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 1.3

F24135-02

Copyright © 2019, 2020, 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

| Glossary | 1 |
|--|---|
| Key terms | 1 |
| Key Acronyms and Abbreviations | 1 |
| How to use this document | 1 |
| Documentation Admonishments | 1 |
| Locate Product Documentation on the Oracle Help Center Site | |
| Customer Training | |
| My Oracle Support | |
| Emergency Response | |
| Installation Procedures | |
| Bare Metal Installation | |
| OCCNE Installation Overview | |
| Frame and Component Overview | |
| Create OCCNE Instance | |
| Installation Prerequisites | |
| Configure Artifact Acquisition and Hosting | |
| Initial Configuration - Prepare a Minimal Boot Strapping Environment | |
| Installation of Oracle Linux 7.5 on Bootstrap Host | |
| Configure the Installer Bootstrap Host BIOS | 2 |
| Configure Top of Rack 93180YC-EX Switches | 2 |
| Configure Addresses for RMS iLOs, OA, EBIPA | 2 |
| Configure Legacy BIOS on Remaining Hosts | 2 |
| Configure Enclosure Switches | 2 |
| Bastion Host Installation | 2 |
| Automated Installation | 2 |
| Virtualized CNE Installation | 2 |
| Prerequisites | 2 |
| Limitation/Expectations | 2 |
| Upload an Image to an OpenStack Environment | 2 |



Bootstrap Host Creation

2-74

| Pre-deployment Configuration | 2-7 |
|---|------------------------------|
| Deploy the OCCNE Virtualized Cluster | 2-8 |
| Post Installation Activities | |
| Post Install Verification | 3- |
| Post-Installation Security Hardening | 3- |
| Artifact Acquisition and Hosting | |
| YUM Repository Configuration | A- |
| HTTP Repository Configuration | A- |
| Docker Image Registry Configuration | A-1 |
| Reference Procedures | |
| Inventory File Preparation | В- |
| Installation Preflight Checklist | B- |
| Inventory File Template | B-2 |
| Configuring Cross-Site Database Replication | B-2 |
| Install Db Monitor service | |
| | B-3 |
| Install Additional Services/Network Functions | |
| Install Additional Services/Network Functions Change MySQL root user password | B-4 |
| | B-3: B-4: B-4: B-4: |



List of Figures

| 1-1 | Example of a Procedure Steps Used in This Document | 1-4 |
|-----|--|-----|
| 2-1 | Frame Overview | 2-2 |
| 2-2 | Host Designations | 2-3 |
| 2-3 | Node Roles | 2-4 |
| 2-4 | Transient Roles | 2-5 |
| 2-5 | OCCNE Installation Overview | 2-6 |
| B-1 | Rackmount ordering | B-7 |



List of Tables

| 1-1 | Key Terms | 1-1 |
|------|---|------|
| 1-2 | Key Acronyms and Abbreviations | 1-2 |
| 1-3 | Admonishments | 1-4 |
| 2-1 | Oracle eDelivery Artifact Acquisition | 2-7 |
| 2-2 | Procedure to configure MetalLB pools and peers | 2-8 |
| 2-3 | Bootstrap Install Procedure | 2-10 |
| 2-4 | Procedure to configure the Installer Bootstrap Host BIOS | 2-17 |
| 2-5 | Procedure to configure Top of Rack 93180YC-EX Switches | 2-22 |
| 2-6 | Procedure to verify Top of Rack 93180YC-EX Switches | 2-32 |
| 2-7 | SNMPv2c configuration for Top of Rack 93180YC-EX Switches | 2-35 |
| 2-8 | Procedure to configure Addresses for RMS iLOs, OA, EBIPA | 2-38 |
| 2-9 | Procedure to configure the Legacy BIOS on Remaining Hosts | 2-47 |
| 2-10 | Procedure to configure enclosure switches | 2-53 |
| 2-11 | Terminology used in Procedure | 2-59 |
| 2-12 | Bastion Installation | 2-60 |
| 2-13 | Automated Installation | 2-68 |
| 2-14 | Upload an Image to an OpenStack Environment | 2-73 |
| 2-15 | Bootstrap Host Creation | 2-74 |
| 2-16 | Pre-deployment Configuration | 2-76 |
| 2-17 | Deploy the OCCNE Virtualized Cluster | 2-86 |
| 3-1 | OCCNE Post Install Verification | 3-2 |
| 3-2 | Credentials | 3-6 |
| 3-3 | Post-Installation Security Hardening | 3-8 |
| A-1 | Procedure to configure OCCNE YUM Repository | A-3 |
| A-2 | Steps to configure OCCNE HTTP Repository | A-5 |
| B-1 | Base Groups | B-2 |
| B-2 | Data Tier Groups | В-3 |
| B-3 | Edit occne:vars | B-5 |
| B-4 | Edit mysqlndb_all_nodes:vars | B-6 |
| B-5 | Enclosure Switch Connections | B-8 |
| B-6 | ToR Switch Connections | B-9 |
| B-7 | Rackmount Server Connections | B-12 |
| B-8 | Complete Site Survey Subnet Table | B-12 |
| B-9 | Complete Site Survey Host IP Table | B-13 |
| B-10 | Complete VM IP Table | B-14 |



| B-11 | Complete OA and Switch IP Table | |
|------|--|------|
| B-12 | ToR and Enclosure Switches Variables Table (Switch Specific) | B-18 |
| B-13 | Complete Site Survey Repository Location Table | B-19 |
| B-14 | OCCNE Variables | B-28 |
| B-15 | Configure replication | B-30 |
| B-16 | Configure Site 2 | B-31 |
| B-17 | Install Db Monitor service | B-39 |
| B-18 | Install Additional Services/Network Functions | B-41 |
| B-19 | Change MySOL root user password | B-42 |



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 1.3 of the 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. |
| Virtualized CNE | A virtualized CNE is a cloud native environment that is deployed on VMs, rather than on bare metal servers. |

Key Acronyms and Abbreviations

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

Table 1-2 Key Acronyms and Abbreviations

| Acronym/ Abbreviation/Term | Definition |
|-------------------------------|--|
| 5G NF | 3GPP 5G Network Function |
| BIOS | Basic Input Output System |
| CLI | Command Line Interface |
| CNE | Cloud Native Environment |
| DB | Database |
| DBMS | Database Management System |
| DHCP(D) | Dynamic Host Configuration Protocol |
| DNS | Domain Name Server |
| EBIPA | Enclosure Bay IP Addressing |
| FQDN | Fully Qualified Domain name |
| 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 |
| vCNE | Virtualized CNE |
| VSP | Virtual Serial Port |
| YUM | Yellowdog Updator, Modified (a Linux Package Manager) |

How to use this document

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

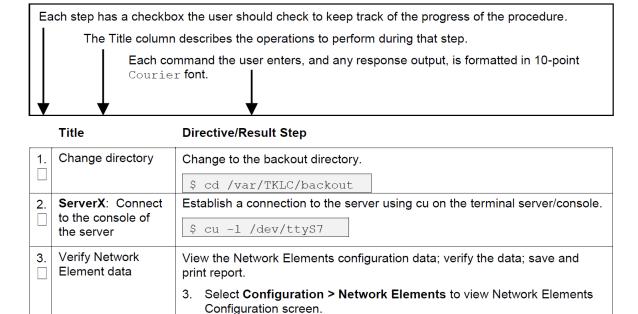
When executing this document for either purpose, there are a few points which help to ensure that the user understands the author's intent. These points are as follows:



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



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 |
|--------|---|
| | Danger: |
| | (This icon and text indicate the possibility of |
| | personal injury.) |
| DANGER | |



Table 1-3 (Cont.) Admonishments

| Icon | Description |
|----------|--|
| <u> </u> | Warning: |
| WARNING | (This icon and text indicate the possibility of equipment damage.) |
| | Caution: |
| CAUTION | (This icon and text indicate the possibility of service interruption.) |

Locate Product Documentation on the Oracle Help Center Site

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

- 1. Access the Oracle Help Center site at http://docs.oracle.com.
- 2. Click Industries.
- 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 Procedures

The installation procedures in this document provision and configure an Oracle Communications Signaling, Network Function Cloud Native Environment (OCCNE). OCCNE 1.3 offers the choice of deployment platform; the CNE can be deployed directly onto dedicated hardware, (referred to as a bare metal CNE), or deployed onto OpenStack-hosted VMs. (referred to as a virtualized CNE).

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

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



In case any procedure requires Linux Shell access, make sure that the shell is with Keepalive to avoid unexpected timeout.

Bare Metal Installation

This section describes the procedure to install OCCNE 1.3.0 onto dedicated bare metal hardware.

OCCNE Installation Overview

Frame and Component Overview

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



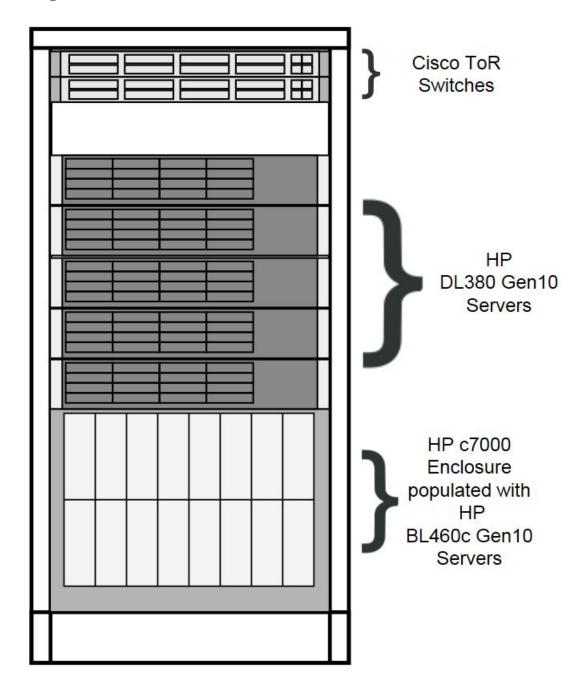
In the installation process, some of the roles of servers change as the installation procedure proceeds.



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.

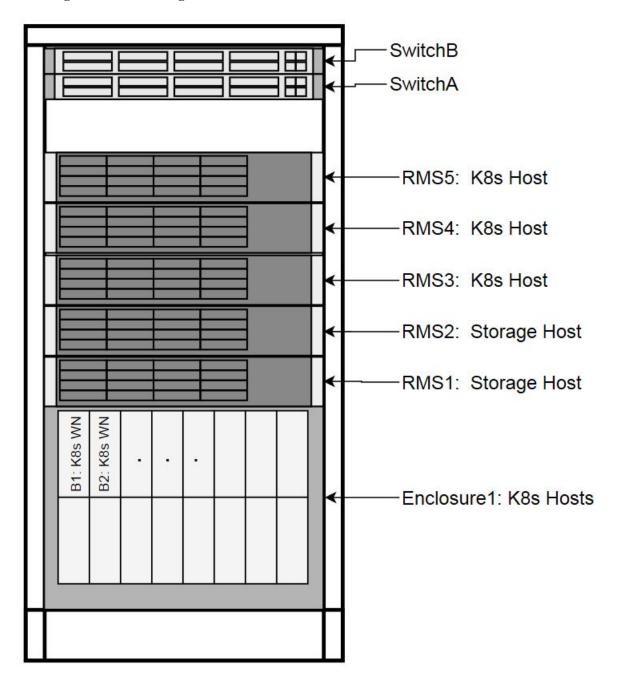
Figure 2-1 Frame Overview



Host Designations

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

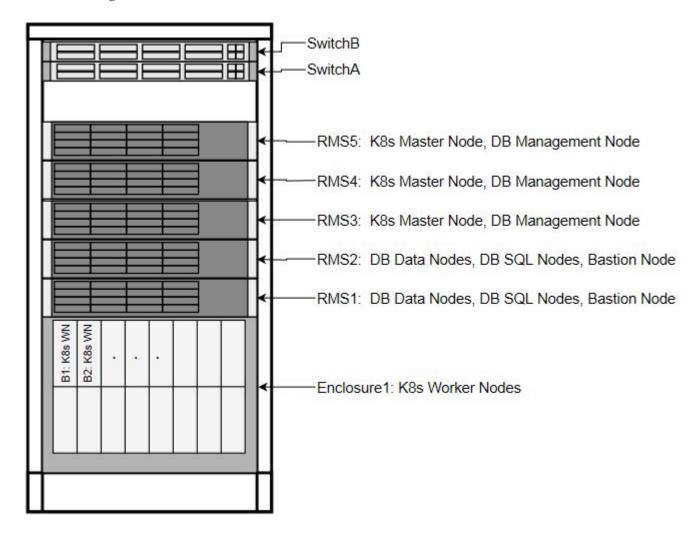
Figure 2-2 Host Designations



Node Roles

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

Figure 2-3 Node Roles

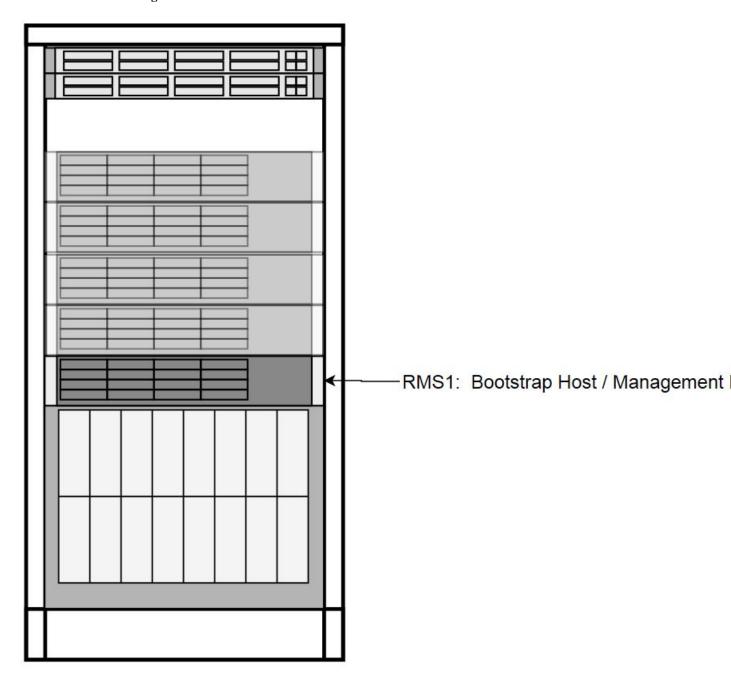


Transient Roles

Transient role is unique in that it has OOB connections to the ToR switches, which 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 2-4 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
Asset Repository
Target
Frame
Local Private Network

Figure 2-5 OCCNE Installation Overview

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

- 1. Check that the hardware is on-site and properly cabled and powered up.
- 2. Pre-assemble the basic ingredients needed to perform a successful install:

a. Identify

- i. Download and stage software and other configuration files using provided manifests. Refer to Artifact Acquisition and Hosting for manifests information.
- ii. Identify the layer 2 (MAC) and layer 3 (IP) addresses for the equipment in the target frame
- iii. Identify the addresses of key external network services (e.g., NTP, DNS, etc.)
- iv. Verify / Set all of the credentials for the target frame hardware to known settings

b. Prepare

- i. Software Repositories: Load the various SW repositories (YUM, Helm, Docker, etc.) using the downloaded software and configuration
- ii. Configuration Files: Populate the hosts inventory file with credentials and layer 2 and layer 3 network information, switch configuration files with assigned IP addresses, and yaml files with appropriate information.

3. Bootstrap the System:

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 RMS2 as an OCCNE Storage Host are installed.
- b. Using the newly constructed MBE, automatically create the first (complete) Management VM on RMS2. This freshly installed Storage Host will include a virtual machine for hosting the Bastion Host.
- c. Using the newly constructed Bastion Host on <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 Installation Procedures) as a temporary Installer Bootstrap Host. The Bootstrap Host is configured with a minimal set of packages needed to configure switches, iLOs, PXE boot environment, and provision RMS2 as an OCCNE Storage Host. A virtual Bastion Host is also provisioned on RMS2. The Bastion Host is then used to provision (and in the case of the Bootstrap Host, re-provision) the remaining OCCNE hosts, install Kubernetes, Database services, and Common Services running within the Kubernetes cluster.

Installation Prerequisites

Complete the procedures outlined in this section before moving on to the Install Procedures section. OCCNE installation procedures require certain artifacts and information to be made available prior to executing installation procedures. This section addresses these prerequisites.

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-1.3.x.tgz | OCCNE Installers (Docker images) | Tar GZ | Docker Registry |
| v980756-01.zip | Zip file of MySQL Cluster Manager 1.4.7+Cluster | Zip of tar file | File repository MySQL Repository Requirements |



Table 2-1 (Cont.) Oracle eDelivery Artifact Acquisition

| Artifact | Description | File Type | Destination Repository |
|----------------|--|----------------------------|--|
| v975367-01.iso | OL7 ISO | ISO | File repository Oracle Linux 7.5 Download Instructions |
| Install Docs | These Install Procedures from OHC | PDFs | N/A |
| Templates | Switch config files, hosts.ini file templates from OHC | Config files (.conf, .ini) | Local media |

Third Party Artifacts

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

Reference Topic Title

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 2-2 Procedure to configure MetalLB pools and peers

| Step # | Procedure | Description |
|--------|--|---|
| 1. | Add BGP peers and address groups | Referring to the data collected in the Preflight Checklist, add BGP peers (ToRswitchA_Platform_IP, ToRswitchB_Platform_IP) and address groups for each address pool. Address-pools list the IP addresses that metalLB is allowed to allocate. |



Table 2-2 (Cont.) Procedure to configure MetalLB pools and peers

| Step # | Procedure | Description |
|--------|----------------------------------|--|
| 2. | Edit the mb_configmap.ya ml file | Edit the mb_configmap.yaml file with the site-specific values found in the Preflight Checklist Note: The name "signaling" is prone to different spellings (UK vs US), therefore pay special attention to how this signaling pool is referenced. configInline: peers: - peer-address: <torswitcha_platform_ip> peer-asn: 64501 my-asn: 64512 - peer-address: <torswitchb_platform_ip> peer-asn: 64501 my-asn: 64512 address-pools: - name: signaling protocol: bgp auto-assign: false addresses: - '<metallb_signal_subnet_ip_range>' - name: oam protocol: bgp auto-assign: false addresses: - '<metallb_oam_subnet_ip_range>'</metallb_oam_subnet_ip_range></metallb_signal_subnet_ip_range></torswitchb_platform_ip></torswitcha_platform_ip> |

Initial Configuration - Prepare a Minimal Boot Strapping Environment

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

Installation of Oracle Linux 7.5 on Bootstrap Host

This procedure outlines the installation steps for installing OL7 onto the OCCNE Installer Bootstrap Host. This host is used to configure the networking throughout the system and install OL7 onto RMS2. The Bootstrap server 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

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

Bootstrap Install Procedure

Table 2-3 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. |
| | | <pre>\$ dd if=<path iso="" to=""> of=<usb device="" path=""> bs=1048576</usb></path></pre> |
| | | 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=1048576</pre> |



 Table 2-3
 (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. | | | | |
| | | 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. The following menu is displayed on the screen. Select Test this media & install Oracle Linux 7.x and hit ENTER . This begins the verification of the media and the boot process. | | | | |
| | | | After the verification reaches 100%, the following Welcome screen is displayed. When prompted for the language to use, select the default setting: English (United States) and hit Continue in the lower left corner. | | | | |
| | | 6. | The INSTALLATION SUMMARY page, is displayed. The following setting are expected: | | | | |
| | | | | 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) and select DONE in the upper right corner. If a dialog appears indicating there is not enough free space (which might mean an OS has already been installed), 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. | | | | |
| | | 7 | 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. | | | | |



 Table 2-3 (Cont.) Bootstrap Install Procedure

| Step # | Procedure | Description |
|--------|-----------|---|
| | | 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 2-3
 (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 2-3
 (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 fsdisk 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 2-3
 (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 status</pre> |
| | | 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 2-3 (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: |
| | | <pre>\$ cd /etc/xinetd.d \$ ls chargen-dgram chargen-stream daytime-dgram daytime-stream discard-dgram discard-stream echo-dgram echo-stream tcpmux-server tftp time- dgram time-stream</pre> |
| | | 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/sda1 is no longer shown as mounted to /media/usb.</pre> |
| | | 11. This procedure is complete. |

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 2-4
 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. |
| | | 1. After making the proper connections for the KVM on the back of the Bootstrap host to have access to the console, the user should reboot the host by momentarily pressing the power button on the front of the Bootstrap host. |
| | | 2. Expose the HP Proliant DL380 Gen10 System Utilities. Once the remote console has been exposed, the system must be reset to force it through the restart process. When the initial window is displayed, hit the F9 key repeatedly. Once the F9 is highlighted at the lower left corner of the remote console, it should eventually bring up the main System Utility. |
| | | 3. The System Utilities screen is exposed in the remote console. |
| 2. | 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. 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 2-4 (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></pre> in the Password field. |
| | | 6. Select F10: Save to save and stay in the utility or select the F12: Save and Exit to save and exit, to complete the current boot process. |



 Table 2-4
 (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 2-4 (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 |
| | | Expose the System Configuration Utility by following Step 1. |
| | | b. Select System Configuration. |
| | | c. Select BIOS/Platform Configuration (RBSU) |
| | | d. Select Virtualization Options This screen displays the settings for the Intel(R) Virtualization Technology (IntelVT), Intel(R) VT-d, and SR-IOV options (Enabled or Disabled). The default values for each option is Enabled. |
| | | e. Select F10: Save to save and stay in the utility or select the F12: Save and Exit to save and exit, to complete the current boot process. |
| 6. | Disable RAID Configurations | Expose the System Configuration Utility by following Step 1. |
| | | 2. Select System Configuration. |
| | | 3. Select Embedded RAID 1 : HPE Smart Array P408i-a SR Gen 10. |
| | | 4. Select Array Configuration. |
| | | 5. Select Manage Arrays. |
| | | 6. Select Array A (or any designated Array Configuration if there are more than one). |
| | | 7. Select Delete Array. |
| | | 8. Select Submit Changes. |
| | | 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 2-4
 (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 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

Procedures

Configuration

Table 2-5 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: OCCNE 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 2-5 (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 Addresses (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</torswitch_cnemanagementnet_vip></cne_management_ip_with_prefix> |
| | | <pre>ipv4.addresses 192.168.2.11/24 \$ nmcli con up mgmtBridge</pre> |
| | | <pre>\$ nmcli con add type team con-name team0 ifname team0 team.runner lacp \$ nmcli con add type team-slave con-name team0-slave-1 ifname eno5 master team0 \$ nmcli con add type team-slave con-name team0-slave-2 ifname eno6 master team0 \$ nmcli con mod team0 ipv4.method manual ipv4.addresses 172.16.3.4/24 \$ nmcli con add con-name team0.4 type vlan id 4 dev team0 \$ nmcli con mod team0.4 ipv4.method manual ipv4.addresses <cne_management_ip_address_with_prefix> ipv4.gateway <torswitch_cnemanagementnet_vip> \$nmcli con up team0.4</torswitch_cnemanagementnet_vip></cne_management_ip_address_with_prefix></pre> |



Table 2-5 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches

| Ston # | Procedure | Description |
|--------|--|--|
| Step # | | Description Classification of the Control of the Co |
| 4. | Edit the /etc/ xinetd.d/tftp file | Edit the /etc/xinetd.d/ffp 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 { |
| 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 2-5 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches

| Step # | Procedure | Description |
|--------|---|--|
| 9. | Copy the ifcfg template files | Copy the ifcfg template files to /tmp directory for later use. \$ cp /media/usb/ifcfg-vlan /tmp \$ cp /media/usb/ifcfg-bridge /tmp |
| 10. | Modify POAP script File. | Modify POAP script File. Make the following change for the first server information: The username and password are the credentials used to login to the Bootstrap host. \$ vi /var/lib/tftpboot/poap_nexus_script.py Host name and user credentials options = { "username": " <username>", "password": "<password>", "hostname": "192.168.2.11", "transfer_protocol": "scp", "mode": "serial_number", "target_system_image": "nxos.9.2.3.bin", } Note: The version nxos.9.2.3.bin is used by default. If different version is to be used, modify the "target_system_image" with new version.</password></username> |
| 11. | 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 |
| 12. | 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. |



Table 2-5 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches

| I | |
|---|---|
| Step # Procedure | Description |
| 13. 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/ftpboot/conf. <switcha number="" serial=""> file, including all the values in the curly braces as following code block. These values are contained at Installation PreFlight checklist: ToR and Enclosure Switches Variables Table (Switch Specific) and Installation PreFlight Checklist: Complete OA and Switch IP Table. Modify these values with the following sed commands, or use an editor such as vi etc. \$ sed -i 's/{switchname}/<switch_name>/' conf.<switcha number="" serial=""> \$ sed -i 's/{admin_password}/<admin_password>/' conf.<switcha number="" serial=""> \$ sed -i 's/{auser_name}/<suser_name>/' conf.<switcha number="" serial=""> \$ sed -i 's/{suser_password}/<suser_password>/' conf.<switcha number="" serial=""> \$ sed -i 's/{ospf_md5_key}/<cospf_md5_key>/' conf.<switcha number="" serial=""> \$ sed -i 's/{ospf_md5_key}/<cospf_area_id>/' conf.<switcha number="" serial=""> \$ sed -i 's/{Sopf_aREA_ID}/<sopf_area_id>/' conf.<switcha number="" serial=""> \$ sed -i 's/{NTPSERVER1}/<ntp_server_1>/' conf.<switcha number="" serial=""> \$ sed -i 's/{NTPSERVER2}/<ntp_server_3>/' conf.<switcha number="" serial=""> \$ sed -i 's/{NTPSERVER3}/<ntp_server_3>/' conf.<switcha number="" serial=""> \$ sed -i 's/{NTPSERVER3}/<ntp_server_4>/' conf.<switcha number="" serial=""> \$ sed -i 's/{NTPSERVER4}/<ntp_server_4>/' conf.<switcha number="" serial=""> Note: If less than 5 ntp servers available, delete the extra ntp server lines such as command: \$ sed -i 's/{NTPSERVER5}/d' conf.<switcha number="" serial=""> Note: different delimiter is used in next two commands due to '/' sign in the variables \$ sed -i 's/{NTPSERVER5}/d' conf.<switcha number="" serial=""> \$ sed -i 's/{NTPSERVER5}/d' conf.<switcha number="" serial=""> \$ sed -i 's#{CNE_Management_SwA_Address}#<torswitcha_cnemanagementnet_ips#g' conf.<switcha="" number="" serial=""> \$ sed -i 's#{CNE_Management_Prefix}#<cnemanagementnet_ips#g' conf.<switcha="" number="" serial=""> \$ sed -i 's#{SUL_replication_SwA_Address}#<torswitcha_sqlreplication_net_ips#g' conf.<switcha="" number="" serial=""> \$ sed -i 's#{SUL_replication_SwA_Address}#<torswitchb_sqlreplication_n< td=""></torswitchb_sqlreplication_n<></torswitcha_sqlreplication_net_ips#g'></cnemanagementnet_ips#g'></torswitcha_cnemanagementnet_ips#g'></switcha></switcha></switcha></switcha></ntp_server_4></switcha></ntp_server_4></switcha></ntp_server_3></switcha></ntp_server_3></switcha></ntp_server_1></switcha></sopf_area_id></switcha></cospf_area_id></switcha></cospf_md5_key></switcha></suser_password></switcha></suser_name></switcha></admin_password></switcha></switch_name></switcha> |



Table 2-5 (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 2-5 (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 2-5 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches

| Step # Proced | e Description |
|---|---|
| 14. Copy the /var tftpbood 93180_ B.cfg in file called /v /tftpboo conf. <s b="" number<="" serial="" td=""><td>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: Complete OA and Switch IP Table.</switcha></td></s> | 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: Complete OA and Switch IP Table.</switcha> |



Table 2-5 (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 2-5 (Cont.) Procedure to configure Top of Rack 93180YC-EX Switches

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

Verification



Table 2-6 Procedure to verify Top of Rack 93180YC-EX Switches

| Step # Procedure | Description |
|--|---|
| 1. After the ToR switches configured, ping the switches from bootstrap server. The switches mgmt0 interfaces are configured with the IP addresses which are in the conf files. | Note: Wait till the device responds. \$ ping 192.168.2.1 PING 192.168.2.1 (192.168.2.1) 56(84) bytes of data. 64 bytes from 192.168.2.1: icmp_seq=1 ttl=255 time=0.419 ms 64 bytes from 192.168.2.1: icmp_seq=2 ttl=255 time=0.496 ms 64 bytes from 192.168.2.1: icmp_seq=3 ttl=255 time=0.573 ms 64 bytes from 192.168.2.1: icmp_seq=4 ttl=255 time=0.573 ms 64 bytes from 192.168.2.1: icmp_seq=4 ttl=255 time=0.535 ms ^C 192.168.2.1 ping statistics 4 packets transmitted, 4 received, 0% packet loss, time 3000ms rtt min/avg/max/mdev = 0.419/0.505/0.573/0.063 ms \$ ping 192.168.2.2 PING 192.168.2.2 (192.168.2.2) 56(84) bytes of data. 64 bytes from 192.168.2.2: icmp_seq=1 ttl=255 time=0.572 ms 64 bytes from 192.168.2.2: icmp_seq=2 ttl=255 time=0.582 ms 64 bytes from 192.168.2.2: icmp_seq=3 ttl=255 time=0.584 ms 64 bytes from 192.168.2.2: icmp_seq=4 ttl=255 time=0.554 ms 67 192.168.2.2 ping statistics 4 packets transmitted, 4 received, 0% packet loss, time 3001ms rtt min/avg/max/mdev = 0.466/0.543/0.582/0.051 ms |



Table 2-6 (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.gnu.org/licenses/old-licenses/library.txt. # |



Table 2-6 (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 m4route-mem minimum 58 maximum 58 limit-resource m6route-mem minimum 8 maximum 8 feature scp-server feature sftp-server cfs eth distribute feature ospf feature bgp feature bgp feature lacp feature lacp feature bfd feature vrrpv3 |
| 4. | Un-mount the Utility USB and remove it: umount / media/usb | Connect or enable customer uplink. |
| 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> |



Table 2-6 (Cont.) Procedure to verify Top of Rack 93180YC-EX Switches

| Step # | Procedure | Description |
|--------|---|---|
| 6. | 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> |

SNMPv2c configuration

Table 2-7 SNMPv2c configuration for Top of Rack 93180YC-EX Switches

| Step # | Procedure | Description |
|--------|--|--|
| 1. | 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. [root@RMS1 ~]# ssh <user_name>@<torswitcha_cnemanagementnet_ip> # configure terminal (config)# snmp-server host <snmp_trap_receiver_address> traps version 2c <snmp_community_string> (config)# snmp-server host <snmp_trap_receiver_address> use-vrf default (config)# snmp-server host <snmp_trap_receiver_address> source-interface Ethernet1/51 (config)# snmp-server enable traps (config)# snmp-server community <snmp_community_string> group network-admin (config)# copy running-config startup-config</snmp_community_string></snmp_trap_receiver_address></snmp_trap_receiver_address></snmp_community_string></snmp_trap_receiver_address></torswitcha_cnemanagementnet_ip></user_name></snmp_community_string></snmp_trap_receiver_address> |



Table 2-7 (Cont.) SNMPv2c configuration for Top of Rack 93180YC-EX Switches

| 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</pre> |



Table 2-7 (Cont.) SNMPv2c configuration for Top of Rack 93180YC-EX Switches

| Step # | Procedure | Description |
|--------|----------------------------------|---|
| 3. | Traffic egress out of cluster | Traffic egress out of cluster, including snmptrap traffic to snmp trap receiver, and traffic goes to signal server: |
| | | <pre>[root@RMS1 ~]# ssh</pre> |
| | | overload add-route ip nat inside source list host-snmptrap interface Ethernet1/51 overload |
| | | interface Vlan3 ip nat inside |
| | | interface Vlan100 ip nat outside |
| | | interface Ethernet1/51 ip nat outside |
| | | interface Ethernet1/52 ip nat outside |
| | | (config)# copy running-config startup-config |
| | | 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 2-8 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/dhcp.conf file should already have been configured in procedure OCCNE 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 2-8 (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 22:07:26; 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 5 2019/03/29 11:14:04; ends 5 2019/03/29 11:14:04; 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; 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; cht 5 2019/03/29 15:56:23; cltt 5 2019/03/29 15:56:23; binding state free; hardware ethernet b8:83:03:47:5f:14; uid "\000\270\203\003G^T\000\000\000\000"; } lease 192.168.20.104 { starts 5 2019/03/29 12:56:23; binding state free; hardware ethernet b8:83:03:47:5e:54; uid "\000\270\203\003G^T\000\000\000\000"; } lease 192.168.20.104 { starts 5 2019/03/29 13:08:21; cht 5 2019/03/29 16:08:21; clt 5 2019/03/29 16:08:21; clt 5 2019/03/29 16:08:21; clt 5 2019/03/29 13:08:21; binding state free; hardware ethernet b8:83:03:47:5e:54; uid "\000\270\203\003Gd\234\000\000\000\000"; |



Table 2-8 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

| | 1 | |
|--------|-----------|--|
| Step # | Procedure | Description |
| | | } |
| | | lease 192.168.20.108 { |
| | | starts 5 2019/03/29 09:57:02; |
| | | ends 5 2019/03/29 21:57:02; |
| | | tstp 5 2019/03/29 21:57:02; |
| | | cltt 5 2019/03/29 09:57:02; |
| | | binding state active; |
| | | next binding state free; |
| | | rewind binding state free; |
| | | hardware ethernet fc:15:b4:la:ea:05; |
| | | uid "\001\374\025\264\032\352\005"; |
| | | client-hostname "OA-FC15B41AEA05"; |
| | | } |
| | | lease 192.168.20.107 { |
| | | starts 5 2019/03/29 12:02:50; |
| | | ends 6 2019/03/30 00:02:50; |
| | | tstp 6 2019/03/30 00:02:50; |
| | | cltt 5 2019/03/29 12:02:50; |
| | | binding state active; |
| | | next binding state free; |
| | | rewind binding state free; |
| | | hardware ethernet 9c:b6:54:80:d7:d7; |
| | | uid "\001\234\266T\200\327\327"; |
| | | client-hostname "SA-9CB65480D7D7"; |
| | | } |
| | | server-duid "\000\001\000\001\$# |
| | | \364\344\270\203\003Gim"; |
| | | lease 192.168.20.107 { |
| | | starts 5 2019/03/29 18:09:47; ends 6 2019/03/30 06:09:47; |
| | | ends 6 2019/03/30 06:09:4/; cltt 5 2019/03/29 18:09:47; |
| | | binding state active; |
| | | next binding state free; |
| | | rewind binding state free; |
| | | hardware ethernet 9c:b6:54:80:d7:d7; |
| | | uid "\001\234\266T\200\327\327"; |
| | | client-hostname "SA-9CB65480D7D7"; |
| | | [} |
| | | lease 192.168.20.108 { |
| | | starts 5 2019/03/29 18:09:54; |
| | | ends 6 2019/03/30 06:09:54; |
| | | cltt 5 2019/03/29 18:09:54; |
| | | binding state active; |
| | | next binding state free; |
| | | rewind binding state free; |
| | | hardware ethernet fc:15:b4:1a:ea:05; |
| | | uid "\001\374\025\264\032\352\005"; |
| | | client-hostname "OA-FC15B41AEA05"; |
| | | [} |
| | | lease 192.168.20.106 { |
| | | starts 5 2019/03/29 18:10:04; |
| | | ends 5 2019/03/29 21:10:04; |
| | | cltt 5 2019/03/29 18:10:04; |
| | | binding state active; |
| | | |



Table 2-8 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

| Step # | Procedure | Description |
|--------|---|---|
| | | <pre>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 18:10:35; cltt 5 2019/03/29 18:10:35; binding state active; next binding state free; rewind binding state free; hardware ethernet b8:83:03:47:64:9c; uid "\000\270\203\003Gd\234\000\000\000"; client-hostname "ILO2M2909004F"; } lease 192.168.20.105 { starts 5 2019/03/29 18:10:40; ends 5 2019/03/29 21:10:40; cltt 5 2019/03/29 18:10:40; binding state active; next binding state free; rewind binding state free; rewind binding state free; hardware ethernet b8:83:03:47:5e:54; uid "\000\270\203\003G^T\000\000\000"; client-hostname "ILO2M29090048";</pre> |
| 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,oemHPE_rc,oemHPE_power,oemHPE_vm status=0 status_tag=COMMAND COMPLETED Tue Apr 2 20:08:30 2019 User added successfully.</pre> |



Table 2-8 (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=NO status=0 status_tag=COMMAND COMPLETED Tue Apr 2 20:04:53 2019 Network settings change applied. Settings change applied, iLO 5 will now be reset. Logged Out: It may take several minutes before you can log back in. CLI session stopped packet_write_wait: Connection to 192.168.20.104 port 22: Broken pipe</pre> |
| 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 hpiLO-> set /map1/enetport1/lanendpt1/ipendpt1 IPv4Address=192.168.20.122 SubnetMask=255.255.255.0 status=0 status_tag=COMMAND COMPLETED Tue Apr 2 20:22:23 2019 Network settings change applied. Settings change applied, iLO 5 will now be reset. Logged Out: It may take several minutes before you can log back in. CLI session stopped packet_write_wait: Connection to 192.168.20.104 port 22: Broken pipe #</new></new></new></pre> |



Table 2-8 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

| Step# | Procedure | Description |
|-------|---|---|
| 8. | Procedure Set EBIPA addresses for InterConnect Bays (Enclosure Switches). After login to OA, set EBIPA addressed for the two enclosure switches. The addresses have to be in the subnet with server team0.2 address in order for TFTP to work. | Description Set address for each enclosure switch, note the last number 1 or 2 is the interconnect bay number. OA-FC15B41AEAO5> 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-FC15B41AEAO5> 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 2-8 (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-FC15B41AEAO5> 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 # 5 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 # 8 to IP address 192.168.20.148 Successfully set device (iLO) bay # 10 to IP address 192.168.20.149 Successfully set device (iLO) bay # 10 to IP address 192.168.20.150 Successfully set device (iLO) bay #10 to IP address 192.168.20.151 Successfully set device (iLO) bay #12 to IP address 192.168.20.152 Successfully set device (iLO) bay #12 to IP address 192.168.20.152 Successfully set device (iLO) bay #12 to IP address 192.168.20.152 Successfully set device (iLO) bay #14 to IP address 192.168.20.153 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 192.168.20.155 Successfully set device (iLO) bay #16 to IP address 192.168.20.155 Successfully set device (iLO) bay #16 to IP address 192.168.20.155 Successfully set device (iLO) bay #16 to IP address 192.168.20.155 Successfully set device (iLO) bay #16 to IP address 192.168.20.155 |



Table 2-8 (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 and OAs. After that, the username and password can be used to access OAs. | OA-FC15B41AEA05> ADD USER <username> New Password: ******* Confirm : ******** User "<username>" created. You may set user privileges with the 'SET USER ACCESS' and 'ASSIGN' commands. OA-FC15B41AEA05> set user access <username> ADMINISTRATOR "<username>" has been given administrator level privileges. OA-FC15B41AEA05> ASSIGN SERVER ALL <username> <username> has been granted access to the valid requested bay(s) OA-FC15B41AEA05> ASSIGN OA <username> <username> has been granted access to the OA.</username></username></username></username></username></username></username></username> |
| 11. | From OA, go to each blade with "connect server server >my number>", add New User for each blade. | OA-FC15B41AEA05> connect server 4 Connecting to bay 4 User:OAtmp-root-5CBF2E61 logged-in to ILO2M290605KP.(192.168.20.144 / FE80::AF1:EAFF:FE89:460) iLO Standard Blade Edition 1.37 at Oct 25 2018 Server Name: Server Power: On hpiLO-> hpiLO-> hpiLO-> create /map1/accounts1 username=root password=TklcRoot group=admin,config,oemHPE_rc,oemHPE_power,oemHPE_vm status=2 status_tag=COMMAND PROCESSING FAILED error_tag=COMMAND SYNTAX ERROR Tue Apr 23 16:18:58 2019 User added successfully. |



Table 2-8 (Cont.) Procedure to configure Addresses for RMS iLOs, OA, EBIPA

| Step # | Procedure | Description |
|--------|---|---|
| 12. | Change to static IP on OA. In order not reply on DHCP and | Note : After the following change, 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. |
| | make the OA address stable, change to static IP. | OA-FC15B41AEA05> SET IPCONFIG STATIC 1 192.168.20.131 255.255.255.0 Static IP settings successfully updated. These setting changes will take effect immediately. |
| | | OA-FC15B41AEA05> SET IPCONFIG STATIC 2 192.168.20.132 255.255.255.0 Static IP settings successfully updated. These setting changes will take effect immediately. OA-FC15B41AEA05> |

Configure Legacy BIOS on Remaining Hosts

These procedures define the steps necessary to configure additional Legacy BIOS for all hosts in OCCNE 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.



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 2-9.
- 3. The procedures listed here apply to both Gen10 DL380 RMSs and Gen10 BL460c Blades in a C7000 enclosure.
- 4. Access to the enclosure blades in these procedures is via the Bootstrap host using SSH on the KVM. This is possible because the prerequisites are complete. If the prerequisites are not completed before executing this procedure, the enclosure blades are only accessible via the KVM connected directly to the active OA. In this case the mouse is not usable and screen manipulations are performed using the keyboard ESC and directional keys.
- 5. This procedure does NOT apply to the Bootstrap Host.



References

- 1. HPE iLO 5 User Guide 1.15
- 2. UEFI System Utilities User Guide for HPE ProLiant Gen10 Servers and HPE Synergy
- 3. UEFI Workload-based Performance and Tuning Guide for HPE ProLiant Gen10 Servers and HPE Synergy
- 4. HPE BladeSystem Onboard Administrator User Guide
- 5. OCCNE Inventory File Preparation

Steps to configure the Legacy BIOS on Remaining Hosts

Table 2-9 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 2-9 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

| Step # | Procedure | Des | cription |
|--------|---|-----|--|
| 2. | Expose the System Utility for an Enclosure Blade | 1. | The blades are maintained via the OAs in the enclosure. Because each blade iLO has already been assigned an IP address from the prerequisites, the blades can each be reached using SSH from the Bootstrap host login shell on the KVM. |
| | | | 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@ <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</blade_ilo_ip_address> |
| | | | 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 2-9 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

| Step # | Procedure | Des | cription |
|--------|---------------------------------------|-----|--|
| 3. | Change over from UEFI Booting Mode to | 1. | Expose the System Utility by following step 1 or 2 depending on the hardware being configured. |
| | Legacy BIOS | 2. | Select System Configuration |
| | Booting Mode | 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 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 |
| | | 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 2-9 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

| Step # | Procedure | Description |
|--------|--|---|
| 4. | Force PXE to boot from the first Embedded FlexibleLOM | Expose the System Utility by following step 1 or 2 depending on the hardware being configured. |
| | | 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 2-9 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

| Step # | Procedure | Des | cription |
|--------|--------------------------------|---|--|
| 6. | Disable RAID Configurations | proc | CNE does not currently support any RAID configuration. Follow this redure to disable RAID settings if the default settings of the System ity 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. | |
| | | 1. | Expose the System Utility by following step 1 or 2 depending on the hardware being configured. |
| | | 2. | Select System Configuration. |
| | | 3. | Select Embedded RAID 1 : HPE Smart Array P408i-a SR Gen 10. |
| | | 4. | Select Array Configuration. |
| | | 5. | Select Manage Arrays. |
| | | 6. | Select Array A (or any designated Array Configuration if there are more than one). |
| | | 7. | Select Delete Array . A warning is displayed indicating the following: |
| | | | Deletes an Array. All the data on the logical drives that are part of deleted array will be lost. Also if the deleted array is the only one on the controller, the controller settings will be erased and its default configuration is restored. |
| | | 8. | Hit ENTER , the changes are submitted and Delete Array Successful is displayed. |
| | | 9. | Hit ENTER to go back to the main menu for the HPE Smart Array. |
| | | 10. | Select F10 if it is desired to save and stay in the utility or select the F12 it is desired to save and exit to continue the current boot process. |



Table 2-9 (Cont.) Procedure to configure the Legacy BIOS on Remaining Hosts

| Step # | Procedure | Description |
|--------|---|---|
| 7. | Enable the Primary and Secondary Boot Devices | This steps provide necessary to configure the primary and secondary bootable devices for a Gen10 Server. |
| | | Note: There can be multiple configurations of hardware drives on the server that include both Hard Drives (HDD) and Solid State Hard Drives (SSD). SSDs are indicated by SATA-SSD ATA in the drive description. The commands below include two HDDs and two SSDs. The SSDs are not to be selected for this configuration. The actual selections may be different based on the hardware being updated. |
| | | Expose the System Utility by following step 1 or 2 depending on the hardware being configured. |
| | | 2. Select System Configuration. |
| | | 3. Select Embedded RAID 1 : HPE Smart Array P408i-a SR Gen 10. |
| | | 4. Select Set Bootable Device(s) for Legacy Boot Mode. If the boot devices are not set then Not Set is displayed for the primary and secondary devices. |
| | | 5. Examine the list of available hardware drives. If one or more HDDs are available, continue with this procedure. Note: A single drive can be set as both the primary and secondary boot device but that is not part of this configuration. |
| | | 6. Select Bootable Physical Drive |
| | | 7. Select Port 1 Box:3 Bay:1 Size:1.8 TB SAS HP EG00100JWJNR. Note: This example includes two HDDs and two SSDs. The actual configuration may be different. |
| | | 8. Select Set as Primary Bootable Device. |
| | | 9. Hit ENTER. Note: There is no need to set the secondary boot device. Leave it as Not Set. |
| | | 10. Hit the ESC key to back out to the System 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 Configure Top of Rack 93180YC-EX Switches has been completed.
- Procedure Configure Addresses for RMS iLOs, OA, EBIPA has been completed.
- The Utility USB is available containing the necessary files as per: Installation PreFlight checklist: Create Utility USB.

Limitations/Expectations



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

References

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

Procedure

Table 2-10 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}/<suser_password>/' 6127xlg_irf.cfg</suser_password></user_name></admin_password></switch_name> |



Table 2-10 (Cont.) Procedure to configure enclosure switches

| Step # | Procedure | Description |
|--------|--------------------------------------|--|
| 3. | 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>ssystem-view System View: return to User View with Ctrl+Z.</hpe> |
| | | <pre>(Note: Run the following commands:) irf member 1 priority 32 interface range Ten-GigabitEthernet 1/0/17 to Ten-GigabitEthernet 1/0/20 shutdown quit irf-port 1/1 port group interface Ten-GigabitEthernet1/0/17 port group interface Ten-GigabitEthernet1/0/18 port group interface Ten-GigabitEthernet1/0/19 port group interface Ten-GigabitEthernet1/0/20 quit interface range Ten-GigabitEthernet 1/0/17 to Ten-GigabitEthernet 1/0/20</pre> |



Table 2-10 (Cont.) Procedure to configure enclosure switches

| Step # | Procedure | Description |
|--------|-------------------------|---|
| | | undo shutdown |
| | | quit |
| | | save |
| | | irf-port-configuration active |
| | | |
| 4. | Access the InterConnect | Access the InterConnect Bay2 6127XLG switch to re-number to IRF 2. |
| | Bay2 6127XLG | OA-FC15B41AEA05> connect interconnect 2 |
| | | |
| | | <hpe><hpe>>system-view</hpe></hpe> |
| | | 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 2-10 (Cont.) Procedure to configure enclosure switches

| Step # | Procedure | Description |
|--------|---|--|
| 5. | Configure the IRF on Bay2 6127XLG | After rebooting, the interfaces will begin with number 2 such as Ten-GigabitEthernet2/0/17, Ten-GigabitEthernet2/1/5. Run the following commands: |
| | switch | 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 2-10 (Cont.) Procedure to configure enclosure switches

| Step # | Procedure | Description | | |
|--------|---|-------------------------|---|---------|
| 7. | Verify the IRF for the 6127XLG switches. | When reboot is finished | 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 | 0/18 | Ten- |
| | | GigabitEthernet2/0 |)/19 | Ten- |
| | | GigabitEthernet2/0 | 0/20 | Ten- |
| | | [HPE] | | |



 Table 2-10 (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 db-2/RMS2 with an operating system and configure it to fulfill the role of Database Host. After the Bastion Host is created, it is used to complete the installation of OCCNE.

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

Table 2-11 Terminology used in Procedure

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



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



Table 2-12 Bastion Installation

| Step # | Procedure | Des | cription |
|--------|--|-----|--|
| 1. | Copy the Necessary Files from the Utility USB to Support the OS | 1. | Login to the Bootstrap Host using the root credentials configured during OS installation of the Bootstrap Host. |
| | | 2. | Create the directories needed on the Installer Bootstrap Host. |
| | Install | | <pre>\$ mkdir -p /var/occne/cluster/ <cluster_short_name>/yum.repos.d</cluster_short_name></pre> |
| | | 3. | Mount the Utility USB. Note: Follow the instructions for mounting a USB in Linux are at: Installation of Oracle Linux 7.x on Bootstrap Host. |
| | | 4. | Copy the hosts.ini file (created using procedure: OCCNE Inventory File Preparation) into the /var/occne/cluster/ <cluster_short_name>/ directory. This hosts.ini file defines the Bastion KVM Host to the Provision Container running the provision image downloaded from the repo.</cluster_short_name> |
| | | | <pre>\$ cp /<path_to_usb>/hosts.ini /var/occne/ cluster/<cluster_short_name>/hosts.ini</cluster_short_name></path_to_usb></pre> |
| | | | <pre>Example: \$ cp /media/usb/hosts.ini /var/occne/cluster/ rainbow/hosts.ini</pre> |
| | | 5. | Edit the /var/occne/cluster/ <cluster_short_name>/hosts.ini file to include the ToR host_net (vlan3) VIP for NTP clock synchronization. Use the ToR VIP address (ToRswitch_Platform_VIP) as defined in procedure: Installation PreFlight Checklist: Complete OA and Switch IP SwitchTable as the NTP source. Update the ntp_server field with the VIP address. Update the occne_repo_host_address to point to this Bootstrap Host internal address. This is being used for PXE booting and accessing the NFS share on the Installer Bootstrap Host (db-1/RMS1).</cluster_short_name> |
| | | | Example (from hosts.sample.ini): |
| | | | ntp_server='172.16.3.1' occne_repo_host_address='172.16.3.4' |
| | | 6. | Copy the customer specific .repo file from the Utility USB to the Installer Bootstrap Host. |
| | | | This is the .repo file created by the customer that provides access to the onsite (within their network) yum repositories needed to complete the full deployment of OCCNE 1.3 onto the Installer Bootstrap Host. It needs to be in two places, one for the local system, and one for the target systems. |
| | | | <pre>\$ cp /<path_to_usb>/</path_to_usb></pre> |



Table 2-12 (Cont.) Bastion Installation

| Step # | Procedure | Description |
|--------|-----------|---|
| | | <pre>Example: \$ cp /media/usb/ol7-mirror.repo /var/occne/ cluster/rainbow/yum.repos.d/ \$ cp -r /var/occne/cluster/rainbow/ yum.repos.d /var/occne/ \$ echo "reposdir=/var/occne/yum.repos.d" >> /etc/yum.conf</pre> |



Table 2-12 (Cont.) Bastion Installation

| Step # | Procedure | Description |
|--------|--|---|
| 2. | Set up the /etc/hosts file for the Central Repo and Verify Access | 1. Add an entry to the /etc/hosts file on the Installer Bootstrap Host to provide a name mapping for the Customer central repository. |
| | 7100035 | \$ vi /etc/hosts |
| | | Example: 10.75.200.217 rainbow-reg |
| | | To Verify: \$ ping <central_repo_name></central_repo_name> |
| | | Example: # ping rainbow-reg PING reg-1 (10.75.200.217) 56(84) bytes of data. 64 bytes from reg-1 (10.75.200.217): icmp_seq=1 ttl=61 time=0.248 ms 64 bytes from reg-1 (10.75.200.217): icmp_seq=2 ttl=61 time=0.221 ms 64 bytes from reg-1 (10.75.200.217): icmp_seq=3 ttl=61 time=0.239 ms |
| | | 2. Verify the repo access execute the following command: |
| | | \$ yum repolist |
| | | Example: \$ yum repolist Loaded plugins: ulninfo repo id repo name status !UEKR5/x86_64 Unbreakable Enterprise Kernel Release 5 for Oracle Linux 7 (x86_64) 80 !addons/x86_64 Oracle Linux 7 Addons |
| | | 91 !developer/x86_64 Packages for creating test and development environments for Oracle Linux 7 226 !developer_EPEL/x86_64 Packages for creating test and development environments for Oracle Linux 7 13,246 !ksplice/x86_64 Ksplice for Oracle Linux 7 (x86_64) |
| | | 393 !latest/x86_64 Oracle Linux 7 Latest (x86_64) |



Table 2-12 (Cont.) Bastion Installation

| Step # | Procedure | Description |
|--------|---|---|
| 3. | Copy the OL7 ISO to the Installer Bootstrap Host | The iso file must be accessible from a Customer Site Specific repository. This file should be accessible because the ToR switch configurations were completed in procedure: Configure Top of Rack 93180YC-EX Switches |
| | | Copy the OL7 ISO file to the /var/occne directory. The example below uses OracleLinux-7.5-x86_64-disc1.iso. If this file was copied to the Utility USB, it can be copied from there into the same directory on the Bootstrap Host. |
| | | Note : If the user copies this ISO from their laptop then they must use an application like WinSCP pointing to the Management Interface IP. |
| | | <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> |
| 4. | Install Packages onto the Installer Bootstrap Host | Use YUM to install necessary packages onto the installer Bootstrap Host. |
| | | \$ yum install docker-engine nfs_utils ansible |
| 5. | Set up access to the Docker Registry on the Installer Bootstrap Host | <pre>1. Copy the docker registry certificate to two places on the Bootstrap Host. Note: How to obtain the docker registry certificate <source/> is not necessarily covered in the procedure. The user can use the instructions at reference 1 to understand this more thoroughly. \$ mkdir -p /var/occne/certificates \$ cp <source/>.crt /var/occne/certificates/ <occne_private_registry>:<occne_private_registry _port="">.crt \$ mkdir -p /etc/docker/certs.d/ <occne_private_registry>:<occne_private_registry _port=""> \$ cp <source/>.crt /etc/docker/certs.d/ <occne_private_registry>:<occne_private_registry _port=""> \$ cp <source/>.crt /etc/docker/certs.d/ <occne_private_registry>:<occne_private_registry _port="">/ca.crt</occne_private_registry></occne_private_registry></occne_private_registry></occne_private_registry></occne_private_registry></occne_private_registry></occne_private_registry></occne_private_registry></pre> |
| | | Example: \$ mkdir -p /var/occne/certificates \$ cp <source/> .crt /var/occne/certificates/ rainbow_reg:5000.crt \$ mkdir -p /etc/docker/certs.d/rainbow_reg:5000 \$ cp <source/> .crt /etc/docker/certs.d/ rainbow_reg:5000/ca.crt 2. Start the docker daemon. \$ systemctl daemon-reload \$ systemctl restart docker \$ systemctl enable docker Verify docker is running: \$ ps -elf grep docker \$ systemctl status docker |



Table 2-12 (Cont.) Bastion Installation

| Step # | Procedure | Description |
|--------|---|--|
| 6. | Setup NFS on the Installer Bootstrap Host | Run the following commands using sudo (assumes nfs-utils has already been installed in procedure: Installation of Oracle Linux 7.x on Bootstrap Host: Install Additional Packages). Note: The IP address used in the echo command is the Platform VLAN IP Address (VLAN 3)of the Bootstrap Host (RMS 1) as given in: Installation PreFlight Checklist: Site Survey Host Table. \$ echo '/var/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 \$ systemctl status nfs-server |
| 7. | 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-disc1.iso /mnt \$ cp /mnt/isolinux/initrd.img /var/occne/pxelinux \$ cp /mnt/isolinux/vmlinuz /var/occne/pxelinux |
| 8. | Verify and Set the PXE Configuration File Permissions on the Installer Bootstrap Host | Each file configured in the step above must be open for read and write permissions. \$ chmod -R 777 /var/occne/pxelinux |
| 9. | Disable DHCP and TFTP on the Installer Bootstrap Host | The TFTP and DHCP services running on the Installer Bootstrap Host may still be running. These services must be disabled. \$ systemctl stop dhcpd \$ systemctl disable dhcpd \$ systemctl stop tftp \$ systemctl disable tftp |



Table 2-12 (Cont.) Bastion Installation

| Step # | Procedure | Description |
|--------|--|---|
| 10. | Disable SELINUX | Set SELINUX to permissive mode. In order to successfully set the SELINUX mode, a reboot of the system is required. The getenforce command is used to determine the status of SELINUX. |
| | | <pre>\$ getenforce active</pre> |
| | | If the output of this command displays "active", change it to "permissive" by editing the /etc/ selinux/config file. |
| | | <pre>\$ vi /etc/selinux/config</pre> |
| | | Change the SELINUX variable to passive: SELINUX=permissive save the file |
| | | Reboot the system: reboot |
| 11. | Generate the SSH private and public keys on Bootstrap Host. | This command generates a private and public key for the cluster. These keys are passed to the Bastion Host and used to communicate to other nodes from that Bastion Host. The public key is passed to each node on OS install. Do not supply a passphrase when it asks for one. Just hit enter. |
| | | Note: The private key (occne_id_rsa) must be copied to a server that going to access the Bastion Host because the Bootstrap Host is repaved. This key is used later in the procedure to access the Bastion Host after it has been created. |
| | | Execute the following commands on the Bootstrap Host: |
| | | <pre>\$ mkdir -m 0700 /var/occne/cluster/ <cluster_short_name>/.ssh \$ ssh-keygen -b 4096 -t rsa -C "occne installer key" -f "/var/occne/cluster/ <cluster_short_name>/.ssh/occne_id_rsa" -q -N ""</cluster_short_name></cluster_short_name></pre> |



Table 2-12 (Cont.) Bastion Installation

| Step # | Procedure | Description | |
|--------|---|--|--|
| 12. | Execute the OS Install and Bastion VM Creation on Bastion KVM Host (RMS2) from the Installer Bootstrap Host | 1. Run the docker commands below to perform the OS install and Bastion Host VM creation on the Bastion KVM Host (RMS2): \$ docker runrmnetwork hostcap- add=NET_ADMIN -v /var/occne/cluster/ <cluster_short_name>/:/host -v /var/occne/:/var/ occne:rw -e "OCCNEARGS= limit= limit= limit= localhost" <image_name>:<image_tag></image_tag></image_name></cluster_short_name> | |
| | | <pre>\$ docker run -itrmnetwork hostcap- add=NET_ADMIN -v /var/occne/cluster/rainbow/:/ host -v /var/occne/:/var/occne:rw -e "OCCNEARGS= limit=bastion-2.rainbow.lab.us.oracle.com,db-2.r ainbow.lab.us.oracle.com,localhost" winterfell: 5000/occne/provision:1.3.0</pre> | |
| | | 2. Verify that Bastion Host VM is installed by logging into RMS2/db-2 and issuing the following command. The <ansible_host> field (which is an IPv4 address) is derived from the hosts.ini file db-2 entry for host_hp_gen_x groups. Note: This command is optional. Had a failure actually occurred, the docker run command would have experienced failures. ssh -i /var/occne/cluster/ <cluster_short_name>/.ssh/occne_id_rsa admusr@<oam_host> \$ sudo virsh list</oam_host></cluster_short_name></ansible_host> | |
| | | Example: ssh -i /var/occne/cluster/ rainbow.lab.us.oracle.com/.ssh/occne_id_rsa admusr@10.75.148.6 \$ sudo virsh list Id Name State 11 bastion-2.rainbow.lab.us.oracle.com running 3. Login to Bastion Host from the Boostrap Host as admusr using the generated key from the /var/occne/cluster/ <cluster_short_name> directory to confirm the VM is set up correctly. The <oam_host> field (which is an IPv4 address) is derived from the hosts.ini file bastion-2 entry for the host_kernel_virtual group. Note: This command is optional. Had a failure actually occurred, the docker run command would have experienced failures.</oam_host></cluster_short_name> | |



Table 2-12 (Cont.) Bastion Installation

| Step # | Procedure | Description |
|--------|-----------|--|
| | | <pre>ssh -i /var/occne/cluster/ <cluster_short_name>/.ssh/occne_id_rsa admusr@<oam_host></oam_host></cluster_short_name></pre> |
| | | Example: ssh -i /var/occne/cluster/rainbow/.ssh/ occne_id_rsa admusr@10.75.148.5 |

Automated Installation

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

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



Table 2-13 Automated Installation

| Step # | Procedure | Description | | | |
|--------|---|--|---|--|--|
| 1. | Setting up for and executing the deploy.sh script on the Bootstrap Host | The deploy.sh script performs the initial pre-configuration of the Bastion host. This includes installing ansible, executing the ansible playbook configBastionHost.yaml to setup the initial files and staging directories on the Bastion Host and executing the pipeline to setup the artifacts directory. The script is executed on the Bootstrap Host using a set of environment variables that can be initialized on the command line along with the deploy.sh script. These variables include the following: | | | |
| | | Name | Comment | Example usage | |
| | | CENTRAL_RE PO | Defines the customer specific repository host name. Note: This would be used in conjunction with CENTRAL_REPO_IP and CENTRAL_REPO_DOCK ER_PORT. | CENTRAL_REPO=custo mer_repo \ CENTRAL_REPO_IP=10. 10.10.10 \ CENTRAL_REPO_DOCK ER_PORT=5000 ./ deploy.sh | |
| | | CENTRAL_RE PO_IP | Defines the customer specific repository IPv4 address. | See above. | |
| | | CENTRAL_RE PO_DOCKER_ PORT | Defines the customer specific repository docker port number. Defaults to 5000. | See above. | |
| | | OCCNE_CLUS TER | Defines the cluster short name. | OCCNE_CLUSTER=rainb ow | |
| | | BASTION_HO ST | Bastion Host "ansible_host" IP address (ex: 172.16.3.100) | BASTION_HOST=172.16. 3.100 | |
| | | OCCNE_VERS ION | The version tag of the image releases | OCCNE_VERSION=1.3.2 | |



Table 2-13 (Cont.) Automated Installation

| Step # | Procedure | Des | cription |
|--------|---|-----|---|
| 2. | Copy necessary files from Utility USB to the Bootstrap Host staging directory | 1. | The MySQL .zip file (ex: V980756-01.zip) must be copied to the staging directory /var/occne directory. This file should be provided from the Utility USB. This file is used for installing the ndb MySQL nodes. |
| | | | <pre>\$ cp /<usb_dev>/db/*.zip /var/occne/*.zip</usb_dev></pre> |
| | | 2. | Copy the configBastionHost.yaml file from the Customer Utility USB to the staging directory /var/occne/. |
| | | | <pre>\$ cp /<usb_dev>/configBastionHost.yaml /var/ occne/.</usb_dev></pre> |
| | | 3. | Copy the deploy.sh script from the Customer Utility USB to the staging directory at /var/occne/ and set the file to executable. |
| | | | <pre>\$ cp /<usb_dev>/deploy.sh /var/occne/. \$ chmod +x /var/occne/deploy.sh</usb_dev></pre> |
| | | 4. | On the Bootstrap Host add the following entry to the /etc/hosts file: |
| | | | \$ sudo vi /etc/hosts |
| | | | Add the following: <bastion_ip_address> <bastion_full_name> Save/write the file</bastion_full_name></bastion_ip_address> |
| | | | Example entry: 172.16.3.100 bastion-2.rainbow.us.labs.oracle.com |



Table 2-13 (Cont.) Automated Installation

| Step# | Procedure | Description |
|-------|----------------|---|
| 3. | Execute Deploy | Execute the deploy.sh script from the /var/occne/ directory with the required parameters set. |
| | | <pre>\$ export CENTRAL_REPO=<customer_specific_repo_name> \$ export CENTRAL_REPO_IP=<customer_specific_repo_ipv4> \$ export OCCNE_CLUSTER=<cluster_short_name> \$ export OCCNE_BASTION= \$ export OCCNE_VERSION=1.3.x \$./deploy.sh</cluster_short_name></customer_specific_repo_ipv4></customer_specific_repo_name></pre> |
| | | Customer Example: \$ export CENTRAL_REPO=central-repo \$ export CENTRAL_REPO_IP=10.10.10.10 \$ export OCCNE_CLUSTER=rainbow \$ export OCCNE_BASTION=bastion-2.rainbow.us.labs.oracle.com \$ export OCCNE_VERSION=1.3.2 \$./deploy.sh |
| | | The command above can be executed like the following: \$ OCCNE_VERSION=1.3.2 CENTRAL_REPO=central-repo CENTRAL_REPO_IP=10.10.10.10 OCCNE_CLUSTER=rainbow OCCNE_BASTION=bastion-2.rainbow.us.labs.oracle.com ./ deploy.sh |
| | | Internal Example: Internally the defaults would be used so the only possible variables that needs to be provided are the OCCNE_CLUSTER, OCCNE_BASTION, and OCCNE_VERSION. |
| | | <pre>\$ OCCNE_VERSION=1.3.2 OCCNE_CLUSTER=rainbow OCCNE_BASTION=bastion-2.rainbow.us.labs.oracle.com ./ deploy.sh</pre> |



Table 2-13 (Cont.) Automated Installation

| Step # | Procedure | Description | |
|--------|--|--|--|
| 4. | Executing Final Deploy on Bastion Host | The following commands are executed from the Bastion Host to complete the Bastion Host configuration and deploy OCCNE on the Bare Metal system. | |
| | | Note : The Bootstrap Host cannot be used to access the Bastion Host as it will be re-paved from execution of this command. | |
| | | 1. Login to the Bastion Host as admusr. The private key that was saved earlier should be used to access the Bastion Host from a server other than the Bootstrap Host using the ssh command. This private key should be copied to the /home/ <user>/.ssh directory on that server as id_rsa using scp (or winSCP from a desktop PC). The permissions of the key must be set to 0600 using the command: chmod 0600 ~/.ssh/id_rsa</user> | |
| | | <pre>\$ ssh -i ~/.ssh/id_rsa admusr@<bastion_host_external_ip_address></bastion_host_external_ip_address></pre> | |
| | | 2. The MetalLB configuration file must be located on the Bastion Host at /var/occne/cluster/ <cluster_name>. Follow the procedure at Populate the MetalLB Configuration File to find an example of this file and instructions on how to configure it. You will have to manually edit (using vi) a file on the Bastion Host and copy the contents of the example MetalLB file (maintaining the formatting) into that file on the Bastion Host.</cluster_name> | |
| | | 3. Execute the following command to complete the deployment of OCCNE from the Bastion Host (excluding re-install on the Bastion Host and its KVM host, which are already setup). This action will re-pave the Bootstrap Host RMS. Note: You can use the following to pipe the output of the pipeline.sh script to a file: tee pipeline\$(date +"%F_%H%M%S").log | |
| | | <pre>\$ export PROV_DEPLOY_ARGS='limit=!</pre> | |
| | | Customer Example: \$ export PROV_DEPLOY_ARGS='limit=! bastion-2.rainbow.lab.us.oracle.com,! db-2.rainbow.lab.us.oracle.com' \$ export OCCNE_VERSION=1.3.2 \$ /var/occne/cluster/rainbow/artifacts/pipeline.sh | |
| | | To save the output from the pipeline.sh script the command can be written as: \$ /var/occne/cluster/rainbow/artifacts/ pipeline.sh tee pipeline\$(date +"%F_%H%M%S").log | |



Table 2-13 Automated Installation

| Step # | Procedure | Description | | |
|--------|---|---|--|--|
| 5. | Update the Bastion KVM Host repo file | Since db-2 was not part of the final OS install and cluster deploy, it's /var/occne/yum.repos.d/*.repo file is not pointing to the Bastion Host as its YUM repo. That file on RMS2/db-2 must be updated so that it now points to the Bastion Host as the repo. After the Bastion Host was created, the .repo file that was copied onto the Bastion Host has the correct settings. That file can just be copied back to RMS2/db-2. | | |
| | | 1. From the Bastion Host, login to the Bastion Host KVM Host, using the occne private key and the internal host IP address for that node (extracted from the hosts.ini file) | | |
| | | <pre>\$ ssh -i /var/occne/cluster/ <cluster_short_name>/.ssh/occne_id_rsa admusr@<bastion_kvm_host_ip_address></bastion_kvm_host_ip_address></cluster_short_name></pre> | | |
| | | 2. Remove the existing .repo files: | | |
| | | <pre>\$ sudo rm /var/occne/yum.repos.d/*.repo</pre> | | |
| | | 3. Copy the Bastion specific .repo file from the Bastion Host to the Bastion KVM Host. Execute this command from the Bastion KVM Host. | | |
| | | <pre>scp admusr@<bastion_ip_address>:/var/occne/ yum.repos.d/*.repo /var/occne/yum.repos.d/</bastion_ip_address></pre> | | |
| | | Example: scp admusr@172.16.3.100:/var/occne/yum.repos.d/ *.repo /var/occne/yum.repos.d/ | | |
| 6. | Change MySQL root user password | Refer to Change MySQL root user password | | |

Virtualized CNE Installation

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

Prerequisites

- The user has access to an existing OpenStack Environment including the OpenStack Desktop.
- 2. The OpenStack Environment is configured with appropriate resource flavors and network resources for resource allocation to the VMs created via this procedure.
- 3. Octavia Load Balancing plugin must be installed on the OpenStack Environment.
- 4. Users must have a pub key that can be configured for logging into the Bootstrap Host. This key should be placed into the customer OpenStack Environment prior to running this procedure using the following:
 - Use the **Import Keytab** on the Launch **Instance**→**Key Pair** dialog or via the **Compute**→**Access and Security**



Limitation/Expectations

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



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

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



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

Upload an Image to an OpenStack Environment

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

Note:

This procedure is executed from the OpenStack Desktop.

Table 2-14 Upload an Image to an OpenStack Environment

| Step # | Procedure | Description |
|--------|--------------------------------|--|
| 1. | Navigate to Images | Go to Compute → Images |
| 2. | Create Image | Select the +Create Image button. This brings up a new dialog. You have to enter a name for the image. Use the same name as was used to create and download the imag (ex: occne_bootstrap-1.3.0.qcow2). |
| 3. | Choose the source of the image | Using the Image Source pull down select: Image File . This will enable a Browse button. Select this button to bring up a Windows Explorer dialog. From the Windows dialog, select the image that was created in the previous procedures. This will insert the image name into the OpenStack Create Image dialog and set the Format for you. |



Table 2-14 (Cont.) Upload an Image to an OpenStack Environment

| Step # | Procedure | Description |
|--------|-----------------|---|
| 4. | Upload Image | Select the Create Image button at the bottom right of the dialog. This will start the image upload process. It will take a while so be patient. You will not be able to actually see the image being uploaded even if you log into another OpenStack instance. |
| 5. | Check the image | When the process is complete, the image should be listed in the Compute → Images screen. Again you will have to use Next to go through all the images and finally get to the image you uploaded depending on how many images are on the system. |

Bootstrap Host Creation

The Bootstrap Host is provisioned to drive the creation of the virtualized cluster using Terraform, the OpenStack Client, and Ansible Playbook(s). A qcow2 image was provided as part of the OSDC download and should be available on the users OpenStack Environment as per the previous section of this document.



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

Table 2-15 Bootstrap Host Creation

| Step # | Procedure |
|--------|---|
| 1. | Login to the OpenStack Environment using your OpenStack credentials, the appropriate domain and project name. |
| 2. | Select Compute → Instances |
| 3. | Select the Launch Instances tab on the upper right. A dialog will appear to configure a VM instance. |
| 4. | Enter an Instance Name (for example: occne- <name>). Leave the Availability Zone and Count set as is.</name> |
| 5. | Select Source on the left hand side of the dialog. A new dialog appears (Note : there might be a long list of available images to choose from) |
| 6. | Make sure the filter pulldown is set to Image |
| 7. | Enter occne_bootstrap in the filter. This will display the occne_bootstrap- <x.y.z>.qcow2 image uploaded in the previous sections of this procedure.</x.y.z> |
| 8. | Select the OCCNE Bootstrap Host image by selecting the "+" on the right side of the image listing. This adds the image as the source for this VM. |



Table 2-15 (Cont.) Bootstrap Host Creation

| Step # | Procedure | | | |
|--------|--|--|--|--|
| 9. | Select Flavor | | | |
| 10. | Enter a string which best describes the flavor being used for this customer specific OpenStack Environment in the search filter. This brings up a new dialog. | | | |
| 11. | Select the appropriate customer specific Flavor (for example: occne_bsh_flavor) by selecting the "+" on the right side of the flavor listings. This adds the resources to the Launch Instance dialog. Note : The BSH image requires a flavor that includes a disk size of 40GB or higher. The RAM size should be 8GB or higher although that is not a restriction. | | | |
| 12. | Select Networks | | | |
| 13. | Enter the appropriate network name as defined by the customer with the OpenStack Environment (example: ext-net) in the search filter. This brings up a new dialog. | | | |
| 14. | Select the appropriate network by selecting the "+" on the right side of the flavor listings. This adds the external network interface to the Launch Instance dialog. | | | |
| 15. | Select Key Pair . This dialog assumes you have already uploaded a public key to to OpenStack (see Prerequisites). | | | |
| 16. | Select the appropriate key by selecting the "+" on the right side of the key pair listings. This adds the public key to the authorized_keys file on the Bootstrap Host. | | | |
| 17. | Select Configuration . This screen allows the user to add configuration data which is used by cloud-init to set on the VM, the initial admusr and hostname/FQDN additions to the /etc/hosts file. Use the following configuration until there is more information available. This must be copied into the Customization Script text box. Make sure the fields marked as <instance_name_from_details_screen> are updated with the instance name you used in step 5 above. Leave the other fields on this dialog in their default setting.</instance_name_from_details_screen> | | | |
| | #cloud-config | | | |
| | <pre>hostname: <instance_name_from_details_screen> fqdn: <instance_name_from_details_screen> system_info:</instance_name_from_details_screen></instance_name_from_details_screen></pre> | | | |
| | <pre>default_user: name: admusr lock_passwd: false write_files: - content: </pre> | | | |
| | 127.0.0.1 localhost localhost4 localhost4.localdomain4 <instance_name_from_details_screen></instance_name_from_details_screen> | | | |
| | <pre>::1 localhost localhost6 localhost6.localdomain6 <instance_name_from_details_screen> path: /etc/hosts owner: root:root</instance_name_from_details_screen></pre> | | | |
| | permissions: '0644' | | | |
| 18. | Select Launch Instance at the lower right side of the initial dialog. This will launch the creation of the VM. This can be observed back at the Compute→Instances screen. | | | |

Pre-deployment Configuration

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

Table 2-16 Pre-deployment Configuration

| Step # | Procedure | Des | scription |
|--------|--|---|--|
| 1. | Obtain the TLS Certificate for OpenStack | Depending on the Customer's environment it is very likely that the customer's OpenStack uses certificates for TLS access to the API. Without this certificate, OpenStack commands will not work. Customer's may have to obtain this certificate before using OpenStack client commands. | |
| | | 1. | Ask the OpenStack admin to provide the required TLS certificate to access the client commands (ex. in an Oracle OpenStack system installed with kolla, the certificate will be located at /etc/kolla/certiifcates/haproxy-ca.crt) |
| | | 2. | Copy the certificate to the Bootstrap Host at location: /etc/pki/ <openstack_release_name>/haproxy-ca.crt (ex. /etc/pki/kolla/haproxy-ca.crt) (If /etc/pki/<openstack_release_name> does not exist, it can be created using command: mkdir -p /etc/pki/<openstack_release_name></openstack_release_name></openstack_release_name></openstack_release_name> |
| | | 3. | Set the environment variable OS_CACERT to the location where the certificate was copied to using the command: export OS_CACERT=/etc/pki/ <openstack_release_name>/haproxy-ca.crt (ex. export OS_CACERT=/etc/pki/kolla/haproxy-ca.crt)</openstack_release_name> |



Table 2-16 (Cont.) Pre-deployment Configuration

| Step # | Procedure | Description | | |
|--------|---------------------------------------|--|--|--|
| 2. | Get the Openstack RC (API v3) File | This file exports a number of environment variables on the Bootstrap Host for the given user which directs the OpenStack Client commands towards the particular OpenStack Environment. It must be copied to the users he directory on the Bootstrap Host so that the OpenStack Client commands can be executed. Note: These instructions may be somewhat different on | | |
| | | OpenStack Desktops. | | |
| | | From the OpenStack Desktop: go to Project → Compute → Access & Security. | | |
| | | 2. Select the API Access | | |
| | | On the right hand side, select Download OpenStack RC File v3. This will download a openrc.sh file prefixed with the OpenStack project name (ex: 5G-openrc.sh). to your PC. | | |
| | | SCP this file (ie. winSCP) to the Bootstrap Host in the / home/admusr directory as .<pre><pre><pre><pre><pre>popenrc.sh</pre></pre></pre></pre></pre> | | |
| | | Note: In order for SCP/winSCP to work properly, the key mentioned in the Prerequisites above must be used to access the Bootstrap Host. It may also be necessary to add the appropriate Security Group Rules to support SSH (Rule: SSH, Remote: CIDR CIDR: 0.0.0.0/0) under the Network → Security Groups page in the OpenStack Environment. Contact the OpenStack Administrator to get the proper rules added if necessary. | | |
| | | 5. Execute the following command: source . <pre>cproject_name</pre> -openrc.sh | | |
| | | 6. Execute the following command to verify the OpenStack Client is working: openstack image list | | |
| 3. | Create SSH Key on Bootstrap Host | Create the keys that will be used to access the other VMs. This command generates the private and public keys that are passed to the Bastion Host and used to communicate to other node from that Bastion Host. Do not supply a passphrase when it asks for one. Just hit enter. Also the private key should be copied to a place for safe keeping should the Bootstrap Host be destroyed. | | |
| | | \$ ssh-keygen -m PEM -t rsa -b 2048 | | |



Table 2-16 (Cont.) Pre-deployment Configuration

| Step # | Procedure | Description | | |
|--------|--------------------------------|--|---|--|
| 4. | Add Files to /tmp Directory | These files must be copied to the directories listed using scp or some other means (ie. winSCP). | | |
| | | 1. | There are three directories on the Bootstrap Host. These three directories are as follows: | |
| | | | a. /tmp/yum.repos.d | |
| | | | b. /tmp/db | |
| | | | c. /tmp/certificates | |
| | | 2. | Within these three directories the user must supply the following mandatory files: | |
| | | | a. Add a customer specific central repo .repo file to the /tmp/yum.repos.d directory which allows for access to the customer specific repo (ex: winterfell-mirror.repo). | |
| | | | b. Add a mysql .zip file. (ex: V980756-01.zip) to the /tmp/db directory. This file is used for installing the ndb MySQL cluster and is downloaded from OSDC. | |
| | | | c. Add a docker registry certificate to the /tmp/ certificates directory for the central docker registry. This file mus tbe copied using the following format: <central_repo_hostname>:<central_repo_p ort="">.crt(ex: winterfell:5000.crt).</central_repo_p></central_repo_hostname> | |



Table 2-16 (Cont.) Pre-deployment Configuration

| Step # | Procedure | Description |
|--------|---|--|
| 5. | Updating the ~.configure/ openstack/ clouds.yaml File | 1. To obtain the values for the authorization fields need in clouds.yaml below, execute the openstack configuration show command (make sure you have sourced the openrc script before executing this command). |
| | | <pre>\$ openstack configuration show</pre> |
| | | + |
| | | + |
| | | -+ |
| | | Field |
| | | Value |
| | | + |
| | | -+ |
| | | api_timeout |
| | | None |
| | | application_catalog_api_version 1 |
| | | auth.auth_url http:// |
| | | thundercloud.us.oracle.com:5000/v3 |
| | | auth.password |
| | | <pre><redacted></redacted></pre> |
| | | auth.project_domain_id 6c3468f1207c4e00bb441746c2046a90 |
| | | auth.project_id |
| | | 811ef89b5f154ab0847be2f7e41117c0 auth.project_name |
| | | OCCNE |
| | | auth.user_domain_name |
| | | LDAP |
| | | auth.username |
| | | John.Doe |
| | | auth_type |
| | | password baremetal_api_version |
| | | 1 |
| | | beta_command |
| | | False |
| | | cacert |
| | | None |
| | | cert |
| | | None compute_api_version |
| | | 3 combace_abt_serviou |
| | | container_api_version |
| | | container_infra_api_version |
| | | 1 |
| | | + |
| | | -+ |



Table 2-16 (Cont.) Pre-deployment Configuration

Step # Procedure Description

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

```
$ openstack network list
```

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

3. Edit the ~/.configure/openstack/clouds.yaml file (using sudo vi ~/.configure/openstack/clouds.yaml) and update following values for your Openstack Environment using the commands listed above (use double quotes where indicated in the example given below):

```
clouds:
  mycloud:
    auth:
      auth_url: <openstack-Identity-api-
url>
      username: <openstack-user-name>
      project_name: <openstack-project-</pre>
name>
      project_id: <openstack-Project ID>
      user_domain_name: <openstack-Domain
Name>
      password: <openstack-user-password>
    region_name: <openstack-region-name-if-
available>
    interface: <openstack-intrface- public/</pre>
private>
    identity_api_version: <openstack-
indentity-api-version>
    loadbalancer:
        lbaas_enabled: true
        subnet_id:
<openstack_lbaas_subnet_id>
        floating_network_id:
<openstack_lbaas_floating_network_id>
```



Table 2-16 (Cont.) Pre-deployment Configuration

| Step # | Procedure | Description |
|--------|-----------|--------------------------------------|
| | | use_octavia: true |
| | | lb_method: ROUND_ROBIN |
| | | Example: |
| | | clouds: |
| | | mycloud: |
| | | auth: |
| | | <pre>auth_url: http://</pre> |
| | | thundercloud.us.oracle.com:5000/v3 |
| | | username: "john.doe" |
| | | <pre>project_name: "OCCNE"</pre> |
| | | <pre>project_id:</pre> |
| | | 811ef89b5f154ab0847be2f7e41117c0 |
| | | user_domain_name: "LDAP" |
| | | password: "johnspw" |
| | | region_name: "RegionOne" |
| | | <pre>interface: "public"</pre> |
| | | <pre>identity_api_version: 3</pre> |
| | | loadbalancer: |
| | | lbaas_enabled: true |
| | | subnet_id: c0e0c185-ed65-4a53- |
| | | a7a3-418277fb9a20 |
| | | floating_network_id: |
| | | e4351e3e-81e3-4a83-bdc1-dde1296690e3 |
| | | use_octavia: true |
| | | |

lb_method: ROUND_ROBIN



Table 2-16 (Cont.) Pre-deployment Configuration

| Step # Procedure | | Description | | | |
|------------------|------------------------------|--|--|--|--|
| 6. | Updating cluster.tfvars File | The fields in the cluster.tfvars file must be configured to adapt to the current customer Openstack Environment. The steps below detail how to collect and set the fields that must be changed. | | | |
| | | From the /var/terraform directory, copy directory occne_example and its contents (cluster.tfvars) using the command below to create a new directory. Change the name of the new directory to include your name to distinguish it from the occne_example directory. Note: The directory name is not used for any special purpose other than to distinguish it from the occne_example directory. It can be called anything. | | | |
| | | <pre>\$ cp -R occne_example occne_<user></user></pre> | | | |
| | | 2. Use the following commands to retrieve the information necessary to configure the cluster.tfvars | | | |
| | | a. The different flavor settings should be set according to the recommendations from Reference 1 in this document. An Admin user of the custome specific OpenStack Environment must add the flavors and provide the UUID of those flavors for configuration into the cluster.tfvars file. The UUID of each specific flavor that is used must be added as the value field of the key/value fields in the cluster.tfvars file. | | | |
| | | b. Once flavors have been added to the OpenStack Environment, the UUID can be retrieved via the following OpenStack Client command from the Bootstrap Host shell: | | | |
| | | <pre>\$ openstack flavor list grep <flavor_name></flavor_name></pre> | | | |
| | | Example: | | | |
| | | <pre>\$ openstack flavor list grep medium 43c7b73b-42a7-4f40-b52e-11f6803fc750 oc-cne.medium</pre> | | | |
| | | <pre>\$ openstack flavor show 43c7b73b-42a7-4f40-b52e-11f6803fc750 +</pre> | | | |
| | | + | | | |
| | | + | | | |



Table 2-16 (Cont.) Pre-deployment Configuration

Step# **Procedure** Description False OS-FLV-EXT-DATA:ephemeral 0 access_project_ids None disk 80 | id 43c7b73b-42a7-4f40-b52e-11f6803fc750 name occne.medium os-flavor-access:is_public | True properties ram 16384 | rxtx_factor 1.0 swap vcpus Use the following command to retrieve the floatingip_pool name and external_net UUID. \$ openstack network list ID Name Subnets 1d25d5ea-77ca-4f56-b364-f53b09292e7b ext-net2 f5c5ee71-8688-466d-a79f-4306e2bf3f6a 668bc488-5307-49ad-9332-24fb0767bb39 test-network 9432b2d5-99c0-43ee-8f8c-4709f38b68d9 903155c7-c3ff-4283-bc2b-f34e8b6e76b0 occne-ebadger-tc1 | ecbadd3ee239-4830-b8c1-5ff94fa64c3a |



Table 2-16 (Cont.) Pre-deployment Configuration

Step # Procedure Description

Navigate to occne_<user> directory and edit the contents of the cluster.tfvars file in the newly created directory:

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

4. The fields in the cluster.tfvars file must be configured to adapt to the current customer Openstack Provider. Initially the cluster_name and network_name should be set as the same value.

```
# <Kubernetes cluster name here>
cluster_name = "<cluster-name>"
# networking
network_name = "<cluster-name>"
```

5. For setting the ntp_server value in the cluster.tfvars file, use the IP Address of your cloud URL. One way of obtaining this is using the ping command on your Bootstrap Host. (For example: ping thundercloud.us.oracle.com)

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



Table 2-16 (Cont.) Pre-deployment Configuration

| Step # | Procedure | Description | |
|--------|-----------------------------------|-------------|--|
| 7. | Updating the clouds.yaml file for | Not file | te: This must be executed after updating the cluster.tfvars |
| | Load Balancers | 1. | Get the subnet_id value by running following terrform commands: |
| | | | <pre>\$ terraform init \$ terraform apply -auto-approve -var- file=occne_<user>/cluster.tfvars</user></pre> |
| | | 2. | After successful run of terraform apply, run following OpenStack command to get the subnet_id. |
| | | | openstack network list |
| | | | ++ ++ |
| | | | ID Name Subnets |
| | | | + |
| | | | + |
| | | | 668bc488-5307-49ad-9332-24fb0767bb39 test-network |
| | | | 9432b2d5-99c0-43ee-8f8c-4709f38b68d9 903155c7-c3ff-4283-bc2b-f34e8b6e76b0 occne-ebadger-tc1 ecbadd3e-e239-4830-b8c1-5ff94fa64c3a |
| | | | 90c160aa-2ef7-47d3-a212-e1790d56c971 ext-net-ipv6 |
| | | | c4a7569b-5448-4add-8c4e-006bbdd984ef cluster1 |
| | | | e4351e3e-81e3-4a83-bdc1-dde1296690e3 ext-net |
| | | | fc36d63f-b30b-4c7f-979f-9b52b614bbd7 occne-mkingre |
| | | | 7631612f-5d22-49be-975c-6e0a9329339b + |
| | | | ++ |

- **3.** Copy the value from subnets column for the name that is same as the cluster_name from cluster.tfvars file.
- **4.** Edit the ~/.configure/openstack/clouds.yaml file (using sudo vi ~/.configure/openstack/clouds.yaml) and update the subnet_id value keeping other values previously configured the same.



Table 2-16 (Cont.) Pre-deployment Configuration

| Step # | Procedure | Description |
|--------|-----------|---|
| | | loadbalancer: |
| | | lbaas_enabled: true |
| | | <pre>subnet_id: <cluster_name_subnet_id></cluster_name_subnet_id></pre> |
| | | Example ((using cluster name: test- |
| | | network): |
| | | loadbalancer: |
| | | lbaas_enabled: true |
| | | subnet_id: |
| | | 9432b2d5-99c0-43ee-8f8c-4709f38b68d9 |
| | | |

Deploy the OCCNE Virtualized Cluster

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

Table 2-17 Deploy the OCCNE Virtualized Cluster

| Step # | Procedure | Description |
|--------|-------------------------------------|--|
| 1. | Deploy OCCNE Virtualized Cluster | Execute the following command from the /var/terraform directory on the Bootstrap Host. This command may take a while to run (can be up to 2 hours depending on the machines its run on). |
| | | <pre>\$ OCCNE_VERSION=1.3.x OCCNE_TFVARS_DIR=occne_<user> CENTRAL_REPO=<central_repo_hostname> CENTRAL_REPO_IP=<central_repo_ipv4_address> ./ deploy.sh</central_repo_ipv4_address></central_repo_hostname></user></pre> |
| | | Example: \$ OCCNE_VERSION=1.3.2 OCCNE_TFVARS_DIR=occne_jdoe CENTRAL_REPO=winterfell CENTRAL_REPO_IP=10.75.216.10 ./deploy.sh |
| 2. | Change MySQL root user password | Refer to Change MySQL root user password |



Post Installation Activities

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

Post Install Verification

Introduction

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

Prerequisities

- 1. Common services has been installed on all nodes hosting the cluster.
- 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 3-1 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_IP: \$KIBANA_LOADBALANCER_PORT http://10.75.182.51:80 Launch the Browser and navigate to http:// \$KIBANA_LOADBALANCER_IP: \$KIBANA_LOADBALANCER_IP: \$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. | |
| 2. | Using Kibana verify Log and Tracer data is stored in Elasticsearch | Navigate to "Management" Tab in Kibana. 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. jaeger-* logstash-* Type logstash* in the index pattern field and wait for few seconds. Verify the "Success" message and index pattern "logstash-YYYY.MM.DD" appeared as highlighted in the bottom red box. Click on "Next step" Select "I don't want to use the Time Filter" and click on "Create index pattern" Ensure the Web page having the indices appear in the main viewer frame Click on "Discover" Tab and you should be able to view raw Log records. Repeat steps 3-6 using "jaeger*" instead of "logstash* to ensure the data is stored in elastic search. | |
| 3. | Verify Elasticsearch cluster health | Navigate to "Dev Tools" in Kibana 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 3-1 (Cont.) OCCNE Post Install Verification

| Step No. | Procedure | Des | cription |
|-------------|---|-----|---|
| 4. | Verify Prometheus Alert manager is accessible | 2. | 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_PORT http://10.75.182.53:80 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 3-1 (Cont.) OCCNE Post Install Verification

| Step No. | Procedure | Des | scription |
|-------------|--|-----|--|
| 5. | Verify metrics are scraped and stored in | | Run below commands to get the load-balancer IP address and port number for Prometheus Server Web Interface. |
| | prometheus server | | <pre># 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}")</pre> |
| | | | <pre># 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}")</pre> |
| | | | <pre># Complete url for accessing prometheus in external browser \$ echo http://\$PROMETHEUS_LOADBALANCER_IP: \$PROMETHEUS_LOADBALANCER_PORT http://10.75.182.54:80</pre> |
| | | 2. | Launch the Browser and navigate to http:// \$PROMETHEUS_LOADBALANCER_IP: \$PROMETHEUS_LOADBALANCER_PORT (e.g.: http:// 10.75.182.54:80 in the example above) received in the output of the above commands. Ensure the Prometheus server GUI is accessible. |
| | | 3. | Select "UP" option from "insert metric at cursor" drop down and click on "Execute" button. |
| | | 4. | Here the entries present under the Element section are scrape endpoints and under the value section its corresponding status(1 for up 0 for down). Ensure all the scrape endpoints have value as 1 (means up and running). |
| 6. | Verify Alerts are configured | 1. | Navigate to alerts tab of Prometheus server GUI or navigate using URL http://\$PROMETHEUS_LOADBALANCER_IP: \$PROMETHEUS_LOADBALANCER_PORT/ alertsFor <prometheus_loadbalancer_ip>and<prometh eus_loadbalancer_port=""></prometh></prometheus_loadbalancer_ip> |
| | | 2. | If below alerts are seen in " Alerts" tab of prometheus GUI, then Alerts are configured properly. |



Table 3-1 (Cont.) OCCNE Post Install Verification

| Step No. | Procedure | Des | scription |
|-------------|--|-----|---|
| 7. | 7. Verify grafana is accessible and change the | | Run below commands to get the load-balancer IP address and port number for Grafana Web Interface. |
| | default password for admin user | | <pre># 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> |
| | | | <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> |
| | | | <pre># Complete url for accessing grafana in external browser \$ echo http://\$GRAFANA_LOADBALANCER_IP: \$GRAFANA_LOADBALANCER_PORT http://10.75.182.55:80</pre> |
| | | 2. | Launch the Browser and navigate to http:// \$GRAFANA_LOADBALANCER_IP: \$GRAFANA_LOADBALANCER_PORT (e.g.: http:// 10.75.182.55:80 in the example above) received in the output of the above commands. Ensure the Prometheus server GUI is accessible. The default username and password is admin/admin for the 1st time access. |
| | | 3. | At first connection to the Grafana dashboard, a 'Change Password' screen will appear. Change the password to the customer provided credentials. |
| | | | Note : Grafana data is not persisted, so if Grafana services restarted for some reason change password screen will appear again. |
| | | 4. | Grafana dashboards are accessed after the changing the default password in the above step. |
| | | 5. | Click on "New dashboard" as marked red below. |
| | | 6. | Click on "Add Query" |
| | | 7. | From " <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. |
| | | 8. | In the Query Section marked in Red below put " sum by(name) ({kubernetes_namespace="occne-infra"}) " and then click any where outside of the textbox and wait for few seconds. Ensure the dashboard appearing in the top section of the page. This link shows all the metrics and number of entries in each metrics over time span originated from kubernetes namespace 'occne-infra. In the add query section we can give any valid promQl query.Example for using the metrics list link above to write a promQL query: sum(\$metricnamefromlist)sum by(kubernetes_pod_name) (\$metricnamefromlist{kubernetes_namespace="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:



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

Table 3-2 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
- HP commands to configure Enclosure switch username and password:https://support.hpe.com/hpsc/doc/public/display?docId=c04763521
- 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
- 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
- 7. See Enclosure switch procedure for initial username/password: Configure Enclosure Switches

Procedure



Table 3-3 Post-Installation Security Hardening

| Stop No. | Duo a aduma | Da | anintian |
|----------|--|-------------|--|
| Step No. | Procedure | Description | |
| 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 3-3 (Cont.) Post-Installation Security Hardening

| Step No. | Procedure | Des | scription |
|----------|-----------|-----|---|
| | | | <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: # copy running-config startup-config |
| | | | # copy running config startap config [#################################### |



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

| Step No. | Procedure | Description |
|----------|--|--|
| 2. | Reset Credentials on the Enclosure Switch | From bastion host, login to the switch with username and password from the procedure: |
| | | [bastion host]# ssh <username>@<switch address="" ip=""> <username>@<switch address="" ip="">'s password: <password></password></switch></username></switch></username> |
| | | ************************************** |
| | | * Without the owner's prior written consent, |
| | | * no decompiling or reverse-engineering shall be allowed. |
| | | <pre><switchname> <switchname>sys System View: return to User View with Ctrl+Z. [switchname]</switchname></switchname></pre> |
| | | 2. Change the password for current username: |
| | | <pre>[switchname]local-user <username> class</username></pre> |
| | | 3. Create new username: |
| | | <pre>[switchname]local-user <newusername> class [manage network] New local user added. [switchname-luser-manage- <newusername>]password simple <newpassword> [switchname-luser-manage- <newusername>]quit [switchname]</newusername></newpassword></newusername></newusername></pre> |
| | | 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. [bastion host]#</switch></switchname></pre> |



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

| Step No. | Procedure | Description |
|----------|-----------|---|
| | | <pre>[bastion host]# ssh <newusername>@<switch address="" ip=""> <newusername>@<switch address="" ip="">'s password: <newpassword></newpassword></switch></newusername></switch></newusername></pre> |
| | | ************************************** |
| | | shall be allowed. *********************************** |
| | | <pre><switchname> <switchname>sys System View: return to User View with Ctrl+Z. [switchname]</switchname></switchname></pre> |
| | | 5. Delete the previous old username if it is not needed: [switchname]undo local-user <username></username> |
| | | class <current class=""></current>Save the above configuration: |
| | | <pre>[switchname]save The current configuration will be written to the device. Are you sure? [Y/N]:y Please input the file name(*.cfg) [flash:/<filename>] (To leave the existing filename unchanged, press the enter key): flash:/<filename> exists, overwrite? [Y/N]:y Validating file. Please wait Saved the current configuration to mainboard device successfully. Slot 1: Save next configuration file successfully. [switchname]</filename></filename></pre> |



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

| Step No. | Procedure | Des | scription |
|----------|---|-----|--|
| 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</oa></username> |
| | | | password: <password></password> |
| | | | HPE BladeSystem Onboard Administrator (C) Copyright 2006-2018 Hewlett Packard Enterprise Development LP |
| | | | Type 'HELP' to display a list of valid commands. Type 'HELP <command/> ' to display detailed information about a specific command. Type 'HELP HELP' to display more detailed information about the help system. |
| | | | 03.345036705701. |
| | | 2. | OA-A45D36FD5FB1> Change the password for current username: |
| | | | OA-A45D36FD5FB1> set password <newpassword></newpassword> |
| | | | Changed password for the " <username>" user account.</username> |



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

| Step No. | Procedure | Des | scription |
|----------|-----------|-----|---|
| | | | |
| | | | OA-A45D36FD5FB1> |
| | | 3. | Add new user: |
| | | | OA-A45D36FD5FB1> add user <newusername></newusername> |
| | | | New Password: <newpassword> Confirm : <newpassword> User "<newusername>" created. You may set user privileges with the 'SET USER ACCESS' and 'ASSIGN' commands.</newusername></newpassword></newpassword> |
| | | | OA-A45D36FD5FB1> set user access <newusername> [ADMINISTRATOR OPERATOR USER]</newusername> |
| | | | " <newusername>" has been given [administrator operator user] level privileges.</newusername> |
| | | 4. | Assign full access to the enclosure for the user: |
| | | | OA-A45D36FD5FB1> assign server all <newusername></newusername> |
| | | | <pre><newusername> has been granted access to the valid requested bay(s</newusername></pre> |
| | | | OA-A45D36FD5FB1> assign interconnect all <newusername></newusername> |
| | | | <pre><newusername> has been granted access to the valid requested bay(s)</newusername></pre> |
| | | | OA-A45D36FD5FB1> assign oa <newusername></newusername> |
| | | | <pre><newusername> has been granted access to the OA.</newusername></pre> |
| | | 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=""></oa></newusername></oa> |
| | | | 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 3-3 (Cont.) Post-Installation Security Hardening

| Step No. | Procedure | Desc | cription |
|----------|-----------|------|---|
| | | | tored and can result in criminal or civil prosecution under applicable law. |
| | | | |
| | | | Firmware Version: 4.85 Built: 04/06/2018 @ 06:14 OA Bay Number: 1 OA Role: Active <newusername>@<oa address="">'s password:<newpassword></newpassword></oa></newusername> |
| | | | 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> |
| | | 6. | Delete previous user if not needed: |
| | | | OA-A45D36FD5FB1> remove user <username></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.</username> |



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

| Step No. | Procedure | Des | scription |
|----------|---|-----------|---|
| 4. | Reset Credentials for the ILO Admin Console | 1. | From bastion host, login to the iLO with username and password from the procedure: [root@winterfell ~]# ssh <username>@<ilo address=""> <username>@<ilo address="">'s password: <password> User:<username> logged-in to(<ilo address=""> / <ipv6 address="">) iLO Advanced 2.61 at Jul 27 2018 Server Name: <server name=""> Server Power: On</server></ipv6></ilo></username></password></ilo></username></ilo></username> |
| | | 2. | hpiLO-> Change current password: |
| | | <i>2.</i> | <pre>hpiLO-> set /map1/accounts1/ <username> password=<newpassword> status=0 status_tag=COMMAND COMPLETED</newpassword></username></pre> |
| | | | 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 3-3 (Cont.) Post-Installation Security Hardening

| Step No. | Procedure | Description |
|----------|---|---|
| | | <pre>[bastion host]# ssh <newusername>@<ilo address=""></ilo></newusername></pre> |
| 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 3-3 (Cont.) Post-Installation Security Hardening

| CA . N. | D I | Description |
|----------|------------------------------------|--|
| 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 | 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

Artifact Acquisition and Hosting

Introduction

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

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

YUM Repository Configuration

Introduction

To perform an installation without the system needing access to the internet, a local YUM mirror must be made of the OL7 latest, epel, and addons repository used by the OS installation.

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

Prerequisites

- Local YUM mirror repository for the OL7 'latest', 'epel', and 'addons' repositories.
 Directions here: HTTPS://WWW.ORACLE.COM/TECHNICAL-RESOURCES/
 ARTICLES/IT-INFRASTRUCTURE/UNBREAKABLE-LINUX-NETWORK.HTML
- 2. Subscribe to following channels while creating the yum mirror from uln:

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

References

Oracle YUM mirroring directions:



https://www.oracle.com/technetwork/articles/servers-storage-admin/yum-reposetup-1659167.html



Procedure Steps

Table A-1 Procedure to configure OCCNE YUM Repository

| g. " | ъ . | 5 |
|--------|---|--|
| Step # | Procedure | Description |
| 1. | Create OL7 repository mirror repo | HTTPS://WWW.ORACLE.COM/TECHNICAL-RESOURCES/ARTICLES/IT-INFRASTRUCTURE/UNBREAKABLE-LINUX-NETWORK.HTML Below is an example of a repository file providing the details on a mirror with the necessary repositories. This repository file would be placed on the OCCNE Bootstrap machine that will setup the OCCNE Bastion Host. (directions on the locations in the installation procedure) |
| | | <pre>[local_ol7_x86_64_UEKR5] name=Unbreakable Enterprise Kernel Release 5 for Oracle Linux 7 (x86_64) baseurl=http://10.75.155.195/yum/OracleLinux/OL7/ UEKR5/\$basearch/ gpgcheck=1 gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY enabled=1 proxy=_none_</pre> |
| | | <pre>[local_ol7_x86_64_latest] name=Oracle Linux 7 Latest (x86_64) baseurl=http://lo.75.155.195/yum/OracleLinux/OL7/latest/ \$basearch/ gpgcheck=1 gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY enabled=1 proxy=_none_</pre> |
| | | <pre>[local_ol7_x86_64_addons] name=Oracle Linux 7 Addons (x86_64) baseurl=http://lo.75.155.195/yum/OracleLinux/OL7/addons/ \$basearch/ gpgcheck=1 gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY enabled=1 proxy=_none_</pre> |
| | | <pre>[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_</pre> |
| | | <pre>[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</pre> |



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

| Step # | Procedure | Description |
|--------|-----------|--|
| экер# | riocedure | 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://lo.75.155.195/yum/OracleLinux/OL7/developer/ EPEL/\$basearch/ gpgcheck=1 gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY enabled=1 |
| | | proxy=_none_ |

HTTP Repository Configuration

Introduction

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

Prerequisites

- 1. Docker is setup and docker commands can be run by the target system.
- 2. HTTP server that is reachable by the target system, Example- Running Nginx in docker container.

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

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

OR

Use the html directory of Apache http server created during setting up yum mirror to perform the tasks listed below. Note: Create new directories for kubernetes binaries and helm charts in html folder



Procedure Steps

Table A-2 Steps to configure OCCNE HTTP Repository

| Stons | Drogodure | Description |
|----------|------------------------------------|---|
| Steps 1. | Retrieve Kubernetes Binaries | Description 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 script will retrieve the proper binaries and place them in a directory named 'binaries' under the command-line specified directory. This 'binaries' directory needs to then be placed on the HTTP server where it can be served up, with the URL identified in the clusters hosts.ini inventory file (see below). deploy/retrieve_k8s_bins.sh #!/bin/bash #################################### |
| | | kube_version='v1.12.5' # k8s_install/ kubespray/roles/download/defaults/main.yaml |
| | | Scar carr - 45ma |



Table A-2 (Cont.) Steps to configure OCCNE HTTP Repository

| Steps | Procedure | Description |
|-------|----------------|--|
| | | <pre>mkdir -p \$1/binaries/\$kube_version wget -P \$1/binaries/\$kube_version https:// storage.googleapis.com/kubernetes-release/ release/\${kubeadm_version}/bin/linux/\$ {image_arch}/kubeadm wget -P \$1/binaries/\$kube_version https:// storage.googleapis.com/kubernetes-release/ release/\${kube_version}/bin/linux/amd64/ hyperkube wget -P \$1/binaries https://github.com/coreos/ etcd/releases/download/\${etcd_version}/etcd-\$ {etcd_version}-linux-amd64.tar.gz wget -P \$1/binaries https://github.com/ containernetworking/plugins/releases/download/ \$cni_version/cni-plugins-\${image_arch}-\$ {cni_version}.tgz</pre> |
| 2. | Run the script | <pre>\$ retrieve_k8s_bins.sh <directoryname></directoryname></pre> |



Table A-2 (Cont.) Steps to configure OCCNE HTTP Repository

| Steps | Procedure | Description | | |
|-------|--------------------------------------|---|--------|--|
| 3. | Retrieve Helm binaries and charts | The Configuration installer requires access to an HTTP server from which it can download the proper version of a set of Helm charts for the common services. To provision an internal HTTP repository one will need to obtain these charts from the internet, and place them at a known location on the internal HTTP server. | | |
| | | deploy/helm_images.txt | | |
| | | ###################################### | | |
| | | П | # | |
| | | # Copyright (c) 2019 Oracle and/or its | | |
| | | affiliates. All rights reserved. # | | |
| | | # | | |
| | | # | | |
| | | ###################################### | | |
| | | # chart-name chart-version | | |
| | | stable/elasticsearch | 1.27.2 | |
| | | stable/elasticsearch-curator | 1.2.1 | |
| | | stable/elasticsearch-exporter | 1.1.2 | |
| | | stable/fluentd-elasticsearch | 2.0.7 | |
| | | stable/grafana | 3.3.8 | |
| | | stable/kibana | 3.0.0 | |
| | | stable/metallb | 0.8.4 | |
| | | stable/prometheus | 8.8.0 | |
| | | stable/prometheus-node-exporter | | |
| | | stable/metrics-server | 2.5.1 | |
| | | incubator/jaeger | 0.8.3 | |
| | | <pre># this one is part of the configure code-base, so not pulled. There is an image associated in the docker image repo. # storage/occne-local/helm/provisioner 2.3.0</pre> | | |
| | | | | |



Table A-2 (Cont.) Steps to configure OCCNE HTTP Repository



Table A-2 (Cont.) Steps to configure OCCNE HTTP Repository

| Steps | Procedure | Description | |
|-------|----------------|---|--|
| | | helm initclient-only helm repo add incubator https://kubernetes- charts-incubator.storage.googleapis.com/ helm repo add kiwigrid https:// kiwigrid.github.io helm repo update | |
| | | <pre>helm repo update # fetch archives of helm charts and place them in the proper directory for the local HTTP repository HELMCHART_DIR=\$1/charts echo "fetching charts" # regular expression to match for valid line (variable as in-line regex is sometimes parsed differently in different bash versions) re='^(\S+)\s+(\S+)' while read line; do if [[\${line} =~ ^'#'(.*)]]; then # comment, just echo it echo "\${BASH_REMATCH[0]}" elif [[\${line} =~ ^'''(.*)]]; then # markdown code delimiter, ignore : elif [[\${line} =~ \${re}]]; then echo "Retrieving chart='\$</pre> | |
| | | <pre>{BASH_REMATCH[1]}' version='\$ {BASH_REMATCH[2]}'" helm fetch \${BASH_REMATCH[1]} version=\${BASH_REMATCH[2]} -d \${HELMCHART_DIR} fi done echo "completed fetching charts" cd \$startdir</pre> | |
| 5. | Run the script | <pre>\$ retrieve_helm.sh <directoryname> <helm_images>.txt</helm_images></directoryname></pre> | |



Steps Procedure Description 6. 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-2 (Cont.) Steps to configure OCCNE HTTP Repository

Docker Image Registry Configuration

Introduction

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

Prerequisites

Docker images for OCCNE 1.3 release must be pulled to the executing system.

- 1. Docker is installed and docker commands can be run
- 2. Make sure docker registry is running
 - \$ dockerps
- 3. If not then creating a local docker registry accessible by the target of the installation

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

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

References

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

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

Provision the registry with the necessary images

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



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

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

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

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

```
docker run --rm -it -v /var/occne/<cluster>/:/host occne/
<configure_install_image_name>:<1.3.0_tag> /getdeps/getdeps
docker run --rm -it -v /var/occne/<cluster>/:/host occne/
<k8s_install_image_name>:<1.3.0_tag> /getdeps/getdeps

Example-
$ docker run --rm -it -v /var/occne/rainbow/:/host occne/configure:1.3.0 /
getdeps/getdeps
$ docker run --rm -it -v /var/occne/rainbow/:/host occne/k8s_install:1.3.0 /
getdeps/getdeps
```

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

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

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

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

Verify the list of repositories in the docker registry

Access endpoint <registryaddress>:<port>/v2/_catalog using a browser

or

from any linux server with curl command available and can query the repo server address, using curl command:

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

Sample:

```
$ {"repositories":["coredns/coredns","docker.elastic.co/elasticsearch/
elasticsearch-oss","docker.elastic.co/kibana/kibana-oss","gcr.io/google-
containers/fluentd-elasticsearch","gcr.io/google-containers/kube-
apiserver","gcr.io/google-containers/kube-controller-manager","gcr.io/google-
containers/kube-proxy","gcr.io/google-containers/kube-scheduler","gcr.io/google-
containers/pause","gcr.io/google_containers/cluster-proportional-autoscaler-
amd64","gcr.io/google_containers/metrics-server-amd64","gcr.io/google_containers/
pause-amd64","gcr.io/kubernetes-helm/tiller","grafana/grafana","jaegertracing/
jaeger-agent","jaegertracing/jaeger-collector","jaegertracing/jaeger-
query","jimmidyson/configmap-reload","justwatch/
elasticsearch_exporter","k8s.gcr.io/addon-resizer","lachlanevenson/k8s-
```



helm", "metallb/controller", "metallb/speaker", "nginx", "prom/alertmanager", "prom/prometheus", "prom/pushgateway", "quay.io/calico/cni", "quay.io/calico/cotl", "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"]}



B

Reference Procedures

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

Inventory File Preparation

Introduction

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

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

Inventory File Overview

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

- 1. A host is defined as a Fully Qualified Domain Name (FQDN). Properties are defined as key=value pairs.
- 2. A property applies to a specific host when it appears on the same line as the host.
- 3. Square brackets define group names. For example host_hp_gen_10defines the group of physical HP Gen10 machines. There is no explicit "end of group" delimiter, rather group definitions end at the next group declaration or the end of the file. Groups can not be nested.
- 4. A property applies to an entire group when it is defined under a group heading not on the same line as a host.
- 5. Groups of groups are formed using the children keyword. For example, the occne:childrencreates an occne group comprised of several other groups.
- 6. Inline comments are not allowed.



Table B-1 Base Groups

Group Name Description/Comments

host_hp_gen_10

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 containers use this IP to communicate with the host
- 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. To get the mac address, login to the above ilo address with ssh, the username and password are the pxe install lights out usr and pxe install lights out passwd, which are created in the Configure Addresses for RMS iLOs, OA, EBIPA process. After login, run command "show / system1/network1/Integrated NICs", Port1NIC MACAddress is for eno1, Port5NIC MACAddress is for eno5.

The default configuration of a node in this group is for a Gen 10 RMS with modules providing boot interfaces at Linux interface identifiers 'eno5' and 'eno6'. For Gen 10 bladesthe boot interfaces are usually 'eno1' and 'eno2' and should be specified by adding the following properties (outlined below).

- pxe config ks nic: The bootable interface to initiate the installation process (for Gen10 blades ='eno1')
- pxe config nic list: The set of interfaces to team together (for Gen10 blades ='eno1,eno2')

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

- ansible host: The IP address for the host's primary interface. The occne containers use this interface to communicate with the host.
- kvm host: The physical host name.fqdn running KVM that hosts this VM host (ex. for guest db-10.icemark.lab.us.oracle.com the kvm host is db-1.icemark.lab.us.oracle.com). Bastion-1 should be on db-1, bastion-2 should be on db-2.
- mac: Always begin with 52-54-00 with the last 3 hex values being unique within the hosts.ini file (ex: mac=52-54-00-c1-8e-38)
- signal host: The Signalling network IPv4 address assigned to the MySQL NDB SQL Nodevirtual machines. : The ILO network IPv4 address assigned to the Bastionhost virtual machines. This IP is manually assigned and should be on the same network as the host_hp_gen_10/9/8 node iLo address except for the final octet. For ex:ample: If the kvm host=db-2 and the iLo field in host_hp_gen_10/9/8 for db-2 is set to 192.168.20.101, this value can be set to 192.168.20.201. (be sure .201 is not being used within that network by executing a ping on that address from the host node...db-1 in this example...the ping should fail).
- oam host: The OAM network IPv4 address assigned to the Bastionhost virtual machines.

kvm hosts:children The list of all physical hosts which will be hosting the virtual hosts. This should be the set data store and kube-master. Do not modify.

occne:children

Do not modify the children of the occne

occne:vars

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



Table B-1 (Cont.) Base Groups

| Group Name | Description/Comments |
|--------------------------|---|
| data_store | The list of Storage Hosts |
| kube-master | The list of Master Node hosts where kubernetes master components run. |
| etcd | The list of hosts that compose the etcd server. Should always be an odd number. This set is the same list of nodes as the kube-master |
| kube-node | The list of Worker Nodes. Worker Nodes are where kubernetes pods run and should be comprised of the bladed hosts. |
| k8s- cluster:children | Do not modify the children of k8s-cluster |
| bastion_hosts | The list of Bastion Hosts name.fqdn (ex: bastion-1.icemark.lab.us.oracle.com) |

Data Tier Groups

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

Table B-2 Data Tier Groups

| Group Name | Description/Comments |
|--------------------------|---|
| mysqlndb_mgm_nodes | The list of MySQL Management nodes. In OCCNE 1.0 this group consists of three virtual machines distributed equally among the kube-masternodes. 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. |
| mysqlndb_data_nodes_ng 0 | 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 NodeId parameter. • NodeId: Parameter must be unique across the MySQL Cluster and have a value between 1 and 48. For Ex: NodeId should be assigned with value 1 and 2 [mysqlndb_data_nodes_ng0] db-6.foo.lab.us.oracle.com NodeId=1 db-7.foo.lab.us.oracle.com NodeId=2 |
| mysqlndb_data_nodes_ng 1 | 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 NodeId parameter. • NodeId: Parameter must be unique across the MySQL Cluster and have a value between 1 and 48. For Ex: NodeId should be assigned with value 3 and 4 [mysqlndb_data_nodes_ng1] db-8.foo.lab.us.oracle.com NodeId=3 db-9.foo.lab.us.oracle.com NodeId=4 |
| mysqlndb_data_nodes | The list of MySQL Data node groups. In OCCNE 1.2 this group consists of 2 groups, each groups consists of two virtual machines distributed equally among the Storage Hosts. |



Table B-2 (Cont.) Data Tier Groups

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

Prerequisites

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

Procedure

OCCNE Cluster Name

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

Geographic and organizational identifiers may be multiple levels deep.

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

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

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

Create host_hp_gen_10 and host_kernel_virtual group lists

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



recommended host name prefix for nodes in the host_kernel_virtual group is "db-x" where x is a number 1 to N. MySQL Cluster nodes should not be differentiated using host names.

Edit occne:vars

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

Table B-3 Edit occne:vars

| occne_cluster_name subnet_ipv4 | Set to the OCCNE Cluster Name generated in step 2.1 above. Set to the subnet of the network used to assign IPs |
|-----------------------------------|--|
| subnet_ipv4 | Set to the subnet of the network used to assign IPs |
| | for OCCNE hosts |
| subnet_cidr | Appears this is not used so does not need to be included. If it does need to be included, set to the cidr notation for the subnet. For example /24 |
| netmask | Set appropriately for the network used to assign IF for OCCNE hosts. |
| proadcast_address | Set appropriately for the network used to assign IF for OCCNE hosts. |
| default_route | Set to the IP of the TOR switch. |
| name_server | 'none' |
| ntp_server | Set to the IP of the TOR switch. |
| occne_repo_host | Set to the hostname of the bootstrap host initially. This defaults to "bootstrap". It can remain as that value or the user can change it to their own specifications but they must adhere to hostname conventions. |
| occne_repo_host_address | Set to the internal (ansible_host) IPv4 address of the occne_repo_host. |
| oxe_install_lights_out_usr | Set to the user name configured for iLO admins or each host in the OCCNE Frame. |
| oxe_install_lights_out_passwd | Set to the password configured for iLO admins on each host in the OCCNE Frame. |
| lo_vlan_id | Set to the VLAN ID of the ILO network For Ex: 2 |
| lo_subnet_ipv4 | Set to the subnet of the ILO network used to assign IPs for Storage hosts |
| lo_subnet_cidr | Set to the cidr notation for the subnet. For example 24 |
| lo_netmask | Set appropriately for the network used to assign ILO IPs for Storage hosts. |
| lo_broadcast_address | Set appropriately for the network used to assign ILO IPs for OCCNE hosts. |
| lo_default_route | Set to the ILO VIP of the TOR switch. |
| mgmt_vlan_id | Set to the VLAN ID of the Management network For Ex: 4 |
| mgmt_subnet_ipv4 | Set to the subnet of the Management network used to assign IPs for Storage hosts |
| mgmt_subnet_cidr | Set to the cidr notation for the Management subne For example 29 |



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

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

Edit mysqlndb_all_nodes:vars

Table B-4 Edit mysqlndb_all_nodes:vars

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

OCCNE Inventory Boilerplate

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

Installation Preflight Checklist

Introduction

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

Prerequisites

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

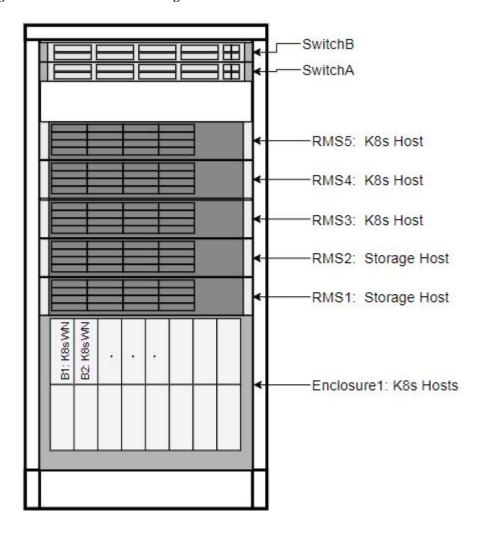


Confirm Hardware Installation

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

Rackmount ordering (frame not to scale)

Figure B-1 Rackmount ordering



Enclosure, ToR, and RMS Connections

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

Enclosure Switch Connections

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



Table B-5 Enclosure Switch Connections

| Switch Port Name/ID | Destination (To) | Cable Type | Module |
|---------------------|---|------------------------|------------|
| (From) | Destination (10) | Cable Type | 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 | 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 | 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 Artifact Acquisition and Hosting are available in repositories accessible from the OCCNE Frame.

Keyboard, Video, Mouse (KVM) Availability

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

Procedure

Complete Site Survey Subnet Table

Table B-8 Complete Site Survey Subnet Table

| SI No. | Network Description | Subnet Allocation | Bitmask | VLAN ID | Gateway Address |
|--------|--|----------------------|---------|------------|--------------------|
| 1 | iLO/OA Network | 192.168.20.0 | 24 | 2 | N/A |
| 2 | Platform Network | 172.16.3.0 | 24 | 3 | 172.16.3.1 |
| 3 | Switch Configuration Network | 192.168.2.0 | 24 | N/A | N/A |
| 4 | Management Network - Bastion Hosts | | 28 | 4 | |
| 5 | Signaling Network - MySQL Replication | | 29 | 5 | |



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

| SI 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 |
| 2 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 values that are prefilled are fixed in the topology and do not need to be changed. Blank values indicate that customer engagement is needed to determine the appropriate value.



| | 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.

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 $6$etqyspJhPUG440V0
$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
@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/
sudoers
echo "proxy=HTTP_PROXY" >> /etc/yum.conf
# Configure keys for admusr
mkdir -m0700 /home/admusr/.ssh/
chown admusr:admusr /home/admusr/.ssh
cat <<EOF >/home/admusr/.ssh/authorized_keys
PUBLIC_KEY
EOF
echo "Configuring SSH..."
cp /etc/ssh/sshd_config /etc/ssh/sshd_config.orig && \
sed -i 's/#Protocol 2/Protocol 2/' /etc/ssh/sshd_config && \
sed -i 's/#LogLevel.*/LogLevel INFO/' /etc/ssh/sshd_config && \
sed -i 's/X11Forwarding yes/X11Forwarding no/' /etc/ssh/sshd_config && \backslash
sed -i 's/#MaxAuthTries.*/MaxAuthTries 4/' /etc/ssh/sshd_config && \
sed -i 's/#IgnoreRhosts.*/IgnoreRhosts yes/' /etc/ssh/sshd_config
```



```
if [ `grep HostBasedAuthentication /etc/ssh/sshd_config | wc -l` -lt 1 ]; then
    echo 'HostBasedAuthentication no' >> /etc/ssh/sshd_config
fi
sed -i 's/#PermitRootLogin.*/PermitRootLogin no/' /etc/ssh/sshd_config && \
sed -i 's/PermitRootLogin.*/PermitRootLogin no/' /etc/ssh/sshd_config && \
sed -i 's/#PermitEmptyPasswords.*/PermitEmptyPasswords no/' /etc/ssh/sshd_config
sed -i 's/#PermitUserEnvironment.*/PermitUserEnvironment no/' /etc/ssh/
sshd_config && \
sed -i 's/PermitUserEnvironment.*/PermitUserEnvironment no/' /etc/ssh/sshd_config
if [ `grep -i 'Ciphers aes128-ctr,aes192-ctr,aes256-ctr' /etc/ssh/sshd_config |
wc -l` -lt 1 ]; then
    echo 'Ciphers aes128-ctr,aes192-ctr,aes256-ctr' >> /etc/ssh/sshd_config
    if [ $? -ne 0 ]; then
        echo " ERROR: echo 1 failed"
    fi
fi
if [ `grep '^MACs' /etc/ssh/sshd_config | wc -l` -lt 1 ]; then
        echo 'MACs hmac-sha2-512-etm@openssh.com,hmac-sha2-256-
etm@openssh.com,umac-128-etm@openssh.com,hmac-sha2-512,hmac-
sha2-256,umac-128@openssh.com' >> /etc/ssh/sshd_config
        if [ $? -ne 0 ]; then
            echo " ERROR: echo 2 failed"
        fi
fi
sed -i 's/#ClientAliveInterval.*/ClientAliveInterval 300/' /etc/ssh/sshd_config
sed -i 's/#ClientAliveCountMax.*/ClientAliveCountMax 0/' /etc/ssh/sshd_config
sed -i 's/#Banner.*/Banner \/etc\/issue.net/' /etc/ssh/sshd_config
```



egrep -q "^(\s*)LoginGraceTime\s+\S+(\s*#.*)?\s*\$" /etc/ssh/sshd_config && sed -ri "s/^(\s*)LoginGraceTime\s+\S+(\s*#.*)?\s*\$/\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

Inventory File Template

The host.ini file contains the inventory used by the various OCCNE deployment containers that will instantiate the OCCNE cluster.

Template example

The inventory is composed of multiple groups (indicated by bracketed strings):

- local: OCCNE ansible use. Do not modify.
- occne: list of servers in the OCCNE cluster that will be installed by the os_install
 container.
- k8s-cluster: list of servers in the kubernetes cluster.
- kube-master: list of servers that will be provisioned as kubernetes master nodes by the k8s install container.
- kube-node: list of servers that will be provisioned as kubernetes worker nodes by the k8s install container.
- etcd: list of servers that will be provisioned as part of kubernetes etcd cluster by the k8s install container.
- data_store: list of servers that will be host the VMs of the MySQL database cluster, os install container will install kvm on them.
- occne:vars: list of occne environment variables. Values for variables are required. See below for description.



OCCNE Variables

Table B-14 OCCNE Variables

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



Table B-14 (Cont.) OCCNE Variables

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

Configuring Cross-Site Database Replication

This section details the procedure for configuring database replication between two instances of OCCNE.

Prerequisites

- 1. Two separate instances of OCCNE must be installed and running.
- 2. Each site has 2 SQL nodes. The SQL nodes in each instance are assigned IDs 55 and 56.
- **3.** There is network connectivity between the 2 sites.

Procedures

In the following procedures, the following shorthand references will be used:

Site 1 = the first installed instance of OCCNE Site 2 = the second installed instance of OCCNE SQL node A = SQL node A at either site SQL node B = SQL node B at either site

Configure replication

This procedure will configure the SQL nodes for replication.



Table B-15 Configure replication

| Step No# | Procedure | Description |
|----------|----------------------|--|
| 1. | Configure SQL node A | At Site 1, login to any management node. From the MySQL Cluster Manager Client execute: |
| | | <pre>mcm> set server-id:mysqld:56=11 occnendbclustera; mcm> set binlog-format:mysqld:56=row occnendbclustera; mcm> set log_bin:mysqld:56=/var/occnedb/binlogs/ mysql-bin.log occnendbclustera; mcm> set relay_log:mysqld:56=/var/occnedb/mysql/ mysql-relay-bin occnendbclustera; mcm> set relay_log_index:mysqld:56=/var/occnedb/ mysql/mysql-relay-bin.index occnendbclustera; mcm> set expire_logs_days:mysqld:56=10 occnendbclustera; mcm> set max_binlog_size:mysqld:56=10 occnendbclustera; mcm> set auto-increment-increment:mysqld:56=2 occnendbclustera; mcm> set auto-increment-offset:mysqld:56=2 occnendbclustera; mcm> set ndb-log-update-as-write:mysqld:56=0 occnendbclustera; mcm> set slave-skip-errors:mysqld:56=1590 occnendbclustera; mcm> set skip-slave-start:mysqld:56=TRUE occnendbclustera;</pre> |
| 2. | Configure SQL node B | At Site 1, login to any management node. From the MySQL Cluster Manager Client execute: |
| | | <pre>mcm> set server-id:mysqld:57=12 occnendbclustera;</pre> |

Configure Site 2

The following steps are performed at Site 2.



Table B-16 Configure Site 2

| Step No# | Procedure | Des | scription |
|----------|-----------------------|-----|---|
| 1. | | 1. | Configure SQL node A At Site 2, login to any management node. From the MySQL Cluster Manager Client execute: |
| | | | <pre>mcm> set server-id:mysqld:56=21 occnendbclustera; mcm> set binlog-format:mysqld:56=row occnendbclustera; mcm> set log_bin:mysqld:56=/var/occnedb/ binlogs/mysql-bin.log occnendbclustera; mcm> set relay_log:mysqld:56=/var/occnedb/ mysql/mysql-relay-bin occnendbclustera; mcm> set relay_log_index:mysqld:56=/var/ occnedb/mysql/mysql-relay-bin.index occnendbclustera; mcm> set expire_logs_days:mysqld:56=10 occnendbclustera; mcm> set max_binlog_size:mysqld: 56=1073741824 occnendbclustera; mcm> set avaic ingreenent ingreenent:mysqld:</pre> |
| | | | <pre>mcm> set auto-increment-increment:mysqld: 56=2 occnendbclustera; mcm> set auto-increment-offset:mysqld:56=2 occnendbclustera; mcm> set ndb-log-update-as-write:mysqld: 56=0 occnendbclustera; mcm> set slave-skip-errors:mysqld:56=1590 occnendbclustera; mcm> set skip-slave-start:mysqld:56=TRUE occnendbclustera;</pre> |
| | | 2. | Configure SQL node B At Site 2, login to any management node. From the MySQL Cluster Manager Client execute: |
| | | | <pre>mcm> set server-id:mysqld:57=12 occnendbclustera;</pre> |
| 2. | Restart all SQL nodes | 1. | Restart SQL nodes on Site 1 At Site 1, login to any management node. From the MySQL Cluster Manager client execute: |
| | | | <pre>mcm> stop process 56 occnendbclustera; mcm> start process 56 occnendbclustera; mcm> stop process 57 occnendbclustera; mcm> start process 57 occnendbclustera;</pre> |
| | | 2. | Restart SQL nodes on Site 2 At Site 2, login to any management node. From the MySQL Cluster Manager client execute: |
| | | | <pre>mcm> stop process 56 occnendbclustera; mcm> start process 56 occnendbclustera; mcm> stop process 57 occnendbclustera; mcm> start process 57 occnendbclustera;</pre> |



Table B-16 (Cont.) Configure Site 2

| Step No# | Procedure | Des | scription |
|----------|--------------------------|-----|---|
| 3. | Create replication users | 1. | Create replication user at Site 1 At Site 1, login to SQL node A. From the MySQL client execute: |
| | | | <pre>mysql> /usr/local/mysql/bin/mysql -h 127.0.0.1 -uroot -p<password> mysql> GRANT REPLICATION SLAVE ON *.* TO 'myslave'@'<public 2="" a="" at="" ip="" node="" of="" site="" sql="">' IDENTIFIED BY '<password>';</password></public></password></pre> |
| | | 2. | Create replication use at Site 2 At Site 2, login to SQL node A. From the MySQL client execute: |
| | | | <pre>mysql> /usr/local/mysql/bin/mysql -h 127.0.0.1 -uroot -p<password> mysql> GRANT REPLICATION SLAVE ON *.* TO 'myslave'@'<ip 1="" a="" at="" node="" of="" site="" sql="">' IDENTIFIED BY '<password>';</password></ip></password></pre> |



Table B-16 (Cont.) Configure Site 2

| Step No# | Procedure | Des | scription |
|----------|------------------------|-----|--|
| 4. | Initialize replication | 1. | Get replication information for Site 2 At Site 2, login to SQL node A. From the MySQL client execute: |
| | | | mysql> SHOW MASTER STATUS; |
| | | | You will see output similar to the following: |
| | | | + |
| | | | ++ |
| | | | File Position Binlog_Do_DB Binlog_Ignore_DB Executed_Gtid_Set |
| | | | + |
| | | | ++ |
| | | | mysql-bin.000003 154 |
| | | | + |
| | | | ++ |
| | | | 1 row in set (0.00 sec) |
| | | | Note : This is just example output; the values you see will be different. |
| | | 2. | Initialize replication at Site 1 At Site 1, login to SQL node A. From the MySQL client execute: |
| | | | <pre>mysql> STOP SLAVE; Query OK, 0 rows affected, 1 warning (0.00 sec)</pre> |
| | | | <pre>mysql> CHANGE MASTER TO master_host='<ip 2="" a="" at="" node="" of="" site="" sql="">', master_port=3306, master_user='myslave', master_password='<password>', master_log_file='<binlog file="">', master_log_pos=<binlog position="">; Query OK, 0 rows affected, 1 warning (0.05 sec)</binlog></binlog></password></ip></pre> |
| | | | <pre>mysql> START SLAVE; Query OK, 0 rows affected (0.00 sec)</pre> |
| | | | Where Sinlog_file> is the value of the File column in the output shown above, and Sinlog_position> is the value of the Position column in the output shown above. |
| | | 3. | Get replication information for Site 1 |



Table B-16 (Cont.) Configure Site 2

| Step No# | Procedure | Description | |
|----------|-----------|---|--------|
| | | At Site 1, login to SQL node A. From the MySQL c execute: | lient |
| | | mysql> SHOW MASTER STATUS; | |
| | | You will see output similar to the following: | |
| | | + | |
| | | ++ | |
| | | File Position Binlog_Do_DB Binlog_Ignore_DB Executed_Gtid_Set | |
| | | + + +++ | |
| | | mysql-bin.000001 580 | |
| | | + | |
| | | ++ | |
| | | 1 row in set (0.00 sec) | |
| | | Note : This is just example output; the values you se be different. | e will |
| | | 4. Initialize replication at Site 2 At Site 2, login to SQL node A. From the MySQL of execute: | client |
| | | <pre>mysql> STOP SLAVE; Query OK, 0 rows affected, 1 warning (0 sec)</pre> | .00 |
| | | <pre>mysql> CHANGE MASTER TO master_host='<ib 1="" a="" at="" node="" site="" sql="">', master_port=3300 master_user='myslave', master_password='<password>', master_log_file='<binlog file="">',</binlog></password></ib></pre> | |

sec)

mysql> START SLAVE;

Where

sinlog_file> is the value of the File column in the output shown above, and

sinlog_position> is the value of the Position column in the output shown above.

Query OK, 0 rows affected, 1 warning (0.05

master_log_pos=<binlog position>;

Query OK, 0 rows affected (0.00 sec)



Table B-16 (Cont.) Configure Site 2

| Step No# | Procedure | Description |
|----------|---------------------------|---|
| 5. | Verify replication status | Verify Site1 At Site 1, login to SQL node A. From the MySQL client execute: |
| | | mysql> SHOW SLAVE STATUS \G; |
| | | ************************************** |
| | | <pre>Slave_IO_State: Waiting for master to send event</pre> |
| | | Master_Host: 10.75.212.246 |
| | | Master_User: myslave |
| | | Master_Port: 3306 |
| | | Connect_Retry: 60 |
| | | Master_Log_File: mysql-bin.000001 |
| | | Read_Master_Log_Pos: 4175 |
| | | Relay_Log_File: mysql-relay-bin.000002 |
| | | Relay_Log_Pos: 3810 |
| | | Relay_Master_Log_File: mysql-bin.000001 |
| | | Slave_IO_Running: Yes |
| | | Slave_SQL_Running: Yes |
| | | Replicate_Do_DB: |
| | | Replicate_Ignore_DB: |
| | | Replicate_Do_Table: |
| | | Replicate_Ignore_Table: |
| | | Replicate_Wild_Do_Table: |
| | | Replicate_Wild_Ignore_Table: |
| | | Last_Errno: 0 |
| | | Last_Error: |
| | | Skip_Counter: 0 |
| | | Exec_Master_Log_Pos: 4175 |
| | | Relay_Log_Space: 4017 |
| | | Until_Condition: None |
| | | Until_Log_File: |
| | | Until_Log_Pos: 0 |
| | | Master_SSL_Allowed: No |
| | | Master_SSL_CA_File: |
| | | Master_SSL_CA_Path: |
| | | Master_SSL_Cert: |
| | | Master_SSL_Cipher: |
| | | Master_SSL_Key: |
| | | Seconds_Behind_Master: 0 |
| | | Master_SSL_Verify_Server_Cert: No |
| | | Last IO Errno: 0 |
| | | Last IO Error: |
| | | Last_SQL_Errno: 0 |
| | | |
| | | Last_SQL_Error: |
| | | Replicate_Ignore_Server_Ids: |
| | | Master_Server_Id: 2155 |
| | | Master_UUID: 41ea3e4f-0a6f-11e9-8d7a- |
| | | fa163fe70331 |
| | | Master_Info_File: /usr/local/mysql/data/ |
| | | master.info |
| | | SQL_Delay: 0 |



Table B-16 (Cont.) Configure Site 2

| Step No# | Procedure | Description |
|----------|-----------|--|
| | | SQL_Remaining_Delay: NULL Slave_SQL_Running_State: Slave has read al relay log; waiting for more updates Master_Retry_Count: 86400 Master_Bind: Last_IO_Error_Timestamp: Last_SQL_Error_Timestamp: Master_SSL_Crl: Master_SSL_Crl: Retrieved_Gtid_Set: Executed_Gtid_Set: Auto_Position: 0 Replicate_Rewrite_DB: |
| | | Channel_Name: Master_TLS_Version: 1 row in set (0.00 sec) |
| | | Verify that the Slave_IO_Running and Slave_SQL_Running fields both have the value "Yes". |
| | | <pre>Slave_IO_Running: Yes Slave_SQL_Running: Yes</pre> |
| | | Verify that the Last_IO_Errno and Last_SQL_Errno fields both have the value "Yes". Verify that the Last_IO_Error and Last_SQL_Error field values are been empty. |
| | | Last_IO_Errno: 0 Last_IO_Error: Last_SQL_Errno: 0 Last_SQL_Error: |
| | | If the values of fields referenced above match the lister values, then replication is working properly at this site |
| | | Verify Site 2 At Site 2, login to SQL node A. From the MySQL clie execute: |
| | | mysql> SHOW SLAVE STATUS \G; |
| | | ************************************** |
| | | Master_Log_File: mysql-bin.000003 Read_Master_Log_Pos: 154 Relay_Log_File: mysql-relay-bin.000002 Relay_Log_Pos: 320 |



Table B-16 (Cont.) Configure Site 2

| Step No# | Procedure | Description |
|----------|-----------|---|
| | | Relay_Master_Log_File: mysql-bin.000003 |
| | | Slave_IO_Running: Yes |
| | | Slave_SQL_Running: Yes |
| | | Replicate_Do_DB: |
| | | Replicate_Ignore_DB: |
| | | Replicate_Do_Table: |
| | | Replicate_Ignore_Table: |
| | | Replicate_Wild_Do_Table: |
| | | Replicate_Wild_Ignore_Table: |
| | | Last Errno: 0 |
| | | Last Error: |
| | | Skip_Counter: 0 |
| | | Exec_Master_Log_Pos: 154 |
| | | Relay_Log_Space: 527 |
| | | Until_Condition: None |
| | | Until_Log_File: |
| | | Until_Log_Pos: 0 |
| | | Master_SSL_Allowed: No |
| | | Master_SSL_CA_File: |
| | | Master_SSL_CA_Path: |
| | | Master_SSL_Cert: |
| | | Master_SSL_Cipher: |
| | | Master_SSL_Key: |
| | | Seconds_Behind_Master: 0 |
| | | Master_SSL_Verify_Server_Cert: No |
| | | Last_IO_Errno: 0 |
| | | Last_IO_Error: |
| | | Last_SQL_Errno: 0 |
| | | Last_SQL_Error: |
| | | Replicate_Ignore_Server_Ids: |
| | | Master_Server_Id: 1155 |
| | | Master_UUID: a9a78193-0a6e-11e9-9dbe- |
| | | fal63fab3df6 |
| | | Master_Info_File: /usr/local/mysql/data/ |
| | | master.info |
| | | SQL Delay: 0 |
| | | SQL_Remaining_Delay: NULL |
| | | Slave SQL Running State: Slave has read a |
| | | relay log; waiting for more updates |
| | | Master_Retry_Count: 86400 |
| | | Master_Bind: |
| | | Last IO Error Timestamp: |
| | | Last_SQL_Error_Timestamp: |
| | | Master_SSL_Crl: |
| | | Master_SSL_Crlpath: |
| | | Retrieved Gtid Set: |
| | | Executed Gtid_Set: |
| | | Auto_Position: 0 |
| | | Replicate_Rewrite_DB: |
| | | Channel Name: |
| | | Channel_Name: Master_TLS_Version: |
| | | Masrer Tro Aerstoll. |



Table B-16 (Cont.) Configure Site 2

| Step No# | Procedure | Description |
|----------|-----------|--|
| | | Verify that the Slave_IO_Running and Slave_SQL_Running fields both have the value "Yes". |
| | | <pre>Slave_IO_Running: Yes Slave_SQL_Running: Yes</pre> |
| | | Verify that the Last_IO_Errno and Last_SQL_Errno fields both have the value "Yes". Verify that the Last_IO_Error and Last_SQL_Error field values are both empty. |
| | | <pre>Last_IO_Errno: 0 Last_IO_Error: Last_SQL_Errno: 0 Last_SQL_Error:</pre> |
| | | If the values of fields referenced above match the listed values, then replication is working properly at this site. |

Install Db Monitor service

The DB Monitor service is used for collecting and monitoring the MySQL NDB Cluster metrics. These Metrics are monitored and are used to raise the alerts when an anomaly occurs or when a potential failure is round the corner.

The below procedure are used to install the db monitor service

Prerequisites

- 1. All procedures in Bare Metal Installation Procedure is complete.
- 2. The host running the docker image must have dockerMySQL NDB Cluster should be installed.
- 3. A docker image named 'occne/db_monitor_svc' must be available in the customer repository.

Limitations and Expectations

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



Procedure

Table B-17 Install Db Monitor service

| Step No # | Procedure | Des | scription |
|-----------|-------------------------------|-----|--|
| 1. | Install Db Monitor Service | 1. | Retrieve the Db Monitor Service docker image from docker registry to bastion host. |
| | | | a. Login to bastion host |
| | | | b. Pull the Db Monitor Service Docker image. |
| | | | <pre>\$ docker pull <docker_common_registry: 5000="">/occne/db_monitor_svc:1.3.0 \$ docker tag <docker_common_registry: 5000="">/occne/db_monitor_svc:1.3.0 <bastion_host:5000>/occne/ db_monitor_svc:1.3.0 \$ docker push <bastion_host:5000>/occne/ db_monitor_svc:1.3.0</bastion_host:5000></bastion_host:5000></docker_common_registry:></docker_common_registry:></pre> |
| | | | For Example: if Docker_Common_registry is winterfell and bastion host is stark-bastion1 |
| | | | <pre>\$ docker pull winterfell:5000/occne/ db_monitor_svc:1.3.0 \$ docker tag winterfell:5000/occne/ db_monitor_svc:1.3.0 stark- bastion1:5000/occne/db_monitor_svc:1.3.0 \$ docker push stark-bastion1:5000/occne/ db_monitor_svc:1.3.0</pre> |
| | | 2. | Install Db Monitor Service using helm |
| | | | a. Update values.yaml file, update repository, tag, public IP of the SQL node (primary_sql_host_ip, secondary_sql_host_ip) |
| | | | <pre>\$ cd <directory db="" for="" monitor="" service=""> # update repository, tag, primary_sql_host_ip, secondary_sql_host_ip \$ vi values.yaml image: repository: <bastion_host:5000>/occne/ db_monitor_svc tag: 1.3.0 pullPolicy: Always</bastion_host:5000></directory></pre> |
| | | | <pre>mysql: primaryhost: "primary_sql_host_ip" secondaryhost: "secondary_sql_host_ip" port: "3306"</pre> |



b. Run the HELM install command to install NF using

HELM chart.

Table B-17 (Cont.) Install Db Monitor service

| Step No # | Procedure | Des | cription |
|-----------|---|-----|---|
| | | | <pre>\$ cd <directory db="" for="" monitor="" service=""> # HELM chart should be present in the repository with the chart name specified in the below command \$ helm installname occne-db-monitor- svcnamespace occne-infra -f values.yaml</directory></pre> |
| | | | Example: \$ cd /home/admusr/db-monitor-service \$ helm installname occne-db-monitor- svcnamespace occne-infra -f values.yaml |
| 2. | Verify Db Monitor Service Installation | 1. | Once Db Monitor Service is installed using the above HELM Install command, Run the below command to check the status of Db Monitor Service. |
| | | | \$ kubectl get podsnamespace=occne-infra |
| | | 2. | To verify if the above command was executed successfully and whether the Db Monitor Service was installed properly, the output shown below can be taken as reference. |
| | | | NAME READY STATUS RESTARTS AGE occne-db-monitor-svc 1/1 Running 0 98m |

Install Additional Services/Network Functions

This assumes the service has docker images located on a docker registry that is reachable by the cluster's bastion, and associated helm charts located at a URL also accessible by the bastion.

Run the following commands from the cluster bastion.



Table B-18 Install Additional Services/Network Functions

| Step No # | Procedure | Description | | | | |
|---|--|---|--|--|--|--|
| 1. | Copy docker images needed for the service into the bastion-host | Create a file docker_images.txt listing the required docker images and tags | | | | |
| | docker registry | dockerRepo:5000/ <servicenameimage></servicenameimage> | | | | |
| | | Example | | | | |
| | | dockerRepo:5000/serviceNameImage2:1.2.2 | | | | |
| | | 2. Load these images into the cluster-local docker registry by running: | | | | |
| | | <pre>\$ /var/occne/cluster/<cluster>/artifacts/ retrieve_docker.sh <central-repo>:<central-repo- docker-port=""> \${HOSTNAME}:5000 < docker_images.txt</central-repo-></central-repo></cluster></pre> | | | | |
| 2. | Copy helm charts needed for the service into the bastion-host helm chart repository | 1. Create a file helm_charts.txt listing the required helm charts and versions: | | | | |
| | | helmRepoName/chart_name <version></version> | | | | |
| | | 2. Add the source helm repository to the cluster-local helm configuration (this need only be done once for each repo): | | | | |
| | | <pre>\$ helm repo add helmRepoName https:// someUrl.oracle.com</pre> | | | | |
| | | 3. Load the chart(s) into the cluster-local helm chart repository by running: | | | | |
| | | <pre>\$ /var/occne/cluster/<cluster>/artifacts/ retrieve_helm.sh http://<central-repo>/occne/ charts /var/www/html/occne /var/occne/cluster/\$ {OCCNE_CLUSTER}/artifacts/ < helm_charts.txt</central-repo></cluster></pre> | | | | |
| 3. Install the service Create a values yaml file on the Bastion Host the needed by the Helm chart | | Create a values.yaml file on the Bastion Host that contains the values needed by the Helm chart | | | | |
| | | To install the service run: | | | | |
| | | <pre>\$ helm installname <release-name>namespace <service-namespace> -f values.yaml <chart_name></chart_name></service-namespace></release-name></pre> | | | | |
| 4. | Updating an already installed service | To update the service run: | | | | |
| | | <pre>\$ helm upgrade -f values.yaml <release-name> <chart_name></chart_name></release-name></pre> | | | | |
| 5. | Removing an already | To remove the service run: | | | | |
| | installed service | <pre>\$ helm del <release-name>purge</release-name></pre> | | | | |

Change MySQL root user password

Following is the procedure to change MySQL root user password.



As part of the installation of the MySQL Cluster, db_install container generates the random password and marked as expired in the MySQL SQL nodes. This password is stored in "/var/occnedb/mysqld_expired.log" file. so we need to login to the each of the MySQL SQL nodes and change the MySQL root user password.

Table B-19 Change MySQL root user password

| Step No.# | Procedure | Description |
|-----------|---|--|
| 1. | Login to MySQL SQL Node VM. | |
| 2. | Login using mysql client | 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 | Enter expired random password stored in "/var/occnedb/mysqld_expired.log" file. \$ mysql -h 127.0.0.1 -uroot -p \$ Enter password: |
| 4. | Change Root Password | Execute the following commands to change the root password: \$ mysql> ALTER USER 'root'@'localhost' IDENTIFIED BY ' <new_password>'; \$ mysql> FLUSH PRIVILEGES; Note: Here 'NEW_PASSWORD' is the password of the mysql root user.</new_password> |

MySQL Repository Requirements

MySQL Cluster Manager is a distributed client/server application consisting of two main components. The MySQL Cluster Manager agent is a set of one or more agent processes that manage NDB Cluster nodes, and the MySQL Cluster Manager client provides a command-line interface to the agent's management functions.

In OCCNE MySQL Cluster Manager 1.4.7 binary distributions that include MySQL NDB Cluster will be used for installing MySQL Cluster Manager 1.4.7 and MySQL NDB Cluster 7.6.8. The complete MySQL NDB Cluster 7.6.8 binary distribution is included in this below software.

MySQL Cluster Binaries

Below binary is used for installation of MySQL Cluster Manager along with the MySQL NDB Cluster, This binary distributions includes MySQL Cluster Manager 1.4.7 and MySQL NDB Cluster 7.6.8 in it. This software will be downloaded from the Oracle Software Delivery Cloud (OSDC) i.e. https://edelivery.oracle.com.

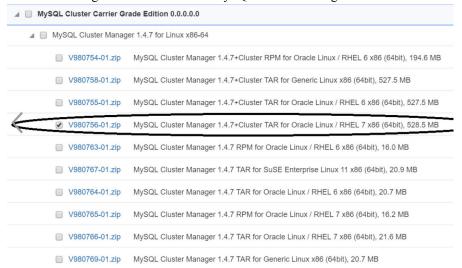
Download MySQL Cluster Manager

- Login/Access https://edelivery.oracle.com Oracle Software Delivery Cloud (OSDC) page, to download MySQL Cluster Manager 1.4.7+Cluster TAR for Oracle Linux / RHEL 7 x86
- 2. Enter "MySQL Cluster Carrier Grade Edition" and click on **Search**.



- 3. "DLP:MySQL Cluster Carrier Grade Edition 0.0.0.0.0 (MySQL Cluster Carrier Grade Edition)" is listed, click on **Add to Cart**.
- 4. Click on **Checkout**, following page will be displayed, deselect "Selected Software".
- Select MySQL Cluster Manager 1.4.7 and Select Platform as "Linux x86-64", Click on "Continue".
- Accept the licence agreement and click on "Continue".
- 7. Select "MySQL Cluster Manager 1.4.7+Cluster TAR for Oracle Linux / RHEL 7 x86 (64bit)" as shown below and Click on "**Download**".

This will install download manager and then provide the path where to download, download manager will download the MySQL Cluster Manager 1.4.7+Cluster software.



8. View the download progress in Download manager, once download is completed, MySQL Cluster Manager software (V980756-01.zip) is downloaded. Once download is completed, V980756-01.zip file is used to install MySQL Cluster Manager 1.4.7(MCM) and MySQL NDB Cluster 7.6.8.

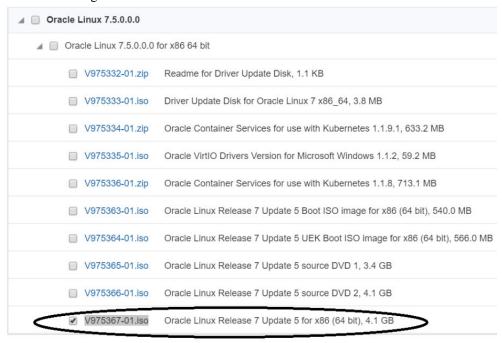
Oracle Linux 7.5 Download Instructions

The procedure to download Oracle Linux 7.5 is explained below:

- Login/Access https://edelivery.oracle.com Oracle Software Delivery Cloud (OSDC) page, to download Oracle Linux 7.5 ISO.
- Enter "Oracle Linux 7.5" and click on Search
 "DLP: Oracle Linux 7.5.0.0.0 (Oracle Linux)" will be listed as shown below, click on
 Add to Cart.
- 3. Click on **Checkout**, following page will be displayed.
- 4. Select "Oracle Linux 7.5.0.0.0" and Select Platform as "x86-64 bit", Click on "Continue".
- 5. Accept the licence agreement and click on "Continue".
- 6. Select "Oracle Linux Release 7 Update 5 for x86 (64 bit), 4.1 GB" as shown below and Click on "**Download**".
 - V975367-01.iso Oracle Linux Release 7 Update 5 for x86 (64 bit), 4.1 GB



This will install download manager and then provide the path where to download, download manager will download the Oracle Linux 7.5 software.



Total 10 distinct files Total Size 14.1 GB

7. View the download progress in Download manager, once download is completed, Oracle Linux 7.5(V975367-01.iso) is downloaded. Once the download is complete the Oracle Linux 7.5(V975367-01.iso) is used for installing host servers and VM creation.

