

Oracle® Private Cloud Appliance

Administrator's Guide for Release 2.4.2

ORACLE®

F23452-05
March 2020

Oracle Legal Notices

Copyright © 2013, 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.

Table of Contents

Preface	vii
1 Concept, Architecture and Life Cycle of Oracle Private Cloud Appliance	1
1.1 What is Oracle Private Cloud Appliance	1
1.2 Hardware Components	2
1.2.1 Management Nodes	4
1.2.2 Compute Nodes	5
1.2.3 Storage Appliance	5
1.2.4 Network Infrastructure	8
1.3 Software Components	12
1.3.1 Oracle Private Cloud Appliance Dashboard	12
1.3.2 Password Manager (Wallet)	13
1.3.3 Oracle VM Manager	13
1.3.4 Operating Systems	13
1.3.5 Databases	13
1.3.6 Oracle Private Cloud Appliance Management Software	17
1.3.7 Oracle Private Cloud Appliance Diagnostics Tool	18
1.4 Provisioning and Orchestration	19
1.4.1 Appliance Management Initialization	19
1.4.2 Compute Node Discovery and Provisioning	19
1.4.3 Server Pool Readiness	20
1.5 High Availability	21
1.6 Oracle Private Cloud Appliance Backup	22
1.7 Oracle Private Cloud Appliance Upgrader	23
2 Monitoring and Managing Oracle Private Cloud Appliance	25
2.1 Connecting and Logging in to the Oracle Private Cloud Appliance Dashboard	26
2.2 Oracle Private Cloud Appliance Accessibility Features	28
2.3 Hardware View	28
2.4 Network Settings	32
2.5 Functional Networking Limitations	35
2.5.1 Network Configuration of Ethernet-based Systems	36
2.5.2 Network Configuration of InfiniBand-based Systems	38
2.6 Network Customization	39
2.6.1 Configuring Custom Networks on Ethernet-based Systems	40
2.6.2 Configuring Custom Networks on InfiniBand-based Systems	44
2.6.3 Deleting Custom Networks	49
2.7 Tenant Groups	50
2.7.1 Design Assumptions and Restrictions	50
2.7.2 Configuring Tenant Groups	51
2.8 Authentication	54
2.9 Health Monitoring	57
3 Updating Oracle Private Cloud Appliance	61
3.1 Before You Start Updating	61
3.1.1 Warnings and Cautions	62
3.1.2 Backup Prevents Data Loss	64
3.2 Using the Oracle Private Cloud Appliance Upgrader	64
3.2.1 Rebooting the Management Node Cluster	65
3.2.2 Installing the Oracle Private Cloud Appliance Upgrader	65
3.2.3 Verifying Upgrade Readiness	66
3.2.4 Executing a Controller Software Update	68
3.3 Upgrading the Virtualization Platform	75
3.4 Upgrading Component Firmware	78

3.4.1	Firmware Policy	78
3.4.2	Install the Current Firmware on All Compute Nodes	79
3.4.3	Upgrading the Operating Software on the Oracle ZFS Storage Appliance	79
3.4.4	Upgrading the Cisco Switch Firmware	83
3.4.5	Upgrading the NM2-36P Sun Datacenter InfiniBand Expansion Switch Firmware	86
3.4.6	Upgrading the Oracle Fabric Interconnect F1-15 Firmware	89
4	The Oracle Private Cloud Appliance Command Line Interface (CLI)	93
4.1	CLI Usage	94
4.1.1	Interactive Mode	94
4.1.2	Single-command Mode	96
4.1.3	Controlling CLI Output	96
4.1.4	Internal CLI Help	98
4.2	CLI Commands	99
4.2.1	add compute-node	99
4.2.2	add network	100
4.2.3	add network-to-tenant-group	101
4.2.4	backup	102
4.2.5	configure vhbases	103
4.2.6	create lock	104
4.2.7	create network	105
4.2.8	create tenant-group	106
4.2.9	create uplink-port-group	107
4.2.10	delete config-error	108
4.2.11	delete lock	109
4.2.12	delete network	110
4.2.13	delete task	111
4.2.14	delete tenant-group	112
4.2.15	delete uplink-port-group	113
4.2.16	deprovision compute-node	114
4.2.17	diagnose	116
4.2.18	get log	119
4.2.19	list	120
4.2.20	remove compute-node	126
4.2.21	remove network	127
4.2.22	remove network-from-tenant-group	128
4.2.23	reprovision	129
4.2.24	rerun	130
4.2.25	set system-property	131
4.2.26	show	133
4.2.27	start	138
4.2.28	stop	139
4.2.29	update appliance	140
4.2.30	update password	140
4.2.31	update compute-node	142
5	Managing the Oracle VM Virtual Infrastructure	145
5.1	Guidelines and Limitations	146
5.2	Logging in to the Oracle VM Manager Web UI	149
5.3	Monitoring Health and Performance in Oracle VM	149
5.4	Creating and Managing Virtual Machines	150
5.5	Managing Virtual Machine Resources	153
5.6	Configuring Network Resources for Virtual Machines	155
5.6.1	Configuring VM Network Resources on Ethernet-based Systems	155
5.6.2	Configuring VM Network Resources on InfiniBand-based Systems	158
5.7	Viewing and Managing Storage Resources	161

5.7.1 Oracle ZFS Storage Appliance ZS7-2	162
5.7.2 Oracle ZFS Storage Appliance ZS5-ES and Earlier Models	162
5.8 Tagging Resources in Oracle VM Manager	163
5.9 Managing Jobs and Events	163
6 Servicing Oracle Private Cloud Appliance Components	165
6.1 Oracle Auto Service Request (ASR)	166
6.1.1 Understanding Oracle Auto Service Request (ASR)	166
6.1.2 ASR Prerequisites	167
6.1.3 Setting Up ASR and Activating ASR Assets	167
6.2 Replaceable Components	168
6.2.1 Rack Components	168
6.2.2 Oracle Server X8-2 Components	169
6.2.3 Oracle Server X7-2 Components	170
6.2.4 Oracle Server X6-2 Components	171
6.2.5 Oracle Server X5-2 Components	172
6.2.6 Sun Server X4-2 Components	173
6.2.7 Sun Server X3-2 Components	173
6.2.8 Oracle ZFS Storage Appliance ZS7-2 Components	174
6.2.9 Oracle ZFS Storage Appliance ZS5-ES Components	175
6.2.10 Oracle ZFS Storage Appliance ZS3-ES Components	176
6.2.11 Sun ZFS Storage Appliance 7320 Components	178
6.2.12 Oracle Switch ES1-24 Components	179
6.2.13 NM2-36P Sun Datacenter InfiniBand Expansion Switch Components	179
6.2.14 Oracle Fabric Interconnect F1-15 Components	180
6.3 Preparing Oracle Private Cloud Appliance for Service	181
6.4 Servicing the Oracle Private Cloud Appliance Rack System	182
6.4.1 Powering Down Oracle Private Cloud Appliance (When Required)	182
6.4.2 Service Procedures for Rack System Components	183
6.5 Servicing an Oracle Server X8-2	184
6.5.1 Powering Down Oracle Server X8-2 for Service (When Required)	184
6.5.2 Service Procedures for Oracle Server X8-2 Components	185
6.6 Servicing an Oracle Server X7-2	186
6.6.1 Powering Down Oracle Server X7-2 for Service (When Required)	186
6.6.2 Service Procedures for Oracle Server X7-2 Components	188
6.7 Servicing an Oracle Server X6-2	188
6.7.1 Powering Down Oracle Server X6-2 for Service (When Required)	189
6.7.2 Service Procedures for Oracle Server X6-2 Components	190
6.8 Servicing an Oracle Server X5-2	191
6.8.1 Powering Down Oracle Server X5-2 for Service (When Required)	191
6.8.2 Service Procedures for Oracle Server X5-2 Components	192
6.9 Servicing a Sun Server X4-2	193
6.9.1 Powering Down Sun Server X4-2 for Service (When Required)	193
6.9.2 Service Procedures for Sun Server X4-2 Components	195
6.10 Servicing a Sun Server X3-2	196
6.10.1 Powering Down Sun Server X3-2 for Service (When Required)	196
6.10.2 Service Procedures for Sun Server X3-2 Components	197
6.11 Servicing the Oracle ZFS Storage Appliance ZS7-2	198
6.11.1 Powering Down the Oracle ZFS Storage Appliance ZS7-2 for Service (When Required)	198
6.11.2 Service Procedures for Oracle ZFS Storage Appliance ZS7-2 Components	200
6.12 Servicing the Oracle ZFS Storage Appliance ZS5-ES	201
6.12.1 Powering Down the Oracle ZFS Storage Appliance ZS5-ES for Service (When Required)	201
6.12.2 Service Procedures for Oracle ZFS Storage Appliance ZS5-ES Components	202

6.13 Servicing the Oracle ZFS Storage Appliance ZS3-ES	203
6.13.1 Powering Down the Oracle ZFS Storage Appliance ZS3-ES for Service (When Required)	204
6.13.2 Service Procedures for Oracle ZFS Storage Appliance ZS3-ES Components	206
6.14 Servicing the Sun ZFS Storage Appliance 7320	207
6.14.1 Powering Down the Sun ZFS Storage Appliance 7320 for Service (When Required) ..	207
6.14.2 Service Procedures for Sun ZFS Storage Appliance 7320 Components	208
6.15 Servicing an Oracle Switch ES1-24	209
6.15.1 Powering Down the Oracle Switch ES1-24 for Service (When Required)	209
6.15.2 Service Procedures for Oracle Switch ES1-24 Components	210
6.16 Servicing an NM2-36P Sun Datacenter InfiniBand Expansion Switch	210
6.16.1 Powering Down the NM2-36P Sun Datacenter InfiniBand Expansion Switch for Service (When Required)	210
6.16.2 Service Procedures for NM2-36P Sun Datacenter InfiniBand Expansion Switch Components	211
6.17 Servicing an Oracle Fabric Interconnect F1-15	211
6.17.1 Powering Down the Oracle Fabric Interconnect F1-15 for Service (When Required) ..	211
6.17.2 Service Procedures for Oracle Fabric Interconnect F1-15 Components	212
7 Troubleshooting	215
7.1 Setting the Oracle Private Cloud Appliance Logging Parameters	215
7.2 Adding Proxy Settings for Oracle Private Cloud Appliance Updates	216
7.3 Configuring Data Center Switches for VLAN Traffic	217
7.4 Changing the Oracle VM Agent Password	218
7.5 Running Manual Pre- and Post-Upgrade Checks in Combination with Oracle Private Cloud Appliance Upgrader	218
7.6 Enabling Fibre Channel Connectivity on a Provisioned Appliance	220
7.7 Restoring a Backup After a Password Change	222
7.8 Enabling SNMP Server Monitoring	224
7.9 Using a Custom CA Certificate for SSL Encryption	225
7.9.1 Creating a Keystore	226
7.9.2 Importing a Keystore	227
7.10 Reprovisioning a Compute Node when Provisioning Fails	228
7.11 Deprovisioning and Replacing a Compute Node	229
7.12 Eliminating Time-Out Issues when Provisioning Compute Nodes	230
7.13 Returning Oracle VM Server Pool to Operation After Network Services Restart	231
7.14 Recovering from Tenant Group Configuration Mismatches	232
7.15 Configure Xen CPU Frequency Scaling for Best Performance	233
Index	235

Preface

This document is part of the documentation set for Oracle Private Cloud Appliance (PCA) Release 2.4. All Oracle Private Cloud Appliance product documentation is available at:

<https://docs.oracle.com/en/engineered-systems/private-cloud-appliance/index.html>.

The documentation set consists of the following items:

Oracle Private Cloud Appliance Release Notes

The release notes provide a summary of the new features, changes, fixed bugs and known issues in Oracle Private Cloud Appliance.

Oracle Private Cloud Appliance Licensing Information User Manual

The licensing information user manual provides information about the various product licenses applicable to the use of Oracle Private Cloud Appliance.

Oracle Private Cloud Appliance Installation Guide

The installation guide provides detailed instructions to prepare the installation site and install Oracle Private Cloud Appliance. It also includes the procedures to install additional compute nodes, and to connect and configure external storage components.

Oracle Private Cloud Appliance Safety and Compliance Guide

The safety and compliance guide is a supplemental guide to the safety aspects of Oracle Private Cloud Appliance. It conforms to Compliance Model No. ESY27.

Oracle Private Cloud Appliance Administrator's Guide

The administrator's guide provides instructions for using the management software. It is a comprehensive guide to how to configure, monitor and administer Oracle Private Cloud Appliance.

Oracle Private Cloud Appliance Quick Start Poster

The quick start poster provides a step-by-step description of the hardware installation and initial software configuration of Oracle Private Cloud Appliance. A printed quick start poster is shipped with each Oracle Private Cloud Appliance base rack, and is intended for data center operators and administrators who are new to the product.

The quick start poster is also available in the documentation set as an HTML guide, which contains alternate text for ADA 508 compliance.

Oracle Private Cloud Appliance Expansion Node Setup Poster

The expansion node setup poster provides a step-by-step description of the installation procedure for an Oracle Private Cloud Appliance expansion node. A printed expansion node setup poster is shipped with each Oracle Private Cloud Appliance expansion node.

The expansion node setup poster is also available in the documentation set as an HTML guide, which contains alternate text for ADA 508 compliance.

Audience

The Oracle Private Cloud Appliance documentation is written for technicians, authorized service providers, data center operators and system administrators who want to install, configure and maintain a private cloud

environment in order to deploy virtual machines for users. It is assumed that readers have experience installing and troubleshooting hardware, are familiar with web and virtualization technologies and have a general understanding of operating systems such as UNIX (including Linux) and Windows.

The Oracle Private Cloud Appliance makes use of Oracle Linux and Oracle Solaris operating systems within its component configuration. It is advisable that administrators have experience of these operating systems at the very least. Oracle Private Cloud Appliance is capable of running virtual machines with a variety of operating systems including Oracle Solaris and other UNIXes, Linux and Microsoft Windows. The selection of operating systems deployed in guests on Oracle Private Cloud Appliance determines the requirements of your administrative knowledge.

Related Documentation

Additional Oracle components may be included with Oracle Private Cloud Appliance depending on configuration. The documentation for such additional components is available as follows:



Note

If your appliance contains components that are not mentioned below, please consult the related documentation list for [Oracle Private Cloud Appliance Release 2.3](#).

- Oracle Rack Cabinet 1242
https://docs.oracle.com/cd/E85660_01/index.html
- Oracle Server X8-2
https://docs.oracle.com/cd/E93359_01/index.html
- Oracle Server X7-2
https://docs.oracle.com/cd/E72435_01/index.html
- Oracle Server X6-2
https://docs.oracle.com/cd/E62159_01/index.html
- Oracle Server X5-2
https://docs.oracle.com/cd/E41059_01/index.html
- Oracle ZFS Storage Appliance ZS7-2
https://docs.oracle.com/cd/F13758_01/index.html
- Oracle ZFS Storage Appliance ZS5-ES
https://docs.oracle.com/cd/E59597_01/index.html
- Oracle Integrated Lights Out Manager (ILOM)
https://docs.oracle.com/cd/E81115_01/index.html
- Oracle Switch ES1-24
https://docs.oracle.com/cd/E39109_01/index.html

- NM2-36P Sun Datacenter InfiniBand Expansion Switch
https://docs.oracle.com/cd/E76424_01/index.html
- Oracle Fabric Interconnect F1-15
https://docs.oracle.com/cd/E38500_01/index.html
- Oracle VM
https://docs.oracle.com/cd/E64076_01/index.html
- Oracle Enterprise Manager Plug-in
<https://docs.oracle.com/en/enterprise-manager/cloud-control/enterprise-manager-cloud-control/13.3.1/empca/index.html>

Feedback

Provide feedback about this documentation at:

<http://www.oracle.com/goto/docfeedback>

Conventions

The following text conventions are used in this document:

Convention	Meaning
boldface	Boldface type indicates graphical user interface elements associated with an action, or terms defined in text or the glossary.
<i>italic</i>	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
<code>monospace</code>	Monospace type indicates commands within a paragraph, URLs, code in examples, text that appears on the screen, or text that you enter.

Document Revision

Document generated on: 2020-03-18 (revision: 1962)

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at

<https://www.oracle.com/corporate/accessibility/>.

Access to Oracle Support

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit

<https://www.oracle.com/corporate/accessibility/learning-support.html#support-tab>.

Chapter 1 Concept, Architecture and Life Cycle of Oracle Private Cloud Appliance

Table of Contents

- 1.1 What is Oracle Private Cloud Appliance 1
- 1.2 Hardware Components 2
 - 1.2.1 Management Nodes 4
 - 1.2.2 Compute Nodes 5
 - 1.2.3 Storage Appliance 5
 - 1.2.4 Network Infrastructure 8
- 1.3 Software Components 12
 - 1.3.1 Oracle Private Cloud Appliance Dashboard 12
 - 1.3.2 Password Manager (Wallet) 13
 - 1.3.3 Oracle VM Manager 13
 - 1.3.4 Operating Systems 13
 - 1.3.5 Databases 13
 - 1.3.6 Oracle Private Cloud Appliance Management Software 17
 - 1.3.7 Oracle Private Cloud Appliance Diagnostics Tool 18
- 1.4 Provisioning and Orchestration 19
 - 1.4.1 Appliance Management Initialization 19
 - 1.4.2 Compute Node Discovery and Provisioning 19
 - 1.4.3 Server Pool Readiness 20
- 1.5 High Availability 21
- 1.6 Oracle Private Cloud Appliance Backup 22
- 1.7 Oracle Private Cloud Appliance Upgrader 23

This chapter describes what Oracle Private Cloud Appliance is, which hardware and software it consists of, and how it is deployed as a virtualization platform.

1.1 What is Oracle Private Cloud Appliance

Responding to the Cloud Challenges

Cloud architectures and virtualization solutions have become highly sophisticated and complex to implement. They require a skill set that no single administrator has had to master in traditional data centers: system hardware, operating systems, network administration, storage management, applications. Without expertise in every single one of those domains, an administrator cannot take full advantage of the features and benefits of virtualization technology. This often leads to poor implementations with sub-optimal performance and reliability, which impairs the flexibility of a business.

Aside from the risks created by technical complexity and lack of expertise, companies also suffer from an inability to deploy new infrastructure quickly enough to suit their business needs. The administration involved in the deployment of new systems, and the time and effort to configure these systems, can amount to weeks. Provisioning new applications into flexible virtualized environments, in a fraction of the time required for physical deployments, generates substantial financial benefits.

Fast Deployment of Converged Infrastructure

Oracle Private Cloud Appliance is an offering that industry analysts refer to as a *Converged Infrastructure Appliance*: an infrastructure solution in the form of a hardware appliance that comes from the factory pre-

configured. It enables the operation of the entire system as a single unit, not a series of individual servers, network hardware and storage providers. Installation, configuration, high availability, expansion and upgrading are automated and orchestrated to an optimal degree. Within a few hours after power-on, the appliance is ready to create virtual servers. Virtual servers are commonly deployed from virtual appliances, in the form of Oracle VM templates (individual pre-configured VMs) and assemblies (interconnected groups of pre-configured VMs).

Modular Implementation of a Complete Stack


With Oracle Private Cloud Appliance, Oracle offers a unique full stack of hardware, software, virtualization technology and rapid application deployment through virtual appliances. All this is packaged in a single modular and extensible product. The minimum configuration consists of a base rack with infrastructure components, a pair of management nodes, and two compute nodes. This configuration can be extended by one compute node at a time. All rack units, whether populated or not, are pre-cabled and pre-configured at the factory in order to facilitate the installation of expansion compute nodes on-site at a later time.

Ease of Use

The primary value proposition of Oracle Private Cloud Appliance is the integration of components and resources for the purpose of ease of use and rapid deployment. It should be considered a general purpose solution in the sense that it supports the widest variety of operating systems, including Windows, and any application they might host. Customers can attach their existing storage or connect new storage solutions from Oracle or third parties.

1.2 Hardware Components

The current Oracle Private Cloud Appliance hardware platform, with factory-installed Controller Software Release 2.4.x, consists of an Oracle Rack Cabinet 1242 base, populated with the hardware components identified in [Figure 1.1](#). Previous generations of hardware components continue to be supported by the latest Controller Software, as described below.

Figure 1.1 Components of an Oracle Private Cloud Appliance Rack**Table 1.1 Figure Legend**

Item	Quantity	Description
A	2	Oracle ZFS Storage Appliance ZS7-2 controller server
B	2	Oracle Server X8-2, used as management nodes
C	2-25	Oracle Server X8-2, used as virtualization compute nodes (Due to the power requirements of the Oracle Server X8-2, if the appliance is equipped with 22kVA PDUs, the maximum number of compute nodes is 22. With 15KVA PDUs the maximum is 13 compute nodes.)
D	2	Cisco Nexus 9336C-FX2 Switch, used as leaf/data switches
E	1	Oracle ZFS Storage Appliance ZS7-2 disk shelf
F	1	Cisco Nexus 9348GC-FXP Switch
G	2	Cisco Nexus 9336C-FX2 Switch, used as spine switches

Support for Previous Generations of Hardware Components

The latest version of the Oracle Private Cloud Appliance Controller Software continues to support all earlier configurations of the hardware platform. These may include the following components:

Table 1.2 Supported Hardware

Component Type	Component Name and Minimum Software Version
Management Nodes	<ul style="list-style-type: none"> Oracle Server X5-2 (release 2.0.3 or newer)

Component Type	Component Name and Minimum Software Version
Compute Nodes	<ul style="list-style-type: none"> • Sun Server X4-2 (release 1.1.3 or newer) • Sun Server X3-2 (since initial release)
	<ul style="list-style-type: none"> • Oracle Server X7-2 (release 2.3.2 or newer) • Oracle Server X6-2 (release 2.2.1 or newer) • Oracle Server X5-2 (release 2.0.3 or newer) • Sun Server X4-2 (release 1.1.3 or newer) • Sun Server X3-2 (since initial release)
Storage Appliance	<ul style="list-style-type: none"> • Oracle ZFS Storage Appliance ZS5-ES (release 2.3.3 or newer) • Oracle ZFS Storage Appliance ZS3-ES (release 1.1.3 or newer) • Sun ZFS Storage Appliance 7320 (since initial release)
InfiniBand Network Hardware	<ul style="list-style-type: none"> • Oracle Fabric Interconnect F1-15 (since initial release) • NM2-36P Sun Datacenter InfiniBand Expansion Switch (since initial release)
Internal Management Switch	<ul style="list-style-type: none"> • Oracle Switch ES1-24 (since initial release)

1.2.1 Management Nodes

At the heart of each Oracle Private Cloud Appliance installation is a pair of management nodes. They are installed in rack units 5 and 6 and form a cluster in active/standby configuration for high availability: both servers are capable of running the same services and have equal access to the system configuration, but one operates as the master while the other is ready to take over the master functions in case a failure occurs. The master management node runs the full set of services required, while the standby management node runs a subset of services until it is promoted to the master role. The master role is determined at boot through OCFS2 Distributed Lock Management on an iSCSI LUN, which both management nodes share on the ZFS Storage Appliance installed inside the rack. Because rack units are numbered from the bottom up, and the bottom four are occupied by components of the ZFS Storage Appliance, the master management node is typically the server in rack unit 5. It is the only server that must be powered on by the administrator in the entire process to bring the appliance online.

For details about how high availability is achieved with Oracle Private Cloud Appliance, refer to [Section 1.5, “High Availability”](#).

When you power on the Oracle Private Cloud Appliance for the first time, you can change the factory default IP configuration of the management node cluster, so that it can be easily reached from your data center network. The management nodes share a Virtual IP, where the management web interface can be accessed. This virtual IP is assigned to whichever server has the *master* role at any given time. During system initialization, after the management cluster is set up successfully, the master management node loads a number of Oracle Linux services, in addition to Oracle VM and its associated MySQL database – including network, sshd, ntpd, iscsi initiator, dhcpd – to orchestrate the provisioning of all system components. During provisioning, all networking and storage is configured, and all compute nodes are discovered, installed and added to an Oracle VM server pool. All provisioning configurations are preloaded at the factory and should not be modified by the customer.

For details about the provisioning process, refer to [Section 1.4, “Provisioning and Orchestration”](#).

1.2.2 Compute Nodes

The compute nodes in the Oracle Private Cloud Appliance constitute the virtualization platform. The compute nodes provide the processing power and memory capacity for the virtual servers they host. The entire provisioning process is orchestrated by the management nodes: compute nodes are installed with Oracle VM Server 3.4.x and additional packages for Software Defined Networking. When provisioning is complete, the Oracle Private Cloud Appliance Controller Software expects all compute nodes in the same rack to be part of the same Oracle VM server pool.

For hardware configuration details of the compute nodes, refer to [Server Components](#) in the Oracle Private Cloud Appliance Installation Guide.

The Oracle Private Cloud Appliance Dashboard allows the administrator to monitor the health and status of the compute nodes, as well as all other rack components, and perform certain system operations. The virtual infrastructure is configured and managed with Oracle VM Manager.

The Oracle Private Cloud Appliance offers modular compute capacity that can be increased according to business needs. The minimum configuration of the base rack contains just two compute nodes, but it can be expanded by one node at a time up to 25 compute nodes. Apart from the hardware installation, adding compute nodes requires no intervention by the administrator. New nodes are discovered, powered on, installed and provisioned automatically by the master management node. The additional compute nodes are integrated into the existing configuration and, as a result, the Oracle VM server pool offers increased capacity for more or larger virtual machines.

As a further expansion option, the Oracle Server X8-2 compute nodes can be ordered with pre-installed fibre channel cards, or equipped with fibre channel cards after installation. Once these compute nodes are integrated in the Oracle Private Cloud Appliance environment, the fibre channel HBAs can connect to standard FC switches and storage hardware in your data center. External FC storage configuration is managed through Oracle VM Manager. For more information, refer to the [Fibre Channel Storage Attached Network](#) section of the *Oracle VM Concepts Guide*.



Caution

When using expansion nodes containing fibre channel cards in a system with InfiniBand-based network architecture, the vHBAs **must** be disabled on those compute nodes.

Because of the diversity of possible virtualization scenarios it is difficult to quantify the compute capacity as a number of virtual machines. For sizing guidelines, refer to the chapter entitled [Configuration Maximums](#) in the *Oracle Private Cloud Appliance Release Notes*.

1.2.3 Storage Appliance

The Oracle Private Cloud Appliance Controller Software continues to provide support for previous generations of the ZFS Storage Appliance installed in the base rack. However, there are functional differences between the Oracle ZFS Storage Appliance ZS7-2, which is part of systems with an Ethernet-based network architecture, and the previous models of the ZFS Storage Appliance, which are part of systems with an InfiniBand-based network architecture. For clarity, this section describes the different storage appliances separately.

1.2.3.1 Oracle ZFS Storage Appliance ZS7-2

The Oracle ZFS Storage Appliance ZS7-2, which consists of two controller servers installed at the bottom of the appliance rack and disk shelf about halfway up, fulfills the role of 'system disk' for the entire appliance. It is crucial in providing storage space for the Oracle Private Cloud Appliance software.

A portion of the disk space, 3TB by default, is made available for customer use and is sufficient for an Oracle VM storage repository with several virtual machines, templates and assemblies. The remaining part of approximately 100TB in total disk space can also be configured as a storage repository for virtual machine resources. Further capacity extension with external storage is also possible.

The hardware configuration of the Oracle ZFS Storage Appliance ZS7-2 is as follows:

- Two clustered storage heads with two 14TB hard disks each
- One fully populated disk chassis with twenty 14TB hard disks
- Four cache disks installed in the disk shelf: 2x 200GB SSD and 2x 7.68TB SSD
- RAID-1 configuration, for optimum data protection, with a total usable space of approximately 100TB

The storage appliance is connected to the management subnet ([192.168.4.0/24](#)) and the storage subnet ([192.168.40.0/24](#)). Both heads form a cluster in active-passive configuration to guarantee continuation of service in the event that one storage head should fail. The storage heads share a single IP in the storage subnet, but both have an individual management IP address for convenient maintenance access. The RAID-1 storage pool contains two projects, named [OVCA](#) and [OVM](#).

The [OVCA](#) project contains all LUNs and file systems used by the Oracle Private Cloud Appliance software:

- LUNs
 - [Locks](#) (12GB) – to be used exclusively for cluster locking on the two management nodes
 - [Manager](#) (200GB) – to be used exclusively as an additional file system on both management nodes
- File systems:
 - [MGMT_ROOT](#) – to be used for storage of all files specific to the Oracle Private Cloud Appliance
 - [Database](#) – placeholder file system for databases
 - [Incoming](#) (20GB) – to be used for FTP file transfers, primarily for Oracle Private Cloud Appliance component backups
 - [Templates](#) – placeholder file system for future use
 - [User](#) – placeholder file system for future use
 - [Yum](#) – to be used for system package updates

The [OVM](#) project contains all LUNs and file systems used by Oracle VM:

- LUNs
 - [iscsi_repository1](#) (3TB) – to be used as Oracle VM storage repository
 - [iscsi_serverpool1](#) (12GB) – to be used as server pool file system for the Oracle VM clustered server pool
- File systems:
 - [nfs_repository1](#) (3TB) – used by [kdump](#); not available for customer use
 - [nfs_serverpool1](#) (12GB) – to be used as server pool file system for the Oracle VM clustered server pool in case NFS is preferred over iSCSI



Caution

If the internal ZFS Storage Appliance contains customer-created LUNs, make sure they are not mapped to the default initiator group. See [Customer Created LUNs Are Mapped to the Wrong Initiator Group](#) in the Oracle Private Cloud Appliance Release Notes.

In addition to offering storage, the ZFS storage appliance also runs the `xinetd` and `ftpsd` services. These complement the Oracle Linux services on the master management node in order to orchestrate the provisioning of all Oracle Private Cloud Appliance system components.

1.2.3.2 Oracle ZFS Storage Appliance ZS5-ES and Earlier Models

The Oracle ZFS Storage Appliance ZS5-ES installed at the bottom of the appliance rack should be considered a 'system disk' for the entire appliance. Its main purpose is to provide storage space for the Oracle Private Cloud Appliance software. A portion of the disk space is made available for customer use and is sufficient for an Oracle VM storage repository with a limited number of virtual machines, templates and assemblies.

The hardware configuration of the Oracle ZFS Storage Appliance ZS5-ES is as follows:

- Two clustered storage heads with two 3.2TB SSDs each, used exclusively for cache
- One fully populated disk chassis with twenty 1.2TB 10000 RPM SAS hard disks
- RAID-Z2 configuration, for best balance between performance and data protection, with a total usable space of approximately 15TB



Note

Oracle Private Cloud Appliance base racks shipped prior to software release 2.3.3 use a Sun ZFS Storage Appliance 7320 or Oracle ZFS Storage Appliance ZS3-ES. Those systems may be upgraded to a newer software stack, which continues to provide support for each Oracle Private Cloud Appliance storage configuration. The newer storage appliance offers the same functionality and configuration, with modernized hardware and thus better performance.

The storage appliance is connected to the management subnet ([192.168.4.0/24](#)) and the InfiniBand (IPoIB) storage subnet ([192.168.40.0/24](#)). Both heads form a cluster in active-passive configuration to guarantee continuation of service in the event that one storage head should fail. The storage heads share a single IP in the storage subnet, but both have an individual management IP address for convenient maintenance access. The RAID-Z2 storage pool contains two projects, named [OVCA](#) and [OVM](#) .

The [OVCA](#) project contains all LUNs and file systems used by the Oracle Private Cloud Appliance software:

- LUNs
 - [Locks](#) (12GB) – to be used exclusively for cluster locking on the two management nodes
 - [Manager](#) (200GB) – to be used exclusively as an additional file system on both management nodes
- File systems:
 - [MGMT_ROOT](#) – to be used for storage of all files specific to the Oracle Private Cloud Appliance
 - [Database](#) – placeholder file system for databases

- `Incoming` (20GB) – to be used for FTP file transfers, primarily for Oracle Private Cloud Appliance component backups
- `Templates` – placeholder file system for future use
- `User` – placeholder file system for future use
- `Yum` – to be used for system package updates

The `OVM` project contains all LUNs and file systems used by Oracle VM:

- LUNs
 - `iscsi_repository1` (300GB) – to be used as Oracle VM storage repository
 - `iscsi_serverpool1` (12GB) – to be used as server pool file system for the Oracle VM clustered server pool
- File systems:
 - `nfs_repository1` (300GB) – to be used as Oracle VM storage repository in case NFS is preferred over iSCSI
 - `nfs_serverpool1` (12GB) – to be used as server pool file system for the Oracle VM clustered server pool in case NFS is preferred over iSCSI



Caution

If the internal ZFS Storage Appliance contains customer-created LUNs, make sure they are not mapped to the default initiator group. See [Customer Created LUNs Are Mapped to the Wrong Initiator Group](#) in the Oracle Private Cloud Appliance Release Notes.

In addition to offering storage, the ZFS storage appliance also runs the `xinetd` and `tftpd` services. These complement the Oracle Linux services on the master management node in order to orchestrate the provisioning of all Oracle Private Cloud Appliance system components.

1.2.4 Network Infrastructure

For network connectivity, Oracle Private Cloud Appliance relies on a physical layer that provides the necessary high-availability, bandwidth and speed. On top of this, several different virtual networks are optimally configured for different types of data traffic. Only the internal administration network is truly physical; the appliance data connectivity uses Software Defined Networking (SDN). The appliance rack contains redundant network hardware components, which are pre-cabled at the factory to help ensure continuity of service in case a failure should occur.

Depending on the exact hardware configuration of your appliance, the physical network layer is either high-speed Ethernet or InfiniBand. In this section, both network architectures are described separately in more detail.

1.2.4.1 Ethernet-Based Network Architecture

Oracle Private Cloud Appliance with Ethernet-based network architecture relies on redundant physical high-speed Ethernet connectivity.

Administration Network

The administration network provides internal access to the management interfaces of all appliance components. These have Ethernet connections to the Cisco Nexus 9348GC-FXP Switch, and all have a predefined IP address in the [192.168.4.0/24](#) range. In addition, all management and compute nodes have a second IP address in this range, which is used for Oracle Integrated Lights Out Manager (ILOM) connectivity.

While the appliance is initializing, the data network is not accessible, which means that the internal administration network is temporarily the only way to connect to the system. Therefore, the administrator should connect a workstation to the reserved Ethernet port 48 in the Cisco Nexus 9348GC-FXP Switch, and assign the fixed IP address [192.168.4.254](#) to the workstation. From this workstation, the administrator opens a browser connection to the web server on the master management node at <https://192.168.4.216>, in order to monitor the initialization process and perform the initial configuration steps when the appliance is powered on for the first time.

Data Network

The appliance data connectivity is built on redundant Cisco Nexus 9336C-FX2 Switches in a leaf-spine design. In this two-layer design, the leaf switches interconnect the rack hardware components, while the spine switches form the backbone of the network and perform routing tasks. Each leaf switch is connected to all the spine switches, which are also interconnected. The main benefits of this network architecture are extensibility and path optimization. An Oracle Private Cloud Appliance rack contains two leaf and two spine switches.

The Cisco Nexus 9336C-FX2 Switch offers a maximum throughput of 100Gbit per port. The spine switches use 5 interlinks (500Gbit); the leaf switches use 2 interlinks (200Gbit) and 2x2 crosslinks to the spines. Each compute node is connected to both leaf switches in the rack, through the `bond1` interface that consists of two 100Gbit Ethernet ports in link aggregation mode. The two storage controllers are connected to the spine switches using 4x40Gbit connections.

For external connectivity, 5 ports are reserved on each spine switch. Four ports are available for custom network configurations; one port is required for the default uplink. This **default external uplink** requires that port 5 on both spine switches is split using a QSFP+-to-SFP+ four way splitter or breakout cable. Two of those four 10GbE SFP+ breakout ports per spine switch, ports 5/1 and 5/2, must be connected to a pair of next-level data center switches, also called top-of-rack or ToR switches.

Software Defined Networking

While the physical data network described above allows the data packets to be transferred, the true connectivity is implemented through Software Defined Networking (SDN). Using VxLAN encapsulation and VLAN tagging, thousands of virtual networks can be deployed, providing segregated data exchange. Traffic can be internal between resources within the appliance environment, or external to network storage, applications, or other resources in the data center or on the internet. SDN maintains the traffic separation of hard-wired connections, and adds better performance and dynamic (re-)allocation. From the perspective of the customer network, the use of VxLANs in Oracle Private Cloud Appliance is transparent: encapsulation and de-encapsulation take place internally, without modifying inbound or outbound data packets. In other words, this design extends customer networking, tagged or untagged, into the virtualized environment hosted by the appliance.

During the initialization process of the Oracle Private Cloud Appliance, several essential default networks are configured:

- The **Internal Storage Network** is a redundant 40Gbit Ethernet connection from the spine switches to the ZFS storage appliance. All four storage controller interfaces are bonded using LACP into one datalink. Management and compute nodes can reach the internal storage over the [192.168.40.0/21](#) subnet on VLAN 3093. This network also fulfills the heartbeat function for the clustered Oracle VM server pool.

- The **Internal Management Network** provides connectivity between the management nodes and compute nodes in the subnet [192.168.32.0/21](#) on VLAN 3092. It is used for all network traffic inherent to Oracle VM Manager, Oracle VM Server and the Oracle VM Agents.
- The **Internal Underlay Network** provides the infrastructure layer for data traffic between compute nodes. It uses the subnet [192.168.64.0/21](#) on VLAN 3091. On top of the internal underlay network, internal VxLAN overlay networks are built to enable virtual machine connectivity where only internal access is required.

One such internal VxLAN is configured in advance: the *default internal VM network*, to which all compute nodes are connected with their `vx2` interface. Untagged traffic is supported by default over this network. Customers can add VLANs of their choice to the Oracle VM network configuration, and define the subnet(s) appropriate for IP address assignment at the virtual machine level.

- The **External Underlay Network** provides the infrastructure layer for data traffic between Oracle Private Cloud Appliance and the data center network. It uses the subnet [192.168.72.0/21](#) on VLAN 3090. On top of the external underlay network, VxLAN overlay networks with external access are built to enable public connectivity for the physical nodes and all the virtual machines they host.

One such public VxLAN is configured in advance: the *default external network*, to which all compute nodes and management nodes are connected with their `vx13040` interface. Both tagged and untagged traffic are supported by default over this network. Customers can add VLANs of their choice to the Oracle VM network configuration, and define the subnet(s) appropriate for IP address assignment at the virtual machine level.

The default external network also provides access to the management nodes from the data center network and allows the management nodes to run a number of system services. The management node external network settings are configurable through the [Network Settings tab](#) in the Oracle Private Cloud Appliance Dashboard. If this network is a VLAN, its ID or tag must be configured in the Network Setup tab of the Dashboard.

For the appliance default networking to be configured successfully, the default external uplink must be in place before the initialization of the appliance begins. At the end of the initialization process, the administrator assigns three reserved IP addresses from the data center (public) network range to the management node cluster of the Oracle Private Cloud Appliance: one for each management node, and an additional Virtual IP shared by the clustered nodes. From this point forward, the Virtual IP is used to connect to the master management node's web server, which hosts both the Oracle Private Cloud Appliance Dashboard and the Oracle VM Manager web interface.



Caution

It is critical that **both** spine Cisco Nexus 9336C-FX2 Switches have **two** 10GbE connections each to a pair of next-level data center switches. For this purpose, a 4-way breakout cable must be attached to port 5 of each spine switch, and 10GbE breakout ports 5/1 and 5/2 must be used as uplinks. Note that ports 5/3 and 5/4 remain unused.

This outbound cabling between the spine switches and the data center network should be crossed or meshed, to ensure optimal continuity of service.

1.2.4.2 InfiniBand-Based Network Architecture

Oracle Private Cloud Appliance with InfiniBand-based network architecture relies on a physical InfiniBand network fabric, with additional Ethernet connectivity for internal management communication.

Ethernet

The Ethernet network relies on two interconnected Oracle Switch ES1-24 switches, to which all other rack components are connected with CAT6 Ethernet cables. This network serves as the appliance management network, in which every component has a predefined IP address in the `192.168.4.0/24` range. In addition, all management and compute nodes have a second IP address in this range, which is used for Oracle Integrated Lights Out Manager (ILOM) connectivity.

While the appliance is initializing, the InfiniBand fabric is not accessible, which means that the management network is the only way to connect to the system. Therefore, the administrator should connect a workstation to the available Ethernet port 19 in one of the Oracle Switch ES1-24 switches, and assign the fixed IP address `192.168.4.254` to the workstation. From this workstation, the administrator opens a browser connection to the web server on the master management node at <http://192.168.4.216>, in order to monitor the initialization process and perform the initial configuration steps when the appliance is powered on for the first time.

InfiniBand

The Oracle Private Cloud Appliance rack contains two NM2-36P Sun Datacenter InfiniBand Expansion Switches. These redundant switches have redundant cable connections to both InfiniBand ports in each management node, compute node and storage head. Both InfiniBand switches, in turn, have redundant cable connections to both Fabric Interconnects in the rack. All these components combine to form a physical InfiniBand backplane with a 40Gbit (Quad Data Rate) bandwidth.

When the appliance initialization is complete, all necessary Oracle Private Cloud Appliance software packages, including host drivers and InfiniBand kernel modules, have been installed and configured on each component. At this point, the system is capable of using software defined networking (SDN) configured on top of the physical InfiniBand fabric. SDN is implemented through the Fabric Interconnects.

Fabric Interconnect

All Oracle Private Cloud Appliance network connectivity is managed through the Fabric Interconnects. Data is transferred across the physical InfiniBand fabric, but connectivity is implemented in the form of Software Defined Networks (SDN), which are sometimes referred to as 'clouds'. The physical InfiniBand backplane is capable of hosting thousands of virtual networks. These Private Virtual Interconnects (PVI) dynamically connect virtual machines and bare metal servers to networks, storage and other virtual machines, while maintaining the traffic separation of hard-wired connections and surpassing their performance.

During the initialization process of the Oracle Private Cloud Appliance, five essential networks, four of which are SDNs, are configured: a storage network, an Oracle VM management network, a management Ethernet network, and two virtual machine networks. Tagged and untagged virtual machine traffic is supported. VLANs can be constructed using virtual interfaces on top of the existing bond interfaces of the compute nodes.

- The **storage network**, technically not software-defined, is a bonded IPoIB connection between the management nodes and the ZFS storage appliance, and uses the `192.168.40.0/24` subnet. This network also fulfills the heartbeat function for the clustered Oracle VM server pool. DHCP ensures that compute nodes are assigned an IP address in this subnet.
- The **Oracle VM management network** is a PVI that connects the management nodes and compute nodes in the `192.168.140.0/24` subnet. It is used for all network traffic inherent to Oracle VM Manager, Oracle VM Server and the Oracle VM Agents.
- The **management Ethernet network** is a bonded Ethernet connection between the management nodes. The primary function of this network is to provide access to the management nodes from the data center network, and enable the management nodes to run a number of system services. Since all compute

nodes are also connected to this network, Oracle VM can use it for virtual machine connectivity, with access to and from the data center network. The management node external network settings are configurable through the [Network Settings tab](#) in the Oracle Private Cloud Appliance Dashboard. If this network is a VLAN, its ID or tag must be configured in the Network Setup tab of the Dashboard.

- The **public virtual machine network** is a bonded Ethernet connection between the compute nodes. Oracle VM uses this network for virtual machine connectivity, where external access is required. Untagged traffic is supported by default over this network. Customers can add their own VLANs to the Oracle VM network configuration, and define the subnet(s) appropriate for IP address assignment at the virtual machine level. For external connectivity, the next-level data center switches must be configured to accept your tagged VLAN traffic.
- The **private virtual machine network** is a bonded Ethernet connection between the compute nodes. Oracle VM uses this network for virtual machine connectivity, where only internal access is required. Untagged traffic is supported by default over this network. Customers can add VLANs of their choice to the Oracle VM network configuration, and define the subnet(s) appropriate for IP address assignment at the virtual machine level.

Finally, the Fabric Interconnects also manage the physical public network connectivity of the Oracle Private Cloud Appliance. Two 10GbE ports on each Fabric Interconnect must be connected to redundant next-level data center switches. At the end of the initialization process, the administrator assigns three reserved IP addresses from the data center (public) network range to the management node cluster of the Oracle Private Cloud Appliance: one for each management node, and an additional Virtual IP shared by the clustered nodes. From this point forward, the Virtual IP is used to connect to the master management node's web server, which hosts both the Oracle Private Cloud Appliance Dashboard and the Oracle VM Manager web interface.



Caution

It is critical that **both** Fabric Interconnects have **two** 10GbE connections each to a pair of next-level data center switches. This configuration with four cable connections provides redundancy and load splitting at the level of the Fabric Interconnects, the 10GbE ports and the data center switches. This outbound cabling should not be crossed or meshed, because the internal connections to the pair of Fabric Interconnects are already configured that way. The cabling pattern plays a key role in the continuation of service during failover scenarios involving Fabric Interconnect outages and other components.

1.3 Software Components

This section describes the main software components the Oracle Private Cloud Appliance uses for operation and configuration.

1.3.1 Oracle Private Cloud Appliance Dashboard

The Oracle Private Cloud Appliance provides its own web-based graphical user interface that can be used to perform a variety of administrative tasks specific to the appliance. The Oracle Private Cloud Appliance Dashboard is an Oracle JET application that is available through the active management node.

Use the Dashboard to perform the following tasks:

- Appliance system monitoring and component identification
- Initial configuration of management node networking data
- Resetting of the global password for Oracle Private Cloud Appliance configuration components

The Oracle Private Cloud Appliance Dashboard is described in detail in [Chapter 2, *Monitoring and Managing Oracle Private Cloud Appliance*](#).

1.3.2 Password Manager (Wallet)

All components of the Oracle Private Cloud Appliance have administrator accounts with a default password. After applying your data center network settings through the Oracle Private Cloud Appliance Dashboard, it is recommended that you modify the default appliance password. The Authentication tab allows you to set a new password, which is applied to the main system configuration components. You can set a new password for all listed components at once or for a selection only.

Passwords for all accounts on all components are stored in a global Wallet, secured with 512-bit encryption. To update the password entries, you use either the Oracle Private Cloud Appliance Dashboard or the Command Line Interface. For details, see [Section 2.8, “Authentication”](#).

1.3.3 Oracle VM Manager

All virtual machine management tasks are performed within Oracle VM Manager, a WebLogic application that is installed on each of the management nodes and which provides a web-based management user interface and a command line interface that allows you to manage your Oracle VM infrastructure within the Oracle Private Cloud Appliance.

Oracle VM Manager is comprised of the following software components:

- **Oracle VM Manager application:** provided as an Oracle WebLogic Server domain and container.
- **Oracle WebLogic Server 12c:** including Application Development Framework (ADF) Release 12c, used to host and run the Oracle VM Manager application
- **MySQL 5.6 Enterprise Edition Server:** for the exclusive use of the Oracle VM Manager application as a management repository and installed on the Database file system hosted on the ZFS storage appliance.

Administration of virtual machines is performed using the Oracle VM Manager web user interface, as described in [Chapter 5, *Managing the Oracle VM Virtual Infrastructure*](#). While it is possible to use the command line interface provided with Oracle VM Manager, this is considered an advanced activity that should only be performed with a thorough understanding of the limitations of Oracle VM Manager running in the context of an Oracle Private Cloud Appliance.

1.3.4 Operating Systems

Hardware components of the Oracle Private Cloud Appliance run their own operating systems:

- Management Nodes: Oracle Linux 6 with UEK R4
- Compute Nodes: Oracle VM Server 3.4.6
- Oracle ZFS Storage Appliance: Oracle Solaris 11

All other components run a particular revision of their respective firmware. All operating software has been selected and developed to work together as part of the Oracle Private Cloud Appliance. When an update is released, the appropriate versions of all software components are bundled. When a new software release is activated, all component operating software is updated accordingly. You should not attempt to update individual components unless Oracle explicitly instructs you to.

1.3.5 Databases

The Oracle Private Cloud Appliance uses a number of databases to track system states, handle configuration and provisioning, and for Oracle VM Manager. All databases are stored on the ZFS storage appliance, and are exported via an NFS file system. The databases are accessible to each management node to ensure high availability.



Caution


Databases must never be edited manually. The appliance configuration depends on them, so manipulations are likely to break functionality.

The following table lists the different databases used by the Oracle Private Cloud Appliance.

Table 1.3 Oracle Private Cloud Appliance Databases

Item	Description
Oracle Private Cloud Appliance Node Database	<p>Contains information on every compute node and management node in the rack, including the state used to drive the provisioning of compute nodes and data required to handle software updates.</p> <p>Type: BerkeleyDB</p> <p>Location: <code>MGMT_ROOT/db/node</code> on the ZFS, accessible via <code>/nfs/shared_storage/db/node</code> on each management node</p>
Oracle Private Cloud Appliance Inventory Database	<p>Contains information on all hardware components appearing in the management network <code>192.168.4.0/24</code>. Components include the management and compute nodes but also switches, fabric interconnects, ZFS storage appliance and PDUs. The stored information includes IP addresses and host names, pingable status, when a component was last seen online, etc. This database is queried regularly by a number of Oracle Private Cloud Appliance services.</p> <p>Type: BerkeleyDB</p> <p>Location: <code>MGMT_ROOT/db/inventory</code> on the ZFS, accessible via <code>/nfs/shared_storage/db/inventory</code> on each management node</p>
Oracle Private Cloud Appliance Netbundle Database	<p>Predefines Ethernet and bond device names for all possible networks that can be configured throughout the system, and which are allocated dynamically.</p> <p>Type: BerkeleyDB</p> <p>Location: <code>MGMT_ROOT/db/netbundle</code> on the ZFS, accessible via <code>/nfs/shared_storage/db/netbundle</code> on each management node</p>
Oracle Private Cloud Appliance DHCP Database	<p>Contains information on the assignment of DHCP addresses to newly detected compute nodes.</p> <p>Type: BerkeleyDB</p> <p>Location: <code>MGMT_ROOT/db/dhcp</code> on the ZFS, accessible via <code>/nfs/shared_storage/db/dhcp</code> on each management node</p>
Cisco Data Network Database	<p>Contains information on the networks configured for traffic through the spine switches, and the interfaces participating in the networks.</p> <p>Used only on systems with Ethernet-based network architecture.</p>

Item	Description
Cisco Management Switch Ports Database	<p>Type: BerkeleyDB</p> <p>Location: <code>MGMT_ROOT/db/cisco_data_network_db</code> on the ZFS, accessible via <code>/nfs/shared_storage/db/cisco_data_network_db</code> on each management node</p> <p>Defines the factory-configured map of Cisco Nexus 9348GC-FXP Switch ports to the rack unit or element to which that port is connected. It is used to map switch ports to machine names.</p> <p>Used only on systems with Ethernet-based network architecture.</p> <p>Type: BerkeleyDB</p> <p>Location: <code>MGMT_ROOT/db/cisco_ports</code> on the ZFS, accessible via <code>/nfs/shared_storage/db/cisco_ports</code> on each management node</p>
Oracle Fabric Interconnect Database	<p>Contains IP and host name data for the Oracle Fabric Interconnect F1-15s.</p> <p>Used only on systems with InfiniBand-based network architecture.</p> <p>Type: BerkeleyDB</p> <p>Location: <code>MGMT_ROOT/db/infrastructure</code> on the ZFS, accessible via <code>/nfs/shared_storage/db/infrastructure</code> on each management node</p>
Oracle Switch ES1-24 Ports Database	<p>Defines the factory-configured map of Oracle Switch ES1-24 ports to the rack unit or element to which that port is connected. It is used to map Oracle Switch ES1-24 ports to machine names.</p> <p>Used only on systems with InfiniBand-based network architecture.</p> <p>Type: BerkeleyDB</p> <p>Location: <code>MGMT_ROOT/db/opus_ports</code> on the ZFS, accessible via <code>/nfs/shared_storage/db/opus_ports</code> on each management node</p>
Oracle Private Cloud Appliance Mini Database	<p>A multi-purpose database used to map compute node hardware profiles to on-board disk size information. It also contains valid hardware configurations that servers must comply with in order to be accepted as an Oracle Private Cloud Appliance component. Entries contain a sync ID for more convenient usage within the Command Line Interface (CLI).</p> <p>Type: BerkeleyDB</p> <p>Location: <code>MGMT_ROOT/db/mini_db</code> on the ZFS, accessible via <code>/nfs/shared_storage/db/mini_db</code> on each management node</p>
Oracle Private Cloud Appliance Monitor Database	<p>Records fault counts detected through the ILOMs of all active components identified in the Inventory Database.</p> <p>Type: BerkeleyDB</p> <p>Location: <code>MGMT_ROOT/db/monitor</code> on the ZFS, accessible via <code>/nfs/shared_storage/db/monitor</code> on each management node</p>

Item	Description
Oracle Private Cloud Appliance Setup Database	<p>Contains the data set by the Oracle Private Cloud Appliance Dashboard setup facility. The data in this database is automatically applied by both the active and standby management nodes when a change is detected.</p> <p>Type: BerkeleyDB</p> <p>Location: <code>MGMT_ROOT/db/setup</code> on the ZFS, accessible via <code>/nfs/shared_storage/db/setup</code> on each management node</p>
Oracle Private Cloud Appliance Task Database	<p>Contains state data for all of the asynchronous tasks that have been dispatched within the Oracle Private Cloud Appliance.</p> <p>Type: BerkeleyDB</p> <p>Location: <code>MGMT_ROOT/db/task</code> on the ZFS, accessible via <code>/nfs/shared_storage/db/task</code> on each management node</p>
Oracle Private Cloud Appliance Synchronization Databases	<p>Contain data and configuration settings for the synchronization service to apply and maintain across rack components. Errors from failed attempts to synchronize configuration parameters across appliance components can be reviewed in the <code>sync_errored_tasks</code> database, from where they can be retried or acknowledged.</p> <p>Synchronization databases are not present by default. They are created when the first synchronization task of a given type is received.</p> <p>Type: BerkeleyDB</p> <p>Location: <code>MGMT_ROOT/db/sync_*</code> on the ZFS, accessible via <code>/nfs/shared_storage/db/sync_*</code> on each management node</p>
Oracle Private Cloud Appliance Update Database	<p>Used to track the two-node coordinated management node update process.</p> <div style="display: flex; align-items: flex-start;"> <div style="margin-right: 10px;">  </div> <div> <p>Note</p> <p>Database schema changes and wallet changes between different releases of the controller software are written to a file. It ensures that these critical changes are applied early in the software update process, before any other appliance components are brought back up.</p> </div> </div> <p>Type: BerkeleyDB</p> <p>Location: <code>MGMT_ROOT/db/update</code> on the ZFS, accessible via <code>/nfs/shared_storage/db/update</code> on each management node</p>
Oracle Private Cloud Appliance Tenant Database	<p>Contains details about all tenant groups: default and custom. These details include the unique tenant group ID, file system ID, member compute nodes, status information, etc.</p> <p>Type: BerkeleyDB</p> <p>Location: <code>MGMT_ROOT/db/tenant</code> on the ZFS, accessible via <code>/nfs/shared_storage/db/tenant</code> on each management node</p>

Item	Description
Oracle VM Manager Database	<p>Used on each management node as the management database for Oracle VM Manager. It contains all configuration details of the Oracle VM environment (including servers, pools, storage and networking), as well as the virtualized systems hosted by the environment.</p> <p>Type: MySQL Database</p> <p>Location: <code>MGMT_ROOT/ovmm_mysql/data/</code> on the ZFS, accessible via <code>/nfs/shared_storage/ovmm_mysql/data/</code> on each management node</p>

1.3.6 Oracle Private Cloud Appliance Management Software

The Oracle Private Cloud Appliance includes software that is designed for the provisioning, management and maintenance of all of the components within the appliance. The controller software, which handles orchestration and automation of tasks across various hardware components, is not intended for human interaction. Its appliance administration functions are exposed through the browser interface and command line interface, which are described in detail in this guide.



Important

All configuration and management tasks must be performed using the Oracle Private Cloud Appliance Dashboard and the Command Line Interface. Do not attempt to run any processes directly without explicit instruction from an Oracle Support representative. Attempting to do so may render your appliance unusable.

Besides the Dashboard and CLI, this software also includes a number of Python applications that run on the active management node. These applications are found in `/usr/sbin` on each management node and some are listed as follows:

- `pca-backup`: the script responsible for performing backups of the appliance configuration as described in [Section 1.6, “Oracle Private Cloud Appliance Backup”](#)
- `pca-check-master`: a script that verifies which of the two management nodes currently has the master role
- `ovca-daemon`: the core provisioning and management daemon for the Oracle Private Cloud Appliance
- `pca-dhcpd`: a helper script to assist the DHCP daemon with the registration of compute nodes
- `pca-diag`: a tool to collect diagnostic information from your Oracle Private Cloud Appliance, as described in [Section 1.3.7, “Oracle Private Cloud Appliance Diagnostics Tool”](#)
- `pca-factory-init`: the appliance initialization script used to set the appliance to its factory configuration. This script does not function as a reset; it is only used for initial rack setup.
- `pca-redirect`: a daemon that redirects HTTP or HTTPS requests to the Oracle Private Cloud Appliance Dashboard described in [Section 1.3.1, “Oracle Private Cloud Appliance Dashboard”](#)
- `ovca-remote-rpc`: a script for remote procedure calls directly to the Oracle VM Server Agent. Currently it is only used by the management node to monitor the heartbeat of the Oracle VM Server Agent.
- `ovca-rpc`: a script that allows the Oracle Private Cloud Appliance software components to communicate directly with the underlying management scripts running on the management node

Many of these applications use a specific Oracle Private Cloud Appliance library that is installed in `/usr/lib/python2.6/site-packages/ovca/` on each management node.

1.3.7 Oracle Private Cloud Appliance Diagnostics Tool

The Oracle Private Cloud Appliance includes a tool that can be run to collect diagnostic data: logs and other types of files that can help to troubleshoot hardware and software problems. This tool is located in `/usr/sbin/` on each management and compute node, and is named `pca-diag`. The data it retrieves, depends on the selected command line arguments:

- `pca-diag`

When you enter this command, without any additional arguments, the tool retrieves a basic set of files that provide insights into the current health status of the Oracle Private Cloud Appliance. You can run this command on all management and compute nodes. All collected data is stored in `/tmp`, compressed into a single tarball (`ovcadiag_<node-hostname>_<ID>_<date>_<time>.tar.bz2`).

- `pca-diag version`

When you enter this command, version information for the current Oracle Private Cloud Appliance software stack is displayed. The `version` argument cannot be combined with any other argument.

- `pca-diag ilom`

When you enter this command, diagnostic data is retrieved, by means of `ipmitool`, through the host's ILOM. The data set includes details about the host's operating system, processes, health status, hardware and software configuration, as well as a number of files specific to the Oracle Private Cloud Appliance configuration. You can run this command on all management and compute nodes. All collected data is stored in `/tmp`, compressed into a single tarball (`ovcadiag_<node-hostname>_<ID>_<date>_<time>.tar.bz2`).

- `pca-diag vmpinfo`



Caution

When using the `vmpinfo` argument, the command must be run from the master management node.

When you enter this command, the Oracle VM diagnostic data collection mechanism is activated. The `vmpinfo3` script collects logs and configuration details from the Oracle VM Manager, and logs and `sosreport` information from each Oracle VM Server or compute node. All collected data is stored in `/tmp`, compressed into **two tarballs**: `ovcadiag_<node-hostname>_<ID>_<date>_<time>.tar.bz2` and `vmpinfo3-<version>-<date>-<time>.tar.gz`.

To collect diagnostic information for a subset of the Oracle VM Servers in the environment, you run the command with an additional `servers` parameter: `pca-diag vmpinfo servers='ovcacn07r1,ovcacn08r1,ovcacn09r1'`

Diagnostic collection with `pca-diag` is possible from the command line of any node in the system. Only the master management node allows you to use all of the command line arguments. Although `vmpinfo` is not available on the compute nodes, running `pca-diag` directly on the compute can help retrieve important diagnostic information regarding Oracle VM Server that cannot be captured with `vmpinfo`.

The `pca-diag` tool is typically run by multiple users with different roles. System administrators or field service engineers may use it as part of their standard operating procedures, or Oracle Support teams may request that the tool be run in a specific manner as part of an effort to diagnose and resolve reported

hardware or software issues. For additional information and instructions, also refer to the section [“Data Collection for Service and Support”](#) in the Oracle Private Cloud Appliance Release Notes.

1.4 Provisioning and Orchestration

As a converged infrastructure solution, the Oracle Private Cloud Appliance is built to eliminate many of the intricacies of optimizing the system configuration. Hardware components are installed and cabled at the factory. Configuration settings and installation software are preloaded onto the system. Once the appliance is connected to the data center power source and public network, the provisioning process between the administrator pressing the power button of the first management node and the appliance reaching its *Deployment Readiness* state is entirely orchestrated by the master management node. This section explains what happens as the Oracle Private Cloud Appliance is initialized and all nodes are provisioned.

1.4.1 Appliance Management Initialization

Boot Sequence and Health Checks

When power is applied to the first management node, it takes approximately five minutes for the server to boot. While the Oracle Linux 6 operating system is loading, an Apache web server is started, which serves a static welcome page the administrator can browse to from the workstation connected to the appliance management network.

The necessary Oracle Linux services are started as the server comes up to runlevel 3 (multi-user mode with networking). At this point, the management node executes a series of system health checks. It verifies that all expected infrastructure components are present on the appliance administration network and in the correct predefined location, identified by the rack unit number and fixed IP address. Next, the management node probes the ZFS storage appliance for a management NFS export and a management iSCSI LUN with OCFS2 file system. The storage and its access groups have been configured at the factory. If the health checks reveal no problems, the `ocfs2` and `o2cb` services are started up automatically.

Management Cluster Setup

When the OCFS2 file system on the shared iSCSI LUN is ready, and the `o2cb` services have started successfully, the management nodes can join the cluster. In the meantime, the first management node has also started the second management node, which will come up with an identical configuration. Both management nodes eventually join the cluster, but the first management node will take an exclusive lock on the shared OCFS2 file system using Distributed Lock Management (DLM). The second management node remains in permanent standby and takes over the lock only in case the first management node goes down or otherwise releases its lock.

With mutual exclusion established between both members of the management cluster, the master management node continues to load the remaining Oracle Private Cloud Appliance services, including `dhcpd`, Oracle VM Manager and the Oracle Private Cloud Appliance databases. The virtual IP address of the management cluster is also brought online, and the Oracle Private Cloud Appliance Dashboard is activated. The static Apache web server now redirects to the Dashboard at the virtual IP, where the administrator can access a live view of the appliance rack component status.

Once the `dhcpd` service is started, the system state changes to *Provision Readiness*, which means it is ready to discover non-infrastructure components.

1.4.2 Compute Node Discovery and Provisioning

Node Manager

To discover compute nodes, the Node Manager on the master management node uses a DHCP server and the node database. The node database is a BerkeleyDB type database, located on the management

NFS share, containing the state and configuration details of each node in the system, including MAC addresses, IP addresses and host names. The discovery process of a node begins with a DHCP request from the ILOM. Most discovery and provisioning actions are synchronous and occur sequentially, while time consuming installation and configuration processes are launched in parallel and asynchronously. The DHCP server hands out pre-assigned IP addresses on the appliance administration network ([192.168.4.0/24](#)). When the Node Manager has verified that a node has a valid service tag for use with Oracle Private Cloud Appliance, it launches a series of provisioning tasks. All required software resources have been loaded onto the ZFS storage appliance at the factory.

Provisioning Tasks

The provisioning process is tracked in the node database by means of status changes. The next provisioning task can only be started if the node status indicates that the previous task has completed successfully. For each valid node, the Node Manager begins by building a PXE configuration and forces the node to boot using Oracle Private Cloud Appliance runtime services. After the hardware RAID-1 configuration is applied, the node is restarted to perform a kickstart installation of Oracle VM Server. Crucial kernel modules and host drivers are added to the installation. At the end of the installation process, the network configuration files are updated to allow all necessary network interfaces to be brought up.

Once the internal management network exists, the compute node is rebooted one last time to reconfigure the Oracle VM Agent to communicate over this network. At this point, the node is ready for Oracle VM Manager discovery.

As the Oracle VM environment grows and contains more and more virtual machines and many different VLANs connecting them, the number of management operations and registered events increases rapidly. In a system with this much activity the provisioning of a compute node takes significantly longer, because the provisioning tasks run through the same management node where Oracle VM Manager is active. There is no impact on functionality, but the provisioning tasks can take several hours to complete. It is recommended to perform compute node provisioning at a time when system activity is at its lowest.

1.4.3 Server Pool Readiness

Oracle VM Server Pool

When the Node Manager detects a fully installed compute node that is ready to join the Oracle VM environment, it issues the necessary Oracle VM CLI commands to add the new node to the Oracle VM server pool. With the discovery of the first node, the system also configures the clustered Oracle VM server pool with the appropriate networking and access to the shared storage. For every compute node added to Oracle VM Manager the IPMI configuration is stored in order to enable convenient remote power-on/off.

Oracle Private Cloud Appliance expects that all compute nodes in one rack initially belong to a single clustered server pool with High Availability (HA) and Distributed Resource Scheduling (DRS) enabled. When all compute nodes have joined the Oracle VM server pool, the appliance is in *Ready* state, meaning virtual machines (VMs) can be deployed.

Expansion Compute Nodes

When an expansion compute node is installed, its presence is detected based on the DHCP request from its ILOM. If the new server is identified as an Oracle Private Cloud Appliance node, an entry is added in the node database with "*new*" state. This triggers the initialization and provisioning process. New compute nodes are integrated seamlessly to expand the capacity of the running system, without the need for manual reconfiguration by an administrator.

Synchronization Service

As part of the provisioning process, a number of configuration settings are applied, either globally or at individual component level. Some are visible to the administrator, and some are entirely internal to the system. Throughout the life cycle of the appliance, software updates, capacity extensions and configuration changes will occur at different points in time. For example, an expansion compute node may have different hardware, firmware, software, configuration and passwords compared to the servers already in use in the environment, and it comes with factory default settings that do not match those of the running system. A synchronization service, implemented on the management nodes, can set and maintain configurable parameters across heterogeneous sets of components within an Oracle Private Cloud Appliance environment. It facilitates the integration of new system components in case of capacity expansion or servicing, and allows the administrator to streamline the process when manual intervention is required. The CLI provides an interface to the exposed functionality of the synchronization service.

1.5 High Availability

The Oracle Private Cloud Appliance is designed for high availability at every level of its component make-up.

Management Node Failover

During the factory installation of an Oracle Private Cloud Appliance, the management nodes are configured as a cluster. The cluster relies on an OCFS2 file system exported as an iSCSI LUN from the ZFS storage to perform the heartbeat function and to store a lock file that each management node attempts to take control of. The management node that has control over the lock file automatically becomes the master or active node in the cluster.

When the Oracle Private Cloud Appliance is first initialized, the `o2cb` service is started on each management node. This service is the default cluster stack for the OCFS2 file system. It includes a node manager that keeps track of the nodes in the cluster, a heartbeat agent to detect live nodes, a network agent for intra-cluster node communication and a distributed lock manager to keep track of lock resources. All these components are in-kernel.

Additionally, the `ovca` service is started on each management node. The management node that obtains control over the cluster lock and is thereby promoted to the master or active management node, runs the full complement of Oracle Private Cloud Appliance services. This process also configures the Virtual IP that is used to access the active management node, so that it is 'up' on the active management node and 'down' on the standby management node. This ensures that, when attempting to connect to the Virtual IP address that you configured for the management nodes, you are always accessing the active management node.

In the case where the active management node fails, the cluster detects the failure and the lock is released. Since the standby management node is constantly polling for control over the lock file, it detects when it has control of this file and the `ovca` service brings up all of the required Oracle Private Cloud Appliance services. On the standby management node the Virtual IP is configured on the appropriate interface as it is promoted to the active role.

When the management node that failed comes back online, it no longer has control of the cluster lock file. It is automatically put into standby mode, and the Virtual IP is removed from the management interface. This means that one of the two management nodes in the rack is always available through the same IP address and is always correctly configured. The management node failover process takes up to 5 minutes to complete.

Oracle VM Management Database Failover

The Oracle VM Manager database files are located on a shared file system exposed by the ZFS storage appliance. The active management node runs the MySQL database server, which accesses the database files on the shared storage. In the event that the management node fails, the standby management node is promoted and the MySQL database server on the promoted node is started so that the service can resume as normal. The database contents are available to the newly running MySQL database server.

Compute Node Failover

High availability (HA) of compute nodes within the Oracle Private Cloud Appliance is enabled through the clustered server pool that is created automatically in Oracle VM Manager during the compute node provisioning process. Since the server pool is configured as a cluster using an underlying OCFS2 file system, HA-enabled virtual machines running on any compute node can be migrated and restarted automatically on an alternate compute node in the event of failure.

The [Oracle VM Concepts Guide](#) provides good background information about the principles of high availability. Refer to the section [How does High Availability \(HA\) Work?](#).

Storage Redundancy

Further redundancy is provided through the use of the ZFS storage appliance to host storage. This component is configured with RAID-1 providing integrated redundancy and excellent data loss protection. Furthermore, the storage appliance includes two storage heads or controllers that are interconnected in a clustered configuration. The pair of controllers operate in an active-passive configuration, meaning continuation of service is guaranteed in the event that one storage head should fail. The storage heads share a single IP in the storage subnet, but both have an individual management IP address for convenient maintenance access.

Network Redundancy

All of the customer-usable networking within the Oracle Private Cloud Appliance is configured for redundancy. Only the internal administrative Ethernet network, which is used for initialization and ILOM connectivity, is not redundant. There are two of each switch type to ensure that there is no single point of failure. Networking cabling and interfaces are equally duplicated and switches are interconnected as described in [Section 1.2.4, “Network Infrastructure”](#).

1.6 Oracle Private Cloud Appliance Backup

The configuration of all components within Oracle Private Cloud Appliance is automatically backed up and stored on the ZFS storage appliance as a set of archives. Backups are named with a time stamp for when the backup is run.

During initialization, a crontab entry is created on each management node to perform a global backup twice in every 24 hours. The first backup runs at 09h00 and the second at 21h00. Only the active management node actually runs the backup process when it is triggered.



Note

To trigger a backup outside of the default schedule, use the [Command Line Interface](#). For details, refer to [Section 4.2.4, “backup”](#).

Backups are stored on the `MGMT_ROOT` file system on the ZFS storage appliance and are accessible on each management node at `/nfs/shared_storage/backups`. When the backup process is triggered, it creates a temporary directory named with the time stamp for the current backup process. The entire

directory is archived in a `*.tar.bz2` file when the process is complete. Within this directory several subdirectories are also created:

- **data_net_switch**: used only on systems with Ethernet-based network architecture; contains the configuration data of the spine and leaf switches
- **mgmt_net_switch**: used only on systems with Ethernet-based network architecture; contains the management switch configuration data
- **nm2**: used only on systems with InfiniBand-based network architecture; contains the NM2-36P Sun Datacenter InfiniBand Expansion Switch configuration data
- **opus**: used only on systems with InfiniBand-based network architecture; contains the Oracle Switch ES1-24 configuration data
- **ovca**: contains all of the configuration information relevant to the deployment of the management nodes such as the password wallet, the network configuration of the management nodes, configuration databases for the Oracle Private Cloud Appliance services, and DHCP configuration.
- **ovmm**: contains the most recent backup of the Oracle VM Manager database, the actual source data files for the current database, and the UUID information for the Oracle VM Manager installation. Note that the actual backup process for the Oracle VM Manager database is handled automatically from within Oracle VM Manager. Manual backup and restore are described in detail in the section entitled [Backing up and Restoring Oracle VM Manager](#), in the [Oracle VM Manager Administration Guide](#).
- **ovmm_upgrade**: contains essential information for each upgrade attempt. When an upgrade is initiated, a time-stamped subdirectory is created to store the `preinstall.log` file with the output from the pre-upgrade checks. Backups of any other files modified during the pre-upgrade process, are also saved in this directory.
- **xsigo**: used only on systems with InfiniBand-based network architecture; contains the configuration data for the Fabric Interconnects.
- **zfsa**: contains all of the configuration information for the ZFS storage appliance

The backup process collects data for each component in the appliance and ensures that it is stored in a way that makes it easy to restore that component to operation in the case of failure¹.

Taking regular backups is standard operating procedure for any production system. The internal backup mechanism cannot protect against full system failure, site outage or disaster. Therefore, you should consider implementing a backup strategy to copy key system data to external storage. This requires what is often referred to as a *bastion host*: a dedicated system configured with specific internal and external connections.

For a detailed description of the backup contents, and for guidelines to export internal backups outside the appliance, refer to the Oracle technical white paper entitled [Oracle Private Cloud Appliance Backup Guide](#).

1.7 Oracle Private Cloud Appliance Upgrader

Together with Oracle Private Cloud Appliance Controller Software Release 2.3.4, a new independent upgrade tool was introduced: the Oracle Private Cloud Appliance Upgrader. It is provided as a separate application, with its own release and update schedule. It maintains the phased approach, where management nodes, compute nodes and other rack components are updated in separate procedures, while at the same time it groups and automates sets of tasks that were previously executed as scripted

¹ Restoration from backup must only be performed by Oracle Service Personnel.

or manual steps. The new design has better error handling and protection against terminal crashes, ssh timeouts or inadvertent user termination. It is intended to reduce complexity and improve the overall upgrade experience.

The Oracle Private Cloud Appliance Upgrader was built as a modular framework. Each module consists of pre-checks, an execution phase such as upgrade or install, and post-checks. Besides the standard interactive mode, modules also provide silent mode for programmatic use, and verify-only mode to run pre-checks without starting the execution phase.

The first module developed within the Oracle Private Cloud Appliance Upgrader framework, is the management node upgrade. With a single command, it guides the administrator through the pre-upgrade validation steps – now included in the pre-checks of the Upgrader –, software image deployment, Oracle Private Cloud Appliance Controller Software update, Oracle Linux operating system Yum update, and Oracle VM Manager upgrade.

For software update instructions, see [Chapter 3, *Updating Oracle Private Cloud Appliance*](#).

For specific Oracle Private Cloud Appliance Upgrader details, see [Section 3.2, “Using the Oracle Private Cloud Appliance Upgrader”](#).

Chapter 2 Monitoring and Managing Oracle Private Cloud Appliance

Table of Contents

2.1 Connecting and Logging in to the Oracle Private Cloud Appliance Dashboard	26
2.2 Oracle Private Cloud Appliance Accessibility Features	28
2.3 Hardware View	28
2.4 Network Settings	32
2.5 Functional Networking Limitations	35
2.5.1 Network Configuration of Ethernet-based Systems	36
2.5.2 Network Configuration of InfiniBand-based Systems	38
2.6 Network Customization	39
2.6.1 Configuring Custom Networks on Ethernet-based Systems	40
2.6.2 Configuring Custom Networks on InfiniBand-based Systems	44
2.6.3 Deleting Custom Networks	49
2.7 Tenant Groups	50
2.7.1 Design Assumptions and Restrictions	50
2.7.2 Configuring Tenant Groups	51
2.8 Authentication	54
2.9 Health Monitoring	57

Monitoring and management of the Oracle Private Cloud Appliance is achieved using the Oracle Private Cloud Appliance Dashboard. This web-based graphical user interface is also used to perform the initial configuration of the appliance beyond the instructions provided in the Quick Start poster included in the packaging of the appliance.



Warning

Before starting the system and applying the initial configuration, read and understand the [Oracle Private Cloud Appliance Release Notes](#). The section [Known Limitations and Workarounds](#) provides information that is critical for correctly executing the procedures in this document. Ignoring the release notes may cause you to configure the system incorrectly. Bringing the system back to normal operation may require a complete factory reset.

The Oracle Private Cloud Appliance Dashboard allows you to perform the following tasks:

- Initial software configuration (and reconfiguration) for the appliance using the Network Environment window, as described in [Section 2.4, “Network Settings”](#).
- Hardware provisioning monitoring and identification of each hardware component used in the appliance, accessed via the Hardware View window described in [Section 2.3, “Hardware View”](#).
- Resetting the passwords used for different components within the appliance, via the Password Management window, as described in [Section 2.8, “Authentication”](#).

The Oracle Private Cloud Appliance Controller Software includes functionality that is currently not available through the Dashboard user interface:

- **Backup**

The configuration of all components within Oracle Private Cloud Appliance is automatically backed up based on a crontab entry. This functionality is not configurable. Restoring a backup requires the intervention of an Oracle-qualified service person. For details, see [Section 1.6, “Oracle Private Cloud Appliance Backup”](#).

- **Update**

The update process is controlled from the command line of the master management node, using the Oracle Private Cloud Appliance Upgrader. For details, see [Section 1.7, “Oracle Private Cloud Appliance Upgrader”](#). For step-by-step instructions, see [Chapter 3, *Updating Oracle Private Cloud Appliance*](#).

- **Custom Networks**

In situations where the default network configuration is not sufficient, the command line interface allows you to create additional networks at the appliance level. For details and step-by-step instructions, see [Section 2.6, “Network Customization”](#).

- **Tenant Groups**

The command line interface provides commands to optionally subdivide an Oracle Private Cloud Appliance environment into a number of isolated groups of compute nodes. These groups of servers are called tenant groups, which are reflected in Oracle VM as different server pools. For details and step-by-step instructions, see [Section 2.7, “Tenant Groups”](#).


2.1 Connecting and Logging in to the Oracle Private Cloud Appliance Dashboard

To open the Login page of the Oracle Private Cloud Appliance Dashboard, enter the following address in a Web browser:

```
https://manager-vip:7002/dashboard
```

Where, *manager-vip* refers to the shared Virtual IP address that you have configured for your management nodes during installation. By using the shared Virtual IP address, you ensure that you always access the Oracle Private Cloud Appliance Dashboard on the active management node.

Figure 2.1 Dashboard Login



The screenshot shows the login interface for the Oracle Private Cloud Appliance Admin Console. It features a blue header with a cloud icon and the text "SIGN IN Oracle PCA Admin Console". Below the header, there are two input fields: "User Name:" and "Password:". A blue "Sign In" button is positioned below the password field. The Oracle logo is located in the bottom right corner, and a copyright notice "COPYRIGHT © 2019, ORACLE AND/OR ITS AFFILIATES. ALL RIGHTS RESERVED." is at the bottom left.



Note

If you are following the installation process and this is your first time accessing the Oracle Private Cloud Appliance Dashboard, the Virtual IP address in use by the master management node is set to the factory default `192.168.4.216`. This is an IP address in the internal appliance management network, which can only be reached if you use a workstation patched directly into the available Ethernet port 48 in the Cisco Nexus 9348GC-FXP Switch.

Systems with an InfiniBand-based network architecture contain a pair Oracle Switch ES1-24 switches instead. If your appliance contains such switches, connected the workstation to Ethernet port 19 in one of them, not both.

The default user name is **admin** and the default password is **Welcome1**. For security reasons, you must set a new password at your earliest convenience.



Important

You must ensure that if you are accessing the Oracle Private Cloud Appliance Dashboard through a firewalled connection, the firewall is configured to allow TCP traffic on the port that the Oracle Private Cloud Appliance Dashboard is using to listen for connections.

Enter your Oracle Private Cloud Appliance Dashboard administration user name in the **User Name** field. This is the administration user name you configured during installation. Enter the password for the Oracle Private Cloud Appliance Dashboard administration user name in the **Password** field.



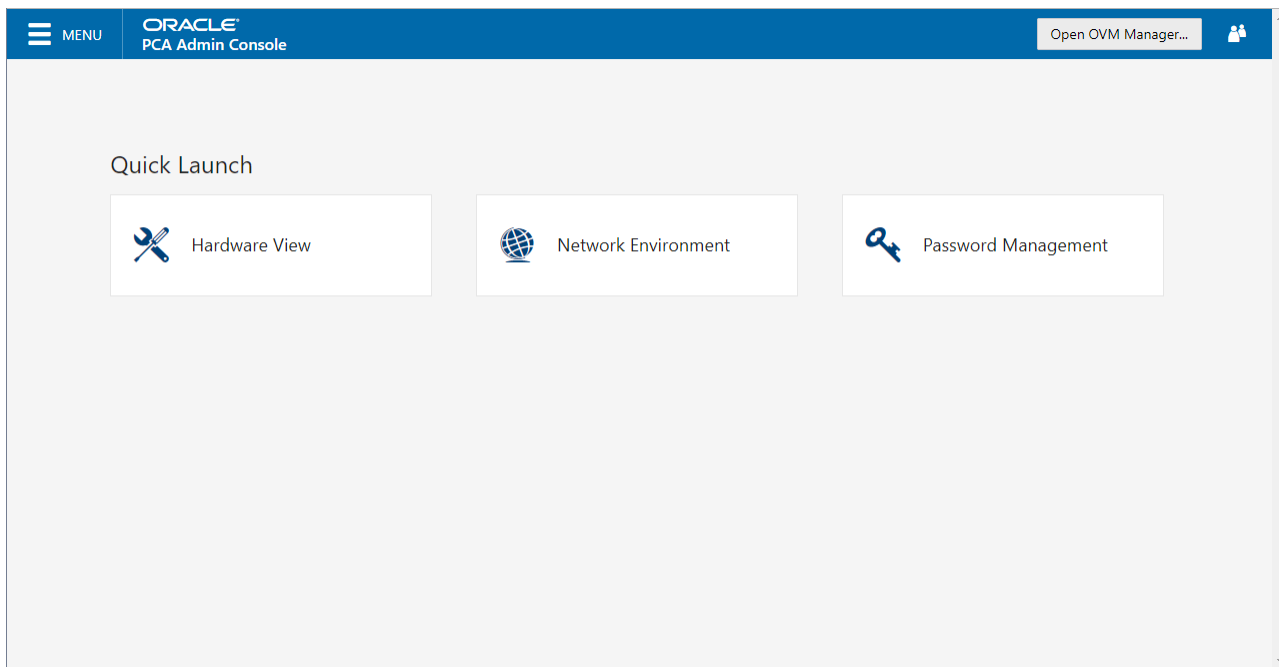
Important

The Oracle Private Cloud Appliance Dashboard makes use of cookies in order to store session data. Therefore, to successfully log in and use the Oracle Private

Cloud Appliance Dashboard, your web browser must accept cookies from the Oracle Private Cloud Appliance Dashboard host.

When you have logged in to the Dashboard successfully, the home page is displayed. The central part of the page contains Quick Launch buttons that provide direct access to the key functional areas.

Figure 2.2 Dashboard Home Page



From every Dashboard window you can always go to any other window by clicking the Menu in the top-left corner and selecting a different window from the list. A button in the header area allows you to open Oracle VM Manager.

2.2 Oracle Private Cloud Appliance Accessibility Features

For detailed accessibility information, refer to the chapter [Documentation Accessibility](#) in the *Oracle Private Cloud Appliance Release Notes*.

2.3 Hardware View

The **Hardware View** window within the Oracle Private Cloud Appliance Dashboard provides a graphical representation of the hardware components as they are installed within the rack. The view of the status of these components is automatically refreshed every 30 seconds by default. You can set the refresh interval or disable it through the Auto Refresh Interval list. Alternatively, a **Refresh** button at the top of the page allows you to refresh the view at any time.

During particular maintenance tasks, such as upgrading management nodes, you may need to disable compute node provisioning temporarily. This **Disable CN Provisioning** button at the top of the page allows you to suspend provisioning activity. When compute node provisioning is suspended, the button text changes to **Enable CN Provisioning** and its purpose changes to allow you to resume compute node provisioning as required.

Rolling over each item in the graphic with the mouse raises a pop-up window providing the name of the component, its type, and a summary of configuration and status information. For compute nodes, the pop-

up window includes a **Reprovision** button, which allows you to restart the provisioning process if the node becomes stuck in an intermittent state or goes into error status before it is added to the Oracle VM server pool. Instructions to reprovision a compute node are provided in [Section 7.10, “Reprovisioning a Compute Node when Provisioning Fails”](#).



Caution

The **Reprovision** button is to be used *only* for compute nodes that fail to complete provisioning. For compute nodes that have been provisioned properly and/or host running virtual machines, the **Reprovision** button is made unavailable to prevent incorrect use, thus protecting healthy compute nodes from loss of functionality, data corruption, or being locked out of the environment permanently.



Caution

Reprovisioning restores a compute node to a clean state. If a compute node was previously added to the Oracle VM environment and has active connections to storage repositories other than those on the internal ZFS storage, the external storage connections need to be configured again after reprovisioning.

Alongside each installed component within the appliance rack, a status icon provides an indication of the *provisioning status* of the component. A status summary is displayed just above the rack image, indicating with icons and numbers how many nodes have been provisioned, are being provisioned, or are in error status. The Hardware View does not provide real-time health and status information about active components. Its monitoring functionality is restricted to the provisioning process. When a component has been provisioned completely and correctly, the Hardware View continues to indicate correct operation even if the component should fail or be powered off. See [Table 2.1](#) for an overview of the different status icons and their meaning.

Table 2.1 Table of Hardware Provisioning Status Icons

Icon	Status	Description
	OK	The component is running correctly and has passed all health check operations. Provisioning is complete.
	Provisioning	The component is running, and provisioning is in progress. The progress bar fills up as the component goes through the various stages of provisioning. Key stages for compute nodes include: HMP initialization actions, Oracle VM Server installation, network configuration, storage setup, and server pool membership.
	Error	The component is not running and has failed health check operations. Component troubleshooting is required and the component may need to be replaced. Compute nodes also have this status when provisioning has failed.



Note

For real-time health and status information of your active Oracle Private Cloud Appliance hardware, after provisioning, consult the Oracle VM Manager or Oracle Enterprise Manager UI.

The Hardware View provides an accessible tool for troubleshooting hardware components within the Oracle Private Cloud Appliance and identifying where these components are actually located within the

rack. Where components might need replacing, the new component must take the position of the old component within the rack to maintain configuration.

Base Rack Front View

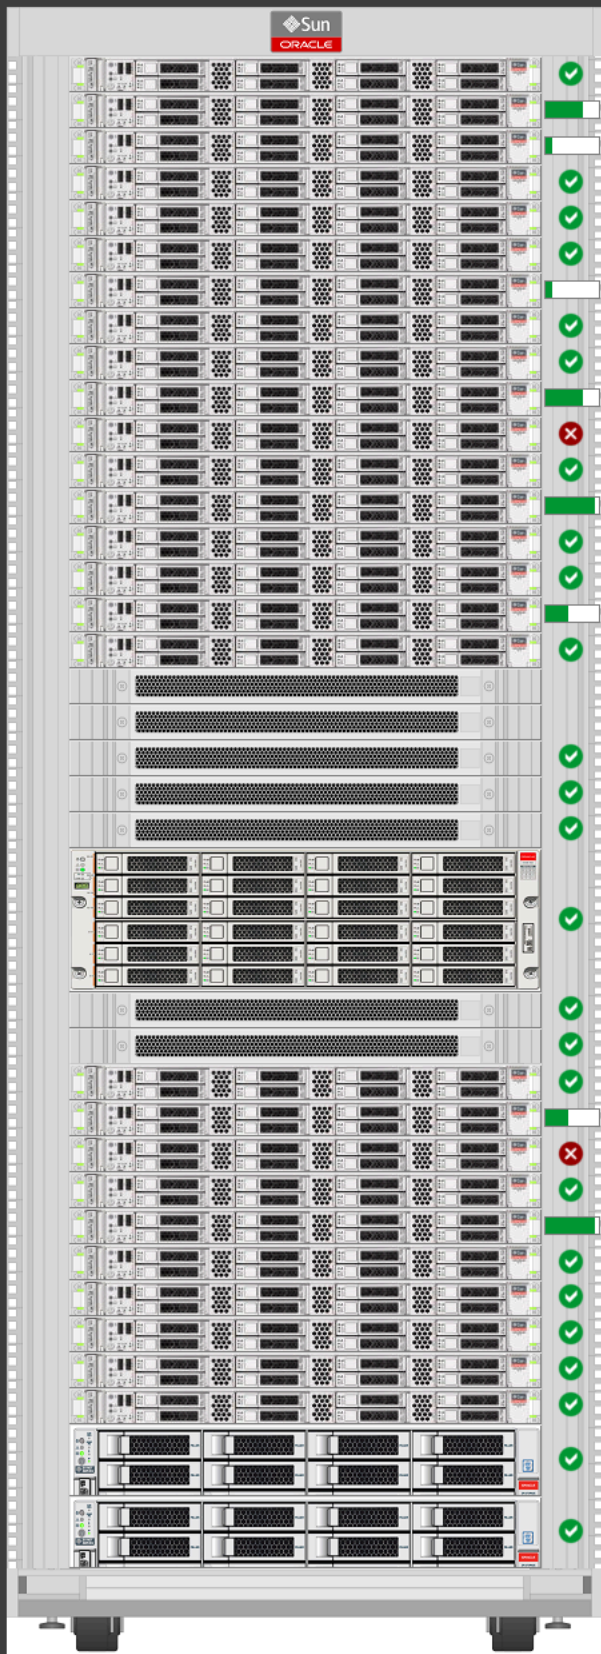
Refresh

Disable CN Provisioning

Auto refresh interval

30 seconds

✔ 7 ⌚ 0 ✖ 0



2.4 Network Settings

The **Network Environment** window is used to configure networking and service information for the management nodes. For this purpose, you should reserve three IP addresses in the public (data center) network: one for each management node, and one to be used as virtual IP address by both management nodes. The virtual IP address provides access to the Dashboard once the software initialization is complete.

To avoid network interference and conflicts, you must ensure that the data center network does not overlap with any of the infrastructure subnets of the Oracle Private Cloud Appliance default configuration. These are the subnets and VLANs you should keep clear:

Subnets:

- 192.168.4.0/24 – internal machine administration network: connects ILOMs and physical hosts
- 192.168.140.0/24 – internal Oracle VM management network: connects Oracle VM Manager, Oracle VM Server and Oracle VM Agents (applies only to the InfiniBand-based architecture)
- 192.168.32.0/21 – internal management network: traffic between management and compute nodes
- 192.168.64.0/21 – underlay network for east/west traffic within the appliance environment
- 192.168.72.0/21 – underlay network for north/south traffic, enabling external connectivity
- 192.168.40.0/21 – storage network: traffic between the servers and the ZFS storage appliance



Note

Each /21 subnet comprises the IP ranges of eight /24 subnets or over 2000 IP addresses. For example: 192.168.32.0/21 corresponds with all IP addresses from 192.168.32.1 to 192.168.39.255.

VLANs:

- 1 – the Cisco default VLAN
- 3040 – the default service VLAN
- 3041-3072 – a range of 31 VLANs reserved for customer VM and host networks
- 3073-3099 – a range reserved for system-level connectivity



Note

VLANs 3090-3093 are already in use for tagged traffic over the /21 subnets listed above.

- 3968-4095 – a range reserved for Cisco internal device allocation

The **Network Environment** window is divided into three tabs: Management Nodes, Data Center Network, and DNS. Each tab is shown in this section, along with a description of the available configuration fields.

You can undo the changes you made in any of the tabs by clicking the Reset button. To confirm the configuration changes you made, enter the Dashboard Admin user password in the applicable field at the bottom of the window, and click Apply Changes.

**Note**

When you click Apply Changes, the configuration settings in all three tabs are applied. Make sure that all required fields in all tabs contain valid information before you proceed.

Figure 2.4 shows the Management Nodes tab. The following fields are available for configuration:

- **Management Node 1:**
 - **IP Address:** Specify an IP address within your datacenter network that can be used to directly access this management node.
 - **Host Name:** Specify the host name for the first management node system.
- **Management Node 2:**
 - **IP Address:** Specify an IP address within your datacenter network that can be used to directly access this management node.
 - **Host Name:** Specify the host name for the second management node system.
- **Management Virtual IP Address:** Specify the shared Virtual IP address that is used to always access the active management node. This IP address must be in the same subnet as the IP addresses that you have specified for each management node.

Figure 2.4 Management Nodes Tab

The screenshot shows the Oracle PCA Admin Console interface. The top navigation bar includes a menu icon, the text 'ORACLE PCA Admin Console', and a button labeled 'Open OVM Manager...'. The main content area is titled 'Network Environment' and contains a sub-header: 'Changes to the Network Environment configuration data requires entering the PCA Admin Password before selecting the 'Apply Changes' button.' Below this, there are three tabs: 'Management Nodes' (selected), 'Data Center Network', and 'DNS'. The 'Management Nodes' tab contains three sections:

- Management Node 1:** IP Address (10.19.1.101) and Host Name (manager1).
- Management Node 2:** IP Address (10.19.1.102) and Host Name (manager2).
- Management Virtual IP Address:** IP Address (10.19.1.100).

 At the bottom, there is a field for the PCA Admin Password (required for changes), a 'Reset' button, and an 'Apply Changes' button.

Figure 2.5 shows the Data Center Network tab. The following fields are available for configuration:

- **Management Network VLAN:** The default configuration does not assume that your management network exists on a VLAN. If you have configured a VLAN on your switch for the management network, you should toggle the slider to the *active* setting and then specify the VLAN ID in the provided field.



Caution

For systems with Ethernet-based network architecture, a management VLAN requires additional configuration steps.

When a VLAN is used for the management network, and VM traffic must be enabled over the same network, you must manually configure a VLAN interface on the vx13040 interfaces of the necessary compute nodes to connect them to the VLAN with the ID in question. For instructions to create a VLAN interface on a compute node, refer to the [Oracle VM documentation](#).

- **Domain Name:** Specify the data center domain that the management nodes belong to.
- **Netmask:** Specify the netmask for the network that the Virtual IP address and management node IP addresses belong to.
- **Default Gateway:** Specify the default gateway for the network that the Virtual IP address and management node IP addresses belong to.
- **NTP:** Specify the NTP server that the management nodes and other appliance components must use to synchronize their clocks to.

Figure 2.5 Data Center Network Tab

The screenshot shows the Oracle PCA Admin Console interface. At the top, there is a blue header with the Oracle logo and 'PCA Admin Console'. A 'MENU' icon is on the left, and 'Open OVM Manager...' with a user icon is on the right. The main content area is titled 'Network Environment' and includes a warning: 'Changes to the Network Environment configuration data requires entering the PCA Admin Password before selecting the 'Apply Changes' button.' Below this, there are three tabs: 'Management Nodes', 'Data Center Network' (which is selected), and 'DNS'. The 'Data Center Network' tab contains the following configuration options:

- Management Network VLAN:** A toggle switch is turned on, and a text field contains the value '20'.
- * Domain Name:** A text field containing 'example.com'.
- * Netmask:** A text field containing '255.255.0.0'.
- * Default Gateway:** A text field containing '10.19.1.1'.
- * NTP:** A text field containing '10.19.1.1'.

At the bottom of the configuration area, there is a password field labeled '* PCA Admin Password(Required for changes)', a 'Reset' button, and a blue 'Apply Changes' button.

Figure 2.6 shows the Data Center Network tab. The following fields are available for configuration:

- **DNS Server 1:** Specify at least one DNS server that the management nodes can use for domain name resolution.
- **DNS Server 2:** Optionally, specify a second DNS server.
- **DNS Server 3:** Optionally, specify a third DNS server.

Figure 2.6 DNS Tab

The screenshot shows the Oracle PCA Admin Console interface for the DNS configuration tab. The header includes the Oracle logo and 'PCA Admin Console'. Below the header, there are three tabs: 'Management Nodes', 'Data Center Network', and 'DNS'. The 'DNS' tab is active. The main content area contains three rows for DNS servers:

Label	Value
* DNS Server 1	217.6.34.47
DNS Server 2	144.20.190.70
DNS Server 3	

At the bottom of the form, there is a password field labeled '* PCA Admin Password(Required for changes)', a 'Reset' button, and an 'Apply Changes' button.

You must enter the current Oracle Private Cloud Appliance *Admin* account password to make changes to any of these settings. Clicking the **Apply Changes** button at the bottom of the page saves the settings that are currently filled out in all three Network Environment tabs, and updates the configuration on each of the management nodes. The *ovca* services are restarted in the process, so you are required to log back in to the Dashboard afterwards.

2.5 Functional Networking Limitations

There are different levels and areas of network configuration in an Oracle Private Cloud Appliance environment. For the correct operation of both the host infrastructure and the virtualized environment it is critical that the administrator can make a functional distinction between the different categories of networking, and knows how and where to configure all of them. This section is intended as guidance to select the suitable interface to perform the main network administration operations.

In terms of functionality, practically all networks operate either at the appliance level or the virtualization level. Each has its own administrative interface: Oracle Private Cloud Appliance Dashboard and CLI on the one hand, and Oracle VM Manager on the other. However, the network configuration is not as clearly separated, because networking in Oracle VM depends heavily on existing configuration at the infrastructure level. For example, configuring a new public virtual machine network in Oracle VM Manager

requires that the hosts or compute nodes have network ports already connected to an underlying network with a gateway to the data center network or internet.

A significant amount of configuration – networking and other – is pushed from the appliance level to Oracle VM during compute node provisioning. This implies that a hierarchy exists; that appliance-level configuration operations must be explored before you consider making changes in Oracle VM Manager beyond the standard virtual machine management.

Network Architecture Differences

Oracle Private Cloud Appliance exists in two different types of network architecture. One is built around a physical InfiniBand fabric; the other relies on physical high speed Ethernet connectivity. While the two implementations offer practically the same functionality, there are visible hardware and configuration differences.

This section is split up by network architecture to avoid confusion. Refer to the subsection that applies to your appliance.

2.5.1 Network Configuration of Ethernet-based Systems

This section describes the Oracle Private Cloud Appliance and Oracle VM network configuration for systems with an Ethernet-based network architecture.

- **Virtual Machine Network**

By default, a fully provisioned Oracle Private Cloud Appliance is ready for virtual machine deployment. In Oracle VM Manager you can connect virtual machines to these networks directly:

- `default_external`, created on the `vx13040` VxLAN interfaces of all compute nodes during provisioning
- `default_internal`, created on the `vx2` VxLAN interfaces of all compute nodes during provisioning

Also, you can create additional VLAN interfaces and VLANs with the *Virtual Machine* role. For virtual machines requiring public connectivity, use the compute nodes' `vx13040` VxLAN interfaces. For internal-only VM traffic, use the `vx2` VxLAN interfaces. For details, see [Section 5.6, “Configuring Network Resources for Virtual Machines”](#).



Note

Do not create virtual machine networks using the `ethx` ports. These are detected in Oracle VM Manager as physical compute node network interfaces, but they are not cabled. Also, the `bondx` ports and default VLAN interfaces (`tun-ext`, `tun-int`, `mgmt-int` and `storage-int`) that appear in Oracle VM Manager are part of the appliance infrastructure networking, and are not intended to be used in VM network configurations.

Virtual machine networking can be further diversified and segregated by means of custom networks, which are described below. Custom networks must be created in the Oracle Private Cloud Appliance CLI. This generates additional VxLAN interfaces equivalent to the default `vx13040` and `vx2`. The custom networks and associated network interfaces are automatically set up in Oracle VM Manager, where you can expand the virtual machine network configuration with those newly discovered network resources.

- **Custom Network**

Custom networks are infrastructure networks you create in addition to the default configuration. These are constructed in the same way as the default private and public networks, but using different compute node network interfaces and terminating on different spine switch ports. Whenever public connectivity is required, additional cabling between the spine switches and the next-level data center switches is required.

Because they are part of the infrastructure underlying Oracle VM, all custom networks must be configured through the Oracle Private Cloud Appliance CLI. The administrator chooses between three types: private, public or host network. For detailed information about the purpose and configuration of each type, see [Section 2.6, “Network Customization”](#).

If your environment has additional tenant groups, which are separate Oracle VM server pools, then a custom network can be associated with one or more tenant groups. This allows you to securely separate traffic belonging to different tenant groups and the virtual machines deployed as part of them. For details, see [Section 2.7, “Tenant Groups”](#).

Once custom networks have been fully configured through the Oracle Private Cloud Appliance CLI, the networks and associated ports automatically appear in Oracle VM Manager. There, additional VLAN interfaces can be configured on top of the new VxLAN interfaces, and then used to create more VLANs for virtual machine connectivity. The host network is a special type of custom public network, which can assume the *Storage* network role and can be used to connect external storage directly to compute nodes.

- **Network Properties**

The network role is a property used within Oracle VM. Most of the networks you configure, have the *Virtual Machine* role, although you could decide to use a separate network for storage connectivity or virtual machine migration. Network roles – and other properties such as name and description, which interfaces are connected, properties of the interfaces and so on – can be configured in Oracle VM Manager, as long as they do not conflict with properties defined at the appliance level.

Modifying network properties of the VM networks you configured in Oracle VM Manager involves little risk. However, you must **not** change the configuration – such as network roles, ports and so on – of the default networks: `eth_management`, `mgmt_internal`, `storage_internal`, `underlay_external`, `underlay_internal`, `default_external`, and `default_internal`. For networks connecting compute nodes, including custom networks, you must use the Oracle Private Cloud Appliance CLI. Furthermore, you cannot modify the functional properties of a custom network: you have to delete it and create a new one with the required properties.

The maximum transfer unit (MTU) of a network interface, standard port or bond, cannot be modified. It is determined by the hardware properties or the SDN configuration, which cannot be controlled from within Oracle VM Manager.

- **VLAN Management**

With the exception of the underlay VLAN networks configured through SDN, and the appliance management VLAN you configure in the Network Settings tab of the Oracle Private Cloud Appliance Dashboard, all VLAN configuration and management operations are performed in Oracle VM Manager. These VLANs are part of the VM networking.

**Tip**

When a large number of VLANs is required, it is good practice not to generate them all at once, because the process is time-consuming. Instead, add (or remove) VLANs in groups of 10.

2.5.2 Network Configuration of InfiniBand-based Systems

This section describes the Oracle Private Cloud Appliance and Oracle VM network configuration for systems with an InfiniBand-based network architecture.

• Virtual Machine Network

By default, a fully provisioned Oracle Private Cloud Appliance is ready for virtual machine deployment. In Oracle VM Manager you can connect virtual machines to these networks directly:

- `vm_public_vlan`, created on the `bond4` interfaces of all compute nodes during provisioning
- `vm_private`, created on the `bond3` interfaces of all compute nodes during provisioning

Also, you can create additional VLAN interfaces and VLANs with the *Virtual Machine* role. For virtual machines requiring public connectivity, use the compute nodes' `bond4` ports. For internal-only VM traffic, use the `bond3` ports. For details, see [Section 5.6, "Configuring Network Resources for Virtual Machines"](#).



Note

Do not create virtual machine networks using the `ethx` ports. These are detected in Oracle VM Manager as physical compute node network interfaces, but they are not cabled. Also, most network interfaces are combined in pairs to form bond ports, and are not intended to be connected individually.

Virtual machine networking can be further diversified and segregated by means of custom networks, which are described below. Custom networks must be created in the Oracle Private Cloud Appliance CLI. This generates additional bond ports equivalent to the default `bond3` and `bond4`. The custom networks and associated bond ports are automatically set up in Oracle VM Manager, where you can expand the virtual machine network configuration with those newly discovered network resources.

• Custom Network

Custom networks are infrastructure networks you create in addition to the default configuration. These are constructed in the same way as the default private and public networks, but using different compute node bond ports and terminating on different Fabric Interconnect I/O ports. Whenever public connectivity is required, additional cabling between the I/O ports and the next-level data center switches is required.

Because they are part of the infrastructure underlying Oracle VM, all custom networks must be configured through the Oracle Private Cloud Appliance CLI. The administrator chooses between three types: private, public or host network. For detailed information about the purpose and configuration of each type, see [Section 2.6, "Network Customization"](#).

If your environment has tenant groups, which are separate Oracle VM server pools, then a custom network can be associated with one or more tenant groups. This allows you to securely separate traffic belonging to different tenant groups and the virtual machines deployed as part of them. For details, see [Section 2.7, "Tenant Groups"](#).

Once custom networks have been fully configured through the Oracle Private Cloud Appliance CLI, the networks and associated ports automatically appear in Oracle VM Manager. There, additional VLAN interfaces can be configured on top of the new bond ports, and then used to create more VLANs for virtual machine connectivity. The host network is a special type of custom public network, which can assume the *Storage* network role and can be used to connect external storage directly to compute nodes.

• Network Properties

The network role is a property used within Oracle VM. Most of the networks you configure, have the *Virtual Machine* role, although you could decide to use a separate network for storage connectivity or virtual machine migration. Network roles – and other properties such as name and description, which interfaces are connected, properties of the interfaces and so on – can be configured in Oracle VM Manager, as long as they do not conflict with properties defined at the appliance level.

Modifying network properties of the VM networks you configured in Oracle VM Manager involves little risk. However, you must **not** change the configuration – such as network roles, ports and so on – of the default networks: `mgmt_public_eth`, `192.168.140.0`, `192.168.40.0`, `vm_public_vlan` and `vm_private`. For networks connecting compute nodes, including custom networks, you must use the Oracle Private Cloud Appliance CLI. Furthermore, you cannot modify the functional properties of a custom network: you have to delete it and create a new one with the required properties.

The maximum transfer unit (MTU) of a network interface, standard port or bond, cannot be modified. It is determined by the hardware properties or the Fabric Interconnect configuration, which cannot be controlled from within Oracle VM Manager.

• VLAN Management

With the exception of the appliance management VLAN, which is configured in the Network Settings tab of the Oracle Private Cloud Appliance Dashboard, all VLAN configuration and management operations are performed in Oracle VM Manager. These VLANs are part of the VM networking.



Tip

When a large number of VLANs is required, it is good practice not to generate them all at once, because the process is time-consuming. Instead, add (or remove) VLANs in groups of 10.

2.6 Network Customization

The Oracle Private Cloud Appliance controller software allows you to add custom networks at the appliance level. This means that certain hardware components require configuration changes to enable the additional connectivity. The new networks are then configured automatically in your Oracle VM environment, where they can be used for isolating and optimizing network traffic beyond the capabilities of the default network configuration. All custom networks, both internal and public, are VLAN-capable.



Warning

Do not modify the network configuration while upgrade operations are running. No management operations are supported during upgrade, as these may lead to configuration inconsistencies and significant repair downtime.



Warning

Custom networks must never be deleted in Oracle VM Manager. Doing so would leave the environment in an error state that is extremely difficult to repair. To avoid downtime and data loss, always perform custom network operations in the Oracle Private Cloud Appliance CLI.



Caution

The following network limitations apply:

- The maximum number of custom external networks is 7 per tenant group or per compute node.
- The maximum number of custom internal networks is 3 per tenant group or per compute node.
- The maximum number of VLANs is 256 per tenant group or per compute node.
- Only one host network can be assigned per tenant group or per compute node.



Caution

When configuring custom networks, make sure that no provisioning operations or virtual machine environment modifications take place. This might lock Oracle VM resources and cause your Oracle Private Cloud Appliance CLI commands to fail.

Creating custom networks requires use of the CLI. The administrator chooses between three types: a network internal to the appliance, a network with external connectivity, or a host network. Custom networks appear automatically in Oracle VM Manager. The internal and external networks take the *virtual machine* network role, while a host network may have the *virtual machine* and *storage* network roles.

The host network is a particular type of external network: its configuration contains additional parameters for subnet and routing. The servers connected to it also receive an IP address in that subnet, and consequently can connect to an external network device. The host network is particularly useful for direct access to storage devices.

Network Architecture Differences

Oracle Private Cloud Appliance exists in two different types of network architecture. One is built around a physical InfiniBand fabric; the other relies on physical high speed Ethernet connectivity. While the two implementations offer practically the same functionality, the configuration of custom networks is different due to the type of network hardware.

This section is split up by network architecture to avoid confusion. Refer to the subsection that applies to your appliance.

2.6.1 Configuring Custom Networks on Ethernet-based Systems

This section describes how to configure custom networks on a system with an Ethernet-based network architecture.

For all networks with external connectivity, the spine Cisco Nexus 9336C-FX2 Switch ports must be specified so that these are reconfigured to route the external traffic. These ports must be cabled to create the physical uplink to the next-level switches in the data center. For detailed information, refer to [Appliance Uplink Configuration](#) in the *Oracle Private Cloud Appliance Installation Guide*.

Creating a Custom Network on an Ethernet-based System

1. Using SSH and an account with superuser privileges, log into the active management node.



Note

The default `root` password is *Welcome1*. For security reasons, you must set a new password at your earliest convenience.

```
# ssh root@10.100.1.101
```

```
root@10.100.1.101's password:
root@ovcamn05r1 ~]#
```

2. Launch the Oracle Private Cloud Appliance command line interface.

```
# pca-admin
Welcome to PCA! Release: 2.4.2
PCA>
```

3. If your custom network requires public connectivity, you need to use one or more spine switch ports. Verify the number of ports available and carefully plan your network customizations accordingly. The following example shows how to retrieve that information from your system:

```
PCA> list network-port
```

Port	Switch	Type	State	Networks
1:1	ovcasw22r1	10G	down	None
1:2	ovcasw22r1	10G	down	None
1:3	ovcasw22r1	10G	down	None
1:4	ovcasw22r1	10G	down	None
2	ovcasw22r1	40G	up	None
3	ovcasw22r1	auto-speed	down	None
4	ovcasw22r1	auto-speed	down	None
5:1	ovcasw22r1	10G	up	default_external
5:2	ovcasw22r1	10G	down	default_external
5:3	ovcasw22r1	10G	down	None
5:4	ovcasw22r1	10G	down	None
1:1	ovcasw23r1	10G	down	None
1:2	ovcasw23r1	10G	down	None
1:3	ovcasw23r1	10G	down	None
1:4	ovcasw23r1	10G	down	None
2	ovcasw23r1	40G	up	None
3	ovcasw23r1	auto-speed	down	None
4	ovcasw23r1	auto-speed	down	None
5:1	ovcasw23r1	10G	up	default_external
5:2	ovcasw23r1	10G	down	default_external
5:3	ovcasw23r1	10G	down	None
5:4	ovcasw23r1	10G	down	None

```
-----
22 rows displayed

Status: Success
```

4. For a custom network with external connectivity, configure an uplink port group with the uplink ports you wish to use for this traffic. Select the appropriate breakout mode

```
PCA> create uplink-port-group MyUplinkPortGroup '1:1 1:2' 10g-4x
Status: Success
```



Note

The port arguments are specified as 'x:y' where *x* is the switch port number and *y* is the number of the breakout port, in case a splitter cable is attached to the switch port. The example above shows how to retrieve that information.

You must set the breakout mode of the uplink port group. When a 4-way breakout cable is used, all four ports must be set to either 10Gbit or 25Gbit. When no breakout cable is used, the port speed for the uplink port group should be either 100Gbit or 40Gbit, depending on connectivity requirements. See [Section 4.2.9, “create uplink-port-group”](#) for command details.

Network ports can not be part of more than one network configuration.

5. Create a new network and select one of these types:

- `rack_internal_network`
- `external_network`
- `host_network`

Use the following syntax:

- For an internal-only network, specify a network name.

```
PCA> create network MyInternalNetwork rack_internal_network
Status: Success
```

- For an external network, specify a network name and the spine switch port group to be configured for external traffic.

```
PCA> create network MyPublicNetwork external_network MyUplinkPortGroup
Status: Success
```

- For a host network, specify a network name, the spine switch ports to be configured for external traffic, the subnet, and optionally the routing configuration.

```
PCA> create network MyHostNetwork host_network MyUplinkPortGroup \
10.10.10 255.255.255.0 10.1.20.0/24 10.10.10.250
Status: Success
```



Note

In this example the additional network and routing arguments for the host network are specified as follows, separated by spaces:

- `10.10.10` = subnet prefix
- `255.255.255.0` = netmask
- `10.1.20.0/24` = route destination (as subnet or IPv4 address)
- `10.10.10.250` = route gateway

The subnet prefix and netmask are used to assign IP addresses to servers joining the network. The optional route gateway and destination parameters are used to configure a static route in the server's routing table. The route destination is a single IP address by default, so you must specify a netmask if traffic could be intended for different IP addresses in a subnet.

When you define a host network, it is possible to enter invalid or contradictory values for the Prefix, Netmask and Route_Destination parameters. For example, when you enter a prefix with "0" as the first octet, the system attempts to configure IP addresses on compute node Ethernet interfaces starting with 0. Also, when the netmask part of the route destination you enter is invalid, the network is still created, even though an exception occurs. When such a poorly configured network is in an invalid state, it cannot be

reconfigured or deleted with standard commands. If an invalid network configuration is applied, use the `--force` option to delete the network.

Details of the create network command arguments are provided in [Section 4.2.7, “create network”](#) in the CLI reference chapter.



Caution

Network and routing parameters of a host network cannot be modified. To change these settings, delete the custom network and re-create it with updated settings.

6. Connect the required servers to the new custom network. You must provide the network name and the names of the servers to connect.

```
PCA> add network MyPublicNetwork ovcacn07r1
Status: Success
PCA> add network MyPublicNetwork ovcacn08r1
Status: Success
PCA> add network MyPublicNetwork ovcacn09r1
Status: Success
```

7. Verify the configuration of the new custom network.

```
PCA> show network MyPublicNetwork
-----
Network_Name      MyPublicNetwork
Trunkmode         None
Description       None
Ports             ['1:1', '1:2']
vNICs            None
Status           ready
Network_Type      external_network
Compute_Nodes     ovcacn07r1, ovcacn08r1, ovcacn09r1
Prefix           None
Netmask          None
Route Destination None
Route Gateway    None
-----
```

Status: Success

As a result of these commands, a VxLAN interface is configured on each of the servers to connect them to the new custom network. These configuration changes are reflected in the **Networking** tab and the **Servers and VMs** tab in Oracle VM Manager.



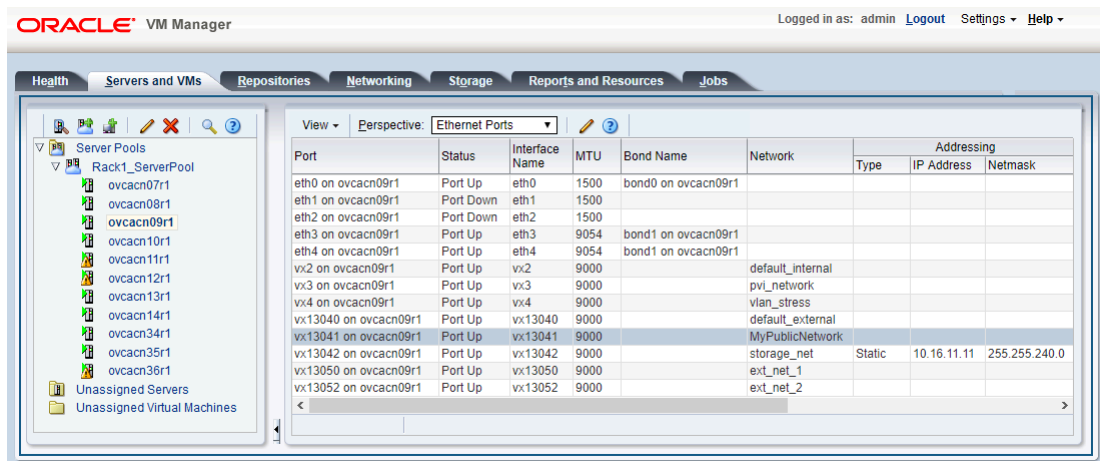
Note

If the custom network is a host network, the server is assigned an IP address based on the prefix and netmask parameters of the network configuration, and the final octet of the server's internal management IP address.

For example, if the compute node with internal IP address 192.168.4.9 were connected to the host network used for illustration purposes in this procedure, it would receive the address 10.10.10.9 in the host network.

Figure 2.7 shows a custom network named *MyPublicNetwork*, which is VLAN-capable and uses the compute node's *vx13041* interface.

Figure 2.7 Oracle VM Manager View of Custom Network Configuration (Ethernet-based Architecture)



- To disconnect servers from the custom network use the *remove network* command.



Warning

Before removing the network connection of a server, make sure that no virtual machines are relying on this network.

When a server is no longer connected to a custom network, make sure that its port configuration is cleaned up in Oracle VM.

```
PCA> remove network MyPublicNetwork ovcacn09r1
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Success
```

2.6.2 Configuring Custom Networks on InfiniBand-based Systems

This section describes how to configure custom networks on a system with an InfiniBand-based network architecture.

For all networks with external connectivity the Fabric Interconnect I/O ports must be specified so that these are reconfigured to route the external traffic. These ports must be cabled to create the physical uplink to the next-level switches in the data center.

Creating a Custom Network on an InfiniBand-based System

- Using SSH and an account with superuser privileges, log into the active management node.



Note

The default `root` password is *Welcome1*. For security reasons, you must set a new password at your earliest convenience.

```
# ssh root@10.100.1.101
root@10.100.1.101's password:
root@ovcamn05r1 ~|#
```

- Launch the Oracle Private Cloud Appliance command line interface.

```
# pca-admin
Welcome to PCA! Release: 2.4.2
PCA>
```

- If your custom network requires public connectivity, you need to use one or more Fabric Interconnect ports. Verify the number of I/O modules and ports available and carefully plan your network customizations accordingly. The following example shows how to retrieve that information from your system:

```
PCA> list network-card --sorted-by Director

Slot   Director          Type                                     State   Number_Of_Ports
-----
      3  ovcasw15r1        sanFc2Port8GbLrCardEthIb              up           2
     18  ovcasw15r1        sanFc2Port8GbLrCardEthIb              up           2
     16  ovcasw15r1        nwEthernet4Port10GbCardEthIb          up           4
      5  ovcasw15r1        nwEthernet4Port10GbCardEthIb          up           4
     17  ovcasw15r1        nwEthernet4Port10GbCardEthIb          up           4
      4  ovcasw15r1        nwEthernet4Port10GbCardEthIb          up           4
     16  ovcasw22r1        nwEthernet4Port10GbCardEthIb          up           4
      5  ovcasw22r1        nwEthernet4Port10GbCardEthIb          up           4
     18  ovcasw22r1        sanFc2Port8GbLrCardEthIb              up           2
     17  ovcasw22r1        nwEthernet4Port10GbCardEthIb          up           4
      4  ovcasw22r1        nwEthernet4Port10GbCardEthIb          up           4
      3  ovcasw22r1        sanFc2Port8GbLrCardEthIb              up           2
-----
12 rows displayed

Status: Success
PCA> list network-port --filter-column Type --filter nwEthernet* --sorted-by State

Port    Director          Type                                     State   Networks
-----
4:4     ovcasw15r1        nwEthernet10GbPort                     down    None
4:3     ovcasw15r1        nwEthernet10GbPort                     down    None
4:2     ovcasw15r1        nwEthernet10GbPort                     down    None
5:4     ovcasw15r1        nwEthernet10GbPort                     down    None
5:3     ovcasw15r1        nwEthernet10GbPort                     down    None
5:2     ovcasw15r1        nwEthernet10GbPort                     down    None
10:4    ovcasw15r1        nwEthernet10GbPort                     down    None
10:3    ovcasw15r1        nwEthernet10GbPort                     down    None
10:2    ovcasw15r1        nwEthernet10GbPort                     down    None
```

```

10:1      ovcasw15r1      nwEthernet10GbPort      down      None
11:4      ovcasw15r1      nwEthernet10GbPort      down      None
11:3      ovcasw15r1      nwEthernet10GbPort      down      None
11:2      ovcasw15r1      nwEthernet10GbPort      down      None
11:1      ovcasw15r1      nwEthernet10GbPort      down      None
4:4       ovcasw22r1      nwEthernet10GbPort      down      None
4:3       ovcasw22r1      nwEthernet10GbPort      down      None
4:2       ovcasw22r1      nwEthernet10GbPort      down      None
5:4       ovcasw22r1      nwEthernet10GbPort      down      None
5:3       ovcasw22r1      nwEthernet10GbPort      down      None
5:2       ovcasw22r1      nwEthernet10GbPort      down      None
10:4      ovcasw22r1      nwEthernet10GbPort      down      None
10:3      ovcasw22r1      nwEthernet10GbPort      down      None
10:1      ovcasw22r1      nwEthernet10GbPort      down      None
11:3      ovcasw22r1      nwEthernet10GbPort      down      None
11:2      ovcasw22r1      nwEthernet10GbPort      down      None
11:1      ovcasw22r1      nwEthernet10GbPort      down      None
4:1       ovcasw15r1      nwEthernet10GbPort      up        mgmt_public_eth, vm_public_vlan
5:1       ovcasw15r1      nwEthernet10GbPort      up        mgmt_public_eth, vm_public_vlan
4:1       ovcasw22r1      nwEthernet10GbPort      up        mgmt_public_eth, vm_public_vlan
5:1       ovcasw22r1      nwEthernet10GbPort      up        mgmt_public_eth, vm_public_vlan
10:2      ovcasw22r1      nwEthernet10GbPort      up        None
11:4      ovcasw22r1      nwEthernet10GbPort      up        None
-----
32 rows displayed

Status: Success
    
```

4. Create a new network and select one of these types:

- [rack_internal_network](#)
- [external_network](#)
- [host_network](#)

Use the following syntax:

- For an internal-only network, specify a network name.

```

PCA> create network MyInternalNetwork rack_internal_network
Status: Success
    
```

- For an external network, specify a network name and the Fabric Interconnect port(s) to be configured for external traffic.

```

PCA> create network MyPublicNetwork external_network '4:2 5:2'
Status: Success
    
```



Note

The port arguments are specified as '*x:y*' where *x* is the I/O module slot number and *y* is the number of the port on that module. The example above shows how to retrieve that information.

I/O ports can not be part of more than one network configuration.

If, instead of using the CLI interactive mode, you create a network in a single CLI command from the Oracle Linux prompt, you must escape the quotation

marks to prevent bash from interpreting them. Add a backslash character before each quotation mark:

```
# pca-admin create network MyPublicNetwork external_network \'4:2 5:2\'
```

- For a host network, specify a network name, the Fabric Interconnect ports to be configured for external traffic, the subnet, and optionally the routing configuration.

```
PCA> create network MyHostNetwork host_network '10:1 11:1' \
10.10.10 255.255.255.0 10.1.20.0/24 10.10.10.250
Status: Success
```



Note

In this example the additional network and routing arguments for the host network are specified as follows, separated by spaces:

- `10.10.10` = subnet prefix
- `255.255.255.0` = netmask
- `10.1.20.0/24` = route destination (as subnet or IPv4 address)
- `10.10.10.250` = route gateway

The subnet prefix and netmask are used to assign IP addresses to servers joining the network. The optional route gateway and destination parameters are used to configure a static route in the server's routing table. The route destination is a single IP address by default, so you must specify a netmask if traffic could be intended for different IP addresses in a subnet.

When you define a host network, it is possible to enter invalid or contradictory values for the Prefix, Netmask and Route_Destination parameters. For example, when you enter a prefix with "0" as the first octet, the system attempts to configure IP addresses on compute node Ethernet interfaces starting with 0. Also, when the netmask part of the route destination you enter is invalid, the network is still created, even though an exception occurs. When such a poorly configured network is in an invalid state, it cannot be reconfigured or deleted with standard commands. If an invalid network configuration is applied, use the `--force` option to delete the network.

Details of the create network command arguments are provided in [Section 4.2.7, "create network"](#) in the CLI reference chapter.



Caution

Network and routing parameters of a host network cannot be modified. To change these settings, delete the custom network and re-create it with updated settings.

5. Connect the required servers to the new custom network. You must provide the network name and the names of the servers to connect.

```
PCA> add network MyPublicNetwork ovcacn07r1
Status: Success
PCA> add network MyPublicNetwork ovcacn08r1
Status: Success
PCA> add network MyPublicNetwork ovcacn09r1
```

Status: Success

6. Verify the configuration of the new custom network.

```
PCA> show network MyPublicNetwork
```

```
-----
Network_Name      MyPublicNetwork
Trunkmode         True
Description       User defined network
Ports             ['4:2', '5:2']
vNICs             ovcacn09r1-eth8, ovcacn07r1-eth8, ovcacn08r1-eth8
Status            ready
Network_Type      external_network
Compute_Nodes     ovcacn07r1, ovcacn08r1, ovcacn09r1
Prefix            None
Netmask           None
Route Destination None
Route Gateway     None
-----
```

Status: Success

As a result of these commands, a bond of two new vNICs is configured on each of the servers to connect them to the new custom network. These configuration changes are reflected in the **Networking** tab and the **Servers and VMs** tab in Oracle VM Manager.



Note

If the custom network is a host network, the server is assigned an IP address based on the prefix and netmask parameters of the network configuration, and the final octet of the server's internal management IP address.

For example, if the compute node with internal IP address 192.168.4.9 were connected to the host network used for illustration purposes in this procedure, it would receive the address 10.10.10.9 in the host network.

Figure 2.8 shows a custom network named *MyPublicNetwork*, which is VLAN-enabled and uses the compute node's *bond5* interface consisting of Ethernet ports (vNICs) *eth8* and *eth8B*.

Figure 2.8 Oracle VM Manager View of Custom Network Configuration (InfiniBand-based Architecture)

Port	Status	Interface Name	Ethernet Ports	Networking		Addressing	
				Network	Type	IP Address	Netmask
bond0 on ovcacn07r1	Port Up	bond0	eth4 on ovcacn07r1, eth4B on ...	192.168.140.0	Static	192.168.140.8	255.255.255.0
bond1 on ovcacn07r1	Port Up	bond1		192.168.40.0	Static	192.168.40.8	255.255.255.0
bond2 on ovcacn07r1	Port Up	bond2	eth5 on ovcacn07r1, eth5B on ...	mgmt_public_eth			
bond3 on ovcacn07r1	Port Up	bond3	eth6 on ovcacn07r1, eth6B on ...	vm_private			
bond4 on ovcacn07r1	Port Up	bond4	eth7 on ovcacn07r1, eth7B on ...	vm_public_vlan			
bond5 on ovcacn07r1	Port Up	bond5	eth8 on ovcacn07r1, eth8B on ...	MyPublicNetwork			

7. To disconnect servers from the custom network use the *remove network* command.



Warning

Before removing the network connection of a server, make sure that no virtual machines are relying on this network.

When a server is no longer connected to a custom network, make sure that its port configuration is cleaned up in Oracle VM.

```
PCA> remove network MyPublicNetwork ovcacn09r1
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Success
```

2.6.3 Deleting Custom Networks

This section describes how to delete custom networks. The procedure is the same for systems with an Ethernet-based and InfiniBand-based network architecture.

Deleting a Custom Network



Caution

Before deleting a custom network, make sure that all servers have been disconnected from it first.

1. Using SSH and an account with superuser privileges, log into the active management node.



Note

The default `root` password is `Welcome1`. For security reasons, you must set a new password at your earliest convenience.

```
# ssh root@10.100.1.101
root@10.100.1.101's password:
root@ovcamn05r1 ~]#
```

2. Launch the Oracle Private Cloud Appliance command line interface.

```
# pca-admin
Welcome to PCA! Release: 2.4.2
PCA>
```

3. Verify that all servers have been disconnected from the custom network. No vNICs or nodes should appear in the network configuration.



Caution

Related configuration changes in Oracle VM must be cleaned up as well.



Note

The command output sample below shows a public network configuration on an Ethernet-based system. The configuration of a public network on an InfiniBand-based system looks slightly different.

```
PCA> show network MyPublicNetwork
```

```
-----
Network_Name      MyPublicNetwork
Trunkmode         None
Description       None
Ports            ['1:1', '1:2']
vNICs            None
Status           ready
Network_Type     external_network
Compute_Nodes    None
Prefix           None
Netmask          None
Route_Destination None
Route_Gateway    None
-----
```

4. Delete the custom network.

```
PCA> delete network MyPublicNetwork
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Success
```



Caution

If a custom network is left in an invalid or error state, and the delete command fails, you may use the `--force` option and retry.

2.7 Tenant Groups

A standard Oracle Private Cloud Appliance environment built on a full rack configuration contains 25 compute nodes. A *tenant group* is a logical subset of a single Oracle Private Cloud Appliance environment. Tenant groups provide an optional mechanism for an Oracle Private Cloud Appliance administrator to subdivide the environment in arbitrary ways for manageability and isolation. The tenant group offers a means to isolate compute, network and storage resources per end customer. It also offers isolation from cluster faults.

2.7.1 Design Assumptions and Restrictions

Oracle Private Cloud Appliance supports a maximum of 8 tenant groups. This number includes the default tenant group, which cannot be deleted from the environment, and must always contain at least one compute node. Therefore, a single custom tenant group can contain up to 24 compute nodes, while the default *Rack1_ServerPool* can contain all 25.

Regardless of tenant group membership, all compute nodes are connected to all of the default Oracle Private Cloud Appliance networks. Custom networks can be assigned to multiple tenant groups. When a compute node joins a tenant group, it is also connected to the custom networks associated with the tenant group. When you remove a compute node from a tenant group, it is disconnected from those custom networks. A synchronization mechanism, built into the tenant group functionality, keeps compute node network connections up to date when tenant group configurations change.

When you reprovision compute nodes, they are automatically removed from their tenant groups, and treated as new servers. Consequently, when a compute node is reprovisioned, or when a new compute node is added to the environment, it is added automatically to *Rack1_ServerPool*. After successful provisioning you can add the compute node to the appropriate tenant group.

2.7.2 Configuring Tenant Groups

The tenant group functionality can be accessed through the Oracle Private Cloud Appliance CLI. With a specific set of commands you manage the tenant groups, their member compute nodes, and the associated custom networks. The CLI initiates a number of Oracle VM operations to set up the server pool, and a synchronization service maintains settings across the members of the tenant group.



Warning

Do not modify the tenant group configuration while upgrade operations are running. No management operations are supported during upgrade, as these may lead to configuration inconsistencies and significant repair downtime.



Caution

You must not modify the server pool in Oracle VM Manager because this causes inconsistencies in the tenant group configuration and disrupts the operation of the synchronization service and the Oracle Private Cloud Appliance CLI. Only server pool policies may be edited in Oracle VM Manager.

If you inadvertently used Oracle VM Manager to modify a tenant group, see [Section 7.14, “Recovering from Tenant Group Configuration Mismatches”](#).



Note

For detailed information about the Oracle Private Cloud Appliance CLI tenant group commands, see [Chapter 4, The Oracle Private Cloud Appliance Command Line Interface \(CLI\)](#).



Note

The command output samples in this section reflect the network configuration on an Ethernet-based system. The network-related properties of a tenant group look slightly different on an InfiniBand-based system.

Creating and Populating a Tenant Group

- Using SSH and an account with superuser privileges, log into the active management node.



Note

The default `root` password is `Welcome1`. For security reasons, you must set a new password at your earliest convenience.

```
# ssh root@10.100.1.101
root@10.100.1.101's password:
root@ovcamn05r1 ~]#
```

- Launch the Oracle Private Cloud Appliance command line interface.

```
# pca-admin
Welcome to PCA! Release: 2.4.2
PCA>
```

- Create the new tenant group.

```
PCA> create tenant-group myTenantGroup
Status: Success

PCA> show tenant-group myTenantGroup
```

```

-----
Name                myTenantGroup
Default             False
Tenant_Group_ID    0004fb00000200008154bf592c8ac33b
Servers            None
State              ready
Tenant_Group_VIP   None
Tenant_Networks    ['storage_internal', 'mgmt_internal', 'underlay_internal', 'underlay_external',
                    'default_external', 'default_internal']
Pool_FileSystem_ID 3600144f0d04414f400005cf529410003
-----
Status: Success

```

The new tenant group appears in Oracle VM Manager as a new server pool. It has a 12GB server pool file system located on the internal ZFS storage appliance.

4. Add compute nodes to the tenant group.

If a compute node is currently part of another tenant group, it is first removed from that tenant group.



Caution

If the compute node is hosting virtual machines, or if storage repositories are presented to the compute node or its current tenant group, removing a compute node from an existing tenant group will fail. If so, you have to migrate the virtual machines and un-present the repositories before adding the compute node to a new tenant group.

```

PCA> add compute-node ovcacn07r1 myTenantGroup
Status: Success

PCA> add compute-node ovcacn09r1 myTenantGroup
Status: Success

```

5. Add a custom network to the tenant group.

```

PCA> add network-to-tenant-group myPublicNetwork myTenantGroup
Status: Success

```

Custom networks can be added to the tenant group as a whole. This command creates synchronization tasks to configure custom networks on each server in the tenant group.



Caution

While synchronization tasks are running, make sure that no reboot or provisioning operations are started on any of the compute nodes involved in the configuration changes.

6. Verify the configuration of the new tenant group.

```

PCA> show tenant-group myTenantGroup
-----
Name                myTenantGroup
Default             False
Tenant_Group_ID    0004fb00000200008154bf592c8ac33b
Servers            ['ovcacn07r1', 'ovcacn09r1']
State              ready
Tenant_Group_VIP   None
Tenant_Networks    ['storage_internal', 'mgmt_internal', 'underlay_internal', 'underlay_external',
                    'default_external', 'default_internal']
-----

```

```
Pool_FileSystem_ID      'default_external', 'default_internal', 'myPublicNetwork']
3600144f0d04414f400005cf529410003
```

```
-----
Status: Success
```

The new tenant group corresponds with an Oracle VM server pool with the same name and has a pool file system. The command output also shows that the servers and custom network were added successfully.

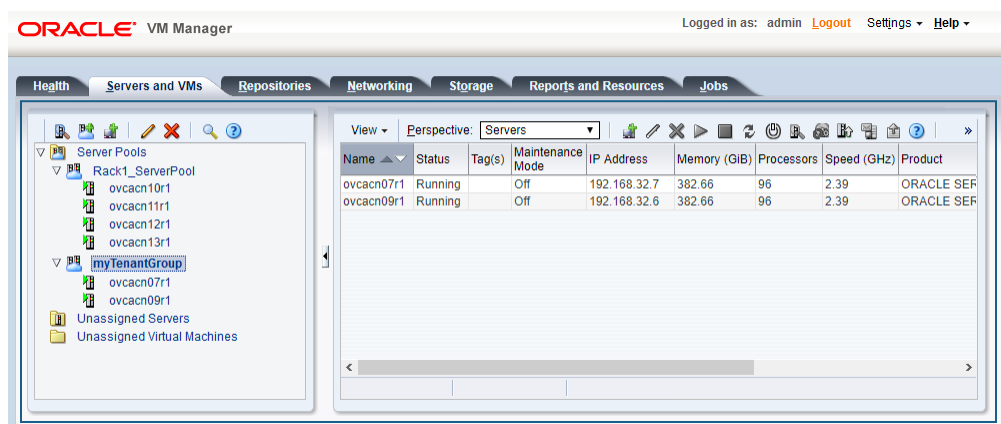
These configuration changes are reflected in the **Servers and VMs** tab in Oracle VM Manager. [Figure 2.9](#) shows a second server pool named *MyTenantGroup*, which contains the two compute nodes that were added as examples in the course of this procedure.



Note

The system does not create a storage repository for a new tenant group. An administrator must configure the necessary storage resources for virtual machines in Oracle VM Manager. See [Section 5.7, “Viewing and Managing Storage Resources”](#).

Figure 2.9 Oracle VM Manager View of New Tenant Group



Reconfiguring and Deleting a Tenant Group

1. Identify the tenant group you intend to modify.

```
PCA> list tenant-group

Name                Default      State
-----
Rack1_ServerPool   True        ready
myTenantGroup       False       ready

2 rows displayed

Status: Success

PCA> show tenant-group myTenantGroup

-----
Name                myTenantGroup
Default             False
Tenant_Group_ID    0004fb00000200008154bf592c8ac33b
Servers             ['ovcacn07r1', 'ovcacn09r1']
State               ready
```

```
Tenant_Group_VIP      None
Tenant_Networks      ['storage_internal', 'mgmt_internal', 'underlay_internal', 'underlay_external',
'default_external', 'default_internal', 'myPublicNetwork']
Pool_Fileystem_ID    3600144f0d04414f400005cf529410003
-----
Status: Success
```

2. Remove a network from the tenant group.

A custom network that has been associated with a tenant group can be removed again. The command results in serial operations, not using the synchronization service, to unconfigure the custom network on each compute node in the tenant group.

```
PCA> remove network-from-tenant-group myPublicNetwork myTenantGroup
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Success
```

3. Remove a compute node from the tenant group.

Use Oracle VM Manager to prepare the compute node for removal from the tenant group. Make sure that virtual machines have been migrated away from the compute node, and that no storage repositories are presented.

```
PCA> remove server ovcach09r1 myTenantGroup
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Success
```

When you remove a compute node from a tenant group, any custom network associated with the tenant group is automatically removed from the compute node network configuration. Custom networks that are not associated with the tenant group are not removed.

4. Delete the tenant group.

Before attempting to delete a tenant group, make sure that all compute nodes have been removed.

```
PCA> delete tenant-group myTenantGroup
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Success
```

When the tenant group is deleted, operations are launched to remove the server pool file system LUN from the internal ZFS storage appliance. The tenant group's associated custom networks are not destroyed.

2.8 Authentication

The **Password Management** window is used to reset the global Oracle Private Cloud Appliance password and to set unique passwords for individual components within the appliance. All actions performed via this tab require that you enter the current password for the Oracle Private Cloud Appliance `admin` user in the

field labelled **Current PCA Admin Password:**. Fields are available to specify the new password value and to confirm the value:

- **Current PCA Admin Password:** You must provide the current password for the Oracle Private Cloud Appliance `admin` user before any password changes can be applied.
- **New Password:** Provide the value for the new password that you are setting.
- **Verify Password:** Confirm the new password and check that you have not mis-typed what you intended.

The window provides a series of check boxes that make it easy to select the level of granularity that you wish to apply to a password change. By clicking **Select All** you can apply a global password to all components that are used in the appliance. This action resets any individual passwords that you may have set for particular components. For stricter controls, you may set the password for individual components by simply selecting the check box associated with each component that you wish to apply a password to.



Caution

When applying a password change, allow 10 minutes for the change to take effect. Do not attempt any further password changes during this 10 minute delay.

- **Select All:** Apply the new password to all components. All components in the list are selected.
- **Oracle VM Manager/PCA admin password:** Set the new password for the Oracle VM Manager and Oracle Private Cloud Appliance Dashboard admin user.
- **Oracle MySQL password:** Set the new password for the ovs user in MySQL used by Oracle VM Manager.
- **Oracle WebLogic Server password:** Set the new password for the weblogic user in WebLogic Server.
- **Oracle Data Network Leaf Switch admin password:** Set the new password for the admin user for the leaf Cisco Nexus 9336C-FX2 Switches.



Note

On InfiniBand-based systems, the list contains three separate password settings for the data network leaf switches, which are NM2-36P Sun Datacenter InfiniBand Expansion Switches:

- The **Leaf Switch root password** check box sets the password for the root user for the NM2-36P Sun Datacenter InfiniBand Expansion Switches.
 - The **Leaf Switch ILOM admin password** check box sets the password for the admin user for the ILOM of the NM2-36P Sun Datacenter InfiniBand Expansion Switches.
 - The **Leaf Switch ILOM operator password** check box sets the password for the operator user for the ILOM of the NM2-36P Sun Datacenter InfiniBand Expansion Switches.
- **Oracle Management Network Switch admin password:** Set the new password for the admin user for the Cisco Nexus 9348GC-FXP Switch.



Note

On InfiniBand-based systems, this setting applies to the root user for the Oracle Switch ES1-24 switches.

- **Oracle Data Network Spine Switch admin password:** Set the new password for the admin user for the *spine* Cisco Nexus 9336C-FX2 Switches.



Note

On InfiniBand-based systems, the list contains three separate password settings for the data network spine switches, which are Oracle Fabric Interconnect F1-15 devices:

- The **Spine Switch admin password** check box sets the password for the admin user for the Oracle Fabric Interconnect F1-15s.
 - The **Spine Switch recovery password** sets the password for recovery operations on the Oracle Fabric Interconnect F1-15s. This password is used in the case of a corruption or when the admin password is lost. The Fabric Interconnects can be booted in 'recovery mode' and this password can be used to access the recovery mode menu.
 - The **Spine Switch root password** check box sets the password for the root user for the Oracle Fabric Interconnect F1-15s.
- **Oracle ZFS Storage root password:** Set the new password for the root user for the ZFS storage appliance.
 - **PCA Management Node root password:** Set the new password for the root user for both management nodes.
 - **PCA Compute Node root password:** Set the new password for the root user for all compute nodes.
 - **PCA Management Node SP/ILOM root password:** Set the new password for the root user for the ILOM on both management nodes.
 - **PCA Compute Node SP/ILOM root password:** Set the new password for the root user for the ILOM on all compute nodes.

Figure 2.10 Password Management

The screenshot shows the Oracle PCA Admin Console interface for password management. The header includes the Oracle logo and 'PCA Admin Console'. The main content area is titled 'Password Management' and has a sub-tab 'Change Passwords'. On the left, there are three password input fields: '* Current PCA Admin Password', '* New Password', and '* Verify New Password'. On the right, under 'Apply New Password To:', there is a 'Select All' checkbox and a list of components with checkboxes, all of which are checked. The components listed are: Oracle VM Manager/PCA Admin password, Oracle MySQL password, Oracle Weblogic Server password, Oracle Data Network Leaf Switch admin password, Oracle Management Network Switch admin password, Oracle Data Network Spine Switch admin password, Oracle ZFS Storage root password, PCA Management Node root password, PCA Compute Node root password, PCA Management Node SP/ILOM root password, and PCA Compute Node SP/ILOM root password. At the bottom right, there are 'Reset' and 'Apply Changes' buttons.

The functionality that is available in the Oracle Private Cloud Appliance Dashboard is equally available via the Oracle Private Cloud Appliance CLI as described in [Section 4.2.30, “update password”](#).



Caution

Passwords of components must not be changed manually as this will cause mismatches with the authentication details stored in the Oracle Private Cloud Appliance Wallet.

2.9 Health Monitoring

The Oracle Private Cloud Appliance Controller Software contains a monitoring service, which is started and stopped with the `ovca` service on the active management node. When the system runs for the first time it creates an *inventory database* and *monitor database*. Once these are set up and the monitoring service is active, health information about the hardware components is updated continuously.

The inventory database is populated with information about the various components installed in the rack, including the IP addresses to be used for monitoring. With this information, the *ping manager* pings all known components every 3 minutes and updates the inventory database to indicate whether a component is pingable and when it was last seen online. When errors occur they are logged in the monitor database. Error information is retrieved from the component ILOMs.

For troubleshooting purposes, historic health status details can be retrieved through the CLI support mode by an **authorized Oracle Field Engineer**. When the CLI is used in support mode, a number of additional commands are available; two of which are used to display the contents of the health monitoring databases.

- Use `show db inventory` to display component health status information from the inventory database.
- Use `show db monitor` to display errors logged in the monitoring database.

The appliance administrator can retrieve current component health status information from the Oracle Linux command line on the master management node, using the Oracle Private Cloud Appliance Health Check utility. The Health Check utility is built on the framework of the Oracle Private Cloud Appliance Upgrader, and is included in the Upgrader package. It detects the appliance network architecture and runs the sets of health checks defined for the system in question.

Checking the Current Health Status of an Oracle Private Cloud Appliance Installation

1. Using SSH and an account with superuser privileges, log in to the active management node.



Note

The default `root` password is *Welcome1*. For security reasons, you must set a new password at your earliest convenience.

```
# ssh root@10.100.1.101
root@10.100.1.101's password:
root@ovcamn05r1 ~]#
```

2. Launch the Health Check utility.

```
# pca_healthcheck
PCA Rack Type: PCA X8_BASE.
Please refer to log file
/nfs/shared_storage/pca_upgrader/log/pca_healthcheck_2019_10_04-12.09.45.log
for more details.
```

After detecting the rack type, the utility executes the applicable health checks.

```
Beginning PCA Health Checks...

Check Management Nodes Are Running                1/24
Check Support Packages                            2/24
Check PCA DBs Exist                               3/24
PCA Config File                                    4/24
Check Shares Mounted on Management Nodes           5/24
Check PCA Version                                  6/24
Check Installed Packages                           7/24
Check for OpenSSL CVE-2014-0160 - Security Update  8/24
Management Nodes Have IPv6 Disabled                9/24
Check Oracle VM Manager Version                   10/24
Oracle VM Manager Default Networks                 11/24
Repositories Defined in Oracle VM Manager          12/24
PCA Services                                       13/24
Oracle VM Server Model                             14/24
Network Interfaces on Compute Nodes                15/24
Oracle VM Manager Settings                         16/24
Check Network Leaf Switch                          17/24
Check Network Spine Switch                         18/24
All Compute Nodes Running                          19/24
Test for ovs-agent Service on Compute Nodes        20/24
Test for Shares Mounted on Compute Nodes           21/24
Check for bash ELSA-2014-1306 - Security Update    22/24
Check Compute Node's Active Network Interfaces      23/24
Checking for xen OVMSA-2014-0026 - Security Update  24/24

PCA Health Checks completed after 2 minutes
```

3. When the health checks have been completed, check the report for failures.

```

Check Management Nodes Are Running           Passed
Check Support Packages                       Passed
Check PCA DBs Exist                          Passed
PCA Config File                              Passed
Check Shares Mounted on Management Nodes     Passed
Check PCA Version                            Passed
Check Installed Packages                     Passed
Check for OpenSSL CVE-2014-0160 - Security Update Passed
Management Nodes Have IPv6 Disabled         Passed
Check Oracle VM Manager Version             Passed
Oracle VM Manager Default Networks          Passed
Repositories Defined in Oracle VM Manager   Passed
PCA Services                                 Passed
Oracle VM Server Model                       Passed
Network Interfaces on Compute Nodes         Passed
Oracle VM Manager Settings                  Passed
Check Network Leaf Switch                   Passed
Check Network Spine Switch                  Failed
All Compute Nodes Running                   Passed
Test for ovs-agent Service on Compute Nodes Passed
Test for Shares Mounted on Compute Nodes     Passed
Check for bash ELSA-2014-1306 - Security Update Passed
Check Compute Node's Active Network Interfaces Passed
Checking for xen OVMSA-2014-0026 - Security Update Passed

-----
Overall Status                               Failed
-----

Please refer to log file
/nfs/shared_storage/pca_upgrader/log/pca_healthcheck_2019_10_04-12.09.45.log
for more details.

```

4. If certain checks have resulted in failures, review the log file for additional diagnostic information. Search for text strings such as "error" and "failed".

```

# grep -inr "failed" /nfs/shared_storage/pca_upgrader/log/pca_healthcheck_2019_10_04-12.09.45.log

726:[2019-10-04 12:10:51 264234] INFO (healthcheck:254) Check Network Spine Switch Failed -
731: Spine Switch ovcasw22r1 North-South Management Network Port-channel check [FAILED]
733: Spine Switch ovcasw22r1 Multicast Route Check [FAILED]
742: Spine Switch ovcasw23r1 North-South Management Network Port-channel check [FAILED]
750:[2019-10-04 12:10:51 264234] ERROR (precheck:148) [Check Network Spine Switch ()] Failed
955:[2019-10-04 12:12:26 264234] INFO (precheck:116) [Check Network Spine Switch ()] Failed

# less /nfs/shared_storage/pca_upgrader/log/pca_healthcheck_2019_10_04-12.09.45.log

[...]
Spine Switch ovcasw22r1 North-South Management Network Port-channel check [FAILED]
Spine Switch ovcasw22r1 OSPF Neighbor Check [OK]
Spine Switch ovcasw22r1 Multicast Route Check [FAILED]
Spine Switch ovcasw22r1 PIM RP Check [OK]
Spine Switch ovcasw22r1 NVE Peer Check [OK]
Spine Switch ovcasw22r1 Spine Filesystem Check [OK]
Spine Switch ovcasw22r1 Hardware Diagnostic Check [OK]
[...]

```

5. Investigate and fix any detected problems. Repeat the health check until the system passes all checks.

Chapter 3 Updating Oracle Private Cloud Appliance

Table of Contents

3.1 Before You Start Updating	61
3.1.1 Warnings and Cautions	62
3.1.2 Backup Prevents Data Loss	64
3.2 Using the Oracle Private Cloud Appliance Upgrader	64
3.2.1 Rebooting the Management Node Cluster	65
3.2.2 Installing the Oracle Private Cloud Appliance Upgrader	65
3.2.3 Verifying Upgrade Readiness	66
3.2.4 Executing a Controller Software Update	68
3.3 Upgrading the Virtualization Platform	75
3.4 Upgrading Component Firmware	78
3.4.1 Firmware Policy	78
3.4.2 Install the Current Firmware on All Compute Nodes	79
3.4.3 Upgrading the Operating Software on the Oracle ZFS Storage Appliance	79
3.4.4 Upgrading the Cisco Switch Firmware	83
3.4.5 Upgrading the NM2-36P Sun Datacenter InfiniBand Expansion Switch Firmware	86
3.4.6 Upgrading the Oracle Fabric Interconnect F1-15 Firmware	89

Due to the nature of the Oracle Private Cloud Appliance – where the term *appliance* is key – an update is a delicate and complicated procedure that deals with different hardware and software components at the same time. It is virtually impossible to automate the entire process, and more importantly it would be undesirable to take the appliance and the virtual environment it hosts out of service entirely for updating. Instead, updates can be executed in phases and scheduled for minimal downtime. The following table explains how an Oracle Private Cloud Appliance update handles different levels or areas of appliance functionality.

Table 3.1 Functional Break-Down of an Appliance Update

Functionality	Physical Location	Description
controller software	management nodes	all components required to set up the management cluster, manage and configure the appliance, and orchestrate compute node provisioning
virtualization platform	compute nodes	all components required to configure the compute nodes and allow virtual machines to be hosted on them
component firmware	infrastructure components	all low-level software components required by the various hardware components for their normal operation as part of the appliance

3.1 Before You Start Updating

Please read and observe the critical information in this section before you begin any procedure to update your Oracle Private Cloud Appliance.

All the software included in a given release of the Oracle Private Cloud Appliance software is tested to work together and should be treated as one package. Consequently, no appliance component should

be updated individually, unless Oracle provides specific instructions to do so. All Oracle Private Cloud Appliance software releases are downloaded as a single large `.iso` file, which includes the items listed above.



Note

The appliance update process must **always** be initiated from the **master management node**.

3.1.1 Warnings and Cautions

Read and understand these warnings and cautions before you start the appliance update procedure. They help you avoid operational issues including data loss and significant downtime.



Minimum Release

In this version of the Oracle Private Cloud Appliance Administrator's Guide, it is assumed that your system is currently **running Controller Software release 2.3.4 or 2.4.1 prior to this software update**.

If your system is currently running an earlier version, please refer to the [Update chapter](#) of the [Administrator's Guide for Release 2.3](#). Follow the appropriate procedures and make sure that your appliance configuration is valid for the Release 2.4.2 update.



No Critical Operations

When updating the Oracle Private Cloud Appliance software, make sure that no provisioning operations occur and that any externally scheduled backups are suspended. Such operations could cause a software update or component firmware upgrade to fail and lead to system downtime.



YUM Disabled

On Oracle Private Cloud Appliance management nodes the YUM repositories have been intentionally disabled and should not be enabled by the customer. Updates and upgrades of the management node operating system and software components must only be applied through the update mechanism described in the documentation.



Firmware Policy

To ensure that your Oracle Private Cloud Appliance configuration remains in a qualified state, take the required firmware upgrades into account when planning the controller software update. For more information, refer to [Section 3.4.1, "Firmware Policy"](#).



No Backup

During controller software updates, backup operations must be prevented. The Oracle Private Cloud Appliance Upgrader disables `crond` and blocks backups.



CA Certificate and Keystore

If you have generated custom keys using `ovmkeytool.sh` in a previous version of the Oracle Private Cloud Appliance software, you must regenerate the keys prior to

updating the Controller Software. For instructions, refer to the support note with [Doc ID 2597439.1](#). See also [Section 7.9.1, “Creating a Keystore”](#).



Proxy Settings

If direct public access is not available within your data center and you make use of proxy servers to facilitate HTTP, HTTPS and FTP traffic, it may be necessary to edit the Oracle Private Cloud Appliance system properties, using the CLI on each management node, to ensure that the correct proxy settings are specified for a download to succeed from the Internet. This depends on the network location from where the download is served. See [Section 7.2, “Adding Proxy Settings for Oracle Private Cloud Appliance Updates”](#) for more information.



Custom LUNs on Internal Storage

If the internal ZFS Storage Appliance contains customer-created LUNs, make sure they are not mapped to the default initiator group. See [Customer Created LUNs Are Mapped to the Wrong Initiator Group](#) in the Oracle Private Cloud Appliance Release Notes.



Oracle VM Availability During Update to Release 2.4.x

When updating the Oracle Private Cloud Appliance Controller Software to Release 2.4.x, Oracle VM Manager is unavailable for the entire duration of the update. The virtualized environment remains functional, but configuration changes and other management operations are not possible.



Compute Node Upgrade ONLY Through Oracle Private Cloud Appliance CLI

Compute nodes cannot be upgraded to the appropriate Oracle VM Server Release 3.4.x with the Oracle VM Manager web UI. You must upgrade them using the `update compute-node` command within the Oracle Private Cloud Appliance CLI.

To perform this CLI-based upgrade procedure, follow the specific instructions in [Section 3.3, “Upgrading the Virtualization Platform”](#).



Do Not Override Oracle VM Global Update Settings

As stated in [Section 5.1, “Guidelines and Limitations”](#), at the start of [Chapter 5, *Managing the Oracle VM Virtual Infrastructure*](#), the settings of the default server pool and custom tenant groups must not be modified through Oracle VM Manager. For compute node upgrade specifically, it is critical that the server pool option "Override Global Server Update Group" remains deselected. The compute node update process must use the repository defined globally, and overriding this will cause the update to fail.



Post-Update Synchronization

Once you have confirmed that the update process has completed, it is advised that you wait a further 30 minutes before starting another compute node or management node software update. This allows the necessary synchronization tasks to complete.

If you ignore the recommended delay between these update procedures there could be issues with further updating as a result of interference between existing and new tasks.

3.1.2 Backup Prevents Data Loss

An update of the Oracle Private Cloud Appliance software stack may involve a complete re-imaging of the management nodes. Any customer-installed agents or customizations are overwritten in the process. Before applying new appliance software, back up all local customizations and prepare to re-apply them after the update has completed successfully.

Oracle Enterprise Manager Plug-in Users

If you use Oracle Enterprise Manager and the Oracle Enterprise Manager Plug-in to monitor your Oracle Private Cloud Appliance environment, always back up the *oraInventory* Agent data to `/nfs/shared_storage` before updating the controller software. You can restore the data after the Oracle Private Cloud Appliance software update is complete.

For detailed instructions, refer to the [Agent Recovery](#) section in the [Oracle Enterprise Manager Plug-in documentation](#).

3.2 Using the Oracle Private Cloud Appliance Upgrader



UPGRADE BOTH MANAGEMENT NODES CONSECUTIVELY

With the Oracle Private Cloud Appliance Upgrader, the two management node upgrade processes are theoretically separated. Each management node upgrade is initiated by a single command and managed through the Upgrader, which invokes the native Oracle VM Manager upgrade mechanisms. However, you must **treat the upgrade of the two management nodes as a single operation**.

During the management node upgrade, the high-availability (HA) configuration of the management node cluster is temporarily broken. To restore HA management functionality and mitigate the risk of data corruption, it is critical that you **start the upgrade of the second management node immediately after a successful upgrade of the first management node**.



NO MANAGEMENT OPERATIONS DURING UPGRADE

The Oracle Private Cloud Appliance Upgrader manages the entire process to upgrade both management nodes in the appliance. Under no circumstances should you perform any management operations – through the Oracle Private Cloud Appliance Dashboard or CLI, or Oracle VM Manager – while the Upgrader process is running, and until **both management nodes** have been successfully upgraded through the Upgrader. Although certain management functions cannot be programmatically locked during the upgrade, they are not supported, and are likely to cause configuration inconsistencies and considerable repair downtime.

Once the upgrade has been successfully completed on **both management nodes**, you can safely execute appliance management tasks and configuration of the virtualized environment.

As of Release 2.3.4, a separate command line tool is provided to manage the Controller Software update process. A version of the Oracle Private Cloud Appliance Upgrader is included in the Controller Software `.iso` image. However, Oracle recommends that you download and install the latest stand-alone version of the Upgrader tool on the management nodes. The Oracle Private Cloud Appliance Upgrader requires only a couple of commands to execute several sets of tasks, which were scripted or manual steps in previous releases. The Upgrader is more robust and easily extensible, and provides a much better overall upgrade experience.

A more detailed description of the Oracle Private Cloud Appliance Upgrader is included in the introductory chapter of this book. Refer to [Section 1.7, “Oracle Private Cloud Appliance Upgrader”](#).

3.2.1 Rebooting the Management Node Cluster

It is advised to reboot both management nodes before starting the appliance software update. This leaves the management node cluster in the cleanest possible state, ensures that no system resources are occupied unnecessarily, and eliminates potential interference from processes that have not completed properly.

Rebooting the Management Node Cluster

1. Using SSH and an account with superuser privileges, log into both management nodes using the IP addresses you configured in the Network Setup tab of the Oracle Private Cloud Appliance Dashboard. If you use two separate consoles you can view both side by side.



Note

The default `root` password is `Welcome1`. For security reasons, you must set a new password at your earliest convenience.

2. Run the command `pca-check-master` on both management nodes to verify which node owns the master role.
3. Reboot the management node that is **NOT** currently the master. Enter `init 6` at the prompt.
4. Ping the machine you rebooted. When it comes back online, reconnect using SSH and monitor system activity to determine when the secondary management node takes over the master role. Enter this command at the prompt: `tail -f /var/log/messages`. New system activity notifications will be output to the screen as they are logged.
5. In the other SSH console, which is connected to the current active management node, enter `init 6` to reboot the machine and initiate management node failover.

The log messages in the other SSH console should now indicate when the secondary management node takes over the master role.

6. Verify that both management nodes have come back online after reboot and that the master role has been transferred to the other manager. Run the command `pca-check-master` on both management nodes.

If this is the case, proceed with the software update steps below.

3.2.2 Installing the Oracle Private Cloud Appliance Upgrader

The Oracle Private Cloud Appliance Upgrader is a separate application with its own release plan, independent of Oracle Private Cloud Appliance. Always download and install the latest version of the Oracle Private Cloud Appliance Upgrader before you execute any verification or upgrade procedures.

Downloading and Installing the Latest Version of the Oracle Private Cloud Appliance Upgrader

1. Log into [My Oracle Support](#) and download the latest version of the Oracle Private Cloud Appliance Upgrader.

The Upgrader can be found under patch ID 30459450, and is included in part 1 of a series of downloadable zip files. Any updated versions of the Upgrader will be made available in the same location.

To obtain the Upgrader package, download this zip file and extract the file `pca_upgrader-<version>.el6.noarch.rpm`.

2. Copy the downloaded `*.rpm` package to the master management node and install it.

```
[root@ovcamn05r1 ~]# pca-check-master
NODE: 192.168.4.3 MASTER: True
root@ovcamn05r1 tmp]# rpm -ivh pca_upgrader-1.2-89.el6.noarch.rpm
Preparing...##### [100%]
1:pca_upgrader##### [100%]
```



Caution

Always download and use the latest available version of the Oracle Private Cloud Appliance Upgrader.

3. If the version of the Oracle Private Cloud Appliance Upgrader you downloaded, is newer than the version shipped in the Controller Software ISO, then upgrade to the newer version. From the directory where the `*.rpm` package was saved, run the command `rpm -U pca_upgrader-1.2-89.el6.noarch.rpm`.
4. Repeat the `*.rpm` upgrade on the second management node.

The Oracle Private Cloud Appliance Upgrader verifies the version of the Upgrader installed on the second management node. Only if the version in the ISO is newer, the package on the second management node is automatically upgraded in the process. If you downloaded a newer version, you must upgrade the package manually on both management nodes.

3.2.3 Verifying Upgrade Readiness

The Oracle Private Cloud Appliance Upgrader has a verify-only mode. It allows you to run all the pre-checks defined for a management node upgrade, without proceeding to the actual upgrade steps. The terminal output and log file report any issues you need to fix before the system is eligible for the next Controller Software update.



Note

The Oracle Private Cloud Appliance Upgrader cannot be stopped by means of a keyboard interrupt or by closing the terminal session. After a keyboard interrupt (`Ctrl+C`) the Upgrader continues to execute all pre-checks. If the terminal session is closed, the Upgrader continues as a background process.

If the Upgrader process needs to be terminated, enter this command `pca_upgrader --kill`.

Verifying the Upgrade Readiness of the Oracle Private Cloud Appliance

1. Go to Oracle VM Manager and make sure that all compute nodes are in *Running* status. If any server is not in Running status, resolve the issue before proceeding. For instructions to correct the compute node status, refer to the support note with [Doc ID 2245197.1](#).
2. Perform the required manual pre-upgrade checks. Refer to [Section 7.5, “Running Manual Pre- and Post-Upgrade Checks in Combination with Oracle Private Cloud Appliance Upgrader”](#) for instructions.
3. Log in to [My Oracle Support](#) and download the required Oracle Private Cloud Appliance software update.

You can find the update by searching for the product name “Oracle Private Cloud Appliance”, or for the Patch or Bug Number associated with the update you need.



Caution

Read the information and follow the instructions in the [readme](#) file very carefully. It is crucial for a successful Oracle Private Cloud Appliance Controller Software update and Oracle VM upgrade.

4. Make the update, a zipped ISO, available on an HTTP or FTP server that is reachable from your Oracle Private Cloud Appliance. Alternatively, if upgrade time is a major concern, you can download the ISO file to the local file system on both management nodes. This reduces the upgrade time for the management nodes, but has no effect on the time required to upgrade the compute nodes or the Oracle VM database.

The Oracle Private Cloud Appliance Upgrader downloads the ISO from the specified location and unpacks it on the management node automatically at runtime.

5. Using SSH and an account with superuser privileges, log in to the **master** management node through its individually assigned IP address, **not** the shared virtual IP.



Note

During the upgrade process, the interface with the shared virtual IP address is shut down. Therefore, you must log in using the individually assigned IP address of the management node.

The default `root` password is `Welcome1`. For security reasons, you must set a new password at your earliest convenience.

6. From the master management node, run the Oracle Private Cloud Appliance Upgrader in verify-only mode. The target of the command must be the *stand-by* management node.



Note

The console output below is an example. You may see a different output, depending on the specific architecture and configuration of your appliance.

```
[root@ovcamn05r1 ~]# pca-check-master
NODE: 192.168.4.3 MASTER: True

root@ovcamn05r1 ~]# pca_upgrader -V -t management -c ovcamn06r1 -g 2.4.2 \
-l http://<path-to-iso>/ovca-2.4.2-b000.iso.zip

PCA Rack Type: PCA X8_BASE.

Please refer to log file
/nfs/shared_storage/pca_upgrader/log/pca_upgrader_2019_09_26-10.02.46.log
for more details.

Beginning PCA Management Node Pre-Upgrade Checks...

Validate the Image Provided                               1/35
Internal ZFSSA Available Space Check                    2/35
MN Disk and Shared Storage Space Check                  3/35
[...]
Oracle VM Minimum Version Check                         32/35
OS Check                                                 33/35
Password Check                                          34/35
```

```

OSA Disabled Check 35/35
PCA Management Node Pre-Upgrade Checks completed after 0 minutes
-----
PCA Management Node Pre-Upgrade Checks Passed
-----
Validate the Image Provided Passed
Internal ZFSSA Available Space Check Passed
[...]
OS Check Passed
Password Check Passed
OSA Disabled Check Passed
-----
Overall Status Passed
-----

```

7. As the verification process runs, check the console output for test progress. When all pre-checks have been completed, a summary is displayed. A complete overview of the verification process is saved in the file `/nfs/shared_storage/pca_upgrader/log/pca_upgrader_<date>-<time>.log`.

Some pre-checks may result in a warning. These warnings are unlikely to cause issues, and therefore do not prevent you from executing the upgrade, but they do indicate a situation that should be investigated. When an upgrade command is issued, warnings do cause the administrator to be prompted whether to proceed with the upgrade, or quit and investigate the warnings first.

8. If pre-checks have failed, consult the log file for details. Fix the reported problems, then execute the verify command again.



Note

If errors related to SSL certificates are reported, check whether these have been re-generated using `ovmkeytool.sh`. This can cause inconsistencies between the information stored in the Wallet and the actual location of your certificate. For detailed information and instructions to resolve the issue, refer to the support note with [Doc ID 2597439.1](#).

9. Repeat this process until no more pre-check failures are reported. When the system passes all pre-checks, it is ready for the Controller Software update.

3.2.4 Executing a Controller Software Update

During a Controller Software update, the virtualized environment does not accept any management operations. After successful upgrade of the management node cluster, the compute nodes must also be upgraded in phases, and additional firmware upgrades on several rack components may be required. When you have planned all these upgrade tasks, and when you have successfully completed the upgrade readiness verification, your environment is ready for a Controller Software update and any additional upgrades.

No upgrade procedure can be executed without completing the pre-checks. Therefore, the upgrade command first executes the same steps as in [Section 3.2.3, “Verifying Upgrade Readiness”](#). After successful verification, the upgrade steps are started.



Note

The console output shown throughout this section is an example. You may see a different output, depending on the specific architecture and configuration of your appliance.



Note

The Oracle Private Cloud Appliance Upgrader cannot be stopped by means of a keyboard interrupt or by closing the terminal session.

After a keyboard interrupt (**Ctrl+C**) the Upgrader continues the current phase of the process. If pre-checks are in progress, they are all completed, but the upgrade phase does not start automatically after successful completion of all pre-checks. If the upgrade phase is in progress at the time of the keyboard interrupt, it continues until upgrade either completes successfully or fails.

If the terminal session is closed, the Upgrader continues as a background process.

If the Upgrader process needs to be terminated, enter this command:
`pca_upgrader --kill.`

Upgrading the Oracle Private Cloud Appliance Controller Software

1. Using SSH and an account with superuser privileges, log in to the **master** management node through its individually assigned IP address, **not** the shared virtual IP.



Note

During the upgrade process, the interface with the shared virtual IP address is shut down. Therefore, you must log in using the individually assigned IP address of the management node.

The default `root` password is `Welcome1`. For security reasons, you must set a new password at your earliest convenience.



NO MANAGEMENT OPERATIONS DURING UPGRADE

Under no circumstances should you perform any management operations – through the Oracle Private Cloud Appliance Dashboard or CLI, or Oracle VM Manager – while the Upgrader process is running, and until **both management nodes** have been successfully upgraded through the Upgrader.

2. From the master management node, run the Oracle Private Cloud Appliance Upgrader with the required upgrade parameters. The target of the command must be the *stand-by* management node.

```
[root@ovcamn05r1 ~]# pca-check-master
NODE: 192.168.4.3 MASTER: True

root@ovcamn05r1 ~]# pca_upgrader -U -t management -c ovcamn06r1 -g 2.4.2 \
-l http://<path-to-iso>/ovca-2.4.2-b000.iso.zip

PCA Rack Type: PCA X8_BASE.

Please refer to log file
/nfs/shared_storage/pca_upgrader/log/pca_upgrader_2019_09_26-12.10.13.log
for more details.

Beginning PCA Management Node Pre-Upgrade Checks...
[...]
```

Warning: The management precheck completed with warnings.
It is safe to continue with the management upgrade from this point
or the upgrade can be halted to investigate the warnings.

```
Do you want to continue? [y/n]: y
```

After successfully completing the pre-checks, the Upgrader initiates the Controller Software update on the other management node. If any errors occur during the upgrade phase, tasks are rolled back and the system is returned to its original state from before the upgrade command.

Rollback works for errors that occur during these steps:

- downloading the ISO
- setting up the YUM repository
- taking an Oracle VM backup
- breaking the Oracle Private Cloud Appliance HA model

```
Beginning PCA Management Node upgrade for ovcamn06r1
Disable PCA Backups                                1/16
Download ISO                                       2/16
Setup Yum Repo                                     3/16
Take OVM Backup                                    4/16
Break PCA HA Model                                 5/16
Place PCA Upgrade Locks                           6/16
Perform Yum Upgrade                                7/16
Reboot                                              8/16
Perform Oracle VM Upgrade                          9/16
Relink to Shared Storage                           10/16
Install PCA UI                                     11/16
Run Post-Upgrade Sync Tasks                       12/16
Remove PCA Upgrade Locks                          13/16
Cleanup ISO Backup Directories                     14/16
Restore PCA Backups                                15/16
Upgrade is complete                                16/16
PCA Management Node upgrade of ovcamn06r1 completed after 43 minutes

Beginning PCA Post-Upgrade Checks...
OVM Manager Cache Size Check                       1/1
PCA Post-Upgrade Checks completed after 2 minutes

-----
PCA Management Node Pre-Upgrade Checks              Passed
-----
Validate the Image Provided                         Passed
Internal ZFSSA Available Space Check               Passed
[...]

-----
PCA Management Node Upgrade                        Passed
-----
Disable PCA Backups                                Passed
[...]
Restore PCA Backups                                Passed
Upgrade is complete                                Passed
[...]

-----
Overall Status                                     Passed
-----
PCA Management Node Pre-Upgrade Checks              Passed
PCA Management Node Upgrade                        Passed
```




Tip

When the ISO is copied to the local file system of both management nodes, the management node upgrade time is just over 1 hour each. The duration of the entire upgrade process depends heavily on the size of the environment: the number of compute nodes and their configuration, the size of the Oracle VM database, etc.

If you choose to copy the ISO locally, replace the location URL in the `pca_upgrader` command with `-l file:///<path-to-iso>/ovca-2.4.2-b000.iso.zip`.

3. Monitor the progress of the upgrade tasks. The console output provides a summary of each executed task. If you need more details on a task, or if an error occurs, consult the log file. You can track the logging activity in a separate console window by entering the command `tail -f /nfs/shared_storage/pca_upgrader/log/pca_upgrader_<date>-<time>.log`. The example below shows several key sections in a typical log file.



Note

Once the upgrade tasks have started, it is no longer possible to perform a rollback to the previous state.

```
# tail -f /nfs/shared_storage/pca_upgrader/log/pca_upgrader_2019_09_26-12.10.13.log

[2019-09-26 12:10:13 44526] INFO (pca_upgrader:59) Starting PCA Upgrader...
[2019-09-26 12:10:13 44526] INFO (validate_rack:29) Rack Type: hardware_blue
[2019-09-26 12:10:13 44526] INFO (pca_upgrader:62) PCA Rack Type: PCA X8_BASE.
[2019-09-26 12:10:13 44526] DEBUG (util:511) the dlm_locks command output is debugfs.ocfs2 1.8.6
[?1034hdebugfs: dlm_locks -f /sys/kernel/debug/o2dlm/ovca/locking_state
Lockres: master                               Owner: 0      State: 0x0
Last Used: 0      ASTs Reserved: 0      Inflight: 0      Migration Pending: No
Refs: 3      Locks: 1      On Lists: None
Reference Map: 1
  Lock-Queue Node Level Conv Cookie          Refs AST BAST Pending-Action
  Granted    0     EX   -1   0:1          2   No  No   None

debugfs: quit
[2019-09-26 12:10:13 44526] INFO (util:520) This node (192.168.4.3) is the master
[2019-09-26 12:10:14 44526] DEBUG (run_util:17) Writing 44633 to /var/run/ovca/upgrader.pid
[2019-09-26 12:10:14 44633] DEBUG (process_flow:37) Execute precheck steps for component: management
[2019-09-26 12:10:25 44633] INFO (precheck:159) [Validate the Image Provided
(Verify that the image exists and is correctly named)] Passed
[2019-09-26 12:10:25 44633] INFO (precheck_utils:471) Checking the existence of default OVMM networks.
[2019-09-26 12:10:25 44633] INFO (precheck_utils:2248) Checking for PCA services.
[2019-09-26 12:10:25 44633] INFO (precheck_utils:1970) Checking if there are multiple tenant groups.
[...]

[2019-09-26 12:10:32 44633] INFO (precheck_utils:1334) Checking hardware faults on host ovcasw16r1.
[2019-09-26 12:10:32 44633] INFO (precheck_utils:1334) Checking hardware faults on host ilom-ovcasn02r1.
[2019-09-26 12:10:32 44633] INFO (precheck_utils:1334) Checking hardware faults on host ilom-ovcamn06r1.
[2019-09-26 12:10:32 44633] INFO (precheck_utils:1334) Checking hardware faults on host ovcacn09r1.
[2019-09-26 12:10:32 44633] INFO (precheck_utils:1334) Checking hardware faults on host ovcacn08r1.
[2019-09-26 12:10:32 44633] INFO (precheck_utils:1334) Checking hardware faults on host ilom-ovcacn08r1.
[2019-09-26 12:10:32 44633] INFO (precheck:159) [Hardware Faults Check (Verifying that there are no
hardware faults on any rack component)] Passed
The check succeeded. There are no hardware faults on any rack component.
[...]

[2019-09-26 12:10:34 44633] INFO (precheck_utils:450) Checking storage...
```

Executing a Controller Software Update

```
Checking server pool...
Checking pool filesystem...
Checking rack repository Rack1-Repository...
[...]

[2019-09-26 12:10:53 44633] INFO (precheck:159) [OS Check (Checking the management nodes
and compute nodes are running the correct Oracle Linux version)] Passed
The check succeeded on the management nodes. The check succeeded on all the compute nodes.
[...]

***** PCA Management Node Pre-Upgrade Checks Summary *****
[2019-09-26 12:11:00 44633] INFO (precheck:112) [Validate the Image Provided
(Verify that the image exists and is correctly named)] Passed
[...]
[2019-09-26 12:11:00 44633] INFO (precheck:112) [OSA Disabled Check
(Checking OSA is disabled on all management nodes and compute nodes)] Passed
The check succeeded on the management nodes. The check succeeded on all the compute nodes.
[2019-09-26 12:11:00 44633] INFO (precheck:113)
***** End of Summary *****

[2019-09-26 12:11:02 44633] DEBUG (process_flow:98) Successfully completed precheck. Proceeding to upgrade.
[2019-09-26 12:11:02 44633] DEBUG (process_flow:37) Execute upgrade steps for component: management
[...]
[2019-09-26 12:25:04 44633] DEBUG (mn_upgrade_steps:237) Verifying ISO image version
[2019-09-26 12:25:04 44633] INFO (mn_upgrade_steps:243) Successfully downloaded ISO
[2019-09-26 12:25:17 44633] INFO (mn_upgrade_utils:136) /nfs/shared_storage/pca_upgrader/scripts/
1.2-89.el6/remote_yum_setup -t 2019_09_26-12.10.13 -f /nfs/shared_storage/pca_upgrader/pca_upgrade.repo
Successfully setup the yum config
[...]

[2019-09-26 12:26:44 44633] DEBUG (mn_upgrade_steps:339) Successfully completed break_ha_model
[2019-09-26 12:26:44 44633] INFO (util:184) Created lock: all_provisioning
[2019-09-26 12:26:44 44633] INFO (util:184) Created lock: database
[2019-09-26 12:26:44 44633] INFO (util:184) Created lock: cn_upgrade
[2019-09-26 12:26:44 44633] INFO (util:184) Created lock: mn_upgrade
[2019-09-26 12:26:44 44633] DEBUG (mn_upgrade_steps:148) Successfully completed place_pca_locks
[2019-09-26 12:26:45 44633] INFO (mn_upgrade_utils:461) Beginning Yum Upgrade
[...]

Setting up Upgrade Process
Resolving Dependencies
[...]
Transaction Summary
=====
Install      2 Package(s)
Upgrade     4 Package(s)

Total download size: 47 M
Downloading Packages:
[...]

[2019-09-26 12:35:40 44633] INFO (util:520) This node (192.168.4.3) is the master
[2019-09-26 12:35:40 44633] INFO (mn_upgrade_utils:706) Beginning Oracle VM upgrade on ovcam06r1
[...]
Oracle VM upgrade script finished with success
STDOUT:
Oracle VM Manager Release 3.4.6 Installer
Oracle VM Manager Installer log file:
/var/log/ovmm/ovm-manager-3-install-2019-09-26-123633.log
Verifying upgrading prerequisites ...
[...]

Running full database backup ...
Successfully backed up database to /u01/app/oracle/mysql/dbbackup/3.4.6_preUpgradeBackup-20190926_123644
Running ovm_preUpgrade script, please be patient this may take a long time ...
Exporting weblogic embedded LDAP users
Stopping service on Linux: ovmcli ...
```

```
Stopping service on Linux: ovmm ...
Exporting core database, please be patient this may take a long time ...
[...]

Installation Summary
-----
Database configuration:
  Database type           : MySQL
  Database host name      : localhost
  Database name           : ovs
  Database listener port  : 1521
  Database user           : ovs

Weblogic Server configuration:
  Administration username : weblogic

Oracle VM Manager configuration:
  Username                : admin
  Core management port     : 54321
  UUID                    : 0004fb0000010000f1b07bf678cf43d6

Passwords:
There are no default passwords for any users. The passwords to use for Oracle VM Manager,
Database, and Oracle WebLogic Server have been set by you during this installation.
In the case of a default install, all passwords are the same.
[...]
Oracle VM Manager upgrade complete.
[...]

Successfully started Oracle VM services
Preparing to install the PCA UI
[...]
Successfully installed PCA UI
[...]

[09/26/2019 12:56:56 34239] DEBUG (complete_postupgrade_tasks:228) Copying switch config files
[09/26/2019 12:56:56 34239] INFO (complete_postupgrade_tasks:231)
[09/26/2019 12:56:56 34239] INFO (update:896) Scheduling post upgrade sync tasks...
[...]

STDOUT: Looking for [tenant group] [Rack1_ServerPool]
Looking for [storage array] [OVCA_ZFSSA_Rack1]
ID: 0004fb0000020000c500310305e98353
Server: ovcacn09r1
Server: ovcacn07r1
Server: ovcacn08r1
Finding the LUN that is used as the heartbeat device in tenant group by page83id.
page83_ID: 3600144f09c52bc4200005d8c891f0003
[...]

Successfully completed PCA management node post upgrade tasks
```

When the upgrade tasks have been completed successfully, the master management node is rebooted, and the upgraded management node assumes the master role. The new master management node's operating system is now up-to-date, and it runs the new Controller Software version and upgraded Oracle VM Manager installation.



Tip

Rebooting the management node is expected to take up to 10 minutes.

To monitor the reboot process and make sure the node comes back online as expected, log in to the rebooting management node ILOM.

```
Broadcast message from root@ovcamn05r1 (pts/2) (Mon Sep 26 14:48:52 2019):
```

Management Node upgrade succeeded. The master manager will be rebooted to initiate failover in one minute.

4. Log into the upgraded management node, which has now become the **master** management node. Use its individually assigned IP address, **not** the shared virtual IP.

```
[root@ovcamn06r1 ~]# pca-check-master
NODE: 192.168.4.4 MASTER: True

[root@ovcamn06r1 ~]# head /etc/ovca-info
==== Begin build info ====
date: 2019-09-30
release: 2.4.2
build: 404
=== Begin compute node info ===
compute_ovm_server_version: 3.4.6
compute_ovm_server_build: 2.4.2-631
compute_rpms_added:
  osc-oracle-s7k-2.1.2-4.el6.noarch.rpm
  ovca-support-2.4.2-137.el6.noarch.rpm
```



NO MANAGEMENT OPERATIONS DURING UPGRADE

Under no circumstances should you perform any management operations – through the Oracle Private Cloud Appliance Dashboard or CLI, or Oracle VM Manager – while the Upgrader process is running, and until **both management nodes** have been successfully upgraded through the Upgrader.

5. From the new master management node, run the Oracle Private Cloud Appliance Upgrader command again. The target of the command must be the *stand-by* management node, which is the original master management node from where you executed the command for the first run.

```
root@ovcamn06r1 ~]# pca_upgrader -U -t management -c ovcamn05r1 -g 2.4.2 \
-l http://<path-to-iso>/ovca-2.4.2-b000.iso.zip

PCA Rack Type: PCA X8_BASE.

Please refer to log file
/nfs/shared_storage/pca_upgrader/log/pca_upgrader_2019_09_26-21.37.22.log
for more details.

Beginning PCA Management Node Pre-Upgrade Checks...
[...]

*****
Warning: The management precheck completed with warnings.
It is safe to continue with the management upgrade from this point
or the upgrade can be halted to investigate the warnings.
*****
Do you want to continue? [y/n]: y

Beginning PCA Management Node upgrade for ovcamn05r1
[...]

-----
Overall Status                                     Passed
-----
PCA Management Node Pre-Upgrade Checks             Passed
PCA Management Node Upgrade                         Passed
PCA Post-Upgrade Checks                            Passed

Broadcast message from root@ovcamn05r1 (pts/2) (Mon Sep 26 23:18:27 2019):
```

```
Management Node upgrade succeeded. The master manager will be rebooted to initiate failover in one minu
```

The upgrade steps are executed the same way as during the first run. When the second management node is rebooted, the process is complete. At this point, both management nodes run the updated Oracle Linux operating system, Oracle Private Cloud Appliance Controller Software, and Oracle VM Manager. The high-availability cluster configuration of the management nodes is restored, and all Oracle Private Cloud Appliance and Oracle VM Manager management functionality is operational again. However, do not perform any management operations until you have completed the required manual post-upgrade checks.

6. Perform the required manual post-upgrade checks on management nodes and compute nodes. Refer to [Section 7.5, “Running Manual Pre- and Post-Upgrade Checks in Combination with Oracle Private Cloud Appliance Upgrader”](#) for instructions.

When the management node cluster upgrade is complete, proceed with the following tasks:

- Compute node upgrades, as described in [Section 3.3, “Upgrading the Virtualization Platform”](#).
- Firmware upgrades, as described in [Section 3.4, “Upgrading Component Firmware”](#).

3.3 Upgrading the Virtualization Platform

Some releases of the Oracle Private Cloud Appliance Controller Software include a new version of Oracle VM, the virtualization platform used in Oracle Private Cloud Appliance. As part of the controller software update, the new Oracle VM Manager Release is automatically installed on both management nodes.

After the controller software update on the management nodes, Oracle VM Manager displays events indicating that the compute nodes are running outdated version of Oracle VM Server. These events are informational and do not prevent any operations, but it is recommended that you upgrade all compute nodes to the new Oracle VM Server Release at your earliest convenience.

The Oracle VM Server upgrade was intentionally decoupled from the automated controller software update process. This allows you to plan the compute node upgrades and the migration or downtime of your virtual machines in steps and outside peak hours. As a result, service interruptions for users of the Oracle VM environment can be minimized or even eliminated. By following the instructions in this section, you also make sure that previously deployed virtual machines remain fully functional when the appliance update to the new software release is complete.

During an upgrade of Oracle VM Server, no virtual machine can be running on a given compute node. VMs using resources on a shared storage repository can be migrated to other running compute nodes. If a VM uses resources local to the compute node you want to upgrade, it must be shut down, and returned to service after the Oracle VM Server upgrade.

When you install Oracle Private Cloud Appliance Controller Software Release 2.4.x, the management nodes are set up to run Oracle VM Manager 3.4.x. Compute nodes cannot be upgraded to the corresponding Oracle VM Server Release with the Oracle VM Manager web UI. You must upgrade them using the `update compute-node` command within the Oracle Private Cloud Appliance CLI.

Upgrading a Compute Node to a Newer Oracle VM Server Release



Caution

Execute this procedure on each compute node *after* the software update on the management nodes has completed successfully.



Caution

If compute nodes are running other packages that are not part of Oracle Private Cloud Appliance, these must be uninstalled before the Oracle VM Server upgrade.

1. Make sure that the appliance software has been updated successfully to the new release.

You can verify this by logging into the master management node and entering the following command in the Oracle Private Cloud Appliance CLI:

```
# pca-admin
Welcome to PCA! Release: 2.4.2
PCA> show version
```

```
-----
Version          2.4.2
Build            000
Date             2019-10-10
-----
```

```
Status: Success
```

Leave the console and CLI connection open. You need to run the update command later in this procedure.

2. Log in to Oracle VM Manager.

For details, see [Section 5.2, “Logging in to the Oracle VM Manager Web UI”](#).

3. Migrate all running virtual machines away from the compute node you want to upgrade.

Information on migrating virtual machines is provided in the Oracle VM Manager User's Guide section entitled [Migrate or Move Virtual Machines](#).

4. Place the compute node in maintenance mode.

Information on using maintenance mode is provided in the Oracle VM Manager User's Guide section entitled [Edit Server](#).

- a. In the **Servers and VMs** tab, select the Oracle VM Server in the navigation pane. Click **Edit Server** in the management pane toolbar.

The **Edit Server** dialog box is displayed.

- b. Select the **Server in Maintenance Mode** check box to place the Oracle VM Server into maintenance mode. Click OK.

The Oracle VM Server is in maintenance mode and ready for servicing.

5. Run the Oracle VM Server update for the compute node in question.
 - a. Return to the open management node console window with active CLI connection.
 - b. Run the `update compute-node` command for the compute nodes you wish to update at this time. Run this command for one compute node at a time.



Warning

Running the `update compute-node` command with multiple servers as arguments is not supported. Neither is running the command concurrently in separate terminal windows.

```
PCA> update compute-node ovcacn09r1
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y

Status: Success
```

This CLI command invokes a validation mechanism, which verifies critical requirements that a compute node must meet to qualify for the Oracle VM Server 3.4.x upgrade. It also ensures that all the necessary packages are installed from the correct source location, and configured properly.

- c. Wait for the command to complete successfully. The update takes approximately 30 minutes for each compute node.

As part of the update procedure, the Oracle VM Server is restarted but remains in maintenance mode.



Warning

If the compute node does not reboot during the update, you must restart it from within Oracle VM Manager.

6. Return to Oracle VM Manager to take the compute node out of maintenance mode.
 - a. In the **Servers and VMs** tab, select the Oracle VM Server in the navigation pane. Click **Edit Server** in the management pane toolbar.

The **Edit Server** dialog box is displayed.

- b. Clear the **Server in Maintenance Mode** check box. Click OK.

The Oracle VM Server rejoins the server pool as a fully functioning member.

7. Repeat this procedure for each compute node in your Oracle Private Cloud Appliance.

The appliance software update is now complete. Next, perform the required post-upgrade verification steps. The procedure for those additional manual verification tasks is documented in the *Post Upgrade* section of the support note with [Doc ID 2242177.1](#).

After successful completion of the post-upgrade verification steps, the Oracle Private Cloud Appliance is ready to resume all normal operations.

3.4 Upgrading Component Firmware

All the software components in a given Oracle Private Cloud Appliance release are designed to work together. As a general rule, no individual appliance component should be upgraded. If a firmware upgrade is required for one or more components, the correct version is distributed inside the Oracle Private Cloud Appliance `.iso` file you downloaded from [My Oracle Support](#). When the image file is unpacked on the internal shared storage, the firmwares are located in this directory: `/nfs/shared_storage/mgmt_image/firmware/`.



Warning

Do not perform any compute node provisioning operations during firmware upgrades.



Caution

For certain services it is necessary to upgrade the Hardware Management Pack after a Controller Software update. For additional information, refer to [Some Services Require an Upgrade of Hardware Management Pack](#) in the *Oracle Private Cloud Appliance Release Notes*.

If a specific or additional procedure to upgrade the firmware of an Oracle Private Cloud Appliance hardware component is available, it appears in this section. For components not listed here, you may follow the instructions provided in the product documentation of the subcomponent. An overview of the documentation for appliance components can be found in the [Preface](#) of this book and on the index page of the Oracle Private Cloud Appliance Documentation Library.

3.4.1 Firmware Policy

To improve Oracle Private Cloud Appliance supportability, reliability and security, Oracle has introduced a standardized approach to component firmware. The general rule remains unchanged: components and their respective firmware are designed to work together, and therefore should not be upgraded separately. However, the firmware upgrades, which are provided as part of the `.iso` file of a given controller software release, are no longer optional.

As part of the test process prior to a software release, combinations of component firmware are tested on all applicable hardware platforms. This allows Oracle to deliver a fully qualified set of firmware for the appliance as a whole, corresponding to a software release. In order to maintain their Oracle Private Cloud Appliance in a qualified state, customers who upgrade to a particular software release, are expected to also install all the qualified firmware upgrades delivered as part of the controller software.

The firmware versions that have been qualified by Oracle for a given release are listed in the *Oracle Private Cloud Appliance Release Notes* for that release. Please refer to the Release Notes for the Oracle Private Cloud Appliance Controller Software release running on your system, and open the chapter *Firmware Qualification*.

Note that the file names shown in the procedures below may not exactly match the file names in the `.iso` image on your system.



Interim Firmware Patches

Oracle periodically releases firmware patches for many products, for example to eliminate security vulnerabilities. It may occur that an important firmware patch is released for a component of Oracle Private Cloud Appliance outside of the normal Controller Software release schedule. When this occurs, the patches go through the

same testing as all other appliance firmware, but they are not added to the qualified firmware list or the installation `.iso` for the affected Controller Software release.

After thorough testing, important firmware patches that cannot be included in the Controller Software `.iso` image are made available to Oracle Private Cloud Appliance users through [My Oracle Support](#).

3.4.2 Install the Current Firmware on All Compute Nodes

To avoid compatibility issues with newer Oracle Private Cloud Appliance Controller Software and Oracle VM upgrades, you should always install the server ILOM firmware included in the ISO image of the current Oracle Private Cloud Appliance software release. When the ISO image is unpacked on the appliance internal storage, the firmware directory can be reached from the management nodes at this location: `/nfs/shared_storage/mgmt_image/firmware/`.

For firmware upgrade instructions, refer to the Administration Guide of the server series installed in your appliance rack.

3.4.3 Upgrading the Operating Software on the Oracle ZFS Storage Appliance

The instructions in this section are specific for a component firmware upgrade as part of the Oracle Private Cloud Appliance.



Caution

During this procedure, the Oracle Private Cloud Appliance services on the management nodes must be halted for a period of time. Plan this upgrade carefully, so that no compute node provisioning, Oracle Private Cloud Appliance configuration changes, or Oracle VM Manager operations are taking place at the same time.



Warning

For firmware upgrades to version 8.7.20 or newer, an intermediate upgrade to version 8.7.14 is required. Version 8.7.14 can then be upgraded to the intended newer version. For additional information, refer to [Oracle ZFS Storage Appliance Firmware Upgrade 8.7.20 Requires A Two-Phased Procedure](#) in the *Oracle Private Cloud Appliance Release Notes*.



Note

Detailed information about software upgrades can be found in the *Oracle ZFS Storage Appliance Customer Service Manual* (document ID: F13771). Refer to the section “Upgrading the Software”.

The Oracle Private Cloud Appliance internal ZFS Storage Appliance contains two clustered controllers in an active/passive configuration. You may disregard the upgrade information for standalone controllers.

Upgrading the ZFS Storage Appliance Operating Software

1. Before initiating the upgrade on the storage controllers, follow the preparation instructions in the *Oracle ZFS Storage Appliance Customer Service Manual*. Refer to the section entitled “Preparing for a Software Upgrade”.
2. Log on to the master management node using SSH and an account with superuser privileges.
3. Unzip the firmware package included in the Oracle Private Cloud Appliance software image.

```
[root@ovcamn05r1 ~]# mkdir /nfs/shared_storage/yum/ak
[root@ovcamn05r1 ~]# cd /nfs/shared_storage/yum/ak
[root@ovcamn05r1 ak]# unzip /nfs/shared_storage/mgmt_image/firmware/storage/AK_NAS/p27357887_20131_Generic.
Archive:  /nfs/shared_storage/mgmt_image/firmware/storage/AK_NAS/p27357887_20131_Generic.zip
extracting: ak-nas-2013.06.05.8.6-1.2.6.4214.1x-nondebug.pkg
inflating: OS8.8.6_Readme.html
```

4. Download the software update package to both storage controllers. Their management IP addresses are 192.168.4.1 and 192.168.4.2.

- a. Log on to one of the storage controllers using SSH and an account with superuser privileges.

```
[root@ovcamn05r1 ~]# ssh root@192.168.4.1
Password:
ovcasn01r1:>
```

- b. Enter the following series of commands to download the software update package from the shared storage directory to the controller.

```
ovcasn01r1:> maintenance system updates download
ovcasn01r1:maintenance system updates download (uncommitted)> \
    set url=http://192.168.4.100/shares/export/Yum/ak/ak-nas-2013.06.05.8.6-1.2.6.4214.1x-nondebug.pkg
    url = http://192.168.4.100/shares/export/Yum/ak/ak-nas-2013.06.05.8.6-1.2.6.4214.1x-non
ovcasn01r1:maintenance system updates download (uncommitted)> set user=root
    user = root
ovcasn01r1:maintenance system updates download (uncommitted)> set password
Enter password:
    password = *****
ovcasn01r1:maintenance system updates download (uncommitted)> commit
Transferred 157M of 484M (32.3%) ...
```

- c. Wait for the package to fully download and unpack before proceeding.

- d. Repeat these steps for the second storage controller.

5. Check the storage cluster configuration and make sure you are logged on to the standby controller.

```
ovcasn02r1:> configuration cluster show
Properties:
    state = AKCS_STRIPPED
    description = Ready (waiting for failback)
    peer_asn = 8a535bd2-160f-c93b-9575-d29d4c86cac5
    peer_hostname = ovcasn01r1
    peer_state = AKCS_OWNER
    peer_description = Active (takeover completed)
```

6. Always upgrade the operating software **first** on the standby controller.

- a. Display the available operating software versions and select the version you downloaded.

```
ovcasn02r1:> maintenance system updates
ovcasn02r1:maintenance system updates> show
Updates:

UPDATE                                     RELEASE DATE      RELEASE NAME      STATUS
ak-nas@2013.06.05.8.3,1-1.1                2019-1-4 09:57:19 OS8.8.3           previous
ak-nas@2013.06.05.8.5,1-1.3                2019-3-30 07:27:20 OS8.8.5           current
ak-nas@2013.06.05.8.6,1-1.4                2019-6-21 20:56:45 OS8.8.6           waiting
ovcasn02r1:maintenance system updates> select ak-nas@2013.06.05.8.6,1-1.4
```

- b. Launch the upgrade process with the selected software version.

```
ovcasn02r1:maintenance system updates> upgrade
```

This procedure will consume several minutes and requires a system reboot upon successful update, but can be aborted with [Control-C] at any time prior to reboot. A health check will validate system readiness before an update is attempted, and may also be executed independently using the check command.

Are you sure? (Y/N) Y

- c. At the end of the upgrade, when the controller has fully rebooted and rejoined the cluster, log back in and check the cluster configuration. The upgraded controller must still be in the state "Ready (waiting for failback)".

```
ovcasn02r1:> configuration cluster show
Properties:
    state = AKCS_STRIPPED
    description = Ready (waiting for failback)
    peer_asn = 8a535bd2-160f-c93b-9575-d29d4c86cac5
    peer_hostname = ovcasn01r1
    peer_state = AKCS_OWNER
    peer_description = Active (takeover completed)
```

7. From the Oracle Private Cloud Appliance master management node, stop the Oracle Private Cloud Appliance services.



Caution

Do not skip this step. Executing the storage controller operating software upgrade while the Oracle Private Cloud Appliance services are running, will result in errors and possible downtime.

```
[root@ovcamn05r1 ~]# service ovca stop
```

8. Upgrade the operating software on the second storage controller.
 - a. Check the storage cluster configuration. Make sure you are logged on to the active controller.

```
ovcasn01r1:> configuration cluster show
Properties:
    state = AKCS_OWNER
    description = Active (takeover completed)
    peer_asn = 34e4292a-71ae-6ce1-e26c-cc38c2af9719
    peer_hostname = ovcasn02r1
    peer_state = AKCS_STRIPPED
    peer_description = Ready (waiting for failback)
```

- b. Display the available operating software versions and select the version you downloaded.

```
ovcasn01r1:> maintenance system updates
ovcasn01r1:maintenance system updates> show
Updates:
UPDATE                                RELEASE DATE                RELEASE NAME                STATUS
ak-nas@2013.06.05.8.3,1-1.1           2019-1-4 09:57:19          OS8.8.3                    previous
ak-nas@2013.06.05.8.5,1-1.3           2019-3-30 07:27:20          OS8.8.5                    current
ak-nas@2013.06.05.8.6,1-1.4           2019-6-21 20:56:45          OS8.8.6                    waiting

ovcasn01r1:maintenance system updates> select ak-nas@2013.06.05.8.6,1-1.4
```

- c. Launch the upgrade process with the selected software version.

```
ovcasn01r1:maintenance system updates> upgrade
This procedure will consume several minutes and requires a system reboot upon successful update, but can be aborted with [Control-C] at any time prior to reboot. A health check will validate system readiness before an update is
```

```
attempted, and may also be executed independently using the check command.
```

```
Are you sure? (Y/N) Y
```

- d. At the end of the upgrade, when the controller has fully rebooted and rejoined the cluster, log back in and check the cluster configuration.

```
ovcasn01r1:> configuration cluster show
Properties:
                state = AKCS_STRIPPED
description = Ready (waiting for failback)
peer_asn = 34e4292a-71ae-6ce1-e26c-cc38c2af9719
peer_hostname = ovcasn02r1
peer_state = AKCS_OWNER
peer_description = Active (takeover completed)
```

The last upgraded controller must now be in the state "Ready (waiting for failback)". The controller that was upgraded first, took over the active role during the upgrade and reboot of the second controller, which held the active role originally.

9. Now that both controllers have been upgraded, verify that all disks are online.

```
ovcasn01r1:> maintenance hardware show
[...]
NAME          STATE  MANUFACTURER  MODEL          SERIAL          RPM  TYPE
chassis-000  1906NMQ803  ok            Oracle         Oracle Storage DE3-24C  1906NMQ803  7200  hdd
disk-000     HDD 0      ok            WDC            W7214A520ORA014T  001851N3VKLT  9JG3VKLT  7200  dat
disk-001     HDD 1      ok            WDC            W7214A520ORA014T  001851N5K85T  9JG5K85T  7200  dat
disk-002     HDD 2      ok            WDC            W7214A520ORA014T  001851N5MPXT  9JG5MPXT  7200  dat
disk-003     HDD 3      ok            WDC            W7214A520ORA014T  001851N5L08T  9JG5L08T  7200  dat
disk-004     HDD 4      ok            WDC            W7214A520ORA014T  001851N42KNT  9JG42KNT  7200  dat
[...]
```

10. Initiate an Oracle Private Cloud Appliance management node failover and wait until all services are restored on the other management node. This helps prevent connection issues between Oracle VM and the ZFS storage.

- a. Log on to the master management node using SSH and an account with superuser privileges.
- b. Reboot the master management node.

```
[root@ovcamn05r1 ~]# pca-check-master
NODE: 192.168.4.3 MASTER: True
[root@ovcamn05r1 ~]# shutdown -r now
```

- c. Log on to the other management node and wait until the necessary services are running.



Note

Enter this command at the prompt: `tail -f /var/log/messages`. The log messages should indicate when the management node takes over the master role.

Verify the status of the services:

```
[root@ovcamn06r1 ~]# service ovca status
Checking Oracle Fabric Manager: Running
MySQL running (70254) [ OK ]
Oracle VM Manager is running...
Oracle VM Manager CLI is running...
tinyproxy (pid 71315 71314 71313 71312 71310 71309 71308 71307 71306 71305 71301) is running...
dhcpd (pid 71333) is running...
snmptrapd (pid 71349) is running...
```

```
log server (pid 6359) is running...
remaster server (pid 6361) is running...
http server (pid 71352) is running...
taskmonitor server (pid 71356) is running...
xmlrpc server (pid 71354) is running...
nodestate server (pid 71358) is running...
sync server (pid 71360) is running...
monitor server (pid 71363) is running...
```

11. When the storage controller cluster has been upgraded, remove the shared storage directory you created to make the unzipped package available.

```
# cd /nfs/shared_storage/yum/ak
# ls
ak-nas-2013.06.05.8.6-1.2.6.4214.1x-nondebug.pkg  OS8.8.6_Readme.html
# rm ak-nas-2013.06.05.8.6-1.2.6.4214.1x-nondebug.pkg  OS8.8.6_Readme.html
rm: remove regular file `ak-nas-2013.06.05.8.6-1.2.6.4214.1x-nondebug.pkg'? yes
rm: remove regular file `OS8.8.6_Readme.html'? yes
# cd ..
# rmdir ak
#
```

3.4.4 Upgrading the Cisco Switch Firmware

The instructions in this section are specific for a component firmware upgrade as part of the Oracle Private Cloud Appliance.



Note

Cisco switches are part of systems with an Ethernet-based network architecture.

Upgrading the Firmware of all Cisco Leaf, Spine, and Management Switches

1. Log on to the master management node using SSH and an account with superuser privileges.
2. Verify that the new Cisco NX-OS firmware image is available on the appliance shared storage. During the Controller Software update, the Oracle Private Cloud Appliance Upgrader copies the file to this location:

```
/nfs/shared_storage/mgmt_image/firmware/ethernet/Cisco/nxos.7.0.3.I7.7.bin
```

3. Log on as admin to the switch you wish to upgrade.

```
root@ovcamn05r1 ~]# ssh admin@ovcasw15r1
User Access Verification
Password:
ovcasw15r1#
```

Please upgrade the switches, one at a time, in this order:

- a. Leaf Cisco Nexus 9336C-FX2 Switches: ovcasw15r1, ovcasw16r1
 - b. Spine Cisco Nexus 9336C-FX2 Switches: ovcasw22r1, ovcasw23r1
 - c. Management Cisco Nexus 9348GC-FXP Switch: ovcasw21r1
4. Copy the firmware file to the bootflash location on the switch.

The copy command for the management switch ovcasw21r1 is slightly different. Select the appropriate option.

- Leaf and Spine switches:

```

ovcasw15r1# copy scp://root@192.168.4.216//nfs/shared_storage/mgmt_image/firmware/ethernet \
/Cisco/nxos.7.0.3.I7.7.bin bootflash:nxos.7.0.3.I7.7.bin vrf management
root@192.168.4.216's password:
nxos.7.0.3.I7.7.bin                               100%  930MB  15.8MB/s   01:01
Copy complete, now saving to disk (please wait)...
Copy complete.

```

- Management switch:

```

ovcasw21r1# copy scp://root@192.168.4.216//nfs/shared_storage/mgmt_image/firmware/ethernet \
/Cisco/nxos.7.0.3.I7.7.bin bootflash:nxos.7.0.3.I7.7.bin vrf default
root@192.168.4.216's password:
nxos.7.0.3.I7.7.bin                               100%  930MB  15.8MB/s   01:01
Copy complete, now saving to disk (please wait)...
Copy complete.

```

5. Save the current running configuration as the startup configuration.

```

ovcasw15r1# copy running-config startup-config
[#####] 100%
Copy complete, now saving to disk (please wait)...
Copy complete.

```

6. Install the Cisco NX-OS software that was copied to the bootflash location. When prompted about the disruptive upgrade, enter *y* to continue with the installation.

```

ovcasw15r1# install all nxos bootflash:nxos.7.0.3.I7.7.bin
Installer will perform compatibility check first. Please wait.
Installer is forced disruptive

Verifying image bootflash:/nxos.7.0.3.I7.7.bin for boot variable "nxos".
[#####] 100% -- SUCCESS

Verifying image type.
[#####] 100% -- SUCCESS

Preparing "nxos" version info using image bootflash:/nxos.7.0.3.I7.7.bin.
[#####] 100% -- SUCCESS

Preparing "bios" version info using image bootflash:/nxos.7.0.3.I7.7.bin.
[#####] 100% -- SUCCESS

Performing module support checks.
[#####] 100% -- SUCCESS

Notifying services about system upgrade.
[#####] 100% -- SUCCESS

Compatibility check is done:
Module  bootable          Impact  Install-type  Reason
-----  -----  -----  -----  -----
      1      yes      disruptive      reset  default upgrade is not hitless

Images will be upgraded according to following table:
Module  Image                Running-Version(pri:alt)  New-Version  Upg-Required
-----  -----  -----  -----  -----
      1     nxos                7.0(3)I7(6)              7.0(3)I7(7)      yes
      1     bios      v05.33(09/08/2018):v05.28(01/18/2018)  v05.38(06/12/2019)  yes

Switch will be reloaded for disruptive upgrade.
Do you want to continue with the installation (y/n)? [n] y

Install is in progress, please wait.

Performing runtime checks.

```

```
[#####] 100% -- SUCCESS

Setting boot variables.
[#####] 100% -- SUCCESS

Performing configuration copy.
[#####] 100% -- SUCCESS

Module 1: Refreshing compact flash and upgrading bios/loader/bootrom.
Warning: please do not remove or power off the module at this time.
[#####] 100% -- SUCCESS

Finishing the upgrade, switch will reboot in 10 seconds.
```

7. Verify that the correct software version is active on the switch.

```
ovcasw15r1# show version
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.
[...]

Software
  BIOS: version 05.38
  NXOS: version 7.0(3)I7(7)
  BIOS compile time: 06/12/2019
  NXOS image file is: bootflash://nxos.7.0.3.I7.7.bin
  NXOS compile time: 8/28/2019 16:00:00 [08/29/2019 00:41:42]
[...]

ovcasw15r1#
```

8. Verify the VPC status.



Note

This step does not apply to the appliance internal management network switch (Cisco Nexus 9348GC-FXP Switch). Proceed to the next step.

Use the command shown below. The output values should match this example.

```
ovcasw15r1# show vpc brief
Legend:
          (*) - local vPC is down, forwarding via vPC peer-link

vPC domain id           : 2
Peer status              : peer adjacency formed ok <---- verify this field
vPC keep-alive status   : peer is alive <----- verify this field
Configuration consistency status : success <----- verify this field
Per-vlan consistency status : success <----- verify this field
Type-2 consistency status : success <----- verify this field
vPC role                 : secondary
Number of vPCs configured : 27
Peer Gateway             : Enabled
Dual-active excluded VLANs : -
Graceful Consistency Check : Enabled
Auto-recovery status     : Disabled
Delay-restore status     : Timer is off.(timeout = 30s)
Delay-restore SVI status : Timer is off.(timeout = 10s)
Operational Layer3 Peer-router : Enabled
```

9. Log out of the switch. The firmware has been upgraded successfully.

10. Proceed to the next Cisco switch in the appliance. Upgrade the switches, one at a time, in this order:

- a. Leaf Cisco Nexus 9336C-FX2 Switches: ovcasw15r1, ovcasw16r1
- b. Spine Cisco Nexus 9336C-FX2 Switches: ovcasw22r1, ovcasw23r1
- c. Management Cisco Nexus 9348GC-FXP Switch: ovcasw21r1

3.4.5 Upgrading the NM2-36P Sun Datacenter InfiniBand Expansion Switch Firmware

The instructions in this section are specific for a component firmware upgrade as part of the Oracle Private Cloud Appliance.



Note

InfiniBand switches are part of systems with an InfiniBand-based network architecture.



Warning

For firmware upgrades to version 2.2.8 or newer, an intermediate upgrade to unsigned version 2.2.7-2 is required. Version 2.2.7-2 can then be upgraded to the intended newer version. For additional information, refer to [NM2-36P Sun Datacenter InfiniBand Expansion Switch Firmware Upgrade 2.2.9-3 Requires A Two-Phased Procedure](#) in the *Oracle Private Cloud Appliance Release Notes*.



Note

Detailed information about firmware upgrades can be found in the *Sun Datacenter InfiniBand Switch 36 Product Notes for Firmware Version 2.2* (document ID: E76431). Refer to the section “Upgrading the Switch Firmware”.



Caution

It is recommended that you back up the current configuration of the NM2-36P Sun Datacenter InfiniBand Expansion Switches prior to performing a firmware upgrade.

Backup and restore instructions are provided in the maintenance and configuration management sections of the *Oracle ILOM Administration Guide* that corresponds with the current ILOM version used in the switch. For example:

- Oracle ILOM 3.0: https://docs.oracle.com/cd/E36265_01/html/E36266/ceiidgfj.html#scrolltoc
- Oracle ILOM 3.2: https://docs.oracle.com/cd/E37444_01/html/E37446/z400371a1482122.html#scrolltoc

Upgrading the InfiniBand Switch Firmware

1. Log on to the master management node using SSH and an account with superuser privileges.
2. Unzip the firmware package included in the Oracle Private Cloud Appliance software image.

```
[root@ovcamn05r1 ~]# mkdir /nfs/shared_storage/yum/nm2
[root@ovcamn05r1 ~]# cd /nfs/shared_storage/yum/nm2
[root@ovcamn05r1 nm2]# unzip /nfs/shared_storage/mgmt_image/firmware/IB_gateway/NM2-36P/p22173626_227_Gener
Archive:  /nfs/shared_storage/mgmt_image/firmware/IB_gateway/NM2-36P/p22173626_227_Generic.zip
inflating: license.txt
```



```

inflating: readme_SUN_DCS_36p_2.1.8-1.txt
creating: SUN_DCS_36p_2.1.8-1/
inflating: SUN_DCS_36p_2.1.8-1/pkey_filter.pl
creating: SUN_DCS_36p_2.1.8-1/SUN_DCS_36p/
inflating: SUN_DCS_36p_2.1.8-1/SUN_DCS_36p/SUN-ILOM-CONTROL-MIB.mib
inflating: SUN_DCS_36p_2.1.8-1/SUN_DCS_36p/SUN-FABRIC-MIB.mib
inflating: SUN_DCS_36p_2.1.8-1/SUN_DCS_36p/sundcs_36p_repository_2.1.8_1.pkg
inflating: SUN_DCS_36p_2.1.8-1/SUN_DCS_36p/SUN-HW-TRAP-MIB.mib
inflating: SUN_DCS_36p_2.1.8-1/SUN_DCS_36p/SUN-DCS-IB-MIB.txt
inflating: SUN_DCS_36p_2.1.8-1/SUN_DCS_36p/SUN-PLATFORM-MIB.mib
inflating: SUN_DCS_36p_2.1.8-1/SUN_DCS_36p/ENTITY-MIB.mib
inflating: SUN_DCS_36p_2.1.8-1/SUN_DCS_36p_2.1.8-1_metadata.xml
inflating: SUN_DCS_36p_2.1.8-1/README_pkey_filter
inflating: SUN_DCS_36p_2.1.8-1_THIRDPARTYLICENSE.pdf

```

3. Log on to one of the InfiniBand switches as root.

```

root@ovcamn05r1 ~]# ssh root@192.168.4.202
root@192.168.4.202's password:
You are now logged in to the root shell.
It is recommended to use ILOM shell instead of root shell.
All usage should be restricted to documented commands and documented config files.
To view the list of documented commands, use "help" at linux prompt.
[root@ilom-ovcasw19r1 ~]#

```

4. Check the master configuration and the state of the SubnetManager.

```

[root@ilom-ovcasw19r1 ~]# getmaster
Local SM not enabled
Last change in Master SubnetManager status detected at: Thu Mar 22 14:29:18 UTC 2018
Master SubnetManager on sm lid 6 sm guid 0x13970201001ba4 : MT25408 ConnectX Mellanox Technologies
Master SubnetManager Activity Count: 348521 Priority: 0

```



Warning

The command output must read Local SM not enabled. If this is not the case, abort this procedure and contact Oracle Support.

5. List the details of the current firmware version.

```

[root@ilom-ovcasw19r1 ~]# version
SUN DCS 36p version: 1.3.3-2
Build time: Feb 19 2013 13:29:01
SP board info:
Manufacturing Date: 2012.06.23
Serial Number: "NCDBJ1073"
Hardware Revision: 0x0007
Firmware Revision: 0x0000
BIOS version: SUN0R100
BIOS date: 06/22/2010

```

6. Connect to the ILOM and start the firmware upgrade procedure. Press "Y" when prompted to load the file.

```

[root@ilom-ovcasw19r1 ~]# spsh
Oracle(R) Integrated Lights Out Manager
Version ILOM 3.0 r47111
Copyright (c) 2012, Oracle and/or its affiliates. All rights reserved.

-> load -source http://192.168.4.1/shares/export/Yum/nm2/ \SUN_DCS_36p_2.1.8-1/SUN_DCS_36p/sundcs_36p_r

Downloading firmware image. This will take a few minutes.
Are you sure you want to load the specified file (y/n)

Setting up environment for firmware upgrade. This will take few minutes.

```

```
Starting SUN DCS 36p FW update
=====
Performing operation: I4 A
=====
I4 fw upgrade from 7.3.0(INI:4) to 7.4.1010(INI:4):
Upgrade started...
Upgrade completed.
INFO: I4 fw upgrade from 7.3.0(INI:4) to 7.4.1010(INI:4) succeeded
=====
Summary of Firmware update
=====
I4 status           : FW UPDATE - SUCCESS
I4 update succeeded on : A
I4 already up-to-date on : none
I4 update failed on   : none
=====
Performing operation: SUN DCS 36p firmware update
=====
SUN DCS 36p upgrade from 1.3.3-2 to 2.1.8-1:
Upgrade started...
Upgrade completed.
INFO: SUN DCS 36p upgrade from 1.3.3-2 to 2.1.8-1 succeeded
Firmware update is complete.
ILOM will be restarted and will take 2 minutes to come up.
You will need to reconnect to Integrated Lights Out Manager.
```

7. Reconnect to the InfiniBand switch to verify that the new firmware is running and to confirm that the SubnetManager remains disabled.

```
root@ovcamn05r1 ~]# ssh root@192.168.4.202
root@192.168.4.202's password:
[root@ilom-ovcasw19r1 ~]# version
SUN DCS 36p version: 2.1.8-1
Build time: Sep 18 2015 10:26:47
SP board info:
Manufacturing Date: 2013.06.15
Serial Number: "NCDBJ1073"
Hardware Revision: 0x0007
Firmware Revision: 0x0000
BIOS version: SUN0R100
BIOS date: 06/22/2010

[root@ilom-ovcasw19r1 ~]# getmaster
Local SM not enabled
```



Warning

The command output must read Local SM not enabled. If this is not the case, abort this procedure and contact Oracle Support.

8. When the first InfiniBand switch has completed the upgrade successfully and has come back online, connect to the other InfiniBand switch, with IP address 192.168.4.203, and execute the same procedure.
9. When both InfiniBand switches have been upgraded, remove the shared storage directory you created to make the unzipped package available.

```
root@ovcamn05r1 ~]# cd /nfs/shared_storage/yum/
root@ovcamn05r1 yum]# ls -al
total 323
drwxr-xr-x  8 root root   8 Mar 26 07:57 .
drwxrwxrwx 31 root root  31 Mar 13 13:38 ..
drwxr-xr-x  2 root root   5 Mar 13 12:04 backup_COMPUTE
drwxr-xr-x  2 root root   5 Mar 13 13:19 current_COMPUTE
drwxr-xr-x  3 root root   6 Mar 26 07:58 nm2
```

```
drwxr-xr-x 3 root root 587 Mar 13 12:04 OVM_3.4.4_1735_server
drwxr-xr-x 3 root root 18 Mar 13 12:03 OVM_3.4.4_1735_transition
drwxr-xr-x 4 root root 9 Mar 13 12:03 OVM_3.4.4_1735_update
root@ovcamn05r1 yum]# rm -rf nm2/
```

3.4.6 Upgrading the Oracle Fabric Interconnect F1-15 Firmware

The instructions in this section are specific for a component firmware upgrade as part of the Oracle Private Cloud Appliance.



Note

Fabric Interconnects are part of systems with an InfiniBand-based network architecture.



Note

Detailed information about firmware upgrades can be found in the *XgOS User's Guide* (document ID: E53170). Refer to the section "System Image Upgrades".



Caution

It is recommended that you back up the current configuration of the Fabric Interconnects prior to performing a firmware upgrade. Store the backup configuration on another server and add a time stamp to the file name for future reference.

For detailed information, refer to the section "Saving and Restoring Your Configuration" in the *XgOS User's Guide* (document ID: E53170).

Upgrading the Fabric Interconnect Firmware

1. Log on to the master management node using SSH and an account with superuser privileges.
2. Copy the firmware package from the Oracle Private Cloud Appliance software image to the Yum repository share.

```
root@ovcamn05r1 ~]# cp /nfs/shared_storage/mgmt_image/firmware/IB_gateway/ \
OFI/xsigo-4.0.12-XGOS.xpf /nfs/shared_storage/yum/
```

3. Log on to one of the Fabric Interconnects as admin.

```
root@ovcamn05r1 ~]# ssh admin@192.168.4.205
Password:
Last login: Thu Oct 15 10:57:23 2015 from 192.168.4.4
        Welcome to XgOS
Copyright (c) 2007-2012 Xsigo Systems, Inc. All rights reserved.
Enter "help" for information on available commands.
Enter the command "show system copyright" for licensing information
admin@ovcasw22r1[xsigo]
```

4. List the details of the current firmware version.

```
admin@ovcasw22r1[xsigo] show system version
Build 3.9.4-XGOS - (buildsys) Thu Mar 19 03:25:26 UTC 2015
admin@ovcasw22r1[xsigo]
```

5. Check the master configuration and the state of the SubnetManager. Optionally run the additional diagnostics command for more detailed information.

```
admin@ovcasw22r1[xsigo] show diagnostics sm-info
- SM is running on          ovcasw22r1
```

```

- SM Lid                39
- SM Guid               0x139702010017b4
- SM key                0x0
- SM priority           0
- SM State              MASTER

admin@ovcasw22r1[xsigo] show diagnostics opensm-param

OpenSM $ Current log level is 0x83
OpenSM $ Current sm-priority is 0
OpenSM $
  OpenSM Version       : OpenSM 3.3.5
  SM State              : Master
  SM Priority           : 0
  SA State              : Ready
  Routing Engine        : minhop
  Loaded event plugins  : <none>

  PerfMgr state/sweep state : Disabled/Sleeping

  MAD stats
  -----
  QP0 MADs outstanding      : 0
  QP0 MADs outstanding (on wire) : 0
  QP0 MADs rcvd             : 6323844
  QP0 MADs sent             : 6323676
  QP0 unicasts sent        : 2809116
  QP0 unknown MADs rcvd    : 0
  SA MADs outstanding       : 0
  SA MADs rcvd             : 120021107
  SA MADs sent             : 120024422
  SA unknown MADs rcvd     : 0
  SA MADs ignored          : 0

  Subnet flags
  -----
  Sweeping enabled         : 1
  Sweep interval (seconds) : 10
  Ignore existing lfts     : 0
  Subnet Init errors       : 0
  In sweep hop 0           : 0
  First time master sweep  : 0
  Coming out of standby    : 0

  Known SMs
  -----
  Port GUID      SM State  Priority
  -----
  0x139702010017b4 Master    0      SELF
  0x139702010017c0 Standby   0

```

6. Start the system upgrade procedure.

```

admin@ovcasw22r1[xsigo] system upgrade
http://192.168.4.1/shares/export/Yum/xsigo-4.0.12-XGOS.xpf forcebaseos
Copying...
#####
[100%]
You have begun to upgrade the system software.
Please be aware that this will cause an I/O service interruption
and the system may be rebooted.
The following software will be installed:
1. XgOS Operating System software including SCP Base OS
2. XgOS Front-panel software

```

```

3. XgOS Common Chassis Management software on IOC
4. XgOS VNIC Manager and Agent software
5. XgOS VN10G and VN10x1G Manager and Agent software
6. XgOS VHBA and VHBA-2 Manager and Agent software
7. XgOS VN10G and VN10x1G Manager and Agent software with Eth/IB Interfaces
8. XgOS VN4x10G and VN2x10G Manager and Agent software with Eth/IB Interfaces
9. XgOS VHBA-3 Manager and Agent software
10. XgOS VHBA 2x 8G FC Manager and Agent software
Are you sure you want to update the software (y/n)? y
Running verify scripts...
Running preunpack scripts...
Installing...
#####
[100%]
Verifying...
#####
[100%]
Running preinstall scripts...
Installing Base OS - please wait...
LABEL=/dev/uba /mnt/usb vfat rw 0 0
Rootfs installation successful
The installer has determined that a reboot of the Base OS is required (HCA driver changed)
The installer has determined that a cold restart of the Director is necessary
Installing package...
Running postinstall scripts...
Installation successful. Please stand by for CLI restart.
admin@iowa[xsigo] Rebooting OS. Please log in again in a couple of minutes...

*****
Xsigo system is being shut down now
*****
Connection to 192.168.4.204 closed.

```

After reboot, it takes approximately 10 minutes before you can log back in. The upgrade resets the admin user's password to the default "admin". It may take several attempts, but login with the default password eventually succeeds.

7. Reconnect to the Fabric Interconnect to change the admin and root passwords back to the setting from before the firmware upgrade.



Note

When you log back in after the firmware upgrade, you may encounter messages similar to this example:

```

Message from syslogd@ovcasw22r1 at Fri Jun 22 09:49:33 2018 ...
ovcasw22r1 systemcontroller[2713]: [EMERG] ims::IMSService [ims::failedloginattempt]
user admin has tried to log on for 5 times in a row without success !!

```

These messages indicate failed login attempts from other Oracle Private Cloud Appliance components. They disappear after you set the passwords back to their original values.

Modify the passwords for users root and admin as follows:

```

admin@ovcasw22r1[xsigo] set system root-password
Administrator's password: admin
New password: myOriginalRootPassword
New password again: myOriginalRootPassword

admin@ovcasw22r1[xsigo] set user admin -password
New password: myOriginalAdminPassword
New password again: myOriginalAdminPassword

```

8. Reconnect to the Fabric Interconnect to verify that the new firmware is running and to confirm that all vNICs and vHBAs are in up/up state.

```

root@ovcamn05r1 ~]# ssh admin@192.168.4.205
admin@ovcasw22r1[xsigo] show system version
Build 4.0.12-XGOS - (buildsys) Fri Jun 22 04:42:35 UTC 2018

admin@ovcasw22r1[xsigo] show diagnostics sm-info
- SM is running on          ovcasw22r1
- SM Lid                    39
- SM Guid                   0x139702010017b4
- SM key                    0x0
- SM priority               0
- SM State                  MASTER

admin@ovcasw22r1[xsigo] show vnic

name                state      mac-addr          ipaddr           if              if-state
-----
eth4.ovcacn08r1    up/up     00:13:97:59:90:11 0.0.0.0/32      mgmt_pvi(64539) up
eth4.ovcacn09r1    up/up     00:13:97:59:90:0D 0.0.0.0/32      mgmt_pvi(64539) up
eth4.ovcacn10r1    up/up     00:13:97:59:90:09 0.0.0.0/32      mgmt_pvi(64539) up
eth4.ovcacn11r1    up/up     00:13:97:59:90:1D 0.0.0.0/32      mgmt_pvi(64539) up
eth4.ovcacn12r1    up/up     00:13:97:59:90:19 0.0.0.0/32      mgmt_pvi(64539) up
[...]
eth7.ovcacn29r1    up/up     00:13:97:59:90:28 0.0.0.0/32      5/1             up
eth7.ovcamn05r1    up/up     00:13:97:59:90:04 0.0.0.0/32      4/1             up
eth7.ovcamn06r1    up/up     00:13:97:59:90:08 0.0.0.0/32      5/1             up
40 records displayed

admin@ovcasw22r1[xsigo] show vhba

name                state      fabric-state      if      if-state      wwnn
-----
vhba03.ovcacn07r1    up/up     down(Down)        12/1    down          50:01:39:71:00:58:B1:0A
vhba03.ovcacn08r1    up/up     down(Down)        3/1     down          50:01:39:71:00:58:B1:08
vhba03.ovcacn09r1    up/up     down(Down)        12/1    down          50:01:39:71:00:58:B1:06
vhba03.ovcacn10r1    up/up     down(Down)        3/1     down          50:01:39:71:00:58:B1:04
[...]
vhba04.ovcacn29r1    up/up     down(Down)        12/2    down          50:01:39:71:00:58:B1:13
vhba04.ovcamn05r1    up/up     down(Down)        3/2     down          50:01:39:71:00:58:B1:01
vhba04.ovcamn06r1    up/up     down(Down)        12/2    down          50:01:39:71:00:58:B1:03
20 records displayed

```

9. When the first Fabric Interconnect has completed the upgrade successfully and has come back online, connect to the other Fabric Interconnect, with IP address 192.168.4.204, and execute the same procedure.

Chapter 4 The Oracle Private Cloud Appliance Command Line Interface (CLI)

Table of Contents

4.1 CLI Usage	94
4.1.1 Interactive Mode	94
4.1.2 Single-command Mode	96
4.1.3 Controlling CLI Output	96
4.1.4 Internal CLI Help	98
4.2 CLI Commands	99
4.2.1 add compute-node	99
4.2.2 add network	100
4.2.3 add network-to-tenant-group	101
4.2.4 backup	102
4.2.5 configure vhbases	103
4.2.6 create lock	104
4.2.7 create network	105
4.2.8 create tenant-group	106
4.2.9 create uplink-port-group	107
4.2.10 delete config-error	108
4.2.11 delete lock	109
4.2.12 delete network	110
4.2.13 delete task	111
4.2.14 delete tenant-group	112
4.2.15 delete uplink-port-group	113
4.2.16 deprovision compute-node	114
4.2.17 diagnose	116
4.2.18 get log	119
4.2.19 list	120
4.2.20 remove compute-node	126
4.2.21 remove network	127
4.2.22 remove network-from-tenant-group	128
4.2.23 reprovision	129
4.2.24 rerun	130
4.2.25 set system-property	131
4.2.26 show	133
4.2.27 start	138
4.2.28 stop	139
4.2.29 update appliance	140
4.2.30 update password	140
4.2.31 update compute-node	142

All Oracle Private Cloud Appliance command line utilities are consolidated into a single command line interface that is accessible via the management node shell by running the `pca-admin` command located at `/usr/sbin/pca-admin`. This command is in the system path for the root user, so you should be able to run the command from anywhere that you are located on a management node. The CLI provides access to all of the tools available in the Oracle Private Cloud Appliance Dashboard, as well as many that do not have a Dashboard equivalent. The design of the CLI makes it possible to script actions that may need to be performed more regularly, or to write integration scripts with existing monitoring and maintenance software not directly hosted on the appliance.

It is important to understand that the CLI, described here, is distinct from the Oracle VM Manager command line interface, which is described fully in the Oracle VM documentation available at http://docs.oracle.com/cd/E64076_01/E64086/html/index.html.

In general, it is preferable that CLI usage is restricted to the active management node. While it is possible to run the CLI from either management node, some commands are restricted to the active management node and return an error if you attempt to run them on the passive management node.

4.1 CLI Usage

The Oracle Private Cloud Appliance command line interface is triggered by running the `pca-admin` command. It can run either in interactive mode (see [Section 4.1.1, “Interactive Mode”](#)) or in single-command mode (see [Section 4.1.2, “Single-command Mode”](#)) depending on whether you provide the syntax to run a particular CLI command when you invoke the command line interpreter.

The syntax when using the CLI is as follows:

```
PCA> Command Command_Target <Arguments> Options
```

where:

- *Command* is the command type that should be initiated. For example `list`;
- *Command_Target* is the Oracle Private Cloud Appliance component or process that should be affected by the command. For example `management-node`, `compute-node`, `task` etc;
- *<Arguments>* consist of positioning arguments related to the command target. For instance, when performing a reprovisioning action against a compute node, you should provide the specific compute node that should be affected as an argument for this command. For example: `reprovision compute-node ovcacn11r1`;
- *Options* consist of options that may be provided as additional parameters to the command to affect its behavior. For instance, the `list` command provides various sorting and filtering options that can be appended to the command syntax to control how output is returned. For example: `list compute-node --filter-column Provisioning_State --filter dead`. See [Section 4.1.3, “Controlling CLI Output”](#) for more information on many of these options.

The CLI includes its own internal help that can assist you with understanding the commands, command targets, arguments and options available. See [Section 4.1.4, “Internal CLI Help”](#) for more information on how to use this help system. When used in interactive mode, the CLI also provides tab completion to assist you with the correct construction of a command. See [Section 4.1.1.1, “Tab Completion”](#) for more information on this.

4.1.1 Interactive Mode

The Oracle Private Cloud Appliance command line interface (CLI) provides an interactive shell that can be used for user-friendly command line interactions. This shell provides a closed environment where users can enter commands specific to the management of the Oracle Private Cloud Appliance. By using the CLI in interactive mode, the user can avail of features like tab completion to easily complete commands correctly. By default, running the `pca-admin` command without providing any additional parameters causes the CLI interpreter to run in interactive mode.

It is possible to identify that you are in a CLI shell running in interactive mode as the shell prompt is indicated by `PCA>`.

Example 4.1 An example of interactive mode usage of the CLI

```
# pca-admin
```



```

Welcome to PCA! Release: 2.4.1
PCA> list management-node

Management_Node  IP_Address  Provisioning_Status  ILOM_MAC  Provisioning_State  Master
-----
ovcamn05r1      192.168.4.3  RUNNING             00:10:e0:e9:1f:c9  running            Yes
ovcamn06r1      192.168.4.4  RUNNING             00:10:e0:e7:26:ad  running            None
-----
2 rows displayed

Status: Success
PCA> exit
#

```

To exit from the CLI when it is in interactive mode, you can use either the `q`, `quit`, or `exit` command, or alternatively use the Ctrl+D key combination.

4.1.1.1 Tab Completion

The CLI supports tab-completion when in interactive mode. This means that pressing the tab key while entering a command can either complete the command on your behalf, or can indicate options and possible values that can be entered to complete a command. Usually you must press the tab key at least twice to effect tab-completion.

Tab-completion is configured to work at all levels within the CLI and is context sensitive. This means that you can press the tab key to complete or prompt for commands, command targets, options, and for certain option values. For instance, pressing the tab key twice at a blank prompt within the CLI automatically lists all possible commands, while pressing the tab key after typing the first letter or few letters of a command automatically completes the command for you. Once a command is specified, followed by a space, pressing the tab key indicates command targets. If you have specified a command target, pressing the tab key indicates other options available for the command sequence. If you press the tab key after specifying a command option that requires an option value, such as the `--filter-column` option, the CLI attempts to provide you with the values that can be used with that option.

Example 4.2 Examples showing tab-completion

```

PCA> <tab>
EOF      backup      create      deprovision  exit      help      q
remove   rerun       shell       start        update    add       configure
delete   diagnose    get         list         quit     reprovision  set
show     stop

PCA> list <tab>
compute-node      lock      mgmt-switch-port  network-port      task
update-task      uplink-port-group  config-error      management-node   network
network-switch   tenant-group      uplink-port

PCA> list com<tab>pute-node

```

The `<tab>` indicates where the user pressed the tab key while in an interactive CLI session. In the final example, the command target is automatically completed by the CLI.

4.1.1.2 Running Shell Commands

It is possible to run standard shell commands while you are in the CLI interpreter shell. These can be run by either preceding them with the `shell` command or by using the `!` operator as a shortcut to indicate that the command that follows is a standard shell command. For example:

```

PCA> shell date
Wed Jun  5 08:15:56 UTC 2019

```

```
PCA> !uptime > /tmp/uptime-today
PCA> !rm /tmp/uptime-today
```

4.1.2 Single-command Mode

The CLI supports 'single-command mode', which allows you to execute a single command from the shell via the CLI and to obtain the output before the CLI exits back to the shell. This is particularly useful when writing scripts that may interact with the CLI, particularly if used in conjunction with the CLI's JSON output mode described in [Section 4.1.3.1, "JSON Output"](#).

To run the CLI in single-command mode, simply include the full command syntax that you wish to execute as parameters to the `pca-admin` command.

An example of single command mode is provided below:

```
# pca-admin list compute-node
Compute_Node  IP_Address  Provisioning_Status  ILOM_MAC  Provisioning_State
-----
ovcacn12r1    192.168.4.8  RUNNING              00:10:e0:e5:e6:d3  running
ovcacn07r1    192.168.4.7  RUNNING              00:10:e0:e6:8d:0b  running
ovcacn13r1    192.168.4.11  RUNNING              00:10:e0:e6:f7:f7  running
ovcacn14r1    192.168.4.9   RUNNING              00:10:e0:e7:15:eb  running
ovcacn10r1    192.168.4.12  RUNNING              00:10:e0:e7:13:8d  running
ovcacn09r1    192.168.4.6   RUNNING              00:10:e0:e6:f8:6f  running
ovcacn11r1    192.168.4.10  RUNNING              00:10:e0:e6:f9:ef  running
-----
7 rows displayed
#
```

4.1.3 Controlling CLI Output

The CLI provides options to control how output is returned in responses to the various CLI commands that are available. These are provided as additional options as the final portion of the syntax for a CLI command. Many of these options can make it easier to identify particular items of interest through sorting and filtering, or can be particularly useful when scripting solutions as they help to provide output that is more easily parsed.

4.1.3.1 JSON Output

JSON format is a commonly used format to represent data objects in a way that is easy to machine-parse but is equally easy for a user to read. Although JSON was originally developed as a way to represent JavaScript objects, parsers are available for a wide number of programming languages, making it an ideal output format for the CLI if you are scripting a custom solution that may need to interface directly with the CLI.

The CLI returns its output for any command in JSON format if the `--json` option is specified when a command is run. Typically this option may be used when running the CLI in single-command mode. An example follows:

```
# pca-admin list compute-node --json
{
  "00:10:e0:e5:e6:ce": {
    "name": "ovcacn12r1",
    "ilom_state": "running",
    "ip": "192.168.4.8",
    "tenant_group_name": "Rack1_ServerPool",
    "state": "RUNNING",
    "networks": "default_external, default_internal",
```

```

    "ilom_mac": "00:10:e0:e5:e6:d3"
  },
  "00:10:e0:e6:8d:06": {
    "name": "ovcacn07r1",
    "ilom_state": "running",
    "ip": "192.168.4.7",
    "tenant_group_name": "Rack1_ServerPool",
    "state": "RUNNING",
    "networks": "default_external, default_internal",
    "ilom_mac": "00:10:e0:e6:8d:0b"
  },
  [...]
  "00:10:e0:e6:f9:ea": {
    "name": "ovcacn11r1",
    "ilom_state": "running",
    "ip": "192.168.4.10",
    "tenant_group_name": "",
    "state": "RUNNING",
    "networks": "default_external, default_internal",
    "ilom_mac": "00:10:e0:e6:f9:ef"
  }
}

```

In some cases the JSON output may contain more information than is displayed in the tabulated output that is usually shown in the CLI when the `--json` option is not used. Furthermore, the keys used in the JSON output may not map identically to the table column names that are presented in the tabulated output.

Sorting and filtering options are currently not supported in conjunction with JSON output, since these facilities can usually be implemented on the side of the parser.

4.1.3.2 Sorting

Typically, when using the `list` command, you may wish to sort information in a way that makes it easier to view items of particular interest. This is achieved using the `--sorted-by` and `--sorted-order` options in conjunction with the command. When using the `--sorted-by` option, you must specify the column name against which the sort should be applied. You can use the `--sorted-order` option to control the direction of the sort. This option should be followed either with `ASC` for an ascending sort, or `DES` for a descending sort. If this option is not specified, the default sort order is ascending.

For example, to sort a view of compute nodes based on the status of the provisioning for each compute node, you may do the following:

```

PCA> list compute-node --sorted-by Provisioning_State --sorted-order ASC

Compute_Node  IP_Address  Provisioning_Status  ILOM_MAC  Provisioning_State
-----
ovcacn08r1    192.168.4.9  RUNNING              00:10:e0:65:2f:b7  dead
ovcacn28r1    192.168.4.10  RUNNING              00:10:e0:62:31:81  initializing_stage_wait_for_hmp
ovcacn10r1    192.168.4.7   RUNNING              00:10:e0:65:2f:cf  initializing_stage_wait_for_hmp
ovcacn30r1    192.168.4.8   RUNNING              00:10:e0:40:cb:59  running
ovcacn07r1    192.168.4.11  RUNNING              00:10:e0:62:ca:09  running
ovcacn26r1    192.168.4.12  RUNNING              00:10:e0:65:30:f5  running
ovcacn29r1    192.168.4.5   RUNNING              00:10:e0:31:49:1d  running
ovcacn09r1    192.168.4.6   RUNNING              00:10:e0:65:2f:3f  running
-----
8 rows displayed

Status: Success

```

Note that you can use tab-completion with the `--sorted-by` option to easily obtain the options for different column names. See [Section 4.1.1.1, “Tab Completion”](#) for more information.

4.1.3.3 Filtering

Some tables may contain a large number of rows that you are not interested in, to limit the output to items of particular interest you can use the filtering capabilities that are built into the CLI. Filtering is achieved using a combination of the `--filter-column` and `--filter` options. The `--filter-column` option must be followed by specifying the column name, while the `--filter` option is followed with the specific text that should be matched to form the filter. The text that should be specified for a `--filter` may contain wildcard characters. If that is not the case, it must be an exact match. Filtering does not currently support regular expressions or partial matches.

For example, to view only the compute nodes that have a Provisioning state equivalent to 'dead', you could use the following filter:

```
PCA> list compute-node --filter-column Provisioning_State --filter dead

Compute_Node  IP_Address  Provisioning_Status  ILOM_MAC  Provisioning_State
-----
ovcacn09r1    192.168.4.10  DEAD                00:10:e0:0f:55:cb  dead
ovcacn11r1    192.168.4.9   DEAD                00:10:e0:0f:57:93  dead
ovcacn14r1    192.168.4.7   DEAD                00:10:e0:46:9e:45  dead
ovcacn36r1    192.168.4.11  DEAD                00:10:e0:0f:5a:9f  dead
-----
4 rows displayed

Status: Success
```

Note that you can use tab-completion with the `--filter-column` option to easily obtain the options for different column names. See [Section 4.1.1.1, "Tab Completion"](#) for more information.

4.1.4 Internal CLI Help

The CLI includes its own internal help system. This is triggered by issuing the `help` command:

```
PCA> help

Documented commands (type help <topic>):
=====
add          create      diagnose   list        rerun      start
backup      delete     get         remove      set         stop
configure   deprovision help        reprovision show        update

Undocumented commands:
=====
EOF exit q quit shell
```

The help system displays all of the available commands that are supported by the CLI. These are organized into 'Documented commands' and 'Undocumented commands'. Undocumented commands are usually commands that are not specific to the management of the Oracle Private Cloud Appliance, but are mostly discussed within this documentation. Note that more detailed help can be obtained for any documented command by appending the name of the command to the `help` query. For example, to obtain the help documentation specific to the `list` command, you can do the following:

```
PCA> help list
Usage: pca-admin list <Command Target> [OPTS]

Command Targets:
  compute-node      List computer node.
  config-error      List configuration errors.
  lock              List lock.
  management-node   List management node.
  mgmt-switch-port  List management switch port.
```

```

network          List active networks.
network-port     List network port.
network-switch   List network switch.
task             List task.
tenant-group     List tenant-group.
update-task      List update task.
uplink-port      List uplink port.
uplink-port-group List uplink port group.

Options:
--json           Display the output in json format.
--less           Display output in the less pagination mode.
--more          Display output in the more pagination mode.
--tee=OUTPUTFILENAME Export output to a file.
--sorted-by=SORTEDBY  Sorting the table by a column.
--sorted-order=SORTEDORDER
                  Sorting order.
--filter-column=FILTERCOLUMN
                  Table column that needs to be filtered.
--filter=FILTER      filter criterion

```

You can drill down further into the help system for most commands by also appending the command target onto your [help](#) query:

```

PCA> help reprovision compute-node
Usage:
reprovision compute-node <compute node name> [options]

Example:
reprovision compute-node ovcacn1lr1

Description:
Reprovision a compute node.

```

Finally, if you submit a help query for something that doesn't exist, the help system generates an error and automatically attempts to prompt you with alternative candidates:

```

PCA> list ta
Status: Failure
Error Message: Error (MISSING_TARGET_000): Missing command target for command: list.
Command targets can be: ['update-task', 'uplink-port-group', 'config-error', 'network',
'lock', 'network-port', 'tenant-group', 'network-switch', 'task', 'compute-node',
'uplink-port', 'mgmt-switch-port', 'management-node'].

```

4.2 CLI Commands

This section describes all of the documented commands available via the CLI.

Note that there are slight differences in the CLI commands available on Ethernet-based systems and InfiniBand-based systems. If you issue a command that is not available on your specific architecture, the command fails.

4.2.1 add compute-node

Adds a compute node to an existing tenant group. To create a new tenant group, see [Section 4.2.8, “create tenant-group”](#).

Syntax

```

add compute-node node tenant-group-name [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]

```

where `tenant-group-name` is the name of the tenant group you wish to add one or more compute nodes to, and `node` is the name of the compute node that should be added to the selected tenant group.

Description

Use the `add compute-node` command to add the required compute nodes to a tenant group you created. If a compute node is currently part of another tenant group, it is first removed from that tenant group. If custom networks are already associated with the tenant group, the newly added server is connected to those networks as well. Use the command `add network-to-tenant-group` to associate a custom network with a tenant group.

Options

The following table shows the available options for this command.

Option	Description
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.3 Adding a Compute Node to a Tenant Group

```
PCA> add compute-node ovcacn09r1 myTenantGroup
Status: Success
```

4.2.2 add network

Connects a server node to an existing network. To create a new custom network, see [Section 4.2.7, “create network”](#).

Syntax

```
add network network-name node [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

where `network-name` is the name of the network you wish to connect one or more servers to, and `node` is the name of the server node that should be connected to the selected network.

Description

Use the `add network` command to connect the required server nodes to a custom network you created. When you set up custom networks between your servers, you create the network first, and then add the required servers to the network. Use the `create network` command to configure additional custom networks.

Options

The following table shows the available options for this command.

Option	Description
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.4 Connecting a Compute Node to a Custom Network

```
PCA> add network MyNetwork ovcacn09r1
Status: Success
```

4.2.3 add network-to-tenant-group

Associates a custom network with an existing tenant group. To create a new tenant group, see [Section 4.2.8, “create tenant-group”](#). To create a new custom network, see [Section 4.2.7, “create network”](#).

Syntax

```
add network-to-tenant-group network-name tenant-group-name [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

where `network-name` is the name of an existing custom network, and `tenant-group-name` is the name of the tenant group you wish to associate the custom network with.

Description

Use the `add network-to-tenant-group` command to connect all member servers of a tenant group to a custom network. The custom network connection is configured when a server joins the tenant group, and unconfigured when a server is removed from the tenant group.



Note

This command involves verification steps that are performed in the background. Consequently, even though output is returned and you regain control of the CLI, certain operations continue to run for some time.

Options

The following table shows the available options for this command.

Option	Description
<code>--json</code>	Return the output of the command in JSON format

Option	Description
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.5 Associating a Custom Network with a Tenant Group

```
PCA> add network-to-tenant-group myPublicNetwork myTenantGroup

Validating servers in the tenant group... This may take some time.

The job for sync all nodes in tenant group with the new network myPublicNetwork has been submitted.
Please look into "/var/log/ovca.log" and "/var/log/ovca-sync.log" to monitor the progress.

Status: Success
```

4.2.4 backup

Triggers a manual backup of the Oracle Private Cloud Appliance.



Note

The backup command can only be executed from the active management node; not from the standby management node.

Syntax

```
backup [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

Description

Use the `backup` command to initiate a backup task outside of the usual cron schedule. The backup task performs a full backup of the Oracle Private Cloud Appliance as described in [Section 1.6, “Oracle Private Cloud Appliance Backup”](#). The CLI command does not monitor the progress of the backup task itself, and exits immediately after triggering the task, returning the task ID and name, its initial status, its progress and start time. This command must only ever be run on the active management node.

You can use the `show task` command to view the status of the task after you have initiated the backup. See [Example 4.44, “Show Task”](#) for more information.

Options

There are no further options for this command.

Examples

Example 4.6 Running a backup task

```
PCA> backup
```


The backup job has been submitted. Use "show task <task id>" to monitor the progress.

```
Task_ID          Status  Progress  Start_Time          Task_Name
-----
3769a13df448a2  RUNNING None      06-05-2019 09:21:36  backup
-----
1 row displayed
Status: Success
```

4.2.5 configure vhas

Configures vHBAs on compute nodes. This command is used only on systems with InfiniBand-based network architecture.

Syntax

```
configure vhas { ALL | node } [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

where `node` is the compute node name for the compute node for which the vHBAs should be configured, and `ALL` refers to all compute nodes provisioned in your environment.

Description

This command creates the default virtual host bus adapters (vHBAs) for fibre channel connectivity, if they do not exist. Each of the four default vHBAs corresponds with a bond on the physical server. Each vHBA connection between a server node and Fabric Interconnect has a unique mapping. Use the `configure vhas` command to configure the virtual host bus adapters (vHBA) on all compute nodes or a specific subset of them.

Options

The following table shows the available options for this command.

Option	Description
<code>ALL node</code>	Configure vHBAs for all compute nodes or for one or more specific compute nodes.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.7 Configuring the vHBAs for Specific Compute Nodes

```
PCA> configure vhas ovcacn11r1 ovcacn14r1
Compute_Node          Status
-----
```

```

ovcacn14r1      Succeeded
ovcacn11r1      Succeeded
-----
2 rows displayed
Status: Success

```

4.2.6 create lock

Imposes a lock on certain appliance functionality.



Caution

Never use locks without consultation or specific instructions from Oracle Support.

Syntax

```

create lock { all_provisioning | cn_upgrade | database | install | manufacturing
| mn_upgrade | provisioning | service } [ --json ] [ --less ] [ --more ] [ --
tee=OUTPUTFILENAME ]

```

Description

Use the `create lock` command to temporarily disable certain appliance-level functions. The lock types are described in the Options.

Options

The following table shows the available options for this command.

Option	Description
<code>all_provisioning</code>	Suspend all management node updates and compute node provisioning. Running tasks are completed and stop before the next stage in the process. A daemon checks for locks every few seconds. Once the lock has been removed, the update or provisioning processes continue from where they were halted.
<code>cn_upgrade</code>	Prevent all compute node upgrade operations.
<code>database</code>	Impose a lock on the databases during the management node update process. The lock is released after the update.
<code>install</code>	Placeholder lock type. Currently not used.
<code>manufacturing</code>	For usage in manufacturing. This lock type prevents the first boot process from initiating between reboots in the factory. As long as this lock is active, the <code>ovca</code> service does not start.
<code>mn_upgrade</code>	Prevent all management node upgrade operations.
<code>provisioning</code>	Prevent compute node provisioning. If a compute node provisioning process is running, it stops at the next stage. A daemon checks for locks every few seconds. Once the lock has been removed, all nodes advance to the next stage in the provisioning process.

Option	Description
<code>service</code>	Placeholder lock type. Behavior is identical to manufacturing lock.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.8 Imposing a Provisioning Lock

```
PCA> create lock provisioning
Status: Success
```

4.2.7 create network

Creates a new custom network, private or public, at the appliance level. See [Section 2.6, “Network Customization”](#) for detailed information.

Syntax

```
create network network-name { rack_internal_network | external_network port-group
| host_network port-group prefix netmask [route-destination gateway] } [ --json ] [ --
less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

where `network-name` is the name of the custom network you wish to create.

If the network type is `external_network`, then the spine switch ports used for public connectivity must also be specified as `port-group`. For this purpose, you must first create an uplink port group. See [Section 4.2.9, “create uplink-port-group”](#) for more information.

If the network type is `host_network`, then additional arguments are expected. The subnet arguments are mandatory; the routing arguments are optional.

- `prefix`: defines the fixed part of the host network subnet, depending on the netmask
- `netmask`: determines which part of the subnet is fixed and which part is variable
- `[route-destination]`: the external network location reachable from within the host network, which can be specified as a single valid IPv4 address or a subnet in CIDR notation.
- `[gateway]`: the IP address of the gateway for the static route, which must be inside the host network subnet

The IP addresses of the hosts or physical servers are based on the prefix and netmask of the host network. The final octet is the same as the corresponding internal management IP address. The routing information from the `create network` command is used to configure a static route on each compute node that joins the host network.

Options

The following table shows the available options for this command.

Option	Description
<code>{ rack_internal_network external_network host_network }</code>	The type of custom network to create. The options are: <ul style="list-style-type: none"> • a network internal to the rack • a network with external connectivity • a network with external connectivity, accessible for physical hosts
<code>external_network port-group</code>	To create a custom network with external connectivity, you must specify the ports on the spine switch as well. The ports must belong to an uplink port group, and you provide the port group name as an argument in this command.
<code>host_network port-group prefix netmask [route-destination gateway]</code>	To create a custom host network, you must specify the ports on the spine switch as with an external network. The ports must belong to an uplink port group, and you provide the port group name as an argument in this command. In addition, the host network requires arguments for its subnet. The routing arguments are optional. All four arguments are explained in the Syntax section above.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.9 Creating an Internal Custom Network

```
PCA> create network MyPrivateNetwork rack_internal_network
Status: Success
```

Example 4.10 Creating a Custom Network with External Connectivity

```
PCA> create network MyPublicNetwork external_network myUplinkPortGroup
Status: Success
```

4.2.8 create tenant-group

Creates a new tenant group. With the tenant group, which exists at the appliance level, a corresponding Oracle VM server pool is created. See [Section 2.7, “Tenant Groups”](#) for detailed information.

Syntax

```
create tenant-group tenant-group-name [ --json ][ --less ][ --more ][ --tee=OUTPUTFILENAME ]
```

where `tenant-group-name` is the name of the tenant group – and server pool – you wish to add to the environment.

Description

Use the `create tenant-group` command to set up a new placeholder for a separate group of compute nodes. The purpose of the tenant group is to group a number of compute nodes in a separate server pool. When the tenant group exists, add the required compute nodes using the `add compute-node` command. If you want to connect all the members of a server pool to a custom network, use the command `add network-to-tenant-group`.

Options

The following table shows the available options for this command.

Option	Description
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.11 Creating a Tenant Group

```
PCA> create tenant-group myTenantGroup
Status: Success
```

4.2.9 create uplink-port-group

Creates a new uplink port group. Uplink port groups define which spine switch ports are used together and in which breakout mode they operate. For detailed information, refer to [Appliance Uplink Configuration](#) in the *Oracle Private Cloud Appliance Installation Guide*. This command is used only on systems with Ethernet-based network architecture.

Syntax

```
create uplink-port-group port-group-name ports { 10g-4x | 25g-4x | 40g | 100g } [ --json ] [ --less ][ --more ][ --tee=OUTPUTFILENAME ]
```

where `port-group-name` is the name of the uplink port group, which must be unique. An uplink port group consists of a list of `ports` operating in one of the available breakout modes.

Description

Use the `create uplink-port-group` command to configure the ports reserved on the spine switches for external connectivity. Port 5 is configured and reserved for the default external network; ports 1-4 can be used for custom external networks. The ports can be used at their full 100Gbit bandwidth, at 40Gbit, or split with a breakout cable into four equal breakout ports: 4x 10Gbit or 4x 25Gbit. The port speed is reflected in the breakout mode of the uplink port group.

Options

The following table shows the available options for this command.

Option	Description
<code>ports</code>	To create an uplink port group, you must specify which ports on the spine switches belong to the port group. Ports must always be specified in adjacent pairs. They are identified by their port number and optionally, separated by a colon, also their breakout port ID. Put the port identifiers between quotes as a space-separated list, for example: <code>'1 2'</code> or <code>'3:1 3:2'</code> .
<code>{ 10g-4x 25g-4x 40g 100g }</code>	Set the breakout mode of the uplink port group. When a 4-way breakout cable is used, all four ports must be set to either 10Gbit or 25Gbit. When no breakout cable is used, the port speed for the uplink port group should be either 100Gbit or 40Gbit, depending on connectivity requirements.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.12 Creating an Uplink Port Group

```
PCA> create uplink-port-group myUplinkPortGroup '3:1 3:2' 10g-4x
Status: Success

PCA> create uplink-port-group myStoragePortGroup '1 2' 40g
Status: Success
```

4.2.10 delete config-error

The `delete config-error` command can be used to delete a failed configuration task from the configuration error database.

Syntax

```
delete config-error id [ --confirm ] [ --force ] [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

where `id` is the identifier for the configuration error that you wish to delete from the database.

Description

Use the `delete config-error` command to remove a configuration error from the configuration error database. This is a destructive operation and you are prompted to confirm whether or not you wish to continue, unless you use the `--confirm` flag to override the prompt.

Once a configuration error has been deleted from the database, you may not be able to re-run the configuration task associated with it. To obtain a list of configuration errors, use the `list config-error` command. See [Example 4.34, “List All Configuration Errors”](#) for more information.

Options

The following table shows the available options for this command.

Option	Description
<code>--confirm</code>	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
<code>--force</code>	Force the command to be executed even if the target is in an invalid state. This option is not risk-free and should only be used as a last resort.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.13 Removing a Configuration Error

```
PCA> delete config-error 87
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Success
```

4.2.11 delete lock

Removes a lock that was previously imposed on certain appliance functionality.

Syntax

```
delete lock { all_provisioning | cn_upgrade | database | install | manufacturing |
mn_upgrade | provisioning | service } [ --confirm ] [ --force ] [ --json ] [ --less ] [ --more ]
[ --tee=OUTPUTFILENAME ]
```

Description

Use the `delete lock` command to re-enable the appliance-level functions that were locked earlier.

Options

The following table shows the available options for this command.

Option	Description
{ <code>all_provisioning</code> <code>cn_upgrade</code> <code>database</code> <code>install</code> <code>manufacturing</code> <code>mn_upgrade</code> <code>provisioning</code> <code>service</code> }	The type of lock to be removed. For a description of lock types, see Section 4.2.6, “create lock” .
<code>--confirm</code>	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
<code>--force</code>	Force the command to be executed even if the target is in an invalid state. This option is not risk-free and should only be used as a last resort.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.14 Unlocking Provisioning

```
PCA> delete lock provisioning
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Success
```

4.2.12 delete network

Deletes a custom network. See [Section 2.6, “Network Customization”](#) for detailed information.

Syntax

```
delete network network-name [ --confirm ] [ --force ] [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```


where `network-name` is the name of the custom network you wish to delete.

Description

Use the `delete network` command to remove a previously created custom network from your environment. This is a destructive operation and you are prompted to confirm whether or not you wish to continue, unless you use the `--confirm` flag to override the prompt.

A custom network can only be deleted after all servers have been removed from it. See [Section 4.2.21, “remove network”](#).

Default Oracle Private Cloud Appliance networks are protected and any attempt to delete them will fail.

Options

The following table shows the available options for this command.

Option	Description
<code>--confirm</code>	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
<code>--force</code>	Force the command to be executed even if the target is in an invalid state. This option is not risk-free and should only be used as a last resort.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.15 Deleting a Custom Network

```
PCA> delete network MyNetwork
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Success
```

Example 4.16 Attempting to Delete a Default Network

```
PCA> delete network default_internal
Status: Failure
Error Message: Error (NETWORK_003): Exception while deleting network: default_internal.
['INVALID_NAME_002: Invalid Network name: default_internal. Name is reserved.']
```

4.2.13 delete task

The `delete` command can be used to delete a task from the database.

Syntax

```
delete task id [ --confirm ] [ --force ] [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

where `id` is the identifier for the task that you wish to delete from the database.

Description

Use the `delete task` command to remove a task from the task database. This is a destructive operation and you are prompted to confirm whether or not you wish to continue, unless you use the `--confirm` flag to override the prompt.

Options

The following table shows the available options for this command.

Option	Description
<code>--confirm</code>	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
<code>--force</code>	Force the command to be executed even if the target is in an invalid state. This option is not risk-free and should only be used as a last resort.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.17 Removing a Task

```
PCA> delete task 341e7bc74f339c
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Success
```

4.2.14 delete tenant-group

Deletes a tenant group. The default tenant group cannot be deleted. See [Section 2.7, “Tenant Groups”](#) for detailed information.

Syntax

```
delete tenant-group tenant-group-name [ --confirm ] [ --force ] [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

where `tenant-group-name` is the name of the tenant group – and server pool – you wish to add to the environment.

Description

Use the `delete tenant-group` command to remove a previously created, non-default tenant group from your environment. All servers must be removed from the tenant group before it can be deleted. When the tenant group is deleted, the server pool file system is removed from the internal ZFS storage.

This is a destructive operation and you are prompted to confirm whether or not you wish to continue, unless you use the `--confirm` flag to override the prompt.

Options

The following table shows the available options for this command.

Option	Description
<code>--confirm</code>	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
<code>--force</code>	Force the command to be executed even if the target is in an invalid state. This option is not risk-free and should only be used as a last resort.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.18 Deleting a Tenant Group

```
PCA> delete tenant-group myTenantGroup
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Success
```

4.2.15 delete uplink-port-group

Deletes an uplink port group. See [Section 4.2.9, “create uplink-port-group”](#) for more information about the use of uplink port groups. This command is used only on systems with Ethernet-based network architecture.

Syntax

```
delete uplink-port-group port-group-name [ --confirm ] [ --force ] [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

where `port-group-name` is the name of the uplink port group you wish to remove from the environment.

Description

Use the `delete uplink-port-group` command to remove a previously created uplink port group from your environment. If the uplink port group is used in the configuration of a network, this network must be deleted before the uplink port group can be deleted. Otherwise the delete command will fail.

This is a destructive operation and you are prompted to confirm whether or not you wish to continue, unless you use the `--confirm` flag to override the prompt.

Options

The following table shows the available options for this command.

Option	Description
<code>--confirm</code>	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
<code>--force</code>	Force the command to be executed even if the target is in an invalid state. This option is not risk-free and should only be used as a last resort.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.19 Deleting an Uplink Port Group

```
PCA> delete uplink-port-group myUplinkPortGroup
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Success
```

4.2.16 deprovision compute-node

Cleanly removes a previously provisioned compute node's records in the various configuration databases. A provisioning lock must be applied in advance, otherwise the node is reprovisioned shortly after deprovisioning.

Syntax

```
deprovision compute-node compute-node-name [ --confirm ] [ --force ] [ --json ] [ --less ]
[ --more ] [ --tee=OUTPUTFILENAME ]
```

where `compute-node-name` is the name of the compute node you wish to remove from the appliance configuration.

Description

Use the `deprovision compute-node` command to take an existing compute node out of the appliance in such a way that it can be repaired or replaced, and subsequently rediscovered as a brand new component. The compute node configuration records are removed cleanly from the system.



Caution

For deprovisioning to succeed, the compute node ILOM password must be the default *Welcome 1*. If this is not the case, the operation may result in an error. This also applies to reprovisioning an existing compute node.

By default, the command does not continue if the compute node contains running VMs. The correct workflow is to impose a provisioning lock before deprovisioning a compute node, otherwise it is rediscovered and provisioned again shortly after deprovisioning has completed. When the appliance is ready to resume its normal operations, release the provisioning lock again. For details, see [Section 4.2.6, “create lock”](#) and [Section 4.2.11, “delete lock”](#).

This is a destructive operation and you are prompted to confirm whether or not you wish to continue, unless you use the `--confirm` flag to override the prompt.

Options

The following table shows the available options for this command.

Option	Description
<code>--confirm</code>	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
<code>--force</code>	Force the command to be executed even if the target is in an invalid state. This option is not risk-free and should only be used as a last resort.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.20 Deprovisioning a Compute Node

```
deprovision compute-node ovcacn29r1
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Shutting down dhcpd: [ OK ]
```

```
Starting dhcpd: [ OK ]
Shutting down dnsmasq: [ OK ]
Starting dnsmasq: [ OK ]

Status: Success
```

4.2.17 diagnose

Performs various diagnostic checks against the Oracle Private Cloud Appliance for support purposes.



Caution

The `diagnose software` command is deprecated. It will be removed in the next release of the Oracle Private Cloud Appliance Controller Software. Diagnostic functions are now available through a separate health check tool. See [Section 2.9, “Health Monitoring”](#) for more information.

The other `diagnose` commands remain functional.

Syntax

```
diagnose { ilom | software | hardware | rack-monitor } [ --force ] [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

The following table describes each possible target of the `diagnose` command.

Command Target	Information Displayed
hardware	<p>The <code>hardware</code> diagnostic has two further options:</p> <ul style="list-style-type: none"> The <code>rack</code> option displays status information for rack components that were pingable at least once in the lifetime of the rack. The command output is real-time information. <p>If required, the results can be filtered by component type (cn, ilom, mn, etc.) Use tab completion to see all component types available.</p> <ul style="list-style-type: none"> The <code>reset</code> option must be followed by a component host name. The command resets the event counters in the monitor database to zero for the component in question. <p>If a component is or was in critical state, the reset command re-enables monitoring for that component.</p>
ilom	<p>The <code>ilom</code> diagnostic checks that the ILOM for each component is accessible on the management network.</p>
leaf-switch (Ethernet-based systems only)	<p>The <code>leaf-switch</code> diagnostic performs health checks on the leaf switches.</p>
leaf-switch-resources (Ethernet-based systems only)	<p>The <code>leaf-switch-resource</code> diagnostic checks the CPU and memory status of each leaf switch.</p>
link-status (Ethernet-based systems only)	<p>The <code>link-status</code> diagnostic returns the status of the leaf switch link ports.</p>

Command Target	Information Displayed
rack-monitor	The <code>rack-monitor</code> diagnostic checks for errors that may have been registered by the monitor service. Optionally these can be filtered per component category. If required, the results can be filtered by component type (cn, ilom, mn, etc.) Use tab completion to see all component types available.
software	The <code>software</code> diagnostic triggers the Oracle Private Cloud Appliance software acceptance tests.
spine-switch (Ethernet-based systems only)	The <code>spine-switch</code> diagnostic performs health checks on the spine switch.
spine-switch-resources (Ethernet-based systems only)	The <code>spine-switch-resource</code> diagnostic checks the CPU and memory status of the spine switch.
switch-logs (Ethernet-based systems only)	The <code>switch-logs</code> diagnostic has two further options: <ul style="list-style-type: none"> • The <code>process</code> option displays information for the processes run on the switches. • The <code>core</code> option displays information about core dumps. <p>Access the switch directly for log details.</p>
uplink-port-statistics (Ethernet-based systems only)	The <code>uplink-port-statistics</code> diagnostic displays north-south data traffic statistics for the spine switches.

Description

Use the `diagnose` command to initiate a diagnostic check of various components that make up Oracle Private Cloud Appliance.

A large part of the diagnostic information is stored in the inventory database and the monitor database. The inventory database is populated from the initial rack installation and keeps a history log of all the rack components. The monitor database stores rack component events detected by the monitor service. Some of the diagnostic commands are used to display the contents of these databases.

Options

The following table shows the available options for this command.

Option	Description
<code>--force</code>	Force the command to be executed even if the target is in an invalid state. This option is not risk-free and should only be used as a last resort.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.

Option	Description
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.
<code>--tests=TESTS</code>	Returns the output of specific tests you designate, rather than running the full set of tests.
<code>--version=VERSION</code>	Defines what version of software the command will run on. The default version is 2.4.2, but you can run the command on other version you specify here.

Examples

Example 4.21 Running the ILOM Diagnostic

```
PCA> diagnose ilom
Checking ILOM health.....please wait..

IP_Address      Status          Health_Details
-----
192.168.4.129   Not Connected   None
192.168.4.128   Not Connected   None
192.168.4.127   Not Connected   None
192.168.4.126   Not Connected   None
192.168.4.125   Not Connected   None
192.168.4.124   Not Connected   None
192.168.4.123   Not Connected   None
192.168.4.122   Not Connected   None
192.168.4.121   Not Connected   None
192.168.4.120   Not Connected   None
192.168.4.101   OK              None
192.168.4.102   OK              None
192.168.4.105   Faulty          Mon Nov 25 14:17:37 2013 Power    PS1 (Power Supply 1)
                  A loss of AC input to a power supply has occurred.
                  (Probability: 100, UUID: 2c1ec5fc-ffa3-c768-e602-ca12b86e3ea1,
                  Part Number: 07047410, Serial Number: 476856F+1252CE027X,
                  Reference Document: http://www.sun.com/msg/SPX86-8003-73)
192.168.4.107   OK              None
192.168.4.106   OK              None
192.168.4.109   OK              None
192.168.4.108   OK              None
192.168.4.112   OK              None
192.168.4.113   Not Connected   None
192.168.4.110   OK              None
192.168.4.111   OK              None
192.168.4.116   Not Connected   None
192.168.4.117   Not Connected   None
192.168.4.114   Not Connected   None
192.168.4.115   Not Connected   None
192.168.4.118   Not Connected   None
192.168.4.119   Not Connected   None
-----
27 rows displayed

Status: Success
```

Example 4.22 Running the Software Diagnostic

```
PCA> diagnose software
PCA Software Acceptance Test runner utility
Test - 01 - OpenSSL CVE-2014-0160 Heartbleed bug Acceptance [PASSED]
```



```

Test - 02 - PCA package Acceptance [PASSED]
Test - 03 - Shared Storage Acceptance [PASSED]
Test - 04 - PCA services Acceptance [PASSED]
Test - 05 - PCA config file Acceptance [PASSED]
Test - 06 - Check PCA DBs exist Acceptance [PASSED]
Test - 07 - Compute node network interface Acceptance [PASSED]
Test - 08 - OVM manager settings Acceptance [PASSED]
Test - 09 - Check management nodes running Acceptance [PASSED]
Test - 10 - Check OVM manager version Acceptance [PASSED]
Test - 11 - OVM server model Acceptance [PASSED]
Test - 12 - Repositories defined in OVM manager Acceptance [PASSED]
Test - 13 - Management Nodes have IPv6 disabled [PASSED]
Test - 14 - Bash Code Injection Vulnerability bug Acceptance [PASSED]
Test - 15 - Check Oracle VM 3.4 xen security update Acceptance [PASSED]
Test - 16 - Test for ovs-agent service on CNs Acceptance [PASSED]
Test - 17 - Test for shares mounted on CNs Acceptance [PASSED]
Test - 18 - All compute nodes running Acceptance [PASSED]
Test - 19 - PCA version Acceptance [PASSED]
Test - 20 - Check support packages in PCA image Acceptance [PASSED]

Status: Success

```

Example 4.23 Running the Leaf-Switch Diagnostic

```

PCA> diagnose leaf-switch

Switch      Health Check Name      Status
-----
ovcasw15r1  CDP Neighbor Check     Passed
ovcasw15r1  Virtual Port-channel check Passed
ovcasw15r1  Management Node Port-channel check Passed
ovcasw15r1  Leaf-Spine Port-channel check Passed
ovcasw15r1  OSPF Neighbor Check    Passed
ovcasw15r1  Multicast Route Check  Passed
ovcasw15r1  Leaf Filesystem Check  Passed
ovcasw15r1  Hardware Diagnostic Check Passed
ovcasw16r1  CDP Neighbor Check     Passed
ovcasw16r1  Virtual Port-channel check Passed
ovcasw16r1  Management Node Port-channel check Passed
ovcasw16r1  Leaf-Spine Port-channel check Passed
ovcasw16r1  OSPF Neighbor Check    Passed
ovcasw16r1  Multicast Route Check  Passed
ovcasw16r1  Leaf Filesystem Check  Passed
ovcasw16r1  Hardware Diagnostic Check Passed
-----
16 rows displayed

Status: Success

```

4.2.18 get log

Retrieves the log files from the selected components and saves them to a directory on the rack's shared storage.



Note

Currently the spine or data switch is the only target component supported with this command.

Syntax

```
get log component [ --confirm ] [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

where `component` is the identifier of the rack component from which you want to retrieve the log files.

Description

Use the `get log` command to collect the log files of a given rack component or set of rack components of a given type. The command output indicates where the log files are saved: this is a directory on the internal storage appliance in a location that both management nodes can access. From this location you can examine the logs or copy them to your local system so they can be included in your communication with Oracle.

Options

The following table shows the available options for this command.

Option	Description
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.24 Collecting the Log Files from the Spine Switch

Note that the CLI uses 'data_switch' as the internal alias for a spine Cisco Nexus 9336C-FX2 Switch.

```
PCA> get log data_switch
Log files copied to: /nfs/shared_storage/incoming
Status: Success
```

4.2.19 list

The `list` command can be used to list the different components and tasks within the Oracle Private Cloud Appliance. The output displays information relevant to each component or task. Output from the `list` command is usually tabulated so that different fields appear as columns for each row of information relating to the command target.

Syntax

```
list { backup-task | compute-node | config-error | lock | management-node | mgmt-switch-
port | network | network-card | network-port | network-switch | ofm-network | opus-port
| server-profile | storage-network | task | tenant-group | update-task | uplink-port |
uplink-port-group | wwpn-info } [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME
] [ [ --sorted-by SORTEDBY | --sorted-order SORTEDORDER ] ] [ [ --filter-column FILTERCOLUMN | --filter
FILTER ] ]
```

where `SORTEDBY` is one of the table column names returned for the selected command target, and `SORTEDORDER` can be either `ASC` for an ascending sort, or `DES` for a descending sort. See [Section 4.1.3.2, “Sorting”](#) for more information.

where *FILTERCOLUMN* is one of the table column names returned for the selected command target, and *FILTER* is the text that you wish to match to perform your filtering. See [Section 4.1.3.3, “Filtering”](#) for more information.

The following table describes each possible target of the `list` command.

Command Target	Information Displayed
backup-task	Displays basic information about all backup tasks.
compute-node	Displays basic information for all compute nodes installed.
config-error	Displays all configuration tasks that were not completed successfully and ended in an error.
lock	Displays all locks that have been imposed.
management-node	Displays basic information for both management nodes.
mgmt-switch-port	Displays connection information about every port in the Oracle Private Cloud Appliance environment belonging to the internal administration or management network. The ports listed can belong to a switch, a server node or any other connected rack component type.
network	Displays all networks configured in the environment.
network-card (InfiniBand-based systems only)	Displays information about the I/O modules installed in the Fabric Interconnects.
network-port	Displays the status of all ports on all I/O modules installed in the networking components.
network-switch (Ethernet-based systems only)	Displays basic information about all switches installed in the Oracle Private Cloud Appliance environment.
ofm-network (InfiniBand-based systems only)	Displays network configuration, read directly from the Oracle Fabric Manager software on the Fabric Interconnects.
opus-port (InfiniBand-based systems only)	Displays connection information about every port of every Oracle Switch ES1-24 in the Oracle Private Cloud Appliance environment.
server-profile (InfiniBand-based systems only)	Displays a list of connectivity profiles for servers, as stored by the Fabric Interconnects. The profile contains essential networking and storage information for the server in question.
storage-network (InfiniBand-based systems only)	Displays a list of known storage clouds. The configuration of each storage cloud contains information about participating Fabric Interconnect ports and server vHBAs.
task	Displays a list of running, completed and failed tasks.
tenant-group	Displays all configured tenant groups. The list includes the default configuration as well as custom tenant groups.

Command Target	Information Displayed
update-task	Displays a list of all software update tasks that have been started on the appliance.
uplink-port (Ethernet-based systems only)	Displays information about spine switch port configurations for external networking.
uplink-port-group (Ethernet-based systems only)	Displays information about all uplink port groups configured for external networking.
wwpn-info (InfiniBand-based systems only)	Displays a list of all World Wide Port Names (WWPNs) for all ports participating in the Oracle Private Cloud Appliance Fibre Channel fabric. In the standard configuration each compute node has a vHBA in each of the four default storage clouds.

Note that you can use tab completion to help you correctly specify the `object` for the different command targets. You do not need to specify an `object` if the command target is `system-properties` or `version`.

Description

Use the `list` command to obtain tabulated listings of information about different components or activities within the Oracle Private Cloud Appliance. The `list` command can frequently be used to obtain identifiers that can be used in conjunction with many other commands to perform various actions or to obtain more detailed information about a specific component or task. The `list` command also supports sorting and filtering capabilities to allow you to order information or to limit information so that you are able to identify specific items of interest quickly and easily.

Options

The following table shows the available options for this command.

Option	Description
{ <code>backup-task</code> <code>compute-node</code> <code>config-error</code> <code>lock</code> <code>management-node</code> <code>network</code> <code>network-card</code> <code>network-port</code> <code>ofm-network</code> <code>opus-port</code> <code>server-profile</code> <code>storage-network</code> <code>task</code> <code>tenant-group</code> <code>update-task</code> <code>wwpn-info</code> }	The command target to list information for.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.
[<code>--sorted-by SORTEDBY</code>]	Sort the table by the values within a particular column in the table, specified by replacing <code>SORTEDBY</code> with the name of the column that should be used to perform the sort.

Option	Description
[<code>--sorted-order SORTEDORDER</code>]	Used to specify the sort order, which can either be <code>ASC</code> for an ascending sort, or <code>DES</code> for a descending sort. You must use the <code>--sorted-by</code> option in conjunction with this option.
[<code>--filter-column FILTERCOLUMN</code>]	Filter the table for a value within a particular column in the table, specified by replacing <code>FILTERCOLUMN</code> with the name of the column that should be used to perform the sort. You must use the <code>--filter</code> option in conjunction with this option.
[<code>--filter FILTER</code>]	The filter that should be applied to values within the column specified by the <code>--filter-column</code> option.

Examples

Example 4.25 List all management nodes

```
PCA> list management-node
```

Management_Node	IP_Address	Provisioning_Status	ILOM_MAC	Provisioning_State	Master
ovcamn05r1	192.168.4.3	RUNNING	00:10:e0:e9:1f:c9	running	None
ovcamn06r1	192.168.4.4	RUNNING	00:10:e0:e7:26:ad	running	Yes

```
-----
2 rows displayed

Status: Success
```

Example 4.26 List all compute nodes

```
PCA> list compute-node
```

Compute_Node	IP_Address	Provisioning_Status	ILOM_MAC	Provisioning_State
ovcacn10r1	192.168.4.7	RUNNING	00:10:e0:65:2f:4b	running
ovcacn08r1	192.168.4.5	RUNNING	00:10:e0:65:2f:f3	initializing_stage_wait_...
ovcacn09r1	192.168.4.10	RUNNING	00:10:e0:62:98:e3	running
ovcacn07r1	192.168.4.8	RUNNING	00:10:e0:65:2f:93	running

```
-----
4 rows displayed

Status: Success
```

Example 4.27 List All Tenant Groups

```
PCA> list tenant-group
```

Name	Default	State
Rack1_ServerPool	True	ready
myTenantGroup	False	ready

```
-----
2 rows displayed

Status: Success
```

Example 4.28 List Appliance Networks

```
PCA> list network
```

Network_Name	Default	Type	Trunkmode	Description
custom_internal	False	rack_internal_network	None	None
default_internal	True	rack_internal_network	None	None
storage_net	False	host_network	None	None
default_external	True	external_network	None	None

```
-----
4 rows displayed
```

```
Status: Success
```

Example 4.29 List the Network Ports Configured on the Spine Cisco Nexus 9336C-FX2 Switches

```
PCA> list network-port
```

Port	Switch	Type	State	Networks
1	ovcasw22r1	40G	up	storage_net
2	ovcasw22r1	40G	up	storage_net
3	ovcasw22r1	auto-speed	down	None
4	ovcasw22r1	auto-speed	down	None
5:1	ovcasw22r1	10G	up	default_external
5:2	ovcasw22r1	10G	down	default_external
5:3	ovcasw22r1	10G	down	None
5:4	ovcasw22r1	10G	down	None
1	ovcasw23r1	40G	up	storage_net
2	ovcasw23r1	40G	up	storage_net
3	ovcasw23r1	auto-speed	down	None
4	ovcasw23r1	auto-speed	down	None
5:1	ovcasw23r1	10G	up	default_external
5:2	ovcasw23r1	10G	down	default_external
5:3	ovcasw23r1	10G	down	None
5:4	ovcasw23r1	10G	down	None

```
-----
16 rows displayed
```

```
Status: Success
```

Example 4.30 List Ports on the Management Cisco Nexus 9348GC-FXP Switch Using a Filter

Note that the CLI uses the internal alias `mgmt-switch-port`. In this example the command displays all internal Ethernet connections from compute nodes to the Cisco Nexus 9348GC-FXP Switch. A wildcard is used in the `--filter` option.

```
PCA> list mgmt-switch-port --filter-column=Hostname --filter=*cn*r1
```

Dest	Dest_Port	Hostname	Key	MGMTSWITCH	RACK	RU	Src_Port	Type
07	Net-0	ovcacn07r1	CISCO-1-5	CISCO-1	1	7	5	compute
08	Net-0	ovcacn08r1	CISCO-1-6	CISCO-1	1	8	6	compute
09	Net-0	ovcacn09r1	CISCO-1-7	CISCO-1	1	9	7	compute
10	Net-0	ovcacn10r1	CISCO-1-8	CISCO-1	1	10	8	compute
11	Net-0	ovcacn11r1	CISCO-1-9	CISCO-1	1	11	9	compute
12	Net-0	ovcacn12r1	CISCO-1-10	CISCO-1	1	12	10	compute
13	Net-0	ovcacn13r1	CISCO-1-11	CISCO-1	1	13	11	compute
14	Net-0	ovcacn14r1	CISCO-1-12	CISCO-1	1	14	12	compute
34	Net-0	ovcacn34r1	CISCO-1-15	CISCO-1	1	34	15	compute
35	Net-0	ovcacn35r1	CISCO-1-16	CISCO-1	1	35	16	compute
36	Net-0	ovcacn36r1	CISCO-1-17	CISCO-1	1	36	17	compute
37	Net-0	ovcacn37r1	CISCO-1-18	CISCO-1	1	37	18	compute
38	Net-0	ovcacn38r1	CISCO-1-19	CISCO-1	1	38	19	compute
39	Net-0	ovcacn39r1	CISCO-1-20	CISCO-1	1	39	20	compute
40	Net-0	ovcacn40r1	CISCO-1-21	CISCO-1	1	40	21	compute
41	Net-0	ovcacn41r1	CISCO-1-22	CISCO-1	1	41	22	compute
42	Net-0	ovcacn42r1	CISCO-1-23	CISCO-1	1	42	23	compute
26	Net-0	ovcacn26r1	CISCO-1-35	CISCO-1	1	26	35	compute
27	Net-0	ovcacn27r1	CISCO-1-36	CISCO-1	1	27	36	compute
28	Net-0	ovcacn28r1	CISCO-1-37	CISCO-1	1	28	37	compute
29	Net-0	ovcacn29r1	CISCO-1-38	CISCO-1	1	29	38	compute
30	Net-0	ovcacn30r1	CISCO-1-39	CISCO-1	1	30	39	compute
31	Net-0	ovcacn31r1	CISCO-1-40	CISCO-1	1	31	40	compute
32	Net-0	ovcacn32r1	CISCO-1-41	CISCO-1	1	32	41	compute
33	Net-0	ovcacn33r1	CISCO-1-42	CISCO-1	1	33	42	compute

```
-----
25 rows displayed
```

```
Status: Success
```

Example 4.31 List All Tasks

```
PCA> list task
```

Task_ID	Status	Progress	Start_Time	Task_Name
376a676449206a	SUCCESS	100	06-06-2019 09:00:01	backup
376ce11fc6c39c	SUCCESS	100	06-06-2019 04:23:41	update_download_image
376a02cf798f68	SUCCESS	100	06-05-2019 21:00:02	backup
376c7c8afcc86a	SUCCESS	100	06-05-2019 09:00:01	backup

```
-----
4 rows displayed
```

```
Status: Success
```

Example 4.32 List Uplink Ports to Configure External Networking

```
PCA> list uplink-port
```

Interface Name	Switch	Status	Admin_Status	PortChannel	Speed
Ethernet1/1	ovcasw22r1	up	up	111	40G
Ethernet1/1	ovcasw23r1	up	up	111	40G
Ethernet1/2	ovcasw22r1	up	up	111	40G
Ethernet1/2	ovcasw23r1	up	up	111	40G
Ethernet1/3	ovcasw22r1	down	down	None	auto
Ethernet1/3	ovcasw23r1	down	down	None	auto
Ethernet1/4	ovcasw22r1	down	down	None	auto
Ethernet1/4	ovcasw23r1	down	down	None	auto
Ethernet1/5/1	ovcasw22r1	up	up	151	10G
Ethernet1/5/1	ovcasw23r1	up	up	151	10G
Ethernet1/5/2	ovcasw22r1	down	up	151	10G
Ethernet1/5/2	ovcasw23r1	down	up	151	10G
Ethernet1/5/3	ovcasw22r1	down	down	None	10G
Ethernet1/5/3	ovcasw23r1	down	down	None	10G
Ethernet1/5/4	ovcasw22r1	down	down	None	10G
Ethernet1/5/4	ovcasw23r1	down	down	None	10G

```
-----
16 rows displayed
```

```
Status: Success
```

Example 4.33 List Uplink Port Groups

```
PCA> list uplink-port-group
```

Port_Group_Name	Ports	Mode	Speed	Breakout_Mode	Enabled	State
default_5_1	5:1 5:2	LAG	10g	10g-4x	True	(up)* Not all ports are up
default_5_2	5:3 5:4	LAG	10g	10g-4x	False	down

```
-----
2 rows displayed
```

```
Status: Success
```

Example 4.34 List All Configuration Errors

```
PCA> list config-error
```

ID	Module	Host	Timestamp
87	Management node password	192.168.4.4	Mon Jun 03 02:45:42 2019

```
54      MySQL management password 192.168.4.216 Mon Jun 03 02:44:54 2019
-----
2 rows displayed
Status: Success
```

4.2.20 remove compute-node

Removes a compute node from an existing tenant group.

Syntax

```
remove compute-node node tenant-group-name [ --confirm ] [ --force ] [ --json ] [ --less ]
[ --more ] [ --tee=OUTPUTFILENAME ]
```

where `tenant-group-name` is the name of the tenant group you wish to remove one or more compute nodes from, and `node` is the name of the compute node that should be removed from the selected tenant group.

Description

Use the `remove compute-node` command to remove the required compute nodes from their tenant group. Use Oracle VM Manager to prepare the compute nodes first: make sure that virtual machines have been migrated away from the compute node, and that no storage repositories are presented. Custom networks associated with the tenant group are removed from the compute node, not from the tenant group.

This is a destructive operation and you are prompted to confirm whether or not you wish to continue, unless you use the `--confirm` flag to override the prompt.

Options

The following table shows the available options for this command.

Option	Description
<code>--confirm</code>	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
<code>--force</code>	Force the command to be executed even if the target is in an invalid state. This option is not risk-free and should only be used as a last resort.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.35 Removing a Compute Node from a Tenant Group

```
PCA> remove compute-node ovcacn09r1 myTenantGroup
```



```
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Success
```

4.2.21 remove network

Disconnects a server node from a network.

Syntax

```
remove network network-name node [ --confirm ] [ --force ] [ --json ] [ --less ] [ --more ] [
--tee=OUTPUTFILENAME ]
```

where `network-name` is the name of the network from which you wish to disconnect one or more servers, and `node` is the name of the server node that should be disconnected from the selected network.

Description

Use the `remove network` command to disconnect server nodes from a custom network you created. In case you want to delete a custom network from your environment, you must first disconnect all the servers from that network. Then use the `delete network` command to delete the custom network configuration. This is a destructive operation and you are prompted to confirm whether or not you wish to continue, unless you use the `--confirm` flag to override the prompt.

Options

The following table shows the available options for this command.

Option	Description
<code>--confirm</code>	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
<code>--force</code>	Force the command to be executed even if the target is in an invalid state. This option is not risk-free and should only be used as a last resort.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.36 Disconnecting a Compute Node from a Custom Network

```
PCA> remove network MyNetwork ovcacn09r1
*****
```

```
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Success
```

4.2.22 remove network-from-tenant-group

Removes a custom network from a tenant group.

Syntax

```
remove network-from-tenant-group network-name tenant-group-name [ --confirm ] [ --force ] [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

where `network-name` is the name of a custom network associated with a tenant group, and `tenant-group-name` is the name of the tenant group you wish to remove the custom network from.

Description

Use the `remove network-from-tenant-group` command to break the association between a custom network and a tenant group. The custom network is unconfigured from all tenant group member servers.

This is a destructive operation and you are prompted to confirm whether or not you wish to continue, unless you use the `--confirm` flag to override the prompt.

Options

The following table shows the available options for this command.

Option	Description
<code>--confirm</code>	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
<code>--force</code>	Force the command to be executed even if the target is in an invalid state. This option is not risk-free and should only be used as a last resort.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.37 Removing a Custom Network from a Tenant Group

```
PCA> remove network-from-tenant-group myPublicNetwork myTenantGroup
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
```

```
Are you sure [y/N]:y
Status: Success
```

4.2.23 reprovision

The `reprovision` command can be used to trigger reprovisioning for a specified compute node within the Oracle Private Cloud Appliance.



Caution

Reprovisioning restores a compute node to a clean state. If a compute node was previously added to the Oracle VM environment and has active connections to storage repositories other than those on the internal ZFS storage, the external storage connections need to be configured again after reprovisioning.

Syntax

```
reprovision { compute-node } node [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ] [ --force ] [ --save-local-repo ]
```

where `node` is the compute node name for the compute node that should be reprovisioned.

Description

Use the `reprovision` command to reprovision a specified compute node. The provisioning process is described in more detail in [Section 1.4, “Provisioning and Orchestration”](#).

The `reprovision` command triggers a task that is responsible for handling the reprovisioning process and exits immediately with status 'Success' if the task has been successfully generated. This does not mean that the reprovisioning process itself has completed successfully. To monitor the status of the reprovisioning task, you can use the `list compute-node` command to check the provisioning state of the servers. You can also monitor the log file for information relating to provisioning tasks. The location of the log file can be obtained by checking the `Log_File` parameter when you run the `show system-properties` command. See [Example 4.43, “Show System Properties”](#) for more information.

Options

The following table shows the available options for this command.

Option	Description
<code>compute-node</code>	The command target to perform the reprovision operation against.
<code>--save-local-repo</code>	Skip the HMP step in the provisioning process in order to save the local storage repository.
<code>--json</code>	Return the output of the command in JSON format.
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.

Option	Description
<code>--force</code>	Force the command to be executed even if the target is in an invalid state. This option is not risk-free and should only be used as a last resort.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.38 Reprovisioning a Compute Node



Caution

Do not force reprovisioning on a compute node with running virtual machines because they will be left in an indeterminate state.

```
PCA> reprovision compute-node ovcacn11r1
The reprovision job has been submitted.
Use "show compute-node <compute node name>" to monitor the progress.
Status: Success
```

4.2.24 rerun

Triggers a configuration task to re-run on the Oracle Private Cloud Appliance.

Syntax

```
rerun { config-task } id [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

where `id` is the identifier for the configuration task that must be re-run.

Description

Use the `rerun` command to re-initiate a configuration task that has failed. Use the `list config-error` command to view the configuration tasks that have failed and the associated identifier that you should use in conjunction with this command. See [Example 4.34, “List All Configuration Errors”](#) for more information.

Options

The following table shows the available options for this command.

Option	Description
<code>config-task</code>	The command target to perform the rerun operation against.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.39 Re-run a configuration task

```
PCA> rerun config-task 84
Status: Success
```

4.2.25 set system-property

Sets the value for a system property on the Oracle Private Cloud Appliance.

Syntax

```
set system-property { ftp_proxy | http_proxy | https_proxy | log_count |
log_file | log_level | log_size | timezone } value [ --json ] [ --less ] [ --more ] [ --
tee=OUTPUTFILENAME ]
```

where `value` is the value for the system property that you are setting.

Description

Use the `set system-property` command to set the value for a system property on the Oracle Private Cloud Appliance.



Important

The `set system-property` command only affects the settings for the management node where it is run. If you change a setting on the active management node, using this command, you should connect to the passive management node and run the equivalent command there as well, to keep the two systems synchronized. This is the only exception where it is necessary to run a CLI command on the passive management node.

You can use the `show system-properties` command to view the values of various system properties at any point. See [Example 4.43, "Show System Properties"](#) for more information.




Important

Changes to system-properties usually require that you restart the service for the change to take effect. To do this, you must run `service ovca restart` in the shell of the active management node after you have set the system property value.

Options

The following table shows the available options for this command.

Option	Description
<code>ftp_proxy</code>	Set the value for the IP address of an FTP Proxy
<code>http_proxy</code>	Set the value for the IP address of an HTTP Proxy
<code>https_proxy</code>	Set the value for the IP address of an HTTPS Proxy
<code>log_count</code>	Set the value for the number of log files that should be retained through log rotation
<code>log_file</code>	Set the value for the location of a particular log file.

Option	Description
	<div style="display: flex; align-items: flex-start;"> <div style="margin-right: 20px;">  </div> <div> <p>Caution</p> <p>Make sure that the new path to the log file exists. Otherwise, the log server stops working.</p> <p>The system always prepends <code>/var/log</code> to your entry. Absolute paths are converted to <code>/var/log/<path></code>.</p> </div> </div> <p>This property can be defined separately for the following log files: backup, cli, diagnosis, monitor, ovca, snmp, and syncservice.</p>
<code>log_level</code>	<p>Set the value for the log level output. Accepted log levels are: CRITICAL, DEBUG, ERROR, INFO, WARNING.</p> <p>This property can be defined separately for the following log files: backup, cli, diagnosis, monitor, ovca, snmp, and syncservice. Use tab completion to insert the log file in the command before the log level value.</p>
<code>log_size</code>	<p>Set the value for the maximum log size before a log is rotated</p>
<code>timezone</code>	<p>Set the time zone for the location of the Oracle Private Cloud Appliance.</p> <p>There are several hundred options, and the selection is case sensitive. It is suggested to use tab completion to find the most accurate setting for your location.</p>
<code>--json</code>	<p>Return the output of the command in JSON format</p>
<code>--less</code>	<p>Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.</p>
<code>--more</code>	<p>Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.</p>
<code>--tee=OUTPUTFILENAME</code>	<p>When returning the output of the command, also write it to the specified output file.</p>

Examples

Example 4.40 Changing the location of the sync service log file

```
PCA> set system-property log_file syncservice sync/ovca-sync.log
Status: Success

PCA> show system-properties
-----
[...]
Backup.Log_File      /var/log/ovca-backup.log
Backup.Log_Level     DEBUG
Cli.Log_File         /var/log/ovca-cli.log
Cli.Log_Level        DEBUG
Sync.Log_File        /var/log/sync/ovca-sync.log
```

```

Sync.Log_Level      DEBUG
Diagnosis.Log_File  /var/log/ovca-diagnosis.log
Diagnosis.Log_Level DEBUG
[...]
-----
Status: Success

```

**Note**

Log configuration through the CLI is described in more detail in [Section 7.1, “Setting the Oracle Private Cloud Appliance Logging Parameters”](#).

Example 4.41 Configuring and unconfiguring an HTTP proxy

```

PCA> set system-property http_proxy http://10.1.1.11:8080
Status: Success

PCA> set system-property http_proxy ''
Status: Success

```

**Note**

Proxy configuration through the CLI is described in more detail in [Section 7.2, “Adding Proxy Settings for Oracle Private Cloud Appliance Updates”](#).

Example 4.42 Configuring the Oracle Private Cloud Appliance Time Zone

```

PCA> set system-property timezone US/Eastern
Status: Success

```

4.2.26 show

The `show` command can be used to view information about particular objects such as tasks, rack layout or system properties. Unlike the `list` command, which applies to a whole target object type, the `show` command displays information specific to a particular target object. Therefore, it is usually run by specifying the command, the target object type and the object identifier.

Syntax

```

show { cloud-wwpn | compute-node | network | rack-layout | server-profile | storage-
network | system-properties | task | tenant-group | version | vhba-info } object [ --json ]
[ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]

```

Where `object` is the identifier for the target object that you wish to show information for. The following table provides a mapping of identifiers that should be substituted for `object`, depending on the command target.

Command Target	Object Identifier
cloud-wwpn (InfiniBand-based systems only)	Storage Network/Cloud Name
compute-node	Compute Node Name
network	Network Name
rack-layout	Rack Architecture or Type
server-profile (InfiniBand-based systems only)	Server Name
storage-network (InfiniBand-based systems only)	Storage Network/Cloud Name
system-properties	(none)

Command Target	Object Identifier
task	Task ID
tenant-group	Tenant Group Name
version	(none)
vhba-info (InfiniBand-based systems only)	Compute Node Name

Note that you can use tab completion to help you correctly specify the `object` for the different command targets. You do not need to specify an `object` if the command target is `system-properties` or `version`.

Description

Use the `show` command to view information specific to a particular target object, identified by specifying the identifier for the object that you wish to view. The exception to this is the option to view `system-properties`, for which no identifier is required.

Frequently, the `show` command may display information that is not available using the `list` command in conjunction with its filtering capabilities.

Options

The following table shows the available options for this command.

Option	Description
<code>cloud-wwpn compute-node network rack-layout server-profile storage-network system-properties task tenant-group version vhba-info</code>	The command target to show information for.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.43 Show System Properties



Note

This command only displays the system properties for the management node where it is run. If the system properties have become unsynchronized across the two management nodes, the information reflected by this command may not apply to both systems. You can run this command on either the active or passive management node if you need to check that the configurations match.


```
PCA> show system-properties
```

```
-----
HTTP_Proxy          None
HTTPS_Proxy         None
FTP_Proxy           None
Log_File            /var/log/ovca.log
Log_Level           DEBUG
Log_Size (MB)       250
Log_Count           5
Timezone            Etc/UTC
Backup.Log_File     /var/log/ovca-backup.log
Backup.Log_Level    DEBUG
Cli.Log_File        /var/log/ovca-cli.log
Cli.Log_Level       DEBUG
Sync.Log_File       /var/log/ovca-sync.log
Sync.Log_Level      DEBUG
Diagnosis.Log_File  /var/log/ovca-diagnosis.log
Diagnosis.Log_Level DEBUG
Monitor.Log_File    /var/log/ovca-monitor.log
Monitor.Log_Level   INFO
Snmp.Log_File       /nfs/shared_storage/logs/ovca_snmptrapd.log
Snmp.Log_Level      DEBUG
-----
```

```
Status: Success
```

Example 4.44 Show Task

```
PCA> show task 341e7bc74f339c
```

```
-----
Task_Name           backup
Status              RUNNING
Progress            70
Start_Time          05-27-2019 09:59:36
End_Time            None
Pid                 1503341
Result              None
-----
```

```
Status: Success
```

Example 4.45 Show Rack Layout

```
PCA> show rack-layout x8-2_base
```

RU	Name	Role	Type	Sub_Type	Units
--	----	----	----	-----	-----
42	ovcacn42r1	compute	compute		[42]
41	ovcacn41r1	compute	compute		[41]
40	ovcacn40r1	compute	compute		[40]
39	ovcacn39r1	compute	compute		[39]
38	ovcacn38r1	compute	compute		[38]
37	ovcacn37r1	compute	compute		[37]
36	ovcacn36r1	compute	compute		[36]
35	ovcacn35r1	compute	compute		[35]
34	ovcacn34r1	compute	compute		[34]
33	ovcacn33r1	compute	compute		[33]
32	ovcacn32r1	compute	compute		[32]
31	ovcacn31r1	compute	compute		[31]
30	ovcacn30r1	compute	compute		[30]
29	ovcacn29r1	compute	compute		[29]
28	ovcacn28r1	compute	compute		[28]
27	ovcacn27r1	compute	compute		[27]
26	ovcacn26r1	compute	compute		[26]

show

```
25 N / A      infrastructure filler [25, 24]
24 N / A      infrastructure filler [25, 24]
23 ovcasw23r1 infrastructure cisco-data cisco4 [23]
22 ovcasw22r1 infrastructure cisco-data cisco3 [22]
21 ovcasw21r1 infrastructure cisco [21]
20 N / A      infrastructure zfs-storage disk-shelf [20, 19, 18, 17]
19 N / A      infrastructure zfs-storage disk-shelf [20, 19, 18, 17]
18 N / A      infrastructure zfs-storage disk-shelf [20, 19, 18, 17]
17 N / A      infrastructure zfs-storage disk-shelf [20, 19, 18, 17]
16 ovcasw16r1 infrastructure cisco-data cisco2 [16]
15 ovcasw15r1 infrastructure cisco-data cisco1 [15]
14 ovcacn14r1 compute compute [14]
13 ovcacn13r1 compute compute [13]
12 ovcacn12r1 compute compute [12]
11 ovcacn11r1 compute compute [11]
10 ovcacn10r1 compute compute [10]
9 ovcacn09r1 compute compute [9]
8 ovcacn08r1 compute compute [8]
7 ovcacn07r1 compute compute [7]
6 ovcamn06r1 infrastructure management management2 [6]
5 ovcamn05r1 infrastructure management management1 [5]
4 ovcasn02r1 infrastructure zfs-storage zfs-head2 [4, 3]
3 ovcasn02r1 infrastructure zfs-storage zfs-head2 [4, 3]
2 ovcasn01r1 infrastructure zfs-storage zfs-head1 [2, 1]
1 ovcasn01r1 infrastructure zfs-storage zfs-head1 [2, 1]
0 ovcapduBr1 infrastructure pdu pdu2 [0]
0 ovcapduAr1 infrastructure pdu pdu1 [0]
```

44 rows displayed

Status: Success

Example 4.46 Show the Configuration Details of the default_external Network

```
PCA> show network default_external
```

```
-----
Network_Name      default_external
Trunkmode         None
Description       None
Ports            ['5:1', '5:2']
vNICs            None
Status           ready
Network_Type      external_network
Compute_Nodes     ovcacn12r1, ovcacn07r1, ovcacn13r1, ovcacn14r1, ovcacn10r1, ovcacn09r1, ovcacn11r1
Prefix           192.168.200.0/21
Netmask          None
Route_Destination None
Route_Gateway    None
-----
```

Status: Success

Example 4.47 Show Details of a Tenant Group

```
PCA> show tenant-group myTenantGroup
```

```
-----
Name              myTenantGroup
Default          False
Tenant_Group_ID  0004fb0000020000155c15e268857a78
Servers          ['ovcacn09r1', 'ovcacn10r1']
State            ready
Tenant_Group_VIP None
Tenant_Networks ['myPublicNetwork']
Pool_Fileystem_ID 3600144f0d29d4c86000057162ecc0001
```

```
-----
Status: Success
```

Example 4.48 Show Details of a Custom Network

```
PCA> show network myHostNetwork
```

```
-----
Network_Name      myHostNetwork
Trunkmode         None
Description       None
Ports            ['1', '2']
vNICs            None
Status           ready
Network_Type     host_network
Compute_Nodes    ovcacn42r1, ovcacn01r2, ovcacn02r2
Prefix           10.10.10
Netmask          255.255.240.0
Route_Destination 10.10.20.0/24
Route_Gateway    10.10.10.250
-----
```

```
Status: Success
```

Example 4.49 Show the WWPNs for a Storage Network

```
PCA> show cloud-wwpn Cloud_A
```

```
-----
Cloud_Name      Cloud_A
WWPN_List       50:01:39:70:00:58:91:1C, 50:01:39:70:00:58:91:1A,
                50:01:39:70:00:58:91:18, 50:01:39:70:00:58:91:16,
                50:01:39:70:00:58:91:14, 50:01:39:70:00:58:91:12,
                50:01:39:70:00:58:91:10, 50:01:39:70:00:58:91:0E,
                50:01:39:70:00:58:91:0C, 50:01:39:70:00:58:91:0A,
                50:01:39:70:00:58:91:08, 50:01:39:70:00:58:91:06,
                50:01:39:70:00:58:91:04, 50:01:39:70:00:58:91:02,
                50:01:39:70:00:58:91:00
-----
```

```
Status: Success
```

Example 4.50 Show the vHBA configuration for a Compute Node

```
PCA> show vhma-info ovcacn10r1
```

vHBA_Name	Cloud	WWNN	WWPN
vhba03	Cloud_C	50:01:39:71:00:58:B1:04	50:01:39:70:00:58:B1:04
vhba02	Cloud_B	50:01:39:71:00:58:91:05	50:01:39:70:00:58:91:05
vhba01	Cloud_A	50:01:39:71:00:58:91:04	50:01:39:70:00:58:91:04
vhba04	Cloud_D	50:01:39:71:00:58:B1:05	50:01:39:70:00:58:B1:05

```
-----
4 rows displayed
```

```
Status: Success
```

Example 4.51 Show Oracle Private Cloud Appliance Version Information

```
PCA> show version
```

```
-----
Version          2.4.1
```

```
Build      819
Date      2019-06-20
-----
```

```
Status: Success
```

4.2.27 start

Starts up a rack component.



Caution

The `start` command is deprecated. It will be removed in the next release of the Oracle Private Cloud Appliance Controller Software.

Syntax

```
start { compute-node CN | management-node MN } [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

where `CN` refers to the name of the compute node and `MN` refers to the name of the management node to be started.

Description

Use the `start` command to boot a compute node or management node. You must provide the host name of the server you wish to start.

Options

The following table shows the available options for this command.

Option	Description
<code>compute-node CN management-node MN</code>	Start either a compute node or a management node. Replace <code>CN</code> or <code>MN</code> respectively with the host name of the server to be started.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.52 Starting a Compute Node

```
PCA> start compute-node ovcacn11r1
Status: Success
```

4.2.28 stop

Shuts down a rack component or aborts a running task.



Caution

The `stop` commands to shut down rack components are deprecated. It will be removed in the next release of the Oracle Private Cloud Appliance Controller Software.

The other `stop` commands, to abort tasks, remain functional.

Syntax

```
stop { compute-node CN | management-node MN | task id | update-task id } [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

where `CN` or `MN` refers to the name of the server to be shut down, and `id` refers to the identifier of the task to be aborted.

Description

Use the `stop` command to shut down a compute node or management node or to abort a running task. Depending on the command target you must provide either the host name of the server you wish to shut down, or the unique identifier of the task you wish to abort. This is a destructive operation and you are prompted to confirm whether or not you wish to continue, unless you use the `--confirm` flag to override the prompt.

Options

The following table shows the available options for this command.

Option	Description
<code>compute-node CN management-node MN</code>	<p>Shut down either a compute node or a management node. Replace <code>CN</code> or <code>MN</code> respectively with the host name of the server to be shut down.</p> <div style="display: flex; align-items: center;"> <div> <p>Caution</p> <p>These options are deprecated.</p> </div> </div>
<code>task id update-task id</code>	<p>Aborts a running task.</p> <p>Use the <code>update-task</code> target type specifically to abort a software update task. It does not take a task ID as an argument, but the management node IP address.</p> <div style="display: flex; align-items: center;"> <div> <p>Caution</p> <p>Stopping an update task is a risky operation and should be used with extreme caution.</p> </div> </div>
<code>--confirm</code>	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
<code>--json</code>	Return the output of the command in JSON format

Option	Description
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.53 Aborting a Task

```
PCA> stop task 341d45b5424c16
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Success
```

4.2.29 update appliance

This command is deprecated. Its functionality is part of the Oracle Private Cloud Appliance Upgrader.



Caution

Release 2.4.1 is for factory installation only. It cannot be used for field updates or upgrade operations on existing appliance environments.

4.2.30 update password

Modifies the password for one or more components within the Oracle Private Cloud Appliance.

Syntax

```
update password { LeafSwitch-admin | MgmtNetSwitch-admin | SpineSwitch-admin | mgmt-
root | mysql-appfw | mysql-ovs | mysql-root | ovm-admin | spCn-root | spMn-root | spZfs-
root | system-root | wls-weblogic | zfs-root } [ PCA-password target-password ] [ --
confirm ] [ --force ] [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

where `PCA-password` is the current password of the Oracle Private Cloud Appliance admin user, and `target-password` is the new password to be applied to the target rack component.

Description

Use the `update password` command to modify the password for one or more components within the Oracle Private Cloud Appliance. This is a destructive operation and you are prompted to confirm whether or not you wish to continue, unless you use the `--confirm` flag to override the prompt.

Optionally you provide the current Oracle Private Cloud Appliance password and the new target component password with the command. If not, you are prompted for the current password of the Oracle Private Cloud Appliance admin user and for the new password that should be applied to the target.

Options

The following table shows the available options for this command.

Option	Description
<code>LeafSwitch-admin</code>	Sets a new password for the <code>admin</code> user on the leaf Cisco Nexus 9336C-FX2 Switches.
<code>MgmtNetSwitch-admin</code>	Sets a new password for the <code>admin</code> user on the Cisco Nexus 9348GC-FXP Switch.
<code>SpineSwitch-admin</code>	Sets a new password for the <code>admin</code> user on the spine Cisco Nexus 9336C-FX2 Switches.
<code>mgmt-root</code>	Sets a new password for the <code>root</code> user on the management nodes.
<code>mysql-appfw</code>	Sets a new password for the <code>appfw</code> user in the MySQL database. The <code>mysql-appfw</code> , <code>mysql-ovs</code> , <code>mysql-root</code> and <code>wls-weblogic</code> passwords are synchronized automatically, because these must always be identical.
<code>mysql-ovs</code>	Sets a new password for the <code>ovs</code> user in the MySQL database. The <code>mysql-appfw</code> , <code>mysql-ovs</code> , <code>mysql-root</code> and <code>wls-weblogic</code> passwords are synchronized automatically, because these must always be identical.
<code>mysql-root</code>	Sets a new password for the <code>root</code> user in the MySQL database. The <code>mysql-appfw</code> , <code>mysql-ovs</code> , <code>mysql-root</code> and <code>wls-weblogic</code> passwords are synchronized automatically, because these must always be identical.
<code>ovm-admin</code>	Sets a new password for the <code>admin</code> user in Oracle VM Manager.
<code>spCn-root</code>	Sets a new password for the <code>root</code> user in the compute node ILOMs.
<code>spMn-root</code>	Sets a new password for the <code>root</code> user in the management node ILOMs.
<code>spZfs-root</code>	Sets a new password for the <code>root</code> user on the ZFS storage appliance as well as its ILOM.
<code>system-root</code>	Sets a new password for the <code>root</code> user on all compute nodes.
<code>wls-weblogic</code>	Sets a new password for the <code>weblogic</code> user in WebLogic Server. The <code>mysql-appfw</code> , <code>mysql-ovs</code> , <code>mysql-root</code> and <code>wls-weblogic</code> passwords are synchronized automatically, because these must always be identical.
<code>zfs-root</code>	Sets a new password for the <code>root</code> user on the ZFS storage appliance as well as its ILOM.
<code>--confirm</code>	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.

Option	Description
<code>--force</code>	Force the command to be executed even if the target is in an invalid state. This option is not risk-free and should only be used as a last resort.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.54 Changing the Oracle VM Manager Administrator Password

```
PCA> update password ovm-admin
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Current PCA Password:
New ovm-admin Password:
Confirm New ovm-admin Password:
Status: Success
```

4.2.31 update compute-node

Updates the Oracle Private Cloud Appliance compute nodes to the Oracle VM Server version included in the Oracle Private Cloud Appliance ISO image.

Syntax

```
update compute-node { node } [ --confirm ] [ --force ] [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]
```

where `node` is the identifier of the compute node that must be updated with the Oracle VM Server version provided as part of the appliance software ISO image. Run this command for one compute node at a time.



Warning

Running the `update compute-node` command with multiple `node` arguments is not supported. Neither is running the command concurrently in separate terminal windows.

Description

Use the `update compute-node` command to install the new Oracle VM Server version on the selected compute node or compute nodes. This is a destructive operation and you are prompted to confirm whether or not you wish to continue, unless you use the `--confirm` flag to override the prompt.

Options

The following table shows the available options for this command.

Option	Description
<code>--confirm</code>	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
<code>--force</code>	Force the command to be executed even if the target is in an invalid state. This option is not risk-free and should only be used as a last resort.
<code>--json</code>	Return the output of the command in JSON format
<code>--less</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>less</code> command on the Linux command line. This option allows both forward and backward navigation through the command output.
<code>--more</code>	Return the output of the command one screen at a time for easy viewing, as with the <code>more</code> command on the Linux command line. This option allows forward navigation only.
<code>--tee=OUTPUTFILENAME</code>	When returning the output of the command, also write it to the specified output file.

Examples

Example 4.55 Upgrade a Compute Node to Oracle VM Server Release 4.2.x

```
PCA> update compute-node ovcacn10r1
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y

Status: Success
```

Chapter 5 Managing the Oracle VM Virtual Infrastructure

Table of Contents

5.1 Guidelines and Limitations	146
5.2 Logging in to the Oracle VM Manager Web UI	149
5.3 Monitoring Health and Performance in Oracle VM	149
5.4 Creating and Managing Virtual Machines	150
5.5 Managing Virtual Machine Resources	153
5.6 Configuring Network Resources for Virtual Machines	155
5.6.1 Configuring VM Network Resources on Ethernet-based Systems	155
5.6.2 Configuring VM Network Resources on InfiniBand-based Systems	158
5.7 Viewing and Managing Storage Resources	161
5.7.1 Oracle ZFS Storage Appliance ZS7-2	162
5.7.2 Oracle ZFS Storage Appliance ZS5-ES and Earlier Models	162
5.8 Tagging Resources in Oracle VM Manager	163
5.9 Managing Jobs and Events	163



Warning

Access to the Oracle VM Manager web user interface, command line interface and web services API is provided without restrictions. The configuration of Oracle Private Cloud Appliance components within Oracle VM Manager is automatic and handled by the Oracle Private Cloud Appliance provisioning process. Altering the configuration of these components directly within Oracle VM Manager is not supported and may result in the malfunction of the appliance.

Here is a non-exhaustive list of critical limitations that are known to be violated regularly, which results in severe system configuration problems and significant downtime:

- **DO NOT rename host names** of compute names or other Oracle Private Cloud Appliance components.
- **DO NOT rename server pools.**
- **DO NOT rename built-in repositories.**
- **DO NOT rename existing networks or modify their properties** (VLAN tag, MTU, and so on), except as documented explicitly in the Oracle Private Cloud Appliance Administrator's Guide.
- **DO NOT add the VM role to the internal management network or internal storage network.**



Warning

The appliance controller software enables customization of networking, external storage connectivity and server pools – known as tenant groups in Oracle Private Cloud Appliance. The resulting Oracle VM configurations also must not be altered within Oracle VM Manager.

Use of Oracle VM Manager in the context of Oracle Private Cloud Appliance should be limited to the management and creation of virtual machines.

Configuring additional storage, creating repositories, and setting up additional networks specifically for the use of virtual machines is possible. However, this should be done carefully, to avoid disrupting the configuration specific to the Oracle Private Cloud Appliance.

Management of virtual machines and your Oracle VM environment is achieved using the Oracle VM Manager Web UI (User Interface). While Oracle VM Manager does provide a command line interface and web services API, use of these on your Oracle Private Cloud Appliance should only be attempted by advanced users with a thorough understanding of Oracle VM and the usage limitations within an Oracle Private Cloud Appliance context.

The information provided in here, is a description of the Oracle VM Manager Web UI within the context of the Oracle Private Cloud Appliance. Where particular actions within the Oracle VM Manager Web UI are referenced, a link to the appropriate section within the Oracle VM Manager User's Guide is provided. The complete Oracle VM Manager User's Guide is available at this URL: https://docs.oracle.com/cd/E64076_01/E64082/html/index.html.



Note

When consulting the Oracle VM documentation directly, keep in mind the limitations imposed by using it within Oracle Private Cloud Appliance. More details about the use of the Oracle VM documentation library can be found in [About the Oracle VM Documentation Library](#).

New users of Oracle VM who want to learn the fundamentals of creating and maintaining a virtualized environment should consult the [Oracle VM Concepts Guide](#). It describes the concepts on which the Oracle VM components and functionality are based, and also links to operational procedures in the Oracle VM Manager User's Guide.

The Oracle VM Manager Web UI is available at the virtual IP address that you configured for your management nodes during installation. This virtual IP address is automatically assigned to whichever management node is currently the master or active node within the cluster. If that management node becomes unavailable, the standby management node is promoted to the active role and takes over the IP address automatically. See [Section 1.5, "High Availability"](#) for more information on management node failover.

The Oracle VM Manager Web UI is configured to listen for HTTPS requests on port 7002.

5.1 Guidelines and Limitations

The Oracle VM Manager Web User Interface is provided without any software limitation to its functionality. Once your appliance has been provisioned, the Oracle VM environment is fully configured and ready to use for the deployment and management of your virtual machines. In this section, the operations that are explicitly not permitted, are presented as guidelines and limitations that should be followed when working within Oracle VM Manager, or executing operations programmatically through the command line interface (CLI) or web services API (WSAPI).

The following actions must not be performed, except if Oracle gives specific instructions to do so.

Do Not:

- attempt to discover, remove, rename or otherwise modify servers or their configuration;

- attempt to modify the NTP configuration of a server;
- attempt to add, remove, rename or otherwise modify server pools or their configuration;
- attempt to change the configuration of server pools corresponding with tenant groups configured through the appliance controller software (except for DRS policy setting);
- attempt to move servers out of the existing server pools;
- attempt to add or modify or remove server processor compatibility groups;
- attempt to modify or remove the existing local disk repositories or the repository named **Rack1-repository**;
- attempt to delete or modify any of the preconfigured default networks, or custom networks configured through the appliance controller software;
- attempt to connect virtual machines to the appliance management network;
- attempt to modify or delete any existing Storage elements that are already configured within Oracle VM, or use the reserved names of the default storage elements – for example `OVCA_ZFSSA_Rack1` – for any other configuration;
- attempt to configure global settings, such as YUM Update, in the **Reports and Resources** tab (except for tags, which are safe to edit);
- attempt to select a non-English character set or language for the operating system, because this is not supported by Oracle VM Manager – see support note with [Doc ID 2519818.1](#).

If you ignore this advice, the Oracle Private Cloud Appliance automation, which uses specific naming conventions to label and manage assets, may fail. Out-of-band configuration changes would not be known to the orchestration software of the Oracle Private Cloud Appliance. If a conflict between the Oracle Private Cloud Appliance configuration and Oracle VM configuration occurs, it may not be possible to recover without data loss or system downtime.



Note

An exception to these guidelines applies to the creation of a *Service VM*. This is a VM created specifically to perform administrative operations, for which it needs to be connected to both the public network and internal appliance networks. For detailed information and instructions, refer to the support note with [Doc ID 2017593.1](#).

There is a known issue with the Oracle Private Cloud Appliance Upgrader, which stops the upgrade process if Service VMs are present. For the appropriate workaround, consult the support note with [Doc ID 2510822.1](#).

Regardless of which interface you use to access the Oracle VM functionality directly, the same restrictions apply. In summary, you may use the Web UI, CLI or WSAPI for the operations listed below.

Use the Oracle VM Interfaces for:

- configuration and management of VM networks, VLAN interfaces and VLANs;
- configuration of VM vNICs and connecting VMs to networks;
- all VM configuration and life cycle management;
- attaching and managing external storage for VM usage;

- compute node IPMI control.

About the Oracle VM Documentation Library

You can find the complete Oracle VM documentation library at this URL: https://docs.oracle.com/cd/E64076_01/index.html.

It is critical that you understand the scope of Oracle VM within the specific context of Oracle Private Cloud Appliance. A major objective of the appliance is to orchestrate or fully automate a number of Oracle VM operations. It also imposes restrictions that do not exist in other Oracle VM environments, on infrastructure aspects such as server hardware, networking and storage configuration. Consequently, some chapters or even entire books in the Oracle VM documentation library are irrelevant to Oracle Private Cloud Appliance customers, or should not be used because they describe procedures that conflict with the way the appliance controller software configures and manages the Oracle VM infrastructure.

This list, which is not meant to be exhaustive, explains which parts of the Oracle VM documentation should not be referenced because the functionality in question is either not supported or managed at the level of the appliance controller software:

- Installation and Upgrade Guide

Oracle Private Cloud Appliance always contains a clustered pair of management nodes with Oracle VM Manager pre-installed. When you power on the appliance for the first time, the compute node provisioning process begins, and one of the provisioning steps is to install Oracle VM Server on the compute nodes installed in the appliance rack. The installation of additional compute nodes and upgrades of the appliance software are orchestrated in a similar way.

- Getting Started Guide

Although the getting started guide is an excellent way to progress through the entire chain of operations from discovering the first Oracle VM Server to the point of accessing a fully operational virtual machine, it does not help the Oracle Private Cloud Appliance user, who only needs Oracle VM Manager in order to create and manage virtual machines.

- Administration Guide

This guide describes a number of advanced system administration tasks, most of which are performed at the level of the virtualization platform. The information in this book may be useful for specific configurations or environments, but we recommend that you consult with Oracle subject matter experts to avoid making changes that adversely affect the Oracle Private Cloud Appliance environment.

- Command Line Interface and Web Services API

The recommended interface to manage the Oracle VM environment within Oracle Private Cloud Appliance is the Oracle VM Manager Web UI. The CLI and WSAPI should be used with care, within the limitations described in the Oracle Private Cloud Appliance documentation. They can be safely used in a programmatic context, for example to automate operations related to the virtual machine life cycle (which includes create, clone, start, stop, migrate VMs, pinning CPUs, uploading templates and ISOs, and so on).

Since Oracle VM Manager is the preferred interface to manage the virtualized environment, this chapter provides links to various sections of the Oracle VM Manager User's Guide in order to help Oracle Private Cloud Appliance users perform the necessary tasks. The book is closely aligned with the structure of the Web UI it describes, and the sections and links in this chapter conveniently follow the same basic outline. Where the Oracle VM Manager functionality overlaps with the default Oracle Private Cloud Appliance configuration the document indicates which operations are safe and which should be avoided.

5.2 Logging in to the Oracle VM Manager Web UI

To open the Login page of the Oracle VM Manager Web UI, enter the following address in a Web browser:

`https://manager-vip:7002/ovm/console`

Where, *manager-vip* refers to the virtual IP address, or corresponding host name, that you have configured for your management nodes during installation. By using the virtual IP address, you ensure that you always access the Oracle VM Manager Web UI on the active management node.



Important

You must ensure that if you are accessing Oracle VM Manager through a firewalled connection, the firewall is configured to allow TCP traffic on the port that Oracle VM Manager is using to listen for connections.

Enter your Oracle VM Manager administration user name in the **Username** field. This is the administration user name you configured during installation. Enter the password for the Oracle VM Manager administration user name in the **Password** field.



Important

The Oracle VM Manager Web UI makes use of cookies in order to store session data. Therefore, to successfully log in and use the Oracle VM Manager Web UI your web browser must accept cookies from the Oracle VM Manager host.

5.3 Monitoring Health and Performance in Oracle VM

The **Health** tab provides a view of the health of the compute nodes and the server pool within your environment. This information complements the Hardware View provided in the Oracle Private Cloud Appliance Dashboard. See [Section 2.3, “Hardware View”](#) for more information.

The **Statistics** subtabs available on the Health tab provides statistical information, including graphs that can be refreshed with short intervals or at the click of a button, for CPU and memory usage and for file system utilization. These statistics can be viewed at a global scale to determine overall usage, or at the detail level of a category of resources or even a single item.

The Server and VM Statistics subtab can display information per server to see the performance of each individual compute node, or per virtual machine to help track the usage and resource requirements for any of the virtual machines within your environment. The File System Statistics subtab displays storage space utilization information, organized by storage location, and allows you to track available space for individual file systems over time.

For detailed information on using the Health tab, please refer to the section entitled [Health Tab](#) in the [Oracle VM Manager User's Guide](#).

In addition to the Health tab you can also monitor the status of many resource categories through the **Info perspective** or **Events perspective**. When you select these perspectives in the Management pane, the type of information displayed depends on the active item in the Navigation pane on the left hand side. Both the Info perspective and the Events perspective are common to many elements within the Oracle VM Manager Web UI.

The following sections in the [Oracle VM Manager User's Guide](#) provide detailed information about both perspectives, using the server pool item as an example:

- the Oracle VM Manager [Info perspective](#)
- the Oracle VM Manager [Events perspective](#)

5.4 Creating and Managing Virtual Machines

The **Servers and VMs** tab is used to create and manage your virtual machines. By default, compute nodes in the base rack of the appliance are listed as belonging to a single server pool called **Rack1_ServerPool**. The configuration of the default server pool must not be altered. There is no need to discover servers, as compute nodes are automatically provisioned and discovered within an Oracle Private Cloud Appliance. Editing the configuration of the server pool, servers and processor compatibility groups is not supported. The primary purpose of this tab within the Oracle Private Cloud Appliance context is to create and manage your virtual machines.

Virtual machines can be created using:

- ISO files in a repository (hardware virtualized only)
- Mounted ISO files on an NFS, HTTP or FTP server (paravirtualized only)
- Virtual machine templates (by cloning a template)
- Existing virtual machines (by cloning a virtual machine)
- Virtual machine assemblies or virtual appliances

Virtual machines require most installation resources to be located in the storage repository, managed by Oracle VM Manager, with the exception of mounted ISO files for paravirtualized guests. See [Section 5.5, "Managing Virtual Machine Resources"](#) for more information on importing these resources into the Oracle Private Cloud Appliance repository.

The following list provides an outline of actions that you can perform in this tab, with links to the relevant documentation within the [Oracle VM Manager User's Guide](#):

- Create a virtual machine

You can create a virtual machine following the instructions provided in the section entitled [Create Virtual Machine](#).

You do not need to create any additional server pools. You need only ensure that your installation media has been correctly imported into the Oracle Private Cloud Appliance repository.

- View virtual machine information and events

You can view information about your virtual machine or access virtual machine events by following the information outlined in the section entitled [View Virtual Machine Events](#).

- Edit a virtual machine

You can edit virtual machine parameters as described in the section entitled [Edit Virtual Machine](#).

- Start a virtual machine

Further information is provided in the section entitled [Start Virtual Machines](#).

- Connect to a virtual machine console

There are two options for virtual machine console connections:

- For more information about the use of the VM console, refer to the section entitled [Launch Console](#).
- For more information about the use of the VM serial console, refer to the section entitled [Launch Serial Console](#).
- Stop a virtual machine
Further information is provided in the section entitled [Stop Virtual Machines](#).
- Kill a virtual machine
Further information is provided in the section entitled [Kill Virtual Machines](#).
- Restart a virtual machine
Further information is provided in the section entitled [Restart Virtual Machines](#).
- Suspend a virtual machine
Further information is provided in the section entitled [Suspend Virtual Machines](#).
- Resume a virtual machine
Further information is provided in the section entitled [Resume Virtual Machine](#).
- Migrate or move a virtual machine between repositories, between servers, and to or from the Unassigned Virtual Machines folder

Further information is provided in the section entitled [Migrate or Move Virtual Machines](#).

It is possible to create alternate repositories if you have extended the system with external storage. If you have an additional repository, this function can be used to move a virtual machine from one repository to another.

Because there is only a single server pool available in a default Oracle Private Cloud Appliance base rack, migration of virtual machines can only be achieved between servers and between a server and the Unassigned Virtual Machines folder. Migration between server pools is possible if you have customized the default configuration by creating tenant groups. See [Section 2.7, “Tenant Groups”](#) for more information.

Modifying Server Processor Compatibility Groups is not permitted.



Caution

Compute nodes of different hardware generations operate within the same server pool but belong to different CPU compatibility groups. By default, live migration between CPU compatibility groups is not supported, meaning that virtual machines must be cold-migrated between compute nodes of different generations.

If *live* migration between compute nodes of different generations is required, it must only be attempted from an older to a newer hardware generation, and never in the opposite direction. To achieve this, the administrator must first create new compatibility groups.

For more information about CPU compatibility groups, please refer to the section entitled [Server Processor Compatibility Perspective](#).

For more information about the Unassigned Virtual Machines folder, refer to the section entitled [Unassigned Virtual Machines Folder](#).

- Control virtual machine placement through anti-affinity groups.

You can prevent virtual machines from running on the same physical host by adding them to an anti-affinity group. This is particularly useful for redundancy and load balancing purposes.

Further information about anti-affinity groups is provided in the section entitled [What are Anti-Affinity Groups?](#) in the *Oracle VM Concepts Guide*.

For instructions to create and manage anti-affinity groups, refer to the section entitled [Anti-Affinity Groups Perspective](#) in the *Oracle VM Manager User's Guide*.

- Clone a virtual machine

Further information is provided in the section entitled [Clone a Virtual Machine or Template](#).

You can create a clone customizer to set up the clone parameters, such as networking, and the virtual disk, and ISO resources. For more information about clone customizers, please refer to the section entitled [Manage Clone Customizers](#).

- Export virtual machines to a virtual appliance

Exporting a virtual appliance lets you reuse virtual machines with other instances of Oracle VM, or with other virtualization environments that support the Open Virtualization Format (OVA). You can export one or more virtual machines to a virtual appliance. Further information is provided in the section entitled [Export to Virtual Appliance](#).

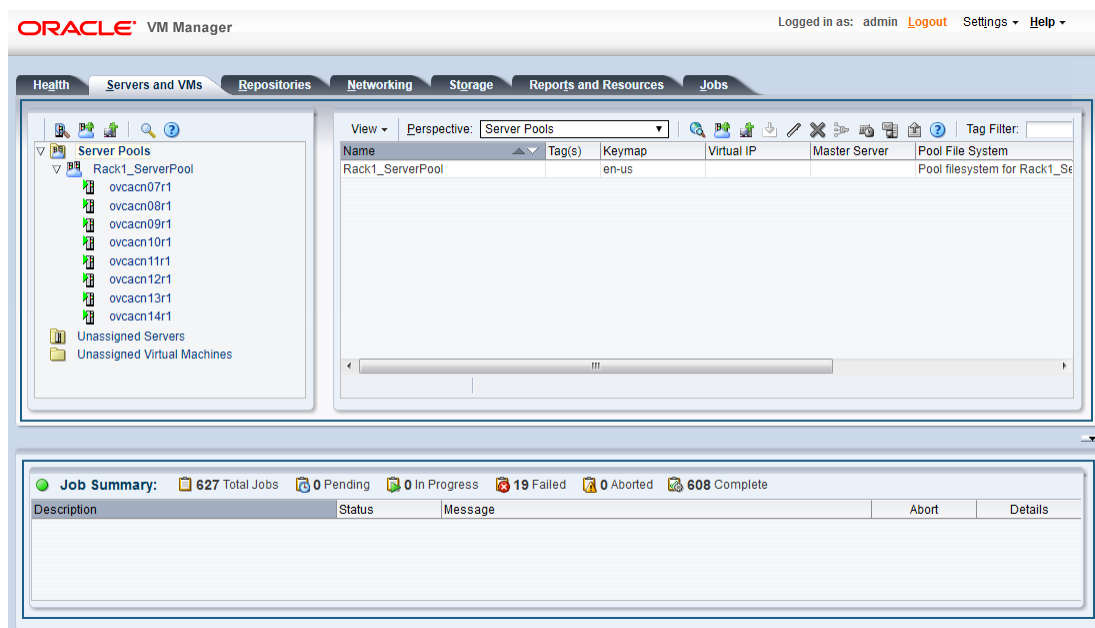
- Send a message to a virtual machine

If you have installed Oracle VM Guest Additions within your virtual machine, you can use the Oracle VM Messaging framework to send messages to your virtual machines to trigger actions within a virtual machine. Refer to the section entitled [Send VM Messages](#) for more information.

- Delete a virtual machine

Further information is provided in the section entitled [Delete Virtual Machines](#).

Figure 5.1 A view of the Servers and VMs tab



5.5 Managing Virtual Machine Resources

The **Repositories** tab provides a view of the Oracle Private Cloud Appliance repository. By default, a shared repository is configured on the internal ZFS storage appliance and named **Rack1-repository**. Additional local repositories are configured using the free disk space of each compute node. None of the default repository configurations may be altered.



Caution

Using local storage on the compute nodes has implications that you should take into account when planning the deployment of your virtual environment. For example:

- Virtual machines with resources in a local storage repository cannot be migrated to another compute node.
- Templates, assemblies and ISOs in local storage repositories cannot be used to create virtual machines on another compute node.
- If a compute node becomes unavailable, its locally stored virtual machines and resources cannot be restored or migrated to another compute node for continued service.
- The virtual machines and resources in local storage repositories are not protected by automatic failover and high-availability mechanisms normally offered by a clustered Oracle VM server pool with shared storage repository.

Additional repositories should be configured using external storage solutions. If the system contains an Oracle ZFS Storage Appliance ZS7-2, extra disk trays can be installed to provide the space for additional repositories. For information about extending the storage capacity of Oracle Private Cloud Appliance, see [Section 5.7, “Viewing and Managing Storage Resources”](#).

The Repositories tab is used to manage virtual machine resources, such as installation media and virtual disks. From this tab, it is possible to create, import or clone Oracle VM templates, virtual appliances and ISO image files. It is also possible to create, modify, or clone virtual disks here. The following list provides an outline of actions that you can perform in this tab, with links to the relevant documentation within the Oracle VM Manager User's Guide:

- Manage Virtual Machine Templates

- Import a template
- Edit a template
- Clone a VM or template
- Move a template
- Manage template clone customizers
- Delete a template

All documentation for these actions can be found in the section entitled [VM Templates Perspective](#).

For specific information about virtual appliances offered through Oracle Technology Network, refer to [Virtual Appliances from Oracle](#).

- Manage Virtual Appliances

- Import a virtual appliance
- Create a VM from a virtual appliance
- Edit a virtual appliance
- Refresh a virtual appliance
- Delete a virtual appliance

All documentation for these actions can be found in the section entitled [Virtual Appliances Perspective](#).

For specific information about virtual appliances offered through Oracle Technology Network, refer to [Virtual Appliances from Oracle](#).

- Manage Virtual Machine ISO Image Files

- Import an ISO
- Edit an ISO
- Clone an ISO
- Delete an ISO

All documentation for these actions can be found in the section entitled [ISOs Perspective](#).

- Manage Virtual Disks

- Create a virtual disk
- Import a virtual disk

- Edit a virtual disk
- Clone a virtual disk
- Delete a virtual disk

All documentation for these actions can be found in the section entitled [Virtual Disks Perspective](#).

- View Virtual Machine Configuration Entries

For more information, refer to the section entitled [VM Files Perspective](#).

Virtual Appliances from Oracle

Through the [Oracle VM product pages](#), you can find several pre-configured Oracle VM Virtual Appliances, which can be downloaded for convenient deployment on Oracle Private Cloud Appliance. These virtual appliances allow users of Oracle Private Cloud Appliance to rapidly set up a typical Oracle product stack within their Oracle VM environment, without having to perform the full installation and configuration process.

For detailed information, including documentation specific to the virtual appliances, refer to the [Oracle VM Virtual Appliances](#) overview page.

For Oracle VM instructions related to virtual appliances, follow the links provided above.

For more general information about the use of virtual appliances and templates, refer to the chapter [Understanding Repositories](#) in the *Oracle VM Concepts Guide*. The most relevant sections are:

- How is a Repository Organized?
- How are Virtual Appliances Managed?

5.6 Configuring Network Resources for Virtual Machines

The **Networking** tab is used to manage networks within the Oracle VM environment running on the Oracle Private Cloud Appliance.



Caution

By default, a number of networks are defined during factory installation. These **must not be altered** as they are required for the correct operation of the Oracle Private Cloud Appliance software layer.

Oracle Private Cloud Appliance exists in two different types of network architecture. One is built around a physical InfiniBand fabric; the other relies on physical high speed Ethernet connectivity. While the two implementations offer practically the same functionality, the configuration of default networks is different due to the type of network hardware. As a result, the procedures to create VLAN networks for virtual machine traffic are different as well.

This section is split up by network architecture to avoid confusion. Refer to the subsection that applies to your appliance.

5.6.1 Configuring VM Network Resources on Ethernet-based Systems

On a system with an Ethernet-based network architecture, default networks are set up as follows:

- [192.168.32.0](#) : the internal management network

This is a private network providing connectivity between the management nodes and compute nodes, using VLAN 3092. It is used for all network traffic inherent to Oracle VM Manager, Oracle VM Server and the Oracle VM Agents.

- [192.168.40.0](#) : the internal storage network

This is a private network used exclusively for traffic to and from the ZFS storage appliance. Both management nodes and compute nodes can reach the internal storage on VLAN 3093. The network also fulfills the heartbeat function for the clustered Oracle VM server pool.

Additionally, two networks are listed with the **VM Network** role:

- [default_external](#)

This default network is the standard choice for virtual machines requiring external network connectivity. It supports both tagged and untagged traffic. For untagged traffic it uses the Oracle VM standard VLAN 1, meaning no additional configuration is required.

If you prefer to use VLANs for your VM networking, configure the additional VLAN interfaces and networks of your choice as follows:



Note

When reprovisioning compute nodes or provisioning newly installed compute nodes, you always need to configure VLANs manually. The VLAN configuration is not applied automatically when the compute node joins an existing server pool.

1. Go to the **Networking** tab and select the **VLAN Interfaces** subtab.

The process for creating VLAN Interfaces is described in detail in the Oracle VM Manager User's Guide in the section entitled [Create VLAN Interfaces](#).

2. Click **Create VLAN Interface**. In the navigation tree of the Create VLAN Interfaces window, select the `vx13040` VxLAN interface of each compute node in the default **Rack1_ServerPool**.
3. In the next step of the wizard, add the VLAN IDs you require. When you complete the wizard, a new VLAN interface for each new VLAN ID is configured on top of each compute node interface you selected.
4. Create a new Oracle VM network with the *VM role*, on the VLAN interfaces for each VLAN tag you created. Each new network should contain the VLAN interfaces associated with a particular VLAN ID; for example all VLAN interfaces with ID 11 on top of a `vx13040` interface.



Tip

You can filter the VLAN interfaces by ID to simplify the selection of the VLAN interfaces participating in the new network.

The process for creating networks with VLAN interfaces is described in the Oracle VM Manager User's Guide in the section entitled [Create New Network](#).



Note

To start using the new network at the VM level, edit the necessary VMs and assign a vNIC to connect to the new network.

5. Configure your data center network accordingly.

- `default_internal`

This default network is intended for virtual machines requiring network connectivity to other virtual machines hosted on the appliance, but *not* external to the appliance. For untagged traffic it uses the Oracle VM standard VLAN 1. To use the VLANs of your choice, configure the additional VLAN interfaces and networks as follows:



Note

When reprovisioning compute nodes or provisioning newly installed compute nodes, you always need to configure VLANs manually. The VLAN configuration is not applied automatically when the compute node joins an existing server pool.

1. Go to the **Networking** tab and select the **VLAN Interfaces** subtab.

The process for creating VLAN Interfaces is described in detail in the Oracle VM Manager User's Guide in the section entitled [Create VLAN Interfaces](#).

2. Click **Create VLAN Interface**. In the navigation tree of the Create VLAN Interfaces window, select the `vx2` VxLAN interface of each compute node in the default **Rack1_ServerPool**.
3. In the next step of the wizard, add the VLAN IDs you require. When you complete the wizard, a new VLAN interface for each new VLAN ID is configured on top of each compute node network port you selected.
4. Create a new VLAN network with the *VM role* for each VLAN tag you added. Each new network should contain the VLAN interfaces associated with a particular VLAN ID; for example all VLAN interfaces with ID 1001 on top of a `vx2` interface.



Tip

You can filter the VLAN interfaces by ID to simplify the selection of the VLAN interfaces participating in the new network.

The process for creating networks with VLAN interfaces is described in the Oracle VM Manager User's Guide in the section entitled [Create New Network](#).

For more information about Oracle Private Cloud Appliance network configuration, see [Section 1.2.4, "Network Infrastructure"](#).



Caution

Do not alter the internal appliance administration network (`192.168.4.0`) connections on the compute nodes or any other rack components. The environment infrastructure depends on the correct operation of this network.

For example, if you configured networking for virtual machines in such a way that they can obtain an IP address in the `192.168.4.0` subnet, IP conflicts and security issues are likely to occur.



Note

If VM-to-VM network performance is not optimal, depending on the type of network load, you could consider increasing the guests' MTU from the default 1500 bytes to

9000. Note that this is a change at the VM level; the compute node interfaces are set to accommodate 9000 bytes already, and must never be modified. Connectivity between VMs and external systems may also benefit from the higher MTU, provided this is supported across the entire network path.

Do not edit or delete any of the networks listed here. Doing so may cause your appliance to malfunction. In an Oracle Private Cloud Appliance context, use the Networking tab to configure and manage Virtual NICs and VLANs for use by your virtual machines.

Figure 5.2 A view of the Networking tab (Ethernet-based Architecture)

The screenshot shows the Oracle VM Manager interface. The top navigation bar includes tabs for Health, Servers and VMs, Repositories, Networking (selected), Storage, Reports and Resources, and Jobs. The main content area is titled 'Networks' and contains a table of network channels. Below the table is a 'Job Summary' section with a table of job status.

Name	ID	Intra-Network Server	Network Channels					Description
			Server Management	Cluster Heartbeat	Live Migrate	Storage	Virtual Machine	
192.168.32.0	10b9cddb10			✓				
192.168.40.0	c0a88c00		✓		✓	✓		
default_external	1039c145f1						✓	
default_internal	10336e5cbd						✓	
nmv_pvl_network	10d3e06a47						✓	
private_vlan1001	10fb124b07						✓	VM private network vlan 1001
private_vlan1002	109b387afa						✓	VM private network vlan 1002
public_vlan0011	105ec61873						✓	VM public network vlan 11
public_vlan0012	1044b1f9bf						✓	VM public network vlan 12
public_vlan0013	1015290b59						✓	VM public network vlan 13
public_vlan0014	1035673979						✓	VM public network vlan 14
public_vlan0015	10012d3835						✓	VM public network vlan 15

Job Summary: 627 Total Jobs 0 Pending 0 In Progress 19 Failed 0 Aborted 608 Complete							
Description	Status	Progress	Message	Timestamp	Duration	Abort	Details
No data to display							

5.6.2 Configuring VM Network Resources on InfiniBand-based Systems

On a system with an InfiniBand-based network architecture, default networks are set up as follows:

- `192.168.140.0` : the management network

This is a private network used exclusively for Oracle VM management traffic. Both management nodes and all compute nodes are connected to this network through their `bond0` interface.

- `192.168.40.0` : the storage network

This is a private IPoIB network used exclusively for traffic to and from the ZFS storage appliance. Both management nodes and both storage controllers are connected to this network through their `bond1` interface.

Additionally, three networks are listed with the **VM Network** role:

- `vm_public_vlan`

This default network is the standard choice for virtual machines requiring external network connectivity. It supports both tagged and untagged traffic. For untagged traffic it uses the Oracle VM standard VLAN 1, meaning no additional configuration is required.

If you prefer to use VLANs for your VM networking, configure the additional VLAN interfaces and networks of your choice as follows:



Note

When reprovisioning compute nodes or provisioning newly installed compute nodes, you always need to configure VLANs manually. The VLAN configuration is not applied automatically when the compute node joins an existing server pool.

1. Go to the **Networking** tab and select the **VLAN Interfaces** subtab.

The process for creating VLAN Interfaces is described in detail in the Oracle VM Manager User's Guide in the section entitled [Create VLAN Interfaces](#).

2. Click **Create VLAN Interface**. In the navigation tree of the Create VLAN Interfaces window, select the `bond4` port of each compute node in the default **Rack1_ServerPool**.
3. In the next step of the wizard, add the VLAN IDs you require. When you complete the wizard, a new VLAN interface for each new VLAN ID is configured on top of each compute node network port you selected.
4. Create a new VLAN network with the *VM role* for each VLAN tag you added. Each new network should contain the VLAN interfaces associated with a particular VLAN ID; for example all VLAN interfaces with ID 11 on top of a `bond4` port.



Tip

You can filter the VLAN interfaces by ID to simplify the selection of the VLAN interfaces participating in the new network.

The process for creating networks with VLAN interfaces is described in the Oracle VM Manager User's Guide in the section entitled [Create New Network](#).

5. Configure your data center network accordingly.

For details, see [Section 7.3, "Configuring Data Center Switches for VLAN Traffic"](#).

- `vm_private`

This default network is intended for virtual machines requiring network connectivity to other virtual machines hosted on the appliance, but *not* external to the appliance. For untagged traffic it uses the

Oracle VM standard VLAN 1. To use the VLANs of your choice, configure the additional VLAN interfaces and networks as follows:



Note

When reprovisioning compute nodes or provisioning newly installed compute nodes, you always need to configure VLANs manually. The VLAN configuration is not applied automatically when the compute node joins an existing server pool.

1. Go to the **Networking** tab and select the **VLAN Interfaces** subtab.

The process for creating VLAN Interfaces is described in detail in the Oracle VM Manager User's Guide in the section entitled [Create VLAN Interfaces](#).

2. Click **Create VLAN Interface**. In the navigation tree of the Create VLAN Interfaces window, select the `bond3` port of each compute node in the default **Rack1_ServerPool**.
3. In the next step of the wizard, add the VLAN IDs you require. When you complete the wizard, a new VLAN interface for each new VLAN ID is configured on top of each compute node network port you selected.
4. Create a new VLAN network with the *VM role* for each VLAN tag you added. Each new network should contain the VLAN interfaces associated with a particular VLAN ID; for example all VLAN interfaces with ID 1001 on top of a `bond3` port.



Tip

You can filter the VLAN interfaces by ID to simplify the selection of the VLAN interfaces participating in the new network.

The process for creating networks with VLAN interfaces is described in the Oracle VM Manager User's Guide in the section entitled [Create New Network](#).

- `mgmt_public_eth`

This network is automatically created during the initial configuration of the appliance. It uses the public network that you configured in the Oracle Private Cloud Appliance Dashboard. The primary function of this network is to provide access to the management nodes from the data center network, and enable the management nodes to run a number of system services. As long as you have not configured this network with a VLAN tag, it may also be used to provide external untagged network access to virtual machines. The subnet associated with this network is the same as your data center network.



Caution

Always use the `vm_public_vlan` network as your first VM network option. The `mgmt_public_eth` is unavailable for VM networking when configured with a management VLAN. When no management VLAN is configured, it is restricted to untagged VM traffic, and should only be considered if the circumstances require it.

For more information about Oracle Private Cloud Appliance network configuration, see [Section 1.2.4, "Network Infrastructure"](#).

**Caution**

Do not alter the internal appliance management network ([192.168.4.0](#)) connections on the compute nodes or any other rack components. The environment infrastructure depends on the correct operation of this network.

For example, if you configured networking for virtual machines in such a way that they can obtain an IP address in the [192.168.4.0](#) subnet, IP conflicts and security issues are likely to occur.

**Note**

If VM-to-VM network performance is not optimal, depending on the type of network load, you could consider increasing the guests' MTU from the default 1500 bytes to 9000. Note that this is a change at the VM level; the compute node interfaces are set to 9000 bytes already, and must never be modified. Connectivity between VMs and external systems may also benefit from the higher MTU, provided this is supported across the entire network path.

Do not edit or delete any of the networks listed here. Doing so may cause your appliance to malfunction. In an Oracle Private Cloud Appliance context, use the Networking tab to configure and manage Virtual NICs and VLANs for use by your virtual machines.

Figure 5.3 A view of the Networking tab (InfiniBand-based Architecture)

The screenshot shows the Oracle VM Manager interface with the Networking tab selected. The main table lists network configurations with columns for Name, ID, Intra-Network Server, Network Channels (Server Management, Cluster Heartbeat, Live Migrate, Storage, Virtual Machine), and Description. Below the table is a Job Summary section showing 627 Total Jobs, 0 Pending, 0 In Progress, 19 Failed, 0 Aborted, and 608 Complete.

Name	ID	Intra-Network Server	Network Channels					Description
			Server Management	Cluster Heartbeat	Live Migrate	Storage	Virtual Machine	
192.168.40.0	10b9cdbb10			✓				
192.168.140.0	c0a88c00		✓		✓	✓		
mgmt_public_eth	10d3e06a47						✓	
private_vlan1001	10fb124b07						✓	VM private network vlan 1001
private_vlan1002	109b387afa						✓	VM private network vlan 1002
public_vlan0011	105ec61873						✓	VM public network vlan 11
public_vlan0012	1044b1f9bf						✓	VM public network vlan 12
public_vlan0013	1015290b59						✓	VM public network vlan 13
public_vlan0014	1035673979						✓	VM public network vlan 14
public_vlan0015	10012d3835						✓	VM public network vlan 15
vm_private	10ec9a04e9						✓	
vm_public_vlan	10f62ad232						✓	

Job Summary: 627 Total Jobs, 0 Pending, 0 In Progress, 19 Failed, 0 Aborted, 608 Complete

Description	Status	Progress	Message	Timestamp	Duration	Abort	Details
No data to display							

5.7 Viewing and Managing Storage Resources

The storage resources underlying the built-in Oracle Private Cloud Appliance ZFS storage repository and the server pool clustering file system are listed under the **Storage** tab within Oracle VM Manager. The internal ZFS storage is listed under the SAN Servers folder. Do not modify or attempt to delete this storage.



Warning

Compute node provisioning relies on the internal ZFS file server and its exported storage. Changing the configuration will cause issues with provisioning and server pool clustering.

There are functional differences between the Oracle ZFS Storage Appliance ZS7-2, which is part of systems with an Ethernet-based network architecture, and the previous models of the ZFS Storage Appliance, which are part of systems with an InfiniBand-based network architecture. For clarity, this section describes the Oracle VM storage resources separately for the different storage appliances.

5.7.1 Oracle ZFS Storage Appliance ZS7-2

The internal ZFS Storage Appliance has sufficient disk space (100TB) for a basic virtualized environment, but the storage capacity for virtual disks and shared file systems can be extended with additional external storage for use within Oracle VM.

Information on expanding your Oracle VM environment with storage repositories located on the external storage is provided in the *Oracle VM Manager User's Guide*. Refer to the section entitled [Storage Tab](#). You are also fully capable of using other networked storage, available on the public network or a custom network, within your own Virtual Machines.

Figure 5.4 A view of the Storage tab (with Oracle ZFS Storage Appliance ZS7-2)

Name	Event Severity	Size (GiB)	Server	Status	Shareable	Description	VM(s)
SUN (1)	Informational	12.0	ovcacn07r1, ovcac...	online	No		
SUN (2)	Informational	3072.0	ovcacn07r1, ovcac...	online	No		

Job Summary: 627 Total Jobs, 0 Pending, 0 In Progress, 19 Failed, 0 Aborted, 608 Complete

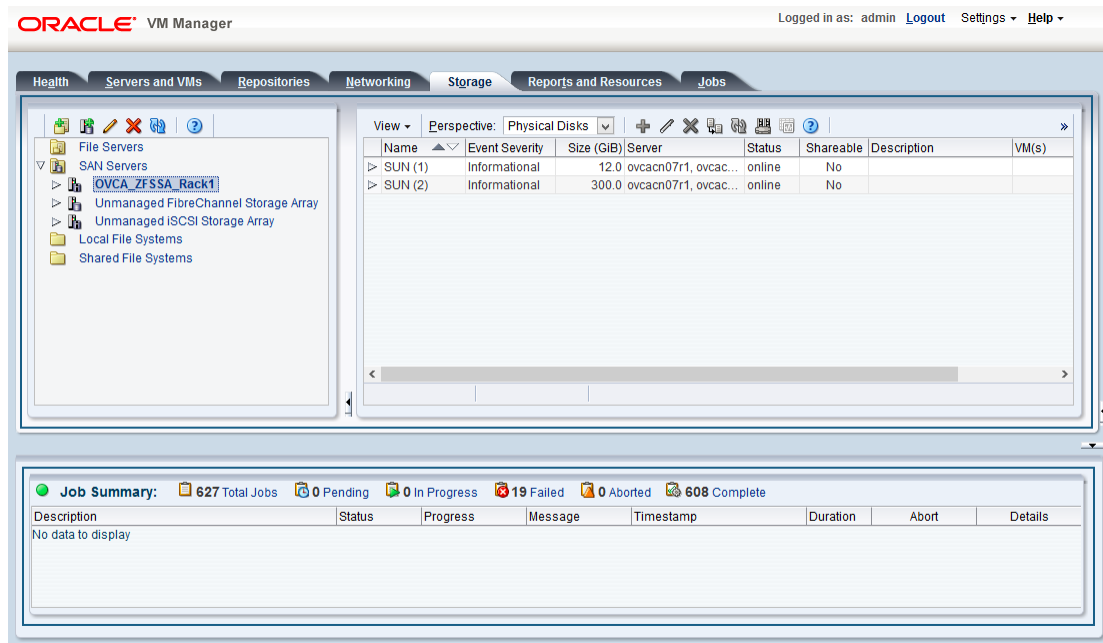
Description	Status	Progress	Message	Timestamp	Duration	Abort	Details
No data to display							

5.7.2 Oracle ZFS Storage Appliance ZS5-ES and Earlier Models

While the storage repository on the internal Oracle ZFS Storage Appliance ZS5-ES and earlier models can be used for a basic virtualized environment and for test purposes, the preferred approach for production environments is to attach additional external storage for use within Oracle VM. The options to extend the storage capacity of an Oracle Private Cloud Appliance are explained in detail in the *Oracle Private Cloud Appliance Installation Guide*: refer to the chapter entitled [Extending Oracle Private Cloud Appliance - Additional Storage](#).

Information on expanding your Oracle VM environment with storage repositories located on the external Fibre Channel or InfiniBand storage is provided in the *Oracle VM Manager User's Guide*. Refer to the section entitled [Storage Tab](#). You are also fully capable of using other networked storage, available on the public network or a custom network, within your own Virtual Machines.

Figure 5.5 A view of the Storage tab (with Oracle ZFS Storage Appliance ZS5-ES)



5.8 Tagging Resources in Oracle VM Manager

The **Reports and Resources** tab is used to configure global settings for Oracle VM and to manage tags, which can be used to identify and group resources. Since many of the global settings such as server update management and NTP configuration are managed automatically within Oracle Private Cloud Appliance, you do not need to edit any settings here. Those configuration changes could cause the appliance to malfunction.

You are able to create, edit and delete tags, by following the instructions in the section entitled [Tags](#).

You can also use this tab to generate XML reports about Oracle VM objects and attributes. For details, refer to the section entitled [Reports](#).

5.9 Managing Jobs and Events

The **Jobs** tab provides a view of the job history within Oracle VM Manager. It is used to track and audit jobs and to help troubleshoot issues within the Oracle VM environment. Jobs and events are described in detail within the Oracle VM Manager User's Guide in the section entitled [Jobs Tab](#).

Since the Recurring Jobs, described in the Oracle VM Manager User's Guide, are all automated and handled directly by the Oracle Private Cloud Appliance, you must not edit any of the settings for recurring jobs.

Chapter 6 Servicing Oracle Private Cloud Appliance Components

Table of Contents

6.1 Oracle Auto Service Request (ASR)	166
6.1.1 Understanding Oracle Auto Service Request (ASR)	166
6.1.2 ASR Prerequisites	167
6.1.3 Setting Up ASR and Activating ASR Assets	167
6.2 Replaceable Components	168
6.2.1 Rack Components	168
6.2.2 Oracle Server X8-2 Components	169
6.2.3 Oracle Server X7-2 Components	170
6.2.4 Oracle Server X6-2 Components	171
6.2.5 Oracle Server X5-2 Components	172
6.2.6 Sun Server X4-2 Components	173
6.2.7 Sun Server X3-2 Components	173
6.2.8 Oracle ZFS Storage Appliance ZS7-2 Components	174
6.2.9 Oracle ZFS Storage Appliance ZS5-ES Components	175
6.2.10 Oracle ZFS Storage Appliance ZS3-ES Components	176
6.2.11 Sun ZFS Storage Appliance 7320 Components	178
6.2.12 Oracle Switch ES1-24 Components	179
6.2.13 NM2-36P Sun Datacenter InfiniBand Expansion Switch Components	179
6.2.14 Oracle Fabric Interconnect F1-15 Components	180
6.3 Preparing Oracle Private Cloud Appliance for Service	181
6.4 Servicing the Oracle Private Cloud Appliance Rack System	182
6.4.1 Powering Down Oracle Private Cloud Appliance (When Required)	182
6.4.2 Service Procedures for Rack System Components	183
6.5 Servicing an Oracle Server X8-2	184
6.5.1 Powering Down Oracle Server X8-2 for Service (When Required)	184
6.5.2 Service Procedures for Oracle Server X8-2 Components	185
6.6 Servicing an Oracle Server X7-2	186
6.6.1 Powering Down Oracle Server X7-2 for Service (When Required)	186
6.6.2 Service Procedures for Oracle Server X7-2 Components	188
6.7 Servicing an Oracle Server X6-2	188
6.7.1 Powering Down Oracle Server X6-2 for Service (When Required)	189
6.7.2 Service Procedures for Oracle Server X6-2 Components	190
6.8 Servicing an Oracle Server X5-2	191
6.8.1 Powering Down Oracle Server X5-2 for Service (When Required)	191
6.8.2 Service Procedures for Oracle Server X5-2 Components	192
6.9 Servicing a Sun Server X4-2	193
6.9.1 Powering Down Sun Server X4-2 for Service (When Required)	193
6.9.2 Service Procedures for Sun Server X4-2 Components	195
6.10 Servicing a Sun Server X3-2	196
6.10.1 Powering Down Sun Server X3-2 for Service (When Required)	196
6.10.2 Service Procedures for Sun Server X3-2 Components	197
6.11 Servicing the Oracle ZFS Storage Appliance ZS7-2	198
6.11.1 Powering Down the Oracle ZFS Storage Appliance ZS7-2 for Service (When Required)	198
6.11.2 Service Procedures for Oracle ZFS Storage Appliance ZS7-2 Components	200
6.12 Servicing the Oracle ZFS Storage Appliance ZS5-ES	201
6.12.1 Powering Down the Oracle ZFS Storage Appliance ZS5-ES for Service (When Required) .	201
6.12.2 Service Procedures for Oracle ZFS Storage Appliance ZS5-ES Components	202
6.13 Servicing the Oracle ZFS Storage Appliance ZS3-ES	203

6.13.1	Powering Down the Oracle ZFS Storage Appliance ZS3-ES for Service (When Required)	204
6.13.2	Service Procedures for Oracle ZFS Storage Appliance ZS3-ES Components	206
6.14	Servicing the Sun ZFS Storage Appliance 7320	207
6.14.1	Powering Down the Sun ZFS Storage Appliance 7320 for Service (When Required)	207
6.14.2	Service Procedures for Sun ZFS Storage Appliance 7320 Components	208
6.15	Servicing an Oracle Switch ES1-24	209
6.15.1	Powering Down the Oracle Switch ES1-24 for Service (When Required)	209
6.15.2	Service Procedures for Oracle Switch ES1-24 Components	210
6.16	Servicing an NM2-36P Sun Datacenter InfiniBand Expansion Switch	210
6.16.1	Powering Down the NM2-36P Sun Datacenter InfiniBand Expansion Switch for Service (When Required)	210
6.16.2	Service Procedures for NM2-36P Sun Datacenter InfiniBand Expansion Switch Components	211
6.17	Servicing an Oracle Fabric Interconnect F1-15	211
6.17.1	Powering Down the Oracle Fabric Interconnect F1-15 for Service (When Required)	211
6.17.2	Service Procedures for Oracle Fabric Interconnect F1-15 Components	212

This chapter explains the service procedures for Oracle Private Cloud Appliance in case a failure occurs. Optionally, you can configure the system with Oracle Auto Service Request (ASR), which generates a service request with Oracle automatically when it detects a hardware malfunction. Certain components of Oracle Private Cloud Appliance are customer-replaceable. These are listed in this chapter, along with the necessary instructions.

6.1 Oracle Auto Service Request (ASR)

Oracle Private Cloud Appliance is qualified for Oracle Auto Service Request (ASR), a software feature for support purposes. It is integrated with My Oracle Support and helps resolve problems faster by automatically opening service requests when specific hardware failures occur. Using ASR is optional: the components must be downloaded, installed and configured in order to enable ASR for your appliance.



Caution

Oracle Auto Service Request (ASR) must be installed by an **authorized Oracle Field Engineer**. Request installation of ASR at the time of system install. Installation at a later date will be a Time and Materials charge.

Oracle is continuously analyzing and improving the ASR fault rules to enhance the Oracle support experience. This includes adding, modifying and removing rules to focus on actionable events from ASR assets while filtering non-actionable events. For up-to-date fault coverage details, please refer to the [Oracle Auto Service Request documentation page](#).

6.1.1 Understanding Oracle Auto Service Request (ASR)

To enable the automated service request feature, the Oracle Private Cloud Appliance components must be configured to send hardware fault telemetry to the ASR Manager software. ASR Manager must be installed on the master management node, which needs an active outbound Internet connection using HTTPS or an HTTPS proxy.

When a hardware problem is detected, ASR Manager submits a service request to Oracle Support Services. In many cases, Oracle Support Services can begin work on resolving the issue before the administrator is even aware the problem exists.

ASR detects faults in the most common hardware components, such as disks, fans, and power supplies, and automatically opens a service request when a fault occurs. ASR does not detect all possible hardware faults, and it is not a replacement for other monitoring mechanisms, such as SMTP and SNMP alerts,

within the customer data center. It is a complementary mechanism that expedites and simplifies the delivery of replacement hardware. ASR should not be used for downtime events in high-priority systems. For high-priority events, contact Oracle Support Services directly.

An email message is sent to both the My Oracle Support email account and the technical contact for Oracle Private Cloud Appliance to notify them of the creation of the service request. A service request may not be filed automatically on some occasions. This can happen because of the unreliable nature of the SNMP protocol or a loss of connectivity to ASR Manager. Oracle recommends that customers continue to monitor their systems for faults and call Oracle Support Services if they do not receive notice that a service request has been filed automatically.

For more information about ASR, consult the following resources:

- Oracle Auto Service Request web page: <https://www.oracle.com/support/premier/auto-service-request.html>.
- Oracle Auto Service Request user documentation: http://docs.oracle.com/cd/E37710_01/index.htm.

6.1.2 ASR Prerequisites

Before you install ASR, make sure that the prerequisites in this section are met.

- Make sure that you have a valid My Oracle Support account.
If necessary, create an account at <https://support.oracle.com>.
- Ensure that the following are set up correctly in My Oracle Support:
 - technical contact person at the customer site who is responsible for Oracle Private Cloud Appliance
 - valid shipping address at the customer site where the Oracle Private Cloud Appliance is located, so that parts are delivered to the site where they must be installed
- Make sure that Oracle Java - JDK 7 (1.7.0_13 or later) or Oracle Java 8 (1.8.0_25 or later) is installed on both management nodes in your Oracle Private Cloud Appliance. Check the version installed on the system by entering the following command at the Oracle Linux prompt: `java -version`.

If the installed version does not comply with the ASR prerequisites, download a compatible Java version, unpack the archive in `/opt/` and install it on both management nodes.



Note

OpenJDK is not supported by ASR.

If necessary, you can download the latest version from the Java SE Downloads page: <http://www.oracle.com/technetwork/java/javase/downloads/>.

- Verify connectivity to the Internet using HTTPS.

For example, try `curl` to test whether you can access <https://support.oracle.com>.

6.1.3 Setting Up ASR and Activating ASR Assets

The necessary packages for ASR Manager must first be downloaded and stored in an installation directory that is accessible from both management nodes. For ASR Manager to work on Oracle Private Cloud Appliance, it must be installed on both management nodes, and failover must be configured so that the ASR Manager role is always fulfilled by the management node that also has the master role.

ASR Manager (ASRM) can be registered as a stand-alone ASRM, pointing directly to My Oracle Support, or as a relay to another ASRM in your network. Even if other systems at your site already use an ASRM, you can choose to register the Oracle Private Cloud Appliance ASRM as stand-alone. This means it communicates directly with the Oracle backend systems, which is the standard registration method.

The Oracle Private Cloud Appliance components that are qualified as ASR assets are the compute nodes and the ZFS Storage Appliance. The two management nodes must not be activated.

The ASR activation mechanism for compute nodes requires operations in two separate locations. First the compute node ILOMs are configured to send SNMP traps to the ASR Manager when a failure occurs. Then the ASR Manager is configured to recognize the ILOMs as assets and accept their input.

The ZFS Storage Appliance runs its own ASR Manager, and relays its ASR data to the Oracle backend systems through the outbound connection of the master management node. To achieve this, Oracle Private Cloud Appliance relies on the `tinyproxy` HTTP and HTTPS proxy daemon. ASR requires `tinyproxy` version 1.8.3 or later to be installed and properly configured on both management nodes.

Detailed installation and configuration instructions are available from My Oracle Support. Refer to the support note with [Doc ID 2560988.1](#).

6.2 Replaceable Components

According to Oracle's Component Replacement Policy, the replaceable components in your system are designated as either field-replaceable units (FRUs) or customer-replaceable units (CRUs).

- A part designated as a FRU must be replaced by an Oracle-qualified service technician.
- A part designated as a CRU can be replaced by a person who is not an Oracle-qualified service technician.

All CRUs and FRUs are listed and identified in this chapter, but the servicing instructions included in this Oracle Private Cloud Appliance Administrator's Guide are focused primarily on CRUs. For FRU replacement, please contact Oracle.

6.2.1 Rack Components

The following table lists the replaceable components of the Oracle Private Cloud Appliance rack.



Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle Private Cloud Appliance components list in the [Oracle System Handbook](#).

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.1 Replaceable Oracle Private Cloud Appliance Rack Components

Component Description	FRU/CRU	Hot-Swap
<i>Oracle Rack Cabinet 1242:</i>		
Jumper Cable C13-C14, 2m	FRU	Yes
Ethernet Cable, Category 5/5E, 10ft, Black	FRU	Yes

Component Description	FRU/CRU	Hot-Swap
Ethernet Cable, Category 5/5E, 10ft, Blue	FRU	Yes
Ethernet Cable, Shielded, Category 5E, 1m, Grey	FRU	Yes
Ethernet Cable, Category 5, 8ft, Black	FRU	Yes
Ethernet Cable, Category 5, 8ft, Green	FRU	Yes
Ethernet Cable, Category 5, 8ft, Yellow	FRU	Yes
Active Optical Cable, Blue, 3m	FRU	Yes
10Gbps QSFP to QSFP Cable, Passive Copper, 3m	FRU	Yes
QSFP28 Cable, 30AWG, Passive Copper, 3m	FRU	Yes
QSFP28 Cable, 30AWG, Passive Copper, 1m	FRU	Yes
1U/2U Screw-Mount Slide Rail Kit	FRU	
1U/2U Cable Management Arm (Snap-in)	FRU	
<i>Power Distribution Units (PDUs):</i>		
15KVA Single-Phase PDU, North America	FRU	Yes
15KVA Three-Phase PDU, North America	FRU	Yes
15KVA Three-Phase PDU, International	FRU	Yes
22KVA Single-Phase PDU, North America	FRU	Yes
22KVA Single-Phase PDU, International	FRU	Yes
24KVA Three-Phase PDU, North America	FRU	Yes
24KVA Three-Phase PDU, International	FRU	Yes

For rack-level component servicing instructions, see [Section 6.4, “Servicing the Oracle Private Cloud Appliance Rack System”](#).

6.2.2 Oracle Server X8-2 Components

The following table lists the replaceable components of the Oracle Server X8-2 compute nodes.



Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle Private Cloud Appliance components list in the [Oracle System Handbook](#).

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.2 Replaceable Oracle Server X8-2 Components

Component Description	FRU/CRU	Hot-Swap
Motherboard Assembly	FRU	No
Quad Counter Rotating Fan Module	CRU	Yes
1-Slot PCI Express Riser Assembly	FRU	No

Component Description	FRU/CRU	Hot-Swap
2-Slot PCI Express Riser Assembly	FRU	No
Type A266 800/1200 Watt AC Input Power Supply	CRU	Yes
Sixteen-core Intel Xeon G-5218 processor (2.3 GHz), 125W	FRU	No
Twenty-four-core Intel Xeon P-8260 processor (2.4 GHz), 165W	FRU	No
CPU Heatsink	FRU	No
2.5" Disk Cage Front Indicator Module	FRU	No
8-Slot 2.5" Disk Backplane Assembly	FRU	No
1.2TB - 10000 RPM SAS-3 Disk Assembly with 1 bracket	CRU	Yes
DDR4 DIMM, 32GB	FRU	No
DDR4 DIMM, 64GB	FRU	No
Dual port 100Gbps Ethernet PCI Express 3.0 Host Channel Adapter (CX-5) (only appliance with Ethernet-based network architecture)	FRU	No
Dual port 80Gbps InfiniBand QDR PCI Express 3.0 Host Channel Adapter M3 (CX-3) (only appliance with InfiniBand-based network architecture)	FRU	No
Dual port 32Gbps Fibre Channel PCI Express 3.0 Host Bus Adapter (optional component)	FRU	No
8-Port 12Gbps SAS-3 RAID PCI Express HBA	FRU	No
System Battery	CRU	No
Cable Kit	FRU	No

For Oracle Server X8-2 component servicing instructions, see [Section 6.5, "Servicing an Oracle Server X8-2"](#).

6.2.3 Oracle Server X7-2 Components

The following table lists the replaceable components of the Oracle Server X7-2 compute nodes.



Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle Private Cloud Appliance components list in the [Oracle System Handbook](#).

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.3 Replaceable Oracle Server X7-2 Components

Component Description	FRU/CRU	Hot-Swap
Motherboard Assembly	FRU	No
Dual Counter Rotating Fan Module	CRU	Yes
1-Slot PCI Express Riser Assembly	FRU	No
2-Slot PCI Express Riser Assembly	FRU	No

Component Description	FRU/CRU	Hot-Swap
A266 1200 Watt AC Input Power Supply	CRU	Yes
Twenty-four-core Intel Xeon P-8160 processor (2.1 GHz), 150W	FRU	No
CPU Heatsink	FRU	No
2.5" Disk Cage Front Indicator Module	FRU	No
4-Slot 2.5" Disk Backplane Assembly	FRU	No
1.2TB - 10000 RPM SAS Disk Assembly with 1 bracket	CRU	Yes
DDR4 DIMM	FRU	No
Dual port 80Gbps InfiniBand QDR PCI Express 3.0 Host Channel Adapter M3 (CX-3)	FRU	No
8-Port 12Gbps SAS-3 RAID PCI Express HBA	FRU	No
System Battery	CRU	No
Cable Kit	FRU	No

For Oracle Server X7-2 component servicing instructions, see [Section 6.6, "Servicing an Oracle Server X7-2"](#).

6.2.4 Oracle Server X6-2 Components

The following table lists the replaceable components of the Oracle Server X6-2 compute nodes.



Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle Private Cloud Appliance components list in the [Oracle System Handbook](#).

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.4 Replaceable Oracle Server X6-2 Components

Component Description	FRU/CRU	Hot-Swap
System Board Assembly	FRU	No
Dual Counter Rotating Fan Module	CRU	Yes
1-Slot PCI Express Riser Assembly	FRU	No
2-Slot PCI Express Riser Assembly	FRU	No
A256 600 Watt AC Input Power Supply	CRU	Yes
Twenty-two-core Intel Xeon processor E5-2699 v4 series (2.2 GHz), 145W	FRU	No
Pre-Greased CPU Heatsink	FRU	No
2.5" Disk Cage Front Indicator Module	FRU	No
4-Slot 2.5" Disk Backplane Assembly	FRU	No
1.2TB - 10000 RPM SAS Disk Assembly with 1 bracket	CRU	Yes

Component Description	FRU/CRU	Hot-Swap
32GB DDR4-2400 Load Reduced DIMM	FRU	No
Dual port 80Gbps InfiniBand QDR PCI Express 3.0 Host Channel Adapter M3 (CX-3)	FRU	No
8GB USB 2.0 Flash Drive	FRU	No
8-Port 12Gbps SAS-3 RAID PCI Express HBA	FRU	No
1U/2U Remote Battery Assembly	CRU	No
Cable Kit	FRU	No

For Oracle Server X6-2 component servicing instructions, see [Section 6.7, “Servicing an Oracle Server X6-2”](#).

6.2.5 Oracle Server X5-2 Components

The following table lists the replaceable components of the Oracle Server X5-2 management and compute nodes.



Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle Private Cloud Appliance components list in the [Oracle System Handbook](#).

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.5 Replaceable Oracle Server X5-2 Components

Component Description	FRU/CRU	Hot-Swap
System Board Assembly	FRU	No
Dual Counter Rotating Fan Module	CRU	Yes
1-Slot PCI Express Riser Assembly	FRU	No
2-Slot PCI Express Riser Assembly	FRU	No
A256 600 Watt AC Input Power Supply	CRU	Yes
Eighteen-core Intel Xeon processor E5-2699 v3 series (2.3 GHz), 145W	FRU	No
Pre-Greased CPU Heatsink	FRU	No
2.5" Disk Cage Front Indicator Module	FRU	No
4-Slot 2.5" Disk Backplane Assembly	FRU	No
1.2TB - 10000 RPM SAS Disk Assembly with 1 bracket	CRU	Yes
32GB DDR4-2133 Load Reduced DIMM	FRU	No
Dual port 80Gbps InfiniBand QDR PCI Express 3.0 Host Channel Adapter M3 (CX-3)	FRU	No
8GB USB 2.0 Flash Drive	FRU	No
8-Port 12Gbps SAS-3 RAID PCI Express HBA	FRU	No

Component Description	FRU/CRU	Hot-Swap
1U/2U Remote Battery Assembly	CRU	No
Cable Kit	FRU	No

For Oracle Server X5-2 component servicing instructions, see [Section 6.8, “Servicing an Oracle Server X5-2”](#).

6.2.6 Sun Server X4-2 Components

The following table lists the replaceable components of the Sun Server X4-2 management and compute nodes.



Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle Private Cloud Appliance components list in the [Oracle System Handbook](#).

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.6 Replaceable Sun Server X4-2 Components

Component Description	FRU/CRU	Hot-Swap
System Board Assembly	FRU	No
Dual Counter Rotating Fan Module	CRU	Yes
1-Slot PCI Express Riser Assembly	FRU	No
2-Slot PCI Express Riser Assembly	FRU	No
A256 600 Watt AC Input Power Supply	CRU	Yes
2.6GHz Intel 8-core Xeon E5-2650, 95W	FRU	No
Pre-Greased CPU Heatsink	FRU	No
2.5" Disk Cage Front Indicator Module	FRU	No
4-Slot 2.5" Disk Backplane Assembly	FRU	No
1.2TB - 10000 RPM SAS Disk Assembly with 1 bracket	CRU	Yes
16GB DDR3-1600 DIMM, 1.35V	FRU	No
Dual port 80Gbps InfiniBand QDR PCI Express 3.0 Host Channel Adapter M3 (CX-3)	FRU	No
4GB USB 2.0 Flash Drive	FRU	No
8-Port 6Gbps SAS-2 RAID PCI Express HBA, B4 ASIC	FRU	No
1U/2U Remote Battery Assembly	CRU	No
Cable Kit	FRU	No

For Sun Server X4-2 component servicing instructions, see [Section 6.9, “Servicing a Sun Server X4-2”](#).

6.2.7 Sun Server X3-2 Components

The following table lists the replaceable components of the Sun Server X3-2 management and compute nodes.



Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle Private Cloud Appliance components list in the [Oracle System Handbook](#).

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.7 Replaceable Sun Server X3-2 Components

Component Description	FRU/CRU	Hot-Swap
System Board Assembly	FRU	No
Dual Counter Rotating Fan Module	CRU	Yes
1-Slot PCI Express Riser Assembly	FRU	No
2-Slot PCI Express Riser Assembly	FRU	No
A256 600 Watt AC Input Power Supply	CRU	Yes
2.2GHz Intel 8-core Xeon E5-2660, 95W	FRU	No
Pre-Greased CPU Heatsink	FRU	No
2.5" Disk Cage Front Indicator Module	FRU	No
4-Slot 2.5" Disk Backplane Assembly	FRU	No
900GB - 10000 RPM SAS Disk Assembly with 1 bracket	CRU	Yes
16GB DDR3-1600 DIMM, 1.35V	FRU	No
Dual 40Gbps InfiniBand 4x QDR PCI Express Low Profile Host Channel Adapter	FRU	No
4GB USB 2.0 Flash Drive	FRU	No
8-Port 6Gbps SAS-2 RAID PCI Express HBA, B4 ASIC	FRU	No
1U/2U Remote Battery Assembly	CRU	No

For Sun Server X3-2 component servicing instructions, see [Section 6.10, "Servicing a Sun Server X3-2"](#).

6.2.8 Oracle ZFS Storage Appliance ZS7-2 Components

The following table lists the replaceable components of the Oracle ZFS Storage Appliance ZS7-2.



Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle Private Cloud Appliance components list in the [Oracle System Handbook](#).

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.8 Replaceable Oracle ZFS Storage Appliance ZS7-2 Components

Component Description	FRU/CRU	Hot-Swap
<i>Oracle ZFS Storage Appliance ZS7-2 Storage Head:</i>		
2.3GHz Intel 18-Core Xeon G-6140, 140W	FRU	No
Pre-greased CPU Heatsink	FRU	No
64GB DDR4 DIMM	FRU	No
7.68TB SAS-3 Disk Assembly	CRU	Yes
1.2TB - 10000 RPM SAS-3 Disk Assembly	CRU	Yes
Fortville dual PCIe 40Gb Ethernet Adapter	FRU	No
2.5" Disk Cage Front Indicator Module	FRU	No
12-Slot 2.5" Disk Backplane Assembly	FRU	No
Interlock Cable, 125mm	FRU	No
Cable Kit	FRU	No
Dual Counter Rotating Fan Module	CRU	Yes
System Board Assembly	FRU	No
3V lithium coin cell battery	CRU	No
Type A266 800/1200 Watt AC Input Power Supply	CRU	Yes
Cluster Heartbeat Assembly	FRU	No
8-Port 12Gbps SAS HBA	FRU	No
4x4 Port 12Gbps SAS-3 PCI Express HBA	FRU	No
<i>Oracle Storage DE3-24C Disk Shelf:</i>		
580 Watt AC Input Power Supply	FRU	Yes
12Gbps SAS-3 I/O Controller Module	FRU	Yes
4RU Chassis Assembly with Midplane	FRU	No
36-Pin Mini SAS3 HD Cable, SFF-8644 to SFF-8644, 3M	FRU	Yes
DE3-24C Mounting Rail Kit	FRU	No
14TB - 7200 RPM SAS-3 Disk Drive Assembly	CRU	Yes
200GB SAS-3 Solid State Drive Assembly	CRU	Yes

For Oracle ZFS Storage Appliance ZS7-2 component servicing instructions, see [Section 6.11, "Servicing the Oracle ZFS Storage Appliance ZS7-2"](#).

6.2.9 Oracle ZFS Storage Appliance ZS5-ES Components

The following table lists the replaceable components of the Oracle ZFS Storage Appliance ZS5-ES.



Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle Private Cloud Appliance components list in the [Oracle System Handbook](#).

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.9 Replaceable Oracle ZFS Storage Appliance ZS5-ES Components

Component Description	FRU/CRU	Hot-Swap
<i>Oracle ZFS Storage Appliance ZS5-ES Storage Head:</i>		
2.3GHz Intel 18-Core Xeon E5-2699 V3, 145W	FRU	No
Pre-greased CPU Heatsink	FRU	No
16GB DDR4-2133 DIMM	FRU	No
3.2TB SAS-3 Solid State Drive Assembly	CRU	Yes
1.2TB - 10000 RPM SAS-3 Disk Assembly	CRU	Yes
Dual 40Gbps InfiniBand 4x QDR PCI Express Low Profile Host Channel Adapter	FRU	No
2-Slot PCI Express Riser Assembly	FRU	No
1-Slot PCI Express Riser Assembly	FRU	No
2.5" Disk Cage Front Indicator Module	FRU	No
8-Slot 2.5" Disk Backplane Assembly	FRU	No
Interlock Cable, 125mm	FRU	No
Cable Kit	FRU	No
Dual Counter Rotating Fan Module	CRU	Yes
System Board Assembly	FRU	No
3V lithium coin cell battery	CRU	No
Type A256 600 Watt AC Input Power Supply	CRU	Yes
Cluster Heartbeat Assembly	FRU	No
8-Port 12Gbps SAS HBA	FRU	No
4x4 Port 12Gbps SAS-3 PCI Express HBA	FRU	No
<i>Oracle Storage DE3-24P Disk Shelf:</i>		
580 Watt AC Input Power Supply	FRU	Yes
12Gbps SAS-3 I/O Controller Module	FRU	Yes
2RU Chassis Assembly with Midplane	FRU	No
36-Pin Mini SAS3 HD Cable, SFF-8644 to SFF-8644, 3M	FRU	Yes
DE3-24P Mounting Rail Kit	FRU	No
1.2TB - 10000 RPM SAS-3 Disk Drive Assembly	CRU	Yes
200GB SAS-3 Solid State Drive Assembly	CRU	Yes

For Oracle ZFS Storage Appliance ZS5-ES component servicing instructions, see [Section 6.12, "Servicing the Oracle ZFS Storage Appliance ZS5-ES"](#).

6.2.10 Oracle ZFS Storage Appliance ZS3-ES Components

The following table lists the replaceable components of the Oracle ZFS Storage Appliance ZS3-ES.



Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle Private Cloud Appliance components list in the [Oracle System Handbook](#).

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.10 Replaceable Oracle ZFS Storage Appliance ZS3-ES Components

Component Description	FRU/CRU	Hot-Swap
<i>Oracle ZFS Storage Appliance ZS3-ES Storage Head:</i>		
2.1GHz Intel 8-Core Xeon E5-2658, 95W	FRU	No
Pre-greased CPU Heatsink	FRU	No
16GB DDR-1600 DIMM, 1.35V	FRU	No
1.6TB SAS Solid State Drive Assembly	CRU	Yes
900GB - 10000 RPM SAS Disk Assembly	CRU	Yes
Dual 40Gbps InfiniBand 4x QDR PCI Express Low Profile Host Channel Adapter	FRU	No
2-Slot PCI Express Riser Assembly	FRU	No
1-Slot PCI Express Riser Assembly	FRU	No
2.5" Disk Cage Front Indicator Module	FRU	No
4-Slot 2.5" Disk Backplane Assembly	FRU	No
Cable Kit	FRU	No
Dual Counter Rotating Fan Module	CRU	Yes
System Board Assembly	FRU	No
3V lithium coin cell battery	CRU	No
Type A256 600 Watt AC Input Power Supply	CRU	Yes
Cluster Heartbeat Assembly	FRU	No
8-Port 6Gbps SAS-2 RAID HBA	FRU	No
8-Port 6Gbps SAS-2 PCI Express HBA (LSI)	FRU	No
<i>Oracle Storage DE2-24P Disk Shelf:</i>		
580 Watt AC Input Power Supply	FRU	Yes
6Gbps SAS-2 I/O Controller Module	FRU	Yes
2RU Chassis Assembly with Midplane	FRU	No
4X Mini SAS Cable, SFF-8088 to SFF-8088, 2M	FRU	Yes
DE2-24P Mounting Rail Kit	FRU	No
900GB 10000 RPM SAS Disk Drive Assembly	CRU	Yes

Component Description	FRU/CRU	Hot-Swap
200GB SAS Solid State Drive Assembly	CRU	Yes

For Oracle ZFS Storage Appliance ZS3-ES component servicing instructions, see [Section 6.13, “Servicing the Oracle ZFS Storage Appliance ZS3-ES”](#).

6.2.11 Sun ZFS Storage Appliance 7320 Components

The following table lists the replaceable components of the Sun ZFS Storage Appliance 7320.



Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle Private Cloud Appliance components list in the [Oracle System Handbook](#).

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.11 Replaceable Sun ZFS Storage Appliance 7320 Components

Component Description	FRU/CRU	Hot-Swap
<i>Sun ZFS 7320 Storage Head:</i>		
2.4GHz Intel Quad-Core Xeon E5620, 12MB, 80W	FRU	No
Xeon Heatsink	FRU	No
8GB Registered DDR3L-1333/DDR3L-1600 DIMM, 1.35V	FRU	No
512GB Solid State Drive SATA-2 Assembly	CRU	Yes
500GB - 10000 RPM SATA Disk Assembly with 1 bracket	CRU	Yes
USB Assembly	FRU	Yes
Dual 40Gbps InfiniBand 4x QDR PCI Express Low Profile Host Channel Adapter	FRU	No
4GB USB 2.0 Flash Drive	FRU	No
1-Slot x8 PCI Express Riser Assembly	FRU	No
1-Slot x16 PCI Express Riser Assembly	FRU	No
Power Distribution Board	FRU	No
8-Slot Disk Backplane, SATA DVD	FRU	No
PDB to System Board Ribbon Cable	FRU	
SFF8087 to SFF8087 Mini-SAS Cable, 690mm	FRU	
6-Pin Fan Power Cable	FRU	
Fan Data Ribbon Cable	FRU	
Bus Bar Set	FRU	
Fan Board Assembly	FRU	
Connector Board Assembly, SATA DVD	FRU	

Component Description	FRU/CRU	Hot-Swap
Fan Module	CRU	Yes
System Board Assembly	FRU	No
3V Lithium Coin Cell Battery	FRU	No
Type A247A 760 Watt AC Input Power Supply	CRU	Yes
Cluster Heartbeat Assembly	FRU	
8-Port 6Gbps SAS-2 RAID HBA	FRU	No
<i>Oracle Storage DE2-24P Disk Shelf:</i>		
580 Watt AC Input Power Supply	CRU	Yes
6Gbps SAS-2 I/O Controller Module	FRU	Yes
2RU Chassis Assembly with Midplane	FRU	No
4X Mini SAS Cable, SFF-8088 to SFF-8088, 2M	FRU	
4X Mini SAS Cable, SFF-8088 to SFF-8088, 0.5M	FRU	
DE2-24P Mounting Rail Kit	FRU	
900GB 10000 RPM SAS Disk Drive Assembly	CRU	Yes
73GB SAS Solid State Drive Assembly	CRU	Yes

For Sun ZFS Storage Appliance 7320 component servicing instructions, see [Section 6.14, “Servicing the Sun ZFS Storage Appliance 7320”](#).

6.2.12 Oracle Switch ES1-24 Components

The following table lists the replaceable components of the Oracle Switch ES1-24.



Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle Private Cloud Appliance components list in the [Oracle System Handbook](#).

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.12 Replaceable Oracle Switch ES1-24 Components

Component Description	FRU/CRU	Hot-Swap
24-Port ES1-24 Switch Assembly	FRU	No
Rear-to-Front Airflow Fan Module	CRU	Yes
Type A247A 760 Watt AC Input Power Supply	CRU	Yes

For Oracle Switch ES1-24 component servicing instructions, see [Section 6.15, “Servicing an Oracle Switch ES1-24”](#).

6.2.13 NM2-36P Sun Datacenter InfiniBand Expansion Switch Components

The following table lists the replaceable components of the Sun Datacenter InfiniBand Expansion Switch NM2-36P.



Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle Private Cloud Appliance components list in the [Oracle System Handbook](#).

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.13 Replaceable NM2-36P Sun Datacenter InfiniBand Expansion Switch Components

Component Description	FRU/CRU	Hot-Swap
Datacenter InfiniBand Switch 36 Subassembly	FRU	No
Type A247A 760 Watt AC Input Power Supply	CRU	Yes
Rear-to-Front Airflow Fan Module	CRU	Yes

For NM2-36P Sun Datacenter InfiniBand Expansion Switch component servicing instructions, see [Section 6.16, “Servicing an NM2-36P Sun Datacenter InfiniBand Expansion Switch”](#).

6.2.14 Oracle Fabric Interconnect F1-15 Components

The following table lists the replaceable components of the Oracle Fabric Interconnect F1-15.



Note

For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle Private Cloud Appliance components list in the [Oracle System Handbook](#).

You access the Oracle System Handbook using this link: https://support.oracle.com/handbook_private/.

Click *Current Systems*, then click your generation of *Oracle Private Cloud Appliance Hardware* to open the main product page in the System Handbook.

Table 6.14 Replaceable Oracle Fabric Interconnect F1-15 Components

Component Description	FRU/CRU	Hot-Swap
F1-15 Power Supply	FRU	Yes
QDR Fabric Board	FRU	No
2U/4U Front Panel G2 (Com-X i7)	FRU	No
F1-15 I/O Management Module	FRU	No
F1-15 Fan Tray	FRU	Yes
Quad Port 10 Gigabit Ethernet (GbE) Module	FRU	Yes
Dual Port 2 x 8 Gigabit Fibre Channel I/O Module	FRU	Yes

Component Description	FRU/CRU	Hot-Swap
F1-15 Chassis without Power Supply, Fan, Fabric Board, Front Panel	FRU	No

For Oracle Fabric Interconnect F1-15 component servicing instructions, see [Section 6.17, “Servicing an Oracle Fabric Interconnect F1-15”](#).

6.3 Preparing Oracle Private Cloud Appliance for Service

This section describes safety considerations and prerequisites for component replacement procedures.

Safety Precautions

For your protection, observe the following safety precautions when servicing your equipment:

- Follow all standard cautions, warnings, and instructions marked on the equipment and described in the following documents:
 - The printed document *Important Safety Information for Sun Hardware Systems (7063567)*
 - The [Oracle Private Cloud Appliance Safety and Compliance Guide \(E88194-02\)](#)
- Follow the safety guidelines described in the [Oracle Private Cloud Appliance Installation Guide \(F23453-04\)](#):
 - [Electrical Power Requirements](#)
 - [Rack-mount Safety Precautions](#)
- Follow the electrostatic discharge safety practices as described in this section.
- Disconnect all power supply cords before servicing components.

Electrostatic Discharge Safety

Devices that are sensitive to electrostatic discharge (ESD), such as motherboards, PCIe cards, drives, processors, and memory cards require special handling.



Caution

Equipment Damage

Take antistatic measures and do not touch components along their connector edges.

- **Use an antistatic wrist strap.**

Wear an antistatic wrist strap and use an antistatic mat when handling components such as drive assemblies, boards, or cards. When servicing or removing rack node components, attach an antistatic strap to your wrist and then to a metal area on the chassis. Then disconnect the power cords from the component. Following this practice equalizes the electrical potentials between you and the component.

An antistatic wrist strap is *not* included in the Oracle Private Cloud Appliance shipment.

- **Use an antistatic mat.**

Place ESD-sensitive components such as the motherboard, memory, and other PCB cards on an antistatic mat.

The following items can be used as an antistatic mat:

- Antistatic bag used to wrap an Oracle replacement part
- An ESD mat (orderable from Oracle)
- A disposable ESD mat (shipped with some replacement parts or optional system components)

6.4 Servicing the Oracle Private Cloud Appliance Rack System

This section provides instructions to service replaceable components (CRUs/FRUs) in the appliance rack. Before starting any service procedure, read and follow the guidelines in [Section 6.3, “Preparing Oracle Private Cloud Appliance for Service”](#).

6.4.1 Powering Down Oracle Private Cloud Appliance (When Required)

Some service procedures may require you to power down the Oracle Private Cloud Appliance. Perform the following steps to manually power down the system.



Caution

Whenever a hardware system must be powered down, make sure that the virtual machines hosted by that system are shut down first. If you power down the appliance with running virtual machines, these will be in an error state when the system is returned to operation.

For details, consult the [Oracle VM Manager User's Guide](#).

- [Stop Virtual Machines](#)
- [Stop Server](#)

Shutting down the Oracle VM environment

1. Log in to Oracle VM Manager and open the Servers and VMs tab.
2. Using the navigation tree, select each virtual machine and click Stop to shut it down gracefully.

If the applications hosted by your VMs require the services and machines to be shut down in a particular order, respect those requirements just like you would with physical machines.

Once the VMs have been shut down, you can proceed to power off the compute nodes.

3. Using the navigation tree, select each compute node and click Stop Server to shut it down gracefully.
4. Using SSH and an account with superuser privileges, log into the active management node at the management virtual IP address. Stop Oracle VM Manager by entering the command `service ovmm stop`.

Powering down the system for service

1. If, at this point, any compute nodes have not shut down properly, press the Power button on the running compute nodes in order to shut them down gracefully.
2. Press the Power button on the management nodes in order to shut them down gracefully.

Once the servers are powered off, you can proceed to power off the storage appliance.

3. Press the Power button on the storage server heads attached to the chassis of the storage device.
4. Toggle the rack Power switches to the Off position.



Note

The Ethernet switches do not have power switches. They power off when power is removed, by way of the power distribution unit (PDU) or at the breaker in the data center.

Returning the system to operation after service or unplanned outage

1. Toggle the power distribution unit (PDU) circuit breakers of both PDUs to the On position.
2. Wait at least two minutes to allow the PDUs to complete their power-on sequence.

The Ethernet switches are powered on with the PDUs.

3. Press the Power button on the storage server heads.

Wait approximately two minutes until the power-on self-test completes, and the Power/OK LED on the front panel lights and remains lit.

4. Press the Power button on the management nodes.

The management node that completes booting first assumes the master role.



Note

Compute nodes do not power on automatically like the internal ZFS Storage Appliance, switches and other components. Make sure that the management nodes and internal storage are up and running, then manually power on the compute nodes.

5. When the management nodes are up, press the Power button on the compute nodes.



Caution

The compute node ILOM policy for automatic power-on is disabled, and must remain disabled, to prevent a server from booting prematurely and disrupting the correct boot order of the appliance components.

When all compute nodes are up, verify the status of all system components in Oracle VM Manager.

If no components are in error state, the appliance is ready to resume normal operation.

6.4.2 Service Procedures for Rack System Components

For parts that are not hot-swappable, power down the Oracle Private Cloud Appliance before starting the service procedure. Generally speaking, hot-swappable components can be serviced without specific additional steps.

Table 6.15 Service Instructions for Rack System Components

Replaceable Part(s)	Hot-Swap	Instructions
Power cables		
Ethernet cables		

Replaceable Part(s)	Hot-Swap	Instructions
Cable management arms (CMAs) (Oracle-qualified service technician only)		For removal and installation of a cable management arm, refer to the Oracle Server X8-2 Installation Guide (part no. E93391) . <ul style="list-style-type: none"> • “Remove the Cable Management Arm” • “Install the Cable Management Arm”
Slide rails (Oracle-qualified service technician only)		To service the slide rails, the server must be removed from the rack. For instructions, refer to the Oracle Server X8-2 Service Manual (part no. E93386) . <ul style="list-style-type: none"> • “Remove the Server From the Rack” • “Reinstall the Server Into the Rack” <p>For slide rail installation instructions, refer to the section Attach the Slide-Rails in the Oracle Server X8-2 Installation Guide (part no. E93391). To remove the slide rails, reverse the installation steps.</p>

6.5 Servicing an Oracle Server X8-2

This section provides instructions to service replaceable components (CRUs/FRUs) in an Oracle Server X8-2 compute node. Before starting any service procedure, read and follow the guidelines in [Section 6.3, “Preparing Oracle Private Cloud Appliance for Service”](#).

6.5.1 Powering Down Oracle Server X8-2 for Service (When Required)

If you need to execute a service procedure that requires the Oracle Server X8-2 to be powered down, follow these instructions:

Placing a compute node into maintenance mode

Before an Oracle Server X8-2 compute node can be powered down, it must be placed into maintenance mode from within Oracle VM Manager. As a result, all virtual machines running on the compute node are automatically migrated to other servers in the Oracle VM server pool, if they are available. Information on maintenance mode is provided in the *Oracle VM Manager User's Guide* section entitled [Edit Server](#).

1. Log in to the Oracle VM Manager Web UI.

For details, refer to the section “[Section 5.2, “Logging in to the Oracle VM Manager Web UI”](#)” in the [Oracle Private Cloud Appliance Administrator's Guide](#).

- a. Enter the following address in a Web browser: `https://manager-vip:7002/ovm/console`.

Replace *manager-vip* with the virtual IP address, or corresponding host name, that you have configured for your management nodes during installation.

- b. Enter the Oracle VM Manager user name and password in the respective fields and click OK.

2. In the **Servers and VMs** tab, select the Oracle VM Server in the navigation pane. Click **Edit Server** in the management pane toolbar.

The **Edit Server** dialog box is displayed.

3. Select the **Maintenance Mode** check box to place the Oracle VM Server into maintenance mode. Click OK.

The Oracle VM Server is in maintenance mode and ready for servicing.

4. When the Oracle Server X8-2 is ready to rejoin the Oracle VM server pool, perform the same procedure and clear the **Maintenance Mode** check box.

Powering down the system

These steps briefly describe the procedure. For detailed instructions, refer to the chapter “[Preparing for Service](#)” in the [Oracle Server X8-2 Service Manual \(part no. E93386\)](#).

1. Power down the server gracefully whenever possible.

The easiest way is to [press and quickly release the Power button](#) on the front panel.

2. Perform immediate shutdown only if the system does not respond to graceful power-down tasks.



Caution

An immediate power down might corrupt system data, therefore, only use this procedure to power down the server after attempting the graceful power down procedure.

3. Disconnect the power cables and data cables from the server.
4. Extend the server to the maintenance position.
5. Most service operations can be performed while the server is in the maintenance position.

However, if necessary, remove the cable management arm (CMA) and pull the server out of the rack.



Caution

The server weighs approximately 15.9 kg (35.0 lb). Two people are required to dismount and carry the chassis.

Returning the system to operation

These steps briefly describe the procedure. For detailed instructions, refer to the chapter “[Returning the Server to Operation](#)” in the [Oracle Server X8-2 Service Manual \(part no. E93386\)](#).

1. If the top cover was removed to service a component, reinstall the top cover on the server.
2. If the server was removed, reinstall it into the rack.
3. Return the server to its normal operational position in the rack, making sure the CMA is correctly installed.
4. Reconnect data cables and power cords.
5. Power on the server.

6.5.2 Service Procedures for Oracle Server X8-2 Components

For parts that are not hot-swappable, power down the Oracle Server X8-2 before starting the service procedure. If the server is in use in the Oracle VM environment, place it in maintenance mode first. This

protects your virtual infrastructure against data corruption, and allows it to remain in service as long as the configuration of your environment allows it.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle Private Cloud Appliance. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.16 Service Procedures for Oracle Server X8-2 Components

Replaceable Part(s)	Hot-Swap	URL
Storage drives	Yes	https://docs.oracle.com/cd/E93359_01/html/E93386/gquak.html#scrolltoc
Fan Modules	Yes	https://docs.oracle.com/cd/E93359_01/html/E93386/gquhg.html#scrolltoc
Power supplies	Yes	https://docs.oracle.com/cd/E93359_01/html/E93386/gqunc.html#scrolltoc
DIMMs (Oracle-qualified service technician only)	No	https://docs.oracle.com/cd/E93359_01/html/E93386/gqvkr.html#scrolltoc
PCI Express risers (Oracle-qualified service technician only)	No	https://docs.oracle.com/cd/E93359_01/html/E93386/gqvft.html#scrolltoc
PCI Express cards (Oracle-qualified service technician only)	No	https://docs.oracle.com/cd/E93359_01/html/E93386/gqvjk.html#scrolltoc
Battery	No	https://docs.oracle.com/cd/E93359_01/html/E93386/gqviw.html#scrolltoc

6.6 Servicing an Oracle Server X7-2

This section provides instructions to service replaceable components (CRUs/FRUs) in an Oracle Server X7-2 compute node. Before starting any service procedure, read and follow the guidelines in [Section 6.3, “Preparing Oracle Private Cloud Appliance for Service”](#).

6.6.1 Powering Down Oracle Server X7-2 for Service (When Required)

If you need to execute a service procedure that requires the Oracle Server X7-2 to be powered down, follow these instructions:

Placing a compute node into maintenance mode

Before an Oracle Server X7-2 compute node can be powered down, it must be placed into maintenance mode from within Oracle VM Manager. As a result, all virtual machines running on the compute node are automatically migrated to other servers in the Oracle VM server pool, if they are available. Information on maintenance mode is provided in the *Oracle VM Manager User's Guide* section entitled [Edit Server](#).

1. Log in to the Oracle VM Manager Web UI.

For details, refer to the section “[Section 5.2, “Logging in to the Oracle VM Manager Web UI”](#)” in the [Oracle Private Cloud Appliance Administrator's Guide](#).

- a. Enter the following address in a Web browser: `https://manager-vip:7002/ovm/console`.

Replace *manager-vip* with the virtual IP address, or corresponding host name, that you have configured for your management nodes during installation.

- b. Enter the Oracle VM Manager user name and password in the respective fields and click OK.
2. In the **Servers and VMs** tab, select the Oracle VM Server in the navigation pane. Click **Edit Server** in the management pane toolbar.

The **Edit Server** dialog box is displayed.

3. Select the **Maintenance Mode** check box to place the Oracle VM Server into maintenance mode. Click OK.

The Oracle VM Server is in maintenance mode and ready for servicing.

4. When the Oracle Server X7-2 is ready to rejoin the Oracle VM server pool, perform the same procedure and clear the **Maintenance Mode** check box.

Powering down the system

These steps briefly describe the procedure. For detailed instructions, refer to the chapter “[Preparing for Service](#)” in the [Oracle Server X7-2 Service Manual \(part no. E72445\)](#).

1. Power down the server gracefully whenever possible.

The easiest way is to [press and quickly release the Power button](#) on the front panel.

2. Perform immediate shutdown only if the system does not respond to graceful power-down tasks.



Caution

An immediate power down might corrupt system data, therefore, only use this procedure to power down the server after attempting the graceful power down procedure.

3. Disconnect the power cables and data cables from the server.
4. Extend the server to the maintenance position.
5. Most service operations can be performed while the server is in the maintenance position.

However, if necessary, remove the cable management arm (CMA) and pull the server out of the rack.



Caution

The server weighs approximately 15.9 kg (35.0 lb). Two people are required to dismount and carry the chassis.

Returning the system to operation

These steps briefly describe the procedure. For detailed instructions, refer to the chapter “[Returning the Server to Operation](#)” in the [Oracle Server X7-2 Service Manual \(part no. E72445\)](#).

1. If the top cover was removed to service a component, reinstall the top cover on the server.
2. If the server was removed, reinstall it into the rack.
3. Return the server to its normal operational position in the rack, making sure the CMA is correctly installed.
4. Reconnect data cables and power cords.
5. Power on the server.

6.6.2 Service Procedures for Oracle Server X7-2 Components

For parts that are not hot-swappable, power down the Oracle Server X7-2 before starting the service procedure. If the server is in use in the Oracle VM environment, place it in maintenance mode first. This protects your virtual infrastructure against data corruption, and allows it to remain in service as long as the configuration of your environment allows it.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle Private Cloud Appliance. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.17 Service Procedures for Oracle Server X7-2 Components

Replaceable Part(s)	Hot-Swap	URL
Storage drives	Yes	http://docs.oracle.com/cd/E72435_01/html/E72445/gquak.html#scrolltoc
Fan Modules	Yes	http://docs.oracle.com/cd/E72435_01/html/E72445/gquhg.html#scrolltoc
Power supplies	Yes	http://docs.oracle.com/cd/E72435_01/html/E72445/gqunc.html#scrolltoc
DIMMs (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E72435_01/html/E72445/gqvkr.html#scrolltoc
PCI Express risers (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E72435_01/html/E72445/gqvft.html#scrolltoc
PCI Express cards (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E72435_01/html/E72445/gqvjk.html#scrolltoc
Battery	No	http://docs.oracle.com/cd/E72435_01/html/E72445/gqviw.html#scrolltoc

6.7 Servicing an Oracle Server X6-2

This section provides instructions to service replaceable components (CRUs/FRUs) in an Oracle Server X6-2 compute node. Before starting any service procedure, read and follow the guidelines in [Section 6.3, "Preparing Oracle Private Cloud Appliance for Service"](#).

6.7.1 Powering Down Oracle Server X6-2 for Service (When Required)

If you need to execute a service procedure that requires the Oracle Server X6-2 to be powered down, follow these instructions:

Placing a compute node into maintenance mode

Before an Oracle Server X6-2 compute node can be powered down, it must be placed into maintenance mode from within Oracle VM Manager. As a result, all virtual machines running on the compute node are automatically migrated to other servers in the Oracle VM server pool, if they are available. Information on maintenance mode is provided in the *Oracle VM Manager User's Guide* section entitled [Edit Server](#).

1. Log in to the Oracle VM Manager Web UI.

For details, refer to the section “[Section 5.2, “Logging in to the Oracle VM Manager Web UI”](#)” in the [Oracle Private Cloud Appliance Administrator's Guide](#).

- a. Enter the following address in a Web browser: `https://manager-vip:7002/ovm/console`.

Replace *manager-vip* with the virtual IP address, or corresponding host name, that you have configured for your management nodes during installation.

- b. Enter the Oracle VM Manager user name and password in the respective fields and click OK.

2. In the **Servers and VMs** tab, select the Oracle VM Server in the navigation pane. Click **Edit Server** in the management pane toolbar.

The **Edit Server** dialog box is displayed.

3. Select the **Maintenance Mode** check box to place the Oracle VM Server into maintenance mode. Click OK.

The Oracle VM Server is in maintenance mode and ready for servicing.

4. When the Oracle Server X6-2 is ready to rejoin the Oracle VM server pool, perform the same procedure and clear the **Maintenance Mode** check box.

Powering down the system

These steps briefly describe the procedure. For detailed instructions, refer to the chapter “[Preparing for Service](#)” in the [Oracle Server X6-2 Service Manual \(part no. E62171\)](#).

1. Power down the server gracefully whenever possible.

The easiest way is to [press and quickly release the Power button](#) on the front panel.

2. Perform immediate shutdown only if the system does not respond to graceful power-down tasks.



Caution

System data may become corrupted during an immediate power down. Use this task only after attempting to power down the server gracefully.

3. Disconnect the power cables and data cables from the server.
4. Extend the server to the maintenance position.
5. Most service operations can be performed while the server is in the maintenance position.

However, if necessary, remove the cable management arm (CMA) and pull the server out of the rack.



Caution

The server weighs approximately 18.1 kg (39.9 lb). Two people are required to dismount and carry the chassis.

Returning the system to operation

These steps briefly describe the procedure. For detailed instructions, refer to the chapter “[Returning the Server to Operation](#)” in the [Oracle Server X6-2 Service Manual \(part no. E62171\)](#).

1. If the top cover was removed to service a component, reinstall the top cover on the server.
2. If the server was removed, reinstall it into the rack.
3. Return the server to its normal operational position in the rack, making sure the CMA is correctly installed.
4. Reconnect data cables and power cords.
5. Power on the server.

6.7.2 Service Procedures for Oracle Server X6-2 Components

For parts that are not hot-swappable, power down the Oracle Server X6-2 before starting the service procedure. If the server is in use in the Oracle VM environment, place it in maintenance mode first. This protects your virtual infrastructure against data corruption, and allows it to remain in service as long as the configuration of your environment allows it.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle Private Cloud Appliance. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.18 Service Procedures for Oracle Server X6-2 Components

Replaceable Part(s)	Hot-Swap	URL
Storage drives	Yes	http://docs.oracle.com/cd/E62159_01/html/E62171/z40000091011460.html#scrolltoc
Fan Modules	Yes	http://docs.oracle.com/cd/E62159_01/html/E62171/z40000091014194.html#scrolltoc
Power supplies	Yes	http://docs.oracle.com/cd/E62159_01/html/E62171/z40000091014153.html#scrolltoc
DIMMs (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E62159_01/html/E62171/z40003f01425075.html#scrolltoc
PCI Express risers (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E62159_01/html/E62171/z40000f91037394.html#scrolltoc

Replaceable Part(s)	Hot-Swap	URL
PCI Express cards (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E62159_01/html/E62171/z40000f91037409.html#scrolltoc
Internal USB flash drives (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E62159_01/html/E62171/z4000a6d1442801.html#scrolltoc
Battery	No	http://docs.oracle.com/cd/E62159_01/html/E62171/z40003f01423753.html#scrolltoc

6.8 Servicing an Oracle Server X5-2

This section provides instructions to service replaceable components (CRUs/FRUs) in an Oracle Server X5-2 management node or compute node. Before starting any service procedure, read and follow the guidelines in [Section 6.3, “Preparing Oracle Private Cloud Appliance for Service”](#).

6.8.1 Powering Down Oracle Server X5-2 for Service (When Required)

If you need to execute a service procedure that requires the Oracle Server X5-2 to be powered down, follow these instructions:



Note

The management nodes are not placed in maintenance mode for servicing. If you need to power down the master management node, bring it offline as described below and wait for the other management node to take over the master role. If you need to power down the secondary management node, no additional steps are required.

Placing a compute node into maintenance mode

Before an Oracle Server X5-2 compute node can be powered down, it must be placed into maintenance mode from within Oracle VM Manager. As a result, all virtual machines running on the compute node are automatically migrated to other servers in the Oracle VM server pool, if they are available. Information on maintenance mode is provided in the *Oracle VM Manager User's Guide* section entitled [Edit Server](#).

1. Log in to the Oracle VM Manager Web UI.

For details, refer to the section “[Section 5.2, “Logging in to the Oracle VM Manager Web UI”](#)” in the [Oracle Private Cloud Appliance Administrator's Guide](#).

- a. Enter the following address in a Web browser: `https://manager-vip:7002/ovm/console`.

Replace *manager-vip* with the virtual IP address, or corresponding host name, that you have configured for your management nodes during installation.

- b. Enter the Oracle VM Manager user name and password in the respective fields and click OK.

2. In the **Servers and VMs** tab, select the Oracle VM Server in the navigation pane. Click **Edit Server** in the management pane toolbar.

The **Edit Server** dialog box is displayed.

3. Select the **Maintenance Mode** check box to place the Oracle VM Server into maintenance mode. Click OK.

The Oracle VM Server is in maintenance mode and ready for servicing.

4. When the Oracle Server X5-2 is ready to rejoin the Oracle VM server pool, perform the same procedure and clear the **Maintenance Mode** check box.

Powering down the system

These steps briefly describe the procedure. For detailed instructions, refer to the chapter “[Preparing for Service](#)” in the [Oracle Server X5-2 Service Manual \(part no. E48320\)](#).

1. Power down the server gracefully whenever possible.

The easiest way is to [press and quickly release the Power button](#) on the front panel.

2. Perform immediate shutdown only if the system does not respond to graceful power-down tasks.



Caution

System data may become corrupted during an immediate power down. Use this task only after attempting to power down the server gracefully.

3. Disconnect the power cables and data cables from the server.
4. Extend the server to the maintenance position.
5. Most service operations can be performed while the server is in the maintenance position.

However, if necessary, remove the cable management arm (CMA) and pull the server out of the rack.



Caution

The server weighs approximately 18.1 kg (39.9 lb). Two people are required to dismount and carry the chassis.

Returning the system to operation

These steps briefly describe the procedure. For detailed instructions, refer to the chapter “[Returning the Server to Operation](#)” in the [Oracle Server X5-2 Service Manual \(part no. E48320\)](#).

1. If the top cover was removed to service a component, reinstall the top cover on the server.
2. If the server was removed, reinstall it into the rack.
3. Return the server to its normal operational position in the rack, making sure the CMA is correctly installed.
4. Reconnect data cables and power cords.
5. Power on the server.

6.8.2 Service Procedures for Oracle Server X5-2 Components

For parts that are not hot-swappable, power down the Oracle Server X5-2 before starting the service procedure. If the server is in use in the Oracle VM environment, place it in maintenance mode first. This

protects your virtual infrastructure against data corruption, and allows it to remain in service as long as the configuration of your environment allows it.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle Private Cloud Appliance. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.19 Service Procedures for Oracle Server X5-2 Components

Replaceable Part(s)	Hot-Swap	URL
Storage drives	Yes	http://docs.oracle.com/cd/E41059_01/html/E48320/z40000091011460.html#scrolltoc
Fan Modules	Yes	http://docs.oracle.com/cd/E41059_01/html/E48320/z40000091014194.html#scrolltoc
Power supplies	Yes	http://docs.oracle.com/cd/E41059_01/html/E48320/z40000091014153.html#scrolltoc
DIMMs (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E41059_01/html/E48320/z40003f01425075.html#scrolltoc
PCI Express risers (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E41059_01/html/E48320/z40000f91037394.html#scrolltoc
PCI Express cards (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E41059_01/html/E48320/z40000f91037409.html#scrolltoc
Internal USB flash drives (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E41059_01/html/E48320/z4000a6d1442801.html#scrolltoc
Battery	No	http://docs.oracle.com/cd/E41059_01/html/E48320/z40003f01423753.html#scrolltoc

6.9 Servicing a Sun Server X4-2

This section provides instructions to service replaceable components (CRUs/FRUs) in a Sun Server X4-2 management node or compute node. Before starting any service procedure, read and follow the guidelines in [Section 6.3, “Preparing Oracle Private Cloud Appliance for Service”](#).

6.9.1 Powering Down Sun Server X4-2 for Service (When Required)

If you need to execute a service procedure that requires the Sun Server X4-2 to be powered down, follow these instructions:



Note

The management nodes are not placed in maintenance mode for servicing. If you need to power down the master management node, bring it offline as described

below and wait for the other management node to take over the master role. If you need to power down the secondary management node, no additional steps are required.

Placing a compute node into maintenance mode

Before a Sun Server X4-2 compute node can be powered down, it must be placed into maintenance mode from within Oracle VM Manager. As a result, all virtual machines running on the compute node are automatically migrated to other servers in the Oracle VM server pool, if they are available. For details, refer to the section “[Placing an Oracle VM Server into Maintenance Mode](#)” in the [Oracle VM Manager User's Guide](#).

1. Log in to the Oracle VM Manager Web UI.

For details, refer to the section “[Section 5.2, “Logging in to the Oracle VM Manager Web UI”](#)” in the [Oracle Private Cloud Appliance Administrator's Guide](#).

- a. Enter the following address in a Web browser: `https://manager-vip:7002/ovm/console`.

Replace *manager-vip* with the virtual IP address, or corresponding host name, that you have configured for your management nodes during installation.

- b. Enter the Oracle VM Manager user name and password in the respective fields and click OK.

2. In the **Servers and VMs** tab, select the Oracle VM Server in the navigation pane. Click **Edit Server** in the management pane toolbar.

The **Edit Server** dialog box is displayed.

3. Select the **Maintenance Mode** check box to place the Oracle VM Server into maintenance mode. Click OK.

The Oracle VM Server is in maintenance mode and ready for servicing.

4. When the Sun Server X4-2 is ready to rejoin the Oracle VM server pool, perform the same procedure and clear the **Maintenance Mode** check box.

Powering down the system

These steps briefly describe the procedure. For detailed instructions, refer to the chapter “[Preparing for Service](#)” in the [Sun Server X4-2 Service Manual \(part no. E38041\)](#).

1. Power down the server gracefully whenever possible.

The easiest way is to [press and quickly release the Power button](#) on the front panel.

2. Perform immediate shutdown only if the system does not respond to graceful power-down tasks.



Caution

System data may become corrupted during an immediate power down. Use this task only after attempting to power down the server gracefully.

3. Disconnect the power cables and data cables from the server.
4. Extend the server to the maintenance position.
5. Most service operations can be performed while the server is in the maintenance position.

However, if necessary, remove the cable management arm (CMA) and pull the server out of the rack.



Caution

The server weighs approximately 18.1 kg (39.9 lb). Two people are required to dismount and carry the chassis.

Returning the system to operation

These steps briefly describe the procedure. For detailed instructions, refer to the chapter “Returning the Server to Operation” in the [Sun Server X4-2 Service Manual \(part no. E38041\)](#).

1. If the top cover was removed to service a component, reinstall the top cover on the server.
2. If the server was removed, reinstall it into the rack.
3. Return the server to its normal operational position in the rack, making sure the CMA is correctly installed.
4. Reconnect data cables and power cords.
5. Power on the server.

6.9.2 Service Procedures for Sun Server X4-2 Components

For parts that are not hot-swappable, power down the Sun Server X4-2 before starting the service procedure. If the server is in use in the Oracle VM environment, place it in maintenance mode first. This protects your virtual infrastructure against data corruption, and allows it to remain in service as long as the configuration of your environment allows it.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle Private Cloud Appliance. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.20 Service Procedures for Sun Server X4-2 Components

Replaceable Part(s)	Hot-Swap	URL
Storage drives	Yes	http://docs.oracle.com/cd/E36975_01/html/E38045/z40000091011460.html#scrolltoc
Fan Modules	Yes	http://docs.oracle.com/cd/E36975_01/html/E38045/z40000091014194.html#scrolltoc
Power supplies	Yes	http://docs.oracle.com/cd/E36975_01/html/E38045/z40000091014153.html#scrolltoc
DIMMs (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E36975_01/html/E38045/z40003f01425075.html#scrolltoc
PCI Express risers (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E36975_01/html/E38045/z40000f91037394.html#scrolltoc

Replaceable Part(s)	Hot-Swap	URL
PCI Express cards (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E36975_01/html/E38045/z40000f91037409.html#scrolltoc
Internal USB flash drives (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E36975_01/html/E38045/z4000a6d1442801.html#scrolltoc
Battery	No	http://docs.oracle.com/cd/E36975_01/html/E38045/z40003f01423753.html#scrolltoc

6.10 Servicing a Sun Server X3-2

This section provides instructions to service replaceable components (CRUs/FRUs) in a Sun Server X3-2 management node or compute node. Before starting any service procedure, read and follow the guidelines in [Section 6.3, “Preparing Oracle Private Cloud Appliance for Service”](#).

6.10.1 Powering Down Sun Server X3-2 for Service (When Required)

If you need to execute a service procedure that requires the Sun Server X3-2 to be powered down, follow these instructions:



Note

The management nodes are not placed in maintenance mode for servicing. If you need to power down the master management node, bring it offline as described below and wait for the other management node to take over the master role. If you need to power down the secondary management node, no additional steps are required.

Placing a compute node into maintenance mode

Before a Sun Server X3-2 compute node can be powered down, it must be placed into maintenance mode from within Oracle VM Manager. As a result, all virtual machines running on the compute node are automatically migrated to other servers in the Oracle VM server pool, if they are available. For details, refer to the section [“Placing an Oracle VM Server into Maintenance Mode”](#) in the [Oracle VM Manager User's Guide](#).

1. Log in to the Oracle VM Manager Web UI.

For details, refer to the section [“Section 5.2, “Logging in to the Oracle VM Manager Web UI””](#) in the [Oracle Private Cloud Appliance Administrator's Guide](#).

- a. Enter the following address in a Web browser: `https://manager-vip:7002/ovm/console`.

Replace *manager-vip* with the virtual IP address, or corresponding host name, that you have configured for your management nodes during installation.

- b. Enter the Oracle VM Manager user name and password in the respective fields and click OK.

2. In the **Servers and VMs** tab, select the Oracle VM Server in the navigation pane. Click **Edit Server** in the management pane toolbar.

The **Edit Server** dialog box is displayed.

3. Select the **Maintenance Mode** check box to place the Oracle VM Server into maintenance mode. Click OK.

The Oracle VM Server is in maintenance mode and ready for servicing.

4. When the Sun Server X3-2 is ready to rejoin the Oracle VM server pool, perform the same procedure and clear the **Maintenance Mode** check box.

Powering down the system

These steps briefly describe the procedure. For detailed instructions, refer to the chapter “[Preparing for Service](#)” in the [Sun Server X3-2 Service Manual \(part no. E22313\)](#).

1. Power down the server gracefully whenever possible.

The easiest way is to [press and quickly release the Power button](#) on the front panel.

2. Perform immediate shutdown only if the system does not respond to graceful power-down tasks.



Caution

System data may become corrupted during an immediate power down. Use this task only after attempting to power down the server gracefully.

3. Extend the server to the maintenance position.
4. Disconnect the power cables and data cables from the server.
5. Most service operations can be performed while the server is in the maintenance position.

However, if necessary, remove the cable management arm (CMA) and pull the server out of the rack.



Caution

The server weighs approximately 18.1 kg (39.9 lb). Two people are required to dismount and carry the chassis.

Returning the system to operation

These steps briefly describe the procedure. For detailed instructions, refer to the chapter “[Returning the Server to Operation](#)” in the [Sun Server X3-2 Service Manual \(part no. E22313\)](#).

1. If the top cover was removed to service a component, reinstall the top cover on the server.
2. If the server was removed, reinstall it into the rack.
3. Reconnect data cables and power cords.
4. Return the server to its normal operational position in the rack, making sure the CMA is correctly installed.
5. Power on the server.

6.10.2 Service Procedures for Sun Server X3-2 Components

For parts that are not hot-swappable, power down the Sun Server X3-2 before starting the service procedure. If the server is in use in the Oracle VM environment, place it in maintenance mode first. This

protects your virtual infrastructure against data corruption, and allows it to remain in service as long as the configuration of your environment allows it.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle Private Cloud Appliance. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.21 Service Procedures for Sun Server X3-2 Components

Replaceable Part(s)	Hot-Swap	URL
Storage drives	Yes	http://docs.oracle.com/cd/E22368_01/html/E27242/z40000091011460.html#scrolltoc
Fan Modules	Yes	http://docs.oracle.com/cd/E22368_01/html/E27242/z40000091014194.html#scrolltoc
Power supplies	Yes	http://docs.oracle.com/cd/E22368_01/html/E27242/z40000091014153.html#scrolltoc
DIMMs (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E22368_01/html/E27242/z40003f01425075.html#scrolltoc
PCI Express risers (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E22368_01/html/E27242/z40000f91037394.html#scrolltoc
PCI Express cards (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E22368_01/html/E27242/z40000f91037409.html#scrolltoc
Internal USB flash drives (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E22368_01/html/E27242/z4000a6d1442801.html#scrolltoc
Battery	No	http://docs.oracle.com/cd/E22368_01/html/E27242/z40003f01423753.html#scrolltoc

6.11 Servicing the Oracle ZFS Storage Appliance ZS7-2

This section provides instructions to service replaceable components (CRUs/FRUs) in the Oracle ZFS Storage Appliance ZS7-2. Before starting any service procedure, read and follow the guidelines in [Section 6.3, “Preparing Oracle Private Cloud Appliance for Service”](#).

6.11.1 Powering Down the Oracle ZFS Storage Appliance ZS7-2 for Service (When Required)

If you need to execute a service procedure that requires the Oracle ZFS Storage Appliance ZS7-2 to be powered down, follow these instructions:

Powering down the storage head/controller

Because the storage controllers are clustered, there is no loss of access to storage when one controller is powered down for service. Performing a graceful shutdown ensures that data is saved and not corrupted,

and that resources are assigned to the other controller in the storage head cluster. Power down a controller for component replacement using one of the following methods:

- Log in to the UI by using the server's IP address in the appliance management network:
 1. In your browser, enter `https://ipaddress:215`.
 2. Log in as root, using the system-wide Oracle Private Cloud Appliance password.
 3. Click the **Power** icon on the left side under *masthead*.
- Alternatively, SSH in to the storage appliance as root, and enter the command `maintenance system poweroff`.

If graceful shutdown as described above is not possible, use the power button:

- Use a pen or non-conducting pointed object to press and release the Power button on the front panel.
- SSH or use a serial connection to log in to the service processor (SP), and then issue the command `stop /SYS`.
- If the server did not respond, initiate an emergency shutdown. Press and hold the Power button for at least four seconds until the Power/OK status indicator on the front panel flashes, indicating that the storage controller is in standby power mode. To completely remove power, disconnect the AC power cords from the rear panel of the storage controller.



Caution

An emergency shutdown causes all applications and files to be closed abruptly without saving. You might corrupt or lose system data, or lose the server configuration (the resources assigned to it) during an immediate power down.



Powering down the disk shelf is not required

All replaceable components in the disk shelf are hot-swappable. The disk shelf itself does not need to be powered down for the replacement of defective components.

However, do not remove a component if you do not have an immediate replacement. The disk shelf must not be operated without all components in place.

Powering on the storage appliance



Caution

The disk shelf must not be operated without all components in place.

1. Connect any storage head power and data cables you removed to service a component.
2. Power on the server by pressing the Power button on the front panel.

If you are not physically located at the system, use either of these ILOM methods instead:

- Log in to the **Oracle ILOM web interface**.
Click Host Management > Power Control, and in the Actions list click **Power On**.
- Log in to the **Oracle ILOM command-line interface (CLI)**.
At the CLI prompt, type the following command: `start /System`.

3. When the controller is powered on and the power-on self-test (POST) code checkpoint tests have completed, the green Power/OK status indicator on the front panel lights and remains lit.
4. If you performed a graceful shutdown earlier, return resources to the server that was just serviced.
 - a. Log into the web UI for the server that was not serviced.
 - b. Go to Configuration > Cluster.
 - c. Click **Failback**.



Note

For information about configuring the clustered servers and attached disk shelves, see the “Oracle ZFS Storage System Administration Guide” for the appropriate software release.

6.11.2 Service Procedures for Oracle ZFS Storage Appliance ZS7-2 Components

For parts that are not hot-swappable, power down the Oracle ZFS Storage Appliance ZS7-2 before starting the service procedure.



Warning

If you need to execute a service procedure that interrupts the connection between virtual machines and their virtual disks, shut down the virtual machines in Oracle VM Manager prior to servicing the storage hardware. Disconnecting a running virtual machine from its disks may cause data corruption.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle Private Cloud Appliance. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.22 Service Procedures for Oracle ZFS Storage Appliance ZS7-2 Components

Replaceable Part(s)	Hot-Swap	URL
Storage head hard drives	Yes	https://docs.oracle.com/cd/F13758_01/html/F13771/gtbno.html#scrolltoc
Disk shelf drives	Yes	https://docs.oracle.com/cd/F13758_01/html/F13771/goxds.html#scrolltoc
Fan modules	Yes	https://docs.oracle.com/cd/F13758_01/html/F13771/gtboxa.html#scrolltoc
Storage head power supplies	Yes	https://docs.oracle.com/cd/F13758_01/html/F13771/gtbon.html#scrolltoc
Disk shelf power supplies	Yes	https://docs.oracle.com/cd/F13758_01/html/F13771/goxbs.html#scrolltoc
Memory modules (Oracle-qualified service technician only)	No	https://docs.oracle.com/cd/F13758_01/html/F13771/gtbou.html#scrolltoc
PCI Express cards	No	https://docs.oracle.com/cd/F13758_01/html/F13771/gtbnz.html#scrolltoc

Replaceable Part(s)	Hot-Swap	URL
(Oracle-qualified service technician only)		
Battery	No	https://docs.oracle.com/cd/F13758_01/html/F13771/gtbwl.html#scrolltoc
Disk shelf I/O modules (Oracle-qualified service technician only)	Yes	https://docs.oracle.com/cd/F13758_01/html/F13771/goxeo.html#scrolltoc
Disk shelf SIM boards (Oracle-qualified service technician only)	Yes	https://docs.oracle.com/cd/F13758_01/html/F13771/goxef.html#scrolltoc

6.12 Servicing the Oracle ZFS Storage Appliance ZS5-ES

This section provides instructions to service replaceable components (CRUs/FRUs) in the Oracle ZFS Storage Appliance ZS5-ES. Before starting any service procedure, read and follow the guidelines in [Section 6.3, “Preparing Oracle Private Cloud Appliance for Service”](#).

6.12.1 Powering Down the Oracle ZFS Storage Appliance ZS5-ES for Service (When Required)

If you need to execute a service procedure that requires the Oracle ZFS Storage Appliance ZS5-ES to be powered down, follow these instructions:

Powering down the storage head/controller

Because the storage controllers are clustered, there is no loss of access to storage when one controller is powered down for service. Performing a graceful shutdown ensures that data is saved and not corrupted, and that resources are assigned to the other controller in the storage head cluster. Power down a controller for component replacement using one of the following methods:

- Log in to the UI by using the server's IP address in the appliance management network:
 1. In your browser, enter `https://ipaddress:215`.
 2. Log in as root, using the system-wide Oracle Private Cloud Appliance password.
 3. Click the **Power** icon on the left side under *masthead*.
- Alternatively, SSH in to the storage appliance as root, and enter the command `maintenance system poweroff`.

If graceful shutdown as described above is not possible, use the power button:

- Use a pen or non-conducting pointed object to press and release the Power button on the front panel.
- SSH or use a serial connection to log in to the service processor (SP), and then issue the command `stop /SYS`.
- If the server did not respond, initiate an emergency shutdown. Press and hold the Power button for at least four seconds until the Power/OK status indicator on the front panel flashes, indicating that the storage controller is in standby power mode. To completely remove power, disconnect the AC power cords from the rear panel of the storage controller.

**Caution**

An emergency shutdown causes all applications and files to be closed abruptly without saving. You might corrupt or lose system data, or lose the server configuration (the resources assigned to it) during an immediate power down.

**Powering down the disk shelf is not required**

All replaceable components in the disk shelf are hot-swappable. The disk shelf itself does not need to be powered down for the replacement of defective components.

However, do not remove a component if you do not have an immediate replacement. The disk shelf must not be operated without all components in place.

Powering on the storage appliance**Caution**

The disk shelf must not be operated without all components in place.

1. Connect any storage head power and data cables you removed to service a component.
2. Power on the server by pressing the Power button on the front panel.

If you are not physically located at the system, use either of these ILOM methods instead:

- Log in to the **Oracle ILOM web interface**.

Click Host Management > Power Control, and in the Actions list click **Power On**.

- Log in to the **Oracle ILOM command-line interface (CLI)**.

At the CLI prompt, type the following command: `start /System`.

3. When the controller is powered on and the power-on self-test (POST) code checkpoint tests have completed, the green Power/OK status indicator on the front panel lights and remains lit.
4. If you performed a graceful shutdown earlier, return resources to the server that was just serviced.
 - a. Log into the web UI for the server that was not serviced.
 - b. Go to Configuration > Cluster.
 - c. Click **Failback**.

**Note**

For information about configuring the clustered servers and attached disk shelves, see the “Oracle ZFS Storage System Administration Guide” for the appropriate software release.

6.12.2 Service Procedures for Oracle ZFS Storage Appliance ZS5-ES Components

For parts that are not hot-swappable, power down the Oracle ZFS Storage Appliance ZS5-ES before starting the service procedure.



Warning

If you need to execute a service procedure that interrupts the connection between virtual machines and their virtual disks, shut down the virtual machines in Oracle VM Manager prior to servicing the storage hardware. Disconnecting a running virtual machine from its disks may cause data corruption.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle Private Cloud Appliance. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.23 Service Procedures for Oracle ZFS Storage Appliance ZS5-ES Components

Replaceable Part(s)	Hot-Swap	URL
Storage head hard drives	Yes	https://docs.oracle.com/cd/E59597_01/html/E59600/gqloy.html#scrolltoc
Disk shelf drives	Yes	https://docs.oracle.com/cd/E79446_01/html/E79459/goxds.html#scrolltoc
Fan modules	Yes	https://docs.oracle.com/cd/E59597_01/html/E59600/gqlib.html#scrolltoc
Storage head power supplies	Yes	https://docs.oracle.com/cd/E59597_01/html/E59600/gqlfa.html#scrolltoc
Disk shelf power supplies	Yes	https://docs.oracle.com/cd/E79446_01/html/E79459/goxbs.html#scrolltoc
Memory modules (Oracle-qualified service technician only)	No	https://docs.oracle.com/cd/E59597_01/html/E59600/gqlgl.html#scrolltoc
PCI Express risers (Oracle-qualified service technician only)	No	https://docs.oracle.com/cd/E59597_01/html/E59600/gqllep.html#scrolltoc
PCI Express cards (Oracle-qualified service technician only)	No	https://docs.oracle.com/cd/E59597_01/html/E59600/gqlkc.html#scrolltoc
Battery	No	https://docs.oracle.com/cd/E59597_01/html/E59600/gqlfu.html#scrolltoc
Disk shelf I/O modules (Oracle-qualified service technician only)	Yes	https://docs.oracle.com/cd/E79446_01/html/E79459/goxeo.html#scrolltoc
Disk shelf SIM boards (Oracle-qualified service technician only)	Yes	https://docs.oracle.com/cd/E79446_01/html/E79459/goxef.html#scrolltoc

6.13 Servicing the Oracle ZFS Storage Appliance ZS3-ES

This section provides instructions to service replaceable components (CRUs/FRUs) in the Oracle ZFS Storage Appliance ZS3-ES. Before starting any service procedure, read and follow the guidelines in [Section 6.3, “Preparing Oracle Private Cloud Appliance for Service”](#).

6.13.1 Powering Down the Oracle ZFS Storage Appliance ZS3-ES for Service (When Required)

If you need to execute a service procedure that requires the Oracle ZFS Storage Appliance ZS3-ES to be powered down, follow these instructions:

Powering down the storage head/controller

Performing a graceful shutdown ensures that data is saved and not corrupted, and that resources are assigned to the other controller in the storage head cluster. This is the preferred method for powering down a controller for component replacement.

1. Ensure that Ethernet cables are connected from your network to the **NET-0** port on the back of each server.
2. Direct your web browser to the server to be serviced by using either the IP address or host name assigned to the NET-0 port as follows: `https://ipaddress:215`.
3. Log in as root, using the system-wide Oracle Private Cloud Appliance password.
4. Go to **Maintenance**, then select **Hardware**.
5. Click the **Show Details** link for the server.
6. Click the Power icon for the server and select **Power off** from the pull-down list.

If graceful shutdown is not possible, use the power button.



Caution

This task forces the main power off. You might corrupt or lose system data, or lose the server configuration (the resources assigned to it) during an immediate power down.

1. Press and quickly release the Power button on the front panel.

This action causes an orderly shutdown of the operating system, and the server enters the standby power mode.

2. If the server did not respond or you need a more immediate shutdown, press and hold the Power button for four seconds.

This forces the main power off and enters the standby power mode immediately. When the main power is off, the Power/OK LED on the front panel begins flashing, indicating that the server is in standby power mode.

If neither graceful shutdown nor emergency shutdown using the power button is possible, for example because you are not physically located at the system, use the ILOM to perform an emergency shutdown. Choose one of the following options:



Caution

This task forces the main power off. You might corrupt or lose system data, or lose the server configuration (the resources assigned to it) during an immediate power down.

- Log in to the **Oracle ILOM web interface**.

In the left pane, click Host Management > Power Control, and in the Actions list click **Immediate Power Off**.

Click Save, and then click OK.

- Log in to the **Oracle ILOM command-line interface (CLI)**.

At the CLI prompt, type the following command: `stop -f /System`.

Powering down the disk shelf

Do not remove a component if you do not have an immediate replacement. The disk shelf must not be operated without all components in place. Powering down or removing all SAS chains from a disk shelf will cause the controllers to panic to prevent data loss. To avoid this, shut down the controllers before decommissioning the shelf.

1. Stop all input and output to and from the disk shelf.
2. Wait approximately two minutes until all disk activity indicators have stopped flashing.
3. Place the power supply on/off switches to the "O" off position.
4. Disconnect the power cords from the external power source.

Powering on the storage appliance

The disk shelf must not be operated without all components in place.

1. Reconnect the disk shelf power and data cables you removed to service a component.
2. Place the power supply on/off switches on the disk shelf to the "I" on position.
3. Wait several minutes until the boot process is complete, at which time the Power LED should be solid green.
4. Connect the storage head power and data cables you removed to service a component.
5. Power on the server by pressing the Power button on the front panel.

If you are not physically located at the system, use either of these ILOM methods instead:

- Log in to the **Oracle ILOM web interface**.

In the left pane, click Host Management > Power Control, and in the Actions list click **Power On**.

- Log in to the **Oracle ILOM command-line interface (CLI)**.

At the CLI prompt, type the following command: `start /System`.

6. Wait approximately two minutes until the power-on self-test (POST) code checkpoint tests have completed, and the Power/OK LED on the front panel lights and remains lit.
7. If you performed a graceful shutdown earlier, return resources to the server that was just serviced.
 - a. Log into the web UI for the server that was not serviced.
 - b. Go to Configuration > Cluster.

c. Click **Failback**.



Note

For information about configuring the clustered servers and attached disk shelves, see the “Oracle ZFS Storage System Administration Guide” for the appropriate software release.

6.13.2 Service Procedures for Oracle ZFS Storage Appliance ZS3-ES Components

For parts that are not hot-swappable, power down the Oracle ZFS Storage Appliance ZS3-ES before starting the service procedure.



Warning

If you need to execute a service procedure that interrupts the connection between virtual machines and their virtual disks, shut down the virtual machines in Oracle VM Manager prior to servicing the storage hardware. Disconnecting a running virtual machine from its disks may cause data corruption.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle Private Cloud Appliance. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.24 Service Procedures for Oracle ZFS Storage Appliance ZS3-ES Components

Replaceable Part(s)	Hot-Swap	URL
Storage head hard drives	Yes	http://docs.oracle.com/cd/E37831_01/html/E48559/z40000091011460.html#scrolltoc
Disk shelf drives	Yes	Refer to the section “Replacing a Drive” on this page: http://docs.oracle.com/cd/E27998_01/html/E48492/maintenance__hardware__procedures__shelf.html#scrolltoc
Fan modules	Yes	http://docs.oracle.com/cd/E37831_01/html/E48559/z40000091014194.html#scrolltoc
Storage head power supplies	Yes	http://docs.oracle.com/cd/E37831_01/html/E48559/z40000091014153.html#scrolltoc
Disk shelf power supplies	Yes	Refer to the section “Replacing a Power Supply” on this page: http://docs.oracle.com/cd/E27998_01/html/E48492/maintenance__hardware__procedures__shelf.html#scrolltoc
Memory modules (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E37831_01/html/E48559/z40003f01425075.html#scrolltoc
PCI Express risers (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E37831_01/html/E48559/z40000f91037394.html#scrolltoc
PCI Express cards	No	http://docs.oracle.com/cd/E37831_01/html/E48559/z40000f91037409.html#scrolltoc

Replaceable Part(s)	Hot-Swap	URL
(Oracle-qualified service technician only)		
Internal USB flash drive	No	http://docs.oracle.com/cd/E37831_01/html/E48559/z4000a6d1442801.html#scrolltoc
(Oracle-qualified service technician only)		
Battery	No	http://docs.oracle.com/cd/E37831_01/html/E48559/z40003f01423753.html#scrolltoc
Disk shelf I/O modules	Yes	Refer to the section “Replacing an I/O Module” on this page: http://docs.oracle.com/cd/E27998_01/html/E48492/maintenance__hardware__procedures__shelf.html#scrolltoc
(Oracle-qualified service technician only)		
Disk shelf SIM boards	Yes	Refer to the section “Replacing a SIM Board” on this page: http://docs.oracle.com/cd/E27998_01/html/E48492/maintenance__hardware__procedures__shelf.html#scrolltoc
(Oracle-qualified service technician only)		

6.14 Servicing the Sun ZFS Storage Appliance 7320

This section provides instructions to service replaceable components (CRUs/FRUs) in the Sun ZFS Storage Appliance 7320. Before starting any service procedure, read and follow the guidelines in [Section 6.3, “Preparing Oracle Private Cloud Appliance for Service”](#).

6.14.1 Powering Down the Sun ZFS Storage Appliance 7320 for Service (When Required)

If you need to execute a service procedure that requires the Sun ZFS Storage Appliance 7320 to be powered down, follow these instructions:

Powering down the storage head/controller

Powering down or removing all SAS chains from a disk shelf will cause the controllers to panic to prevent data loss. To avoid this, shut down the controllers before decommissioning the shelf.

1. Log in to the BUI.
2. Click the Power icon on the left side of the masthead.

If the BUI is not accessible, select one of the following options:



Note

In a configuration with clustered storage heads, always shut down the standby head before the active head.

- SSH into the appliance and issue the `maintenance system poweroff` command.
- SSH or serial console into the service processor (SP) and issue the `stop /SYS` command.
- Use a pen or non-conducting pointed object to press and release the Power button on the front panel.



Caution

To initiate emergency shutdown during which all applications and files will be closed abruptly without saving, press and hold the power button for at least four

seconds until the Power/OK status indicator on the front panel flashes, indicating that the storage controller is in standby power mode.

Powering down the disk shelf

Do not remove a component if you do not have an immediate replacement. The disk shelf must not be operated without all components in place. Powering down or removing all SAS chains from a disk shelf will cause the controllers to panic to prevent data loss. To avoid this, shut down the controllers before decommissioning the shelf.

1. Stop all input and output to and from the disk shelf.
2. Wait approximately two minutes until all disk activity indicators have stopped flashing.
3. Place the power supply on/off switches to the "O" off position.
4. Disconnect the power cords from the external power source.

Powering on the storage appliance

The disk shelf must not be operated without all components in place.

1. Reconnect the disk shelf power and data cables you removed to service a component.
2. Place the power supply on/off switches on the disk shelf to the "I" on position.
3. Wait several minutes until the boot process is complete, at which time the Power LED should be solid green.
4. Connect the storage head power cables and wait approximately two minutes until the Power/OK LED on the front panel next to the Power button lights and remains lit.

6.14.2 Service Procedures for Sun ZFS Storage Appliance 7320 Components

For parts that are not hot-swappable, power down the Sun ZFS Storage Appliance 7320 before starting the service procedure.



Warning

If you need to execute a service procedure that interrupts the connection between virtual machines and their virtual disks, shut down the virtual machines in Oracle VM Manager prior to servicing the storage hardware. Disconnecting a running virtual machine from its disks may cause data corruption.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle Private Cloud Appliance. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.25 Service Procedures for Sun ZFS Storage Appliance 7320 Components

Replaceable Part(s)	Hot-Swap	URL
Storage head HDDs or SSDs	Yes	http://docs.oracle.com/cd/E28317_01/html/E38247/maintenance__hardware__details__7x20.html#maintenance__hardware__details__7x20
Disk shelf drives	Yes	Refer to the section "Replacing a Drive" on this page: http://docs.oracle.com/cd/E27998_01/html/E48492/maintenance__hardware__procedures__shelf.html#scrolltoc

Replaceable Part(s)	Hot-Swap	URL
Fan modules	Yes	http://docs.oracle.com/cd/E28317_01/html/E38247/maintenance__hardware__details__7x20.html#maintenance__har
Storage head power supplies	Yes	http://docs.oracle.com/cd/E28317_01/html/E38247/maintenance__hardware__details__7x20.html#maintenance__har
Disk shelf power supplies	Yes	Refer to the section “Replacing a Power Supply” on this page: http://docs.oracle.com/cd/E27998_01/html/E48492/maintenance__hardware__procedures__shelf.html#scrolltoc
Memory modules (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E28317_01/html/E38247/maintenance__hardware__details__7x20.html#maintenance__har
PCI Express risers and cards (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E28317_01/html/E38247/maintenance__hardware__details__7x20.html#maintenance__har
Battery	No	http://docs.oracle.com/cd/E28317_01/html/E38247/maintenance__hardware__details__7x20.html#maintenance__har
System indicator boards (Oracle-qualified service technician only)	Yes	http://docs.oracle.com/cd/E26765_01/html/E26399/maintenance__hardware__details__7x20.html#maintenance__har
Disk shelf I/O modules (Oracle-qualified service technician only)	Yes	Refer to the section “Replacing an I/O Module” on this page: http://docs.oracle.com/cd/E27998_01/html/E48492/maintenance__hardware__procedures__shelf.html#scrolltoc
Disk shelf SIM boards (Oracle-qualified service technician only)	Yes	Refer to the section “Replacing a SIM Board” on this page: http://docs.oracle.com/cd/E27998_01/html/E48492/maintenance__hardware__procedures__shelf.html#scrolltoc

6.15 Servicing an Oracle Switch ES1-24

This section provides instructions to service replaceable components (CRUs/FRUs) in an Oracle Switch ES1-24. Before starting any service procedure, read and follow the guidelines in [Section 6.3, “Preparing Oracle Private Cloud Appliance for Service”](#).

6.15.1 Powering Down the Oracle Switch ES1-24 for Service (When Required)

If you need to execute a service procedure that requires the Oracle Switch ES1-24 to be powered down, follow these instructions:

Powering down the switch

1. To power down an individual power supply, remove its power cord.
2. To power down the switch, remove the power cords from both power supplies.

Returning the switch to operation

1. Reconnect the power cords to both power supplies.
2. Verify that the switch has power by checking the status LEDs.

The AC LED lights green to indicate the power supply is connected to line power. A moment later, the OK LED lights green to indicate the power supply is fully operational.

6.15.2 Service Procedures for Oracle Switch ES1-24 Components

For parts that are not hot-swappable, power down the Oracle Switch ES1-24 before starting the service procedure.



Warning

Internal Ethernet connectivity is affected while the component is out of service. Please take the necessary precautions.



Caution

When replacing the entire switch assembly, begin by saving the configuration from the existing component, so that you can restore the configuration after replacement.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle Private Cloud Appliance. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.26 Service Procedures for Oracle Switch ES1-24 Components

Replaceable Part(s)	Hot-Swap	URL
Power supplies	Yes	http://docs.oracle.com/cd/E39109_01/html/E39116/z40000349112.html#scrolltoc
Fan module	Yes	http://docs.oracle.com/cd/E39109_01/html/E39116/z40000369112.html#scrolltoc

6.16 Servicing an NM2-36P Sun Datacenter InfiniBand Expansion Switch

This section provides instructions to service replaceable components (CRUs/FRUs) in a NM2-36P Sun Datacenter InfiniBand Expansion Switch. Before starting any service procedure, read and follow the guidelines in [Section 6.3, "Preparing Oracle Private Cloud Appliance for Service"](#).

6.16.1 Powering Down the NM2-36P Sun Datacenter InfiniBand Expansion Switch for Service (When Required)

If you need to execute a service procedure that requires the NM2-36P Sun Datacenter InfiniBand Expansion Switch to be powered down, follow these instructions:

Powering down the switch

1. To power down an individual power supply, remove its power cord.
2. To power down the switch, remove the power cords from both power supplies.

Returning the switch to operation

1. Reconnect the power cords to both power supplies.

2. Verify that the switch has power by checking the status LEDs.

The AC LED lights green to indicate the power supply is connected to line power. A moment later, the OK LED lights green to indicate the power supply is fully operational.

6.16.2 Service Procedures for NM2-36P Sun Datacenter InfiniBand Expansion Switch Components

For parts that are not hot-swappable, power down the NM2-36P Sun Datacenter InfiniBand Expansion Switch before starting the service procedure.



Caution

InfiniBand connectivity may be affected while the component is out of service. Please take the necessary precautions.



Caution

When replacing the entire switch assembly, begin by saving the configuration from the existing component, so that you can restore the configuration after replacement.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle Private Cloud Appliance. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.27 Service Procedures for NM2-36P Sun Datacenter InfiniBand Expansion Switch Components

Replaceable Part(s)	Hot-Swap	URL
Power supplies	Yes	http://docs.oracle.com/cd/E26698_01/html/E26434/z40001f49112.html#scrolltoc
Fans	Yes	http://docs.oracle.com/cd/E26698_01/html/E26434/z40001f59112.html#scrolltoc
Data cables	Yes	http://docs.oracle.com/cd/E26698_01/html/E26434/z40001f69112.html#scrolltoc

6.17 Servicing an Oracle Fabric Interconnect F1-15

This section provides instructions to service replaceable components (CRUs/FRUs) in an Oracle Fabric Interconnect F1-15. Before starting any service procedure, read and follow the guidelines in [Section 6.3, "Preparing Oracle Private Cloud Appliance for Service"](#).

6.17.1 Powering Down the Oracle Fabric Interconnect F1-15 for Service (When Required)

If you need to execute a service procedure that requires the Fabric Interconnect to be powered down, follow these instructions:

Powering down the Oracle Fabric Interconnect F1-15

1. Press the Power button to power down the Fabric Interconnect gracefully.
2. Wait for the Status LED to switch off, indicating that the component has been powered down successfully.

Returning the Oracle Fabric Interconnect F1-15 to operation

1. Press the Power button to power on the Fabric Interconnect.

The Status LED blinks green, indicating that the system control processor is booting.

2. Wait until the Status LED is solid green.

This indicates that the system control processor has finished booting and the Fabric Interconnect is ready for operation.

6.17.2 Service Procedures for Oracle Fabric Interconnect F1-15 Components

For parts that are not hot-swappable, power down the Oracle Fabric Interconnect F1-15 before starting the service procedure.



Caution

Management, storage, VM and external network connectivity may be affected while the Fabric Interconnect or an I/O module is out of service. Please take the necessary precautions.



Caution

When replacing the entire switch assembly, begin by saving the configuration from the existing component, so that you can restore the configuration after replacement.

Generally speaking, hot-swappable components can be serviced without specific additional steps for Oracle Private Cloud Appliance. Follow the applicable procedure in the Service Manual. The following table provides links to each service procedure and indicates whether parts are hot-swappable or require the component to be taken offline and powered down.

Table 6.28 Service Procedures for Oracle Fabric Interconnect F1-15 Components

Replaceable Part(s)	Hot-Swap	URL
Power supplies (Oracle-qualified service technician only)	Yes	http://docs.oracle.com/cd/E38500_01/html/E50997/z40004411020156.html#scrolltoc
Fan modules (Oracle-qualified service technician only)	Yes	http://docs.oracle.com/cd/E38500_01/html/E50997/z40004411020136.html#scrolltoc
Fabric board (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E38500_01/html/E50997/z40004411020657.html#scrolltoc
Management module (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E38500_01/html/E50997/z40004411020369.html#scrolltoc http://docs.oracle.com/cd/E38500_01/html/E50997/z40004411020375.html#scrolltoc
I/O modules	Yes	http://docs.oracle.com/cd/E38500_01/html/E50997/z40004411020323.html#scrolltoc

Replaceable Part(s)	Hot-Swap	URL
(Oracle-qualified service technician only)		http://docs.oracle.com/cd/E38500_01/html/E50997/z400037d1022426.html#scrolltoc
Front panel assembly, including system control processor	No	http://docs.oracle.com/cd/E38500_01/html/E50997/z40004411020496.html#scrolltoc
(Oracle-qualified service technician only)		

Chapter 7 Troubleshooting

Table of Contents

7.1 Setting the Oracle Private Cloud Appliance Logging Parameters	215
7.2 Adding Proxy Settings for Oracle Private Cloud Appliance Updates	216
7.3 Configuring Data Center Switches for VLAN Traffic	217
7.4 Changing the Oracle VM Agent Password	218
7.5 Running Manual Pre- and Post-Upgrade Checks in Combination with Oracle Private Cloud Appliance Upgrader	218
7.6 Enabling Fibre Channel Connectivity on a Provisioned Appliance	220
7.7 Restoring a Backup After a Password Change	222
7.8 Enabling SNMP Server Monitoring	224
7.9 Using a Custom CA Certificate for SSL Encryption	225
7.9.1 Creating a Keystore	226
7.9.2 Importing a Keystore	227
7.10 Reprovisioning a Compute Node when Provisioning Fails	228
7.11 Deprovisioning and Replacing a Compute Node	229
7.12 Eliminating Time-Out Issues when Provisioning Compute Nodes	230
7.13 Returning Oracle VM Server Pool to Operation After Network Services Restart	231
7.14 Recovering from Tenant Group Configuration Mismatches	232
7.15 Configure Xen CPU Frequency Scaling for Best Performance	233

This chapter describes how to resolve a number of common problem scenarios.

7.1 Setting the Oracle Private Cloud Appliance Logging Parameters

When troubleshooting or if you have a support query open, you may be required to change the logging parameters for your Oracle Private Cloud Appliance. The settings for this are contained in `/etc/ovca.conf`, and can be changed using the CLI.

The following instructions must be followed for each of the two management nodes in your environment.

Changing the Oracle Private Cloud Appliance Logging Parameters for a Management Node

1. Gain command line access to the management node. Usually this is achieved using SSH and logging in as the root user with the global Oracle Private Cloud Appliance password.
2. Use the CLI, as described in [Chapter 4, The Oracle Private Cloud Appliance Command Line Interface \(CLI\)](#), to view or modify your appliance log settings. The CLI safely reads and edits the `/etc/ovca.conf` file, to prevent the possibility of configuration file corruption.

- To view the current values for the configurable settings in the configuration file run the CLI as follows:

```
# pca-admin show system-properties
```

- To change the log level:

```
# pca-admin set system-property log_level service LEVEL
```

The `service` argument is the log file category to which the new log level applies. The following services can be specified: backup, cli diagnosis, monitor, ovca, snmp, syncservice.

The `LEVEL` value is one of the following: `DEBUG`, `INFO`, `WARNING`, `ERROR`, `CRITICAL`.

- To change the log file size:

```
# pca-admin set system-property log_size SIZE
```

Where *SIZE*, expressed in MB, is a number from 1 to 512.

- To change the number of backup log files stored:

```
# pca-admin set system-property log_count COUNT
```

Where *COUNT* is a number of files ranging from 0 to 100.

- To change the location where log files are stored:

```
# pca-admin set system-property log_file service PATH
```

Where *PATH* is the new location where the log file for the selected *service* is to be stored. The following services can be specified: backup, cli, diagnosis, monitor, ovca, snmp, and syncservice.



Caution

Make sure that the new path to the log file exists. Otherwise, the log server stops working.

The system always prepends `/var/log` to your entry. Absolute paths are converted to `/var/log/PATH`.

During management node upgrades, the log file paths are reset to the default values.

3. The new log level setting only takes effect after a management node has been rebooted or the service has been restarted by running the `service ovca restart` command on the active management node shell.

7.2 Adding Proxy Settings for Oracle Private Cloud Appliance Updates

If your data center does not provide unlimited internet access and has a proxy server in place to control HTTP, HTTPS or FTP traffic, you may need to configure your management nodes to be able to access external resources; for example for the purpose of performing software updates.

The following instructions must be followed for each of the two management nodes in your environment.

Adding Proxy Settings for a Management Node

1. Gain command line access to the management node. Usually this is achieved using SSH and logging in as the root user with the global Oracle Private Cloud Appliance password.
2. Use the CLI, as described in [Chapter 4, The Oracle Private Cloud Appliance Command Line Interface \(CLI\)](#), to view or modify your proxy settings. The CLI safely reads and edits the `/etc/ovca.conf` file, to prevent the possibility of configuration file corruption.
 - To view the current values for the configurable settings in the configuration file run the CLI as follows:

```
# pca-admin show system-properties
```

- To set an HTTP proxy:

```
# pca-admin set system-property http_proxy http://IP:PORT
```

Where *IP* is the IP address of your proxy server, and *PORT* is the TCP port on which it is listening.



Caution

If your proxy server expects a user name and password, these should be provided when the proxy service is accessed. Do not specify credentials as part of the proxy URL, because this implies that you send sensitive information over a connection that is not secure.

- To set an HTTPS proxy:

```
# pca-admin set system-property https_proxy https://IP:PORT
```

- To set an FTP proxy:

```
# pca-admin set system-property ftp_proxy ftp://IP:PORT
```

3. Setting any single parameter automatically rewrites the configuration file and the proxy settings become active immediately.

7.3 Configuring Data Center Switches for VLAN Traffic



Warning

This section applies **only** to systems with an InfiniBand-based network architecture. The configuration described in this section is valid for the outbound connections through the Oracle Fabric Interconnect F1-15s .

The Oracle Private Cloud Appliance network infrastructure supports the use of VLANs by default. For this purpose, the Oracle Fabric Interconnect F1-15s are set to trunking mode to allow tagged data traffic.



Caution

Do not configure any type of link aggregation group (LAG) across the 10GbE ports: LACP, network/interface bonding or similar methods to combine multiple network connections are not supported.

To provide additional bandwidth to the environment hosted by the Oracle Private Cloud Appliance, create custom networks. For detailed information, see [Section 2.6, “Network Customization”](#).

You may implement VLANs for logical separation of different network segments, or to define security boundaries between networks with different applications – just as you would with physical servers instead of virtual machines.

However, to allow virtual machines hosted by the Oracle Private Cloud Appliance to communicate with systems external to the appliance, you must update the configuration of your next-level data center switches accordingly.

- The switch ports on the receiving end of the outbound appliance connections must be part of each VLAN used within the Oracle Private Cloud Appliance environment.
- The same ports must also be part of the network(s) connecting the external systems that your virtual machines need to access. For example, WAN connectivity implies that virtual machines are able to

reach the public gateway in your data center. As an alternative to VLAN tagging, Layer 3 routing can be used to connect to the Oracle Private Cloud Appliance.

7.4 Changing the Oracle VM Agent Password

The password of the Oracle VM Agent cannot be modified in the Authentication tab of the Oracle Private Cloud Appliance Dashboard, nor with the `update password` command of the Oracle Private Cloud Appliance CLI. If you need to change the agent password, use Oracle VM Manager.

Instructions to change the Oracle VM Agent password can be found at the following location: [Change Oracle VM Agent Passwords on Oracle VM Servers](#) in the *Oracle VM Manager User's Guide for Release 3.4*.

7.5 Running Manual Pre- and Post-Upgrade Checks in Combination with Oracle Private Cloud Appliance Upgrader

Controller software updates must be installed using the Oracle Private Cloud Appliance Upgrader. While the Upgrader tool automates a large number of prerequisite checks, there are still some tasks that must be performed manually before and after the upgrade process. The manual tasks are listed in this section. For more detailed information, please refer to the support note with [Doc ID 2442664.1](#) for Controller Software release 2.3.4, or support note [Doc ID 2605884.1](#) for Controller Software release 2.4.2.

Start by running the Oracle Private Cloud Appliance Upgrader in verify-only mode. The steps are described in [Section 3.2.3, "Verifying Upgrade Readiness"](#). Fix any issues reported by the Upgrader and repeat the verification procedure until all checks complete without errors. Then, proceed to the manual pre-upgrade checks.

Performing Manual Pre-Upgrade Checks

1. Verify the WebLogic password.

On the master Management Node, run the following commands:

```
# cd /u01/app/oracle/ovm-manager-3/bin
# ./ovm_admin --listusers
```

Enter the WebLogic password when prompted. If the password is incorrect, the `ovm_admin` command fails and exits with return code `1`. If the password is correct, the command lists the users and exits with return code of `0`. In the event of an incorrect password, login to the Oracle Private Cloud Appliance web interface and change the `wls-weblogic` password to the expected password.

2. Check that no external storage LUNs are connected to the management nodes.

Verify that none of your external storage LUNs are visible from either management node. For more details, refer to the support note with [Doc ID 2148589.1](#).

If your system is InfiniBand-based and there are no Fibre Channel cards installed in the Fabric Interconnects, you can skip this check.

3. Check for customized `inet` settings on the management nodes.

Depending on the exact upgrade path you are following, `xinetd` may be upgraded. In this case, modified settings are automatically reset to default. Make a note of your custom `inet` settings and verify them after the upgrade process has completed. These setting changes are stored in the file `/etc/postfix/main.cf`.

4. Register the number of objects in the MySQL database.

As the root user on the master management node, download and run the script `number_of_jobs_and_objects.sh`. It is attached to the support note with [Doc ID 2442664.1](#) for Controller Software release 2.3.4, or support note [Doc ID 2605884.1](#) for Controller Software release 2.4.2. It returns the number of objects and the number of jobs in the database. Make a note of these numbers.

5. Verify management node failover.

Reboot the master management node to ensure that the standby management node is capable of taking over the master role.

6. Check the NFS protocol used for the internal ZFS Storage Appliance.

On both management nodes, run the command `nfsstat -m`. Each mounted share should use the NFSv4 protocol.

7. Check the file `/etc/yum.conf` on both management nodes.

If a proxy is configured for YUM, comment out or remove that line from the file.

When you have submitted your system to all pre-upgrade checks and you have verified that it is ready for upgrade, execute the controller software update. The steps are described in [Section 3.2.4, "Executing a Controller Software Update"](#). After successfully upgrading the controller software, proceed to the manual post-upgrade checks for management nodes and compute nodes.

Performing Manual Post-Upgrade Checks on the Management Nodes

1. Check the names of the Unmanaged Storage Arrays.

If the names of the Unmanaged Storage Arrays are no longer displayed correctly after the upgrade, follow the workaround documented in the support note with [Doc ID 2244130.1](#).

2. Check for errors and warnings in Oracle VM.

In the Oracle VM Manager web UI, verify that none of these occur:

- Padlock icons against compute nodes or storage servers
- Red error icons against compute nodes, repositories or storage servers
- Yellow warning icons against compute nodes, repositories or storage servers

3. Check the status of all components in the Oracle Private Cloud Appliance Dashboard.

Verify that a green check mark appears to the right of each hardware component in the Hardware View, and that no red error icons are present.

4. Check networks.

Verify that all networks – factory default and custom – are present and correctly configured.

Performing Manual Post-Upgrade Checks on the Compute Nodes

1. Change the `min_free_kbytes` setting on all compute nodes.

Refer to the support note with [Doc ID 2314504.1](#). Apply the corresponding steps and reboot the compute node after the change has been made permanent.

2. Check that the `fm` package is installed on all compute nodes.

Run the command `rpm -q fm`. If the package is not installed, run the following command:

```
# chkconfig ipmi on; service ipmi start; LFMA_UPDATE=1 /usr/bin/yum install fm -q -y --nogpgcheck
```

3. Perform a virtual machine test.

Start a test virtual machine and verify that networks are functioning. Migrate the virtual machine to a compatible compute node to make sure that live migration works correctly.

7.6 Enabling Fibre Channel Connectivity on a Provisioned Appliance



Warning

This section applies **only** to systems with an InfiniBand-based network architecture. The configuration described in this section is valid for the I/O modules in the Oracle Fabric Interconnect F1-15s .

However, for Oracle Server X8-2 and newer compute nodes, Fibre Channel connectivity through the Fabric Interconnects is **not** supported. Instead, you must use the (optional) physical FC HBA expansion cards. Refer to the section [Extending Storage Capacity of Ethernet-based Systems](#) in the *Oracle Private Cloud Appliance Installation Guide*.

If you ordered an Oracle Private Cloud Appliance without factory-installed Fibre Channel I/O modules and you decide to add external Fibre Channel storage at a later time, when the rack has already been provisioned, your installation must meet these requirements:

- The Oracle Private Cloud Appliance controller software must be at Release 2.1.1 or later.
- A total of four Fibre Channel I/O modules must be installed in slots 3 and 12 of each Oracle Fabric Interconnect F1-15.
- Storage clouds and vHBAs must be configured manually.

Installation information for the optional Fibre Channel I/O modules can be found in the section entitled [Installing Optional Fibre Channel I/O Modules](#) in the *Oracle Private Cloud Appliance Installation Guide*. This section provides detailed CLI instructions to configure the storage clouds and vHBAs associated with Fibre Channel connectivity.

Configuring Storage Clouds and vHBAs for Fibre Channel Connectivity

1. Using SSH and an account with superuser privileges, log into the master management node.



Note

The data center IP address used in this procedure is an example.

```
# ssh root@10.100.1.101
root@10.100.1.101's password:
[root@ovcamn05r1 ~]#
```

2. Launch the Oracle Private Cloud Appliance CLI in interactive mode.

```
# pca-admin
Welcome to PCA! Release: 2.3.2
PCA>
```

3. Verify that no storage clouds or vHBAs exist yet.

```
PCA> list storage-network
```

```
Network_Name          Description
-----
0 rows displayed
Status: Success
```

```
PCA> list wwpn-info
```

```
WWPN                  vHBA      Cloud_Name      Server      Type  Alias
-----
0 rows displayed
Status: Success
```

4. Configure the vHBAs on both management nodes.

```
PCA> configure vhbases ovcamn05r1 ovcamn06r1
```

```
Compute_Node          Status
-----
ovcamn05r1            Succeeded
ovcamn06r1            Succeeded
-----
2 rows displayed
Status: Success
```

5. Verify that the clouds have been configured.

```
PCA> list storage-network
```

```
Network_Name          Description
-----
Cloud_A               Default Storage Cloud ru22 port1 - Do not delete or modify
Cloud_B               Default Storage Cloud ru22 port2 - Do not delete or modify
Cloud_C               Default Storage Cloud ru15 port1 - Do not delete or modify
Cloud_D               Default Storage Cloud ru15 port2 - Do not delete or modify
-----
4 rows displayed
Status: Success
```

6. If the 4 storage clouds have been configured correctly, configure the vHBAs on all compute nodes.

```
PCA> configure vhbases ALL
```

```
Compute_Node          Status
-----
ovcacn07r1            Succeeded
ovcacn08r1            Succeeded
[...]
ovcacn36r1            Succeeded
ovcacn37r1            Succeeded
-----
20 rows displayed
Status: Success
```

7. Verify that all clouds and vHBAs have been configured correctly.

```
PCA> list wwpn-info
```

```
WWPN                  vHBA      Cloud_Name      Server      Type  Alias
-----
```

```

50:01:39:70:00:4F:91:00  vhba01  Cloud_A      ovcamn05r1  MN   ovcamn05r1-Cloud_A
50:01:39:70:00:4F:91:02  vhba01  Cloud_A      ovcamn06r1  MN   ovcamn06r1-Cloud_A
50:01:39:70:00:4F:91:04  vhba01  Cloud_A      ovcacn07r1  CN   ovcacn07r1-Cloud_A
50:01:39:70:00:4F:91:06  vhba01  Cloud_A      ovcacn08r1  CN   ovcacn08r1-Cloud_A
[...]
50:01:39:70:00:4F:F1:05  vhba04  Cloud_D      ovcacn35r1  CN   ovcacn35r1-Cloud_D
50:01:39:70:00:4F:F1:03  vhba04  Cloud_D      ovcacn36r1  CN   ovcacn36r1-Cloud_D
50:01:39:70:00:4F:F1:01  vhba04  Cloud_D      ovcacn37r1  CN   ovcacn37r1-Cloud_D
-----
88 rows displayed

Status: Success

```

```

PCA> show storage-network Cloud_A

-----
Network_Name      Cloud_A
Description       Default Storage Cloud ru22 port1 - Do not delete or modify
Ports             ovcasw22r1:12:1, ovcasw22r1:3:1
vHBAs             ovcacn07r1-vhba01, ovcacn08r1-vhba01, ovcacn10r1-vhba01, [...]
-----
Status: Success

PCA> show storage-network Cloud_B

-----
Network_Name      Cloud_B
Description       Default Storage Cloud ru22 port2 - Do not delete or modify
Ports             ovcasw22r1:12:2, ovcasw22r1:3:2
vHBAs             ovcacn07r1-vhba02, ovcacn08r1-vhba02, ovcacn10r1-vhba02, [...]
-----
Status: Success

PCA> show storage-network Cloud_C

-----
Network_Name      Cloud_C
Description       Default Storage Cloud ru15 port1 - Do not delete or modify
Ports             ovcasw15r1:12:1, ovcasw15r1:3:1
vHBAs             ovcacn07r1-vhba03, ovcacn08r1-vhba03, ovcacn10r1-vhba03, [...]
-----
Status: Success

PCA> show storage-network Cloud_D

-----
Network_Name      Cloud_D
Description       Default Storage Cloud ru15 port2 - Do not delete or modify
Ports             ovcasw15r1:12:2, ovcasw15r1:3:2
vHBAs             ovcacn07r1-vhba04, ovcacn08r1-vhba04, ovcacn10r1-vhba04, [...]
-----
Status: Success

```

The system is now ready to integrate with external Fibre Channel storage. For detailed information and instructions, refer to the section entitled [Adding External Fibre Channel Storage](#) in the *Oracle Private Cloud Appliance Installation Guide*.

7.7 Restoring a Backup After a Password Change

If you have changed the password for Oracle VM Manager or its related components Oracle WebLogic Server and Oracle MySQL database, and you need to restore the Oracle VM Manager from a backup that was made prior to the password change, the passwords will be out of sync. As a result of this password mismatch, Oracle VM Manager cannot connect to its database and cannot be started, so you must first make sure that the passwords are identical.

**Note**

The steps below are not specific to the case where a password change occurred after the backup. They apply to any restore operation.

As of Release 2.3.1, which includes Oracle VM Manager 3.4.2, the database data directory cleanup is built into the restore process, so that step can be skipped.

Resolving Password Mismatches when Restoring Oracle VM Manager from a Backup

1. Create a manual backup of the Oracle VM Manager MySQL database to prevent inadvertent data loss. On the command line of the active management node, run the following command:

- Release 2.2.x and older:

```
# /u01/app/oracle/ovm-manager-3/bin/createBackup.sh -n ManualBackup1
```

- Release 2.3.1 and newer:

```
# /u01/app/oracle/ovm-manager-3/ovm_tools/bin/BackupDatabase -w
INFO: Backup started to:
      /u01/app/oracle/mysql/dbbackup/ManualBackup-20190524_102412
```

2. In the Oracle Private Cloud Appliance Dashboard, change the Oracle MySQL database password back to what it was at the time of the backup.
3. On the command line of the active management node, as `root` user, stop the Oracle VM Manager and MySQL services, and then delete the MySQL data.

```
# service ovmm stop
# service ovmm_mysql stop
# cd /u01/app/oracle/mysql/data
# rm -rf appfw ibdata ib_logfile* mysql mysqld.err ovs performance_schema
```

4. As `oracle` user, restore the database from the selected backup.

- Release 2.2.x and older:

```
# su oracle
$ bash /u01/app/oracle/ovm-manager-3/ovm_shell/tools/RestoreDatabase.sh BackupToBeRestored
INFO: Expanding the backup image...
INFO: Applying logs to the backup snapshot...
INFO: Restoring the backup...
INFO: Success - Done!
INFO: Log of operations performed is available at:
      /u01/app/oracle/mysql/dbbackup/BackupToBeRestored/Restore.log
```

- Release 2.3.1 and newer:

```
# su oracle
$ bash /u01/app/oracle/ovm-manager-3/ovm_tools/bin/RestoreDatabase.sh BackupToBeRestored
INFO: Expanding the backup image...
INFO: Applying logs to the backup snapshot...
INFO: Restoring the backup...
INFO: Success - Done!
INFO: Log of operations performed is available at:
      /u01/app/oracle/mysql/dbbackup/BackupToBeRestored/Restore.log
```

5. As `root` user, start the MySQL and Oracle VM Manager services.

```
$ su root
# service ovmm_mysql start
# service ovmm start
```

After both services have restarted successfully, the restore operation is complete.

7.8 Enabling SNMP Server Monitoring

For troubleshooting or hardware monitoring, it may be useful to enable SNMP on the servers in your Oracle Private Cloud Appliance. While the tools for SNMP are available, the protocol is not enabled by default. This section explains how to enable SNMP with the standard Oracle Linux and additional Oracle Private Cloud Appliance Management Information Bases (MIBs).

Enabling SNMP on the Management Nodes

1. Using SSH and an account with superuser privileges, log into the management node.



Note

The data center IP address used in this procedure is an example.

```
# ssh root@10.100.1.101
root@10.100.1.101's password:
[root@ovcamn05r1 ~]#
```

2. Locate the necessary `rpm` packages in the mounted directory `/nfs/shared_storage/mgmt_image/Packages`, which resides in the `MGMT_ROOT` file system on the ZFS storage appliance. The following packages are part of the Oracle Private Cloud Appliance ISO image:

- `net-snmp-5.5-60.0.1.el6.x86_64.rpm`
- `net-snmp-libs-5.5-60.0.1.el6.x86_64.rpm`
- `net-snmp-utils-5.5-60.0.1.el6.x86_64.rpm`
- `ovca-snmp-0.9-3.el6.x86_64.rpm`
- `lm_sensors-libs-3.1.1-17.el6.x86_64.rpm`

3. Install these packages by running the following command:

```
# rpm -ivh ovca-snmp-0.9-3.el6.x86_64.rpm net-snmp-libs-5.5-49.0.1.el6.x86_64.rpm \
net-snmp-5.5-49.0.1.el6.x86_64.rpm lm_sensors-libs-3.1.1-17.el6.x86_64.rpm \
net-snmp-utils-5.5-49.0.1.el6.x86_64.rpm
```

4. Create an SNMP configuration file: `/etc/snmp/snmpd.conf`.

This is a standard sample configuration:

```
rocommunity public
syslocation MyDataCenter
dlmod ovca /usr/lib64/ovca-snmp/ovca.so
```

5. Enable the `snmpd` service.

```
# service snmpd start
```

6. If desired, enable the `snmpd` service on boot.

```
# chkconfig snmpd on
```

7. Open the SNMP ports on the firewall.

```
# iptables -I INPUT -p udp -m udp --dport 161 -j ACCEPT
```

```
# iptables -I INPUT -p udp -m udp --dport 162 -j ACCEPT
# iptables-save > /etc/sysconfig/iptables
```

SNMP is now ready for use on this management node. Besides the standard Oracle Linux MIBs, these are also available:

- ORACLE-OVCA-MIB::ovcaVersion
- ORACLE-OVCA-MIB::ovcaSerial
- ORACLE-OVCA-MIB::ovcaType
- ORACLE-OVCA-MIB::ovcaStatus
- ORACLE-OVCA-MIB::nodeTable

Usage examples:

```
# snmpwalk -v 1 -c public -O e 130.35.70.186 ORACLE-OVCA-MIB::ovcaVersion
# snmpwalk -v 1 -c public -O e 130.35.70.111 ORACLE-OVCA-MIB::ovcaStatus
# snmpwalk -v 1 -c public -O e 130.35.70.111 ORACLE-OVCA-MIB::nodeTable
```

8. Repeat this procedure on the second management node.

Enabling SNMP on the Compute Nodes



Note

On Oracle Private Cloud Appliance compute nodes, `net-snmp`, `net-snmp-utils` and `net-snmp-libs` are already installed at the factory, but the SNMP service is not enabled or configured.

1. Using SSH and an account with superuser privileges, log into the compute node. It can be accessed through the appliance internal management network.

```
ssh root@192.168.4.5
root@192.168.4.5's password:
[root@ovcacn27r1 ~]#
```

2. Create an SNMP configuration file: `/etc/snmp/snmpd.conf` and make sure this line is included:

```
rocommunity public
```

3. Enable the `snmpd` service.

```
# service snmpd start
```

SNMP is now ready for use on this compute node.

4. If desired, enable the `snmpd` service on boot.

```
# chkconfig snmpd on
```

5. Repeat this procedure on all other compute nodes installed in your Oracle Private Cloud Appliance environment.

7.9 Using a Custom CA Certificate for SSL Encryption

By default, Oracle Private Cloud Appliance and Oracle VM Manager use a self-signed SSL certificate for authentication. While it serves to provide SSL encryption for all HTTP traffic, it is recommended that

you obtain and install your own custom trusted certificate from a well-known and recognized Certificate Authority (CA).

Both the Oracle Private Cloud Appliance Dashboard and the Oracle VM Manager web interface run on Oracle WebLogic Server. The functionality to update the digital certificate and keystore is provided by the Oracle VM Key Tool in conjunction with the Java Keytool in the JDK. The tools are installed on the Oracle Private Cloud Appliance management nodes.

7.9.1 Creating a Keystore

If you do not already have a third-party CA certificate, you can create a new keystore. The keystore you create contains one entry for a private key. After you create the keystore, you generate a certificate signing request (CSR) for that private key and submit the CSR to a third-party CA. The CA then signs the CSR and returns a signed SSL certificate and a copy of the CA certificate, which you then import into your keystore.

Creating a Keystore with a Custom CA Certificate

1. Using SSH and an account with superuser privileges, log into the management node.



Note

The data center IP address used in this procedure is an example.

```
# ssh root@10.100.1.101
root@10.100.1.101's password:
[root@ovcamn05r1 ~]#
```

2. Go to the security directory of the Oracle VM Manager WebLogic domain.

```
# cd /u01/app/oracle/ovm-manager-3/domains/ovm_domain/security
```

3. Create a new keystore. Transfer ownership to user *oracle* in the user group *dba*.

```
# /u01/app/oracle/java/bin/keytool -genkeypair -alias ca -keyalg RSA -keysize 2048 \
-keypass Welcome1 -storetype jks -keystore mykeystore.jks -storepass Welcome1
# chown oracle.dba mykeystore.jks
```

4. Generate a certificate signing request (CSR). Transfer ownership to user *oracle* in the user group *dba*.

```
# /u01/app/oracle/java/bin/keytool -certreq -alias ca -file pcakey.csr \
-keypass Welcome1 -storetype jks -keystore mykeystore.jks -storepass Welcome1
# chown oracle.dba pcakey.csr
```

5. Submit the CSR file to the relevant third-party CA for signing.

6. For the signed files returned by the CA, transfer ownership to user *oracle* in the user group *dba*.

```
# chown oracle.dba ca_cert_file
# chown oracle.dba ssl_cert_file
```

7. Import the signed CA certificate into the keystore.

```
# /u01/app/oracle/java/bin/keytool -importcert -trustcacerts -noprompt -alias ca \
-file ca_cert_file -storetype jks -keystore mykeystore.jks -storepass Welcome1
```

8. Import the signed SSL certificate into the keystore.

```
# /u01/app/oracle/java/bin/keytool -importcert -trustcacerts -noprompt -alias ca \
-file ssl_cert_file -keypass Welcome1 -storetype jks -keystore mykeystore.jks \
-storepass Welcome1
```

- Use the **setsslkey** command to configure the system to use the new keystore.

```
# /u01/app/oracle/ovm-manager-3/ovm_upgrade/bin/ovmkeytool.sh setsslkey
Path for SSL keystore: /u01/app/oracle/ovm-manager-3/domains/ovm_domain/security/mykeystore.jks
Keystore password:
Alias of key to use as SSL key: ca
Key password:
Updating keystore information in WebLogic
Oracle MiddleWare Home (MW_HOME): [/u01/app/oracle/Middleware]
WebLogic domain directory: [/u01/app/oracle/ovm-manager-3/domains/ovm_domain]
Oracle WebLogic Server name: [AdminServer]
WebLogic username: [weblogic]
WebLogic password: [*****]
WLST session logged at: /tmp/wlst-session5820685079094897641.log
```

- Configure the client certificate login.

```
# /u01/app/oracle/ovm-manager-3/bin/configure_client_cert_login.sh \
/u01/app/oracle/ovm-manager-3/domains/ovm_domain/security/pcakey.crt
```

- Test the new SSL configuration by logging into the Oracle Private Cloud Appliance Dashboard. From there, proceed to Oracle VM Manager with the button "Login to OVM Manager". The browser now indicates that your connection is secure.

7.9.2 Importing a Keystore

If you already have a CA certificate and SSL certificate, use the SSL certificate to create a keystore. You can then import that keystore into Oracle Private Cloud Appliance and configure it as the SSL keystore.



Caution

If you have generated custom keys using `ovmkeytool.sh` in a previous version of the Oracle Private Cloud Appliance software, you must regenerate the keys prior to updating the Controller Software. For instructions, refer to the support note with [Doc ID 2597439.1](#).

Importing a Keystore with an Existing CA and SSL Certificate

- Using SSH and an account with superuser privileges, log into the management node.



Note

The data center IP address used in this procedure is an example.

```
# ssh root@10.100.1.101
root@10.100.1.101's password:
[root@ovcamn05r1 ~]#
```

- Import the keystore.

```
# /u01/app/oracle/java/bin/keytool -importkeystore -noprompt \
-srckeystore existing_keystore.jks -srcstoretype source_format -srcstorepass Welcome1
-destkeystore mykeystore.jks -deststoretype jks -deststorepass Welcome1
```

- Use the **setsslkey** command to configure the system to use the new keystore.

```
# /u01/app/oracle/ovm-manager-3/ovm_upgrade/bin/ovmkeytool.sh setsslkey
Path for SSL keystore: /u01/app/oracle/ovm-manager-3/domains/ovm_domain/security/mykeystore.jks
Keystore password:
Alias of key to use as SSL key: ca
Key password:
Updating keystore information in WebLogic
```

```
Oracle MiddleWare Home (MW_HOME): [/u01/app/oracle/Middleware]
WebLogic domain directory: [/u01/app/oracle/ovm-manager-3/domains/ovm_domain]
Oracle WebLogic Server name: [AdminServer]
WebLogic username: [weblogic]
WebLogic password: [*****]
WLST session logged at: /tmp/wlst-session5820685079094897641.log
```

4. Configure the client certificate login.

```
# /u01/app/oracle/ovm-manager-3/bin/configure_client_cert_login.sh /path/to/cacert
```

Where `/path/to/cacert` is the absolute path to the CA certificate.

5. Test the new SSL configuration by logging into the Oracle Private Cloud Appliance Dashboard. From there, proceed to Oracle VM Manager with the button "Login to OVM Manager". The browser now indicates that your connection is secure.

7.10 Reprovisioning a Compute Node when Provisioning Fails

Compute node provisioning is a complex orchestrated process involving various configuration and installation steps and several reboots. Due to connectivity fluctuations, timing issues or other unexpected events, a compute node may become stuck in an intermittent state or go into error status. The solution is to re provision the compute node.



Warning

Reprovisioning is to be applied *only* to compute nodes that fail to complete provisioning.


For correctly provisioned and running compute nodes, reprovisioning functionality is blocked in order to prevent incorrect use that could lock compute nodes out of the environment permanently or otherwise cause loss of functionality or data corruption.

Reprovisioning a Compute Node when Provisioning Fails

1. Log in to the Oracle Private Cloud Appliance Dashboard.
2. Go to the **Hardware View** tab.
3. Roll over the compute nodes that are in Error status or have become stuck in the provisioning process.

A pop-up window displays a summary of configuration and status information.

Figure 7.1 Compute Node Information and Reprovision Button in Hardware View



4. If the compute node provisioning is incomplete and the server is in error status or stuck in an intermittent state for several hours, click the **Reprovision** button in the pop-up window.
5. When the confirmation dialog box appears, click OK to start reprovisioning the compute node.

If compute node provisioning should fail after the server was added to the Oracle VM server pool, additional recovery steps could be required. The cleanup mechanism associated with reprovisioning may be unable to remove the compute node from the Oracle VM configuration. For example, when a server is in locked state or owns the server pool master role, it must be unconfigured manually. In this case you need to perform operations in Oracle VM Manager that are otherwise not permitted. You may also need to power on the compute node manually.

Removing a Compute Node from the Oracle VM Configuration

1. Log into the Oracle VM Manager user interface.

For detailed instructions, see [Section 5.2, “Logging in to the Oracle VM Manager Web UI”](#).

2. Go to the **Servers and VMs** tab and verify that the server pool named `Rack1_ServerPool` does indeed contain the compute node that fails to provision correctly.
3. If the compute node is locked due to a running job, abort it in the **Jobs** tab of Oracle VM Manager.

Detailed information about the use of jobs in Oracle VM can be found in the [Oracle VM Manager User's Guide](#). Refer to the section entitled [Jobs Tab](#).

4. Remove the compute node from the Oracle VM server pool.

Refer to the section entitled [Edit Server Pool](#) in the Oracle VM Manager User's Guide. When editing the server pool, move the compute node out of the list of selected servers. The compute node is moved to the Unassigned Servers folder.

5. Delete the compute node from Oracle VM Manager.

Refer to the [Oracle VM Manager User's Guide](#) and follow the instructions in the section entitled [Delete Server](#).

When the failing compute node has been removed from the Oracle VM configuration, return to the Oracle Private Cloud Appliance Dashboard, to reprovision it. If the compute node is powered off and reprovisioning cannot be started, power on the server manually.

7.11 Deprovisioning and Replacing a Compute Node

When a defective compute node needs to be replaced or repaired, or when a compute node is retired in favor of a newer model with higher capacity and better performance, it is highly recommended that you deprovision the compute node before removing it from the appliance rack. Deprovisioning ensures that all configuration entries for a compute node are removed cleanly, so that no conflicts are introduced when a replacement compute node is installed.

Deprovisioning a Compute Node for Repair or Replacement

1. Log into the Oracle VM Manager user interface.

For detailed instructions, see [Section 5.2, “Logging in to the Oracle VM Manager Web UI”](#).

2. Migrate all virtual machines away from the compute node you wish to deprovision. If any VMs are running on the compute node, the deprovision command fails.

- Using SSH and an account with superuser privileges, log into the active management node, then launch the Oracle Private Cloud Appliance command line interface.

```
# ssh root@10.100.1.101
root@10.100.1.101's password:
root@ovcamn05r1 ~]# pca-admin
Welcome to PCA! Release: 2.4.2
PCA>
```

- Lock provisioning to make sure that the compute node cannot be reprovisioned immediately after deprovisioning.

```
PCA> create lock provisioning
Status: Success
```

- Deprovision the compute node you wish to remove. Repeat for additional compute nodes, if necessary.

```
PCA> deprovision compute-node ovcacn29r1
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Shutting down dhcpd: [ OK ]
Starting dhcpd: [ OK ]
Shutting down dnsmasq: [ OK ]
Starting dnsmasq: [ OK ]

Status: Success
```

- When the necessary compute nodes have been deprovisioned successfully, release the provisioning lock. The appliance resumes its normal operation.

```
PCA> delete lock provisioning
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Success
```

When the necessary repairs have been completed, or when the replacement compute nodes are ready, install the compute nodes into the rack and connect the necessary cables. The controller software detects the new compute nodes and automatically launches the provisioning process.

7.12 Eliminating Time-Out Issues when Provisioning Compute Nodes

The provisioning process is an appliance level orchestration of many configuration operations that run at the level of Oracle VM Manager and the individual Oracle VM Servers or compute nodes. As the virtualized environment grows – meaning there are more virtual machines, storage paths and networks –, the time required to complete various discovery tasks increases exponentially.

The maximum task durations have been configured to reliably accommodate a standard base rack setup. At a given point, however, the complexity of the existing configuration, when replicated to a large number of compute nodes, increases the duration of tasks beyond their standard time-out. As a result, provisioning failures occur.

Because many provisioning tasks have been designed to use a common time-out mechanism, this problem cannot be resolved by simply increasing the global time-out. Doing so would decrease the overall performance of the system. To overcome this issue, additional code has been implemented to allow a

finer-grained definition of time-outs through a number of settings in a system configuration file: `/var/lib/ovca/ovca-system.conf`.

If you run into time-out issues when provisioning additional compute nodes, it may be possible to resolve them by tweaking specific time-out settings in the configuration. Depending on which job failures occur, changing the `storage_refresh_timeout`, `discover_server_timeout` or other parameters could allow the provisioning operations to complete successfully. These changes would need to be applied on both management nodes.

Please contact your Oracle representative if your compute nodes fail to provision due to time-out issues. Oracle product specialists can analyse these failures for you and recommend new time-out parameters accordingly.

7.13 Returning Oracle VM Server Pool to Operation After Network Services Restart



Warning

This section applies **only** to systems with an InfiniBand-based network architecture. The use of the `bond0` interface described in this section is inherent to the network design based on the use of Oracle Fabric Interconnect F1-15s .

When network services are restarted on the master management node, the connection to the Oracle VM management network (`bond0`) is lost. By design, the `bond0` interface is not brought up automatically on boot, so that the virtual IP of the management cluster can be configured on the correct node, depending on which management node assumes the master role. While the master management node is disconnected from the Oracle VM management network, the Oracle VM Manager user interface reports that the compute nodes in the server pool are offline.

The management node that becomes the master, runs the Oracle VM services necessary to bring up the `bond0` interface and configure the virtual IP within a few minutes. It is expected that the compute nodes in the Oracle VM server pool return to their normal online status in the Oracle VM Manager user interface. If the master management node does not reconnect automatically to the Oracle VM management network, bring the `bond0` interface up manually from the Oracle Linux shell.



Warning

Execute this procedure *ONLY* when so instructed by Oracle Support. This should only be necessary in rare situations where the master management node fails to connect automatically. You should never manually disconnect or restart networking on any node.

Manually Reconnecting the Master Management Node to the Oracle VM Management Network

1. Using SSH and an account with superuser privileges, log into the disconnected master management node on the appliance management network.

```
# ssh root@192.168.4.3
root@192.168.4.3's password:
[root@ovcamn05r1 ~]#
```

2. Check the configuration of the `bond0` interface.

If the interface is down, the console output looks similar to this:

```
# ifconfig bond0
bond0      Link encap:Ethernet  HWaddr 00:13:97:4E:B0:02
```

```
BROADCAST MULTICAST MTU:1500 Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:0
RX bytes:0 (0.0 B) TX bytes:0 (0.0 B)
```

3. Bring the `bond0` interface up.

```
# ifconfig bond0 up
```

4. Check the configuration of the `bond0` interface again.

When the interface reconnects successfully to the Oracle VM management network, the console output looks similar to this:

```
# ifconfig bond0
bond0      Link encap:Ethernet  HWaddr 00:13:97:4E:B0:02
          inet addr:192.168.140.4  Bcast:192.168.140.255  Mask:255.255.255.0
          inet6 addr: fe80::213:97ff:fe4e:b002/64 Scope:Link
          UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500 Metric:1
          RX packets:62191 errors:0 dropped:0 overruns:0 frame:0
          TX packets:9183 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:4539474 (4.33 MB)  TX bytes:1853641 (1.77 MB)
```

7.14 Recovering from Tenant Group Configuration Mismatches

Tenant groups are essentially Oracle VM server pools, created and managed at the appliance level, with support for automatic custom network configuration across all pool members. The tenant groups appear in Oracle VM Manager, where the administrator could modify the server pool, but such operations are not supported in Oracle Private Cloud Appliance and cause configuration mismatches.

If you have inadvertently modified the configuration of a tenant group in Oracle VM Manager, follow the instructions in this section to correct the inconsistent state of your environment.



Caution

If the operations described below do not resolve the issue, it could be necessary to reprovision the affected compute nodes. This can result in downtime and data loss.

Adding a Server to a Tenant Group

If you try to add a server to a pool or tenant group using Oracle VM Manager, the operation succeeds. However, the newly added server is not connected to the custom networks associated with the tenant group because the Oracle Private Cloud Appliance controller software is not aware that a server has been added.

To correct this situation, first remove the server from the tenant group again in Oracle VM Manager. Then add the server to the tenant group again using the correct method, which is through the Oracle Private Cloud Appliance CLI. See [Section 2.7.2, “Configuring Tenant Groups”](#).

As a result, Oracle VM Manager and Oracle Private Cloud Appliance are in sync again.

Removing a Server from a Tenant Group

If you try to remove a server from a pool or tenant group using Oracle VM Manager, the operation succeeds. However, the Oracle Private Cloud Appliance controller software is not aware that a server has been removed, and the custom network configuration associated with the tenant group is not removed from the server.

At this point, Oracle Private Cloud Appliance assumes that the server is still a member of the tenant group, and any attempt to remove the server from the tenant group through the Oracle Private Cloud Appliance CLI results in an error:

```
PCA> remove server ovcacn09r1 myTenantGroup
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y

Status: Failure
Error Message: Error (SERVER_001): Exception while trying to
remove the server ovcacn09r1 from tenant group myTenantGroup.
ovcacn09r1 is not a member of the Tenant Group myTenantGroup.
```

To correct this situation, use Oracle VM Manager to add the previously removed server to the tenant group again. Then use the Oracle Private Cloud Appliance CLI to remove the server from the tenant group. See [Section 2.7.2, “Configuring Tenant Groups”](#). After the `remove server` command is applied successfully, the server is taken out of the tenant group, custom network configurations are removed, and the server is placed in the Unassigned Servers group in Oracle VM Manager. As a result, Oracle VM Manager and Oracle Private Cloud Appliance are in sync again.

7.15 Configure Xen CPU Frequency Scaling for Best Performance

The Xen hypervisor offers a mechanism to balance performance and power consumption through CPU frequency scaling. Known as the Current Governor, this mechanism can lower power consumption by throttling the clock speed when a CPU is idle.

Certain versions of Oracle VM Server have the Current Governor set to `ondemand` by default, which dynamically scales the CPU clock based on the load. Oracle recommends that on Oracle Private Cloud Appliance compute nodes you run the Current Governor with the `performance` setting. Particularly if you find that systems are not performing as expected after an upgrade of Oracle VM Server, make sure that the Current Governor is configured correctly.

To verify the Current Governor setting of a compute node, log in using SSH and enter the following command at the Oracle Linux prompt:

```
]# xenpm get-cpufreq-para
cpu id          : 0
affected_cpus   : 0
cpuinfo frequency : max [2301000] min [1200000] cur [2301000]
scaling_driver  : acpi-cpufreq
scaling_avail_gov : userspace performance powersave ondemandcurrent_governor : performance
scaling_avail_freq : *2301000 2300000 2200000 2100000 2000000 1900000 1800000 1700000 1600000 1500000 14
scaling frequency : max [2301000] min [1200000] cur [2301000]
turbo mode      : enabled
[...]
```

The command lists all CPUs in the compute node. If the `current_governor` parameter is set to anything other than `performance`, you should change the Current Governor configuration.

To set performance mode manually, enter this command: `xenpm set-scaling-governor performance`.

To make this setting persistent, add it to the `grub.cfg` file.

1. Add the xen cpu frequency setting to the `/etc/default/grub` template file, as shown in this example:

```
GRUB_CMDLINE_XEN="dom0_mem=max:6144M allowsuperpage dom0_vcpus_pin dom0_max_vcpus=20 cpufreq=xen:perform
```

2. Rebuild `grub.cfg` by means of the following command:

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

Index

Oracle Server X8-2replaceable components, 169

A

accessibility, 28

appliance

architecture, 1

backup, 22

hardware, 1, 2

high availability, 21

overview, 1, 1

provisioning, 1, 19

software, 1, 12

software update, 23

upgrader tool, 23

ASR

introduction, 166

prerequisites, 167

setup and asset activation, 167

understanding ASR, 166

authentication, 54

Auto Service Request

introduction, 166

prerequisites, 167

setup and asset activation, 167

understanding ASR, 166

B

backup, 22, 25

restore after password change, 222

C

certificate

SSL encryption with custom CA certificate, 225

change password

Oracle VM Agent, 218

compute node

deprovisioning, 229

failover, 22

provisioning, 19

provisioning failure, 228

replacement or repair, 229

server information, 150

time-outs when provisioning, 230

upgrade Oracle VM Server, 75

upgrade virtualization platform, 75

compute nodes

hardware, 5

configuration, 25

DNS, 32

initial setup, 32

logging, 215

management network, 32

management node, 32

NTP, 32

proxy, 216

public network, 32

contrast, 28

CRU servicing

NM2-36P Sun Datacenter InfiniBand Expansion Switch, 210

Oracle Fabric Interconnect F1-15, 211

Oracle Server X5-2, 191

Oracle Server X6-2, 188

Oracle Server X7-2, 186

Oracle Server X8-2, 184

Oracle Switch ES1-24, 209

Oracle ZFS Storage Appliance ZS3-ES, 203

Oracle ZFS Storage Appliance ZS5-ES, 201

Oracle ZFS Storage Appliance ZS7-2, 198

rack parts, 182

Sun Server X3-2, 196

Sun Server X4-2, 193

Sun ZFS Storage Appliance 7320, 207

custom networking, 39

configure on Ethernet architecture, 40

configure on InfiniBand architecture, 44

create and delete custom network, 39

D

Dashboard, 25, 32, 54

Hardware View, 28

login, 26

software, 12

database

failover, 21

databases

software, 14

deprovision compute node, 229

E

electrostatic discharge safety, 181

ES1-24 switch

component servicing instructions, 209

external storage

enable fibre channel on provisioned appliance, 220

F

Fabric Interconnect F1-15

component servicing instructions, 211

replaceable components, 180

failover

compute node, 22

database, 22

management node, 21

- networking, 22
- storage, 22

fibre channel

- enable on provisioned appliance, 220

font size, 28

H

hardware, 2

- compute nodes, 5
- health monitoring, 57
- identifying, 28
- management nodes, 4
- monitoring, 28
- networking, 8
- status, 28
- storage, 5
- view, 28

health monitoring, 57

high availability, 21

I

InfiniBand switch

- component servicing instructions, 210
- replaceable components, 180

initial setup, 25

J

jobs and events, 163

L

logging

- configure, 215

M

management node

- failover, 21
- initial setup, 32
- provisioning, 19
- update appliance controller software, 64

management nodes

- hardware, 4

monitoring, 57

- enabling SNMP, 224
- hardware, 25, 28, 149
- Oracle VM, 149
- Oracle VM Events perspective, 149
- Oracle VM Info perspective, 149
- virtual machine, 149

N

network

- configuration, 25

- functional limitations, 35
- limitations Ethernet architecture, 36
- limitations InfiniBand architecture, 38
- monitoring, 25

network customization, 39

- configure on Ethernet architecture, 40
- configure on InfiniBand architecture, 44
- create and delete custom network, 39

Network Environment, 32

networking

- Ethernet architecture, 8
- failover, 22
- hardware, 8
- InfiniBand architecture, 10
- proxy, 216
- server pool offline after network restart, 231
- using VLANs, 217

NM2-36P InfiniBand switch

- component servicing instructions, 210
- replaceable components, 180

O

operating systems

- software, 13

Oracle Server X5-2

- component servicing instructions, 191
- replaceable components, 172

Oracle Server X6-2

- component servicing instructions, 188
- replaceable components, 171

Oracle Server X7-2

- component servicing instructions, 186
- replaceable components, 170

Oracle Server X8-2

- component servicing instructions, 184

Oracle VM, 145

- Events, 163
- Events perspective, 149
- health, 149
- Info perspective, 149
- Jobs, 163
- limitations, 146
- login, 149
- monitoring, 149
- networking, 155
- Reports and Resources, 163
- repositories, 153
- server pool offline after network restart, 231
- Servers and VMs, 150
- storage, 161
- tagging, 163

Oracle VM Agent

- change password, 218

Oracle VM Manager
 adding expansion nodes, 20
 restore backup after password change, 222
 server pool, 20
 server pool mismatch with tenant group, 232
 software, 13

Oracle ZFS Storage Appliance ZS3-ES
 component servicing instructions, 203

Oracle ZFS Storage Appliance ZS5-ES, 7
 component servicing instructions, 201

Oracle ZFS Storage Appliance ZS7-2, 5
 component servicing instructions, 198

P

password, 54
 change Oracle VM Agent password, 218

password change
 failure restoring backup, 222

password management, 54

password manager, 13

power off procedure
 servicing, 182

provisioning, 19
 compute node discovery, 19
 compute node failure, 228
 eliminate time-out issues, 230
 expansion node, 20
 initialization, 19
 server pool configuration, 20

proxy
 configure, 216

R

rack
 component servicing instructions, 182
 replaceable components, 168

replaceable components
 NM2-36P Sun Datacenter InfiniBand Expansion Switch, 179
 Oracle Fabric Interconnect F1-15, 180
 Oracle Server X5-2, 172
 Oracle Server X6-2, 171
 Oracle Server X7-2, 170
 Oracle Server X8-2, 169
 Oracle Switch ES1-24, 179
 Oracle ZFS Storage Appliance ZS3-ES, 176
 Oracle ZFS Storage Appliance ZS5-ES, 175
 Oracle ZFS Storage Appliance ZS7-2, 174
 rack infrastructure, 168
 Sun Server X3-2, 173
 Sun Server X4-2, 173
 Sun ZFS Storage Appliance 7320, 178

reset, 54

S

safety
 electrostatic discharge, 181
 service precautions, 181

screen reader, 28

server pool
 offline after network restart, 231
 tenant group configuration mismatch, 232

service, 166
 ASR prerequisites, 167
 ASR setup and asset activation, 167
 Auto Service Request, 166
 electrostatic discharge, 181
 preparations, 181
 replaceable components, 168
 safety precautions, 181
 servicing Fabric Interconnect parts, 211
 servicing NM2-36P Sun Datacenter InfiniBand Expansion Switch parts, 210
 servicing Oracle Server X5-2 parts, 191
 servicing Oracle Server X6-2 parts, 188
 servicing Oracle Server X7-2 parts, 186
 servicing Oracle Server X8-2 parts, 184
 servicing Oracle Switch ES1-24 parts, 209
 servicing Oracle ZFS Storage Appliance ZS3-ES parts, 203
 servicing Oracle ZFS Storage Appliance ZS5-ES parts, 201
 servicing Oracle ZFS Storage Appliance ZS7-2 parts, 198
 servicing rack parts, 182
 servicing Sun Server X3-2 parts, 196
 servicing Sun Server X4-2 parts, 193
 servicing Sun ZFS Storage Appliance 7320 parts, 207
 understanding ASR, 166

servicing the system
 powering down, 182

SNMP
 enable, 224

software, 12
 Dashboard, 12
 databases, 13
 operating systems, 13
 Oracle VM Manager, 13
 wallet, 13

software update, 23, 61
 update appliance controller software, 64
 upgrade existing compute node with new Oracle VM Server, 75
 upgrader tool, 23, 64

SSL
 using custom certificate, 225

storage

- adding, 161
- configuration, 161
- enable fibre channel on provisioned appliance, 220
- failover, 22
- hardware, 5
- Storage Appliance
 - component servicing instructions, 198, 201, 203, 207
 - replaceable components, 174, 175, 177, 178
- Sun Server X3-2
 - component servicing instructions, 196
 - replaceable components, 174
- Sun Server X4-2
 - component servicing instructions, 193
 - replaceable components, 173
- Sun ZFS Storage Appliance 7320
 - component servicing instructions, 207
- switch ES1-24
 - component servicing instructions, 209
 - replaceable components, 179

T

- tagging resources, 163
- tenant group
 - recovering from configuration mismatches, 232

U

- update, 61
 - update appliance controller software, 64
 - upgrade Oracle VM Server, 75
 - upgrade virtualization platform, 75
 - upgrader tool, 23, 64
- upgrader tool, 23
 - installation, 65
 - pre-checks, 66
 - update appliance controller software, 64, 68
 - upgrade readiness, 66
 - usage, 64

V

- virtual machine
 - clone, 150
 - clone customizer, 153
 - console, 150
 - create, 150
 - delete, 150
 - import, 150
 - installation media, 153
 - ISOs, 153
 - management, 145, 150
 - messaging, 150
 - migrate, 150
 - resources, 153
 - templates, 153

- virtual appliances, 153
- virtual disks, 153
- VNICs, 155
- VLAN
 - enabling VLAN traffic, 217

W

- wallet
 - software, 13

Z

- ZFS Storage Appliance
 - component servicing instructions, 198, 201, 203, 207
 - replaceable components, 174, 175, 177, 178