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





Oracle Communications Oracle Communications Signaling, Cloud Native Environment (OC-CNE) Installation Guide, Release 1.2.0

F21617-01

Copyright © 2019, Oracle and/or its affiliates. All rights reserved.

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 installed on the hardware, and/or documentation, delivered to U.S. Government end users are "commercial computer software" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, use, duplication, disclosure, modification, and adaptation of the programs, including any operating system, integrated software, any programs installed on the hardware, and/or documentation, shall be subject to license terms and license restrictions applicable to the programs. 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 Xeon 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, Opteron, the AMD logo, and the AMD Opteron 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 Introduction

	1-1
Key terms	1-1
Key Acronyms and Abbreviations	1-2
Overview	1-3
OCCNE Installation Overview	1-3
Frame and Component Overview	1-4
Frame Overview	1-4
Host Designations	1-5
Node Roles	1-6
Transient Roles	1-7
Create OCCNE Instance	1-8
How to use this document	1-10
Documentation Admonishments	1-11
Locate Product Documentation on the Oracle Help Center Site	1-12
Customer Training	1-12
My Oracle Support	1-12
My Gracie Support	
Emergency Response	1-13
•	
Emergency Response	
Emergency Response Installation Prerequisites	1-13
Emergency Response Installation Prerequisites Obtain Site Data and Verify Site Installation	2-1
Emergency Response Installation Prerequisites Obtain Site Data and Verify Site Installation Configure Artifact Acquisition and Hosting	2-1 2-1
Installation Prerequisites Obtain Site Data and Verify Site Installation Configure Artifact Acquisition and Hosting Oracle eDelivery Artifact Acquisition	2-1 2-1 2-1 2-1
Installation Prerequisites Obtain Site Data and Verify Site Installation Configure Artifact Acquisition and Hosting Oracle eDelivery Artifact Acquisition Third Party Artifacts	2-1 2-1 2-1 2-1
Installation Prerequisites Obtain Site Data and Verify Site Installation Configure Artifact Acquisition and Hosting Oracle eDelivery Artifact Acquisition Third Party Artifacts Install Procedure	2-1 2-1 2-1 2-1
Installation Prerequisites Obtain Site Data and Verify Site Installation Configure Artifact Acquisition and Hosting Oracle eDelivery Artifact Acquisition Third Party Artifacts Install Procedure Initial Configuration - Prepare a Minimal Boot Strapping Environment	2-1 2-1 2-1 2-1
Installation Prerequisites Obtain Site Data and Verify Site Installation Configure Artifact Acquisition and Hosting Oracle eDelivery Artifact Acquisition Third Party Artifacts Install Procedure Initial Configuration - Prepare a Minimal Boot Strapping Environment Installation of Oracle Linux 7.5 on Bootstrap Host	2-1 2-1 2-1 2-1 3-1 3-3



Configure Legacy BIOS on Remaining Hosts	3-41
Configure Enclosure Switches	3-47
Bastion Host Installation	3-53
Install Host OS onto RMS2 from the Installer Bootstrap Host (RMS1)	3-54
Configuration of the Bastion Host	3-62
Software Installation Procedures - Automated Installation	3-70
Oracle Linux OS Installer	3-70
Database Tier Installer	3-76
OCCNE Kubernetes Installer	3-82
Populate the MetalLB Configuration	3-85
OCCNE Automated Initial Configuration	3-85
NF Installation in the cluster	3-89
Post Installation Activities	
Post Install Verification	4-1
Post-Installation Security Hardening	4-6
Artifacts	
Repository Artifacts	A-1
Docker Repository Requirements	A-13
OCCNE YUM Repository Configuration	A-14
OCCNE HTTP Repository Configuration	A-16
OCCNE Docker Image Registry Configuration	A-17
Reference Procedures	
Inventory File Template	B-1
Inventory File Preparation	B-2
OCCNE Artifact Acquisition and Hosting	B-12
Installation Preflight Checklist	B-12
Installation Use Cases and Repository Requirements	B-33
Topology Connection Tables	B-50
Network Redundancy Mechanisms	B-55
Install VMs for MySQL Nodes and Management Server	B-62



List of Figures

1-1	Frame Overview	1-5
1-2	Host Designations	1-6
1-3	Node Roles	1-7
1-3	Transient Roles	1-8
1-4	OCCNE Installation Overview	1-9
1-6	Example of a Procedure Steps Used in This Document Installation Process	1-11
3-1		3-2
B-1	Rackmount ordering	B-13
B-2	Frame reference	B-34
B-3	Setup the Notebook and USB Flash Drive	B-36
B-4	Setup the Management Server	B-37
B-5	Management Server Unique Connections	B-38
B-6	Configure OAs	B-39
B-7	Configure the Enc. Switches	B-39
B-8	OceanSpray Download Path	B-40
B-9	Install OS on CNE Nodes - Server boot instruction	B-41
B-10	Install OS on CNE Nodes - Server boot process	B-42
B-11	Update OS on CNE Nodes - Ansible	B-43
B-12	Update OS on CNE Nodes - Yum pull	B-44
B-13	Harden the OS	B-45
B-14	Create the Guest	B-46
B-15	Install the Cluster on CNE Nodes	B-47
B-16	Install the Cluster on CNE Nodes - Pull in Software	B-48
B-17	Execute Helm on Master Node	B-49
B-18	Master Node Pulls from Repositories	B-50
B-19	Blade Server NIC Pairing	B-56
B-20	Rackmount Server NIC Pairing	B-56
B-21	Logical Switch View	B-58
B-22	OAM Uplink View	B-59
B-23	Top of Rack Customer Uplink View	B-60
B-24	OAM and Signaling Separation	B-61
B-25	MySQL Cluster Topology	B-63



List of Tables

1-1	Key Terms	1-1
1-2	Key Acronyms and Abbreviations	1-2
1-3	Admonishments	1-11
2-1	Oracle eDelivery Artifact Acquisition	2-1
3-1	Bootstrap Install Procedure	3-4
3-2	Procedure to configure the Installer Bootstrap Host BIOS	3-11
3-3	Procedure to configure Top of Rack 93180YC-EX Switches	3-16
3-4	Procedure to verify Top of Rack 93180YC-EX Switches	3-26
3-5	Procedure to configure SNMP Trap	3-30
3-6	Procedure to configure Addresses for RMS iLOs, OA, EBIPA	3-33
3-7	Procedure to configure the Legacy BIOS on Remaining Hosts	3-42
3-8	Procedure to configure enclosure switches	3-48
3-9	Procedure to install the OL7 image onto the RMS2 via the installer bootstrap host	3-55
3-10	Procedure to configure Bastion Host	3-63
3-11	Procedure to run the auto OS-installer container	3-71
3-12	OCCNE Database Tier Installer	3-77
3-13	Uninstall MySQL Cluster Manager and MySQL NDB Cluster	3-82
3-14	Procedure to install OCCNE Kubernetes	3-83
3-15	Procedure to configure MetalLB pools and peers	3-85
3-16	Procedure to install common services	3-86
3-17	Procedure to install NF in OCCNE Cluster	3-90
4-1	OCCNE Post Install Verification	4-1
4-2	Credentials	4-6
4-3		4-8
A-1	OL YUM Repository Requirements	A-1
A-2	Docker Repository Requirements	A-13
A-3	Steps to configure OCCNE HTTP Repository	A-16
A-4	Steps to configure OCCNE Docker Image Registry	A-18
B-1	Base Groups	B-3
B-2	Data Tier Groups	B-4
B-3	Procedure for OCCNE Inventory File Preparation	B-5
B-4	occne:vars	B-6
B-5	Enclosure Switch Connections	B-13
B-6	ToR Switch Connections	B-15
B-7	Rackmount Server Connections	B-18



B-8	Complete Site Survey Subnet Table	B-18
B-9	Complete Site Survey Host IP Table	B-19
B-10	Complete VM IP Table	B-20
B-11	Complete OA and Switch IP Table	B-21
B-12	ToR and Enclosure Switches Variables Table (Switch Specific)	B-24
B-13	Complete Site Survey Repository Location Table	B-25
B-14	Enclosure Switch Connections	B-51
B-15	ToR Switch Connections	B-52
B-16	Management Server Connections	B-55
B-17	Procedure to install VMs for MySQL Nodes and Management Server	B-66



1

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.

This document applies to version 1.2 of OCCNE installation procedure.

Glossary

Key terms

This table below lists terms used in this document.

Table 1-1 Key Terms

Term	Definition
Host	A computer running an instance of an operating system with an IP address. Hosts can be virtual or physical. The HP DL380 Gen10 Rack Mount Servers and BL460c Gen10 Blades are physical hosts. KVM based virtual machines are virtual hosts. Hosts are also referred to as nodes, machines, or computers.
Database Host	The Database (DB) Host is a physical machine that hosts guest virtual machines which in turn 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.
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.



Table 1-1 (Cont.) Key Terms

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.

Key Acronyms and Abbreviations

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

Table 1-2 Key Acronyms and Abbreviations

Acronym/	Definition
Abbreviation/Term	
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
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



Table 1-2 (Cont.) Key Acronyms and Abbreviations

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

Overview

OCCNE Installation Overview

The installation procedures in this document provision and configure an OCCNE. Using Oracle partners, the customer purchases the required hardware which is then configured and prepared for installation by Oracle Consulting.



To aid with the provisioning, installation, and configuration of OCCNE, a collection of container-based utilities are used to automate much of the initial setup. These utilities are based on tools such as PXE, the Kubespray project, and Ansible:

- PXE helps reliably automate provisioning the hosts with a minimal operating system.
- Kubespray helps reliably install a base Kubernetes cluster, including all dependencies (like etcd), using the Ansible provisioning tool.
- Ansible is used to deploy and manage a collection of operational tools (Common Services)
 provided by open source third party products such as Prometheus, Grafana, ElasticSearch
 and Kibana.
- Common services and functions such as load balancers and ingress controllers are deployed, configured, and managed as Helm packages.

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.



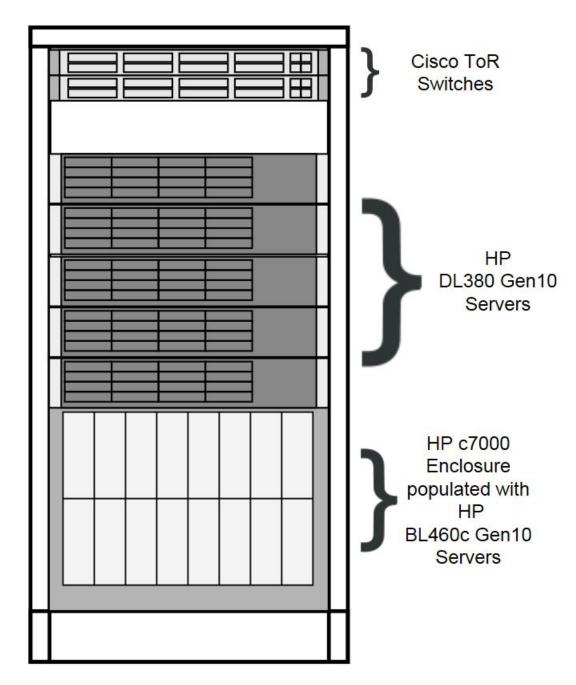
In the installation process, some of the roles of servers change as the installation procedure proceeds.

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.



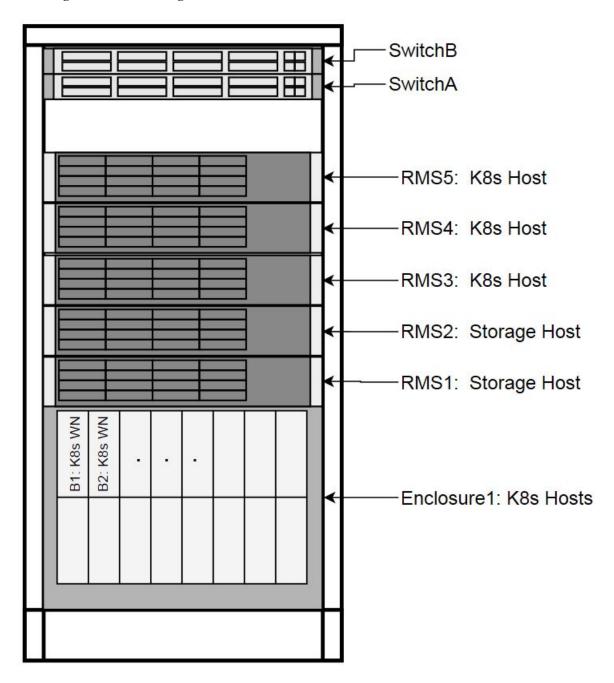
Figure 1-1 Frame Overview



Host Designations

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

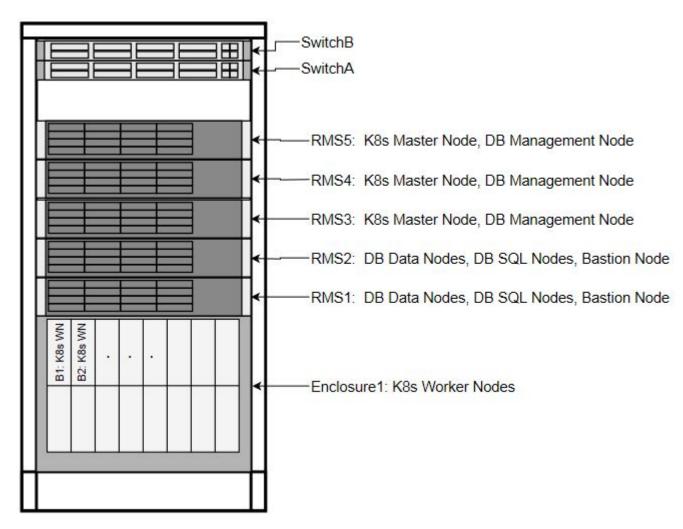
Figure 1-2 Host Designations



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 1-3 Node Roles

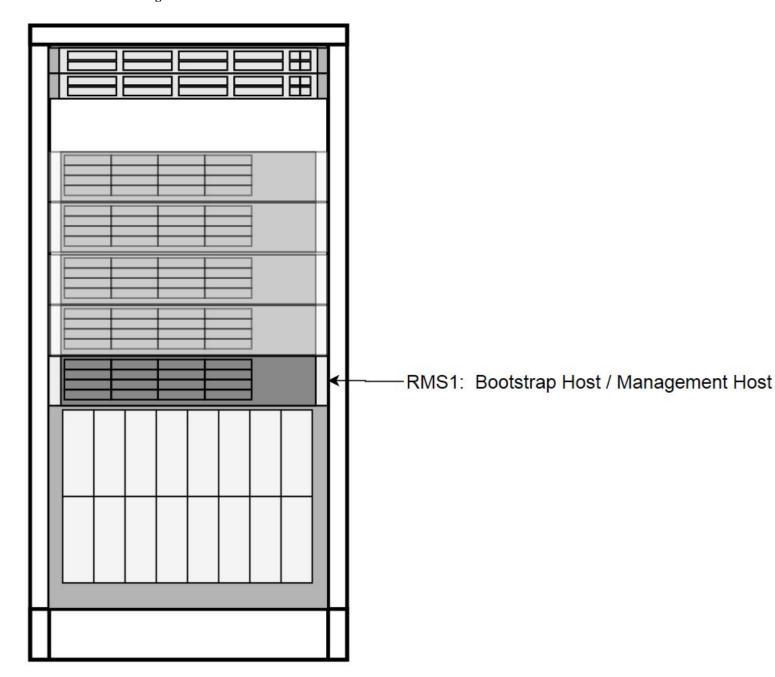


Transient Roles

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 1-4 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:

Bastion Host

Creates

Creates

Creates

Creates

DB Node

Figure 1-5 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 Artifacts 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 - <u>RMS1</u> - the transient Installer Bootstrap Host. In this procedure,



- a minimal set of packages needed to configure switches, iLOs, PXE boot environment, and provision <u>RMS2</u> 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 <u>RMS2</u>, 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 Figure 1-5) 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

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.

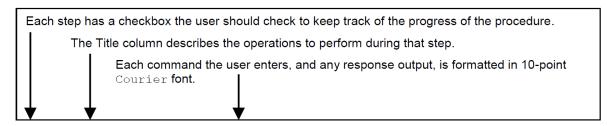
When executing this document for either purpose, there are a few points which help to ensure that the user understands the author's intent. These points are as follows:

- 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 1-6 Example of a Procedure Steps Used in This Document



	Title	Directive/Result Step	
1.	Change directory	Change to the backout directory.	
		\$ cd /var/TKLC/backout	
2.	ServerX: Connect	Establish a connection to the server using cu on the terminal server/console.	
	to the console of the server	\$ cu -l /dev/ttyS7	
3.	Verify Network Element data	View the Network Elements configuration data; verify the data; save and print report.	
		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 1-3 Admonishments

Icon	Description
110	Danger:
	(This icon and text indicate the possibility of personal injury.)
DANGER	
\wedge .	Warning:
<u>/4</u> `	(This icon and text indicate the possibility of equipment damage.)
WARNING	
	Caution:
	(This icon and text indicate the possibility of
CAUTION	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.
- 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.



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. This section addresses these prerequisites.

Obtain Site Data and Verify Site Installation

Execute the Installation Preflight Checklist procedure to obtain site survey data (IP address allocations, repository locations, etc), verify the frame configuration, and obtain important files used in installation procedures.

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 2-1 Oracle eDelivery Artifact Acquisition

Artifact	Description	File Type	Destination Repository
occne-images-1.2.0.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
v975367-01.iso	OL7 ISO	ISO	File repository
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 Artifacts.



3

Install Procedure

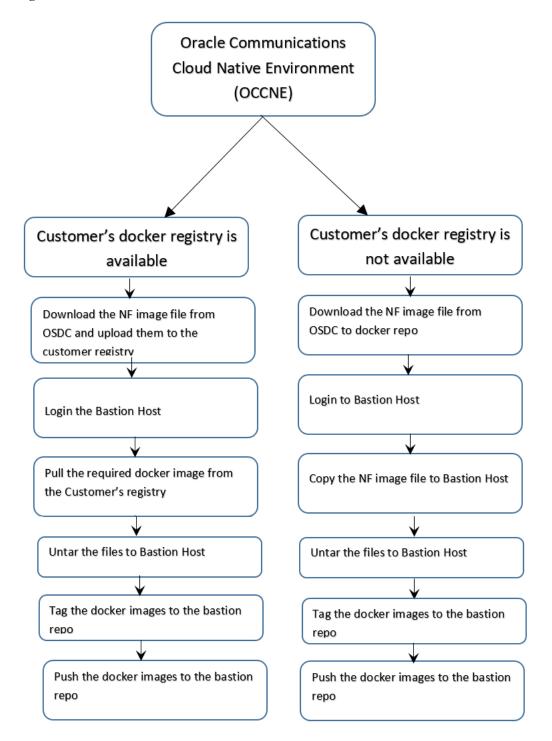
Initial Configuration - Prepare a Minimal Boot Strapping Environment

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.

Following flow chart describes the steps to install NFs on OCCNE:



Figure 3-1 Installation Process





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

- 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

- 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 3-1
 Bootstrap Install Procedure

Step #	Procedure	Description
1.	Create Bootable USB Media	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.
		\$ dd -if= <path iso="" to=""> -of=<usb device="" path=""> -bs=1m</usb></path>
		Example (assuming the USB is on /dev/sdf and the iso file is at /var/occne)
		<pre>\$ dd -if=/var/occne/OracleLinux-7.5-x86_64- discl.iso -of=/dev/sdf -bs=1m</pre>



 Table 3-1 (Cont.) Bootstrap Install Procedure

Step #	Procedure	Des	cription
2.	Install OL7 on the Installer Bootstrap Host.	1.	Connect a Keyboard, Video, and Mouse (KVM) into the Installer Bootstrap Host's monitor and USB ports.
	Bootstap 110st.	2.	Plug the USB flash drive containing the bootable iso into an available USB port on the Bootstrap host (usually in the front panel).
		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.
		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.
		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.
		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.
			a. LANGUAGE SUPPORT: English (United States)
			b. KEYBOARD : English (US)
			c. INSTALLATION SOURCE: Local Media
			d. SOFTWARE SELECTION: Minimal Install
			INSTALLATION DESTINATION should display No disks selected. Select INSTALLATION DESTINATION to indicate the drive to install the OS on.
			Select the first HDD drive (in this case that would be the first one listed or the 1.6 TB disk) and select DONE in the upper right corner.
			If the server has already been installed a red banner at the bottom of the page may indicate there is an error condition. Selecting that banner causes a dialog to appear indicating there is not enough free space (which might mean an OS has already been installed). In the dialog it may show both 1.6 TB HDDs as claimed or just the one.
			If only one HDD is displayed (or it could be both 1.6 TB drives selected, select the Reclaim space button. Another dialog appears. Select the Delete all button and the Reclaim space button again. Select DONE to return to the INSTALLATION SUMMARY screen.
			If the disk selection dialog appears (after selecting the red banner at the bottom of the page), this implies a full installation of the RMS has already been performed (usually this is because the procedure



 Table 3-1 (Cont.) Bootstrap Install Procedure

Step #	Procedure	Description
		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.
		7. Select DONE . This returns to the INSTALLATION SUMMARY page.
		8. At the INSTALLATION SUMMARY screen, select Begin Installation. The CONFIGURATION screen is displayed.
		9. At the CONFIGURATION screen, select ROOT PASSWORD.
		Enter a root password appropriate for this installation. It is good practice to use a customer provided secure password to minimize the host being compromised during installation.
		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.



 Table 3-1 (Cont.) Bootstrap Install Procedure

Step #	Procedure	Description
3.	Install Additional Packages.	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: • dnsmasq • dhcp • xinetd • tftp-server • dos2unix • nfs-utils
		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 no 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:
		\$ lsblk sdd
		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. Note: The output is shortened here for display purposes.
		\$ 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



 Table 3-1 (Cont.) Bootstrap Install Procedure

Step #	Procedure	Description
		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.
		<pre>\$ fdisk -1 /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</pre>
		Device Boot Start End Blocks Id System /dev/sda1 * 0 8929279 4464640 0 Empty /dev/sda2 3076 20503 8714 ef EFI (FAT-12/16/32)
		In the example output above, the /dev/sda2 partition is the EFI boot partition. Therefore the install media files are on /dev/sda1. Use the mount command to mount the install media file system. The same command without any options is used to verify the device is mounted to /media/usb.
		<pre>\$ mount /dev/sda1 /media/usb</pre>
		<pre>\$ mount /dev/sda1 on /media/usb type iso9660 (ro,relatime,nojoliet,check=s,map=n,blocksize=2048)</pre>
		5. Create a yum config file to install packages from local install media. Create a repo file /etc/yum.repos.d/Media.repo with the following information:



 Table 3-1 (Cont.) Bootstrap Install Procedure

Step #	Procedure	Description
		<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>
		6. Disable the default public yum repo. This is done by renaming the current .repo file to end with something other than .repo. Adding .disabled to the end of the file name is standard.
		Note : This can be left in this state as the Installer Bootstrap Host is re-paved in a later procedure.
		<pre>\$ mv /etc/yum.repos.d/public-yum-ol7.repo /etc/ yum.repos.d/public-yum-ol7.repo.disabled</pre>
		7. Use the yum repolist command to check the repository configuration. The output of yum repolist should look like the example below. Verify there no errors regarding un-reachable yum repos.
		<pre>\$ yum repolist Loaded plugins: langpacks, ulninfo repo id repo name</pre>
		status ol7_base_media Oracle Linux 7 Base Media 5,134
		repolist: 5,134
		8. Use yum to install the additional packages from the USB repo.
		<pre>\$ yum install dnsmasq \$ yum install dhcp \$ yum install xinetd \$ yum install tftp-server \$ yum install dos2unix \$ yum install nfs-utils</pre>
		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 /etc/xinetd.d directory. Installation/Verification does not include actually starting any of the services. Service configuration/starting is performed in a later procedure.
		Verify dhcp is installed:
		\$ cd /etc/dhcp \$ ls dhclient.d dhclient-exit-hooks.d dhcpd6.conf dhcpd.conf scripts
		Verify xinetd is installed:



Table 3-1 (Cont.) Bootstrap Install Procedure

Step #	Procedure	Description
		\$ cd /etc/xinetd.d \$ ls chargen-dgram chargen-stream daytime-dgram daytime-stream discard-dgram discard-stream echo-dgram echo-stream tcpmux-server time-dgram time-stream
		Verify tftp is installed: \$ cd /etc/xinetd.d
		<pre>\$ 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>
		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.
		<pre>\$ umount /media/usb</pre>
		<pre>\$ mount Verify that /dev/sdal 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 3-2
 Procedure to configure the Installer Bootstrap Host BIOS

Step #	Procedure	Description
1.	Expose the System Configuration Utility	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.
		 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.	Change over from UEFI Booting Mode to Legacy BIOS Booting Mode	 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. Expose the System Configuration Utility by following Step 1. Select System Configuration. Select BIOS/Platform Configuration (RBSU). 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.
		 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 3-2 (Cont.) Procedure to configure the Installer Bootstrap Host BIOS

Step #	Procedure	Description
3.	Adding a New User Account	This procedure provides the steps required to add a new user account to the server iLO 5 interface.
		Note : This user must match the pxe_install_lights_out_usrfields as provided in the hosts inventory files created using the template: OCCNE Inventory File Preparation.
		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 <pre>password> in the Password field.</pre>
		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 3-2 (Cont.) Procedure to configure the Installer Bootstrap Host BIOS

Step #	Procedure	Description
4.	Force PXE to boot from the first Embedded FlexibleLOM HPE Ethernet 10Gb 2- port Adapter	During host PXE, the DHCP DISCOVER requests from the hosts must be broadcast over the 10Gb port. This 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.
		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.
		 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 3-2 (Cont.) Procedure to configure the Installer Bootstrap Host BIOS

Step #	Procedure	Description
5.	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 Workload Profiles.
		1. Verifying Default Settings
		a. Expose the System Configuration Utility by following Step1.
		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. Disable RAID Configurations		1. Expose the System Configuration Utility by following Step 1.
	Configurations	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 3-2 (Cont.) Procedure to configure the Installer Bootstrap Host BIOS

Step #	Procedure	Description
7.	Enable the Primary Boot Device	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.
		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.
		 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.
		 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.
		 Select F10: Save to save and stay in the utility or select the F12: Save and Exit to save and exit, to complete the current boot process.
8.	Configure the iLO 5 Static IP Address	When configuring the Bootstrap host, the static IP address for the iLO 5 must be configured.
		Note : This procedure requires a reboot after completion.
		1. Expose the System Configuration Utility by following Step 1.
		2. Select System Configuration.
		3. Select iLO 5 Configuration Utility.
		4. Select Network Options.
		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.



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

Prerequisites

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

Limitations/Expectations

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

References

- https://github.com/datacenter/nexus9000/blob/master/nx-os/poap/poap.py
- https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus5000/sw/configuration/guide/cli/CLIConfigurationGuide/Licensing.html

Procedures

Configuration

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

Step #	Procedure	Description
1.	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.	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 ls /media/usb command. Note: Instructions for mounting the USB can be found in: Installation of Oracle Linux 7.5 on Bootstrap Server: Install Additional Packages. Only steps 2 and 3 need to be followed in that procedure.



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

Step #	Procedure	Description
3	Create bridge interface	Create bridge interface to connect both management ports and setup the management bridge to support switch initialization. Note: <cne_management_ip_with_prefix> is from Installation PreFlight Checklist: Complete Site Survey Host IP Table. Row 1 CNE Management IP Addresess (VLAN 4) column. <torswitch_cnemanagementnet_vip> is from Installation PreFlight Checklist: Complete OA and Switch IP Table. \$ nmcli con add con-name mgmtBridge type bridge ifname mgmtBridge \$ nmcli con add type bridge-slave ifname eno2 master mgmtBridge \$ nmcli con add type bridge-slave ifname eno3 master mgmtBridge \$ nmcli con mod mgmtBridge ipv4.method manual ipv4.addresses 192.168.2.11/24 \$ nmcli con up mgmtBridge \$ nmcli con add type team con-name team0 ifname team0 team.runner lacp \$ nmcli con add type team-slave con-name team0-slave-1 ifname eno5 master team0 \$ nmcli con add type team-slave con-name team0-slave-2 ifname eno6 master team0 \$ nmcli con mod team0 ipv4.method manual ipv4.addresses 172.16.3.4/24 \$ nmcli con mod team0 ipv4.method manual ipv4.addresses 172.16.3.4/24 \$ nmcli con mod team0.4 ipv4.method manual ipv4.addresses <ne_management_ip_address_with_prefix> ipv4.gateway <torswitch_cnemanagementnet_vip> \$ nmcli con up team0.4</torswitch_cnemanagementnet_vip></ne_management_ip_address_with_prefix></torswitch_cnemanagementnet_vip></cne_management_ip_with_prefix>



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

Step #	Procedure	Description
4.	Edit the /etc/ xinetd.d/tftp file	Edit the /etc/xinetd.d/tfip file to enable TFTP service. Change the disable option to no, if it is set to yes. \$ vi /etc/xinetd.d/tftp # default: off # description: The tftp server serves files using the trivial file transfer \ # protocol. The tftp protocol is often used to boot diskless \ # workstations, download configuration files to network-aware printers, \ # and to start the installation process for some operating systems. service tftp { socket_type
5.	Enable tftp on the Bootstrap host.	<pre>\$ systemctl start tftp \$ systemctl enable tftp Verify tftp is active and enabled: \$ systemctl status tftp \$ ps -elf grep tftp</pre>
6.	Copy the dhcpd.conf file	Copy the dhcpd.conf file from the Utility USB in Installation PreFlight checklist: Create the dhcpd.conf File to the /etc/dhcp/ directory. \$ cp /media/usb/dhcpd.conf /etc/dhcp/
7.	Restart and enable dhcpd service.	<pre># /bin/systemctl restart dhcpd.service # /bin/systemctl enable dhcpd.service Use the systemctl status dhcpd command to verify active and enabled. # systemctl status dhcpd</pre>
8.	Copy the switch configuration and script files	Copy the switch configuration and script files from the Utility USB to directory /var/lib/tfipboot/. \$ 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/.



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

Step #	Procedure	Description
9.	Modify POAP script File.	Modify POAP script File. Change Username and password credentials used to login to the Bootstrap host.
		<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", }</password></username></pre>
		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.
10.	Modify POAP script file	Modify POAP script file md5sum by executing the md5Poap.sh script from the Utility USB created from Installation PreFlight checklist: Create the md5Poap Bash Script. # cd /var/lib/tftpboot/ # /bin/bash md5Poap.sh
11.	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 3-3 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
12	Copy the /var/lib/ tftpboot/ 93180_switch A.cfg into a file called /var/lib/ tftpboot/ conf. <switch a="" number="" serial=""></switch>	Modify the switch specific values in the /var/lib/tftpboot/conf. <switcha number="" serial=""> file, including all the values in the curly braces as following code block.</switcha>



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

Step #	Procedure	Description
		<pre>nNet_IP>#g' conf.<switcha number="" serial=""> # sed -i 's#{SQL_replication_Prefix}#<sqlreplicationnet_prefix>#g' conf.<switcha number="" serial=""> # ipcalc -n <torswitcha_sqlreplicationnet_ip <sqlreplicationnet_prefix=""> awk -F'=' '{print \$2}' # sed -i 's/{SQL_replication_Subnet}/<output as="" command="" from="" ipcalc="" sql_replication_subnet="">/' conf.<switcha number="" serial=""></switcha></output></torswitcha_sqlreplicationnet_ip></switcha></sqlreplicationnet_prefix></switcha></pre>
		<pre># sed -i 's/{CNE_Management_VIP}/ <torswitch_cnemanagementnet_vip>/g' conf.<switcha number="" serial=""> # sed -i 's/{SQL_replication_VIP}/ <torswitch_sqlreplicationnet_vip>/g' conf.<switcha number="" serial=""> # sed -i 's/{OAM_UPLINK_CUSTOMER_ADDRESS}/ <torswitcha_oam_uplink_customer_ip>/' conf.<switcha number="" serial=""></switcha></torswitcha_oam_uplink_customer_ip></switcha></torswitch_sqlreplicationnet_vip></switcha></torswitch_cnemanagementnet_vip></pre>
		<pre># sed -i 's/{OAM_UPLINK_SwA_ADDRESS}/</pre>
		<pre># ipcalc -n <torswitcha_sqlreplicationnet_ip> awk - F'=' '{print \$2}' # sed -i 's/{MySQL_Replication_SUBNET}/<output above="" appended="" command="" from="" ipcalc="" prefix="" the="" with="">/' conf.<switcha number="" serial=""></switcha></output></torswitcha_sqlreplicationnet_ip></pre>
		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 number="" serial=""></switcha></nxos_version>
		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.



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

Step #	Procedure	Description
		<pre># sed -i 's/{Allow_Access_Server}/<allow_access_server>/' conf.<switcha number="" serial=""></switcha></allow_access_server></pre>



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

Step # Procedure	Description
Step # Copy the /var/lib/ tftpboot/ 93180_switch B.cfg into a file called /var/lib /tftpboot/ conf. <switch b="" number="" serial=""></switch>	Modify the switch specific values in the /var/lib/tftpboot/conf. <switcha number="" serial=""> file, including: hostname, username/password, oam_uplink IP address, signaling_uplink IP address, access-list ALLOW_5G_XSI_LIST permit address, prefix-list ALLOW_5G_XSI. These values are contained at Installation PreFlight checklist: ToR and Enclosure Switches Variables Table and Installation PreFlight Checklist:</switcha>



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

Step #	Procedure	Description
		<pre>nNet_IP>#g' conf.<switchb number="" serial=""> # sed -i 's#{SQL_replication_Prefix}#<sqlreplicationnet_prefix>#g' conf.<switchb number="" serial=""> # ipcalc -n <torswitchb_sqlreplicationnet_ip <sqlreplicationnet_prefix=""> awk -F'=' '{print \$2}' # sed -i 's/{SQL_replication_Subnet}/<output as="" command="" from="" ipcalc="" sql_replication_subnet="">/' conf.<switchb number="" serial=""></switchb></output></torswitchb_sqlreplicationnet_ip></switchb></sqlreplicationnet_prefix></switchb></pre>
		<pre># sed -i 's/{CNE_Management_VIP}/ <torswitch_cnemanagementnet_vip>/' conf.<switchb number="" serial=""> # sed -i 's/{SQL_replication_VIP}/ <torswitch_sqlreplicationnet_vip>/' conf.<switchb number="" serial=""> # sed -i 's/{OAM_UPLINK_CUSTOMER_ADDRESS}/ <torswitchb_oam_uplink_customer_ip>/' conf.<switchb number="" serial=""></switchb></torswitchb_oam_uplink_customer_ip></switchb></torswitch_sqlreplicationnet_vip></switchb></torswitch_cnemanagementnet_vip></pre>
		<pre># sed -i 's/{OAM_UPLINK_SwA_ADDRESS}/</pre>
		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 number="" serial=""></switchb></nxos_version>
		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 number="" serial=""></switchb></allow_access_server>



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

Step #	Procedure	Description
<u> </u>		
14.	Generate the md5 checksum	Generate the md5 checksum for each conf file in /var/lib/tftpboot and copy that into a new file called conf.<switcha b="" number="" serial="">.md5</switcha> .
		<pre>\$ md5sum conf.<switcha number="" serial=""> > conf.<switcha number="" serial="">.md5 \$ md5sum conf.<switchb number="" serial=""> > conf.<switchb number="" serial="">.md5</switchb></switchb></switcha></switcha></pre>
		Serial number>.md5
15.	Verify	Make sure the file permissions are set as given below.
	the /var/lib/ tftpboot directory has the correct files.	Note : The ToR switches are constantly attempting to find and execute the poap_nexus_script.py script which uses tftp to load and install the configuration files.
	11100.	# ls -l /var/lib/tftpboot/
		total 1305096 -rw-rr 1 root root 7161 Mar 25 15:31
		conf. <switcha number="" serial=""></switcha>
		-rw-rr 1 root root 51 Mar 25 15:31 conf. <switcha number="" serial="">.md5</switcha>
		-rw-rr- 1 root root 7161 Mar 25 15:31
		conf. <switchb number="" serial=""></switchb>
		-rw-rr 1 root root 51 Mar 25 15:31 conf. <switchb number="" serial="">.md5</switchb>
		-rwxr-xr-x. 1 root root 75856 Mar 25 15:32 poap_nexus_script.py
16.	Disable	
	firewalld.	<pre>\$ systemctl stop firewalld \$ systemctl disable firewalld</pre>
		To verify: \$ systemctl status firewalld
		Once this is complete, the ToR Switches will attempt to boot from the tftpboot files automatically. Eventually the verification steps can be executed below. It may take about 5 minutes for this to complete.
17.	Un-mount the Utility USB	Un-mount the Utility USB and remove it: umount /media/usb

Verification



Table 3-4 Procedure to verify Top of Rack 93180YC-EX Switches

1. After the ToR switches configured. # ping 192.168.2.1	ep# Procedure D	escription
switches ping 192.168.2.1	<u> </u>	•
Switches from bootstrap 64 bytes from 192.168.2.1: icmp_seq=1 ttl=255 time: ms 64 bytes from 192.168.2.1: icmp_seq=2 ttl=255 time: ms 64 bytes from 192.168.2.1: icmp_seq=3 ttl=255 time: ms 64 bytes from 192.168.2.1: icmp_seq=4 ttl=255 time: ms 64 bytes from 192.168.2.1: icmp_seq=4 ttl=255 time: ms 64 bytes from 192.168.2.1: icmp_seq=4 ttl=255 time: ms 65 bytes from 192.168.2.1: icmp_seq=4 ttl=255 time: ms 66 bytes from 192.168.2.1: icmp_seq=4 ttl=255 time: ms 67 bytes from 192.168.2.2: icmp_seq=4 ttl=255 time: ms 68 bytes from 192.168.2.2: icmp_seq=1 ttl=255 time: ms 69 bytes from 192.168.2.2: icmp_seq=2 ttl=255 time: ms 69 bytes from 192.168.2.2: icmp_seq=4 ttl=255 time: ms 69 bytes from 192.168.2	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. After the ToR switches mgmted for the switches ms for the	ote: Wait till the device responds. ping 192.168.2.1 (192.168.2.1) 56(84) bytes of data. bytes from 192.168.2.1: icmp_seq=1 ttl=255 time=0.419 bytes from 192.168.2.1: icmp_seq=2 ttl=255 time=0.496 bytes from 192.168.2.1: icmp_seq=3 ttl=255 time=0.573 bytes from 192.168.2.1: icmp_seq=4 ttl=255 time=0.535 bytes from 192.168.2.1: icmp_seq=4 ttl=255 time=0.535 192.168.2.1 ping statistics packets transmitted, 4 received, 0% packet loss, time 000ms tt min/avg/max/mdev = 0.419/0.505/0.573/0.063 ms ping 192.168.2.2 (192.168.2.2) 56(84) bytes of data. bytes from 192.168.2.2: icmp_seq=1 ttl=255 time=0.572 bytes from 192.168.2.2: icmp_seq=2 ttl=255 time=0.582 bytes from 192.168.2.2: icmp_seq=2 ttl=255 time=0.582 bytes from 192.168.2.2: icmp_seq=3 ttl=255 time=0.554 bytes from 192.168.2.2: icmp_seq=4 ttl=255 time=0.554



Table 3-4 (Cont.) Procedure to verify Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
2.	Attempt to ssh to the switches with the username/ password provided in the conf files.	# ssh plat@192.168.2.1 The authenticity of host '192.168.2.1 (192.168.2.1)' can't be established. RSA key fingerprint is SHA256:jEPSMHRNg9vejiLcEvw5qprjgt +4ua9jucUBhktH520. RSA key fingerprint is MD5:02:66:3a:c6:81:65:20:2c:6e:cb: 08:35:06:c6:72:ac. Are you sure you want to continue connecting (yes/no)? yes Warning: Permanently added '192.168.2.1' (RSA) to the list of known hosts. User Access Verification Password: Cisco Nexus Operating System (NX-OS) Software TAC support: http://www.cisco.com/tac Copyright (C) 2002-2019, Cisco and/or its affiliates. All rights reserved. The copyrights to certain works contained in this software are owned by other third parties and used and distributed under their own licenses, such as open source. This software is provided "as is," and unless otherwise stated, there is no warranty, express or implied, including but not limited to warranties of merchantability and fitness for a particular purpose. Certain components of this software are licensed under the GNU General Public License (GPL) version 2.0 or GNU General Public License (LGPL) Version 2.1 or 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://opensource.org/licenses/gpl-3.0.html and http://www.opensource.org/licenses/library.txt. #



Table 3-4 (Cont.) Procedure to verify Top of Rack 93180YC-EX Switches

Step # Procedure	Description
3. Verify the running-config has all expected configurations in the conf file using the show running-config command.	# show running-config !Command: show running-config !Running configuration last done at: Mon Apr 8 17:39:38 2019 !Time: Mon Apr 8 18:30:17 2019 version 9.2(3) Bios:version 07.64 hostname 12006-93108A vdc 12006-93108A id 1 limit-resource vlan minimum 16 maximum 4094 limit-resource vrf minimum 2 maximum 4096 limit-resource port-channel minimum 0 maximum 511 limit-resource u4route-mem minimum 248 maximum 248 limit-resource u6route-mem minimum 96 maximum 96 limit-resource m6route-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



Table 3-4 (Cont.) Procedure to verify Top of Rack 93180YC-EX Switches

Verify license on the switches	In case some of the above for switches and at least NXOS license not installed or too I file, following Licensing do install license key. Then run factory default. The switches # show license Example output: # show license MDS20190215085542979. SERVER this_host ANY VENDOR cisco INCREMENT NXOS_ADVANT. uncounted \ VENDOR_STRING	_ADVANTAC ow level, contacument mention "write erase" s will go to PC lic: AGE_XF cisc	GE level act vendo oned in r and "relo OAP con	license is "In use". If or for correct license key eference section to oad" to set back to figuration again.
	Example output: # show license MDS20190215085542979. SERVER this_host ANY VENDOR cisco INCREMENT NXOS_ADVANT. uncounted \	AGE_XF cisc	o 1.0 p	permanent
	<pak></pak> " # show license usage	AD-XF D22412J2F \ ileID>20190 1 <td>\ 2150855 D> \</td> <td></td>	\ 2150855 D> \	
	# show license usage Feature Date Comments NXOS_ADVANTAGE_M4 Unused - NXOS_ADVANTAGE_XF never - NXOS_ESSENTIALS_GF Unused	No Ye	Count s -	Status Expiry In use
		# show license usage Example output: # show license usage Feature Date Comments NXOS_ADVANTAGE_M4 Unused - NXOS_ADVANTAGE_XF never - NXOS_ESSENTIALS_GF Unused -	<pre></pre>	Example output: # show license usage Feature Ins Lic Date Comments Count NXOS_ADVANTAGE_M4 No - Unused - NXOS_ADVANTAGE_XF Yes - never - NXOS_ESSENTIALS_GF No - Unused - Unused



Table 3-4 (Cont.) Procedure to verify Top of Rack 93180YC-EX Switches

Step #	Procedure	Description
5.	Verify the RMS1 can ping the CNE_Manage ment 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</torswitch_cnemanagementnet_vip></torswitch_cnemanagementnet_vip></torswitch_cnemanagementnet_vip></torswitch_cnemanagementnet_vip></torswitch_cnemanagementnet_vip></torswitch_cnemanagementnet_vip></pre>
6.	Enable customer uplink	Connect or enable customer uplink.
7.	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 ~]#</root></cne_management_ip_address></cne_management_ip_address></pre>

SNMP Trap Configuration

Table 3-5 Procedure to configure SNMP Trap

Step #	Procedure	Description
1.	SNMPv2c Configuration	When SNMPv2c configuration is needed, ssh to the two switches, run the following commands:
		These values <snmp_trap_receiver_address>and <snmp_community_string> are from Installation Preflight Checklist</snmp_community_string></snmp_trap_receiver_address>
		<pre>[root@RMS1 ~]# ssh</pre>



Table 3-5 (Cont.) Procedure to configure SNMP Trap

Step #	Procedure	Description
2.	Restrict direct access to ToR switches	In order to restrict direct access to ToR switches, IP access list is created and applied on the uplink interfaces, the following commands are needed on ToR switches:
		<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 {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_SWA_Address}/32 deny ip any {SQL_replication_SWA_Address}/32 deny ip any {SQL_replication_SWB_Address}/32 deny ip any {OAM_UPLINK_SWB_ADDRESS}/32 deny ip any {OAM_UPLINK_SWB_ADDRESS}/32 deny ip any {SIGNAL_UPLINK_SWB_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</torswitcha_cnemanagementnet_ip></user_name></pre>



Table 3-5 (Cont.) Procedure to configure SNMP Trap

Step #	Procedure	Description
3.	Traffic egress	Traffic egress out of cluster, including snmptrap traffic to SNMP trap receiver, and traffic goes to signal server:
		<pre>[root@RMS1 ~]# ssh</pre>
		ip access-list host-sigserver 10 permit ip 172.16.3.0/24 <signal server="">/32</signal>
		<pre>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-snmptrap interface Ethernet1/51 overload</pre>
		interface Vlan3 ip nat inside
		interface Ethernet1/51 ip nat outside
		interface Ethernet1/52 ip nat outside
		Run the same commands on ToR switchB

Configure Addresses for RMS iLOs, OA, EBIPA

Introduction

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

Prerequisites

Procedure OCCNE Configure Top of Rack 93180YC-EX Switches has been completed.

Limitations/Expectations

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



References

HPE BladeSystem Onboard Administrator User Guide

Steps to configure Addresses for RMS iLOs, OA, EBIPA

Table 3-6 Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
1.	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.	Subnet and conf file address	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.



Table 3-6 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
3.	Display the dhcpd leases file at /var/lib/ dhcpd/ dhcpd.leases. The DHCPD lease file will display the DHCP addresses for all RMS iLOs, Enclosure OAs.	# cat /var/lib/dhcpd/dhcpd.leases # The format of this file is documented in the dhcpd.leases(5) manual page. # This lease file was written by isc-dhcp-4.2.5 lease 192.168.20.101 { starts 4 2019/03/28 22:05:26; ends 4 2019/03/28 22:07:26; 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; cltt 4 2019/03/28 22:07:28; cltt 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; cltt 4 2019/03/28 23:03:29; cltt 4 2019/03/28 23:03:29; cltt 4 2019/03/28 23:03:29; cltt 4 2019/03/28 23:05:16; binding state free; hardware ethernet 48:df:37:7a:40:40; } lease 192.168.20.106 { starts 5 2019/03/29 11:14:04; cltt 5 2019/03/29 11:14:04; 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\000"; } lease 192.168.20.105 { starts 5 2019/03/29 15:56:23; cltt 5 2019/03/29 16:08:21; cltt 5 2019/03/29 16:08:21; clt 5 2019/03/29 16:08:21; clt 5 2019/03/29 16:08:21; cltt 5 2019/03/29 16:08:21; cltt 5 2019/03/29 16:08:21; cltt 5 2019/03/29 13:08:21; binding state free; hardware ethernet b8:83:03:47:5e:54; uid "\000\270\203\003G^129 13:08:21; cltt 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\0003Gd234\000\000\000\0000; clt 4



Table 3-6 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step # Procedure Description	
Procedure Description	La:ea:05; 52\005"; AEA05"; 30:d7:d7; 327"; DD7D7"; \$# 7; La:ea:05; 52\005"; AEA05";

Table 3-6 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
		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 18:10:40; cltt 5 2019/03/29 18:10:40; binding state active; next binding state free; rewind binding state free; rewind binding state free; rewind binding state free; hardware ethernet b8:83:03:47:5e:54; uid "\000\270\203\003G^T\000\000\000"; client-hostname "ILO2M29090048";
4.	Access RMS iLO from the DHCP address with default Administrator password. From the above dhcpd.leases file, find the IP address for the iLO name, the default username is Administrator, the password is on the label which can be pulled out from front of server.	Note: The DNS Name on the pull-out label. The DNS Name on the pull-out label should be used to match the physical machine with the iLO IP since the same default DNS Name from the pull-out label is displayed upon logging in to the iLO command line interface, as shown in the example below. # ssh Administrator@192.168.20.104 Administrator@192.168.20.104's password: User:Administrator logged-in to ILO2M2909004F.labs.nc.tekelec.com(192.168.20.104 / FE80::BA83:3FF:FE47:649C) iLO Standard 1.37 at Oct 25 2018 Server Name: Server Power: On
5.	Create RMS iLO new user. Create new user with customized username and password.	<pre>hpiLO-> create /map1/accounts1 username=root password=TklcRoot group=admin,config,oemHP_rc,oemHP_power,oemHP_vm status=0 status_tag=COMMAND COMPLETED Tue Apr 2 20:08:30 2019 User added successfully.</pre>



Table 3-6 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
6.	Disable the DHCP before able to setup static IP. Setup static failed before DHCP is disabled.	<pre>hpilo-> set /map1/dhcpendpt1 EnabledState=N0 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>
7.	Setup RMS iLO static IP address. After a while after previous step, can login back with the same address(which is static IP now) and new username/ password. If don't want to use the same address, go to next step to change the IP address.	<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(192.168.20.104 / FE80::BA83:3FF:FE47:649C) iLO Standard 1.37 at Oct 25 2018 Server Name: Server Power: On </new></new></new></pre> <pre> </pre> <pre> </pre> <pre> </pre> <pre> </pre> <pre> <pre> </pre> <pre> <pre> </pre> <pre> <pre> <pre> <pre> </pre> <pre> <pre> <pre> <pre> <pre> </pre> <pre> <pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre></pre>



Table 3-6 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
8.	Set EBIPA addresses for InterConnect Bays (Enclosure Switches).	From bootstrap server, login to OA, set EBIPA addressed for the two enclosure switches. The addresses have to be in the subnet with server team 0.2 address in order for TFTP to work. Set address for each enclosure switch, note the
		last number 1 or 2 is the interconnect bay number.
		OA-FC15B41AEA05> set ebipa interconnect 192.168.20.133 255.255.255.0 1 Entering anything other than 'YES' will result in
		the command not executing. It may take each interconnect several minutes to acquire the new settings.
		Are you sure you want to change the IP address for the specified interconnect bays? yes
		Successfully set 255.255.255.0 as the netmask for interconnect bays.
		Successfully set interconnect bay # 1 to IP address 192.168.20.133 For the IP addresses to be assigned EBIPA must be
		enabled. OA-FC15B41AEA05> set ebipa interconnect
		192.168.20.134 255.255.255.0 2 Entering anything other than 'YES' will result in
		the command not executing. It may take each interconnect several minutes to acquire the new settings.
		Are you sure you want to change the IP address for the specified interconnect bays? yes
		Successfully set 255.255.255.0 as the netmask for interconnect bays. Successfully set interconnect bay # 2 to IP
		address 192.168.20.134 For the IP addresses to be assigned EBIPA must be
		enabled.



Table 3-6 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
9.	Set EBIPA addresses for Blade Servers. Set EBIPA addressed for all the blade servers. The addresses are in the same subnet with first server team0.2 address and enclosure switches.	OA-FC15B41AEA05> set ebipa server 192.168.20.141 255.255.255.0 1-16 Entering anything other than 'YES' will result in the command not executing. Changing the IP address for device (iLO) bays that are enabled causes the iLOs in those bays to be reset. Are you sure you want to change the IP address for the specified device (iLO) bays? YES Successfully set 255.255.255.0 as the netmask for device (iLO) bays. Successfully set device (iLO) bay # 1 to IP address 192.168.20.141 Successfully set device (iLO) bay # 2 to IP address 192.168.20.142 Successfully set device (iLO) bay # 3 to IP address 192.168.20.143 Successfully set device (iLO) bay # 4 to IP address 192.168.20.144 Successfully set device (iLO) bay # 5 to IP address 192.168.20.145 Successfully set device (iLO) bay # 6 to IP address 192.168.20.146 Successfully set device (iLO) bay # 7 to IP address 192.168.20.147 Successfully set device (iLO) bay # 8 to IP address 192.168.20.147 Successfully set device (iLO) bay # 9 to IP address 192.168.20.148 Successfully set device (iLO) bay # 10 to IP address 192.168.20.150 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 192.168.20.151 Successfully set device (iLO) bay #13 to IP address 192.168.20.153 Successfully set device (iLO) bay #15 to IP address 192.168.20.154 Successfully set device (iLO) bay #15 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 #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>



Table 3-6 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

Step #	Procedure	Description
10.	Add New User for OA.	Create new user, set access level as ADMINISTRATOR, and assign access to all blades, all enclosure switches and OAs. After that, the username and password can be used to access OAs.
		OA-FC15B41AEA05> ADD USER <username> New Password: ******* Confirm : ******* User "<username>" created. You may set user privileges with the 'SET USER ACCESS' and 'ASSIGN' commands.</username></username>
		OA-FC15B41AEA05> set user access <username> ADMINISTRATOR</username>
		" <username>" has been given administrator level privileges.</username>
		OA-FC15B41AEA05> ASSIGN SERVER ALL <username></username>
		<pre><username> has been granted access to the valid requested bay(s)</username></pre>
		OA-FC15B41AEA05> ASSIGN INTERCONNECT ALL <username></username>
		<pre><username> has been granted access to the valid requested bay(s)</username></pre>
		OA-FC15B41AEA05> ASSIGN OA <username></username>
		<pre><username> has been granted access to the OA.</username></pre>
11.	From OA, go to each blade with "connect	OA-FC15B41AEA05> connect server 4
	server bay number>", add New User for each blade.	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->
		<pre>hpiLO-> create /map1/accounts1 username=root password=TklcRoot group=admin,config,oemHPE_rc,oemHPE_power,oemHPE_vm</pre>
		status=2 status_tag=COMMAND PROCESSING FAILED error_tag=COMMAND SYNTAX ERROR Tue Apr 23 16:18:58 2019 User added successfully.



Step #	Procedure	Description
12.	Change to static IP on OA. In order not reply on DHCP and make the OA address stable, change to static IP.	Note: After the following change, on the active OA (could be the bayl OA or bay2 OA), the OA session will be stuck due to the address change, make another server session ready to ssh with the new IP address and new root user. The change on the standby OA will not stuck the OA session. OA-FC15B41AEA05> SET IPCONFIG STATIC 1 192.168.20.131 255.255.255.0 Static IP settings successfully updated. These setting changes will take effect immediately. OA-FC15B41AEA05> SET IPCONFIG STATIC 2 192.168.20.132 255.255.255.0 Static IP settings successfully updated. These setting changes will take effect immediately. OA-FC15B41AEA05>
		OA-FC15B41AEAU5>

Table 3-6 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

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 3-7.
- 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 3-7 Procedure to configure the Legacy BIOS on Remaining Hosts

Step #	Procedure	Description
1.	Expose the System Configuration Utility on a RMS Host	 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. 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. The System Utilities screen is exposed in the remote console.



Table 3-7 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

Step #	Procedure	Des	cription
2.	Expose the System Utility for an Enclosure Blade	em Utility n Enclosure	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.
			a. SSH to the blade using the iLO IP address and the root user and password. This brings up the HP iLO prompt.
			\$ ssh root@ <black 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</black>
			hpiLO->
			b. Use VSP to connect to the blade remote console.
			hpiLO->vsp
			c. Power cycle the blade to bring up the System Utility for that blade.
			Note : The System Utility is a text based version of that exposed on the RMS via the KVM. The user must use the directional (arrow) keys to manipulate between selections, ENTER key to select, and ESC to go back from the current selection.
			d. Access the System Utility by hitting ESC 9.
		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
			a. Expose the System Utility by following step 1 or 2 depending on the hardware being configured.
			b. Select System Configuration
			c. Select BIOS/Platform Configuration (RBSU)
			d. Select Virtualization Options This view displays the settings for the Intel(R) Virtualization Technology (IntelVT), Intel(R) VT-d, and SR-IOV options (Enabled or Disabled). The default values for each option is Enabled.
			e. Select F10 if it is desired to save and stay in the utility or select the F12 if it is desired to save and exit to continue the current boot process.



Table 3-7 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

Step #	Procedure	Des	cription													
3.	Change over from UEFI Booting Mode to Legacy BIOS Booting Mode	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:													
			Boot Mode changes require a system reboot in order to take effect. Changing the Boot Mode can impact the ability of the server to boot the installed operating system. An operating system is installed in the same mode as the platform during the installation. If the Boot Mode does not match the operating system installation, the system cannot boot. The following features require that the server be configured for UEFI Mode: Secure Boot, IPv6 PXE Boot, Boot > 2.2 TB Disks in AHCI SATA Mode, and Smart Array SW RAID.													
			Hit the ENTER key and two selections appear: UEFI Mode(highlighted) and Legacy BIOS Mode													
		6.	Use the down arrow key to select $\bf Legacy~BIOS~Mode$ 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 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.													
			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													
		1					1								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.



Table 3-7 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

Step #	Procedure	Description
4.	Force PXE to boot from the first Embedded	Expose the System Utility by following step 1 or 2 depending on the hardware being configured.
	FlexibleLOM	2. Select System Configuration.
	HPE Ethernet 10Gb 2-port	3. Select BIOS/Platform Configuration (RBSU).
	Adapter	4. Select Boot Options . This menu defines the boot mode.
		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.
		6. Select Legacy BIOS Boot Order In the default view, the 10Gb Embedded FlexibleLOM 1 Port 1 is at the bottom of the list.
		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.
		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.
5.	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 Workload Profiles.
		Verifying Default Settings
		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 3-7 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

Step #	Procedure	cription	
6.	Disable RAID Configurations	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).	
		Note : There may be more than one RAID Array set up. This procedure should be repeated for any RAID configuration.	
		Expose the System Utility by following the hardware being configured.	step 1 or 2 depending on
		Select System Configuration.	
		Select Embedded RAID 1 : HPE Smar 10.	t Array P408i-a SR Gen
		Select Array Configuration.	
		Select Manage Arrays.	
		Select Array A (or any designated Arrare more than one).	ay Configuration if there
		Select Delete Array . A warning is displated following:	ayed indicating the
		Deletes an Array. All the data drives that are part of deleted Also if the deleted array is th controller, the controller sett and its default configuration i	array will be lost. e only one on the ings will be erased
		Hit ENTER , the changes are submitted a Successful is displayed.	and Delete Array
		Hit ENTER to go back to the main men	ı for the HPE Smart Array.
		Select F10 if it is desired to save and start F12 it is desired to save and exit to continuous.	



Table 3-7 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

Step #	Procedure	Description
7.	Enable the Primary and	This steps provide necessary to configure the primary and secondary bootable devices for a Gen10 Server.
	Secondary Boot Devices	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.
		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.
		 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.
		 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 Utilitiesmenu.
		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 OCCNE Configure Top of Rack 93180YC-EX Switches has been completed.
- Procedure OCCNE 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 3-8 Procedure to configure enclosure switches

Step #	Procedure	Description
1.	Copy the 6127XLG configuration file	Copy the 6127XLG configuration file from the Utility USB (See Installation PreFlight checklist: Create the OA 6127XLG Switch Configuration File) to the /var/lib/tftpboot directory on the Installer Bootstrap Host and verify it exists and the permissions. \$ cp /media/usb/6127xlg_irf.cfg /var/lib/tftpboot/6127xlg_irf.cfg \$ ls -l /var/lib/tftpboot/ total 1305096
		-rw-rr 1 root root 311 Mar 25 08:41 6127xlg_irf.cfg
2.	Modify the switch specific values in the /var/lib/ tftpboot/ 6127xlg_irf.cfg file.	These values are contained at Installation PreFlight checklist: Create the OA 6127XLG Switch Configuration File from column Enclosure_Switch. \$ cd /var/lib/tftpboot \$ sed -i 's/{switchname}/ <switch_name>/' 6127xlg_irf.cfg \$ sed -i 's/{admin_password}/<admin_password>/' 6127xlg_irf.cfg \$ sed -i 's/{user_name}/<user_name>/' 6127xlg_irf.cfg \$ sed -i 's/{user_password}/<user_password>/' 6127xlg_irf.cfg</user_password></user_name></admin_password></switch_name>



Table 3-8 (Cont.) Procedure to configure enclosure switches

Step #	Procedure	Description
Step # 3.	Procedure Access the InterConnect Bay1 6127XLG	Access the InterConnect Bayl 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: [HPE] [HPE] save The current configuration will be written to the device. Are you sure? [Y/N]: Before pressing ENTER you must choose 'YES' or 'NO'[Y/N]:y Please input the file name(*.cfg)[flash:/startup.cfg] (To leave the existing filename unchanged, press the enter key): User can leave this default startup.cfg unchanged, or change to another name. The cfg file will be used for next reboot. \$ ssh <0a username>@<0a address> If it shows standby, ssh to the other OA address. OA-FC15B41AEA05> connect interconnect 1 <hpe>system-view System View: return to User View with Ctrl+Z. (Note: Run the following commands:) irf member 1 priority 32 interface range Ten-GigabitEthernet 1/0/17 to Ten-</hpe>
		GigabitEthernet 1/0/20 shutdown
		quit
		irf-port 1/1
		port group interface Ten-GigabitEthernet1/0/17
		port group interface Ten-GigabitEthernet1/0/18
		port group interface Ten-GigabitEthernet1/0/19
		port group interface Ten-GigabitEthernet1/0/20
		quit
		interface range Ten-GigabitEthernet 1/0/17 to Ten-GigabitEthernet 1/0/20



Table 3-8 (Cont.) Procedure to configure enclosure switches

Step #	Procedure	Description
		undo shutdown
		quit
		save
		irf-port-configuration active
4.	Access the	Access the InterConnect Bay2 6127XLG switch to re-number to IRF 2.
	InterConnect Bay2 6127XLG	OA-FC15B41AEA05> connect interconnect 2
		<hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre><hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre></hpre>
		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</hpe>
		Start to check configuration with next startup configuration file, please waitDONE!
		This command will reboot the device. Continue? [Y/N]:Y
		Now rebooting, please wait
		System is starting



Table 3-8 (Cont.) Procedure to configure enclosure switches

Step #	Procedure	Description
5.	Configure the IRF on Bay2 6127XLG switch	After rebooting, the interfaces will begin with number 2 such as Ten-GigabitEthernet2/0/17, Ten-GigabitEthernet2/1/5. Run the following commands:
		system-view
		interface range Ten-GigabitEthernet 2/0/17 to Ten-GigabitEthernet 2/0/20
		shutdown
		quit
		irf-port 2/2
		port group interface Ten-GigabitEthernet2/0/17
		port group interface Ten-GigabitEthernet2/0/18
		port group interface Ten-GigabitEthernet2/0/19
		port group interface Ten-GigabitEthernet2/0/20
		quit
		interface range Ten-GigabitEthernet 2/0/17 to Ten-GigabitEthernet 2/0/20
		undo shutdown
		quit
		save
		irf-port-configuration active
6.	Run "reboot" command on both switches	<pre><hpe>reboot Start to check configuration with next startup configuration file, please waitDONE! This command will reboot the device. Continue? [Y/N]:Y Now rebooting, please wait</hpe></pre>
		System is starting



Table 3-8 (Cont.) Procedure to configure enclosure switches

Step #	Procedure	Description		
7.	Verify the IRF for the 6127XLG switches.		d, verify IRF is working with both mem o switches, which form IRF to act as on	
		<hpe>system-view</hpe>		
		System View: retur	n to User View with Ctrl+Z.	
		[HPE]display irf o	configuration	
		MemberID NewID Port2	IRF-Port1	IRF-
		1 1	Ten-GigabitEthernet1/0/17	disable
			Ten-GigabitEthernet1/0/18	
			Ten-GigabitEthernet1/0/19	
			Ten-GigabitEthernet1/0/20	
		2 2 GigabitEthernet2/0	disable 0/17	Ten-
		GigabitEthernet2/0	/18	Ten-
		GigabitEthernet2/0	1/19	Ten-
		GigabitEthernet2/0	7/20	Ten-
		[HPE]		



Table 3-8 (Cont.) Procedure to configure enclosure switches

Step #	Procedure	Description
8.	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::</hpe></pre>
		<pre><hpe>system-view System View: return to User View with Ctrl+Z. [HPE]configuration replace file flash:/startup.cfg Current configuration will be lost, save current configuration? [Y/N]:N Now replacing the current configuration. Please wait Succeeded in replacing current configuration with the file flash:/startup.cfg. [<switch_name>]save flash:/startup.cfg The current configuration will be saved to flash:/ startup.cfg. Continue? [Y/N]:Y flash:/startup.cfg exists, overwrite? [Y/N]:Y Now saving current configuration to the device. Saving configuration flash:/startup.cfg.Please wait Configuration is saved to device successfully. [<switch_name>]</switch_name></switch_name></hpe></pre>

Bastion Host Installation

This section outlines the use of the Installer Bootstrap Host to provision RMS2 with an operating system and configure it to fulfill the role of Database Host. Subsequently, steps are provided to provision virtual machines that run MySQL services, DBMS, and serve the role as

Bastion Host. After the Bastion Host is provisioned, it is used to complete the installation of OCCNE.

Install Host OS onto RMS2 from the Installer Bootstrap Host (RMS1)

Introduction

These procedures provide the steps required to install the OL7 image onto the RMS2 via the Installer Bootstrap Host using a occne/provision container. Once completed, RMS2 includes all necessary rpm updates and tools necessary to Install the Bastion Host. All commands are executed from RMS1.

Prerequisites

- All procedures in OCCNE Initial Configuration are complete.
- The Utility USB is available containing the necessary files as mentioned in Installation PreFlight checklist.

Limitations and Expectations

All steps are executable from a SSH application (putty) connected laptop accessible via the Management Interface.



Procedures

 $Table \ 3-9 \quad Procedure \ to \ install \ the \ OL7 \ image \ onto \ the \ RMS2 \ via \ the \ installer \ bootstrap \ host$

Step #	Procedure	Description
1.	Copy the Necessary Files from the Utility USB to Support the OS Install	This procedure is used to provide the steps for copying all supporting files from the Utility USB to the appropriate directories so that the Provision Container successfully installs OL7 onto RMS2. Note: The <code>cluster_name</code> field is derived from the <code>occne_cluster_name</code> field in the hosts.ini file.
		1. Create the directories needed on the Installer Bootstrap Host.
		<pre>\$ mkdir /var/occne \$ mkdir /var/occne/<cluster_name> \$ mkdir /var/occne/<cluster_name>/yum.repos.d</cluster_name></cluster_name></pre>
		2. Mount the Utility USB. Note: Instructions for mounting a USB in Linux are at: Installation of Oracle Linux 7.5 on Bootstrap Host: Install Additional Packages. Only follow steps 1-4 to mount the USB.
		3. Copy the hosts.ini file (created using procedure: OCCNE Inventory File Preparation) into the /var/occne/ <cluster_name>/ directory. This hosts.ini file defines RMS2 to the Provision Container running the provision image downloaded from the repo.</cluster_name>
		<pre>\$ cp /media/usb/hosts.ini /var/occne/ <cluster_name>/hosts.ini</cluster_name></pre>
		4. Update the hosts.ini file to include the ToR host_net (vlan3) VIP for NTP clock synchronization. Use the ToR VIP address as defined in procedure: Installation PreFlight Checklist: Complete OA and Switch IP SwitchTable as the NTP source.
		<pre>\$ vim /var/occne/<cluster_name>/hosts.ini</cluster_name></pre>
		Update the ntp_server field with the VIP address.
		5. Copy the customer specific ol7-mirror.repo and the docker-ce-stable repo on the Utility USB to the Installer Bootstrap Host.
		This is the .repo file created by the customer that provides access to the onsite (within their network) repositories needed to complete the full deployment of OCCNE 1.2 and to install docker-ce onto the Installer Bootstrap Host.
		<pre>\$ cp /media/usb/ol7-mirror.repo /var/occne/ <cluster_name>/yum.repos.d/ol7-mirror.repo \$ cp /media/usb/ol7-mirror.repo /etc/yum.repos.d/ ol7-mirror.repo \$ cp /media/usb/docker-ce-stable.repo /etc/ yum.repos.d/docker-ce-stable.repo</cluster_name></pre>
		6. If still enabled from procedure: OCCNE Installation of Oracle Linux 7.5 on Bootstrap Host, the /etc/yum.repos.d/Media.repo is to be disabled.



Table 3-9 $\,$ (Cont.) Procedure to install the OL7 image onto the RMS2 via the installer bootstrap host

Step #	Procedure	Description
		<pre>\$ mv /etc/yum.repos.d/Media.repo /etc/yum.repos.d/ Media.repo.disable</pre>
2.	Copy the OL7 ISO to the Installer Bootstrap Host	The iso file should be accessible from a Customer Site Specific repository. This file should be accessible because the ToR switch configurations were completed in procedure: OCCNE Configure Top of Rack 93180YC-EX Switches.
		Copy from RMS1, the OL7 ISO file to the /var/occne directory. The example below uses OracleLinux-7.5-x86_64-disc1.iso. Note : If the user copies this ISO from their laptop then they must use an application like WinSCP pointing to the Management Interface IP.
		<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</path_to_iso></site_specific_address></usr></pre>
3.	Install Docker onto the Installer Bootstrap Host	Use YUM to install docker-ce onto the installer Bootstrap Host. YUM should use the existing <customer_specific_repo_file>.repo in the /etc/yum.repos.d directory.</customer_specific_repo_file>
		\$ yum install docker-ce-18.06.1.ce-3.el7.x86_64



 $\begin{tabular}{ll} Table 3-9 & (Cont.) \end{tabular} Procedure to install the OL7 image onto the RMS2 via the installer bootstrap host \\ \end{tabular}$

Step #	Procedure	Des	cription
4.	Set up access to the Docker Registry on the Installer Bootstrap Host	1.	Add an entry to the /etc/hosts file on the Installer Bootstrap Host to provide a name mapping for the docker registry using the hosts.ini file fields occne_private_registry and occne_private_registry_address in OCCNE Inventory File Preparation. <occne_private_registry_address> <occne_private_registry> Example:10.75.200.217 reg-1</occne_private_registry></occne_private_registry_address>
		2.	Create the /etc/docker/daemon.json file on the Installer Bootstrap Host. Add an entry for the insecure-registries for the docker registry.
			<pre>\$ mkdir /etc/docker \$ vi /etc/docker/daemon.json Enter the following:</pre>
			{
			<pre>"insecure-registries": ["<occne_private_registry>:<occne_private_registry _port="">"]</occne_private_registry></occne_private_registry></pre>
			}
			Example:
			cat /etc/docker/daemon.json
			{
			"insecure-registries": ["reg-1:5000"]
			}
			To Verify:
			ping <occne_private_registry></occne_private_registry>
			Example:
			# ping reg-1
			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
		3.	Create the docker service http-proxy.conf file.

Table 3-9 $\,$ (Cont.) Procedure to install the OL7 image onto the RMS2 via the installer bootstrap host

Step #	Procedure	Description
		<pre>\$ mkdir -p /etc/systemd/system/docker.service.d/</pre>
		<pre>\$ vi /etc/systemd/system/docker.service.d/http- proxy.conf</pre>
		Add the following:
		[Service]
		<pre>Environment="NO_PROXY=<occne_private_registry_addr ess="">,<occne_private_registry>, 127.0.0.1,localhost"</occne_private_registry></occne_private_registry_addr></pre>
		Example:
		[Service]
		Environment="NO_PROXY=10.75.200.217,reg-1,127.0.0.1,localhost"
		4. Start the docker daemon
		<pre>\$ systemctl daemon-reload \$ systemctl restart docker \$ systemctl enable docker</pre>
		Verify docker is running: \$ ps -elf grep docker \$ systemctl status docker
5.	Setup NFS on the Installer Bootstrap Host	Run the following commands (assumes nfs-utils has already been installed in procedure: OCCNE Installation of Oracle Linux 7.5 on Bootstrap Host: Install Additional Packages).
		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: Complete Site Survey Host Table.
		<pre>\$ echo'/var/occne 172.16.3.4/24(ro,no_root_squash)'>> /etc/exports \$ systemctl start nfs-server \$ systemctl enable nfs-server Verify nfs is running: \$ ps -elf grep nfs</pre>
		\$ systemctl status nfs-server
6.	Set up the Boot Loader on the Installer Bootstrap Host	Execute the following commands: \$ mkdir -p /var/occne/pxelinux \$ mount -t iso9660 -o loop /var/occne/ OracleLinux-7.5-x86_64-discl.iso /mnt \$ cp /mnt/isolinux/initrd.img /var/occne/pxelinux
		\$ cp /mnt/isolinux/vmlinuz /var/occne/pxelinux



Table 3-9 $\,$ (Cont.) Procedure to install the OL7 image onto the RMS2 via the installer bootstrap host

Step #	Procedure	Description
7.	Verify and Set the PXE Configuration File Permissions on the Installer Bootstrap Host	Each file configured in the step above must be open for read and write permissions. \$ chmod 777 /var/occne/pxelinux \$ chmod 777 /var/occne/pxelinux/vmlinuz \$ chmod 777 /var/occne/pxelinux/initrd.img
8.	Disable DHCP and TFTP on the Installer Bootstrap Host	The TFTP and DHCP services running on the Installer Bootstrap Host may still be running. These services must be disabled. \$ systemctl stop dhcpd \$ systemctl disable dhcpd \$ systemctl stop tftp \$ systemctl disable tftp
9.	Disable SELINUX	SELINUX must be set to permissive mode. In order to successfully set the SELINUX mode, a reboot of the system is required. The getenforce command is used to determine the status of SELINUX. \$ getenforce active If the output of this command displays active, change it to permissive by editing the /etc/selinux/config file. \$ vi /etc/selinux/config Change the SELINUX variable to passive: SELINUX=permissive save the file Reboot the system: reboot



Table 3-9 $\,$ (Cont.) Procedure to install the OL7 image onto the RMS2 via the installer bootstrap host

Step #	Procedure	Description
10.	Execute the OS Install on RMS2 from the Installer Bootstrap Host	1. Run the docker commands below to perform the OS install. docker runrmnetwork hostcap- add=NET_ADMIN -v /var/occne/ <cluster_name>/:/host -v /var/occne/:/var/occne:rw -e "OCCNEARGS= limit <db-2_node_name>, localhostskip-tags "datastore,vms_provision,yum_configure" <image_name>:<image_tag> docker runrmnetwork hostcap- add=NET_ADMIN -v /var/occne/<cluster_name>/:/host -v /var/occne/:/var/occne:rw -e "OCCNEARGS= limit <db-2_node_name>,localhosttags yum_configure" <image_name>:<image_tag> docker runrmnetwork hostcap- add=NET_ADMIN -v /var/occne/<cluster_name>/:/host</cluster_name></image_tag></image_name></db-2_node_name></cluster_name></image_tag></image_name></db-2_node_name></cluster_name>
		-v /var/occne/:/var/occne:rw -e "OCCNEARGS= limit <db-2_node_name>,localhosttags datastore" <image_name>:<image_tag> docker runrmnetwork hostcap- add=NET_ADMIN -v /var/occne/<cluster_name>/:/host -v /var/occne/:/var/occne:rw -e "OCCNEARGS= limit <db-2_node_name>,localhosttags vms_provision" <image_name>:<image_tag></image_tag></image_name></db-2_node_name></cluster_name></image_tag></image_name></db-2_node_name>
		Example: docker run -itrmnetwork hostcap- add=NET_ADMIN -v /var/occne/ rainbow.lab.us.oracle.com/:/host -v /var/ occne/:/var/occne:rw -e "OCCNEARGS=limit db-2.rainbow.lab.us.oracle.com,localhostskip- tags "datastore,vms_provision,yum_configure"" 10.75.200.217:5000/occne/provision:1.2.0 docker run -itrmnetwork hostcap- add=NET ADMIN -v /var/occne/
		add=NET_ADMIN -V /Var/occne/ rainbow.lab.us.oracle.com/:/host -v /var/ occne/:/var/occne:rw -e "OCCNEARGS=limit db-2.rainbow.lab.us.oracle.com,localhosttags yum_configure" 10.75.200.217:5000/occne/provision: 1.2.0 docker run -itrmnetwork hostcap- add=NET_ADMIN -v /var/occne/ rainbow.lab.us.oracle.com/:/host -v /var/ occne/:/var/occne:rw -e "OCCNEARGS=limit db-2.rainbow.lab.us.oracle.com,localhosttags datastore" 10.75.200.217:5000/occne/provision: 1.2.0



Table 3-9 (Cont.) Procedure to install the OL7 image onto the RMS2 via the installer bootstrap host

Step #	Procedure	Description
		docker runrmnetwork hostcap- add=NET_ADMIN -v /var/occne/ <cluster_name>/:/host -v /var/occne/:/var/occne:rw -e "OCCNEARGS= limit db-2.rainbow.lab.us.oracle.com,localhost tags vms_provision" 10.75.200.217:5000/occne/ provision:1.2.0 2. Verify that Bastion host vm is installed by logging into RMS2.</cluster_name>
		Run the following in case of failure to Login ONLY: \$ virsh list
		Output Example- Id Name State
		 11 bastion-1.rainbow.lab.us.oracle.com running
		<pre>\$ virsh destroy bastion-1.rainbow.lab.us.oracle.com \$ virsh undefine bastion-1.rainbow.lab.us.oracle.com \$ virsh list</pre>
		Note: virsh list after destroy and undefine should not return any name in the list
		3. Execute the virt-install command on RMS2
		<pre>virt-installname bastion_hostmemory 8192 vcpus 2metadata description="Bastion Host" \</pre>
		disk path=/var/lib/libvirt/ images/bastion_host.qcow2,size=300 \network bridge=teambr0 network bridge=vlan2-brnetwork bridge=vlan4-brgraphics none
		After the VM creation completes, the login prompt appears which allows the user to login to the Bastion Host.



Configuration of the Bastion Host

Introduction

This procedure details the steps necessary to configure the Bastion Host onto RMS2 during initial installation. This VM is used for host provisioning, MySQL Cluster, and installing the BL460c hosts with kubernetes and the common services.

Prerequisites

- Procedure Install Host OS onto RMS2 from the Installer Bootstrap Host (RMS1)has been completed.
- 2. All the release docker images for 1.2.0 (k8s, configure etc) and kubespray base image occne/kubespray:2.10.3.1 should be pulled on to Bastion host.
- 3. All the hosts servers details are captured in Inventory File Preparation.
- Host names and IP Address, network information assigned to this VM is captured in the Installation Preflight Checklist
- 5. Yum repository mirror is setup and accessible by Bastion host.
- 6. Http server is setup and has kubernetes binaries, helm charts on a server with address that is accessible by Bastion Host.
- 7. Docker registry is setup to an address that is reachable by the Bastion host.
- 8. It is assumed that an apache http server (as part of the mirror creation) is created outside of bastion host that supports yum mirror, helm charts and Kubernetes Binaries. (This can be different so directories to copy static content to Bastion host must be verified before starting the rsync procedure).

Limitations and Expectations

All steps are executable from a SSH application (putty) connected laptop accessible via the Management Interface.

References

- Oracle YUM mirroring directions: https://www.oracle.com/technetwork/articles/servers-storage-admin/yum-repo-setup-1659167.html
- 2. https://docs.docker.com/registry/deploying/
- 3. https://computingforgeeks.com/how-to-configure-ntp-server-using-chrony-on-rhel-8/

Procedure

These procedures detail the steps required to configure the existing Bastion Host (Management VM).



Table 3-10 Procedure to configure Bastion Host

Step #	Procedure	Description
1.	Create the /var/ occne/ <cluster_na me=""> directory on the Bastion Host</cluster_na>	Create the directory using the occne_cluster_name variable contained in the hosts.ini file. \$ mkdir /var/occne \$ mkdir /var/occne/ <cluster_name></cluster_name>
2.	Copy the host.ini file to the /var/ occne/ <cluster_na me=""> directory</cluster_na>	Copy the hosts.ini file (created using procedure: OCCNE Inventory File Preparation) into the /var/occne/ <cluster_name>/ directory from RMS1 (this procedure assumes the same hosts.ini file is being used here as was used to install the OS onto RMS2 from RMS1. If not then the hosts.ini file must be retrieved from the Utility USB mounted onto RMS2 and copied from RMS2 to the Bastion Host). This hosts.ini file defines each host to the OS Installer Container running the osinstall image downloaded from the repo. \$ scp root@172.16.3.4:/var/occne/<cluster_name>/hosts.ini /var/occne/<cluster_name>/hosts.ini file requires a "/" to be added to the entry for the occne_helm_images_repo. vim (or use vi) and edit the hosts.ini file and add the "/"to the occne_helm_images_repo entry. occne_helm_images_repo='bastion-1:5000 -> occne_helm_images_repo='bastion-1:5000/</cluster_name></cluster_name></cluster_name>
3.	Check and Disable Firewall	Check the status of the firewall. If active then disable it. \$ systemctl status firewalld \$ systemctl stop firewalld \$ systemctl disable firewalld To verify: \$ systemctl status firewalld



Table 3-10 (Cont.) Procedure to configure Bastion Host

Step #	Procedure	Description	
4.	D:	Create the local YUM repo mirror file in etc/yum.repos.d and add the docker repo mirror. Follow procedure: OCCNE Artifact Acquisition and Hosting	
		2. Disable the public repo	
	Registry on Bastion Host VM	<pre>\$ mv /etc/yum.repos.d/public-yum-ol7.repo /etc/yum.repos.d/ public-yum-ol7.repo.disabled</pre>	
		Install necessary packages from the yum mirror on Bastion Host	
		<pre>\$ yum install rsync \$ yum install createrepo yum-utils \$ yum install docker-ce-18.06.1.ce-3.el7.x86_64 \$ yum install nfs-utils \$ yum install httpd \$ yum install chrony -y \$ yum install curl -y \$ yum install nghttp2 -y</pre>	
		3. Copy the yum mirror contents from the remote server where the yum mirror is deployed, this can be done in the following way: Get the ip address of the yum mirror from the yum repo file.	
		Create an apache http server on Bastion host.	
		<pre>\$ systemctl start httpd \$ systemctl enable httpd \$ systemctl status httpd</pre>	
		Retrieve the latest rpm's from the yum mirror to /var/www/yum on the Bastion host using reposync: Run following repo sync commands to get latest packages on Bastion host	
		<pre>\$ reposync -g -l -d -mrepoid=local_ol7_x86_64_addons newest-onlydownload-metadatadownload_path=/var/www/ html/yum/OracleLinux/OL7/ \$ reposync -g -l -d -mrepoid=local_ol7_x86_64_UEKR5 newest-onlydownload-metadatadownload_path=/var/www/ html/yum/OracleLinux/OL7/ \$ reposync -g -l -d -mrepoid=local_ol7_x86_64_developernewest-onlydownload-metadatadownload_path=/var/www/ html/yum/OracleLinux/OL7/ \$ reposync -g -l -d -m repoid=local_ol7_x86_64_developer_EPELnewest-only download-metadatadownload_path=/var/www/html/yum/</pre>	
		OracleLinux/OL7/ \$ reposync -g -l -d -mrepoid=local_ol7_x86_64_ksplice newest-onlydownload-metadatadownload_path=/var/www/ html/yum/OracleLinux/OL7/	
		<pre>\$ reposync -g -l -d -mrepoid=local_o17_x86_64_latest newest-onlydownload_metadatadownload_path=/var/www/ html/yum/OracleLinux/OL7/</pre>	
		After the above execution, you will be able to see the directory structure in with all the repo id's in /var/www/html/yum/OracleLinux/OL7/. Rename the repositories in OL7/ directory:	



Table 3-10 (Cont.) Procedure to configure Bastion Host

Step #	Procedure	Description
		Note : download_path can be changed according to the folder structure required. Change the names of the copied over folders to match the base url.
		<pre>\$ cd /var/www/html/yum/OracleLinux/OL7/ \$ mv local_o17_x86_64_addons addons \$ mv local_o17_x86_64_UEKR5 UEKR5 \$ mv local_o17_x86_64_developer developer \$ mv local_o17_x86_64_developer_EPEL developer_EPEL \$ mv local_o17_x86_64_ksplice ksplice \$ mv local_o17_x86_64_latest latest</pre>
		Run following createrepo commands to create repo data for each repository channel on Bastion host yum mirror:
		<pre>\$ createrepo -v /var/www/html/yum/OracleLinux/OL7/addons \$ createrepo -v /var/www/html/yum/OracleLinux/OL7/UEKR5 \$ createrepo -v /var/www/html/yum/OracleLinux/OL7/developer \$ createrepo -v /var/www/html/yum/OracleLinux/OL7/ developer_EPEL \$ createrepo -v /var/www/html/yum/OracleLinux/OL7/ksplice \$ createrepo -v /var/www/html/yum/OracleLinux/OL7/latest</pre>
		5. Get Docker-ce and gpg key from mirror Execute the following rsync command:
		<pre>\$ rsync -avzh <login-username>@<ip address="" of="" repo="" server="">:<centos directory="" folder="" path=""> /var/www/html/yum/</centos></ip></login-username></pre>
		6. Create http repository configuration to retrieve Kubernetes binaries and helm binaries/charts on Bastion Host if on a different server Given that the Kubernetes binaries have been created outside of bastion host as part of the procedure of setting up artifacts and repositories, kubernetes/helm binaries and helm charts have to be copied using rsync command to the bastion host. Example below should copy all of the contents from a folder to the static content render folder of the http server on the bastion host:
		<pre>\$ rsync -avzh <login-username>@<ip address="" of="" repo="" server="">:<copy address="" directory="" from=""> /var/www/html</copy></ip></login-username></pre>
		Note : Above is an example directory for an apache folder, if there is another http server running, the directory may be different
		7. Setup Helm and initiate on Bastion Host Get Helm version onto the repo server which is accessible by Bastion Host
		<pre>\$ mkdir helmtar \$ cd helmtar \$ wget https://storage.googleapis.com/kubernetes-helm/helm- v2.9.1-linux-amd64.tar.gz</pre>
		Login to Bastion Host and run the following commands:
		Create a temporary directory on the bastion host. It does not really matter where this directory is created but it must have read/write/execute privileges.



Table 3-10 (Cont.) Procedure to configure Bastion Host

Step #	Procedure	Description
		<pre>\$ mkdir /var/occne/<cluster_name>/tmp1 \$ rsync -avzh <login-username>@<ip address="" of="" repo="" server="">:helmtar /var/occne/<cluster_name>/tmp1 \$ cd /var/occne/<cluster_name>/tmp1/helmtar \$ tar -xvf helm-v2.9.1-linux-amd64.tar.gz \$ rm -f helm-v2.9.1-linux-amd64.tar.gz \$ mv linux-amd64 helm \$ cd helm</cluster_name></cluster_name></ip></login-username></cluster_name></pre>
		<pre># Run the following command in the charts directory of the http server on bastion host to create index.yaml file so that helm chart can be initialized \$./helm repo index <path_to_helm_charts_directory_bastion_host> # initialize helm \$./helm initclient-onlystable-repo-url <bastion_host_occne_helm_stable_repo_url></bastion_host_occne_helm_stable_repo_url></path_to_helm_charts_directory_bastion_host></pre>



Table 3-10 (Cont.) Procedure to configure Bastion Host

Step #	Procedure	Description		
5.	Create a docker registry on Bastion Host	1.	Pull registry Image from Docker registry onto Bastion host to run a registry locally Add the registry IP and port to the /etc/docker/daemon.json file. Create the file if not currently existing.	
			<pre>"insecure-registries" : ["<server_docker_registry_address>:<port>"] }</port></server_docker_registry_address></pre>	
		2.	Start docker: Start the docker daemon.	
			<pre>\$ systemctl daemon-reload \$ systemctl restart docker \$ systemctl enable docker</pre>	
			Verify docker is running: \$ ps -elf grep docker \$ systemctl status docker	
			While creating the docker registry on a server outside of the bastion host, there is no tag added to the registry image and the image is also not added to the docker registry repository of that server. Manually tag the registry image and push it as one of the repositories on the docker registry server:	
			<pre>\$ docker tag registry:<tag> docker_registry_address>:<port>/registry:<tag></tag></port></tag></pre>	
		3.	Push the tagged registry image customer to docker registry repository on server accessible by Bastion Host:	
			<pre>\$ docker push <docker_registry_address>:<port>/ registry:<tag></tag></port></docker_registry_address></pre>	
		4.	Login into Bastion host and pull the registry image onto Bastion Host from customer registry setup on server outside of bastion host	
			<pre>\$ docker pullall-tags <docker_registry_address>:<port>/ registry</port></docker_registry_address></pre>	
		5.	Run Docker registry on Bastion Host	
			<pre>\$ docker run -d -p 5000:5000restart=alwaysname registry registry:<tag></tag></pre>	
			This runs the docker registry local to Bastion host on port 5000.	
		6.	Get docker images from docker registry to Bastion Host docker registry Pull all the docker images from Docker Repository Requirements to the local Bastion Host repository:	
			<pre>\$ docker pullall-tags <docker_registry_address>:<port>/ <image_names_from_attached_list></image_names_from_attached_list></port></docker_registry_address></pre>	
			Note : If following error is encountered during the pull of images "net/http: request canceled (Client.Timeout exceeded while awaiting headers)" from the internal docker registry, edit http-proxy.conf and add the docker registry address to NO_PROXY environment variable	



 Table 3-10
 (Cont.) Procedure to configure Bastion Host

Step #	Procedure	Description		
		7. Tag Images		
		<pre>\$ docker tag <docker_registry_address>:<port>/ <imagename>:<tag> <bastion_host_docker_registry_address>:<port>/ <image_names_from_attached_list></image_names_from_attached_list></port></bastion_host_docker_registry_address></tag></imagename></port></docker_registry_address></pre>		
		Example:		
		<pre>\$ docker tag 10.75.207.133:5000/jaegertracing/jaeger- collector:1.9.0 10.75.216.125:5000/jaegertracing/jaeger- collector</pre>		
		8. Push the images to local Docker Registry created on the Bastion host Create a daemon.json file in /etc/docker directory and add the following to it:		
		{ "insecure-registries" : [" <bastion_host_docker_registry_address>:<port>"] }</port></bastion_host_docker_registry_address>		
		Restart docker:		
		<pre>\$ systemctl daemon-reload \$ systemctl restart docker \$ systemctl enable docker</pre>		
		To verify: \$ systemctl status docker \$ docker push <bastion_host_docker_registry_address>:<port>/ <image_names_from_attached_list></image_names_from_attached_list></port></bastion_host_docker_registry_address>		
6.	Setup NFS	Run the following commands:		
	on the Bastion Host	<pre>\$ echo '/var/occne 172.16.3.100/24(ro,no_root_squash)' >> /etc/ exports \$ systemctl start nfs-server \$ systemctl enable nfs-server</pre>		
		Verify nfs is running: \$ ps -elf grep nfs \$ systemctl status nfs-server		



Table 3-10 (Cont.) Procedure to configure Bastion Host

Step	Procedure	Description
Step# 7.	Setup the Bastion Host to clock off the ToR Switch	The ToR acts as the NTP source for all hosts. Update the chrony.conf file with the source NTP server by adding the VIP address of the ToR switch from: OCCNE 1.0 Installation PreFlight Checklist: Complete OA and Switch IP SwitchTable as the NTP source. \$ vim /etc/chrony.conf Add the following line at the end of the file: server 172.16.3.1 chrony was installed in the first step of this procedure. Enable the service. \$ systemctl enablenow chronyd \$ systemctl status chronyd chrony was installed in the first step of this procedure. Enable the service. \$ systemctl enablenow chronyd\$ systemctl status chronyd Execute the chronyc sources -v command to display the current status of NTP on the Bastion Host. The S field should be set to * indicating NTP sync. \$ chronyc sources -v 210 Number of sources = 1 Source mode '^' = server, '=' = peer, '#' = local clock. / Source state '*' = current synced, '+' = combined , '-' = not combined,
		Change field: ntp_server=' <tor switch="" vip="">'</tor>



Software Installation Procedures - Automated Installation

Using either of the management hosts created in the previous procedures, the installer performs a sequence of procedures to complete the automated installation of the CNE environment.

The following procedures provide the necessary steps to install the different images required during OCCNE Software Installation.

Oracle Linux OS Installer

This procedure provides the steps required to install the OL7 image onto all hosts via the Bastion Host using a occne/provision container. Once completed, all hosts include all necessary rpm updates and tools necessary to run the k8-install procedure.

Prerequisites:

- 1. All procedures in OCCNE Installation of the Bastion Host are complete.
- 2. The Utility USB is available containing the necessary files as per: Installation PreFlight checklist: Miscellaneous Files.

Limitations and Expectations

All steps are executable from a SSH application (putty) connected laptop accessible via the Management Interface.

References

https://docs.ansible.com/ansible/latest/user_guide/intro_patterns.html



Table 3-11 Procedure to run the auto OS-installer container

Step #	Procedure	Description			
1.	Initial Configuration on the Bastion Host to Support the OS	Note : The cluster_name field is derived from the hosts.ini file field:			
	Install	1. Log into the Bastion Host using the IP supplied from: Installation PreFlight Checklist: Complete VM IP Table			
		2. Create the directories needed on the Bastion Host.			
		<pre>\$ mkdir /var/occne \$ mkdir /var/occne/<cluster_name> \$ mkdir /var/occne/<cluster_name>/yum.repos.d</cluster_name></cluster_name></pre>			
		3. Copy the hosts.ini file (created using procedure: Inventory File Preparation) into the /var/occne/ <cluster_name>/ directory from RMS1 (this procedure assumes the same hosts.ini file is being used here as was used to install the OS onto RMS2 from RMS1. If not then the hosts.ini file must be retrieved from the Utility USB mounted onto RMS2 and copied from RMS2 to the Bastion Host).</cluster_name>			
		This hosts.ini file defines each host to the OS Installer(Provision) Container running the provision image downloaded from the repo.			
		<pre>\$ scp root@172.16.3.4:/var/occne/<cluster_name> /var/ occne/<cluster_name>/hosts.ini</cluster_name></cluster_name></pre>			
		4. Update the repository fields in the hosts.ini file to reflect the changes from procedure: OCCNE Configuration of the Bastion Host. The fields listed must reflect the new Bastion Host IP (172.16.3.100) and the names of the repositories.			
		<pre>\$ vim /var/occne/<cluster_name>/hosts.ini</cluster_name></pre>			
		Update the following fields with the new values from the configuration of the Bastion Host. ntp_server			
		occne_private_registry occne_private_registry_address occne_private_registry_port occne_k8s_binary_repo occne_helm_stable_repo_url occne_helm_images_repo docker_rh_repo_base_url docker_rh_repo_gpgkey			
		Example: ntp_server='172.16.3.1' occne_private_registry=registry occne_private_registry_address='10.75.207.133' occne_private_registry_port=5000 occne_k8s_binary_repo='http://10.75.207.133/ binaries/' occne_helm_stable_repo_url='http://10.75.207.133/ helm/'			
		<pre>occne_private_registry_port=5000 occne_k8s_binary_repo='http://10.75.207.133/ binaries/' occne_helm_stable_repo_url='http://10.75.207.1</pre>			



Table 3-11 (Cont.) Procedure to run the auto OS-installer container

Step #	Procedure	Description			
		docker_rh_repo_base_url=http://10.75.207.133/yum/centos/7/updates/x86_64/docker_rh_repo_gpgkey=http://10.75.207.133/yum/centos/RPM-GPG-CENTOS			
		Comment out the fields under ilo network configuration, management network configuration and signalling network configuration for mysql ndb replication in hosts.ini			
		Only keep values for ilo_vlanid, mgmnt_vlan_id and signal_vlan_id and comment all other variables			
2.	Copy the OL7 ISO to the Bastion Host The iso file is normally accessible from a Customer Site Specific repository. It is accessible because the ToR switch configuration: completed in procedure: OCCNE Configure Top of Rack 93180 Switches. For this procedure the file has already been copied to to occne directory on RMS2 and can be copied to the same director Bastion Host.				
		Copy from RMS2, the OL7 ISO file to the /var/occne directory. The example below uses OracleLinux-7.5-x86 64-disc1.iso.			
		Note: If the user copies this ISO from their laptop then they must use an application like WinSCP pointing to the Management Interface IP.			
		\$ scp root@172.16.3.5:/var/occne/OracleLinux-7.5-x86_64-disc1.iso /var/occne/OracleLinux-7.5-x86_64-disc1.iso			
3.	Set up the Boot	Execute the following commands:			
	Loader on the Bastion Host	Note : The iso can be unmounted after the files have been copied if the user wishes to do so using the command: umount /mnt.			
		<pre>\$ mkdir -p /var/occne/pxelinux \$ mount -t iso9660 -o loop /var/occne/OracleLinux-7.5- x86_64-disc1.iso /mnt \$ cp /mnt/isolinux/initrd.img /var/occne/pxelinux \$ cp /mnt/isolinux/vmlinuz /var/occne/pxelinux</pre>			
4.	Verify and Set the PXE Configuration File Permissions on	Each file configured in the step above must be open for read and write permissions. \$ chmod 777 /var/occne/pxelinux \$ chmod 777 /var/occne/pxelinux/vmlinuz			
	the Bastion Host	\$ chmod 777 /var/occne/pxelinux/initrd.img			



Table 3-11 (Cont.) Procedure to run the auto OS-installer container

Step #	Procedure	Des	cription
5.			The customer specific .repo files on the bastion_host must be copied to the /var/occne/ <cluster_name> /yum.repos.d directory and updated to reflect the URL to the bastion host. This file is transferred to /etc/ yum.repos.d directory on the host by ansible after the host has been installed but before the actual yum update is performed.</cluster_name>
			<pre>\$ cp /etc/yum.repos.d/<*.repo> /var/occne/ <cluster_name>/yum.repos.d/.</cluster_name></pre>
		2.	Edit each .repo file in the /var/occne/ <cluster_name>/yum.repos.d directory and update the baseurl IP of the repo to reflect the IP of the bastion_host.</cluster_name>
			<pre>\$ vim /var/occne/<cluster_name>/yum.repos.d/ <repo_name>.repo</repo_name></cluster_name></pre>
			Example:
			<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/</pre>
			<pre>gpgcheck=1 gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY enabled=1 proxy=_none_</pre>
			Change the IP address of the baseurl IP: 10.75.155.195 to the bastion host ip: 172.16.3.100.
			The URL may have to change based on the configuration of the customer repos. That cannot be indicated in this procedure.



Table 3-11 (Cont.) Procedure to run the auto OS-installer container

Step #	Procedure	Description
6.	Execute the OS Install on the Hosts from the Bastion Host	This step requires executing docker run for four different Ansible tags. Note: The <image_name>:<image_tag> represent the images in the docker image registry accessible by Bastion host. Run the docker command below to create a container running bash. This command must include the -it option and the bash executable at the end of the command. After execution of this command the user prompt will be running within the container. docker runrmnetwork hostcap-add=NET_ADMIN - v /var/occne/<cluster_name>/:/host -v /var/occne/:/var/occne:rw -e "OCCNEARGS=skip-tags datastore,vms_provision,yum_configure" <image_name>:<image_tag> docker runrmnetwork hostcap-add=NET_ADMIN - v /var/occne/<cluster_name>/:/host -v /var/occne/:/var/occne:rw -e "OCCNEARGS=tags yum_configure" <image_name>:<image_tag> Example: docker run -itrmnetwork hostcap-add=NET_ADMIN - v /var/occne/rainbow.lab.us.oracle.com/:/host -v /var/occne/:/var/occne/:/var/occne/:/var/occne/:var/occne/:var/occne/:var/occne/:var/occne/:var/occne/:var/occne/:var/occne/:var/occne/provision:1.2.0 docker run -itrmnetwork hostcap-add=NET_ADMIN - v /var/occne/rainbow.lab.us.oracle.com/:/host -v /var/occne/:var/occne/rainbow.lab.us.oracle.com/:/host -v /var/occne/:var/occne/rainbow.lab.us.oracle.com/:/host -v /var/occne/:var/occne/rainbow.lab.us.oracle.com/:/host -v /var/occne/:/var/occne/rainbow.lab.us.oracle.com/:/host -v /var/occne/:/var/occne/rainbow</image_tag></image_name></cluster_name></image_tag></image_name></cluster_name></image_tag></image_name>



Table 3-11 (Cont.) Procedure to run the auto OS-installer container

Step #	Procedure	Description
7.	Update	Execute the following steps on the Master Nodes
	Network configuration	1. cd /etc/sysconfig/network-scripts
	on Master node	
		2. Edit file ifcfg-team0 (using vi).
		3. Comment out the field BRIDGE using the "#" char.
		4. Save and exit out of ifcfg-team0.
		5. Edit file ifcfg-teambr0.
		6. Capture the following lines from the ifcfg-teambr0 file and then comment these out:
		a. IPADDR= <address></address>
		b. GATEWAY= <address></address>
		c. DNS1= <address></address>
		d. PREFIX= <number></number>
		7. Edit ifcfg-team0 and insert these lines at the end of the ifcfg-team0 file.
		8. Once done restart the network on all the master node by executing the command: service network restart.
		9. Check if all the master nodes are reachable from bastion host using the command: ssh -i /var/occne/ <cluster_name>/.ssh/occne_id_rsa admusr@172.16.3.x</cluster_name>
		10. Edit file /etc/ssh/sshd_config and set value for UseDNS to no
		11. Save and restart ssh service by executing the command: service sshd restart
8.	Update Network configuration	Execute the following steps on the Worker Nodes
		1. cd /etc/sysconfig/network-scripts
	on Worker	2. Edit file ifcfg-team0 and uncomment the GATEWAY field.
	node	3. Save and exit the file.
		4. Restart the network by executing the command: service network restart
		5. Check if all the worker nodes are reachable from bastion host using the command: ssh -i /var/occne/ <cluster_name>/.ssh/occne_id_rsa admusr@172.16.3.x</cluster_name>
		6. Edit /etc/ssh/sshd_config file and set value for UseDNS to no
		7. Save and restart ssh service by executing the command: service sshd restart



Step # **Procedure** Description 9. Re-instantiate Run the following commands on RMS1 host OS: the management \$ sudo su link bridge on \$ nmcli con add con-name mgmtBridge type bridge RMS1 ifname mgmtBridge \$ nmcli con add type bridge-slave ifname eno2 master \$ 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 Verify access to the ToR switches management ports. \$ ping 192.168.2.1 \$ ping 192.168.2.2

Table 3-11 (Cont.) Procedure to run the auto OS-installer container

Database Tier Installer

This procedure documents the steps for installing the MySQL Cluster on VM's. Here VM's will be created manually using the virt-install CLI tool; MySQL Cluster will be installed using the db-install docker container.

For Installing the MySQL Cluster on these VM's requires an use of the an inventory file (hosts.ini) where all the MySQL node IP Address are configured. This Inventory file provides the db-install docker container with all the necessary information about the MySQL cluster.

MySQL Cluster will be installed using the MySQL Cluster Manager binary release which includes MySQL NDB Cluster version. Download MySQL Cluster Manager version as specified in the Installation Preflight Checklist.

In OCCNE platform, all the NF's will need a database to store application data, so MySQL Cluster is installed for storing all the application and config data for NF's. For installing MySQL Cluster, VM's will be created in kubernetes master nodes and Database Servers as configured in the Inventory File Template file.

Prerequisites

Below are list of prerequisites required for creating the VM's and installing the MySQL Cluster.

- 1. VM's needed for installing the MySQL Cluster will be created as part of the VM creation procedures Install VMs for MySQL Nodes and Management Server.
- SSH keys generated during the host provisioning in /var/occne/<cluster_name> directory, these SSH keys will be configured in these VM's as part of the OCCNE Install VMs for MySQL Nodes and Management Server, so that db-install container can install these VM's with the MySQL Cluster software.
- 3. The host running the docker image must have docker installed.
- 4. A defined and installed site hosts.ini inventory file should also be present..
- 5. Download MySQL Cluster Manager software as specified in Installation PreFlight Checklist and place it in the /var/occne directory in bastion host(Management VM).



Limitations and Expectations

- 1. db-install container will deploy MySQL cluster in these VM's as per configuration provided in the Inventory File Preparation file.
- 2. The steps below will install different MySQL Cluster nodes(Management nodes, Data nodes and SQL nodes) in these VM's.

References

- 1. MySQL NDB Cluster: https://dev.mysql.com/doc/refman/5.7/en/mysql-cluster.html
- 2. MySQL Cluster Manager: https://dev.mysql.com/doc/mysql-cluster-manager/1.4/en/

Steps to perform OCCNE Database Tier Installer

Table 3-12 OCCNE Database Tier Installer

Step #	Procedure	Description
1	Login in to the Management Node	Login in to the Management Node using the IP address noted in the Installation PreFlight Checklist



Table 3-12 (Cont.) OCCNE Database Tier Installer

Step #	Procedure	Description			
2	Check if MySQL Cluster node VM's are created	Check if MySQL Cluster node VM's are created in the Kuberentes Master nodes and Storage Hosts.			
П		Check MySQL Management Nodes are created in the Kubernetes master nodes			
		a. Login to kubernetes Master node			
		<pre>\$ ssh -i .ssh/occne_id_rsa admusr@10.72.216.XXX \$ sudo su</pre>			
		 Check if MySQL Management Node is installed in this Kubernetes Master Node. 			
		\$\$ virsh listall Id Name State			
		db-3.rainbow.lab.us.oracle.com running			
		Check in all other Kubernetes Master nodes.			
		2. Check MySQL Data nodes and MySQL SQL Nodes are installed in Storage Hosts			
		a. Login to Storage Host			
		<pre>\$ ssh -i .ssh/occne_id_rsa admusr@10.72.216.XXX \$ sudo su</pre>			
		b. Check if MySQL Data nodes and SQL Nodes are installed in this Kubernetes Master Node.			
		\$ virsh listall Id Name State			
		 Id Name State			
		6 db-6.rainbow.lab.us.oracle.com running 7 db-8.rainbow.lab.us.oracle.com running 8 db-10.rainbow.lab.us.oracle.com running Check in other Storage hosts.			
3	Create MySQL Cluster Node VM	If VM's are not created then follow below steps to create MySQL Cluster Node VM's in kubernetes Master nodes and Storage Hosts. Execute Install VMs for MySQL Nodes and Management Server for Installing VM's.			
4	Configure occne_mysqlndb_ DataMemory variable in the hosts.ini file	Check the /var/occne/ <cluster_name> directory which has been created during the os install procedure, as specified in the Oracle Linux OS Installer, this directory consists of the hosts.ini inventory file and SSH keys generated during the os-install, which will be used by db-install container to install MySQL Cluster. Configure occne_mysqlndb_DataMemory variable in the hosts.ini file as</cluster_name>			
		documented in Inventory File Preparation, value for this variable can be obtained from the Install VMs for MySQL Nodes and Management Server.			



Table 3-12 (Cont.) OCCNE Database Tier Installer

Step #	Procedure	Description			
5	Note down the db install container	Note down the db install container name as specified in the manifest.			
	name	Container Name	db_install_container_na me	db_install:0.1.0-beta.3	
		Note: This container will be used in the next step while running the db install container which will install MySQL Cluster			
6	Configure the MySQL NDB Cluster with predefined database	To configure the MySQL NDB Cluster with predefined database, table and configurations related to different NF's, the .sql scripts can be cop in to the initdb_once and initdb directories. The .sql scripts in the initdb_once directory will be deleted after db_install container execute the .sql scripts. \$ sudo su \$ mkdir -p /var/occne/ <cluster_name>/initdb_once \$ mkdir -p /var/occne/<cluster_name>/initdb \$ cp <source_path>/initdb_once/*.sql /var/occne/ <cluster_name>/initdb_once/ \$ cp <source_path>/initdb/*.sql /var/occne/ <cluster_name>/initdb/</cluster_name></source_path></cluster_name></source_path></cluster_name></cluster_name>			
		Example:			
<pre>\$ sudo su \$ mkdir -p /var/occne/rainbow/initdb_or \$ mkdir -p /var/occne/rainbow/initdb \$ cp /var/occne/nfscripts/initdb_once/' <cluster_name>/initdb_once/ \$ cp /var/occne/nfscripts/initdb/*.sql <cluster_name>/initdb/</cluster_name></cluster_name></pre>				/*.sql /var/occne/	
			ts should be copied in to the bastion host, so the can create the databases, tables and configurations in uster.		



Table 3-12 (Cont.) OCCNE Database Tier Installer

Step #	Procedure	Description
7	Run db-install container	The db-install container will install MySQL Cluster on VM's configured in the host.ini inventory file. All the above steps should be performed before running the db-install container. Replace <customer_repo_location> and <db_install_container_name> in below docker command and docker-compose.yaml file.</db_install_container_name></customer_repo_location>
		1. Using docker
		<pre>\$ docker run -itnetwork hostcap-add=NET_ADMIN \ -v /var/occne/<cluster_name>:/host \ -v /var/occne:/var/occne:rw \ <customer_repo_location>/<db_install_container_name></db_install_container_name></customer_repo_location></cluster_name></pre>
		For Example:
		<pre>\$ docker run -itnetwork hostcap-add=NET_ADMIN</pre>
		-v /var/occne/rainbow/:/host \ -v /var/occne:/var/occne:rw \ reg-1:5000/db_install:1.2.0
		2. Using docker-compose
		a. Create a docker-compose.yaml file in the /var/occne/ <cluster_name> directory.</cluster_name>
		Using docker-compose
		<pre>\$ vi docker-compose.yaml db_install_<cluster_name>: net: host stdin_open: true tty: true image: <customer_repo_location>/ <db_install_container_name> container_name: <cluster_name>_db_installer cap_add:</cluster_name></db_install_container_name></customer_repo_location></cluster_name></pre>
		Note : In above docker-compose.yaml file cluster_name should be replaced with the cluster directory name.
		b. Run the docker-compose yaml file Running docker-compose
		<pre>\$ docker-compose runrm db_install_<cluster_name></cluster_name></pre>
		For example: If the directory name created as OccneCluster then cluster_name should be replaced with "OccneCluster".
		<pre>\$ docker-compose runrm db_install_<cluster_name></cluster_name></pre>



Table 3-12 OCCNE Database Tier Installer

Step #	Procedure	Description
		db_install container will take around 5 to 10 mins for installing the MySQL Cluster nodes in these VM's, After db_install container is completed MySQL DB is installed in the VM's as configured in the hosts.ini file.
8	Test the MySQL	Test the MySQL Cluster by executing the following command:
	Cluster	<pre>\$ docker run -itnetwork hostcap-add=NET_ADMIN \ -v /var/occne/<cluster_name>:/host \ -v /var/occne:/var/occne:rw \ <customer_repo_location>/<db_install_container_name> \ /test/cluster_test 0</db_install_container_name></customer_repo_location></cluster_name></pre>
		For Example:
		<pre>\$ docker run -itnetwork hostcap-add=NET_ADMIN \ -v /var/occne/rainbow:/host \ -v /var/occne:/var/occne:rw \ reg-1:5000/db_install:1.2.0 \ /test/cluster_test</pre>
9	Login to the each of the MySQL SQL nodes and change the MySQL root user password	As part of the installation of the MySQL Cluster, db_install container will generate 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.
		1. Login to MySQL SQL Node VM.
		2. Login to mysql client as a root user.
		\$ sudo su \$ mysql -h 127.0.0.1 -uroot -p
		3. Enter expired random password for mysql root user stored in the /var/occnedb/mysqld_expired.log file:
		<pre>\$ mysql -h 127.0.0.1 -uroot -p Enter password:</pre>
		4. Change Root Password:
		<pre>\$ mysql> ALTER USER 'root'@'localhost' IDENTIFIED BY '<new_password>'; \$ mysql> FLUSH PRIVILEGES;</new_password></pre>
		Perform this step for all the remaining SQL nodes.
		Note: Here NEW_PASSWORD is the password of the mysql root user.

Uninstall MySQL Cluster Manager and MySQL NDB Cluster



Table 3-13 Uninstall MySQL Cluster Manager and MySQL NDB Cluster

Step	Procedure	Description
1	Login in to the Bastion host	Login in to the Bastion host, IP address of this bastion host is noted in the Installation PreFlight Checklist
2	Run db_install	Run db_install docker container
		<pre>\$ docker run -itnetwork hostcap-add=NET_ADMIN \ -v /var/occne/<cluster_name>:/host \ -v /var/occne:/var/occne:rw \ <customer_repo_location>/ <db_install_container_name> /bin/bash</db_install_container_name></customer_repo_location></cluster_name></pre>
		For Example:
		<pre>\$ docker run -itnetwork hostcap-add=NET_ADMIN \ -v /var/occne/rainbow/:/host \ -v /var/occne:/var/occne:rw \ reg-1:5000/db_install:1.2.0 /bin/bash</pre>
3	Uninstall MYSQL	Un-install MySQL NDB Cluster.
	NDB Cluster	<pre>\$ ansible-playbook -i /host/hosts.inibecomebecome- user=root \</pre>
		private-key /host/.ssh/occne_id_rsa /occne/cleanup.yaml
		<pre>\$ exit</pre>

OCCNE Kubernetes Installer

These procedures provide the steps required to install the K8's image onto all hosts via the Bastion Host using a occne/k8s_install container. Once completed, configure procedure can be run.

Prerequisites

- 1. All the hosts servers where this VM is created are captured in Inventory File Preparation.
- Kubespray base image should be available on bastion host with tag 2.10.3.1 docker pull cgbu-dev-docker.dockerhub-den.oraclecorp.com/occne/kubespray:2.10.3.1
- **3.** Retag Kubespray base image to winterfell:5000/occne/kubespray:2.10.3.1 on bastion host Example-

docker tag cgbu-dev-docker.dockerhub-den.oraclecorp.com/occne/kubespray:2.10.3.1 winterfell:5000/occne/kubespray:2.10.3.1



For executing k8s install in this release we need to retag kubespray base image to winterfell:5000/occne/kubespray:2.10.3.1 on bastion host but in future releases this will change.



- 4. Host names and IP Address, network information assigned to this VM is captured in the Installation PreFlight Checklist
- 5. Cluster Inventory File and SSH Keys are present in the cluster_name folder in var/occne directory
- **6.** A docker image for 'k8s_install' must be available in the docker registry accessible by Bastion host. Installation Procedure

Limitations and Expectations

All steps are executable from a SSH application (putty) connected laptop accessible via the Management Interface.

Steps to Perform OCCNE Kubernetes Installer

Table 3-14 Procedure to install OCCNE Kubernetes

Step #	Procedure	Description
	Initial Configuration on the Bastion Host to Support the Kubernetes Install	1. Log into the Bastion Host using the IP supplied from: Installation PreFlight Checklist: Complete VM IP Table 2. Verify the entries in the hosts.ini file for occne_private_registry, occne_private_registry_address, occne_private_registry_port and occne_k8s_binary_repo are correct. The fields listed must reflect the new Bastion Host IP and the names of the repositories correctly.
2	Execute the Kubernetes Install on the Hosts from the Bastion Host	Note: The cluster_name field is derived from the hosts.ini file field: occne_cluster_name. The <image_name>:<image_tag> represent the images in the Bastion Host docker image registry as set up in procedure: Configuration of the Bastion Host.</image_tag></image_name>
		Start the k8s_install container with a bash shell: Example: \$ docker runrmnetwork hostcap-add=NET_ADMIN - v /var/occne/ <cluster_name>/:/host -v /var/occne/:/var/ occne:rw -e "OCCNEARGS=<k8s_args>" <docker_registry>/ <image_name>:<image_tag> For example: \$ docker runrm -itnetwork hostcap-add=NET_ADMIN -v /var/occne/rainbow.lab.us.oracle.com/:/host -v /var/ occne/:/var/occne:rw -e OCCNEARGS="-vv" 10.75.200.217:5000/k8s_install:1.2.0</image_tag></image_name></docker_registry></k8s_args></cluster_name>



Table 3-14 (Cont.) Procedure to install OCCNE Kubernetes

Step #	Procedure	Description
3	Update the \$PATH Environment Variable to access the kubectl command from the kubectl.sh script	On the Bastion Host, edit the /root /.bashrc file file. Update the PATH variable in that file. %% On the bastion Host edit file /root/.bash_profile. # .bash_profile # Get the aliases and functions if [-f ~/.bashrc]; then . ~/.bashrc file # User specific environment and startup programs PATH=\$PATH:\$HOME/bin export PATH
		<pre>%% Update the following to the PATH variable: PATH=\$PATH:/var/occne/<cluster_name>/artifacts %% Save the file and source the .bash_profile file: source /root/.bash_profile %% Execute the following to verify the \$PATH has been updated. echo \$PATH /usr/local/sbin:/usr/local/bin:/sbin:/bin:/usr/ sbin:/usr/bin:/var/occne/rainbow.lab.us.oracle.com/ artifacts</cluster_name></pre>
		<pre>%% Make sure the permissions on the /var/occne/ rainbow.lab.us.oracle.com/artifacts/kubectl.sh and /var/ occne/rainbow.lab.us.oracle.com/artifacts/kubectl files are set correctly: -rwxr-xr-x. 1 root root 248122280 May 30 18:23 kubectl -rwxr-xr-x. 1 root root 112 May 30 18:44 kubectl.sh %% If not run the following command: chmod +x kubectl</pre>
4	Run Kubernetes Cluster Tests	For verification of k8s installation, run docker command in the k8s_install /test/cluster_test.



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 3-15 Procedure to configure MetalLB pools and peers

mb_configmap.ya the Preflight Checklist Note: The name "signaling" is prone to different spellings (UK vs US)	Step #	Procedure	Description
mb_configmap.ya ml file mb_configmap.ya ml file mb_configmap.ya ml file Note: The name "signaling" is prone to different spellings (UK vs US therefore pay special attention to how this signaling pool is referenced configInline: peers:	1.	and address	peers (ToRswitchA_Platform_IP, ToRswitchB_Platform_IP) and address groups for each address pool. Address-pools list the IP
addresses: - ' <metallb_oam_subnet_ip_range>'</metallb_oam_subnet_ip_range>	2.	mb_configmap.ya	Note: The name "signaling" is prone to different spellings (UK vs US), therefore pay special attention to how this signaling pool is referenced. configInline: peers:

OCCNE Automated Initial Configuration

Introduction

Common Services typically refers to the collection of various components deployed to OCCNE. The Common services are major functions in action which are able to perform logging, tracing, and metric collection of the cluster. To monitor the cluster and to raise alerts when an anomaly occurs or when a potential failure is round the corner. The below procedure are used to install the common services.



Prerequisites

- 1. All procedures in Kubernetes Installer is complete.
- 2. The bastion host running the docker image must have docker installed.
- 3. A defined and installed site hosts in file should also be present. Check Inventory File Preparation for instructions for developing this file.
- 4. A defined and installed site mb_configmap.yaml file should also be present in the same directory as hosts.ini. Check Populate the MetalLB Configuration File for instructions for developing this file.
- 5. A docker image named 'occne/configure' must be available in the customer repository. Installation Procedure
- **6.** SNMP trap receiver should be setup in cluster reachable network (Optional: Required only for SNMP support)

Limitations and Expectations

All steps are executable from a SSH application (putty) connected laptop accessible via the Management Interface.

Procedure Steps

Table 3-16 Procedure to install common services

Step #	Procedure	Description
1.	Configure variables	The following variables should be configured according to customer needs, these variables can be modified to point to customer-specific repositories and other configurations needing to be done in the hosts.ini files as documented in the Inventory File Template
		 <occne_helm_stable_repo_url></occne_helm_stable_repo_url> <occne_helm_images_repo></occne_helm_images_repo>
		3. <occne_snmp_notifier_destination> (Optional: Only for SNMP Support): Update this parameter with the address of SNMP trap receiver in the format <ip address="">:<port>, where <ip address=""> is IP address of host running SNMP trap receiver and <port> is port of SNMP trap receiver service. For Ex: "127.0.0.1:162" is Default.</port></ip></port></ip></occne_snmp_notifier_destination>



Table 3-16 (Cont.) Procedure to install common services

Step #	Procedure	Description
2.	Run configure image	Run the configure image using the below command. After "configure:" keyword put the tag of your latest pulled image. \$ docker runrm -v / <path_to_cluster>/<cluster_name>:/ host <customer-provided_repository_location>/occne/ configure:<release_tag> Example:</release_tag></customer-provided_repository_location></cluster_name></path_to_cluster>
		\$ docker runrmnetwork hostcap-add=NET_ADMIN - v /var/occne/rainbow.lab.us.oracle.com/:/host -v /var/ occne/:/var/occne:rw -e "OCCNEARGS=limit host_hp_gen_10[0:7],localhost" 10.75.200.217:5000/ configure:1.2.0 Note: Replace the <release_tag> after "configure:" image name with the latest build tag.</release_tag>
3.	Verify the services	After the above command successfully completes, the services deployed and exposed can be verified by using the below commandon the Bastion Host. \$ kubectl.sh get servicenamespace= <namespace_value> or \$ kubectl.sh get serviceall-namespaces Example: kubectl.sh get serviceall-namespaces kubectl.sh get servicenamespace=default kubectl.sh get servicenamespace=kube-system kubectl.sh get srvicesnamespace=occne-infra</namespace_value>



Table 3-16 (Cont.) Procedure to install common services

Step #	Procedure	Description								
4.	Sample Output		To verify if the above command was executed successfully and whether the services was installed properly, the output shown below can be taken as reference.							
		NAMESPACE NAME TYPE PORT(S) AGE	CLUSTER-IP	EXTERNAL-IP						
		default kubernetes	10.233.0.1	<none></none>	443 /					
		TCP,9153/TCP	10.233.0.3	<none></none>	53/UDP,53/ 2d22h					
		kube-system deploy 10.233.49.220		Cl 44134/	usterIP					
		client	occne-elastic-el ClusterIP		15d 32.176					
		<none> TCP occne-infra</none>	9200/ occne-elastic-el	asticsearch-	2d					
		discovery <none></none>	ClusterIP 9300/		2d					
		occne-infra	occne-elastic-exusterIP 10.2							
		TCP occne-infra grafana	occne-	I	2d LoadBalancer					
		_	10.75.207.165		2d					
		kibana 10.233.35.35			oadBalancer					
		TCP occne-infra server	occne-metrics-	ClusterIF	2d					
		10.233.33.55 TCP occne-infra	<none> occne-prometheus</none>	443/	2d					
		alertmanager 10.233.17.109 TCP	_	LoadBalancer 80:30894/	2d					
		occne-infra headless	occne-prometheus ClusterIP	-alertmanager None						
			80/TCP,6783/ occne-prometheus		2d					
		metrics	ClusterIP	None						



Table 3-16 (Cont.) Procedure to install common services

Step #	Procedure	Description			
		<none></none>	80/		0.1
		TCP occne-infra	occne-prome		2d
		exporter <none></none>	9100/	ClusterIP	10.233.32.132
		TCP occne-infra	oggne-prome	theug-	2d
		pushgateway	occiie-broille	ClusterI	P
		10.233.12.53 TCP	<none></none>	9091/	 2d
		occne-infra	occne-prome	theus-	
		server		LoadBala	ncer
		10.233.46.125	10.75.207	.164 80:3146	•
		TCP			2d
		occne-infra agent	occne-trace	r-jaeger- ClusterIP	10.233.47.176
		<none> TCP</none>	5775/UDP,	6831/UDP,6832/ 2d	
		occne-infra collector		ClusterIP	
		<none></none>		,14268/TCP,941 2d	1/
		occne-infra query 10.75.207.162			r 10.233.16.217
		TCP	00.31020/		99
5. 🗆	Remove configurati	If a node fail during removed using the		n deployment, the	configuration can be
	on, if node fails.	v /var/occne/<	cluster_nam CCNEARGS=	e>/:/host -v / limit <host_fi< td=""><td>add=NET_ADMIN - var/occne/:/var/ lter>,localhost</td></host_fi<>	add=NET_ADMIN - var/occne/:/var/ lter>,localhost
			ainbow.lab. cne:rw -e " [0:7],local	us.oracle.com/ OCCNEARGS=li hosttags re	

NF Installation in the cluster

Introduction

This section explains the steps to be followed by NF teams to access OCCNE cluster for installing NF applications in OCCNE cluster.

Procedure Steps



Table 3-17 Procedure to install NF in OCCNE Cluster

Steps	Procedure	Description
1.	Access Bastion Host of OCCNE cluster	Login to the Bastion Host with user account - "admusr" using the IP supplied from: Installation Preflight Checklist
2.	Configure Bastion Host to access the cluster	1. The cluster details/configs are present in /var/ occne/ folder. Modify the bashrc file using the below command.
		<pre># su to root user \$ sudo su</pre>
		<pre># cd to /root directory \$ cd /root</pre>
		<pre># edit bashrc file to add cluster details \$ vi .bashrc</pre>
		2. Add the below 2 lines to the above bashre file by substituting the cluster_name with the required cluster.
		<pre># Add below cluster details to bashrc file \$ export KUBECONFIG=/var/occne/ <cluster_name>/artifacts/admin.conf</cluster_name></pre>
		3. Source bashre file to reflect the changes made. Command to source bashre file:
		\$source.bashrc
3.	Install NF applications using HELM	Run the HELM install command to install NF using HELM chart. Note: The following command is a sample install command, for complete NF installation procedure, refer to specific NF Installation Guide.
		<pre># cd to the application directory \$ cd <directory application="" for="" nf=""> # HELM chart should be present in the repository with the chart name specified in the below command \$ helm install <chart_name>name <nf_name>namespace <namespace_value></namespace_value></nf_name></chart_name></directory></pre>
		Example: \$ cd /home/admusr/ocscp/ocscp- pkg-1.2.0.0.0 \$ /usr/local/bin/helm install -f ocscp_values.yamlname ocscp ocscp namespace ocscp



Table 3-17 (Cont.) Procedure to install NF in OCCNE Cluster

Steps	Procedure	Des	scription		
4. Verif	Verify NF Installation	1.	Once NF application is installed using the above HELM Install command, Run the below command to check the status of NF application. Verify the NF application is deployed correctly:		ow oplication.
			<pre>\$ kubectl.sh g namespace=<nam< pre=""></nam<></pre>	=	
		2.	successfully and w	ove command was evhether the NF Appl the output shown be Output:	ication was
			NAME READY STA AGE <nf_name>1/</nf_name>		
		3.		eployed and expose the below command oyed correctly:	
			<pre>\$ kubectl.sh g namespace=<nam< pre=""></nam<></pre>		
		4.	successfully and w	ove command was evhether the NF servi the output shown be	ce was
			NAME TYPE EXTERNAL-IP <nf_name></nf_name>	CLUSTER-IP PORT(S)	AGE
			ClusterIP <none></none>	10.233.26.162 9200/TCP	8d
			Example: ocscp worker	scp-	
			LoadBal 10.233.28.222 8001:31323/TCP	10.75.207.193	
			TCP ocscp svc	scpc-config-	2d
			LoadBalancer 10.75.207.166 TCP	10.233.41.199 8081:30067/	
			2d	l	
			-	scpc-	
			pilot		0 152
			ClusterI		
			<none> 15007/TCP,1501 TCP 2d</none>	15003/TCP,150 0/TCP,8080/TCP,	
				scpc-soothsayer	-



 Table 3-17 (Cont.) Procedure to install NF in OCCNE Cluster

Steps	Procedure	Description	
		ClusterIP	10.233.7.220
		<none></none>	8084/TCP,8080/TCP,
		8082/TCP,808	3/TCP,8888/
		TCP	2d



4

Post Installation Activities

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.

Prerequisities

- 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.	Run the commands to get the load-balancer IP address and port number for Kibana Web Interface.	# LoadBalancer ip address of the kibana service is retrieved with below command \$ export KIBANA_LOADBALANCER_IP=\$(kubectl get services occne-kibananamespace 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-kibananamespace occne-infra -o jsonpath="{.spec.ports[*].port}") # Complete url for accessing kibana in external browser \$ echo http://\$KIBANA_LOADBALANCER_IP: \$KIBANA_LOADBALANCER_PORT http://10.75.182.51:80 Launch the Browser and navigate to http:// \$KIBANA_LOADBALANCER_IP: \$KIBANA_LOADBALANCER_PORT(e.g.: http://10.75.182.51:80 in the example above) received in the output of the above commands.



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

Step No.	Procedure	Description
2.	Using Kibana verify Log and	1. Navigate to "Management" Tab in Kibana.
	Tracer data is stored in Elasticsearch	2. Click on "Index Patterns". You should be able to see the two patterns as below which confirms Log and Tracer data been stored in Elastic-Search successfully.
		a. jaeger-*
		b. logstash-*
		3. Type logstash* in the index pattern field and wait for few seconds.
		4. Verify the "Success" message and index pattern "logstash-YYYY.MM.DD" appeared as highlighted in the bottom red box . Click on " Next step "
		5. Select "I don't want to use the Time Filter" and click on "Create index pattern"
		6. Ensure the Web page having the indices appear in the main viewer frame
		7. Click on " Discover " Tab and you should be able to view raw Log records.
		8. Repeat steps 3-6 using "jaeger*" instead of "logstash* to ensure the data is stored in elastic search.
3.	Verify Elasticsearch	1. Navigate to "Dev Tools" in Kibana
	cluster health	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.	Verify Prometheus Alert manager is accessible	1. Run below commands to get the load-balancer IP address and port number for Prometheus Alert Manager Web Interface. # LoadBalancer ip address of the alertmanager service is retrieved with below command \$ export ALERTMANAGER_LOADBALANCER_IP=\$(kubectl get services occne-prometheus-alertmanagernamespace 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-alertmanagernamespace occne-infra -o jsonpath="{.spec.ports[*].port}") # Complete url for accessing alertmanager in external browser \$ echo http://\$ALERTMANAGER_LOADBALANCER_IP: \$ALERTMANAGER_LOADBALANCER_IP: \$ALERTMANAGER_LOADBALANCER_PORT http://10.75.182.53:80
		2. Launch the Browser and navigate to http:// \$ALERTMANAGER_LOADBALANCER_IP: \$ALERTMANAGER_LOADBALANCER_PORT (e.g.: http:// 10.75.182.53:80 in the example above) received in the output of the above commands. Ensure the AlertManager GUI is accessible.



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

Step No.	Procedure	Description
5.	Verify metrics are scraped and stored in prometheus server	 Run below commands to get the load-balancer IP address and port number for Prometheus Server Web Interface. # LoadBalancer ip address of the prometheus service is retrieved with below command \$ export PROMETHEUS_LOADBALANCER_PORT=\$(kubectl get services occne-prometheus-servernamespace 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-servernamespace occne-infra -o jsonpath="{.status.loadBalancer.ingress[*].ip}") # Complete url for accessing prometheus in external browser \$ echo http://\$PROMETHEUS_LOADBALANCER_IP: \$PROMETHEUS_LOADBALANCER_PORT http://10.75.182.54:80 Launch the Browser and navigate to http:// \$PROMETHEUS_LOADBALANCER_IP: \$PROMETHEUS_LOADBALANCER_PORT (e.g.: http://10.75.182.54:80 in the example above) received in the output of the above commands. Ensure the Prometheus server GUI is accessible. Select "UP" option from "insert metric at cursor" drop down and click on "Execute" button. 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).
6.	Verify Alerts are configured	 Navigate to alerts tab of Prometheus server GUI or navigate using URL http://\$PROMETHEUS_LOADBALANCER_IP: \$PROMETHEUS_LOADBALANCER_PORT/ alertsFor<prometheus_loadbalancer_ip>and<prometh eus_loadbalancer_port=""></prometh></prometheus_loadbalancer_ip> 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	Desc	cription
7.	Verify grafana is accessible and change the default password for admin user		Run below commands to get the load-balancer IP address and port number for Grafana Web Interface. # LoadBalancer ip address of the grafana service is retrieved with below command \$ export GRAFANA_LOADBALANCER_IP=\$(kubectl get services occne-grafananamespace occne-infra -o jsonpath="{.status.loadBalancer.ingress[*].ip}")
			<pre># LoadBalancer port number of the grafana service is retrieved with below command \$ export GRAFANA_LOADBALANCER_PORT=\$(kubectl get services occne-grafananamespace occne-infra -o jsonpath="{.spec.ports[*].port}")</pre>
			# Complete url for accessing grafana in external browser \$ echo http://\$GRAFANA_LOADBALANCER_IP: \$GRAFANA_LOADBALANCER_PORT http://10.75.182.55:80
			Launch the Browser and navigate to http:// \$GRAFANA_LOADBALANCER_IP: \$GRAFANA_LOADBALANCER_PORT (e.g.: http:// 10.75.182.55:80 in the example above) received in the output of the above commands. Ensure the Prometheus server GUI is accessible. The default username and password is admin/admin for the 1st time access.
			At first connection to the Grafana dashboard, a 'Change Password' screen will appear. Change the password to the customer provided credentials.
			Note : Grafana data is not persisted, so if Grafana services restarted for some reason change password screen will appear again.
			Grafana dashboards are accessed after the changing the default password in the above step.
		5.	Click on "New dashboard" as marked red below.
		6.	Click on "Add Query"
			From " <i>Queries to</i> " 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.
			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="occne-infra"})For more details about promQl please follow the link.



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:

Table 4-2 Credentials

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

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

Prerequisites

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

Limitations and Expectations

The focus of this procedure is to secure the various credentials used or created during the install procedure. There are additional security audits that the CNE operator should perform such as



scanning repositories for vulnerabilities, monitoring the system for anomalies, regularly checking security logs. These are outside the scope of this post-installation procedure.

References

- 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
- See ToR switch procedure for initial username/password configuration: Configure Top of Rack 93180YC-EX Switches
- See procedure to configure initial iLO/OA username/password: Configure Addresses for RMS iLOs, OA, EBIPA
- See Enclosure switch procedure for initial username/password: Configure Enclosure Switches

Procedure



Table 4-3

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



Table 4-3 (Cont.)

Step No.	Procedure	Description
		 <switch name="">#</switch>
		5. Delete the previous old username if it is not needed:
		<pre># # configure Enter configuration commands, one per line. End with CNTL/Z. (config)# no username <username> (config)#exit #</username></pre>
		6. Change the enable secret when needed:
		<pre># (config)# enable secret <newenablepassword> (config)# exit #</newenablepassword></pre>
		7. Save the above configuration:
		<pre># copy running-config startup-config [####################################</pre>



Table 4-3 (Cont.)

G. N		_	
Step No.	Procedure		scription
2.	Reset Credentials on the Enclosure Switch	1.	From bastion host, login to the switch with username and password from the procedure:
			<pre>[bastion host]# ssh <username>@<switch address="" ip=""> <username>@<switch address="" ip="">'s</switch></username></switch></username></pre>
			password: <password></password>

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

			<switchname></switchname>
			<pre><switchname>sys</switchname></pre>
			System View: return to User View with
			Ctrl+Z. [switchname]
		2.	Change the password for current username:
			[switchname]local-user <username> class <current class=""></current></username>
			[switchname-luser-manage-
			<pre><username>]password simple <newpassword></newpassword></username></pre>
			<pre>[switchname-luser-manage-<username>]quit [switchname]</username></pre>
		3.	Create new username:
			[switchname]local-user <newusername> class [manage network]</newusername>
			New local user added.
			[switchname-luser-manage- <newusername>]password simple</newusername>
			<newpassword></newpassword>
			[switchname-luser-manage- <newusername>]quit</newusername>
			[switchname]
		4.	Exit from the switch and login with the new username and password to verify the new change works:
			<pre><switchname>quit Connection to <switch address="" ip=""> closed.</switch></switchname></pre>
			[bastion host]#



Table 4-3 (Cont.)

Step No.	Procedure	De	scription
			<pre>[bastion host]# ssh <newusername>@<switch address="" ip=""></switch></newusername></pre>
			<pre><newusername>@<switch address="" ip=""> <newusername>@<switch address="" ip="">'s</switch></newusername></switch></newusername></pre>
			password: <newpassword></newpassword>
			password: \newpassword>

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

			<switchname></switchname>
			<switchname>sys</switchname>
			System View: return to User View with
			Ctrl+Z.
			[switchname]
		5.	Delete the previous old username if it is not needed:
			[switchname]undo local-user <username></username>
			class <current class=""></current>
		6.	Save the above configuration:
			[switchname]save
			The current configuration will be
			written to the device. Are you sure? [Y/
			N]:y
			Please input the file name(*.cfg)
			[flash:/ <filename>]</filename>
			(To leave the existing filename
			unchanged, press the enter key):
			flash:/ <filename> exists, overwrite? [Y/</filename>
			N]:y
			Validating file. Please wait
			Saved the current configuration to
			mainboard device successfully.
			Slot 1:
			Save next configuration file
			successfully.
			[and balance]

[switchname]



Table 4-3 (Cont.)

Step No.	Procedure	Des	cription
3.	Reset Credentials for the OA Admin Console	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)
			[bastion host]# ssh <username>@<oa address=""></oa></username>
			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.
			Firmware Version: 4.85 Built: 04/06/2018 @ 06:14 OA Bay Number: 1 OA Role: Active <username>@<oa address="">'s password:<password></password></oa></username>
			HPE BladeSystem Onboard Administrator (C) Copyright 2006-2018 Hewlett Packard Enterprise Development LP
			Type 'HELP' to display a list of valid commands. Type 'HELP <command/> ' to display detailed information about a specific command. Type 'HELP HELP' to display more detailed information about the help system.
			OA-A45D36FD5FB1>
		2.	Change the password for current username:
			OA-A45D36FD5FB1> set password <newpassword></newpassword>



Changed password for the "<username>"

user account.

Table 4-3 (Cont.)

Step No. Procedure Description

OA-A45D36FD5FB1>

3. Add new user:

OA-A45D36FD5FB1> add user <newusername>

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

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

- "<newusername>" has been given [administrator|operator|user] level privileges.
- **4.** Assign full access to the enclosure for the user:

OA-A45D36FD5FB1> assign server all <newusername>

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

OA-A45D36FD5FB1> assign interconnect all <newusername>

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

OA-A45D36FD5FB1> assign oa <newusername>

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

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

OA-A45D36FD5FB1> exit

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

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



Table 4-3 (Cont.)

Step No.	Procedure	Description
		tored and can result in criminal or civil prosecution under applicable law
		Firmware Version: 4.85
		Built: 04/06/2018 @ 06:14
		OA Bay Number: 1
		OA Role: Active
		<newusername>@<oa address="">'s</oa></newusername>
		password: <newpassword></newpassword>

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

Type 'HELP' to display a list of valid commands.

Type 'HELP <command>' to display detailed information about a specific command

Type 'HELP' to display more detailed information about the help system.

OA-A45D36FD5FB1>

6. Delete previous user if not needed:

OA-A45D36FD5FB1> remove user <username>

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

Are you sure you want to remove testuser1? yes

User "<username>" removed.



Table 4-3 (Cont.)

Step No.	Procedure	Des	cription
4.	Reset Credentials for the ILO Admin Console	1.	From bastion host, login to the iLO with username and password from the procedure:
			<pre>[root@winterfell ~]# ssh <username>@<ilo address=""> <username>@<ilo address="">'s password: <password> User:<username> logged-in to(<ilo address=""> / <ipv6 address="">)</ipv6></ilo></username></password></ilo></username></ilo></username></pre>
			iLO Advanced 2.61 at Jul 27 2018 Server Name: <server name=""> Server Power: On</server>
			hpiLO->
		2.	Change current password:
			hpiLO-> set /map1/accounts1/ <username> password=<newpassword></newpassword></username>
			status=0 status_tag=COMMAND COMPLETED Tue Aug 20 13:27:08 2019
			hpiLO->
		3.	Create new user:
			<pre>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</newpassword></newusername></pre>
			User added successfully.
		4.	Exit from the iLO and login with the new username and password to verify the new change works:
			hpiLO-> exit
			status=0 status_tag=COMMAND COMPLETED Tue Aug 20 13:30:52 2019
			CLI session stopped Received disconnect from <ilo address=""> port 22:11: Client Disconnect Disconnected from <ilo address=""> port 22</ilo></ilo>



Table 4-3 (Cont.)

Step No.	Procedure	Description
		<pre>[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:</server></ipv6></ilo></newusername></newpassword></ilo></newusername></ilo></newusername></pre>
		<pre></pre> <pre></pre> <pre></pre> <pre></pre> <pre></pre> <pre></pre> <pre></pre> <pre>/*hpiLO-> delete /mapl/accountsl/ <username></username></pre>
		status=0 status_tag=COMMAND COMPLETED Tue Aug 20 13:59:04 2019
		User deleted successfully.
5.	Reset Credentials for the root account on Each and Every Server	Login to each and every server in the cluster (ssh admusr@cluster_host) and perform the following command:
		sudo passwd root
6.	Reset (or Delete) Credentials for the admusr account on Each and Every Server	Login to each and every server in the cluster (ssh admusr@cluster_host) and perform the following command: sudo passwd -l admusr
7.	Reset Credentials for the MySQL Accounts	See Database Tier Installer for details on how to reset the DB Account Passwords.



Table 4-3 (Cont.)

Step No.	Procedure	Description
8.	Regenerate / Redistribute SSH Keys	Log into the Bastion Host VM and generate a new cluster-wide keypair by perform the following:
Credentials for the admusr Account	Credentials for the admusr Account	ssh-keygen -b 4096 -t rsa -C "New SSH Key" -f .ssh/new_occne_id_rsa -q -N ""
		Now, for each and every server in the cluster, perform these actions:
		<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/</pre>
		<pre># install the key ssh admusr@cluster_host "cat .ssh/ new_occne_id_rsa.pub >> .ssh/ authorized_keys" # done</pre>
		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:
		<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</pre>
		<pre># 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>
	The new private key (new_occne_id_rsa) should also be copied to any secondary Bastion Host VM, and possibly copied off site and securely saved.	



A

Artifacts

The following appendices outline procedures referenced by one or more install procedures. These procedures may be conditionally executed based on customer requirements or to address certain deployment environments.

Repository Artifacts

OL YUM Repository Requirements

The following manifest includes the current list of RPMs that have been tested with a fully configured system.

Table A-1 OL YUM Repository Requirements

Num	RPM Name/Version	
1	nspr-4.19.0-1.el7_5.x86_64	
2	gpgme-1.3.2-5.el7.x86_64	
3	zlib-1.2.7-18.el7.x86_64	
4	elfutils-libelf-0.172-2.el7.x86_64	
5	numactl-libs-2.0.9-7.el7.x86_64	
6	perl-macros-5.16.3-294.el7_6.x86_64	
7	quota-nls-4.01-17.el7.noarch	
8	GeoIP-1.5.0-13.el7.x86_64	
9	lz4-1.7.5-2.0.1.el7.x86_64	
10	qrencode-libs-3.4.1-3.el7.x86_64	
11	ncurses-libs-5.9-14.20130511.el7_4.x86_64	
12	sg3_utils-libs-1.37-17.el7.x86_64	
13	libss-1.42.9-13.el7.x86_64	
14	info-5.1-5.el7.x86_64	
15	libselinux-utils-2.5-14.1.el7.x86_64	
16	openssl-libs-1.0.2k-16.0.1.el7.x86_64	
17	sed-4.2.2-5.el7.x86_64	
18	python-libs-2.7.5-76.0.1.el7.x86_64	
19	polkit-pkla-compat-0.1-4.el7.x86_64	
20	libdb-5.3.21-24.el7.x86_64	
21	glib2-2.56.1-2.el7.x86_64	
22	iputils-20160308-10.el7.x86_64	
23	libcap-ng-0.7.5-4.el7.x86_64	
24	rhnlib-2.5.65-8.0.1.el7.noarch	
25	libffi-3.0.13-18.el7.x86_64	
26	python-firewall-0.5.3-5.el7.noarch	



Table A-1 (Cont.) OL YUM Repository Requirements

Num	RPM Name/Version	
27	which-2.20-7.el7.x86 64	
28	python-perf-3.10.0-957.5.1.el7.x86 64	
29	expat-2.1.0-10.el7_3.x86_64	
30	bind-libs-9.9.4-73.el7 6.x86 64	
31		
32	libidn-1.28-4.el7.x86_64 nss-pem-1.0.3-5.el7.x86_64	
33	xmlrpc-c-1.32.5-1905.svn2451.el7.x86 64	
34	libssh2-1.4.3-12.el7.x86 64	
35	rpm-4.11.3-35.el7.x86_64	
36	dmraid-1.0.0.rc16-28.el7.x86_64	
37	_	
	gdbm-1.10-8.el7.x86_64	
38	rpm-python-4.11.3-35.el7.x86_64	
39	usb_modeswitch-data-20170806-1.el7.noarch	
40	perl-Pod-Perldoc-3.20-4.el7.noarch	
41	rhn-client-tools-2.0.2-24.0.5.el7.x86_64	
42	perl-Pod-Usage-1.63-3.el7.noarch	
43	dmidecode-3.1-2.el7.x86_64	
44	perl-Storable-2.45-3.el7.x86_64	
45	libteam-1.27-5.el7.x86_64	
46	perl-constant-1.27-2.el7.noarch	
47	libsmartcols-2.23.2-59.el7.x86_64	
48	perl-Socket-2.010-4.el7.x86_64	
49	util-linux-2.23.2-59.el7.x86_64	
50	perl-PathTools-3.40-5.el7.x86_64	
51	kmod-20-23.0.1.el7.x86_64	
52	systemd-219-62.0.4.el7_6.5.x86_64	
53	polkit-0.112-18.0.1.el7_6.1.x86_64	
54	plymouth-0.8.9-0.31.20140113.0.1.el7.x86_64	
55	libreport-python-2.1.11-42.0.1.el7.x86_64	
56	yum-plugin-ulninfo-0.2-13.el7.noarch	
57	boost-system-1.53.0-27.el7.x86_64	
58	grub2-tools-2.02-0.76.0.3.el7.1.x86_64	
59	cronie-anacron-1.4.11-20.el7_6.x86_64	
60	virt-what-1.18-4.el7.x86_64	
61	libnfnetlink-1.0.1-4.el7.x86_64	
62	policycoreutils-2.5-29.0.1.el7_6.1.x86_64	
63	keyutils-libs-1.5.8-3.el7.x86_64	
64	dracut-network-033-554.0.3.el7.x86_64	
65	desktop-file-utils-0.23-1.el7.x86_64	
66	libreport-cli-2.1.11-42.0.1.el7.x86 64	
67	abrt-dbus-2.1.11-52.0.1.el7.x86 64	
68	pinfo-0.6.10-9.el7.x86 64	
	r	



Table A-1 (Cont.) OL YUM Repository Requirements

Num	DDM Nama (Namian	
	RPM Name/Version	
69	hunspell-en-GB-0.20121024-6.el7.noarch	
70	abrt-addon-pstoreoops-2.1.11-52.0.1.el7.x86_64	
71	acl-2.2.51-14.el7.x86_64	
72	lvm2-libs-2.02.180-10.0.1.el7_6.3.x86_64	
73	quota-4.01-17.el7.x86_64	
74	libdwarf-20130207-4.el7.x86_64	
75	kernel-3.10.0-957.5.1.el7.x86_64	
76	setuptool-1.19.11-8.el7.x86_64	
77	make-3.82-23.el7.x86_64	
78	teamd-1.27-5.el7.x86_64	
79	usbutils-007-5.el7.x86_64	
80	libxcb-1.13-1.el7.x86_64	
81	lshw-B.02.18-12.el7.x86_64	
82	libmodman-2.0.1-8.el7.x86_64	
83	kmod-kvdo-6.1.1.125-5.0.1.el7.x86_64	
84	psacct-6.6.1-13.el7.x86_64	
85	boost-date-time-1.53.0-27.el7.x86_64	
86	vim-common-7.4.160-5.el7.x86_64	
87	biosdevname-0.7.3-1.el7.x86_64	
88	hardlink-1.0-19.el7.x86_64	
89	vim-enhanced-7.4.160-5.el7.x86_64	
90	smartmontools-6.5-1.el7.x86_64	
91	libXau-1.0.8-2.1.el7.x86_64	
92	sssd-client-1.16.2-13.el7_6.5.x86_64	
93	aic94xx-firmware-30-6.el7.noarch	
94	NetworkManager-tui-1.12.0-8.el7_6.x86_64	
95	p11-kit-trust-0.23.5-3.el7.x86_64	
96	libreport-plugin-mailx-2.1.11-42.0.1.el7.x86_64	
97	libdrm-2.4.91-3.el7.x86_64	
98	gzip-1.5-10.el7.x86_64	
99	microcode ctl-2.1-47.0.2.el7.x86 64	
100	bash-completion-2.1-6.el7.noarch	
101	shared-mime-info-1.8-4.el7.x86 64	
102	at-3.1.13-24.el7.x86_64	
103	blktrace-1.0.5-8.el7.x86_64	
104	libxml2-python-2.9.1-6.0.1.el7 2.3.x86 64	
105	dracut-config-rescue-033-554.0.3.el7.x86 64	
106	rhn-setup-2.0.2-24.0.5.el7.x86 64	
107	ntsysv-1.7.4-1.el7.x86_64	
108	bind-utils-9.9.4-73.el7 6.x86 64	
109	nano-2.3.1-10.el7.x86 64	
110	e2fsprogs-1.42.9-13.el7.x86 64	
110	0215p10g5-1.42.7-13.017.X00_04	



Table A-1 (Cont.) OL YUM Repository Requirements

RPM Name/Version
man-db-2.6.3-11.el7.x86_64
perl-Pod-Escapes-1.04-294.el7_6.noarch
setserial-2.17-33.el7.x86_64
python-slip-0.4.0-4.el7.noarch
strace-4.12-9.el7.x86_64
pyxattr-0.5.1-5.el7.x86_64
rfkill-0.4-10.el7.x86_64
iwl1000-firmware-39.31.5.1-999.1.el7.noarch
iwl135-firmware-18.168.6.1-999.1.el7.noarch
iwl6050-firmware-41.28.5.1-999.1.el7.noarch
iwl6000g2a-firmware-17.168.5.3-999.1.el7.noarch
iwl2030-firmware-18.168.6.1-999.1.el7.noarch
grub2-common-2.02-0.76.0.3.el7.1.noarch
grub2-pc-modules-2.02-0.76.0.3.el7.1.noarch
glibc-common-2.17-260.0.15.el7_6.3.x86_64
passwd-0.79-4.el7.x86_64
pygpgme-0.3-9.el7.x86_64
filesystem-3.2-25.el7.x86_64
basesystem-10.0-7.0.1.el7.noarch
libpipeline-1.2.3-3.el7.x86_64
kernel-uek-firmware-4.1.12-112.16.4.el7uek.noarch
gpm-libs-1.20.7-5.el7.x86_64
ncurses-base-5.9-14.20130511.el7_4.noarch
libutempter-1.1.6-4.el7.x86_64
libdaemon-0.14-7.el7.x86_64
pcre-8.32-17.el7.x86_64
xz-libs-5.2.2-1.el7.x86_64
libxml2-2.9.1-6.0.1.el7_2.3.x86_64
popt-1.13-16.el7.x86_64
bzip2-libs-1.0.6-13.el7.x86_64
readline-6.2-10.el7.x86_64
libattr-2.4.46-13.el7.x86_64
libacl-2.2.51-14.el7.x86_64
gawk-4.0.2-4.el7_3.1.x86_64
dbus-glib-0.100-7.el7.x86_64
libgpg-error-1.12-3.el7.x86_64
libusbx-1.0.21-1.el7.x86_64
cpio-2.11-27.el7.x86_64
libtar-1.2.11-29.el7.x86_64
json-c-0.11-4.el7_0.x86_64
sqlite-3.7.17-8.el7.x86_64
lua-5.1.4-15.el7.x86_64



Table A-1 (Cont.) OL YUM Repository Requirements

Num	RPM Name/Version
153	usermode-1.111-5.el7.x86_64
154	groff-base-1.22.2-8.el7.x86_64
155	hunspell-1.3.2-15.el7.x86_64
156	dmraid-events-1.0.0.rc16-28.el7.x86_64
157	perl-parent-0.225-244.el7.noarch
158	usb_modeswitch-2.5.1-1.el7.x86_64
159	perl-podlators-2.5.1-3.el7.noarch
160	python-slip-dbus-0.4.0-4.el7.noarch
161	kernel-3.10.0-862.el7.x86_64
162	perl-Encode-2.51-7.el7.x86_64
163	python-hwdata-1.7.3-4.el7.noarch
164	perl-threads-1.87-4.el7.x86_64
165	perl-Filter-1.49-3.el7.x86_64
166	perl-Time-HiRes-1.9725-3.el7.x86_64
167	perl-threads-shared-1.43-6.el7.x86_64
168	perl-Time-Local-1.2300-2.el7.noarch
169	perl-Carp-1.26-244.el7.noarch
170	perl-File-Path-2.09-2.el7.noarch
171	perl-Pod-Simple-3.28-4.el7.noarch
172	fxload-2002_04_11-16.el7.x86_64
173	pciutils-libs-3.5.1-3.el7.x86_64
174	alsa-tools-firmware-1.1.0-1.el7.x86_64
175	libmnl-1.0.3-7.el7.x86_64
176	python-pyudev-0.15-9.el7.noarch
177	libnl3-cli-3.2.28-4.el7.x86_64
178	plymouth-scripts-0.8.9-0.31.20140113.0.1.el7.x86_64
179	p11-kit-0.23.5-3.el7.x86_64
180	libedit-3.0-12.20121213cvs.el7.x86_64
181	rhnsd-5.0.13-10.0.1.el7.x86_64
182	libnl-1.1.4-3.el7.x86_64
183	yum-rhn-plugin-2.0.1-10.0.1.el7.noarch
184	newt-0.52.15-4.el7.x86_64
185	sysvinit-tools-2.88-14.dsf.el7.x86_64
186	libestr-0.1.9-2.el7.x86_64
187	yajl-2.0.4-4.el7.x86_64
188	libtiff-4.0.3-27.el7_3.x86_64
189	hostname-3.13-3.el7.x86_64
190	lzo-2.06-8.el7.x86_64
191	libnetfilter_conntrack-1.0.6-1.el7_3.x86_64
192	iproute-4.11.0-14.el7.x86_64
193	boost-thread-1.53.0-27.el7.x86_64
194	less-458-9.el7.x86_64



Table A-1 (Cont.) OL YUM Repository Requirements

Num	RPM Name/Version
195	libdb-utils-5.3.21-24.el7.x86_64
196	bzip2-1.0.6-13.el7.x86_64
197	libpng-1.5.13-7.el7_2.x86_64
198	mozjs17-17.0.0-20.el7.x86_64
199	libconfig-1.4.9-5.el7.x86_64
200	libproxy-0.4.11-11.el7.x86_64
201	gmp-6.0.0-15.el7.x86_64
202	libverto-0.2.5-4.el7.x86_64
203	pth-2.0.7-23.el7.x86_64
204	libyaml-0.1.4-11.el7_0.x86_64
205	lsscsi-0.27-6.el7.x86_64
206	libtasn1-4.10-1.el7.x86_64
207	cracklib-2.9.0-11.el7.x86_64
208	pkgconfig-0.27.1-4.el7.x86_64
209	newt-python-0.52.15-4.el7.x86_64
210	cyrus-sasl-lib-2.1.26-23.el7.x86_64
211	pygobject2-2.28.6-11.el7.x86_64
212	cracklib-dicts-2.9.0-11.el7.x86_64
213	pam-1.1.8-22.el7.x86_64
214	python-augeas-0.5.0-2.el7.noarch
215	python-iniparse-0.4-9.el7.noarch
216	gettext-0.19.8.1-2.el7.x86_64
217	python-gobject-base-3.22.0-1.el7_4.1.x86_64
218	python-chardet-2.2.1-1.el7_1.noarch
219	python-configobj-4.7.2-7.el7.noarch
220	PyYAML-3.10-11.el7.x86_64
221	m2crypto-0.21.1-17.el7.x86_64
222	fipscheck-lib-1.4.1-6.el7.x86_64
223	xmlrpc-c-client-1.32.5-1905.svn2451.el7.x86_64
224	bash-4.2.46-31.el7.x86_64
225	nss-util-3.36.0-1.1.el7_6.x86_64
226	libselinux-2.5-14.1.el7.x86_64
227	audit-libs-2.8.4-4.el7.x86_64
228	libcom_err-1.42.9-13.el7.x86_64
229	augeas-libs-1.4.0-6.el7_6.1.x86_64
230	libstdc++-4.8.5-36.0.1.el7.x86_64
231	perl-libs-5.16.3-294.el7_6.x86_64
232	libsemanage-2.5-14.el7.x86_64
233	file-libs-5.11-35.el7.x86_64
234	findutils-4.5.11-6.el7.x86_64
235	setup-2.8.71-10.el7.noarch
236	ethtool-4.8-9.el7.x86_64



Table A-1 (Cont.) OL YUM Repository Requirements

Num	RPM Name/Version
237	libjpeg-turbo-1.2.90-6.el7.x86_64
238	dyninst-9.3.1-2.el7.x86_64
239	e2fsprogs-libs-1.42.9-13.el7.x86_64
240	kmod-libs-20-23.0.1.el7.x86_64
241	vim-minimal-7.4.160-5.el7.x86_64
242	ca-certificates-2018.2.22-70.0.el7_5.noarch
243	coreutils-8.22-23.0.1.el7.x86_64
244	libblkid-2.23.2-59.el7.x86_64
245	python-2.7.5-76.0.1.el7.x86_64
246	libmount-2.23.2-59.el7.x86_64
247	grubby-8.28-25.0.1.el7.x86_64
248	pyOpenSSL-0.13.1-4.el7.x86_64
249	python-dmidecode-3.12.2-3.el7.x86_64
250	python-urlgrabber-3.10-9.el7.noarch
251	python-ethtool-0.8-7.el7.x86_64
252	sos-3.6-13.0.1.el7_6.noarch
253	gdb-7.6.1-114.el7.x86_64
254	openssl-1.0.2k-16.0.1.el7.x86_64
255	libtirpc-0.2.4-0.15.el7.x86_64
256	oracle-logos-70.0.3-4.0.9.el7.noarch
257	nss-3.36.0-7.1.el7_6.x86_64
258	nss-tools-3.36.0-7.1.el7_6.x86_64
259	libcurl-7.29.0-51.el7.x86_64
260	rpm-libs-4.11.3-35.el7.x86_64
261	openldap-2.4.44-21.el7_6.x86_64
262	rpm-build-libs-4.11.3-35.el7.x86_64
263	yum-3.4.3-161.0.1.el7.noarch
264	oraclelinux-release-el7-1.0-5.el7.x86_64
265	iptables-1.4.21-28.el7.x86_64
266	libfprint-0.8.2-1.el7.x86_64
267	kernel-tools-libs-3.10.0-957.5.1.el7.x86_64
268	iw-4.3-2.el7.x86_64
269	ipset-libs-6.38-3.el7_6.x86_64
270	lm_sensors-libs-3.4.0-6.20160601gitf9185e5.el7.x86_64
271	procps-ng-3.3.10-23.el7.x86_64
272	device-mapper-1.02.149-10.0.1.el7_6.3.x86_64
273	device-mapper-libs-1.02.149-10.0.1.el7_6.3.x86_64
274	dracut-033-554.0.3.el7.x86_64
275	elfutils-libs-0.172-2.el7.x86_64
276	dbus-libs-1.10.24-12.0.1.el7.x86_64
277	dbus-1.10.24-12.0.1.el7.x86_64
278	satyr-0.13-15.el7.x86_64



Table A-1 (Cont.) OL YUM Repository Requirements

N	DDM V. AV.
Num	RPM Name/Version
279	initscripts-9.49.46-1.0.1.el7.x86_64
280	device-mapper-event-libs-1.02.149-10.0.1.el7_6.3.x86_64
281	libreport-2.1.11-42.0.1.el7.x86_64
282	grub2-tools-minimal-2.02-0.76.0.3.el7.1.x86_64
283	libstoragemgmt-python-1.6.2-4.el7.noarch
284	libstoragemgmt-python-clibs-1.6.2-4.el7.x86_64
285	cronie-1.4.11-20.el7_6.x86_64
286	dhcp-libs-4.2.5-68.0.1.el7_5.1.x86_64
287	selinux-policy-3.13.1-229.0.3.el7_6.9.noarch
288	dhclient-4.2.5-68.0.1.el7_5.1.x86_64
289	kexec-tools-2.0.15-21.0.3.el7.x86_64
290	xdg-utils-1.1.0-0.17.20120809git.el7.noarch
291	grub2-pc-2.02-0.76.0.3.el7.1.x86_64
292	libreport-web-2.1.11-42.0.1.el7.x86_64
293	abrt-2.1.11-52.0.1.el7.x86_64
294	abrt-python-2.1.11-52.0.1.el7.x86_64
295	abrt-addon-vmcore-2.1.11-52.0.1.el7.x86_64
296	abrt-tui-2.1.11-52.0.1.el7.x86_64
297	device-mapper-event-1.02.149-10.0.1.el7_6.3.x86_64
298	lvm2-2.02.180-10.0.1.el7_6.3.x86_64
299	NetworkManager-1.12.0-8.el7_6.x86_64
300	fprintd-0.8.1-2.el7.x86_64
301	openssh-clients-7.4p1-16.el7.x86_64
302	abrt-addon-python-2.1.11-52.0.1.el7.x86_64
303	authconfig-6.2.8-30.el7.x86_64
304	elfutils-0.172-2.el7.x86_64
305	abrt-cli-2.1.11-52.0.1.el7.x86_64
306	kernel-uek-4.1.12-112.16.4.el7uek.x86_64
307	libsss_nss_idmap-1.16.2-13.el7_6.5.x86_64
308	pciutils-3.5.1-3.el7.x86_64
309	kernel-uek-4.1.12-124.26.1.el7uek.x86_64
310	kbd-legacy-1.15.5-15.el7.noarch
311	vim-filesystem-7.4.160-5.el7.x86 64
312	rsync-3.1.2-4.el7.x86_64
313	libX11-common-1.6.5-2.el7.noarch
314	gdk-pixbuf2-2.36.12-3.el7.x86_64
315	firewalld-0.5.3-5.el7.noarch
316	vdo-6.1.1.125-3.el7.x86 64
317	ntpdate-4.2.6p5-28.0.1.el7.x86_64
318	abrt-console-notification-2.1.11-52.0.1.el7.x86 64
319	fprintd-pam-0.8.1-2.el7.x86 64
320	grub2-2.02-0.76.0.3.el7.1.x86 64
L	8-3-2 -10 2 01/0101011111100_01



Table A-1 (Cont.) OL YUM Repository Requirements

Num	RPM Name/Version
321	sysstat-10.1.5-17.e17.x86 64
322	rpcbind-0.2.0-47.el7.x86_64
323	tuned-2.10.0-6.el7.noarch
324	yum-langpacks-0.4.2-7.el7.noarch
325	rng-tools-6.3.1-3.el7.x86_64
326	chrony-3.2-2.0.1.el7.x86_64
327	cyrus-sasl-plain-2.1.26-23.el7.x86_64
328	irqbalance-1.0.8-1.el7.x86_64
329	mtr-0.85-7.el7.x86_64
330	mdadm-4.1-rc1_2.el7.x86_64
331	btrfs-progs-4.9.1-1.0.2.el7.x86_64
332	hwdata-0.252-9.1.el7.x86_64
333	tcsh-6.18.01-15.el7.x86_64
334	ledmon-0.90-1.el7.x86_64
335	cryptsetup-2.0.3-3.el7.x86_64
336	hunspell-en-0.20121024-6.el7.noarch
337	kernel-tools-3.10.0-957.5.1.el7.x86_64
338	rhn-check-2.0.2-24.0.5.el7.x86_64
339	bc-1.06.95-13.el7.x86_64
340	systemtap-runtime-3.3-3.el7.x86_64
341	unzip-6.0-19.el7.x86_64
342	gobject-introspection-1.56.1-1.el7.x86_64
343	time-1.7-45.el7.x86_64
344	libselinux-python-2.5-14.1.el7.x86_64
345	xfsprogs-4.5.0-18.0.1.el7.x86_64
346	alsa-lib-1.1.6-2.el7.x86_64
347	mariadb-libs-5.5.60-1.el7_5.x86_64
348	rdate-1.4-25.el7.x86_64
349	sg3_utils-1.37-17.el7.x86_64
350	bridge-utils-1.5-9.el7.x86_64
351	iprutils-2.4.16.1-1.el7.x86_64
352	libgomp-4.8.5-36.0.1.el7.x86_64
353	scl-utils-20130529-19.el7.x86_64
354	dosfstools-3.0.20-10.el7.x86_64
355	rootfiles-8.1-11.el7.noarch
356	iwl3160-firmware-22.0.7.0-999.1.el7.noarch
357	ivtv-firmware-20080701-26.el7.noarch
358	iwl7265-firmware-22.0.7.0-999.1.el7.noarch
359	iwl105-firmware-18.168.6.1-999.1.el7.noarch
360	iwl7260-firmware-22.0.7.0-999.1.el7.noarch
361	iwl100-firmware-39.31.5.1-999.1.el7.noarch
362	man-pages-3.53-5.el7.noarch
L	F0 0 0 1



Table A-1 (Cont.) OL YUM Repository Requirements

Num	RPM Name/Version
363	iwl5000-firmware-8.83.5.1 1-999.1.el7.noarch
364	iwl6000-firmware-9.221.4.1-999.1.el7.noarch
365	iwl6000g2b-firmware-17.168.5.2-999.1.el7.noarch
	words-3.0-22.el7.noarch
366	
367	iwl4965-firmware-228.61.2.24-999.1.el7.noarch
368	gpg-pubkey-ec551f03-53619141
369	iwl2000-firmware-18.168.6.1-999.1.el7.noarch
370	libgcc-4.8.5-36.0.1.el7.x86_64
371	NetworkManager-config-server-1.12.0-8.el7_6.noarch
372	redhat-release-server-7.6-4.0.1.el7.x86_64
373	libreport-filesystem-2.1.11-42.0.1.el7.x86_64
374	kbd-misc-1.15.5-15.el7.noarch
375	nss-softokn-freebl-3.36.0-5.0.1.el7_5.x86_64
376	glibc-2.17-260.0.15.el7_6.3.x86_64
377	libsepol-2.5-10.el7.x86_64
378	python-pycurl-7.19.0-19.el7.x86_64
379	langtable-0.0.31-3.el7.noarch
380	libuuid-2.23.2-59.el7.x86_64
381	libndp-1.2-7.el7.x86_64
382	langtable-data-0.0.31-3.el7.noarch
383	oraclelinux-release-7.6-1.0.15.el7.x86_64
384	libpcap-1.5.3-11.el7.x86_64
385	perl-5.16.3-294.el7_6.x86_64
386	ustr-1.0.4-16.el7.x86_64
387	file-5.11-35.el7.x86_64
388	sgpio-1.2.0.10-13.el7.x86_64
389	nss-softokn-3.36.0-5.0.1.el7_5.x86_64
390	jasper-libs-1.900.1-33.el7.x86_64
391	freetype-2.8-12.el7_6.1.x86_64
392	tar-1.26-35.el7.x86_64
393	chkconfig-1.7.4-1.el7.x86_64
394	krb5-libs-1.15.1-37.el7_6.x86_64
395	grep-2.20-3.el7.x86 64
396	shadow-utils-4.1.5.1-25.el7.x86 64
397	libcap-2.22-9.el7.x86 64
398	linux-firmware-20181031-999.1.git1baa3486.el7.noarch
399	crontabs-1.11-6.20121102git.el7.noarch
400	python-linux-procfs-0.4.9-4.el7.noarch
401	dbus-python-1.1.1-9.el7.x86_64
402	libgcrypt-1.5.3-14.el7.x86 64
403	python2-futures-3.1.1-5.el7.noarch
404	libnl3-3.2.28-4.el7.x86 64
101	1101113 3.2.20 1.017.A00_0T



Table A-1 (Cont.) OL YUM Repository Requirements

Num	RPM Name/Version
405	bind-libs-lite-9.9.4-73.el7 6.x86 64
406	os-prober-1.58-9.0.1.el7.x86 64
407	tcp wrappers-libs-7.6-77.el7.x86 64
408	binutils-2.27-34.base.0.1.el7.x86_64
409	_
410	openssh-7.4p1-16.el7.x86_64
411	diffutils-3.3-4.el7.x86_64
412	nss-sysinit-3.36.0-7.1.el7_6.x86_64
	xz-5.2.2-1.el7.x86_64
413	curl-7.29.0-51.el7.x86_64
414	hunspell-en-US-0.20121024-6.el7.noarch
415	gnupg2-2.0.22-5.el7_5.x86_64
416	perl-HTTP-Tiny-0.033-3.el7.noarch
417	yum-utils-1.1.31-50.0.1.el7.noarch
418	perl-Text-ParseWords-3.29-4.el7.noarch
419	pixman-0.34.0-1.el7.x86_64
420	libpciaccess-0.14-1.el7.x86_64
421	libsss_idmap-1.16.2-13.el7_6.5.x86_64
422	perl-Exporter-5.68-3.el7.noarch
423	ipset-6.38-3.el7_6.x86_64
424	perl-Scalar-List-Utils-1.27-248.el7.x86_64
425	kpartx-0.4.9-123.el7.x86_64
426	perl-File-Temp-0.23.01-3.el7.noarch
427	cryptsetup-libs-2.0.3-3.el7.x86_64
428	ebtables-2.0.10-16.el7.x86_64
429	perl-Getopt-Long-2.40-3.el7.noarch
430	systemd-libs-219-62.0.4.el7_6.5.x86_64
431	alsa-firmware-1.0.28-2.el7.noarch
432	elfutils-default-yama-scope-0.172-2.el7.noarch
433	plymouth-core-libs-0.8.9-0.31.20140113.0.1.el7.x86_64
434	libassuan-2.1.0-3.el7.x86_64
435	systemd-sysv-219-62.0.4.el7_6.5.x86_64
436	python-gudev-147.2-7.el7.x86_64
437	libunistring-0.9.3-9.el7.x86_64
438	abrt-libs-2.1.11-52.0.1.el7.x86_64
439	slang-2.2.4-11.el7.x86_64
440	libstoragemgmt-1.6.2-4.el7.x86_64
441	jansson-2.10-1.el7.x86_64
442	NetworkManager-libnm-1.12.0-8.el7_6.x86_64
443	jbigkit-libs-2.0-11.el7.x86_64
444	libaio-0.3.109-13.el7.x86_64
445	dhcp-common-4.2.5-68.0.1.el7_5.1.x86_64
446	device-mapper-persistent-data-0.7.3-3.el7.x86 64
<u> </u>	** * —



Table A-1 (Cont.) OL YUM Repository Requirements

Num	RPM Name/Version
447	grub2-tools-extra-2.02-0.76.0.3.el7.1.x86_64
448	libreport-plugin-ureport-2.1.11-42.0.1.el7.x86_64
449	pm-utils-1.4.1-27.el7.x86_64
450	abrt-addon-kerneloops-2.1.11-52.0.1.el7.x86_64
451	tcp_wrappers-7.6-77.el7.x86_64
452	abrt-addon-xorg-2.1.11-52.0.1.el7.x86_64
453	attr-2.4.46-13.el7.x86_64
454	wpa_supplicant-2.6-12.el7.x86_64
455	pinentry-0.8.1-17.el7.x86_64
456	systemd-python-219-62.0.4.el7_6.5.x86_64
457	openssh-server-7.4p1-16.el7.x86_64
458	abrt-addon-ccpp-2.1.11-52.0.1.el7.x86_64
459	mlocate-0.26-8.el7.x86_64
460	ncurses-5.9-14.20130511.el7_4.x86_64
461	kernel-uek-firmware-4.1.12-124.26.1.el7uek.noarch
462	snappy-1.1.0-3.el7.x86_64
463	firewalld-filesystem-0.5.3-5.el7.noarch
464	libX11-1.6.5-2.el7.x86_64
465	libseccomp-2.3.1-3.el7.x86_64
466	kbd-1.15.5-15.el7.x86_64
467	NetworkManager-team-1.12.0-8.el7_6.x86_64
468	libsysfs-2.1.0-16.el7.x86_64
469	selinux-policy-targeted-3.13.1-229.0.3.el7_6.9.noarch
470	tcpdump-4.9.2-3.el7.x86_64
471	audit-2.8.4-4.el7.x86_64
472	crda-3.18_2018.05.31-4.el7.x86_64
473	xfsdump-3.1.7-1.el7.x86_64
474	net-tools-2.0-0.24.20131004git.el7.x86_64
475	python-decorator-3.4.0-3.el7.noarch
476	libgudev1-219-62.0.4.el7_6.5.x86_64
477	parted-3.1-29.0.1.el7.x86_64
478	libpwquality-1.2.3-5.el7.x86_64
479	sudo-1.8.23-3.el7.x86_64
480	zip-3.0-11.el7.x86_64
481	python-six-1.9.0-2.el7.noarch
482	libcroco-0.6.12-4.el7.x86_64
483	ed-1.9-4.el7.x86_64
484	gettext-libs-0.19.8.1-2.el7.x86_64
485	logrotate-3.8.6-17.el7.x86_64
486	traceroute-2.0.22-2.el7.x86_64
487	yum-metadata-parser-1.1.4-10.el7.x86_64
488	wget-1.14-18.el7.x86_64



Table A-1 (Cont.) OL YUM Repository Requirements

Num	RPM Name/Version
489	uname26-1.0-1.el7.x86_64
490	python-kitchen-1.1.1-5.el7.noarch
491	lsof-4.87-6.el7.x86_64
492	pyliblzma-0.5.3-11.el7.x86_64
493	libfastjson-0.99.4-3.el7.x86_64
494	langtable-python-0.0.31-3.el7.noarch
495	man-pages-overrides-7.6.2-1.el7.x86_64
496	python-schedutils-0.4-6.el7.x86_64
497	emacs-filesystem-24.3-22.el7.noarch
498	iwl3945-firmware-15.32.2.9-999.1.el7.noarch
499	redhat-indexhtml-7-13.0.1.el7.noarch
500	mailx-12.5-19.el7.x86_64
501	iwl5150-firmware-8.24.2.2-999.1.el7.noarch
502	rsyslog-8.24.0-34.el7.x86_64
503	fipscheck-1.4.1-6.el7.x86_64
504	bind-license-9.9.4-73.el7_6.noarch
505	tzdata-2018i-1.el7.noarch
506	libuser-0.60-9.el7.x86_64

Docker Repository Requirements

The following manifest includes the current list of docker containers used by OCCNE Common services that have been tested with a fully configured system.

Table A-2 Docker Repository Requirements

Num	Docker Name/Version
1	docker.elastic.co/elasticsearch/elasticsearch-oss:6.7.0
2	quay.io/pires/docker-elasticsearch-curator:5.5.4
3	justwatch/elasticsearch_exporter:1.0.2
4	gcr.io/google-containers/fluentd-elasticsearch:v2.3.2
5	grafana/grafana:6.1.6
6	appropriate/curl:latest
7	busybox:1.31.0
8	docker.elastic.co/kibana/kibana-oss:6.7.0
9	metallb/controller:v0.7.3
10	metallb/speaker:v0.7.3
11	prom/alertmanager:v0.18.0
12	jimmidyson/configmap-reload:v0.2.2
13	quay.io/coreos/kube-state-metrics:v1.6.0
14	prom/node-exporter:v0.17.0
15	prom/prometheus:v2.11.1



Num	Docker Name/Version
16	prom/pushgateway:v0.8.0
17	quay.io/prometheus/node-exporter:v0.17.0
18	jaegertracing/example-hotrod:latest
19	jaegertracing/jaeger-cassandra-schema:latest
20	jaegertracing/jaeger-agent:latest
21	jaegertracing/jaeger-collector:latest
22	jaegertracing/jaeger-query:latest
23	jaegertracing/spark-dependencies:latest

quay.io/external_storage/local-volume-provisioner:v2.2.0

Table A-2 (Cont.) Docker Repository Requirements

OCCNE YUM Repository Configuration

To perform an installation without internet access, create a local YUM mirror with the OL7 latest, epel, and addons repositories used by the "OS installation" process. Additionally a local repository is needed to hold the version of the docker-ce RPM used by the "Kubernetes installer" process. Repository files will need to be created to reference these local YUM repositories, and placed on the necessary machines (those which run the OCCNE installation Docker instances).

Pre-requisites

- Local YUM mirror repository for the OL7 'latest', 'epel', and 'addons' repositories.
 Directions here: https://www.oracle.com/technetwork/articles/servers-storage-admin/yum-repo-setup-1659167.html
- 2. Local YUM repository holding the required docker-ce RPM
- 3. Subscribe to following channels while creating the yum mirror from uln:

```
[017_x86_64_UEKR5]
[017_x86_64_ksplice]
[017_x86_64_latest]
[017_x86_64_addons]
[017_x86_64_developer]
```

For reference, view the yum mirroring instructions here: https://www.oracle.com/technetwork/articles/servers-storage-admin/yum-repo-setup-1659167.html

1. Create OL7 repository mirror repo.

Below is an example of a repository file providing the details on a mirror with the necessary repositories. This repository file would be placed on the machine that will run the OCCNE deployment containers.

```
/etc/yum.repos.d/ol7-mirror.repo

[local_ol7_x86_64_UEKR5]
name=Unbreakable Enterprise Kernel Release 5 for Oracle Linux 7 (x86_64)
baseurl=http://l0.75.155.195/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/$basearch/
qpqcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY
enabled=1
proxy=_none_
[local_ol7_x86_64_addons]
name=Oracle Linux 7 Addons (x86_64)
baseurl=http://10.75.155.195/yum/OracleLinux/OL7/addons/$basearch/
qpqcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY
enabled=1
proxy=_none_
[local_ol7_x86_64_ksplice]
name=Ksplice for Oracle Linux 7 (x86_64)
baseurl=http://10.75.155.195/yum/OracleLinux/OL7/ksplice/$basearch/
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY
enabled=1
proxy=_none_
[local_ol7_x86_64_developer]
name=Packages for creating test and development environments for Oracle
Linux 7 (x86_64)
baseurl=http://10.75.155.195/yum/OracleLinux/OL7/developer/$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_
```

2. Create Docker CE repository repo.

Below is an example of a repository file providing the details on a repository with the necessary docker-ce package.

```
/etc/yum.repos.d/docker-ce-stable.repo
[local_docker-ce-stable]
name=Docker CE Stable (x86_64)
baseurl=http://10.75.155.195/yum/centos/7/updates/$basearch/
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY
enabled=1
```

The docker RPM's placed into the repository directory should be docker-ce-18.09.5-3.el7.x86_64, docker-ce-cli-18.09.5-3.el7.x86_64.rpm and containerd.io-1.2.4-3.1.el7.x86_64.rpm and use create repo to generate repodata for new rpm's. The RPMs can be downloaded from: https://download.docker.com/linux/centos/7/

x86_64/stable and gpg-key for the rpm can be downloaded from: https://download.docker.com/linux/centos/gpg.

OCCNE 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 procedure should be completed before starting this procedure.
- 2. Docker is setup and docker commands can be run by the target system.
- 3. HTTP server that is reachable by the target system, Example- Running Nginx in docker container.

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

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

OR

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



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

Procedure Steps

Table A-3 Steps to configure OCCNE HTTP Repository

Steps	Procedure	Description
1.	Retrieve Kubernetes Binaries	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. 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 URL identified in the clusters hosts.ini inventory file (see below).
		<pre>\$ sh /var/occne/<cluster>/artifacts/ k8s_retrieve_bin.sh /var/www/html</cluster></pre>
		Example: \$ sh /var/occne/rainbow/artifacts/k8s_retrieve_bin.sh /var/www/html



Steps Procedure Description 2. Retrieve Helm The Configuration installer requires access to an HTTP server binaries and charts 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. \$ sh /var/occne/<cluster>/artifacts/ retrieve_helm.sh /var/www/html <helm path> < /var/occne/<cluster>/artifacts/ config_helm_charts.txt Example: \$ sh /var/occne/rainbow/artifacts/ retrieve_helm.sh /var/www/html ./ < /var/occne/</pre> rainbow/artifacts/config_helm_charts.txt Update inventory The hosts.ini inventory file for the cluster needs to have a few file with URLs variables set in the [occne:vars] section to direct the installation logic to the repository directories populated above. In this example the http server is winterfell on port 8082. **Note**: The helm repo has a trailing / the k8s repo does NOT. hosts.ini [occne:vars] occne_k8s_binary_repo='http://winterfell:8082/ binaries' occne_helm_stable_repo_url='http://winterfell: 8082/charts/'

Table A-3 (Cont.) Steps to configure OCCNE HTTP Repository

OCCNE 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

- 1. Docker images for OCCNE 1.2.0 release should be pulled to the executing system.
- 2. Docker is installed and docker commands can be run
- 3. Make sure docker registry is running by registry name provided
 - \$ docker ps
- 4. If not then creating a local docker registry accessible by the target of the installation
 - \$ docker run -d -p <port>:<port> --restart=always --name <registryname>
 registry:2



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

References

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

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

Procedure Steps

Table A-4 Steps to configure OCCNE Docker Image Registry

	ı	
Steps	Procedure	Description
1.	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 k8s install and configure dependencies
		Configure the registry:
		<pre>docker runrm -itnetwork hostcap- add=NET_ADMIN -v /var/occne/<cluster>/:/host occne/ <configure_install_image_name>:<1.2.0_tag> /getdeps/ getdeps</configure_install_image_name></cluster></pre>
		Generate k8s install:
		docker runrm -itnetwork hostcap- add=NET_ADMIN -v /var/occne/ <cluster>/:/host occne/ <k8s_install_image_name>:<1.2.0_tag> /getdeps/getdeps</k8s_install_image_name></cluster>
		Example
		docker runrm -itnetwork hostcap- add=NET_ADMIN -v /var/occne/rainbow/:/host occne/ configure:1.2.0 /getdeps/getdeps docker runrm -itnetwork hostcap- add=NET_ADMIN -v /var/occne/rainbow/:/host occne/ k8s_install:1.2.0 /getdeps/getdeps
		Once the above command is successfully executed, go to /var/occne/ <cluster and="" are="" artifacts="" directory="" execute:<="" file="" in="" k8s_docker_images.txt="" retrieve_docker.sh="" script="" td="" that="" the="" there="" verify=""></cluster>
		<pre>sh /var/occne/<cluster>/artifacts/retrieve_docker.sh docker.io <registryaddress:port> < /var/occne/ <cluster>/artifacts/k8s_docker_images.txt</cluster></registryaddress:port></cluster></pre>
		Once the above command is successfully executed, go to /var/occne/ <cluster and="" are="" artifacts="" directory="" retrieve_docker.sh<br="" that="" there="" verify="">script and config_docker_images.txt file in the directory and execute:</cluster>
		<pre>sh /var/occne/<cluster>/artifacts/retrieve_docker.sh docker.io <registryaddress:port> < /var/occne/ <cluster>/artifacts/config_docker_images.txt</cluster></registryaddress:port></cluster></pre>



Table A-4 (Cont.) Steps to configure OCCNE Docker Image Registry

Steps	Procedure	Description
Steps 2.	Procedure Verify the list of repositories in the docker registry	Description Access endpoint <dockerregistryhostip>:<dockerregistyport>/v2/ _catalog using a browser or from any linux server with curl command available and can ping the repo server address, using curl command \$ curl http://dockerregistryhostip:5000/v2/_catalog Sample Result: \$ {"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- 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-</dockerregistyport></dockerregistryhostip>
		<pre>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"]}</pre>



Table A-4 (Cont.) Steps to configure OCCNE Docker Image Registry

Steps	Procedure	Description
3.	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:
		hosts.ini
		[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

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

Variable	Definitions
occne_cluster_name	k8s cluster name
nfs_host	IP address OS install nfs host (host running the os_install container)
nfs_path	path to mounted OS install media on nfs host. This should always be set to /var/occne/
subnet_ipv4	subnet of IP addresses available for hosts in the OCCNE cluster
subnet_cidr	subnet_ipv4 in cidr notation format
netmask	subnet_ipv4 netmask
broadcast_address	broadcast address on the OCCNE cluster on which pxe server will listen
default_route	default router in the OCCNE cluster
next_server	IP address of TFTP server used for pxe boot (host running the os_install container)
name_server	DNS name server for the OCCNE cluster
ntp_server	NTP server for the OCCNE cluster



Variable	Definitions
http_proxy	HTTP Proxy server
https_proxy	HTTPS Proxy server
occne_private_registry	OCCNE private docker registry
occne_private_registry_add ress	OCCNE private docker registry address
occne_private_registry_port	OCCNE private docker registry port
metallb_peer_address	address of the BGP router peer that metalLB connects to
metallb_default_pool_proto col	protocol used to metalLB to announce allocated IP address
metallb_default_pool_addre sses	range of IP address to be allocated by metalLB from the default pool
pxe_install_lights_out_usr	ILO user
pxe_install_lights_out_pass wd	ILO user password
pxe_config_metrics_persist _size	(optional) Logical volume size for Metrics persistent storage, will override default of $500\mbox{G}$
pxe_config_es_data_persist _size	(optional) Logical volume size for ElasticSearch data persistent storage, will override default of 500G
pxe_config_es_master_pers ist_size	(optional) Logical volume size for ElasticSearch master persistent storage, will override default of 500G

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 section 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.

- A host is defined as a Fully Qualified Domain Name (FQDN). Properties are defined as key=value pairs.
- A property applies to a specific host when it appears on the same line as the host.
- 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
- A property applies to an entire group when it is defined under a group heading not on the same line as a host.



- Groups of groups are formed using the children keyword. For example, the [occne:children] creates an occne group comprised of several other groups.
- Inline comments are not allowed

The OCCNE Inventory file is composed of several groups:

Table B-1 Base Groups

Num	Group Name	Description/Comments
1.	host_hp_gen _10	 The list of all physical hosts in the OCCNE cluster. Each host in this group must also have several properties defined (outlined below) ansible_host: The IP address for the host's teamed primary interface. The occne/provision container uses this IP to configure a static IP for a pair of teamed interfaces when the hosts are provisioned. ilo: 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. mac: The MAC address of the host's network bootable interface. This is typically eno5 for Gen10 RMS hardware and eno1 for Gen10 bladed hardware. MAC addresses must use all lowercase alphanumeric values with a dash as the separator
		 ilo_ansible_host: The ILO IP address assigned to the Storage Hosts. oam_ansible_host: The OAM IP address assigned to the Storage Hosts.
2.	host_kernel_ virtual	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 property. • ilo_ansible_host: The ILO IP address assigned to the Bastion host virtual machines.
		 oam_ansible_host: The OAM IP address assigned to the Bastion host virtual machines.
		 signal_ansible_host: The Signalling IP address assigned to the MySQL NDB SQL Node virtual machines.
3.	occne:childr en	Do not modify the children of the occue group
4.	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.
5.	data_store	The list of Storage Hosts
6.	kube-master	The list of Master Node hosts where kubernetes master components run.
7.	etcd	The list of hosts that compose the etcd server. Should always be an odd number. This set to the same list of nodes as the kube-master group.
8.	kube-node	The list of Worker Nodes. Worker Nodes are where kubernetes pods run and should be comprised of the bladed hosts.
9.	k8s- cluster:childr en	Do not modify the children of k8s-cluster
10.	bastion_host s	The list of Bastion Hosts



Table B-1 (Cont.) Base Groups

Num	Group Name	Description/Comments
11.	skip_kernel_ virtual	The list of Virtual Machines to be skipped while creating the other Virtual machines specified in the "host_kernel_virtual" host group.
		For Ex: Here if the bastion-2.foo.lab.us.oracle.com is already created then we can skip this VM and create only bastion-1 VM.
		[host_kernel_virtual]
		bastion-1.foo.lab.us.oracle.com ansible_host=10.75.216.xx mac=xx-xx-xx-xx-xx host_hp_gen_blade=db-1.foo.lab.us.oracle.com ilo_ansible_host=10.75.xxx.xx oam_ansible_host=10.75.xxx.xx
		bastion-2.foo.lab.us.oracle.com ansible_host=10.75.216.xx mac=xx-xx-xx-xx-xx host_hp_gen_blade=db-2.foo.lab.us.oracle.com ilo_ansible_host=10.75.xxx.xx oam_ansible_host=10.75.xxx.xx
		[skip_kernel_virtual]
		bastion-2.foo.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 across the MySQL cluster. Additional parameter range limitations are outlined below.

Table B-2 Data Tier Groups

Num	Group Name	Description/Comments
1.	mysqlndb_mgm_node s	The list of MySQL Management nodes. In OCCNE 1.2 this group consists of three virtual machines distributed equally among the kubemaster nodes. These nodes must have a NodeId parameter defined: NodeId: Parameter must be unique across the MySQL Cluster and have a value between 49 and 255.
2.	mysqlndb_data_nodes _ng0	 The list of MySQL Data nodes, In OCCNE 1.2 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 Nodeld parameters. Nodeld: Parameter must be unique across the MySQL Cluster and have a value between 1 and 48. For Ex: Nodeld should be assigned with value 1 and 2
		-
		[mysqlndb_data_nodes_ng0]
		db-7.foo.lab.us.oracle.com NodeId=2
3.	mysqlndb_data_nodes _ng1	The list of MySQL Data nodes, In OCCNE 1.2 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. For Ex: NodeId should be assigned with value 3 and 4 [mysqlndb_data_nodes_ng0] db-8.foo.lab.us.oracle.com NodeId=3
		db-9.foo.lab.us.oracle.com NodeId=4
3.		The list of MySQL Data nodes, In OCCNE 1.2 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. For Ex: NodeId should be assigned with value 3 and 4 [mysqlndb_data_nodes_ng0] db-8.foo.lab.us.oracle.com NodeId=3



Table B-2 (Cont.) Data Tier Groups

Num	Group Name	Description/Comments
4.	mysqlndb_data_nodes	The list of MySQL Data node groups. In OCCNE 1.2 this group consists of 2 groups, each groups consists of two virtual machines distributed equally among the Storage Hosts.
5.	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.
		• NodeId: Parameter must be unique across the MySQL Cluster and have a value between 49 and 255.
6.	mysqlndb_all_nodes: children	Do not modify the children of the mysqlndb_all_nodes group.
7.	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

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 Steps

Table B-3 Procedure for OCCNE Inventory File Preparation

Step #	Procedure	Description
1.	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: 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



 Table B-3
 (Cont.) Procedure for OCCNE Inventory File Preparation

Step #	Procedure	Description
2.	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.
3.	Edit occne:vars	Edit the values in the occne:vars group to reflect site specific data. Values in the occne:vars group are defined in Table B-4.
4.	Edit mysqlndb_al l_nodes:vars	Edit the values: occne_mysqlndb_NoOfReplicas: Number of Replicas with in the MySQL NDB Cluster. For Ex: 2 occne_mysqlndb_DataMemory: Size of Data Memory(RAM) assigned to each MySQL Data Nodes. For Ex: 12G

Table B-4 occne:vars

Num	Var Name	Description
1	occne_cluster_name	Set to the OCCNE Cluster Name generated in step 2.1 above.
2	nfs_host	Set to the IP of the bastion host.
3	nfs_path	Set to the location of the nfs root created on the bastion host.
4	subnet_ipv4	Set to the subnet of the network used to assign IPs for OCCNE hosts
5	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
6	netmask	Set appropriately for the network used to assign IPs for OCCNE hosts.
7	broadcast_address	Set appropriately for the network used to assign IPs for OCCNE hosts.
8	default_route	Set to the IP of the TOR switch.
9	next_server	Set to the IP of the bastion host.
10	name_server	Set to the IP of the bastion host.
11	ntp_server	Set to the IP of the TOR switch.
12	http_proxy	Set the http proxy.
13	https_proxy	Set the https proxy.



Table B-4 (Cont.) occne:vars

Num	Var Name	Description
14	occne_private_registry	Set to the non-FQDN of the docker registry used by worker nodes to pull docker images from. NOTE: It's ok if this name is not in DNS, or if DNS is not available. The IP and Port settings are used to configure this registry on each host, placing the name and IP in each host's /etc/ hosts file, ensuring the name resolves to an IP.
15	occne_private_registry_address	Set to the IP of the docker registry above.
16	occne_private_registry_port	Set to the Port of the docker registry above
17	metallb_peer_address	Not used
18	metallb_default_pool_protocol	Not used
19	metallb_default_pool_addresses	Not used
20	pxe_install_lights_out_usr	Set to the user name configured for iLO admins on each host in the OCCNE Frame.
21	pxe_install_lights_out_passwd	Set to the password configured for iLO admins on each host in the OCCNE Frame.
22	occne_k8s_binary_repo	Set the internal IP of bastion-1 and the port configured.
23	helm_stable_repo_url	Set to the url of the local helm repo.
24	occne_helm_stable_repo_url	Set to the url of the local helm repo.
25	occne_helm_images_repo	Set to the url where images referenced in helm charts reside.
26	docker_rh_repo_base_url	Set to the URL of the repo containing the docker RPMs.
27	docker_rh_repo_gpgkey	Set to the URL of the gpgkey in the docker yum repo.
28	ilo_vlan_id	Set to the VLAN ID of the ILO network For Ex: 2
29	ilo_subnet_ipv4	Set to the subnet of the ILO network used to assign IPs for Storage hosts
30	ilo_subnet_cidr	Set to the cidr notation for the subnet. For example 24
31	ilo_netmask	Set appropriately for the network used to assign ILO IPs for Storage hosts.
32	ilo_broadcast_address	Set appropriately for the network used to assign ILO IPs for OCCNE hosts.



Table B-4 (Cont.) occne:vars

Num	Var Name	Description
33	ilo_default_route	Set to the ILO VIP of the TOR switch.
34	mgmt_vlan_id	Set to the VLAN ID of the Management network For Ex: 4
35	mgmt_subnet_ipv4	Set to the subnet of the Management network used to assign IPs for Storage hosts
36	mgmt_subnet_cidr	Set to the cidr notation for the Management subnet. For example 29
37	mgmt_netmask	Set appropriately for the network used to assign Management IPs for Storage hosts.
38	mgmt_broadcast_address	Set appropriately for the network used to assign Management IPs for Storage hosts.
39	mgmt_default_route	Set to the Management VIP of the TOR switch.
40	signal_vlan_id	Set to the VLAN ID of the Signaling network For Ex: 5
41	signal_subnet_ipv4	Set to the subnet of the Signaling network used to assign IPs for Storage hosts
42	signal_subnet_cidr	Set to the cidr notation for the Signaling subnet. For example 29
43	signal_netmask	Set appropriately for the network used to assign Signaling IPs for Storage hosts and MySQL SQL Node VM's.
44	signal_broadcast_address	Set appropriately for the network used to assign Signaling IPs for Storage hosts and MySQL SQL Node VM's.
45	signal_default_route	Set to the Signaling VIP of the TOR switch.
46	mysql_bastion_node_ram	Size of the RAM assigned to the bastion hosts. For EX: 8192 (in MB)
47	mysql_bastion_node_vcpus	Number of cpus assigned to the bastion hosts. For Ex: 4
48	mysql_bastion_node_disk_size	Size of the Disk assigned to the bastion host. For Ex: 300 (means 300 GB)
49	mysql_mgm_node_ram	Size of the RAM assigned to the MySQL MGM Nodes. For EX: 8192 (in MB)
50	mysql_mgm_node_vcpus	Number of cpus assigned to the MySQL MGM Nodes. For Ex: 8



Table B-4 (Cont.) occne:vars

Num	Var Name	Description
51	mysql_mgm_node_disk_size	Size of the Disk assigned to the MySQL MGM Nodes. For Ex: 200 (means 200 GB)
52	mysql_data_node_ram	Size of the RAM assigned to the MySQL DATA Nodes. For EX: 32768 (in MB)
53	mysql_data_node_vcpus	Number of cpus assigned to the MySQL DATA Nodes. For Ex: 12
54	mysql_data_node_disk_size	Size of the Disk assigned to the MySQL DATA Nodes. For Ex: 600 (means 600 GB)
55	mysql_sql_node_ram	Size of the RAM assigned to the MySQL SQL Nodes. For EX: 16384 (in MB)
56	mysql_sql_node_vcpus	Number of cpus assigned to the MySQL SQL Nodes. For Ex: 8
57	mysql_sql_node_disk_size	Size of the Disk assigned to the MySQL SQL Nodes. For Ex: 600 (means 600 GB)
58	occne_snmp_notifier_destination	Set to the address of SNMP trap receiver. For Ex: "127.0.0.1:162"

OCCNE Inventory Boilerplate

- # This is a list of all of the nodes in the targeted deployment system with the
- # IP address to use for Ansible control during deployment.
- # For bare metal hosts, the IP of the ILO is used for driving reboots.
- $\mbox{\tt\#}$ Host MAC addresses is used to identify nodes during PXE-boot phase of the
- # os_install process.
- # MAC addresses must be lowercase and delimited with a dash "-"

[host_hp_gen_10]

k8s-2.foo.lab.us.oracle.com ansible_host=10.75.216.xx ilo=10.75.216.xx mac=xx-xx-



```
xx-xx-xx
k8s-3.foo.lab.us.oracle.com ansible host=10.75.216.xx ilo=10.75.216.xx mac=xx-xx-
k8s-4.foo.lab.us.oracle.com ansible_host=10.75.216.xx ilo=10.75.216.xx mac=xx-xx-
xx-xx-xx-xx
k8s-5.foo.lab.us.oracle.com ansible_host=10.75.216.xx ilo=10.75.216.xx mac=xx-xx-
xx-xx-xx-xx
k8s-6.foo.lab.us.oracle.com ansible_host=10.75.216.xx ilo=10.75.216.xx mac=xx-xx-
xx-xx-xx
k8s-7.foo.lab.us.oracle.com ansible_host=10.75.216.xx ilo=10.75.216.xx mac=xx-xx-
db-1.foo.lab.us.oracle.com ansible_host=10.75.216.xx ilo=10.75.216.xx mac=xx-xx-
db-2.foo.lab.us.oracle.com ansible host=10.75.216.xx ilo=10.75.216.xx mac=xx-xx-
xx-xx-xx
[host_kernel_virtual]
db-3.foo.lab.us.oracle.com ansible host=10.75.216.xx mac=xx-xx-xx-xx-xx
db-4.foo.lab.us.oracle.com ansible_host=10.75.216.xx mac=xx-xx-xx-xx-xx
db-5.foo.lab.us.oracle.com ansible_host=10.75.216.xx mac=xx-xx-xx-xx-xx
db-6.foo.lab.us.oracle.com ansible_host=10.75.216.xx mac=xx-xx-xx-xx-xx
db-7.foo.lab.us.oracle.com ansible_host=10.75.216.xx mac=xx-xx-xx-xx-xx
db-8.foo.lab.us.oracle.com ansible_host=10.75.216.xx mac=xx-xx-xx-xx-xx
db-9.foo.lab.us.oracle.com ansible_host=10.75.216.xx mac=xx-xx-xx-xx-xx
db-10.foo.lab.us.oracle.com ansible_host=10.75.216.xx mac=xx-xx-xx-xx-xx
# Node grouping of which nodes are in the occne system
[occne:children]
host_hp_gen_10
host_kernel_virtual
k8s-cluster
data_store
# Variables that define the OCCNE environment and specify target configuration.
[occne:vars]
occne_cluster_name=foo.lab.us.oracle.com
nfs_host=10.75.216.xx
nfs_path=/var/occne
subnet_ipv4=10.75.216.0
subnet_cidr=/25
netmask=255.255.255.128
broadcast_address=10.75.216.127
default_route=10.75.216.1
next_server=10.75.216.114
name_server='10.75.124.245,10.75.124.246'
ntp_server='10.75.124.245,10.75.124.246'
http_proxy=http://www-proxy.us.oracle.com:80
https_proxy=http://www-proxy.us.oracle.com:80
occne_private_registry=bastion-1
occne_private_registry_address='10.75.216.xx'
occne_private_registry_port=5000
metallb_peer_address=10.75.216.xx
metallb_default_pool_protocol=bgp
metallb_default_pool_addresses='10.75.xxx.xx/xx'
pxe_install_lights_out_usr=root
pxe_install_lights_out_passwd=TklcRoot
occne_k8s_binary_repo='http://bastion-1:8082/binaries'
helm_stable_repo_url='http://<bastion-1 IP addr>:<port>/charts/stable/'
occne_helm_stable_repo_url='http://<bastion-1 IP addr>:<port>/charts/stable/'
```



```
occne_helm_images_repo='bastion-1:5000'/
docker_rh_repo_base_url=http://<bastion-1 IP addr>/yum/centos/7/updates/x86_64/
docker_rh_repo_gpgkey=http://<bastion-1 IP addr>/yum/centos/RPM-GPG-CENTOS
# Node grouping of which nodes are in the occne data_store
[data_store]
db-1.foo.lab.us.oracle.com
db-2.foo.lab.us.oracle.com
# Node grouping of which nodes are to be Kubernetes master nodes (must be at
least 2)
[kube-master]
k8s-1.foo.lab.us.oracle.com
k8s-2.foo.lab.us.oracle.com
k8s-3.foo.lab.us.oracle.com
# Node grouping specifying which nodes are Kubernetes etcd data.
# An odd number of etcd nodes is required.
[etcd]
k8s-1.foo.lab.us.oracle.com
k8s-2.foo.lab.us.oracle.com
k8s-3.foo.lab.us.oracle.com
# Node grouping specifying which nodes are Kubernetes worker nodes.
# A minimum of two worker nodes is required.
[kuhe-node]
k8s-4.foo.lab.us.oracle.com
k8s-5.foo.lab.us.oracle.com
k8s-6.foo.lab.us.oracle.com
k8s-7.foo.lab.us.oracle.com
# Node grouping of which nodes are to be in the OC-CNE Kubernetes cluster
[k8s-cluster:children]
kube-node
kube-master
# The following node groupings are for MySQL NDB cluster
# installation under control of MySQL Cluster Manager
# NodeId should be unique across the cluster, each node should be assigned with
# the unique NodeId, this id will control which data nodes should be part of
# different node groups. For Management nodes, NodeId should be between 49 to
# 255 and should be assigned with unique NodeId with in MySQL cluster.
[mysqlndb_mqm_nodes]
db-3.foo.lab.us.oracle.com NodeId=49
db-4.foo.lab.us.oracle.com NodeId=50
# For data nodes, NodeId should be between 1 to 48, NodeId will be used to
# group the data nodes among different Node Groups.
[mysqlndb_data_nodes]
db-5.foo.lab.us.oracle.com NodeId=1
db-6.foo.lab.us.oracle.com NodeId=2
db-7.foo.lab.us.oracle.com NodeId=3
db-8.foo.lab.us.oracle.com NodeId=4
```



OCCNE 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.

- A local YUM repository 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
- A local HTTP repository is needed to hold Kubernetes binaries and Helm charts
- A local Docker registry 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 for OS installation
- A copy of the for database nodes

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



Enclosure1: K8s Hosts

Rackmount ordering (frame not to scale)

SwitchB
SwitchA

RMS5: K8s Host
RMS4: K8s Host
RMS3: K8s Host
RMS2: Storage Host
RMS1: Storage Host

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

18 8

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



Table B-5 (Cont.) Enclosure Switch Connections

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



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
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
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



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
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.hpe.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 Required Artifacts Are Accessible

Ensure artifacts listed in the Artifacts 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 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		29	4	
5	Signaling Network - MySQL Replication		29	5	



Table B-8 (Cont.) Complete Site Survey Subnet Table

SI No.	Network Description	Subnet Allocation	Bitmask	VLAN ID	Gateway Address
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 B-9 Complete Site Survey Host IP Table

Sl No.	Component/ Resource	Platform VLAN IP Address (VLAN 3)	iLO VLAN IP Address (VLAN 2)	CNE Management IP Address (VLAN 4)	Device iLO IP Address	MAC of Primar y NIC	Notes
1	RMS 1 Host IP	172.16.3.4	192.168.20.1 1		192.168.20. 121	Eno5:	
2	RMS 2 Host IP	172.16.3.5	192.168.20.1 2		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:	



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

Sl No.	Component/ Resource	Platform VLAN IP Address (VLAN 3)	iLO VLAN IP Address (VLAN 2)	CNE Management IP Address (VLAN 4)	Device iLO IP Address	MAC of Primar y NIC	Notes
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:	
15	Enclosure 1 Bay 10 Host IP	172.16.3.20	N/A	N/A	192.168.20. 150	Eno1:	
16	Enclosure 1 Bay 11 Host IP	172.16.3.21	N/A	N/A	192.168.20. 151	Eno1:	
17	Enclosure 1 Bay 12 Host IP	172.16.3.22	N/A	N/A	192.168.20. 152	Eno1:	
18	Enclosure 1 Bay 13 Host IP	172.16.3.23	N/A	N/A	192.168.20. 153	Eno1:	
19	Enclosure 1 Bay 14 Host IP	172.16.3.24	N/A	N/A	192.168.20. 154	Eno1:	
20	Enclosure 1 Bay 15 Host IP	172.16.3.25	N/A	N/A	192.168.20. 155	Eno1:	
21	Enclosure 1 Bay 16 Host IP	172.16.3.26	N/A	N/A	192.168.20. 156	Eno1:	

Complete VM IP Table

Table B-10 Complete VM IP Table

SI No.	Component/ Resource	Platform VLAN IP Address (VLAN 3)	iLO VLAN IP Address (VLAN 2)	CNE Management IP Address (VLAN 4)	SQL Replication IP Address(VLA N 5)	Notes
1	Bastion Host 1	172.16.3.100	192.168.20.10 0		N/A	
2	Bastion Host 2	172.16.3.101	192.168.20.10 1		N/A	
3	MySQL SQL Node 1	172.16.3.102	N/A	N/A		



Table B-10 (Cont.) Complete VM IP Table

SI No.	Component/ Resource	Platform VLAN IP Address (VLAN 3)	iLO VLAN IP Address (VLAN 2)	CNE Management IP Address (VLAN 4)	SQL Replication IP Address(VLA N 5)	Notes
4	MySQL SQL Node 2	172.16.3.103	N/A	N/A		

Complete OA and Switch IP Table

Table B-11 Complete OA and Switch IP Table

S 1 N 0	Procedure Reference Variable Name	Description	IP Address	VLA N ID	Notes
1	N/A	Enclosure 1 IObay1	192.168.2 0.133	N/A	
2	N/A	Enclosure 1 IObay2	192.168.2 0.134	N/A	
3	N/A	Enclosure 1 OA1	192.168.2 0.131	N/A	
4	N/A	Enclosure 1 OA2	192.168.2 0.132	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_CNEManagement Net_IP	Bastion Host Network		4	Address needs to be without prefix length, such as 10.25.100.2
9	ToRswitchB_CNEManagement Net_IP	Bastion Host Network		4	Address needs to be without prefix length, such as 10.25.100.3
1 0	ToRswitch_CNEManagementN et_VIP	Bastion Host Network Default Gateway		4	No prefix length, address only for VIP
1	CNEManagementNet_Prefix	Bastion Host Network Prefix Length		4	number only such as 29
1 2	ToRswitchA_SQLreplicationNe t_IP	SQL Replication Network		5	Address needs to be with prefix length, such as 10.25.200.2



Table B-11 (Cont.) Complete OA and Switch IP Table

S 1 N 0	Procedure Reference Variable Name	Description	IP Address	VLA N ID	Notes
1 3	ToRswitchB_SQLreplicationNe t_IP	SQL Replication Network		5	Address needs to be with prefix length, such as 10.25.200.3
1 4	ToRswitch_SQLreplicationNet_ VIP	SQL Replication Network Default Gateway		5	No prefix length, address only for VIP
1 5	SQLreplicationNet_Prefix	SQL Replication Network Prefix Length		5	number only such as 28
1 6	ToRswitchA_oam_uplink_custo mer_IP	ToR Switch A OAM uplink route path to customer network		N/A	No prefix length in address, static to be /30
1 7	ToRswitchA_oam_uplink_IP	ToR Switch A OAM uplink IP		N/A	No prefix length in address, static to be /30
1 8	ToRswitchB_oam_uplink_custo mer_IP	ToR Switch B OAM uplink route path to customer network		N/A	No prefix length in address, static to be /30
1 9	ToRswitchB_oam_uplink_IP	ToR Switch B OAM uplink IP		N/A	No prefix length in address, static to be /30
2 0	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
2	ToRswitchA_signaling_uplink_ IP	ToR Switch A Signaling uplink IP		N/A	No prefix length in address, static to be /30
2 2	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
2 3	ToRswitchB_signaling_uplink_ IP	ToR Switch B Signaling uplink IP		N/A	No prefix length in address, static to be /30
2 4	ToRswitchA_mngt_IP	ToR Switch A Out of Band Management IP	192.168.2. 1	N/A	
2 5	ToRswitchB_mngt_IP	ToR Switch A Out of Band Management IP	192.168.2. 2	N/A	
6	MetalLB_Signal_Subnet_With_ Prefix	ToR Switch route provisioning for metalLB		N/A	From Section 2.1
7	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



Table B-11 (Cont.) Complete OA and Switch IP Table

S I N o	Procedure Reference Variable Name	Description	IP Address	VLA N ID	Notes
2 8	MetalLB_OAM_Subnet_With_ Prefix	ToR Switch route provisioning for metalLB		N/A	From Section 2.1
9	MetalLB_OAM_Subnet_IP_Ra nge	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
3 0	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 trouble shooting 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.
3	SNMP_Trap_Receiver_Address	IP address of the SNMP trap receiver			
3 2	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 B-12 ToR and Enclosure Switches Variables Table (Switch Specific)

	Key/ Vairable Name	ToR_S witch A Value	ToR _Sw itch B Val ue	Enclosu re_Swit ch1 Value	Enclosure_Switch2 Value	Notes
1	switch_nam e				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.
2	admin_pass word					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_passwo rd					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.</user_name>
5	ospf_md5_k ey			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_versio n			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).



- 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.
- 6. Copy the occne-ks.cfg.j2.new file from OHC into 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.



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.



This file includes some variables that must be updated when used in procedure: Installation of the Bastion Host.



The steps to update those variables are contained in that procedure.

#version=DEVEL

```
# System authorization information
auth --enableshadow --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 $6$etgyspJhPUG440VO
$0FqnB.agxmnDqb.Bh0sSLhq7..t37RwUZr7SlVmIBvMmWVOUjb2DJJ2f4VlrW9RdfVi.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..t37RwUZr7SlVmIBvMmWVoUjb2DJJ2f4VlrW9RdfVi.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
@{\tt 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
\label{local-control} sed --in-place 's/^{k*}(%wheel)s\\+ALL=(ALL))s\\+NOPASSWD:\s\\+ALL\)/\l' / etc/
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 && \backslash
sed -i 's/X11Forwarding yes/X11Forwarding no/' /etc/ssh/sshd_config && \
sed -i 's/#MaxAuthTries.*/MaxAuthTries 4/' /etc/ssh/sshd_config && \backslash
\verb|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*\$/\lLoginGraceTime 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

Installation Use Cases and Repository Requirements

Goals

- Identify the parameters and use case for initial setup and sustained support of an On-prem CNE
- Identify what components need access to software repositories, and how they will be accessed

Background and strategic fit

The installation process will assume a software delivery model of "Indirect Internet Connection". This model allows for a more rapid time to market for initial deployment and security update of the CNE. However, this model creates situations during the install process that require careful explanation and walk-through. Thus, the need for this page.

Requirements

- Installer notebooks may be used to access resources; however, the following limitations will need to be considered:
 - The installer notebook may not arrive on site with Oracle IP, such as source code or install tools
 - The installer notebook may not have customer sensitive material stored on it, such as access credentials



- Initial install may require trained personnel to be on site; however, DR of any individual component should not require trained software personnel to be local to the installing device
 - Physical rackmounting and cabling of replacement equipment should be performed by customer or contractor personnel; but software configuration and restoration of services should not require personnel to be sent to site.
- Oracle Linux Yum repository, Docker registry, and Helm repository is configured and available to the CNE frame for installation activities. Oracle will define what artifacts need to be in these repositories. It will be the customer responsibility to pull the artifacts into repositories reachable by the OCCNE frame.

User Interaction and Design

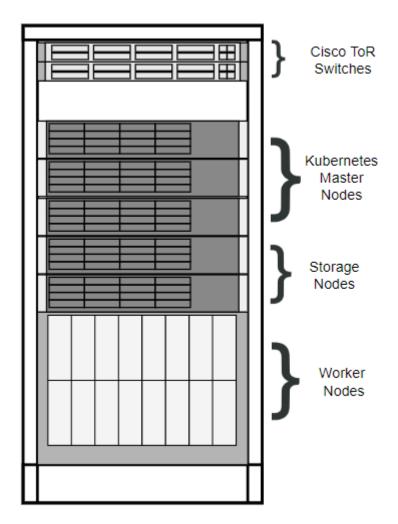
This section walks through expected installation steps for the CNE given the selected software delivery model.

CNE Overview

CNE Frame Overview

For reference regarding installation practices, it is useful to understand the hardware layout involved with the CNE deployment.

Figure B-2 Frame reference





Problem Statement

A solution is needed to initialize the frame with an OS, a Kubernetes cluster, and a set of common services for 5G NFs to be deployed into. How the frame is brought from manufacturing default state to configured and operational state is the topic of this page.

Manufacturing Default State characteristics/assumptions:

- Frame components are "racked and stacked", with power and network connections in place
- Frame ToR switches are not connected to the customer network until they are configured (alternatively, the links can be disabled from the customer side)
- An installer is on-site
- An installer has a notebook and a USB flash drive with which to configure at the first server in the frame
- An installer's notebook has access to the repositories setup by the customer

CNE Installation Preparation

Setting up the Notebook

The installer notebook is considered to be an Oracle asset. As such, it will have limitations applied as mentioned above. The notebook will be used to access the customer instantiated repositories to pull down the OL iso and apply it to a USB flash drive. Steps involved in creating the bootable USB drive will be dependent upon the OS on the notebook (for example, Rufus can be used for a Windows PC, or "dd" command can be used for a Linux PC).



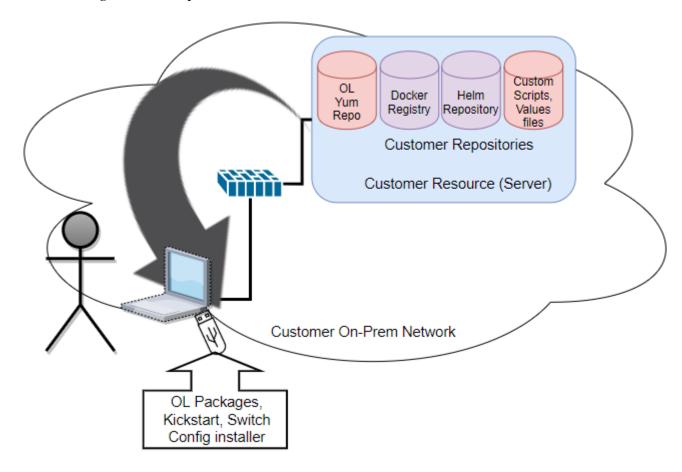


Figure B-3 Setup the Notebook and USB Flash Drive

CNE Installation - Setup the Management Server and Switches

Install OS on a "Bootstrap" Server

The 1st RMS in the frame will be temporarily used as a bootstrap server, whereby a manual method of initial OS install will be applied to start a "standard" process for installing the frame. The activity performed by this "bootstrap" server should be minimized to get to a standard "inframe configuration platform" as soon as possible. The bootstrap server should be re-paved to an "official" configuration as soon as possible. This means the "bootstrap" server will facilitate the configuration of the ToR switches, and the configuration of a Management VM. Once these two items have been completed, and the management VM is accessible from outside the frame, the "bootstrap" server will have fulfilled its purpose and can then be re-paved.

The figure below is very busy with information. Here are the key takeaways:

- The ToR switch uplinks are disabled or disconnected, as the ToR is not yet configured.
 This prevents nefarious network behavior due to redundant connections to unconfigured switches.
 - Until the ToR switches are configured, there is no connection to the customer repositories.
- The red server is special in that it has connections to ToR out of band interfaces (not shown).
- The red server is installed via USB flash drive and local KVM (Keyboard, Video, Mouse).

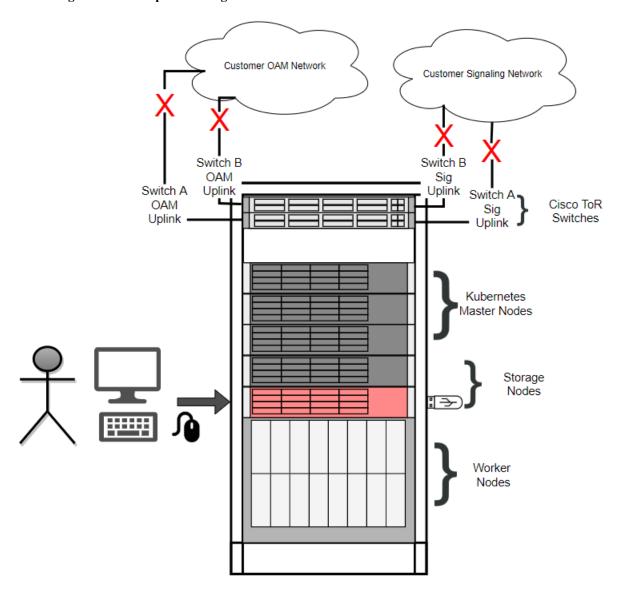


Figure B-4 Setup the Management Server

Setup the Switches

Setup Switch Configuration Services

Configure DHCP, tftp, and network interfaces to support ToR switch configuration activities. For the initial effort of CNE 1.0, this process is expected to be manual, without the need for files to be delivered to the field. Reference configuration files will be made available through documentation. If any files are needed from internet sources, they will be claimed as a dependency in the customer repositories and will be delivered by USB to the bootstrap server, similar to the OL iso.

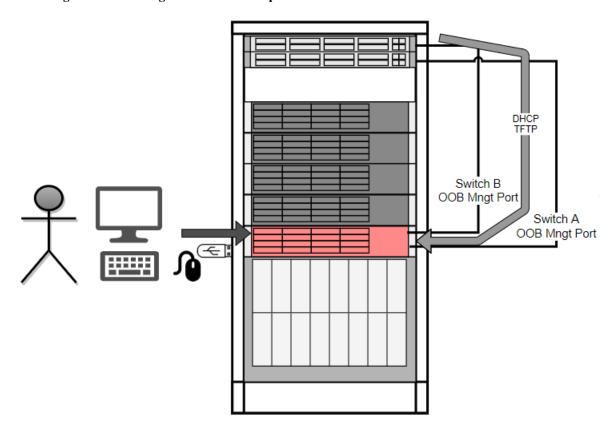


Figure B-5 Management Server Unique Connections

Configure the Enclosure Access

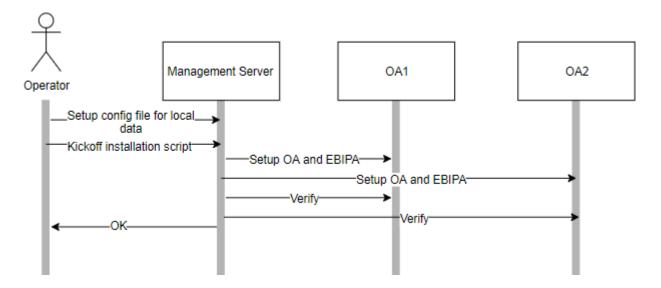
Using the Enclosure Insight Display, configure an IP address for the enclosure.

Configure the OA EBIPA

From the management server, use an automated method, manual procedure, or configuration file to push configuration to the OA, in particular, the EBIPA information for the Compute and IO Bays' management interfaces.



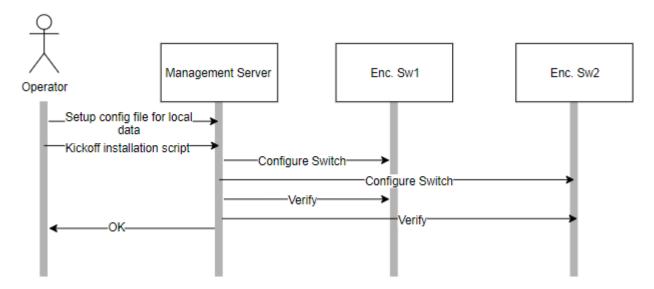
Figure B-6 Configure OAs



Configure the Enclosure Switches

Update switch configuration templates and/or tools with site specific information. Using the switch installation scripts or templates, push the configuration to the switches.

Figure B-7 Configure the Enc. Switches



Engage Customer Downlinks to Frame

At this point, the management server and switches are configured and can be joined to the customer network. Enable the customer uplinks.

Setup Installation Tools

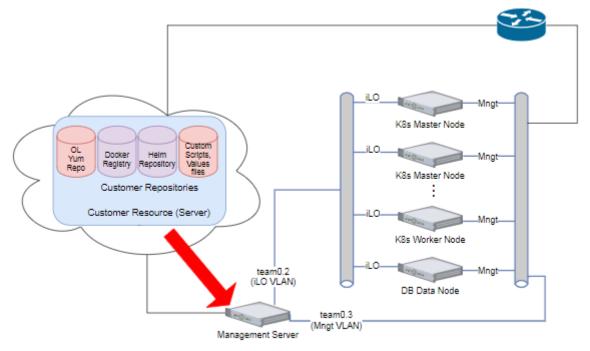
With all frame networking assets (ToR and Enclosure switches) configured and online, the rest of the frame can be setup from the management server.



Install OceanSpray Tools

Install the OceanSpray solution on the Management Server: Host OS Provisioner, Kubespray Installer, Configurator (Helm installer). This will require the management server to pull from the customer-provided docker registry.

Figure B-8 OceanSpray Download Path



Bring down OceanSpray Installer Containers and OL ISO

Configure site specific details in configuration files

Where appropriate, update configuration files with site specific data (hosts.ini, config maps, etc).

Install the Host OS on All Compute Nodes

Perform Host OS installations

Run Host OS Provisioner against all compute nodes (Master nodes, worker nodes, DB nodes).

Ansible Interacts with Server iLOs to perform PXe boot

Over an iLO network, Ansible will communicate to server iLOs to instruct the servers to reboot looking for a network boot option. In the below figure, note that the iLO network is considered to be on a private local network. This isn't a functional requirement, however, it does limit attack vectors. The only potential reason known to the author to make this public would be for sending alarms or telemetry to external NMS stations. We are expecting to feed this same telemetry and alerts into the Cluster. Thus, the iLOs are intending to stay private.



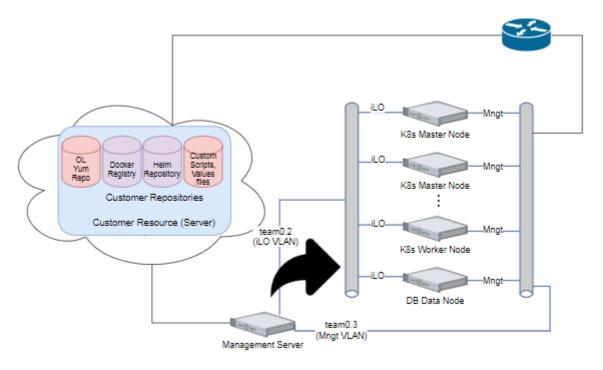


Figure B-9 Install OS on CNE Nodes - Server boot instruction

Ansible playbooks instruct servers to Pxe boot through iLO interaction

Servers Install Host OS

Servers boot by sending DHCP request out available NIC list. The broadcasts out the 10GE NICs are answered by the management server host OS provisioner setup. The management server provides the DHCP address, a boot loader, kickstart file and an OL ISO via NFS (a change in a future release should move this operation to HTTP).

At the end of this installation process, the servers should reboot.

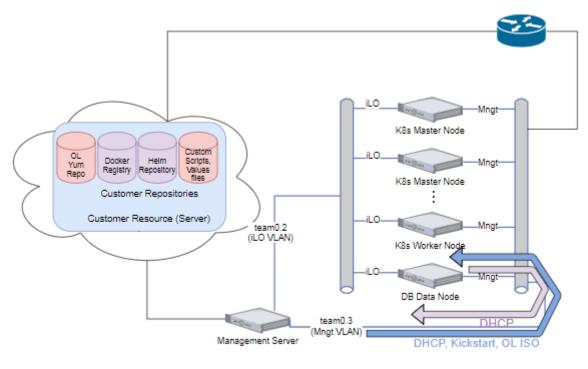


Figure B-10 Install OS on CNE Nodes - Server boot process

Servers boot looking for DHCP and an image

Package Update

At this point, server's host OS is installed, hopefully from the latest OL release. If this was done from a released ISO, then this step involves updating to the latest Errata. If the previous step already involved grabbing the latest package offering, then this step is already taken care of.

Ansible triggers servers to do a Yum update

Ansible playbooks interact with servers to instruct them to perform a Yum update.



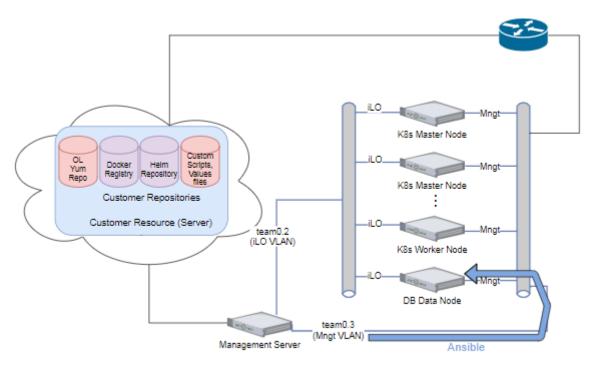


Figure B-11 Update OS on CNE Nodes - Ansible

Servers instructed to do a Yum update

Servers do a Yum update

Up to this point, the host OS management network could have been a private network without access to the outside world. At this point, the servers have to reach out to the defined repositories to access the Yum repository. Implementation can choose to either provide public addresses on the host OS instances, or a NAT function can be employed on the routers to hide the host OS network topology. If a NAT is used, it is expected to be a 1 to n NAT, rather than a 1 to 1. Further, ACLs can be added to prevent any other type of communication in or out of the frame on this network.

At the end of this installation process, the servers should reboot.



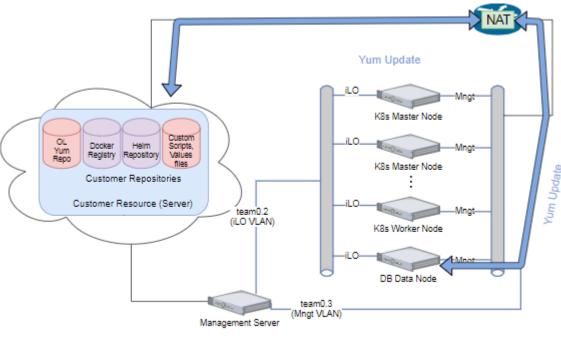


Figure B-12 Update OS on CNE Nodes - Yum pull

Servers perform a Yum update Servers instructed to do a Yum update

Harden the OS

Ansible instructs the servers to run a script to harden the OS.



OL Dooker Helm Scripts,
Registry Repository Values
Gustomer Repositories
Customer Resource (Server)

Team0.2
(iLO VLAN)

Was Worker Node

iLO Mngt

K8s Master Node

iLO Mngt

K8s Worker Node

iLO Mngt

K8s Worker Node

Rem0.3

Mngt

Nanagement Server

Run Hardening Script

Figure B-13 Harden the OS

Servers perform a Yum update

Install VMs as Needed

Some hosts in the CNE solution are to have VMs to address certain functionality, such as the DB service. The management server has a dual role of hosting the configuration aspects as well as hosting a DB data node VM. The K8s master nodes are to host a DB management node VM. This section shows the installation process for this activity.

Create the Guests

Ansible creates the guests on the target hosts.



OL Dooker Helm Scripts, Repository Repositories
Customer Resource (Server)

Customer Resource (Server)

Iteam0.2

(iLO VLAN)

Management Server

Run Hardening Script

Figure B-14 Create the Guest

Servers perform a Yum update

Install the Guest OS

Following a similar process to sections 2.5.1-2.5.3, the VM OS is installed, updated, and hardened. The details of how this is done is slightly different than the host OS, as an iLO connection is not necessary; however, they are similar enough that they won't be detailed here.

Install MySQL

Execute Ansible Playbooks from DB Installer Container

Show simple picture of Ansible touching the DB nodes.

Install Kubernetes on CNE Nodes

Customize Configuration Files

If needed, customize site-specific or deployment specific files.

Run Kubespray Installer

For each master and worker node, install the cluster.

Ansible/Kubespray Reaches Out to Servers to Perform Install



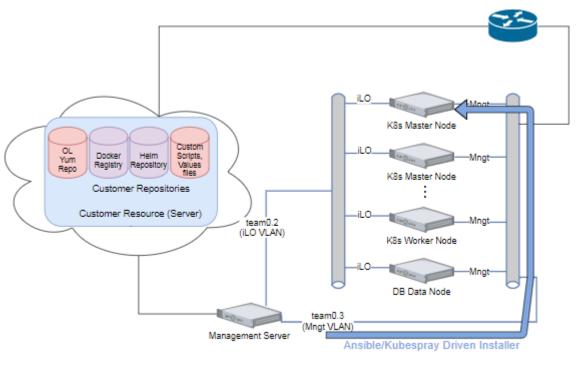


Figure B-15 Install the Cluster on CNE Nodes

Servers instructed to install Kubernetes

Servers Reach Out to Repos and Install Software

This is the 2nd instance where the host OS interfaces need to reach a distant repo. Thus, another NAT traversal is needed. Any ACL restricting access in/out of the solution needs to account for this traffic.



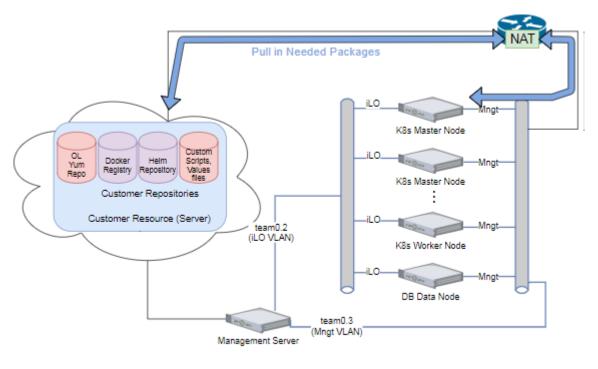


Figure B-16 Install the Cluster on CNE Nodes - Pull in Software

Servers instructed to install Kubernetes

Configure Common Services on CNE Cluster

Customize Site or Deployment Specific Files

If needed, customize site or deployment specific files, such as values files.

Run Configurator on Kubernetes Nodes

Install the Common Services using Helm install playbooks. Kubernetes will ensure appropriate distribution of all Common Services in the cluster.

Ansible Connects to K8s to Run Helm

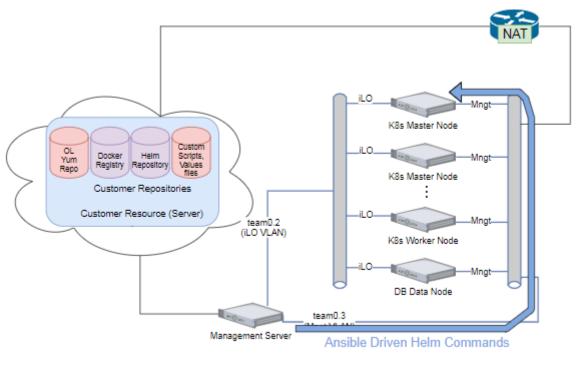


Figure B-17 Execute Helm on Master Node

Kubernetes Helm Commands

Helm Pulls Needed Items from Repositories

In this step, the Cluster IP sources the communication to pull from Helm repositories and Docker registries to install the needed services. The Values files used by Helm are provided in the Configurator container.



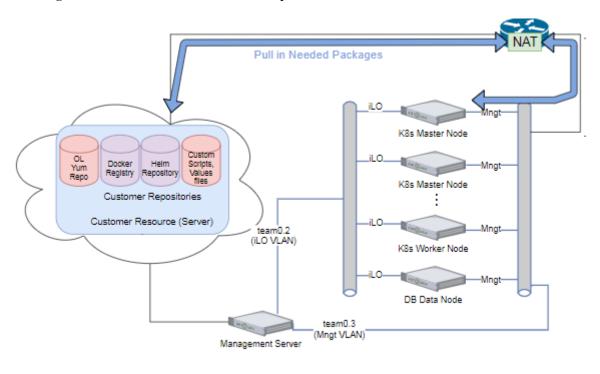


Figure B-18 Master Node Pulls from Repositories

Servers instructed to install Kubernetes

Topology Connection Tables

Enclosure Connections

Blade Server Connections

The HP BL460 Gen10 Blade Server will have a base set of NICs that connect internally to the enclosure's IO bays 1 and 2. These connections are "hard-wired" and won't be documented here. Since no additional NICs are planned beyond the base pair of NICs, no further declarations of network connectivity on the blades needs mentioning. Blade specifications are linked here:

https://www.hpe.com/us/en/product-catalog/servers/proliant-servers/pip.specifications.hpe-proliant-bl460c-gen10-server-blade.1010025832.html.

OA Connections

The Enclosure's Onboard Administrator (OA) will be deployed as a redundant pair, with each connecting with 1GE copper connection to the respective ToR switches' SFP+ ports.

Topology Connections

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-14 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
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			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



Table B-14 (Cont.) Enclosure Switch Connections

Switch Port Name/ID (From)	Destination (To)	Cable Type	Module Required
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-15 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



Table B-15 (Cont.) ToR Switch Connections

Switch Port Name/ID (From)	From Switch 1 to Destination	From Switch 2 to Destination	Cable Type	Module Required
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
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



Table B-15 (Cont.) ToR Switch Connections

Switch Port Name/ID (From)	From Switch 1 to Destination	From Switch 2 to Destination	Cable Type	Module Required
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
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

 $Server\ quick specs\ can\ be\ found\ here:\ https://h20195.www2.hpe.com/v2/getdocument.aspx?\ docname=a00008180enw$

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-16 Management 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)

Network Redundancy Mechanisms

Network Redundancy Mechanisms

This section is intended to cover the redundancy mechanisms used within the CNE Platform. With all links, BFD should be investigated as an optimal failure detection strategy.

Server Level Redundancy

The blade server hardware will be configured with a base pair of 10GE NICs that map to Enclosure IO bays 1 and 2. IO bays 1 and 2 will be equipped with 10GE switches. The blade server OS configuration must pair the base pair of NICs in an active/active configuration.



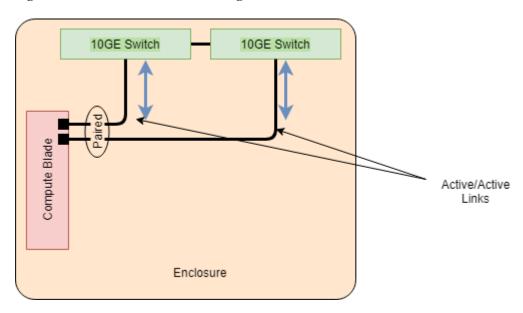
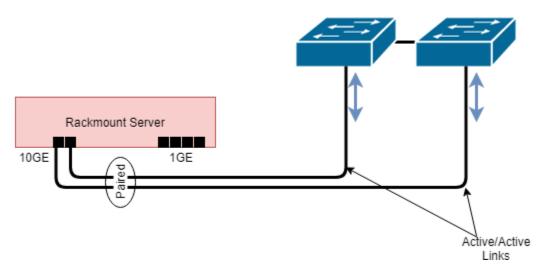


Figure B-19 Blade Server NIC Pairing

The rackmount servers (RMS) will be configured with a base quad 1GE NICs that will be mostly unused (except for switch management connections on the management server). The RMS will also be equipped with a 10GE Flex LOM (FLOM) card. The FLOM NIC ports will be connected to the ToR switches. The RMS OS configuration must pair the FLOM NIC ports in an active/active configuration.

Figure B-20 Rackmount Server NIC Pairing



Production Use-Case

The production environment will use LACP mode for active/active NIC pairing. This is so the NICs can form one logical interface, using a load-balancing algorithm involving a hash of source/dest MAC or IP pairs over the available links. For this to work, the upstream switches need to be "clustered" as a single logical switch. LACP mode will not work if the upstream switches are operating as independent switches (not sharing the same switching fabric). The

current projected switches to be used in the solution are capable of a "clustering" technology, such as HP's IRF, and Cisco's vPC.

Lab Use-Case

Some lab infrastructure will be able to support the production use-case. However, due to its dependence on switching technology, and the possibility that much of the lab will not have the dependent switch capabilities, the NIC pairing strategy will need to support an active/active mode that does not have dependence on switch clustering technologies (adaptive load balancing, round-robin), active/standby that does not have dependence on switch clustering technologies, or a simplex NIC configuration for non-redundant topologies.

Enclosure Switch Redundancy

To support LACP mode of NIC teaming, the Enclosure switches will need to be clustered together as one logical switch. This will involve the switches to be connected together with the Enclosure's internal pathing between the switches. Below is an Enclosure switch interconnect table for reference.

Each blade server will form a 2x10GE Link Aggregation Group (LAG) to the upstream enclosure switches. Up to 16 blades will be communicating through these enclosure switches. Without specific projected data rates, the enclosure uplinks to the Top of Rack (ToR) switches will be sized to a 4x10GE LAG each. Thus, with the switches logically grouped together, an 8x10GE LAG will be formed to the ToR.

For simplicity's sake, the figure below depicts a single black line for each connection. This black line may represent one or more links between the devices. Consult the interconnect tables for what the connection actually represents.



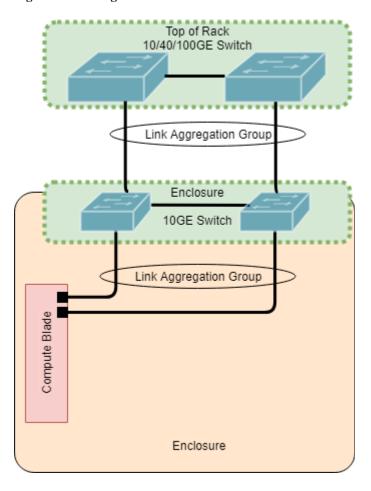


Figure B-21 Logical Switch View

ToR Switch Redundancy

The ToR switch is intended to be deployed as a single logical unit, using a switch clustering technology. If the ToR switch does not support switch clustering, then, minimally, the switch must support Virtual Port channeling technology that allows the downlinks to both enclosure switches to be logically grouped into a common port channel. This is to simplify design, increase throughput, and shorten failover time in the event of a switch failure. Connections upstream to the customer network will be dependent on customer requirements. For the initial launch of CNE, the uplink requirement will vary depending on OAM or Signaling uplink requirements.

OAM Uplink Redundancy

The OAM network uplinks are expected to use static routing in the first deployment of CNE for a targeted customer. This means that the first ToR switch will have a default route to a specific customer OAM router interface, while the other ToR switch has a default route to a different customer OAM router interface. If the ToR switches are clustered together as one logical switch, then this behavior should still work in certain failure scenarios, but more testing would be needed. For example, if the link to OAM subnet 1 were to go down, then the switch physical or virtual interface to OAM subnet 1 should go down, and thus be removed from the route table. If, however, the switch link doesn't go down, but the customer OAM router interface



were to become unreachable, the static default route to OAM subnet 1 would still be active, as there is no intelligence at play to converge to a different default route. This is an area in need of further exploration and development.

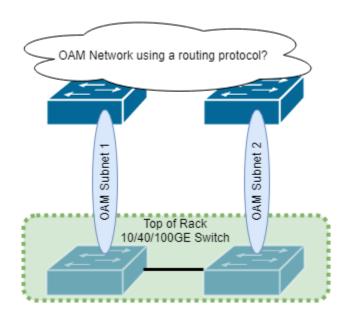


Figure B-22 OAM Uplink View

Signaling Uplink Redundancy

The signaling network uplinks are expected to use OSPF routing protocol with the customer switches to determine optimal and available route paths to customer signaling router interfaces. This implementation will require tuning with the customer network for optimal performance.

If the ToR switches are able to cluster together as one logical switch, then there is no need for an OSPF relationship between the ToR switches. In this case, they would share a common route table and have two possible routes out to the customer router interfaces. If, however, they do not cluster together as a single logical unit, then there would be an OSPF relationship between the two to share route information.



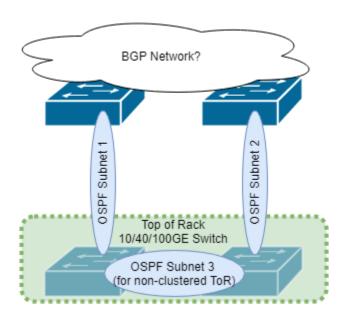


Figure B-23 Top of Rack Customer Uplink View

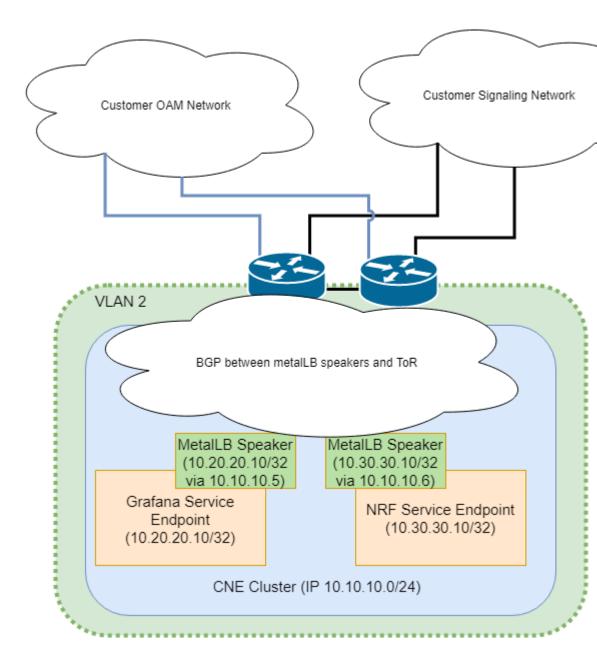
OAM and Signaling Separation

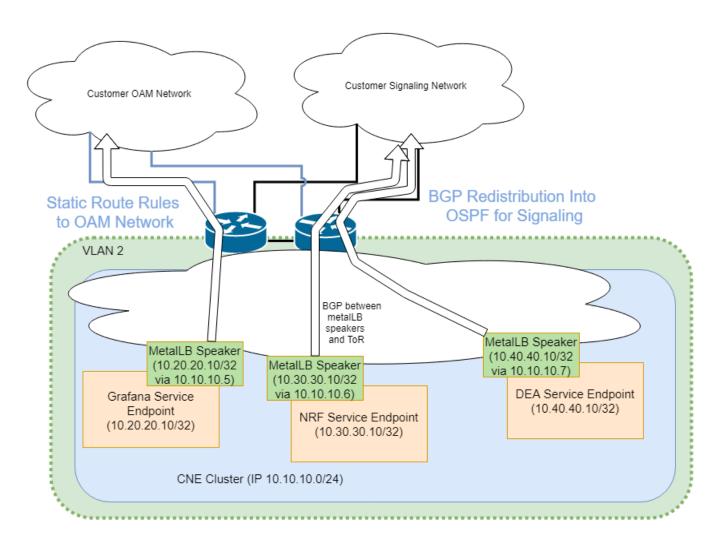
Cloud Native networks are typically flat networks, consisting of a single subnet within a cluster. The Kubernetes networking technology is natively a single network IP space for the cluster. This presents difficulty when deploying in telecom networks that still maintain a strict OAM and Signaling separation in the customer infrastructure. The OC-CNE will provide metalLB load-balancer with BGP integration to the ToR switches as a means to address this problem. Each service end-point will configure itself to use a specific pool of addresses configured within metalLB. There will be two address pools to choose from, OAM and Signaling. As each service is configured with an address from a specific address pool, BGP will share a route to that service over the cluster network. At the ToR, some method of advertising these address pools to OAM and signaling paths will be needed. OAM service endpoints will likely be addressed through static route provisioning. Signaling service endpoints will likely redistribute just one address pool (signaling) into the OSPF route tables.

OAM and Signaling configurations and test results are in the following page:



Figure B-24 OAM and Signaling Separation





OAM type common services, such as EFK, Prometheus, and Grafana will have their service endpoints configured from the OAM address pool. Signaling services, like the 5G NFs, will have their service endpoints configured from the signaling address pool.

Install VMs for MySQL Nodes and Management Server

The following procedure creates Virtual Machines (VM's) for installing MySQL Cluster nodes (Management nodes, Data nodes and SQL nodes) and creating Bastion Host VM on the each Storage Host, install Oracle Linux 7.5 on each VM. This procedure requires all the network information for creating the VM's in different host servers like k8 Master Nodes, Storage Hosts. Here VM's are created using the virt-install CLI tool using the provision docker container and MySQL Cluster is installed using the db-install docker container as outlined in the Database Tier Installer.

After all the hosts are provisioned using the provision container, this procedure is used for creating the VM's in kubernetes Master nodes and Storage Hosts.

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.
- 2. The MySQL Cluster Manager client provides a command-line interface to the agent's management functions.

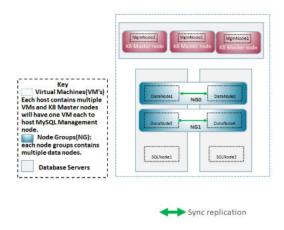
MySQL Cluster Manager binary distributions that include MySQL NDB Cluster are used for installing MySQL Cluster Manager and MySQL NDB Cluster.

Steps for downloading the MySQL Cluster Manager from Oracle Software Delivery Cloud (OSDC) is found in Installation Preflight Checklist

MySQL Cluster Topology

In OCCNE 1.2, MySQL Cluster is installed as shown below in each cluster.

Figure B-25 MySQL Cluster Topology



Virtual Machine

MySQL Cluster is installed on Virtual machines, so the number of VM's required in the Storage Hosts and K8 Master Nodes are as shown below. Each k8 master node is used to create 1 VM for installing the MySQL Management node, so there are 3 MySQL management nodes in the MySQL Cluster. In each storage nodes, 4 VM's are created, i.e. 2 VM's for data nodes, 1 VM for SQL nodes and 1 VM for Management node VM.

No Of MySQL Management Nodes: 3

No Of Data Nodes: 4 No of SQL nodes: 2

No Of Bastion Hosts: 2

Below table shows VM's Created in Host servers:

Host Server	No Of VM's	Node Name	
K8 Master node 1	1	MySQL management node 1	
K8 Master node 2	1	MySQL management node 2	



Host Server	No Of VM's	Node Name
K8 Master node 3	1	MySQL management node 3
Storage Host 1	4	2 MySQL Data Node, 1 MySQL SQL node, 1 Bastion Host VM
Storage Host 2	4	2 MySQL Data Node,1 MySQL SQL node, 1 Bastion Host VM

VM Profile for Management and MySQL Cluster Nodes:

Node Type	RAM	HDD	vcpus	No Of Nodes
MySQL Management Node	8GB	200GB	4	3
MySQL Data Node	50GB	800GB	10	4
MySQL SQL Node	16GB	600GB	10	2
Bastion Host	8GB	300GB	4	2

System Details

IP Address, host names for VM's, Network information for creating the VM's are captured in Installation Preflight Checklist. Configuration details of these VM's are in Inventory file: Inventory File Preparation

Prerequisites

- 1. All the hosts servers where VM's are created are captured in Inventory File Preparation, The kubernetes master nodes are mentioned under [kube-master] and Storage Hosts are mentioned under [data store].
- All Hosts should be provisioned using provision docker container as defined and installed site hosts.ini file.
- 3. Host names and IP Address, network information assigned to these VM's should be captured in the Installation Preflight Checklist
- 4. Bastion Host should be installed in Storage Host(RMS2) and configured in Storage Host (RMS2).
- 5. SSH keys configured in host servers by provision container is stored in Bastion Host (RMS2).
- VM's should be created before installing the kubernetes in k8's Master and kube worker nodes.

Limitations and Expectations

- 1. Both Storage Hosts will have one Management server VM, where Docker is installed, All the host servers are provisioned from this Management server VM.
- Once both storage nodes and host servers are provisioned using the provision container, VM's are created on kubernetes master and DB storage nodes.



References

- 1. https://linux.die.net/man/1/virt-install
- 2. https://linuxconfig.org/how-to-create-and-manage-kvm-virtual-machines-from-cli
- 3. https://www.cyberciti.biz/faq/kvm-install-centos-redhat-using-kickstart-ks-cfg/
- 4. https://opensource.com/business/16/9/linux-users-guide-lvm

Add bridge interface in all the hosts

Create a bridge interface on the Team(team0) interface for creating VM's.



Note:

Below steps should be performed to create the bridge interface (teambr0 and vlan5-br) in each storage hosts and bridge interface(teambr0) in each kubernetes Master nodes one at a time.



Table B-17 Procedure to install VMs for MySQL Nodes and Management Server

Step #	Procedure	Des	scription
1.	Create VMs	1.	Login to the Bastion Host in RMS2.
Ш			\$ sudo su
		2.	Configure inventory file with all the VM's details as described in Inventory File Preparation
		3.	Update skip_kernel_virtual group in the inventory file with the bastion hosts.
			<pre>\$ vi /var/occne/<cluster_name>/hosts.ini [skip_kernel_virtual] bastion-1.delta.lab.us.oracle.com bastion-2.delta.lab.us.oracle.com</cluster_name></pre>
		4.	Create VM's for installing MySQL NDB Cluster. Update bastion hosts in the skip_kernel_virtual group in hosts.ini inventory file.
			<pre>\$ vi /var/occne/<cluster_name>/hosts.ini # List of other VM's which are already created. [skip_kernel_virtual] bastion-1.rainbow.lab.us.oracle.com bastion-2.rainbow.lab.us.oracle.com</cluster_name></pre>
			Run the docker command below to create a provision container running bash
			<pre>\$ docker runrmnetwork hostcap- add=NET_ADMIN -v /var/occne/<cluster_name>/:/ host -v /var/occne/:/var/occne:rw -e "OCCNEARGS=limit <limit_filter>" <image_name>:<image_tag> /bin/bash</image_tag></image_name></limit_filter></cluster_name></pre>
			Example:
			<pre>\$ docker run -itrmnetwork hostcap- add=NET_ADMIN -v /var/occne/ rainbow.lab.us.oracle.com/:/host -v /var/ occne/:/var/occne:rw -e "OCCNEARGS=" 10.75.200.217:5000/provision:1.2.0 /bin/bash</pre>
			Run below command to create VM's
			<pre>\$ ansible-playbook -i /host/hosts.ini vms_provision.yaml \${OCCNEARGS}skip-tags "datastore"</pre>
		5.	Execute security playbook on the VM's
			<pre>\$ docker run -itrmnetwork hostcap- add=NET_ADMIN -v /var/occne/ rainbow.lab.us.oracle.com/:/host -v /var/ occne/:/var/occne:rw -e "OCCNEARGS=limit mysqlndb_all_nodes" 10.75.200.217:5000/provision: 1.2.0 ./run_security</pre>



Table B-17 (Cont.) Procedure to install VMs for MySQL Nodes and Management Server

Step #	Procedure	Description
Packages nodes and Storage H below steps to install Hosts.		By default KVM packages are installed in the Kubernetes Master nodes and Storage Host, so if KVM packages are not installed, follow below steps to install KVM packages in Master nodes and Storage Hosts. Execute provision container to install KVM Packages.
		<pre>\$ docker runrmnetwork hostcap-add=NET_ADMIN -v /var/occne/<cluster_name>/:/host -v /var/ occne/:/var/occne:rw -e 'OCCNEARGS=limit <limit_filter>tags datastore' <image_name>:<image_tag></image_tag></image_name></limit_filter></cluster_name></pre>
		Example:
		<pre>docker run -itrmnetwork hostcap- add=NET_ADMIN -v /var/occne/<cluster_name>/:/host - v /var/occne/:/var/occne:rw -e 'OCCNEARGS=limit host_hp_gen_10:!db-2.delta.lab.us.oracle.comtags datastore' 10.75.200.217:5000/provision:1.2.0</cluster_name></pre>

