to go back to the main menu for the HPE Smart Array. 10. Select F10 if it is desired to save and stay in the utility or select the F12 it is desired to save and exit to continue the current boot process.

Table 3-9 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

Step #	Procedure	Description
7. <input type="checkbox"/>	Enable the Primary and Secondary Boot Devices	<p>This steps provide necessary to configure the primary and secondary bootable devices for a Gen10 Server.</p> <p>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.</p> <ol style="list-style-type: none"> 1. Expose the System Utility by following step 1 or 2 depending on the hardware being configured. 2. Select System Configuration. 3. Select Embedded RAID 1 : HPE Smart Array P408i-a SR Gen 10. 4. 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. 5. Examine the list of available hardware drives. If one or more HDDs are available, continue with this procedure. Note: A single drive can be set as both the primary and secondary boot device but that is not part of this configuration. 6. Select Bootable Physical Drive 7. 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. 8. Select Set as Primary Bootable Device. 9. Hit ENTER. Note: There is no need to set the secondary boot device. Leave it as Not Set. 10. Hit the ESC key to back out to the System Utilities menu. 11. 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**Table 3-10 Procedure to configure enclosure switches**

Step #	Procedure	Description
1. <input type="checkbox"/>	Copy the 6127XLG configuration file	<p>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.</p> <pre>\$ cp /media/usb/6127xlg_irf.cfg /var/lib/tftpboot/6127xlg_irf.cfg</pre> <pre>\$ ls -l /var/lib/tftpboot/</pre> <pre>total 1305096</pre> <pre>-rw-r--r--. 1 root root 311 Mar 25 08:41 6127xlg_irf.cfg</pre>
2. <input type="checkbox"/>	Modify the switch specific values in the /var/lib/tftpboot/6127xlg_irf.cfg file.	<p>These values are contained at Installation PreFlight checklist : Create the OA 6127XLG Switch Configuration File from column Enclosure_Switch.</p> <pre>\$ cd /var/lib/tftpboot</pre> <pre>\$ sed -i 's/{switchname}/<switch_name>/' 6127xlg_irf.cfg</pre> <pre>\$ sed -i 's/{admin_password}/<admin_password>/' 6127xlg_irf.cfg</pre> <pre>\$ sed -i 's/{user_name}/<user_name>/' 6127xlg_irf.cfg</pre> <pre>\$ sed -i 's/{user_password}/<user_password>/' 6127xlg_irf.cfg</pre>

Table 3-10 (Cont.) Procedure to configure enclosure switches

Step #	Procedure	Description
3. <input type="checkbox"/>	Access the InterConnect Bay1 6127XLG	<p>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. Note: When trying to save the config, the following prompt is received:</p> <pre>[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.</pre> <pre>\$ ssh <oa username>@<oa address></pre> <p>If it shows standby, ssh to the other OA address.</p> <pre>OA-FC15B41AEA05> connect interconnect 1</pre> <pre>....</pre> <pre><HPE>system-view</pre> <p>System View: return to User View with Ctrl +Z.</p> <p>(Note: Run the following commands:)</p> <pre>irf member 1 priority 32</pre> <pre>interface range Ten-GigabitEthernet 1/0/17 to Ten-GigabitEthernet 1/0/20</pre> <pre>shutdown</pre> <pre>quit</pre> <pre>irf-port 1/1</pre> <pre>port group interface Ten- GigabitEthernet1/0/17</pre>

Table 3-10 (Cont.) Procedure to configure enclosure switches

Step #	Procedure	Description
		<pre>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</pre>

Table 3-10 (Cont.) Procedure to configure enclosure switches

Step #	Procedure	Description
4. <input type="checkbox"/>	Access the InterConnect Bay2 6127XLG	<p>Access the InterConnect Bay2 6127XLG switch to re-number to IRF 2.</p> <pre> 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... </pre>

Table 3-10 (Cont.) Procedure to configure enclosure switches

Step #	Procedure	Description
		System is starting...
5. <input type="checkbox"/>	Configure the IRF on Bay2 6127XLG switch	<p>After rebooting, the interfaces will begin with number 2 such as Ten-GigabitEthernet2/0/17, Ten-GigabitEthernet2/1/5. Run the following commands:</p> <pre>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</pre>

Table 3-10 (Cont.) Procedure to configure enclosure switches

Step #	Procedure	Description
6. <input type="checkbox"/>	Run "reboot" command on both switches	<pre> <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... </pre>

Table 3-10 (Cont.) Procedure to configure enclosure switches

Step #	Procedure	Description
7. <input type="checkbox"/>	Verify the IRF for the 6127XLG switches.	<p>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.</p> <pre> <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] </pre>

Table 3-10 (Cont.) Procedure to configure enclosure switches

Step #	Procedure	Description
8. <input type="checkbox"/>	Configure the IRF switch with predefined configuration file.	<pre> <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. </pre>

Table 3-10 (Cont.) Procedure to configure enclosure switches

Step #	Procedure	Description
		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 3-11 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](#).

Table 3-12 Bastion Installation

Step #	Procedure	Description
1. <input type="checkbox"/>	Copy the Necessary Files from the Utility USB to Support the OS Install	<ol style="list-style-type: none"> 1. Login to the Bootstrap Host using the root credentials configured during OS installation of the Bootstrap Host. 2. Create the directories needed on the Installer Bootstrap Host. <pre>\$ mkdir -p /var/ocne/cluster/ <cluster_short_name>/yum.repos.d</pre> 3. 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. 4. 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. <pre>\$ cp /<path_to_usb>/hosts.ini /var/ ocne/cluster/<cluster_short_name>/ hosts.ini</pre> <p>Example:</p> <pre>\$ cp /media/usb/hosts.ini /var/ocne/ cluster/rainbow/hosts.ini</pre> 5. 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 occne_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). <p>Example (from hosts.sample.ini):</p> <pre>ntp_server='172.16.3.1' ocne_repo_host_address='172.16.3.4'</pre> 6. Follow procedure Populate the MetallB Configuration File to copy the MetallB

Table 3-12 (Cont.) Bastion Installation

Step #	Procedure	Description
		<p>configuration file into the <code>/var/occne/cluster/<cluster_short_name>/</code> directory and configure it.</p> <pre>\$ cp /<path_to_usb>/mb_configmap.yaml /var/occne/cluster/<cluster_short_name>/mb_configmap.yaml</pre> <p>Example:</p> <pre>\$ cp /media/usb/mb_configmap.yaml /var/occne/cluster/rainbow/mb_configmap.yaml</pre> <p>7. Copy the customer specific <code>.repo</code> file from the Utility USB to the Installer Bootstrap Host.</p> <p>This is the <code>.repo</code> 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.</p> <pre>\$ 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</pre> <p>Example:</p> <pre>\$ 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</pre>

Table 3-12 (Cont.) Bastion Installation

Step #	Procedure	Description
2. <input type="checkbox"/>	Set up the /etc/hosts file for the Central Repo and Verify Access	<p>1. Add an entry to the /etc/hosts file on the Installer Bootstrap Host to provide a name mapping for the Customer Central Repository.</p> <pre>\$ vi /etc/hosts</pre> <p>Example: 10.75.200.217 rainbow-reg</p> <p>To Verify: \$ ping <central_repo_name></p> <p>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</p> <p>2. Verify the repo access execute the following command:</p> <pre>\$ yum repolist</pre> <p>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 Ksplce for</p>

Table 3-12 (Cont.) Bastion Installation

Step #	Procedure	Description
		<pre>Oracle Linux 7 (x86_64) 393 !latest/x86_64 Oracle Linux 7 Latest (x86_64) 5,401 repolist: 19,437</pre>
<p>3.</p> <input type="checkbox"/>	<p>Copy the OL7 ISO to the Installer Bootstrap Host</p>	<p>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</p> <p>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.</p> <p>Note: If the user copies this ISO from their laptop then they must use an application like WinSCP pointing to the Management Interface IP.</p> <pre>\$ scp <usr>@<site_specific_address>:/ <path_to_iso>/OracleLinux-7.5-x86_64- disc1.iso /var/occne/OracleLinux-7.5- x86_64-disc1.iso</pre>
<p>4.</p> <input type="checkbox"/>	<p>Install Packages onto the Installer Bootstrap Host</p>	<p>Use YUM to install necessary packages onto the installer Bootstrap Host.</p> <pre>\$ yum install docker-engine nfs_utils ansible</pre>

Table 3-12 (Cont.) Bastion Installation

Step #	Procedure	Description
5. <input type="checkbox"/>	Set up access to the Docker Registry on the Installer Bootstrap Host	<ol style="list-style-type: none"> <li data-bbox="860 352 1466 913"> 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. <pre data-bbox="906 562 1445 913"> \$ 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 </pre> <p data-bbox="906 945 1015 970">Example:</p> <pre data-bbox="906 976 1356 1192"> \$ 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 </pre> <li data-bbox="860 1228 1466 1522"> Start the docker daemon. <pre data-bbox="906 1297 1258 1386"> \$ systemctl daemon-reload \$ systemctl restart docker \$ systemctl enable docker </pre> <p data-bbox="906 1423 1242 1449">Verify docker is running:</p> <pre data-bbox="906 1455 1242 1522"> \$ ps -elf grep docker \$ systemctl status docker </pre>

Table 3-12 (Cont.) Bastion Installation

Step #	Procedure	Description
6. <input type="checkbox"/>	Setup NFS on the Installer Bootstrap Host	<p>Run the following commands using sudo (assumes nfs-utils has already been installed in procedure: Installation of Oracle Linux 7.x on Bootstrap Host : Install Additional Packages).</p> <p>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.</p> <pre>\$ echo '/var/ocne 172.16.3.4/24(ro,no_root_squash)' >> /etc/ exports \$ systemctl start nfs-server \$ systemctl enable nfs-server</pre> <p>Verify nfs is running:</p> <pre>\$ ps -elf grep nfs \$ systemctl status nfs-server</pre>
7. <input type="checkbox"/>	Set up the Boot Loader on the Installer Bootstrap Host	<p>Execute the following commands:</p> <pre>\$ mkdir -p /var/ocne/pxelinux \$ mount -t iso9660 -o loop /var/ocne/ OracleLinux-7.5-x86_64-discl.iso /mnt \$ cp /mnt/isolinux/initrd.img /var/ocne/ pxelinux \$ cp /mnt/isolinux/vmlinuz /var/ocne/ pxelinux</pre>
8. <input type="checkbox"/>	Verify and Set the PXE Configuration File Permissions on the Installer Bootstrap Host	<p>Each file configured in the step above must be open for read and write permissions.</p> <pre>\$ chmod -R 777 /var/ocne/pxelinux</pre>
9. <input type="checkbox"/>	Disable DHCP and TFTP on the Installer Bootstrap Host	<p>The TFTP and DHCP services running on the Installer Bootstrap Host may still be running. These services must be disabled.</p> <pre>\$ systemctl stop dhcpd \$ systemctl disable dhcpd \$ systemctl stop tftp \$ systemctl disable tftp</pre>

Table 3-12 (Cont.) Bastion Installation

Step #	Procedure	Description
10. <input type="checkbox"/>	Disable SELINUX	<p>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.</p> <pre>\$ getenforce active</pre> <p>If the output of this command displays "active", change it to "permissive" by editing the /etc/selinux/config file.</p> <pre>\$ vi /etc/selinux/config</pre> <p>Change the SELINUX variable to passive: SELINUX=permissive save the file</p> <p>Reboot the system: reboot</p>
11. <input type="checkbox"/>	Generate the SSH private and public keys on Bootstrap Host.	<p>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. Just hit enter.</p> <p>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.</p> <p>Execute the following commands on the Bootstrap Host:</p> <pre>\$ 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 ""</pre>

Table 3-12 (Cont.) Bastion Installation

Step #	Procedure	Description
12. <input type="checkbox"/>	Execute the OS Install and Bastion VM Creation on Bastion KVM Host (RMS2) from the Installer Bootstrap Host	<p>1. Run the docker commands below to perform the OS install and Bastion Host VM creation on the Bastion KVM Host (RMS2):</p> <pre>\$ 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></pre> <p>Example:</p> <pre>\$ 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</pre> <p>2. 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.</p> <p>Note: This command is optional. Had a failure actually occurred, the docker run command would have experienced failures.</p> <pre>ssh -i /var/occne/cluster/<cluster_short_name>/ssh/occne_id_rsa admusr@<oam_host></pre> <pre>\$ sudo virsh list</pre> <p>Example:</p> <pre>ssh -i /var/occne/cluster/rainbow.lab.us.oracle.com/ssh/occne_id_rsa admusr@10.75.148.6</pre> <pre>\$ sudo virsh list</pre> <pre> Id Name State -----</pre>

Table 3-12 (Cont.) Bastion Installation

Step #	Procedure	Description
		<pre>----- 11 bastion-2.rainbow.lab.us.oracle.com running</pre> <p>3. Login to Bastion Host from the Bootstrap Host as admusr using the generated key from the <code>/var/occne/cluster/<cluster_short_name></code> directory to confirm the VM is set up correctly. The <code><oam_host></code> field (which is an IPv4 address) is derived from the <code>hosts.ini</code> file <code>bastion-2</code> entry for the <code>host_kernel_virtual</code> group.</p> <p>Note: This command is optional. Had a failure actually occurred, the docker run command would have experienced failures.</p> <pre>ssh -i /var/occne/cluster/ <cluster_short_name>/.ssh/occne_id_rsa admusr@<oam_host></pre> <p>Example:</p> <pre>ssh -i /var/occne/cluster/rainbow/.ssh/ occne_id_rsa admusr@10.75.148.5</pre>

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.

Table 3-13 Automated Installation

Step #	Procedure	Description																					
1. <input type="checkbox"/>	Setting up for and executing the deploy.sh script on the Bootstrap Host	<p>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:</p> <table border="1"> <thead> <tr> <th>Name</th> <th>Comment</th> <th>Example usage</th> </tr> </thead> <tbody> <tr> <td>CENTRAL_REPO</td> <td>Defines the customer specific repository host name. Note: This would be used in conjunction with CENTRAL_REPO_IP and CENTRAL_REPO_DOCKER_PORT.</td> <td>CENTRAL_REPO=customer_repo \ CENTRAL_REPO_IP=10.10.10.10 \ CENTRAL_REPO_DOCKER_PORT=5000 ./deploy.sh</td> </tr> <tr> <td>CENTRAL_REPO_IP</td> <td>Defines the customer specific repository IPv4 address.</td> <td>See above.</td> </tr> <tr> <td>CENTRAL_REPO_DOCKER_PORT</td> <td>Defines the customer specific repository docker port number. Defaults to 5000.</td> <td>See above.</td> </tr> <tr> <td>OCCNE_CLUSTER</td> <td>Defines the cluster short name.</td> <td>OCCNE_CLUSTER=rainbow</td> </tr> <tr> <td>OCCNE_BASTION</td> <td>Bastion Host full name</td> <td>OCCNE_BASTION=bastion-2.rainbow.us.labs.oracle.com</td> </tr> <tr> <td>OCCNE_VERSION</td> <td>The version tag of the image releases</td> <td>OCCNE_VERSION=1.3.2</td> </tr> </tbody> </table>	Name	Comment	Example usage	CENTRAL_REPO	Defines the customer specific repository host name. Note: This would be used in conjunction with CENTRAL_REPO_IP and CENTRAL_REPO_DOCKER_PORT.	CENTRAL_REPO=customer_repo \ CENTRAL_REPO_IP=10.10.10.10 \ CENTRAL_REPO_DOCKER_PORT=5000 ./deploy.sh	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=1.3.2
Name	Comment	Example usage																					
CENTRAL_REPO	Defines the customer specific repository host name. Note: This would be used in conjunction with CENTRAL_REPO_IP and CENTRAL_REPO_DOCKER_PORT.	CENTRAL_REPO=customer_repo \ CENTRAL_REPO_IP=10.10.10.10 \ CENTRAL_REPO_DOCKER_PORT=5000 ./deploy.sh																					
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=1.3.2																					

Table 3-13 (Cont.) Automated Installation

Step #	Procedure	Description
2. <input type="checkbox"/>	Copy necessary files from Utility USB to the Bootstrap Host staging directory	<p>1. 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.</p> <pre>\$ cp /<usb_dev>/db/*.zip /var/occne/*.zip</pre> <p>2. Copy the configBastionHost.yaml file from the Customer Utility USB to the staging directory /var/occne/.</p> <pre>\$ cp /<usb_dev>/configBastionHost.yaml /var/occne/.</pre> <p>3. Copy the deploy.sh script from the Customer Utility USB to the staging directory at /var/occne/ and set the file to executable.</p> <pre>\$ cp /<usb_dev>/deploy.sh /var/occne/. \$ chmod +x /var/occne/deploy.sh</pre>

Table 3-13 (Cont.) Automated Installation

Step #	Procedure	Description
3. <input type="checkbox"/>	Execute Deploy	<p>Execute the deploy.sh script from the /var/occne/ directory with the required parameters set.</p> <pre>\$ export CENTRAL_REPO=<customer_specific_repo_name> \$ export CENTRAL_REPO_IP=<customer_specific_repo_ip 4> \$ export OCCNE_CLUSTER=<cluster_short_name> \$ export OCCNE_BASTION=<bastion_full_name> \$ \$./deploy.sh</pre> <p>Customer Example:</p> <pre>\$ 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.ora cle.com \$./deploy.sh</pre> <p>The command above can be executed like the following:</p> <pre>\$ CENTRAL_REPO=central-repo CENTRAL_REPO_IP=10.10.10.10 OCCNE_CLUSTER=rainbow OCCNE_BASTION=bastion-2.rainbow.us.labs.ora cle.com ./deploy.sh</pre> <p>Internal Example:</p> <p>Internally the defaults would be used so the only possible variables that needs to be provided are the OCCNE_CLUSTER, OCCNE_BASTION, and OCCNE_VERSION.</p> <pre>\$ OCCNE_CLUSTER=rainbow OCCNE_BASTION=bastion-2.rainbow.us.labs.ora cle.com ./deploy.sh</pre>

Table 3-13 (Cont.) Automated Installation

Step #	Procedure	Description
4. <input type="checkbox"/>	Executing Final Deploy on Bastion Host	<p>The following commands are executed from the Bastion Host to complete the Bastion Host configuration and deploy OCCNE on the Bare Metal system.</p> <p>Note: The Bootstrap Host cannot be used to access the Bastion Host as it will be re-paved from execution of this command.</p> <ol style="list-style-type: none"> 1. 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: <code>chmod 0600 ~/.ssh/id_rsa</code> <pre>\$ ssh -i ~/.ssh/id_rsa admusr@<bastion_external_ip_address></pre> <ol style="list-style-type: none"> 2. 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. <pre>\$ export PROV_DEPLOY_ARGS='--limit=! <bastion_full_name>,! <bastion_kvm_host_full_name>' \$ export DB_ARGS='--extra- vars=mate_site_db_replication_ip=<remote _site_db_replication_service_ip>,occne_m ysqlndb_cluster_id=<mysqlndb_cluster_ide ntifier>' \$ /var/occne/cluster/ <cluster_short_name>/artifacts/ pipeline.sh</pre> <p>Customer Example:</p> <pre>\$ export PROV_DEPLOY_ARGS='--limit=! bastion-2.rainbow.lab.us.oracle.com,! db-2.rainbow.lab.us.oracle.com' \$ export DB_ARGS='--extra- vars=mate_site_db_replication_ip=10.75.1 82.88,occne_mysqlndb_cluster_id=2' \$ /var/occne/cluster/rainbow/artifacts/ pipeline.sh</pre> <p>To save the output from the pipeline.sh script the command can be written as:</p>

Table 3-13 Automated Installation

Step #	Procedure	Description
		<pre>\$ /var/occne/cluster/rainbow/artifacts/ pipeline.sh tee pipeline\$(date +%F_%H %M%S").log</pre> <p>Note: While Installing on the First Site ignore 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.</p>
5. <input type="checkbox"/>	Update the Bastion KVM Host repo file	<p>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.</p> <ol style="list-style-type: none"> 1. 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) <pre>\$ ssh -i /var/occne/cluster/ <cluster_short_name>/.ssh/occne_id_rsa admusr@<bastion_kvm_host_ip_address></pre> 2. Remove the existing .repo files: <pre>\$ sudo rm /var/occne/yum.repos.d/*.repo</pre> 3. Copy the Bastion specific .repo file from the Bastion Host to the Bastion KVM Host. Execute this command from the Bastion KVM Host. <pre>scp <bastion_short_name>:/var/occne/ yum.repos.d/*.repo /var/occne/ yum.repos.d/</pre> <p>Example:</p> <pre>scp bastion-2:/var/occne/yum.repos.d/ *.repo /var/occne/yum.repos.d/</pre>
6. <input type="checkbox"/>	Change MySQL root user password	Refer to Change MySQL root user password

Virtualized CNE Installation

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

Prerequisites

1. The user has 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 via this procedure.
3. Octavia Load Balancing plugin must be installed on the OpenStack Environment.
4. Users must have a pub 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 Keytab** on the Launch **Instance** → **Key Pair** dialog or via the **Compute** → **Access and Security**

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. All necessary images, binaries, etc have been downloaded from Oracle OSDC prior to executing this procedure and these resources are available for use in this procedure.
3. The customer has made available a central repository for all images, binaries, helm charts, etc, prior to executing this procedure.

 **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.

Upload an Image to an OpenStack Environment

This is the process of uploading the qcow2 image. The image is provided via OSDC.

 **Note:**

This procedure is executed from the OpenStack Desktop.

Table 3-14 Upload an Image to an OpenStack Environment

Step #	Procedure	Description
1. <input type="checkbox"/>	Navigate to Images	Go to Compute Images
2. <input type="checkbox"/>	Create Image	Select the +Create Image button. This brings up a new dialog. You have to enter a name for the image. Use the same name as was used to create and download the imag (ex: occne_bootstrap-1.3.0.qcow2).
3. <input type="checkbox"/>	Choose the source of the image	Using the Image Source pull down select: Image File . This will enable a Browse button. Select this button to bring up a Windows Explorer dialog. From the Windows dialog, select the image that was created in the previous procedures. This will insert the image name into the OpenStack Create Image dialog and set the Format for you.
4. <input type="checkbox"/>	Upload Image	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.
5. <input type="checkbox"/>	Check the image	When the process is complete, the image should be listed in the Compute Images screen. Again you will have to use Next to go through all the images and finally get to the image you uploaded depending on how many images are on the system.

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.

Table 3-15 Bootstrap Host Creation

Step #	Procedure
1. <input type="checkbox"/>	Login to the OpenStack Environment using your OpenStack credentials, the appropriate domain and project name.
2. <input type="checkbox"/>	Select Compute Instances
3. <input type="checkbox"/>	Select the Launch Instances tab on the upper right. A dialog will appear to configure a VM instance.
4. <input type="checkbox"/>	Enter an Instance Name (for example: occne-<name>). Leave the Availability Zone and Count set as is.
5. <input type="checkbox"/>	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. <input type="checkbox"/>	Make sure the filter pulldown is set to Image
7. <input type="checkbox"/>	Enter occne_bootstrap in the filter. This will display the occne_bootstrap-<x.y.z>.qcow2 image uploaded in the previous sections of this procedure.
8. <input type="checkbox"/>	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. <input type="checkbox"/>	Select Flavor
10. <input type="checkbox"/>	Enter a string which best describes the flavor being used for this customer specific OpenStack Environment in the search filter. This brings up a new dialog.
11. <input type="checkbox"/>	Select the appropriate customer specific Flavor (for example: occne_bsh_flavor) 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. <input type="checkbox"/>	Select Networks
13. <input type="checkbox"/>	Enter the appropriate network name as defined by the customer with the OpenStack Environment (example: ext-net) in the search filter. This brings up a new dialog.
14. <input type="checkbox"/>	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. <input type="checkbox"/>	Select Key Pair . This dialog assumes you have already uploaded a public key to to OpenStack (see Prerequisites).
16. <input type="checkbox"/>	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.

Table 3-15 (Cont.) Bootstrap Host Creation

Step #	Procedure
17. <input type="checkbox"/>	<p>Select Configuration. This screen allows the user to add configuration data which is used by cloud-init to set on the VM, the initial admusr and hostname/FQDN additions to the /etc/hosts file. Use the following configuration until there is more information available. 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 in their default setting.</p> <pre>#cloud-config hostname: <instance_name_from_details_screen> fqdn: <instance_name_from_details_screen> system_info: default_user: name: admusr 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'</pre>
18. <input type="checkbox"/>	<p>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 ComputeInstances screen.</p>

Pre-deployment Configuration

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

Table 3-16 Pre-deployment Configuration

Step #	Procedure	Description
1. <input type="checkbox"/>	Obtain the TLS Certificate for OpenStack	<p>Depending on the Customer's environment it is very likely that the customer's OpenStack uses certificates for TLS access to the API. Without this certificate, OpenStack commands will not work. Customer's may have to obtain this certificate before using OpenStack client commands.</p> <ol style="list-style-type: none">1. Contact the OpenStack admin to provide the required TLS certificate to access the client commands (ex. in an Oracle OpenStack system installed with kolla, the certificate will be located at <code>/etc/kolla/certificates/haproxy-ca.crt</code>)2. Copy the certificate to the Bootstrap Host at location: <code>/etc/pki/<OpenStack_release_name>/haproxy-ca.crt</code> (ex. <code>/etc/pki/kolla/haproxy-ca.crt</code>) (If <code>/etc/pki/<OpenStack_release_name></code> does not exist, it can be created using command: <code>mkdir -p /etc/pki/<OpenStack_release_name></code>)3. Set the environment variable <code>OS_CACERT</code> to the location where the certificate was copied to using the command: <code>export OS_CACERT=/etc/pki/<OpenStack_release_name>/haproxy-ca.crt</code> (ex. <code>export OS_CACERT=/etc/pki/kolla/haproxy-ca.crt</code>)

Table 3-16 (Cont.) Pre-deployment Configuration

Step #	Procedure	Description
2. <input type="checkbox"/>	Get the Openstack RC (API v3) File	<p>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.</p> <p>Note: These instructions may be somewhat different on OpenStack Desktops.</p> <ol style="list-style-type: none"> 1. From the OpenStack Desktop: go to Project Compute Access & Security. 2. Select the API Access tab 3. On the right hand side, select Download OpenStack RC File v3. This will download a openrc.sh file prefixed with the OpenStack project name (ex: 5G-openrc.sh). to your PC. 4. 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. 5. Execute the following command: <code>source .<project_name>-openrc.sh</code> 6. Execute the following command to verify the OpenStack Client is working: <code>openstack image list</code>
3. <input type="checkbox"/>	Create SSH Key on Bootstrap Host	<p>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 to 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.</p> <pre>\$ ssh-keygen -m PEM -t rsa -b 2048</pre>

Table 3-16 (Cont.) Pre-deployment Configuration

Step #	Procedure	Description
4. <input type="checkbox"/>	Add Files to /tmp Directory	<p>These files must be copied to the directories listed using scp or some other means (ie. winSCP).</p> <ol style="list-style-type: none">1. There are three directories on the Bootstrap Host. These three directories are as follows:<ol style="list-style-type: none">a. /tmp/yum.repos.db. /tmp/dbc. /tmp/certificates2. Within these three directories the user must supply the following mandatory files:<ol style="list-style-type: none">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).

Table 3-16 (Cont.) Pre-deployment Configuration

Step #	Procedure	Description
5. <input type="checkbox"/>	Updating the ~/.configure/ openstack/clouds.yaml File	<p>1. To obtain the values for the authorization fields need in clouds.yaml below, execute the <code>openstack configuration show</code> command (make sure you have sourced the <code>openrc</code> script before executing this command).</p> <pre> \$ openstack configuration show +-----+ +-----+ +-----+ Field +-----+ Value +-----+ +-----+ +-----+ api_timeout None application_catalog_api_version 1 auth.auth_url http://thundercloud.us.oracle.com: 5000/v3 auth.password <redacted> auth.project_domain_id 6c3468f1207c4e00bb441746c2046a90 auth.project_id 811ef89b5f154ab0847be2f7e41117c0 auth.project_name OCCNE auth.user_domain_name LDAP auth.username John.Doe auth_type password baremetal_api_version 1 </pre>

Table 3-16 (Cont.) Pre-deployment Configuration

Step #	Procedure	Description
		 beta_command False
		 cacert None
		 cert None
		 compute_api_version 3
		 container_api_version 1
		 container_infra_api_version 1
		 +-----+ +-----+ ----+

2. To get the floating_network_id value for the load balancer configuration in clouds.yaml below, execute the following commands:

```
$ openstack network list
```

Get the network ID and the subnets id for floating ip address network name:

Example from the openstack command cli list:

```
+-----+
+-----+
+-----+
| ID
| Name          |
Subnets          |
+-----+
+-----+
+-----+
| e4351e3e-81e3-4a83-bdc1-dde1296690e3
| ext-net      | c0e0c185-
ed65-4a53-a7a3-418277fb9a20 |
```

Table 3-16 (Cont.) Pre-deployment Configuration

Step #	Procedure	Description
3.		<p>Edit the <code>~/configure/openstack/clouds.yaml</code> file (using <code>sudo vi ~/configure/openstack/clouds.yaml</code>) and update following values for your Openstack Environment using the commands listed above (use double quotes where indicated in the example given below):</p> <p>Note: The <code>subnet_id</code> field is automatically obtained and populated in the <code>clouds.yaml</code> file. There is no need to configure that value.</p> <pre> clouds: mycloud: auth: auth_url: <openstack-Identity-api-url> username: <openstack-user-name> project_name: <openstack-project-name> project_id: <openstack-Project ID> user_domain_name: <openstack-Domain Name> password: <openstack-user-password> region_name: <openstack-region-name-if-available> interface: <openstack-interface-public/private> identity_api_version: <openstack-identity-api-version> loadbalancer: lbaas_enabled: true subnet_id: <openstack_lbaas_subnet_id> floating_network_id: <openstack_lbaas_floating_network_id> use_octavia: true lb_method: ROUND_ROBIN </pre> <p>Example:</p> <pre> clouds: mycloud: auth: auth_url: http:// thundercloud.us.oracle.com:5000/v3 username: "john.doe" project_name: "OCCNE" project_id: </pre>

Table 3-16 (Cont.) Pre-deployment Configuration

Step #	Procedure	Description
		811ef89b5f154ab0847be2f7e41117c0 user_domain_name: "LDAP" password: "johnspw" region_name: "RegionOne" interface: "public" identity_api_version: 3 loadbalancer: lbaas_enabled: true subnet_id: c0e0c185-ed65-4a53- a7a3-418277fb9a20 floating_network_id: e4351e3e-81e3-4a83-bdc1-dde1296690e3 use_octavia: true lb_method: ROUND_ROBIN

Table 3-16 (Cont.) Pre-deployment Configuration

Step #	Procedure	Description
6. <input type="checkbox"/>	Updating cluster.tfvars File	<p>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.</p> <p>Note: Image to use for bastion, masters, standalone etcd instances, and nodes image = "OracleLinux-7.5-x86_64". An Admin user of the customer specific OpenStack Environment must upload this image 'OracleLinux-7.5-x86_64.qcow2' to the openstack env and it should be accessible under image list.</p> <p>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.</p> <ol style="list-style-type: none"> 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 to distinguish it from the occne_example directory. <p>Note: The directory name is not used for any special purpose other than to distinguish it from the occne_example directory. It can be called anything.</p> <pre>\$ cp -R occne_example occne_<user></pre> Use the following commands to retrieve the information necessary to configure the cluster.tfvars <ol style="list-style-type: none"> The different flavor settings should be set according to the recommendations from Reference 1 in this document. An Admin user of the customer specific OpenStack Environment must add the flavors and provide the UUID of those flavors for configuration into the cluster.tfvars file. The UUID of each specific flavor that is used must be added as the value field of the key/value fields in the cluster.tfvars file. Once flavors have been added to the OpenStack Environment, the UUID can be retrieved via the following OpenStack Client command from the Bootstrap Host shell: <pre>\$ openstack flavor list grep <flavor_name></pre> <p>Example:</p> <pre>\$ openstack flavor list grep medium</pre>

Table 3-16 (Cont.) Pre-deployment Configuration

Step #	Procedure	Description
		<pre> 43c7b73b-42a7-4f40- b52e-11f6803fc750 oc- cne.medium 16384 80 0 2 True </pre>
		<pre> \$ openstack flavor show 43c7b73b-42a7-4f40-b52e-11f6803fc750 +-----+ +-----+ ---+ Field Value +-----+ +-----+ ---+ OS-FLV-DISABLED:disabled False OS-FLV-EXT-DATA:ephemeral 0 access_project_ids None disk 80 id 43c7b73b-42a7-4f40- b52e-11f6803fc750 name oc- cne.medium os-flavor-access:is_public True properties ram 16384 rxtx_factor 1.0 swap </pre>

Table 3-16 (Cont.) Pre-deployment Configuration

Step #	Procedure	Description
		<pre> vcpus 2 +-----+ +-----+ ---+ </pre>
c.	Use the following command to retrieve the floatingip_pool name and external_net UUID.	<pre> \$ 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 </pre>

Table 3-16 (Cont.) Pre-deployment Configuration

Step #	Procedure	Description
		<pre> e4351e3e-81e3-4a83-bdc1- dde1296690e3 ext-net c0e0c185-ed65-4a53- a7a3-418277fb9a20 fc36d63f- b30b-4c7f-979f-9b52b614bbd7 occne- mkingre 7631612f-5d22-49be-975c-6e0a9329339b +-----+ ---+----- +-----+ ---+ </pre>
3.		<p>Navigate to <code>occne_<user></code> directory and edit the contents of the <code>cluster.tfvars</code> file in the newly created directory:</p> <pre> \$ vi occne_<user>/cluster.tfvars </pre>
4.		<p>The fields in the <code>cluster.tfvars</code> file must be configured to adapt to the current customer Openstack Provider. Initially the <code>cluster_name</code> and <code>network_name</code> should be set as the same value.</p> <pre> # <Kubernetes cluster name here> cluster_name = "<cluster-name>" # networking network_name = "<cluster-name>" </pre>
5.		<p>For setting the <code>ntp_server</code> value in the <code>cluster.tfvars</code> 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 <code>thundercloud.us.oracle.com</code>)</p> <pre> \$ 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 </pre>
6.		<p>If your deployment requires a specific Availability Zone other than the default availability zone called <code>nova</code>, please make the following changes in the <code>cluster.tfvars</code> file. This value can be added just after</p>

Table 3-16 (Cont.) Pre-deployment Configuration

Step #	Procedure	Description
		<p>the cluster_name field. If you wish to use nova then do not add these changes.</p> <pre>az_list = ["<your_availability_zone_name>"]</pre> <p>Example:</p> <pre># Authorizing_Zone az_list = ["foobar"]</pre>
7.	Obtain Mate Site DB Replication Service Load Balancer IP	<p>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.</p> <p>Note: If this is a single deploy and or a mated site with this being the first site deployed, this step can be skipped.</p> <ol style="list-style-type: none"> 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. <pre>[cloud-user@arjun1-bastion-1 ~]\$ 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</pre> <p>In the above example we could see "10.75.182.88" as the Mate Site DB Replication Service Load Balancer IP.</p>

Deploy the OCCNE Virtualized Cluster

The execution of the following command does all the work to deploy the VMs in the OpenStack Environment, configure the Bastion Host, and deploy and configure the Kubernetes clusters.

Table 3-17 Deploy the OCCNE Virtualized Cluster

Step #	Procedure	Description
1. <input type="checkbox"/>	Deploy OCCNE Virtualized Cluster	<p>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).</p> <pre>\$ OCCNE_TFVARS_DIR=occne_<user> CENTRAL_REPO=<central_repo_hostname> CENTRAL_REPO_IP=<central_repo_ipv4_address> OCCNE_PIPELINE_ARGS='DB_ARGS=--extra- vars='\''{"mate_site_db_replication_ip": "'< db_replication_service_load_balancer_ip>'", "occne_mysqlndb_cluster_id":<mysqlndb_clust er_idenfifier>}'\'' ./deploy.sh</pre> <p>Example:</p> <pre>\$ OCCNE_TFVARS_DIR=occne_arjun CENTRAL_REPO=winterfell CENTRAL_REPO_IP=10.75.216.10 OCCNE_PIPELINE_ARGS='DB_ARGS=--extra- vars='\''{"mate_site_db_replication_ip": "'1 0.75.182.88'", "occne_mysqlndb_cluster_id": 2}'\'' ./deploy.sh</pre> <p>Note: While Installing on the First Site ignore OCCNE_PIPELINE_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 on second site.</p> <p>Note: The release version defaults to the 1.4.0 GA release. If there is a reason to use a different version, it can be specified by setting the OCCNE_VERSION=<release> variable on the command line.</p>
2. <input type="checkbox"/>	Change MySQL root user password	Refer to Change MySQL root user password

4

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 has 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.

Table 4-1 OCCNE Post Install Verification

Step No.	Procedure	Description
1. <input type="checkbox"/>	Run the commands to get the load-balancer IP address and port number for Kibana Web Interface.	<pre># 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_LOAD BALANCER_PORT http://10.75.182.51:80 Launch the Browser and navigate to http://\$KIBANA_LOADBALANCER_IP:\$KIBANA_LO ADBALANCER_PORT(e.g.: http://10.75.182.51:80 in the example above) received in the output of the above commands.</pre>

Table 4-1 (Cont.) OCCNE Post Install Verification

Step No.	Procedure	Description
2. <input type="checkbox"/>	Using Kibana verify Log and Tracer data is stored in Elasticsearch	<ol style="list-style-type: none"> 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. <ol style="list-style-type: none"> 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.
3. <input type="checkbox"/>	Verify Elasticsearch cluster health	<ol style="list-style-type: none"> 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.

Table 4-1 (Cont.) OCCNE Post Install Verification

Step No.	Procedure	Description
4. <input type="checkbox"/>	Verify Prometheus Alert manager is accessible	<p>1. Run below commands to get the load-balancer IP address and port number for Prometheus Alert Manager Web Interface.</p> <pre># 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</pre> <p>2. Launch the Browser and navigate to <code>http://\$ALERTMANAGER_LOADBALANCER_IP:\$ALERTMANAGER_LOADBALANCER_PORT</code> (e.g.: <code>http://10.75.182.53:80</code> in the example above) received in the output of the above commands. Ensure the AlertManager GUI is accessible.</p>

Table 4-1 (Cont.) OCCNE Post Install Verification

Step No.	Procedure	Description
5. <input type="checkbox"/>	Verify metrics are scraped and stored in prometheus server	<p>1. Run below commands to get the load-balancer IP address and port number for Prometheus Server Web Interface.</p> <pre># 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:\$PROM ETHEUS_LOADBALANCER_PORT http://10.75.182.54:80</pre> <p>2. Launch the Browser and navigate to <code>http://\$PROMETHEUS_LOADBALANCER_IP:\$PROMETHEUS_LOADBALANCER_PORT</code> (e.g.: <code>http://10.75.182.54:80</code> in the example above) received in the output of the above commands. Ensure the Prometheus server GUI is accessible.</p> <p>3. Select "UP" option from "insert metric at cursor" drop down and click on "Execute" button.</p> <p>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).</p>

Table 4-1 (Cont.) OCCNE Post Install Verification

Step No.	Procedure	Description
6. <input type="checkbox"/>	Verify Alerts are configured	<ol style="list-style-type: none">1. Navigate to alerts tab of Prometheus server GUI or navigate using URL http://\$PROMETHEUS_LOADBALANCER_IP:\$PROMETHEUS_LOADBALANCER_PORT/alertsFor<PROMETHEUS_LOADBALANCER_IP>and<PROMETHEUS_LOADBALANCER_PORT>2. If below alerts are seen in " Alerts" tab of prometheus GUI, then Alerts are configured properly.

Table 4-1 (Cont.) OCCNE Post Install Verification

Step No.	Procedure	Description
7. <input type="checkbox"/>	Verify grafana is accessible and change the default password for admin user	<ol style="list-style-type: none"> Run below commands to get the load-balancer IP address and port number for Grafana Web Interface. <pre># 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</pre> Launch the Browser and navigate to <code>http://\$GRAFANA_LOADBALANCER_IP:\$GRAFANA_LOADBALANCER_PORT</code> (e.g.: <code>http://10.75.182.55:80</code> 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. At first connection to the Grafana dashboard, a 'Change Password' screen will appear. Change the password to the customer provided credentials. <p>Note: Grafana data is not persisted, so if Grafana services restarted for some reason change password screen will appear again.</p> Grafana dashboards are accessed after the changing the default password in the above step. Click on "New dashboard" as marked red below. Click on "Add Query"

Table 4-1 (Cont.) OCCNE Post Install Verification

Step No.	Procedure	Description
		<p>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.</p> <p>8. In the Query Section marked in Red below put " sum by(__name__) ({kubernetes_namespace="occne-infra"}) " and then click any where outside of the textbox and wait for few seconds. Ensure the dashboard appearing in the top section of the page. This link shows all the metrics and number of entries in each metrics over time span originated from kubernetes namespace 'occne-infra. In the add query section we can give any valid promQL query.Example for using the metrics list link above to write a promQL query: sum(\$metricnamefromlist)sum by(kubernetes_pod_name) (\$metricnamefromlist{kubernetes_namespace="occ ne-infra"})For more details about promQL please follow the link.</p>

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 4-2 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

Table 4-2 (Cont.) Credentials

Credential Name	Type	Associated Resource	Initial Setting	Credential Rotation
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
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

Table 4-3 Post-Installation Security Hardening

Step No.	Procedure	Description
1. <input type="checkbox"/>	Reset Credentials on the TOR Switch	<p>1. From bastion host, login to the switch with username and password from the procedure</p> <pre>[bastion host]# ssh <username>@<switch IP address> User Access Verification Password: <password></pre> <pre>Cisco Nexus Operating System (NX-OS) Software TAC support: http://www.cisco.com/tac <switch name>#</pre> <p>2. Change the password for current username:</p> <pre># # configure Enter configuration commands, one per line. End with CNTL/Z. (config)# username <username> password <newpassword> (config)#exit #</pre> <p>3. Create new username:</p> <pre># # 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 #</pre> <p>4. Exit from the switch and login with the new username and password to verify the new change works:</p> <pre># exit Connection to <switch IP address> closed. [bastion host]#</pre>

Table 4-3 (Cont.) Post-Installation Security Hardening

Step No.	Procedure	Description
		<pre>[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>#</pre> <p>5. Delete the previous old username if it is not needed:</p> <pre># # configure Enter configuration commands, one per line. End with CNTL/Z. (config)# no username <username> (config)#exit #</pre> <p>6. Change the enable secret when needed:</p> <pre># (config)# enable secret <newenablepassword> (config)# exit #</pre> <p>7. Save the above configuration:</p> <pre># copy running-config startup-config [#####] #] 100% Copy complete, now saving to disk (please wait)... Copy complete. #</pre>

Table 4-3 (Cont.) Post-Installation Security Hardening

Step No.	Procedure	Description
2. <input type="checkbox"/>	Reset Credentials on the Enclosure Switch	<ol style="list-style-type: none"> <li data-bbox="857 380 1448 1045"> 1. From bastion host, login to the switch with username and password from the procedure: <pre data-bbox="906 478 1448 1045"> [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. * ***** ***** </pre> <li data-bbox="857 1050 1448 1491"> 2. Change the password for current username: <pre data-bbox="906 1308 1448 1491"> [switchname]local-user <username> class <current class> [switchname-luser-manage- <username>]password simple <newpassword> [switchname-luser-manage-<username>]quit [switchname] </pre> <li data-bbox="857 1495 1448 1896"> 3. Create new username: <pre data-bbox="906 1598 1448 1896"> [switchname]local-user <newusername> class [manage network] New local user added. [switchname-luser-manage- <newusername>]password simple <newpassword> [switchname-luser-manage- <newusername>]quit [switchname] </pre>

Table 4-3 (Cont.) Post-Installation Security Hardening

Step No.	Procedure	Description
		<p>4. Exit from the switch and login with the new username and password to verify the new change works:</p> <pre> <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] </pre> <p>5. Delete the previous old username if it is not needed:</p> <pre> [switchname]undo local-user <username> class <current class> </pre> <p>6. Save the above configuration:</p> <pre> [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>] </pre>

Table 4-3 (Cont.) Post-Installation Security Hardening

Step No.	Procedure	Description
		<pre>(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]</pre>

Table 4-3 (Cont.) Post-Installation Security Hardening

Step No.	Procedure	Description
3. <input type="checkbox"/>	Reset Credentials for the OA Admin Console	<p>1. 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)</p> <pre>[bastion host]# ssh <username>@<OA address></pre> <p>----- ----- 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 monitored and can result in criminal or civil prosecution under applicable law. ----- -----</p> <pre>Firmware Version: 4.85 Built: 04/06/2018 @ 06:14 OA Bay Number: 1 OA Role: Active <username>@<OA address>'s password:<password></pre> <p>HPE BladeSystem Onboard Administrator (C) Copyright 2006-2018 Hewlett Packard Enterprise Development LP</p> <p>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.</p> <pre>OA-A45D36FD5FB1></pre>

Table 4-3 (Cont.) Post-Installation Security Hardening

Step No.	Procedure	Description
		<p>2. Change the password for current username:</p> <pre>OA-A45D36FD5FB1> set password <newpassword></pre> <p>Changed password for the "<username>" user account.</p> <pre>OA-A45D36FD5FB1></pre> <p>3. Add new user:</p> <pre>OA-A45D36FD5FB1> add user <newusername></pre> <p>New Password: <newpassword> Confirm : <newpassword> User "<newusername>" created. You may set user privileges with the 'SET USER ACCESS' and 'ASSIGN' commands.</p> <pre>OA-A45D36FD5FB1> set user access <newusername> [ADMINISTRATOR OPERATOR USER]</pre> <p>"<newusername>" has been given [administrator operator user] level privileges.</p> <p>4. Assign full access to the enclosure for the user:</p> <pre>OA-A45D36FD5FB1> assign server all <newusername></pre> <p><newusername> has been granted access to the valid requested bay(s)</p> <pre>OA-A45D36FD5FB1> assign interconnect all <newusername></pre> <p><newusername> has been granted access to the valid requested bay(s)</p> <pre>OA-A45D36FD5FB1> assign oa <newusername></pre> <p><newusername> has been granted access to the OA.</p>

Table 4-3 (Cont.) Post-Installation Security Hardening

Step No.	Procedure	Description
		<p>5. Exit from the OA and login with the new username and password to verify the new change works:</p> <pre>OA-A45D36FD5FB1> exit</pre> <p>Connection to <OA address> closed. [bastion host]# ssh <newusername>@<OA address></p> <p>----- ----- 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 monitored and can result in criminal or civil prosecution under applicable law. ----- -----</p> <pre>Firmware Version: 4.85 Built: 04/06/2018 @ 06:14 OA Bay Number: 1 OA Role: Active <newusername>@<OA address>'s password:<newpassword></pre> <p>HPE BladeSystem Onboard Administrator (C) Copyright 2006-2018 Hewlett Packard Enterprise Development LP</p> <p>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.</p>

Table 4-3 (Cont.) Post-Installation Security Hardening

Step No.	Procedure	Description
		<p>OA-A45D36FD5FB1></p> <p>6. Delete previous user if not needed:</p> <p>OA-A45D36FD5FB1> remove user <username></p> <p>Entering anything other than 'YES' will result in the command not executing.</p> <p>Are you sure you want to remove testuser1? yes</p> <p>User "<username>" removed.</p>

Table 4-3 (Cont.) Post-Installation Security Hardening

Step No.	Procedure	Description
<p>4.</p> <p><input type="checkbox"/></p>	<p>Reset Credentials for the ILO Admin Console</p>	<ol style="list-style-type: none"> <li data-bbox="781 380 1370 850"> <p>1. From bastion host, login to the iLO with username and password from the procedure:</p> <pre data-bbox="829 478 1321 787">[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-></pre> <li data-bbox="781 890 1370 1199"> <p>2. Change current password:</p> <pre data-bbox="829 957 1268 1140"></>hpiLO-> set /map1/accounts1/ <username> password=<newpassword> status=0 status_tag=COMMAND COMPLETED Tue Aug 20 13:27:08 2019 </>hpiLO-></pre> <li data-bbox="781 1245 1370 1619"> <p>3. Create new user:</p> <pre data-bbox="829 1312 1357 1556"></>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.</pre> <li data-bbox="781 1665 1370 1843"> <p>4. Exit from the iLO and login with the new username and password to verify the new change works:</p> <pre data-bbox="829 1755 1029 1843"></>hpiLO-> exit status=0</pre>

Table 4-3 (Cont.) Post-Installation Security Hardening

Step No.	Procedure	Description
		<pre> 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-> 5. 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. </pre>
5. <input type="checkbox"/>	Reset Credentials for the root account on Each and Every Server	<p>Login to each and every server in the cluster (ssh admusr@cluster_host) and perform the following command:</p> <pre> sudo passwd root </pre>
6. <input type="checkbox"/>	Reset (or Delete) Credentials for the admusr account on Each and Every Server	<p>Login to each and every server in the cluster (ssh admusr@cluster_host) and perform the following command:</p> <pre> sudo passwd -l admusr </pre>

Table 4-3 (Cont.) Post-Installation Security Hardening

Step No.	Procedure	Description
7. <input type="checkbox"/>	Reset Credentials for the MySQL Accounts	See Database Tier Installer for details on how to reset the DB Account Passwords.
8. <input type="checkbox"/>	Regenerate / Redistribute SSH Keys Credentials for the admusr Account	<p>Log into the Bastion Host VM and generate a new cluster-wide keypair by perform the following:</p> <pre>ssh-keygen -b 4096 -t rsa -C "New SSH Key" -f .ssh/new_occne_id_rsa -q -N ""</pre> <p>Now, for each and every server in the cluster, perform these actions:</p> <pre># 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</pre> <p>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:</p> <pre># 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</pre> <p>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.</p>

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.

YUM Repository Configuration

Introduction

To perform an installation without the system needing access to the internet, a local YUM mirror must be made of the OL7 latest, epel, and addons repository used 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_ksplice]  
[ol7_x86_64_latest]  
[ol7_x86_64_addons]  
[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

Table A-1 Procedure to configure OCCNE YUM Repository

Step #	Procedure	Description
1. <input type="checkbox"/>	Create OL7 repository mirror repo	<p>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)</p> <pre> [local_ol7_x86_64_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_ [local_ol7_x86_64_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_ [local_ol7_x86_64_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_ [local_ol7_x86_64_ksplince] name=Ksplince for Oracle Linux 7 (x86_64) baseurl=http://10.75.155.195/yum/ OracleLinux/OL7/ksplince/\$basearch/ gpgcheck=1 gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY enabled=1 proxy=_none_ [local_ol7_x86_64_developer] name=Packages for creating test and development environments for Oracle Linux </pre>

Table A-1 (Cont.) Procedure to configure OCCNE YUM Repository

Step #	Procedure	Description
		<pre>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_ [local_ol7_x86_64_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_</pre>

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 is setup and docker commands can be run by the target system.
2. 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

Table A-2 Steps to configure OCCNE HTTP Repository

Steps	Procedure	Description
1. <input type="checkbox"/>	Retrieve Kubernetes Binaries	<p>The Kubernetes installer requires access to an HTTP server 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.</p> <p>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 (see below).</p> <pre>\$ sh /var/occne/<cluster>/artifacts/k8s_retrieve_bin.sh /var/www/html</pre> <p>Example:</p> <pre>\$ sh /var/occne/rainbow/artifacts/k8s_retrieve_bin.sh /var/www/html</pre>
2. <input type="checkbox"/>	Retrieve Helm binaries and charts	<p>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 command below.</p> <pre>\$ sh /var/occne/<cluster>/artifacts/retrieve_helm.sh /var/www/html <helm path> < /var/occne/<cluster>/artifacts/config_helm_charts.txt</pre> <p>Example</p> <pre>\$ sh /var/occne/rainbow/artifacts/retrieve_helm.sh /var/www/html ./ < /var/occne/rainbow/artifacts/config_helm_charts.txt</pre>

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 is installed and docker commands can be run
2. Make sure docker registry is running

```
$ dockerps
```

3. 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>:<1.4.x_tag> /getdeps/getdeps  
$ docker run --rm -it -v /var/occne/<cluster>/:/host occne/
```

```
<k8s_install_image_name>:<1.4.x_tag> /getdeps/getdeps
```

Example-

```
$ docker run --rm -it -v /var/ocne/rainbow:/host ocne/configure:1.4.0 /  
getdeps/getdeps  
$ docker run --rm -it -v /var/ocne/rainbow:/host ocne/k8s_install:  
1.4.0 /getdeps/getdeps
```

Once the above command is successfully executed, go to `/var/ocne/<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/ocne/<cluster>/artifacts/retrieve_docker.sh docker.io  
<registryaddress>:<port> < /var/ocne/<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 [occne:vars] section to direct the installation logic to the registry, these variables need to be set to the your docker registry configuration:

```
...  
[occne:vars]  
...  
occne_private_registry=winterfell  
occne_private_registry_address='10.75.216.114'  
occne_private_registry_port=5002  
occne_helm_images_repo='winterfell:5002'  
...
```

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"> • ansible_host:The IP address for the host's primary interface. The occne containers use this interface to communicate with the host. • kvm_host: The physical host name.fqdn running KVM that hosts this VM host (ex. for guest db-10.icemark.lab.us.oracle.com the kvm_host is db-1.icemark.lab.us.oracle.com). Bastion-1 should be on db-1, bastion-2 should be on db-2. • mac: Always begin with 52-54-00 with the last 3 hex values being unique within the hosts.ini file (ex: mac=52-54-00-c1-8e-38) • signal_host: The Signaling network IPv4 address assigned to the MySQL NDB SQL Nodevirtual machines. : The ILO network IPv4 address assigned to the Bastionhost virtual machines. This IP is manually assigned and should be on the same network as the host_hp_gen_10/9/8 node iLo address except for the final octet. For ex:ample: If the kvm_host=db-2 and the iLo field in host_hp_gen_10/9/8 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). • ilo_host: 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 host_hp_gen_10/9/8 node iLo address except for the final octet. For example: If the kvm_host=db-2 and the iLo field in host_hp_gen_10/9/8 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). • oam_host: 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 data_store and kube-master. 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 kube-master

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 Ex: 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 Ex: 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 1.0 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 Ex: 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 Signalling network For Ex: 5
signal_subnet_ipv4	Set to the subnet of the Signalling network used to assign IPs for Storage hosts
signal_subnet_cidr	Set to the cidr notation for the Signalling subnet. For example 29
signal_netmask	Set appropriately for the network used to assign Signalling IPs for Storage hosts and MySQL SQL Node VM's.
signal_broadcast_address	Set appropriately for the network used to assign Signalling IPs for Storage hosts and MySQL SQL Node VM's.
signal_default_route	Set to the Signalling VIP of the TOR switch.
occne_snmp_notifier_destination	Set to the address of SNMP trap receiver. For Ex: "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 Ex: 2
2	occne_mysqlndb_DataMemory	Size of Data Memory(RAM) assigned to each MySQL Data Nodes. For Ex: 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_CNEManagementNetwork_IP	Bastion Host Network		4	Address needs to be without prefix length, such as 10.25.100.2
9	ToRswitchB_CNEManagementNetwork_IP	Bastion Host Network		4	Address needs to be without prefix length, such as 10.25.100.3
10	ToRswitch_CNEManagementNetwork_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_management_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_mgmt_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.yaml 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	Enclos ure_S witch1 Value	Enclosure_Swit ch2 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=NAMESERVERIPS --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..t37RwUZr7S1VmIBvMmWV
VoUjb2DJJ2f4VlrW9RdfVi.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..t37RwUZr7S1VmIBvMmWV
oUjb2DJJ2f4VlrW9RdfVi.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/occne-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 generated in step 2.1 above.
<code>subnet_ipv4</code>	Set to the subnet of the network used to assign IPs for OCCNE hosts
<code>subnet_cidr</code>	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 <code>/24</code>
<code>netmask</code>	Set appropriately for the network used to assign IPs for OCCNE hosts.

Table B-14 (Cont.) OCCNE Variables

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

Table B-15 Install Additional Services/Network Functions

Step No #	Procedure	Description
1. <input type="checkbox"/>	Copy docker images needed for the service into the bastion-host docker registry	<p>1. Create a file <code>docker_images.txt</code> listing the required docker images and tags</p> <pre>dockerRepo:5000/<serviceNameImage></pre> <p>Example</p> <pre>dockerRepo:5000/serviceNameImage2:1.2.2</pre> <p>2. Load these images into the cluster-local docker registry by running:</p> <pre>\$ /var/occne/cluster/<cluster>/artifacts/retrieve_docker.sh <<docker_images.txt</pre>

Table B-15 (Cont.) Install Additional Services/Network Functions

Step No #	Procedure	Description
2. <input type="checkbox"/>	Copy helm charts needed for the service into the bastion-host helm chart repository	<ol style="list-style-type: none"> 1. Create a file helm_charts.txt listing the required helm charts and versions: helmRepoName/chart_name <version> 2. 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 3. Load the chart(s) into the cluster-local helm chart repository by running: \$ /var/ocne/cluster/<cluster>/artifacts/retrieve_helm.sh <<helm_charts.txt
3. <input type="checkbox"/>	Install the service	<p>Create a values.yaml file on the Bastion Host that contains the values needed by the Helm chart</p> <p>To install the service run:</p> <pre>\$ helm install --name <release-name> --namespace <service-namespace> -f values.yaml <chart_name></pre>

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/mysqld_expired.log" file. so we need to login to the each of the MySQL SQL nodes and change the MySQL root user password.

Table B-16 Change MySQL root user password

Step No.#	Procedure	Description
1.	Login to MySQL SQL Node VM.	

Table B-16 (Cont.) Change MySQL root user password

Step No.#	Procedure	Description
2.	Login using mysql client	<p>Login to mysql client as a root user</p> <pre>\$ sudo su \$ mysql -h 127.0.0.1 -uroot -p</pre>
3.	Enter expired random password for mysql root user stored in the "/var/occnedb/mysqlid_expired.log" file	<p>Enter expired random password stored in "/var/occnedb/mysqlid_expired.log" file.</p> <pre>\$ mysql -h 127.0.0.1 -uroot -p \$ Enter password:</pre>
4.	Change Root Password	<p>Execute the following commands to change the root password:</p> <pre>\$ mysql> ALTER USER 'root'@'localhost' IDENTIFIED BY '<NEW_PASSWORD>'; \$ mysql> FLUSH PRIVILEGES;</pre> <p>Note: Here 'NEW_PASSWORD' is the password of the mysql root user.</p>

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.

MySQL Cluster Carrier Grade Edition 0.0.0.0.0	
MySQL Cluster Manager 1.4.7 for Linux x86-64	
<input type="checkbox"/>	V980754-01.zip MySQL Cluster Manager 1.4.7+Cluster RPM for Oracle Linux / RHEL 6 x86 (64bit), 194.6 MB
<input type="checkbox"/>	V980758-01.zip MySQL Cluster Manager 1.4.7+Cluster TAR for Generic Linux x86 (64bit), 527.5 MB
<input type="checkbox"/>	V980755-01.zip MySQL Cluster Manager 1.4.7+Cluster TAR for Oracle Linux / RHEL 6 x86 (64bit), 527.5 MB
<input checked="" type="checkbox"/>	V980756-01.zip MySQL Cluster Manager 1.4.7+Cluster TAR for Oracle Linux / RHEL 7 x86 (64bit), 528.5 MB
<input type="checkbox"/>	V980763-01.zip MySQL Cluster Manager 1.4.7 RPM for Oracle Linux / RHEL 6 x86 (64bit), 16.0 MB
<input type="checkbox"/>	V980767-01.zip MySQL Cluster Manager 1.4.7 TAR for SuSE Enterprise Linux 11 x86 (64bit), 20.9 MB
<input type="checkbox"/>	V980764-01.zip MySQL Cluster Manager 1.4.7 TAR for Oracle Linux / RHEL 6 x86 (64bit), 20.7 MB
<input type="checkbox"/>	V980765-01.zip MySQL Cluster Manager 1.4.7 RPM for Oracle Linux / RHEL 7 x86 (64bit), 16.2 MB
<input type="checkbox"/>	V980766-01.zip MySQL Cluster Manager 1.4.7 TAR for Oracle Linux / RHEL 7 x86 (64bit), 21.6 MB
<input type="checkbox"/>	V980769-01.zip MySQL Cluster Manager 1.4.7 TAR for Generic Linux x86 (64bit), 20.7 MB

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/Access <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.

<input type="checkbox"/>	Oracle Linux 7.5.0.0.0	
<input type="checkbox"/>	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.