Oracle® Private Cloud Appliance

Administrator's Guide for Release 2.2

ORACLE

Table of Contents

Preface

This document is part of the documentation library for Oracle Private Cloud Appliance (PCA) Release 2.2, which is available at:

http://docs.oracle.com/cd/E71897_01.

The documentation library 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 PCA.

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 PCA.

Oracle Private Cloud Appliance Installation Guide

The installation guide provides detailed instructions to prepare the installation site and install Oracle PCA. 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 PCA. It conforms to Compliance Model No. OCA1A.

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 PCA.

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 PCA. A printed quick start poster is shipped with each Oracle PCA 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 library 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 PCA expansion node. A printed expansion node setup poster is shipped with each Oracle PCA expansion node.

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

Audience

The Oracle Private Cloud Appliance (PCA) 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 PCA 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 PCA 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 PCA determines the requirements of your administrative knowledge.

Related Documentation

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

- Sun Rack II 1042 and 1242

  http://docs.oracle.com/cd/E19844-01/index.html

- Oracle Server X6-2

  http://docs.oracle.com/cd/E62159_01/index.html

- Oracle Server X5-2

  http://docs.oracle.com/cd/E41059_01/index.html

- Sun Server X4-2

  http://docs.oracle.com/cd/E36975_01/index.html

- Sun Server X3-2

  http://docs.oracle.com/cd/E22368_01/index.html

- Oracle ZFS Storage Appliance ZS3-ES

  http://docs.oracle.com/cd/E37831_01/index.html

- Sun ZFS Storage Appliance 7320

  http://docs.oracle.com/cd/E28317_01/index.html

- Oracle Switch ES1-24

  http://docs.oracle.com/cd/E39109_01/index.html

- NM2-36P Sun Datacenter InfiniBand Expansion Switch

  http://docs.oracle.com/cd/E36265_01/index.html

- Oracle Fabric Interconnect F1-15

  http://docs.oracle.com/cd/E38500_01/index.html

- Oracle Integrated Lights Out Manager (ILOM) 3.1

  http://docs.oracle.com/cd/E24707_01/index.html

- Oracle VM

  http://docs.oracle.com/cd/E35328_01/index.html

- Oracle Enterprise Manager Plug-in

  http://docs.oracle.com/cd/E63000_01/EMPCA/toc.htm

Access to Oracle Support

Oracle customers have access to electronic support through My Oracle Support. For information, visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=info or visit http://www.oracle.com/pls/topic/lookup?ctx=acc&id=trs if you are hearing impaired.

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.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.
`monospace`	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: 2017-10-24 (revision: 1407)

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

Table of Contents

This chapter describes what Oracle Private Cloud Appliance (PCA) 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 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 (PCA), Oracle offers a unique full stack of hardware, software, virtualization technology and rapid application deployment through assemblies. 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 PCA 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 provide storage solutions from Oracle or third parties.

1.2 Hardware Components

The Oracle Private Cloud Appliance consists of a Sun Rack II 1242 base, populated with the hardware components identified in Figure 1.1.

Note

Oracle Private Cloud Appliance base racks are factory installed with the latest supported generation of server nodes. Earlier generations of the Oracle PCA server architecture continue to be supported by the Oracle PCA controller software. If additional nodes have been installed, or if nodes have been replaced, an Oracle PCA rack can be populated with a mix of supported servers. Besides the marginal performance increase offered by a newer server architecture, there is no functional difference between the server generations within the Oracle PCA environment.

Currently supported server architectures are:

- Oracle Server X6-2: software release 2.2.1 (compute node only)

- Oracle Server X5-2: software release 2.0.3 or newer

- Sun Server X4-2: software release 1.1.3 or newer

- Sun Server X3-2: since initial release

When you order expansion compute nodes, you receive the latest available generation.

The Oracle PCA Controller Software must be upgraded to the correct version to support the hardware installed in your environment.

Note

The initial version of the Oracle Private Cloud Appliance, the X3-2 base rack, was shipped with the Sun ZFS Storage Appliance 7320. This hardware component has since been replaced with the Oracle ZFS Storage Appliance ZS3-ES. Software support for the Oracle ZFS Storage Appliance ZS3-ES is available as of Release 1.1.3. All software releases continue to support the Sun ZFS Storage Appliance 7320.

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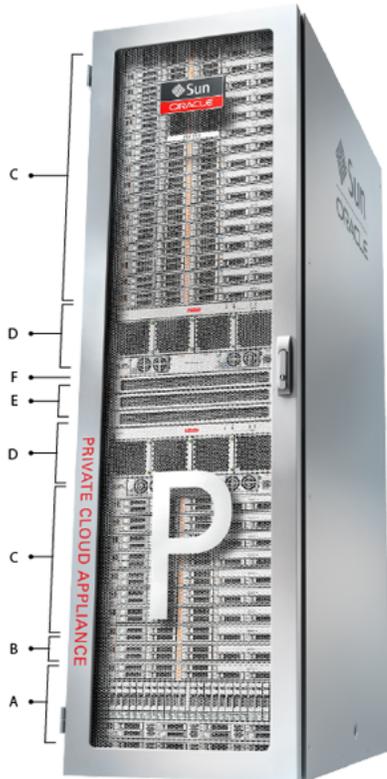
Table 1.1 Figure Legend

Item	Quantity	Description
A	1	Either Oracle ZFS Storage Appliance ZS3-ES or Sun ZFS Storage Appliance 7320
B	2	Either Oracle Server X5-2, Sun Server X4-2, or Sun Server X3-2, used as management nodes
C	2-25	Either Oracle Server X6-2, Oracle Server X5-2, Sun Server X4-2, or Sun Server X3-2, used as virtualization compute nodes (Due to the higher power requirements of the Oracle Server X6-2 and Oracle Server X5-2, if the appliance is equipped with 22kVA PDUs, the maximum number of compute nodes is 23.)
D	2	Oracle Fabric Interconnect F1-15
E	2	NM2-36P Sun Datacenter InfiniBand Expansion Switch
F	2	Oracle Switch ES1-24

1.2.1 Management Nodes

At the heart of each Oracle PCA 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 at the bottom of the rack. Because rack units are numbered from the bottom up, and the bottom four are occupied by 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 PCA, refer to Section 1.5, "High Availability".

When you power on the Oracle PCA 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 6 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 PCA 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 3.2.11 and additional packages for InfiniBand and Software Defined Networking. When provisioning is complete, the Oracle PCA 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 Oracle Server X6-2, Oracle Server X5-2, Sun Server X4-2 and Sun Server X3-2 compute nodes, refer to Server Components in the Oracle Private Cloud Appliance Installation Guide. Different generations of servers may be mixed within the same installation. In such configurations the version of the Oracle PCA controller software must support the most recent server model installed. Compute nodes of different hardware generations operate within the same server pool but belong to different CPU compatibility groups. Since live migration between CPU compatibility groups is not supported, virtual machines have to be cold-migrated between compute nodes of different generations. An exception to this rule can be implemented, but only if the migration occurs from an older model to a newer model. For more information about CPU compatibility groups, please refer to the section "Server Processor Compatibility Groups" in the Oracle VM User's Guide.

The Oracle PCA 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 PCA 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.

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 ZFS Storage Appliance ZS3-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 PCA 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 ZS3-ES is as follows:

- Two clustered storage heads with two 1.6TB SSDs each, used exclusively for cache and logging

- One fully populated disk chassis with twenty 900GB SATA hard disks

- RAID-Z2 configuration, for best balance between performance and data protection, with a total usable space of 11.3TB

> **Note**
>
> Oracle PCA Release 1.0 base racks, which may be upgraded to the Release 2.2 software stack, use a Sun ZFS Storage Appliance 7320. It offers the same performance, functionality, and configuration, but its storage heads use smaller SSDs. The disk shelf and its disks are identical in both models.

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 PCA 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 PCA

  - `Database` – to be used for all system databases

  - `Incoming` (20GB) – to be used for FTP file transfers, primarily for Oracle PCA component backups

  - `Templates` – placeholder file system for future use

  - `User` – placeholder file system for future use

  - `Yum` – placeholder file system for future use

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 PCA system components.

1.2.4 Network Infrastructure

The Oracle Private Cloud Appliance relies on a combination of Ethernet connectivity and an InfiniBand network fabric. 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.

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 PCA 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 PCA 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 PCA 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 PCA, five essential SDNs are configured: a storage network, an Oracle VM management network, a management Ethernet network, and two VLAN-enabled virtual machine networks.

- The **storage network** 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 PCA 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. VLAN 1 is automatically configured for 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. VLAN 1 is automatically configured for 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 PCA. 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 PCA: 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 PCA 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 PCA 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 PCA Dashboard is a WebLogic application that is available via the active management node. It is installed on top of the Oracle WebLogic Server 11*g* that is packaged with Oracle VM Manager, as described in Section 1.3.3, "Oracle VM Manager".

Use the Dashboard to perform the following tasks:

• Appliance system monitoring and component identification

• Monitoring and identifying physical network connections

• Initial configuration of management node networking data

• Resetting of the global password for Oracle PCA configuration components

The Oracle PCA 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 PCA 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 PCA 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 PCA.

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 11*g***: including Application Development Framework (ADF) Release 11*g*, used to host and run the Oracle VM Manager application

- **MySQL 5.5 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 4, *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 UEK2

- Compute Nodes: Oracle Linux 5 with UEK2

- Oracle ZFS Storage Appliance ZS3-ES: 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 PCA. 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 PCA 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.

The following table lists the different databases used by the Oracle PCA.

Table 1.2 Oracle Private Cloud Appliance Databases

Item	Description
Oracle PCA Node Database	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.
	Type: BerkeleyDB
	Location: `MGMT_ROOT/db/node` on the ZFS, accessible via `/nfs/shared_storage/db/node` on each management node
Oracle Fabric Interconnect Database	Contains IP and host name data for the Oracle Fabric Interconnect F1-15s.
	Type: BerkeleyDB

Item	Description
	Location: `MGMT_ROOT/db/infrastructure` on the ZFS, accessible via `/nfs/shared_storage/db/infrastructure` on each management node
Oracle PCA Netbundle Database	Predefines Ethernet and bond device names for all possible networks that can be configured throughout the system, and which are allocated dynamically. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/netbundle` on the ZFS, accessible via `/nfs/shared_storage/db/netbundle` on each management node
Oracle Switch ES1-24 Ports Database	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. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/opus_ports` on the ZFS, accessible via `/nfs/shared_storage/db/opus_ports` on each management node
Oracle PCA DHCP Database	Contains information on the assignment of DHCP addresses to newly detected compute nodes. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/dhcp` on the ZFS, accessible via `/nfs/shared_storage/db/dhcp` on each management node
Oracle PCA Mini Database	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 PCA component. Entries contain a sync ID for more convenient usage within the Command Line Interface (CLI). **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/mini_db` on the ZFS, accessible via `/nfs/shared_storage/db/mini_db` on each management node
Oracle PCA Setup Database	Contains the data set by the Oracle PCA Dashboard setup facility. The data in this database is automatically applied by both the active and standby management nodes when a change is detected. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/setup` on the ZFS, accessible via `/nfs/shared_storage/db/setup` on each management node
Oracle PCA Task Database	Contains state data for all of the asynchronous tasks that have been dispatched within the Oracle PCA. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/task` on the ZFS, accessible via `/nfs/shared_storage/db/task` on each management node

Item	Description
Oracle PCA Synchronization Databases	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 `sync_errored_tasks` database, from where they can be retried or acknowledged. Synchronization databases are not present by default. They are created when the first synchronization task of a given type is received. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/sync_*` on the ZFS, accessible via `/nfs/shared_storage/db/sync_*` on each management node
Oracle PCA Update Database	Used to track the two-node coordinated management node update process. **Note** 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. **Type:** BerkeleyDB **Location:** `MGMT_ROOT/db/update` on the ZFS, accessible via `/nfs/shared_storage/db/update` on each management node
Oracle VM Manager Database	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. **Type:** MySQL Database **Location:** `MGMT_ROOT/ovmm_mysql/data/` on the ZFS, accessible via `/nfs/shared_storage/ovmm_mysql/data/` on each management node

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 PCA 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 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 PCA

- `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 PCA, 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

- `pca-redirect`: a daemon that redirects HTTP or HTTPS requests to the Oracle PCA 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 PCA software components to communicate directly with the underlying management scripts running on the management node

- `pca-vm-tool`: an internal script to facilitate virtual machine deployment.

Many of these applications use a specific Oracle PCA 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 PCA. 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 PCA 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 ILOM of each component in the rack. 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 PCA 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`. Running `pca-diag` from different locations can be particularly helpful in debugging InfiniBand fabric issues.

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 PCA 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 management 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 PCA services, including dhcpd, Oracle VM Manager and the Oracle PCA databases. The virtual IP address of the management cluster is also brought online, and the Oracle PCA Dashboard is started within WebLogic. 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 management network (`192.168.4.0/24`). When the Node Manager has verified that a node has a valid service tag for use with Oracle PCA, 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 PCA 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 for InfiniBand and IO Director (Fabric Interconnect) support are added to the installation. At the

end of the installation process, the network configuration files are updated to allow all necessary network interfaces and bonds to be brought up.

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

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, access to the shared storage, and a virtual IP. For every compute node added to Oracle VM Manager the IPMI configuration is stored in order to enable convenient remote power-on/off.

Oracle PCA expects that all compute nodes in one rack 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 PCA 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 and software 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 PCA 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 PCA, the management nodes are configured as a cluster. The cluster relies on an OCFS2 file system exported as a 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 PCA 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 PCA 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 PCA 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 PCA 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.

Storage Redundancy

Further redundancy is provided through the use of the ZFS storage appliance to host storage. This component is configured to support RAID-Z2 providing integrated redundancy with a fault tolerance of up to two failed drives with zero data loss. 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 PCA 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 is 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 3.2.4, "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:

- **nm2**: contains the NM2-36P Sun Datacenter InfiniBand Expansion Switch configuration data

- **opus**: 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 PCA 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 and is described in detail in the section entitled Oracle VM Manager MySQL Backup, in the Oracle VM Installation and Upgrade Guide.

- **xsigo**: contains the configuration data for the Fabric Interconnects.

- **zfssa**: 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 failure [1].

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 a machine with connections to the internal appliance networks as well as the external (public) data center network.

---

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

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 and Recovery Guide*.

1.7 Oracle Private Cloud Appliance Software Update

Oracle Private Cloud Appliance includes the tools required to perform a full software update of the appliance. On the master management node command line, commands are issued to download a new version of the software stack from a specified URI and update all components to the latest version. For step-by-step instructions, see Section 2.10, "Update".

An Oracle PCA software stack image is provided in the form of a zip file, approximately 6 GB in size. The zip file contains an automatically built ISO with the latest software updates from Oracle.

The software image can be downloaded using the FTP, HTTP or HTTPS protocols [2] . Once the image has been downloaded it is extracted into the `mgmt_image`, `hmp_image` and `ovm_image` subdirectories of `/nfs/shared_storage`, which resides in the `MGMT_ROOT` file system on the ZFS storage appliance.

From the command line, you are able to *activate* a downloaded software image. Activation starts a coordinated process on the management nodes that relies on management node failover to succeed.

The activation process can be summarized as follows:

- The current master begins the process by creating a database entry to include a message containing the version number of the software that you are updating to; a time stamp and the node name of the inactive management node.

- The current master reboots the inactive management node and provisions it with the new software image in the same way that a compute node is provisioned at start-up. The current software on the inactive management node is wiped and the new software is installed directly from the ISO stored in `/nfs/shared_storage/mgmt_image` on the storage appliance.

- After the install, the management node is rebooted and starts its *ovca* service. This causes it to re-enter the remaster code and it finalizes its own database entry.

- The current master detects the database update indicating that the installation has completed and the newly installed management node is ready to resume service as the new master. It creates its own database entry to indicate that it is now being updated, and then reboots.

- The newly installed management node is promoted to the cluster master and then proceeds to perform the provisioning required to install the new software image on the rebooting management node.

- When installation is complete, the management node is rebooted and comes up in standby mode.

---

[2] It may be necessary to edit the Oracle PCA system properties, using the CLI on each management node, to ensure that the correct proxy settings are specified for a download to succeed. See Section 7.1, "Adding Proxy Settings for Oracle Private Cloud Appliance Updates" for more information.

Chapter 2 Monitoring and Managing Oracle Private Cloud Appliance

Table of Contents

Monitoring and management of the Oracle Private Cloud Appliance (PCA) is achieved using the Oracle PCA 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 PCA Dashboard allows you to perform the following tasks:

- Initial software configuration (and reconfiguration) for the appliance using the Network Setup tab, as described in Section 2.5, "Network Settings".

- Hardware monitoring and identification of each hardware component used in the appliance, accessed via the Hardware View tab described in Section 2.3, "Hardware View".

- Physical network monitoring and identification via the Network View tab described in Section 2.4, "Network View".

- Reconfiguration of appliance software, such as resetting the global password used for different components within the appliance, via the Authentication tab, as described in Section 2.8, "Authentication".

The Oracle PCA software includes functionality that is currently not available through the Dashboard user interface:

- **Backup**

  The configuration of all components within Oracle PCA 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. For details, see Section 1.7, "Oracle Private Cloud Appliance Software Update". For step-by-step instructions, see Section 2.10, "Update".

- **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 PCA 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 PCA 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 PCA Dashboard on the active management node.

> **Note**
>
> If you are following the installation process and this is your first time accessing the Oracle PCA 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 19 in one of the Oracle Switch ES1-24 switches. The default user name is **admin** and the default password is **Welcome1**.

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

Enter your Oracle PCA Dashboard administration user name in the **Username** field. This is the administration user name you configured during installation. Enter the password for the Oracle PCA Dashboard administration user name in the **Password** field.

> **Important**
>
> The Oracle PCA Dashboard makes use of cookies in order to store session data. Therefore, to successfully login and use the Oracle PCA Dashboard your web browser must accept cookies from the Oracle PCA Dashboard host.

2.2 Oracle PCA Dashboard Accessibility Features

The Oracle PCA Dashboard is compliant with the Americans with Disabilities Action (ADA) requirements. As part of the effort to make Oracle products, services, and supporting documentation accessible and usable to the disabled community, the Oracle PCA Dashboard allows you to configure the following accessibility features:

- Support for Screen Reader.

- Support for High Contrast.

- Support for Large Fonts.

You can enable these features when logging in to the Oracle PCA Dashboard or you can set or change the accessibility options once you have logged in.

To enable accessibility options when logging in:

1. On the login page of the Oracle PCA Dashboard, click the arrow to expand **Accessibility Options**.

2. Select one or more check box from the following accessibility options:

   - `I use a screen reader.`

   - `I use high contrast.`

   - `I use large fonts.`

To set or change accessibility options while you are logged in:

1. In the toolbar that runs along the top of the links at the top of the Oracle PCA Dashboard, a dropdown menu titled **Settings** allows you to change accessibility options once you are logged in. To make a change, click on the **Settings** link.

2. From the drop-down list, select from the following accessibility options:

   - `I use a screen reader.`

   - `I use high contrast.`

   - `I use large fonts.`

Changes are implemented immediately.

2.3 Hardware View

The **Hardware View** within the Oracle PCA 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 static at the time that the page is loaded. 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 7.12, "A Compute Node Fails to Complete Provisioning".

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. 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, internal InfiniBand configuration, network and 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 PCA hardware, after provisioning, consult the Oracle VM Manager or Oracle Enterprise Manager UI.

If either of the Fabric Interconnects is clicked on in the Hardware View, the view is changed to the I/O Ports tab in the Network View. See Section 2.4, "Network View" for more information.

The Hardware View provides an accessible tool for troubleshooting hardware components within the Oracle PCA 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
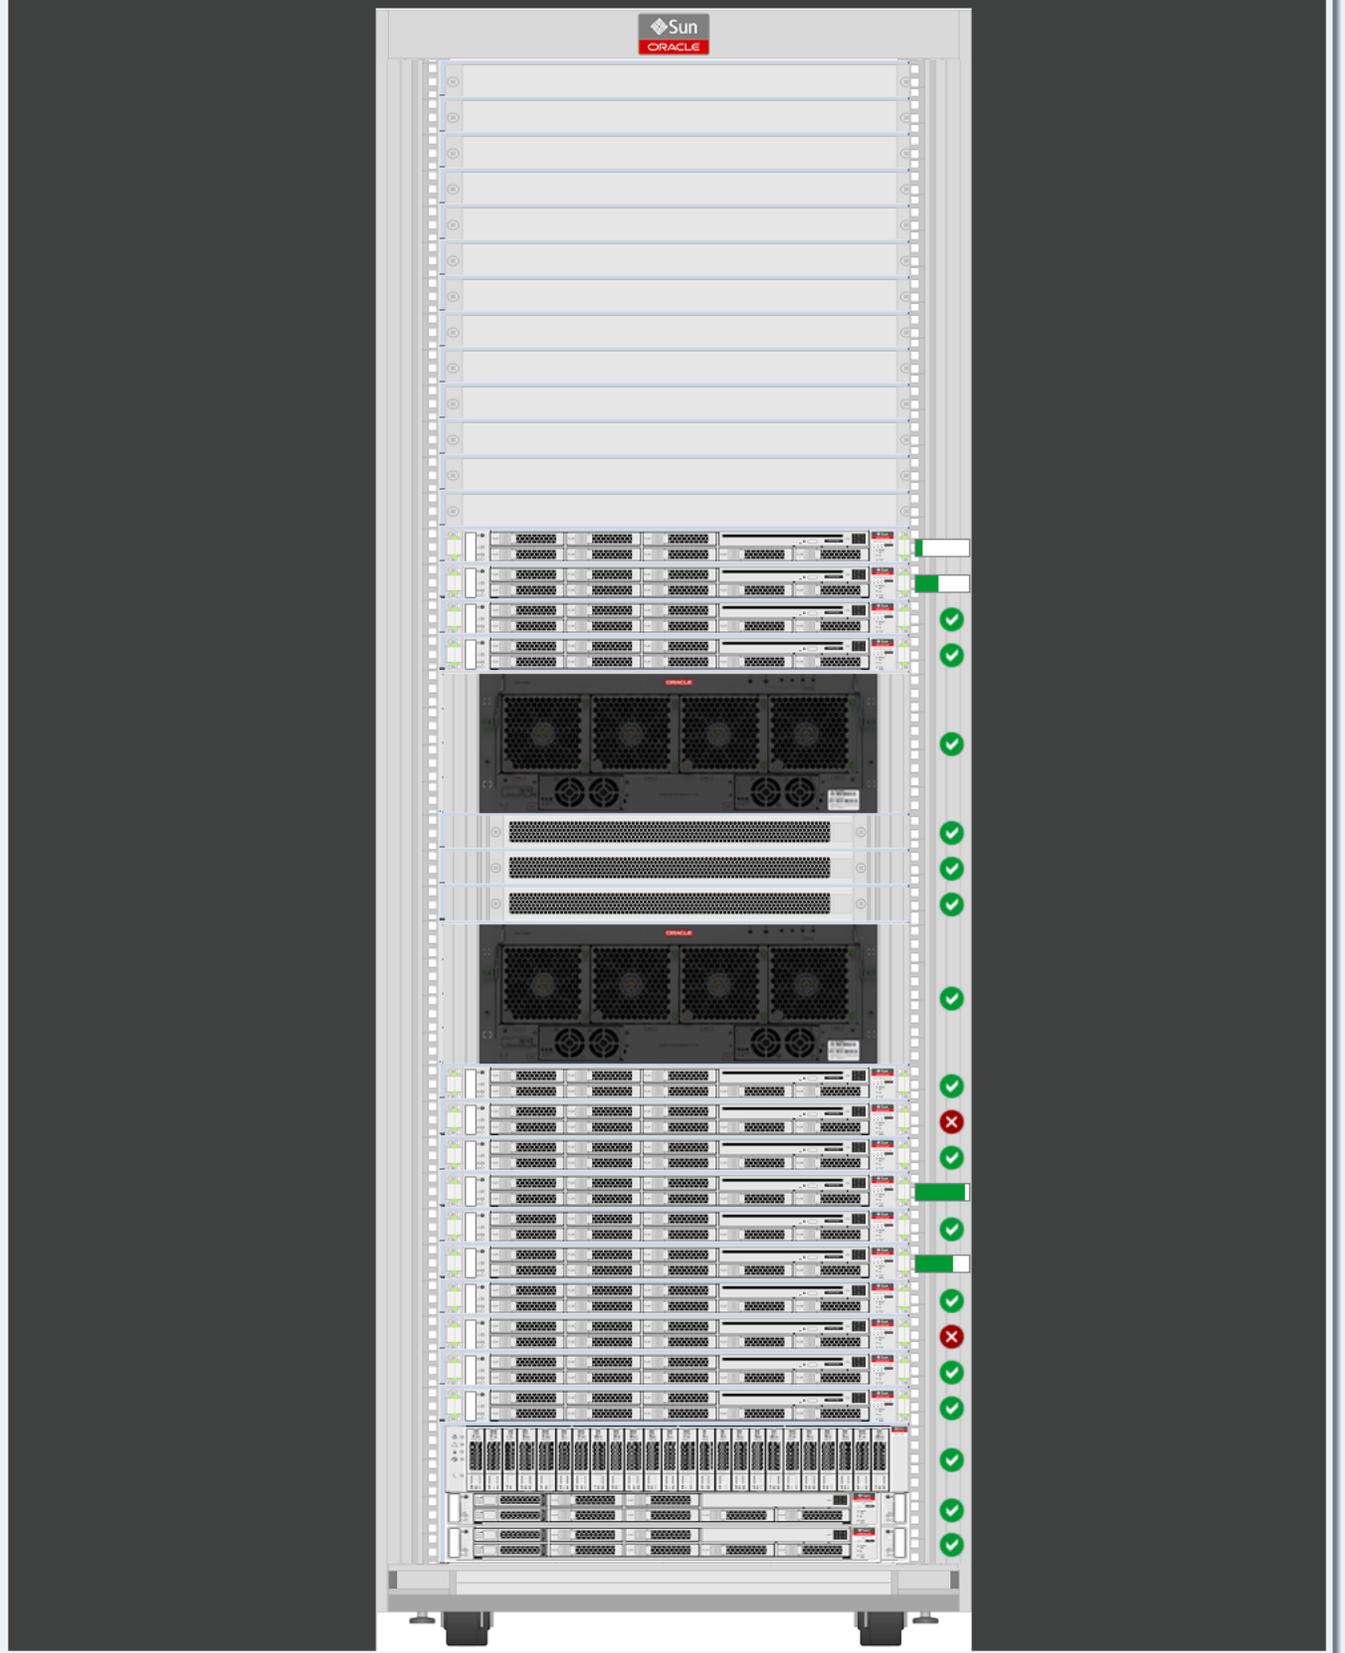
2.4 Network View

The **Network View** provides a graphical representation of the networking hardware within the Oracle PCA rack. Specifically, the I/O modules and their network ports on the are shown on a diagram depicting the back panel of the Fabric Interconnect.

Figure 2.2 Network View - IO Module Ports
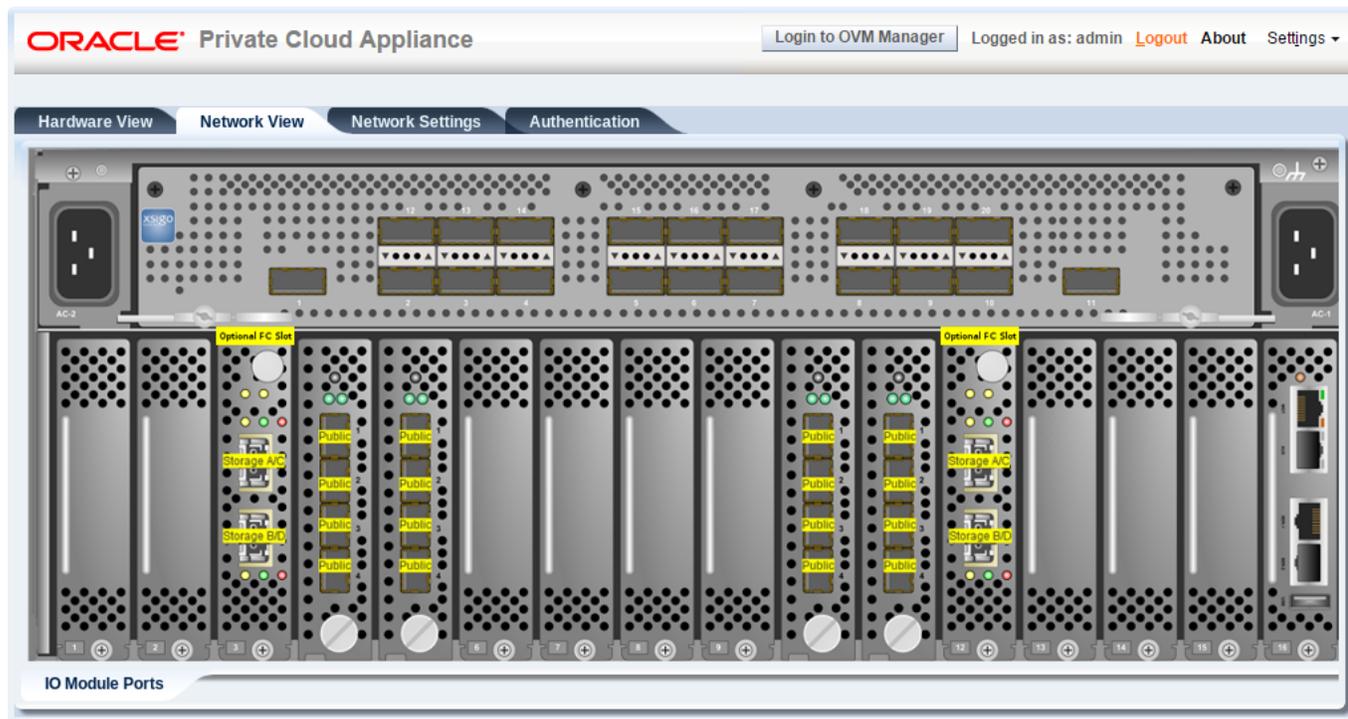
Note that the Fibre Channel (FC) ports on the PCIe Slots 3 and 12 are reserved for use to connect external FC storage to the Oracle PCA . These ports are labeled to indicate the storage clouds that are related to each FC port and can be cabled to one or more external FC switches in the manner described in Adding External Fibre Channel Storage in the *Oracle Private Cloud Appliance Installation Guide*.

2.5 Network Settings

The **Network Settings** tab is used to configure networking and service information for the management nodes. 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.

  - **Hostname:** 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.

  - **Hostname:** Specify the host name for the second management node system.

- **Management Node Cluster:**

  - **Management Network VLAN:**

    - **No VLAN:** By default this radio button is set. The default configuration does not assume that your management network exists on a VLAN.

    - **VLAN Enabled:** If you have configured a VLAN on your switch for the management network, you should set this radio button and then specify the VLAN ID for the VLAN that the management network must use in the provided field.

  - **Virtual IP:** 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.

  - **Domain Name:** Specify the domain that each of the management node systems 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.

  - **DNS:** Specify at least one DNS server that the management nodes can use for domain name resolution.

- **Current PCA Password:** You must enter the current Oracle PCA password to make changes to any of these settings.

Clicking the **OK** button at the bottom of the page saves the settings that are currently displayed on the page 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.

Figure 2.3 Network Settings

2.6 Network Customization

The Oracle PCA controller software allows you to add custom networks at the appliance level. This means that the Fabric Interconnects and other hardware components are reconfigured to enable the additional connectivity. These 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-enabled.

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 PCA CLI.

> **Caution**
>
> The following network limitations apply:
>
> - The maximum number of custom external networks is 7.
>
> - The maximum number of custom internal networks is 3.
>
> - The maximum number of VLANs is a total of 128 across all internal, external, default, and custom networks combined.

> **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 PCA CLI commands to fail.

Creating custom networks requires use of the CLI in this software release. 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 either the *virtual machine* or *storage* network role.

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.

> **Caution**
>
> For the additional routing capabilities with a host network, compute nodes must be upgraded with the Oracle VM Server version included in the Oracle PCA Controller Software Release 2.2.1.

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

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

> **Note**
>
> The default `root` password is *Welcome1*.

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

2. Launch the Oracle PCA command line interface.

```
# pca-admin
Welcome to PCA! Release: 2.2.1
PCA>
```

3. 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.

- 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.
>
> Details of the create network command arguments are provided in Section 3.2.7, "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 2.4 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 2.4 Oracle VM Manager View of Custom Network Configuration
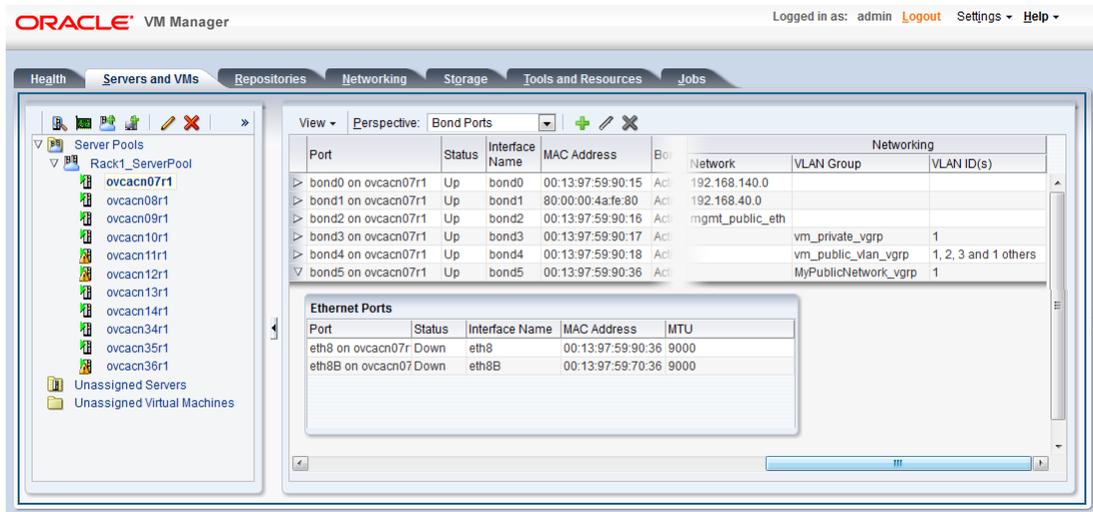
7. To disconnect servers from the custom network use the *remove network* command.

> **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
```

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*.

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

2. Launch the Oracle PCA command line interface.

```
# pca-admin
```

```
Welcome to PCA! Release: 2.2.1
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.

```
PCA> show network MyPublicNetwork

----------------------------------------
Network_Name          MyPublicNetwork
Trunkmode             True
Description           User defined network
Ports                 ['4:2', '5:2']
vNICs                 None
Status                ready
Network_Type          external_network
Compute_Nodes         None
----------------------------------------

Status: Success
```

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 PCA environment could span up to three racks containing up to 85 compute nodes: 25 compute nodes in the base rack and 30 in each of two expansion racks. A *tenant group* is a logical subset of a single Oracle PCA environment. Tenant groups provide an optional mechanism for an Oracle PCA 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 customer. It also offers isolation from cluster faults.

2.7.1 Design Assumptions and Restrictions

Up to Oracle PCA Controller Software Release 2.2.1, a factory-installed appliance base rack has all compute nodes configured as part of a single Oracle VM server pool, named *Rack1_ServerPool*. As of release 2.2.1, or during a controller software update to release 2.2.1, this server pool becomes the default tenant group.

Oracle PCA supports a maximum of 8 tenant groups. This number includes the default tenant group, which cannot be deleted from the environment. A single tenant group can contain up to 32 compute nodes. When expansion racks are added to the Oracle PCA environment, the provisioning system does not add those

compute nodes to a tenant group, but to the Oracle VM Unassigned Servers list instead. After provisioning has completed, you can add the new compute nodes in the expansion rack(s) to the tenant group of your choice.

Regardless of tenant group membership, all compute nodes are connected to all of the default Oracle PCA 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 the a compute node from the base rack is reprovisioned, it is added automatically to *Rack1_ServerPool*. When a compute node in an expansion rack is reprovisioned, it is placed in the Unassigned Servers list. After successful reprovisioning 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 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.

> **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 PCA 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 PCA CLI tenant group commands, see Chapter 3, *The Oracle Private Cloud Appliance Command Line Interface (CLI)*.

Creating and Populating a Tenant Group

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

   > **Note**
   >
   > The default `root` password is *Welcome1*.

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

2. Launch the Oracle PCA command line interface.

   ```
   # pca-admin
   Welcome to PCA! Release: 2.2.1
   PCA>
   ```

3. Create the new tenant group.

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

PCA> show tenant-group myTenantGroup

----------------------------------------
Name                myTenantGroup
Default             False
Tenant Group ID     0004fb0000020000155c15e268857a78
Servers             None
State               ready
Tenant Group VIP    192.168.140.131
Networks            None
Pool Filesystem ID  3600144f0d29d4c86000057162ecc0001
----------------------------------------
```

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, and is assigned a virtual IP address in the `192.168.140.0/24` subnet. The IP address is accessible once a member server of the tenant group has been assigned the server pool master role.

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

   Removing a compute node from an existing tenant group will fail if the server is hosting virtual machines, or if storage repositories are presented. If so, you have to migrate the virtual machines and unpresent the repositories before adding the compute node to a new tenant group.

```
PCA> add server ovcacn09r1 myTenantGroup
Status: Success

PCA> add server ovcacn10r1 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.

6. Verify the configuration of the new tenant group.

```
PCA> show tenant-group myTenantGroup

----------------------------------------
Name                myTenantGroup
Default             False
Tenant Group ID     0004fb0000020000155c15e268857a78
Servers             ['ovcacn09r1', 'ovcacn10r1']
State               ready
Tenant Group VIP    192.168.140.131
Networks            ['myPublicNetwork']
Pool Filesystem ID  3600144f0d29d4c86000057162ecc0001
----------------------------------------
```

```
Status: Success
```

The new tenant group corresponds with an Oracle VM server pool with the same name and has a virtual IP an 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 2.5 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 4.7, "Viewing and Managing Storage Resources".

Figure 2.5 Oracle VM Manager View of New Tenant Group
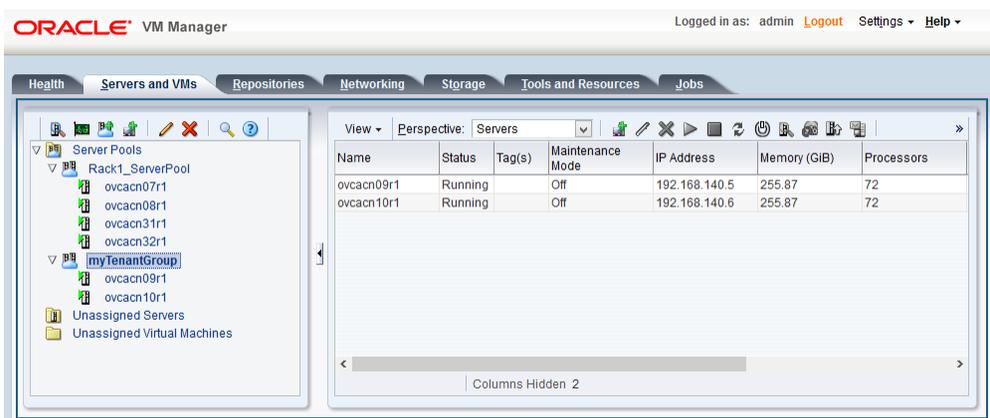
Reconfiguring and Deleting a Tenant Group

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      0004fb0000020000155c15e268857a78
Servers              ['ovcacn09r1', 'ovcacn10r1']
State                ready
Tenant Group VIP     192.168.140.131
Networks             ['myPublicNetwork']
Pool Filesystem ID   3600144f0d29d4c86000057162ecc0001
```

```
---------------------------------------

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 ovcacn09r1 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 free the virtual IP address for re-use, and 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 **Authentication** tab is used to reset the global Oracle PCA 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 PCA 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 PCA 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 screen provides a series of checkboxes that make it easy to select the level of granularity that you wish to apply to a password change. By selecting **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 checkbox associated with each component that you wish to apply a password to. The checkbox labels are described in the following list:

- **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 PCA 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 NM2-36 Switch root password**: Set the new password for the root user for the NM2-36P Sun Datacenter InfiniBand Expansion Switches.

- **Oracle NM2-36 Switch ILOM admin password**: Set the new password for the admin user for the ILOM for the NM2-36P Sun Datacenter InfiniBand Expansion Switches.

- **Oracle NM2-36 Switch ILOM operator password**: Set the new password for the operator user for the ILOM for the NM2-36P Sun Datacenter InfiniBand Expansion Switches.

- **Oracle ES1-24 Switch root password**: Set the new password for the root user for the Oracle Switch ES1-24 switches.

- **Oracle Fabric Interconnect admin password**: Set the new password for the admin user for the Oracle Fabric Interconnect F1-15 switches.

- **Oracle Fabric Interconnect recovery password**: Set the new password for recovery operations on both Oracle Fabric Interconnect F1-15 devices. This password is used in the case of a corruption or when the admin password is lost. The Oracle Fabric Interconnect F1-15 can be booted in 'recovery mode' and this password can be used to access the recovery mode menu.

- **Oracle Fabric Interconnect root password**: Set the new password for the root user for the Fabric Interconnects.

- **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 2.6 Authentication

The functionality that is available in the Oracle PCA Dashboard is equally available via the Oracle PCA CLI as described in Section 3.2.27, "update password".

⚠️ **Caution**

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

2.9 Health Monitoring

The Oracle PCA 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 through the Oracle PCA CLI at any time by means of the `diagnose` command.

Checking the Current Health Status of an Oracle PCA Installation

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

   > **Note**
   >
   > The default `root` password is *Welcome1*.

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

2. Launch the Oracle PCA command line interface.

   ```
   # pca-admin
   Welcome to PCA! Release: 2.1.1
   PCA>
   ```

3. Check the current status of the rack components by querying their ILOMs.

   ```
   PCA> diagnose ilom
   Checking ILOM health............please wait..

   IP_Address      Status           Health_Details
   ----------      ------           --------------
   192.168.4.129   Not Connected
   192.168.4.128   Not Connected
   192.168.4.127   Not Connected
   192.168.4.126   Not Connected
   192.168.4.125   Not Connected
   192.168.4.124   Not Connected
   192.168.4.123   Not Connected
   192.168.4.122   Not Connected
   192.168.4.121   Not Connected
   192.168.4.120   Not Connected
   192.168.4.101   OK
   192.168.4.102   OK
   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
   192.168.4.106   OK
   192.168.4.109   OK
   192.168.4.108   OK
   192.168.4.112   OK
   192.168.4.113   Not Connected
   192.168.4.110   OK
   192.168.4.111   OK
   192.168.4.116   Not Connected
   192.168.4.117   Not Connected
   192.168.4.114   Not Connected
   192.168.4.115   Not Connected
   192.168.4.118   Not Connected
   192.168.4.119   Not Connected
   -----------------
   27 rows displayed

   Status: Success
   ```

4. Verify that the Oracle PCA controller software is fully operational.

   ```
   PCA> diagnose software
   PCA Software Acceptance Test runner utility
   ```

```
Test -      701 - OpenSSL CVE-2014-0160 Heartbleed bug Acceptance [PASSED]
Test -      785 - PCA package Acceptance [PASSED]
Test -     1083 - Mgmt node xsigo network interface Acceptance [PASSED]
Test -      787 - Shared Storage Acceptance [PASSED]
Test -      973 - Simple connectivity Acceptance [PASSED]
Test -     1078 - Test for ovs-agent service on CNs Acceptance [PASSED]
Test -     1079 - Test for shares mounted on CNs Acceptance [PASSED]
Test -     1080 - ovs-log check Acceptance [PASSED]
Test -      788 - PCA services Acceptance [PASSED]
Test -      789 - PCA config file Acceptance [PASSED]
Test -     1300 - All compute nodes running Acceptance [PASSED]
Test -     1318 - Check support packages in PCA image Acceptance [PASSED]
Test -      928 - Repositories defined in OVM manager Acceptance [PASSED]
Test -     1107 - Compute node xsigo network interface Acceptance [PASSED]
Test -     1316 - PCA version Acceptance [PASSED]
Test -     1117 - Network interfaces check Acceptance [PASSED]
Test -      824 - OVM manager settings Acceptance [PASSED]
Test -      927 - OVM server model Acceptance [PASSED]
Test -      925 - PCA log Acceptance [PASSED]
Test -      926 - Networks defined in OVM manager for CNs Acceptance [PASSED]
Test -      822 - Compute node network interface Acceptance [PASSED]
Status: Success
```

5. Close the CLI.

```
PCA> exit
```

2.10 Update

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 PCA update handles different levels or areas of appliance functionality.

Table 2.2 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

> **Tip**
>
> My Oracle Support provides a pre-upgrade checklist with additional information and
> guidance from our product support specialists. The checklist is stored under Doc ID
> 2138663.1.

All the software included in a given release of the Oracle PCA software is tested to work together and
should be treated as one package, so no appliance component should be updated individually. All Oracle
PCA software releases are downloaded as a single large `.iso` file, which includes the items listed above.
Some releases include major updates for one or more appliance components. These are critical stepping

stones in a successful upgrade path so they must not be skipped. This is explained in Section 2.10.1, "Mandatory Updates".

> **Warning**
>
> **Oracle Enterprise Manager Plug-in Users**
>
> If you use Oracle Enterprise Manager and the Oracle Enterprise Manager Plug-in to monitor your Oracle PCA 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 PCA software update is complete.
>
> For detailed instructions, refer to the Agent Recovery section in the Oracle Enterprise Manager Plug-indocumentation.

> **Warning**
>
> An update of the Oracle PCA 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.

> **Warning**
>
> When updating the Oracle PCA software, make sure that no provisioning operations occur and that any externally scheduled backups are suspended. Such operations could prevent a software update or component firmware upgrade to fail and lead to system downtime.

2.10.1 Mandatory Updates

It is expected that not every Oracle PCA system is updated with each minor release. Updates are often selected because they fulfill certain requirements of the environment where the appliance is deployed. However, certain updates bring such significant changes to the product that they must be applied if they are on the path between a system's current software and the intended newer version. This section describes the mandatory updates and the impact they have on the Oracle PCA going forward. A more detailed description of each release can be found in the *Oracle Private Cloud Appliance Release Notes*.

Release 1.1.3

If your Oracle PCA is currently running Controller Software Release 1.0.2, please upgrade to Release 1.1.3 first. Critical changes include:

- Oracle VM performance tuning

- storage connectivity improvements

- support for a new generation of hardware: Sun Server X4-2 and Oracle ZFS Storage Appliance ZS3-ES

- OpenSSL upgrade to eliminate the 'heartbleed bug' (CVE-2014-0160)

> **Warning**
>
> If you are using a VLAN for your management network and you are upgrading from a previous release where you manually configured your environment, these network changes can adversely affect the upgrade process. Any manual configuration

> to allow the management network to run on a VLAN must be reverted on both management nodes, before the software process is started.

To update the appliance software to Release 1.1.3 you must first execute the procedure in Section 2.10.2, "Updating the Oracle PCA Controller Software". After successful completion, perform the additional steps described in these Troubleshooting sections:

- Section 7.5, "Verifying and Re-applying Oracle VM Manager Tuning after Software Update"

- Section 7.6, "Upgrading Existing Compute Node Configuration from Release 1.0.2"

Release 2.0.1

If your Oracle PCA is currently running Controller Software Release 1.1.3, please upgrade to Release 2.0.1 first. Critical changes include:

- external storage support

- appliance management over a VLAN

- Oracle VM upgrade to Release 3.2.8

> **Warning**
>
> If you are using a VLAN for your management network and you are upgrading from a previous release where you manually configured your environment, these network changes can adversely affect the upgrade process. Any manual configuration to allow the management network to run on a VLAN must be reverted on both management nodes, before the software process is started.
>
> As of Release 2.0.1, the Oracle PCA Dashboard provides an option to automatically configure the management network for a VLAN. After the upgrade is completed, you need to set this configuration option to enable VLAN support for the management network. See Section 2.5, "Network Settings" for more information on setting this configuration option.

To update the appliance software to Release 2.0.1 you must first execute the procedure in Section 2.10.2, "Updating the Oracle PCA Controller Software". After successful completion, perform the additional procedure described in Section 2.10.3, "Upgrading the Virtualization Platform".

Release 2.0.5

If your Oracle PCA is currently running Controller Software Release 2.0.1-2.0.4, please upgrade to Release 2.0.5 first. Critical changes include:

- product name change from Oracle Virtual Compute Appliance (VCA) to Oracle Private Cloud Appliance (PCA)

- Oracle Linux kernel upgrade

- support for a new generation of hardware: Oracle Server X5-2 (expansion compute nodes and base racks)

- Oracle VM upgrade to Release 3.2.9

To update the appliance software to Release 2.0.5 you must first execute the procedure in Section 2.10.2, "Updating the Oracle PCA Controller Software". After successful completion, perform the additional procedure described in Section 2.10.3, "Upgrading the Virtualization Platform".

Release 2.2.1

If your Oracle PCA is currently running Controller Software Release 2.0.5, you can upgrade to Release 2.1.1 or Release 2.2.1. Critical changes include:

- In Release 2.1.1:

  - new network customization and monitoring functionality

  - firmware upgrades for network components

  - automated ZFS storage appliance firmware upgrade, including network access to both controllers for easier maintenance

  - support for compute node memory extension kits

  - Oracle VM upgrade to Release 3.2.10

- In Release 2.2.1:

  - support for multiple server pools (tenant groups)

  - support for expansion racks

  - extended network customization functionality

  - support for a new generation of hardware: Oracle Server X6-2 (compute node only)

  - Oracle VM upgrade to Release 3.2.11

To update the appliance software to Release 2.2.1 you must first execute the procedure in Section 2.10.2, "Updating the Oracle PCA Controller Software". After successful completion, perform the additional procedure described in Section 2.10.3, "Upgrading the Virtualization Platform".

2.10.2 Updating the Oracle PCA Controller Software

Updates of the Oracle PCA controller software are performed from the Command Line Interface of the master management node. Software updates are a three-phase process. First, a zipped ISO containing the updated software must be downloaded from My Oracle Support and made available on an HTTP or FTP server. From there, the ISO is downloaded to the Oracle PCA internal storage appliance. When the download is complete and the software is unpacked in the appropriate directories, the update is activated and applied to each affected component.

> **Caution**
>
> 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 PCA 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 7.1, "Adding Proxy Settings for Oracle Private Cloud Appliance Updates" for more information.

> **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.

Optional: Rebooting the Management Node Cluster

Although not strictly necessary, it may be useful 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.

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 PCA Dashboard. If you use two
   separate consoles you can view both side by side.

   > **Note**
   >
   > The default `root` password is *Welcome1*.

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.

Before You Begin the Oracle Private Cloud Appliance Software Update

Read and understand this additional information before you start updating the Oracle PCA software. It
explains which commands to use depending on your specific environment and how to monitor the update
process in case it takes longer than expected or does not complete at all.

> **Update Commands**
>
> Current update commands are part of the Oracle PCA CLI. If your appliance is
> running Controller Software Release 2.0.5 or earlier, you must use the `pca-`
> `updater` or `ovca-updater` commands available from the management node
> command prompt. Add `--help` to display usage options.
>
> For detailed update instructions when running Release 2.0.5 or earlier, refer
> to the *Oracle Private Cloud Appliance Administrator's Guide* Release 2.1.1.
> All documentation for Release 2.1 can be found at http://docs.oracle.com/cd/
> E60765_01/.

Monitoring the Update Process

In addition to the command line method described in the update procedure, you can use additional terminal windows to monitor the progress of the software update. Specifically, if you have an active ssh connection to management node 1, you can watch the update run on the ILOM of management node 2, and vice versa.

1. Open a terminal session and ssh into the active management node. From there, open another terminal session and connect to the secondary management node ILOM, which is updated first when you start the update. You can access the ILOMs over the internal appliance Ethernet network (`192.168.4.0/24`).

   Note

   The internal IP addresses are assigned as follows:

   - The internal host IP and ILOM IP of management node 1 are: `192.168.4.3` and `192.168.4.103`.

   - The internal host IP and ILOM IP of management node 2 are: `192.168.4.4` and `192.168.4.104`.

   ```
   ssh root@10.100.1.101
   root@10.100.1.101's password:

   root@ovcamn05r1 ~]# ssh 192.168.4.104
   Password:

   Oracle(R) Integrated Lights Out Manager
   Version 3.2.4.20.a r94217
   Copyright (c) 2014, Oracle and/or its affiliates. All rights reserved.
   Hostname: ilom-ovcamn06r1

   -> start /SP/console
   Are you sure you want to start /SP/console (y/n)? y
   Serial console started.  To stop, type ESC (
   ```

2. Start the ILOM console.

   ```
   -> start /SP/console
   Are you sure you want to start /SP/console (y/n)? y
   Serial console started.  To stop, type ESC (
   ```

   Messages from the BIOS and from the Oracle Linux and Oracle PCA installations appear in the console. Several reboots occur during the update process. Toward the end of the process a message appears that indicates the system is ready for customer use on the next reboot. At this point your terminal sessions are disconnected. You can log on to the other management node, which has taken over the master role, and follow the second management node update by connecting to its ILOM.

   If you were connected to `192.168.4.3` and `192.168.4.104` at first, then connect to `192.168.4.4` and `192.168.4.103`. Depending on which management node held the master role before the update, you may need to connect in reverse order.

Updating the Oracle Private Cloud Appliance Software

1. Log into My Oracle Support and download the required Oracle PCA 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.

2. Make the update, a zipped ISO, available on an HTTP or FTP server that is reachable from your Oracle PCA.

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

> **Note**
>
> The default `root` password is *Welcome1*.

4. Connect to the management node using its IP address in the data center network, as you configured it in the Network Setup tab of the Oracle PCA Dashboard. For details, see Section 2.5, "Network Settings".

> **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 ~]#
```

5. Launch the Oracle PCA command line interface.

```
# pca-admin
Welcome to PCA! Release: 2.1.1
PCA>
```

6. Download the ISO to your Oracle PCA. Confirm that you want to start the download.

```
PCA> update appliance get_image http://myserver.org/images/pca-2.2.1-b999.iso.zip
************************************************************
 WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
************************************************************
Are you sure [y/N]:y

The update job has been submitted. Use "show task <task id>" to monitor the progress.

Task_ID          Status   Progress Start_Time           Task_Name
-------          ------   -------- ----------           ---------
333dcc8b617f74   RUNNING  None      05-17-2016 09:06:29  update_download_image
--------------
1 row displayed

Status: Success
```

7. Check the progress of the ISO download. When the download is finished, proceed with the next step.

```
PCA> show task 333dcc8b617f74

--------------------------------------
Task_Name           update_download_image
Status              SUCCESS
Progress            100
Start_Time          05-17-2016 09:06:29
End_Time            05-17-2016 09:13:11
Pid                 459257
Result              None
--------------------------------------
```

```
Status: Success
```

> **Note**
>
> After download, the image is unpacked and the files are copied to the `/nfs/shared_storage` directory, which is an NFS mount from the appliance internal storage on both management nodes.

8. When the download has completed successfully, activate it by launching the update process.

```
PCA> update appliance install_image
************************************************************
 WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
************************************************************
Are you sure [y/N]:y
Status: Success
```

Once you issue this command, the update process is initiated as described in Section 1.7, "Oracle Private Cloud Appliance Software Update".

9. Check the progress of the software update.

```
PCA> list update-task
Mgmt_Node_IP    Update_Started        Update_Ended         Elapsed    Update status
-----------     --------------        -----------          -------    -------------
192.168.4.4     05-17-2016 13:08:09   05-17-2016 13:49:04  0:40:55    Succeeded
192.168.4.3     05-17-2016 13:55:41   05-17-2016 14:37:10  0:41:29    Succeeded
----------------
2 rows displayed
Status: Success
```

> **Note**
>
> At a certain point during the update process, the active management node is rebooted. As a result, the SSH connection is lost. In addition, this may cause the Dashboard to become unresponsive temporarily, and you may be required to log back in.

When the master management node reboots, the secondary (updated) management node assumes the master role. The original master management node is then also updated and becomes the backup management node.

The software update process is automated to a certain degree, triggering a number of configuration tasks as it progresses. If any of these tasks should fail, the system writes entries in the error database and attempts to restart them every five minutes over a period of 30 minutes. At any given time the administrator can use the CLI to check for configuration task errors and rerun them if necessary. For details about these particular commands, see Section 3.2.21, "rerun".

> **Caution**
>
> 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.

2.10.3 Upgrading the Virtualization Platform

Some releases of the Oracle PCA Controller Software include a new version of Oracle VM. As part of the controller software update, the new Oracle VM Manager Release is automatically installed on both management nodes, and a new YUM repository configuration is pushed to all compute nodes to point them to the packages of the new Oracle VM Server Release. As a result, all compute nodes in the Oracle PCA rack are marked as "out of date" in Oracle VM Manager. The *Server Out Of Date* event is informational and does 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.

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. At that stage, the YUM repository configuration of all compute nodes has been updated to point to the new Oracle VM Server packages on the Oracle PCA internal storage.

Note

A large part of this procedure must be executed from Oracle VM Manager. An overview of its usage in the context of Oracle PCA, and links to detailed instructions in the Oracle VM documentation, are provided in this Administrator's Guide, in Chapter 4, *Managing the Oracle VM Virtual Infrastructure*.

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 PCA CLI:

   ```
   # pca-admin
   Welcome to PCA! Release: 2.2.1

   PCA> show version

   ---------------------------------------
   Version            2.2.1
   Build              999
   Date               2016-05-17
   ---------------------------------------

   Status: Success
   ```

2. Log in to Oracle VM Manager.

   For details, see Section 4.2, "Logging in to the Oracle VM Manager Web UI".

3. Identify the master server in the server pool.

   a. In the **Servers and VMs** tab, select the server pool in the navigation pane.

   b. In the management pane, select the **Info** perspective from the list in the toolbar.

      A configuration overview of the server pool is displayed.

   c. Locate the *Master Server* setting and write down which server is the pool master.

      > **Caution**
      >
      > The server pool master server should be the last Oracle VM Server to go through this upgrade procedure. This allows you to upgrade the compute nodes one by one without causing the master role to be transferred to another node in the server pool each time. This eliminates potential resource locks and errors related to transferring the master role, and saves a significant amount of time.

4. 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 User's Guide section entitled Migrating Virtual Machines.

5. Place the compute node in maintenance mode.

   Information on maintenance mode is provided in the Oracle VM User's Guide section entitled Placing an Oracle VM Server into 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. 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.

6. Run the Oracle VM Server update on the compute node in question.

   Information on the Oracle VM Server update functionality is provided in the Oracle VM User's Guide section entitled Updating and Upgrading Oracle VM Servers.

   a. In the **Servers and VMs** tab, select the server pool in the navigation pane.

   b. In the management pane, select the **Servers** perspective from the list in the toolbar.

      A table is displayed with all servers in the server pool.

   c. Select the appropriate compute node in the management pane table and click **Update Server** in the management pane toolbar.

      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.

7. 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 **Maintenance Mode** check box. Click OK.

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

8. Repeat this procedure for each compute node in your Oracle PCA.

Caution

The server pool master server should be the last Oracle VM Server to go through this upgrade procedure. This allows you to upgrade the compute nodes one by one without causing the master role to be transferred to another node in the server pool each time. This eliminates potential resource locks and errors related to transferring the master role, and saves a significant amount of time.

The appliance software update is now complete. The Oracle PCA is ready to resume all normal operations.

2.10.4 Upgrading Component Firmware

All the software components in a given Oracle PCA 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 PCA `.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.

If a specific or additional procedure to upgrade the firmware of an Oracle PCA 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 PCA Documentation Library.

2.10.4.1 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

Detailed information about firmware upgrades can be found in the *Sun Datacenter InfiniBand Switch 36 Product Notes for Firmware Version 2.0* (document ID: E26702). Refer to the section "Upgrade the Switch Firmware (CLI)".

Upgrading the InfiniBand Switch 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 PCA software image to the Yum repository share.

```
root@ovcamn05r1 ~]# cp
/nfs/shared_storage/mgmt_image/firmware/IB_gateway/NM2-36P/SUN_DCS_36p_2.1.5-1/ \
SUN_DCS_36p/sundcs_36p_repository_2.1.5_1.pkg /nfs/shared_storage/yum/
```

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:
Last login: Thu Oct 15 09:01:25 2015 from 192.168.4.3
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
20150930 03:41:46 Master SubnetManager on sm lid 39 sm guid 0x139702010017b4 :
MT25408 ConnectX Mellanox Technologies
```

> ⊗ **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: 2.1.2-2
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/sundcs_36p_repository_2.1.5_1.pkg

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 2.0.8-1:
Upgrade started...
Upgrade completed.
INFO: SUN DCS 36p upgrade from 1.3.3-2 to 2.0.8-1 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: 2.1.5-1
Build time: Oct  6 2014 10:35:15
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
20150930 03:41:46 Master SubnetManager on sm lid 39 sm guid 0x139702010017b4 :
MT25408 ConnectX Mellanox Technologies
```

> **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.

2.10.4.2 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**
>
> Detailed information about firmware upgrades can be found in the *XgOS User's Guide* (document ID: E53170). Refer to the section "System Image Upgrades".

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 PCA software image to the Yum repository share.

```
root@ovcamn05r1 ~]# cp /nfs/shared_storage/mgmt_image/firmware/IB_gateway/ \
OFI/xsigo-3.9.4-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

OpenSM $
admin@ovcasw22r1[xsigo]
```

6. Start the system upgrade procedure.

```
admin@ovcasw22r1[xsigo] system upgrade
http://192.168.4.1/shares/export/Yum/xsigo-3.9.4-XGOS.xpf
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
Are you sure you want to update the software (y/n)? y
Running preunpack scripts...
Installing...
################################################################
[100%]
Verifying...
################################################################
[100%]
Running preinstall scripts...
Installing package...
Running postinstall scripts...
Installation successful. Please stand by for CLI restart.
admin@iowa[xsigo]
XgOS CLI is restarting - This might take a couple of minutes...
*01:00
System services are available again. Restarting the CLI now.
```

7. 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 3.9.4-XGOS  - (buildsys) Thu Mar 19 03:25:26 UTC 2015

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
```

8. 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 3 The Oracle Private Cloud Appliance Command Line Interface (CLI)

Table of Contents

All Oracle Private Cloud Appliance (PCA) 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 PCA 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/E35328_01/E35336/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.

3.1 CLI Usage

The Oracle PCA command line interface is triggered by running the `pca-admin` command. It can run either in interactive mode (see Section 3.1.1, "Interactive Mode") or in single-command mode (see Section 3.1.2, "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 PCA 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 3.1.3, "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 3.1.4, "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 3.1.1.1, "Tab Completion" for more information on this.

3.1.1 Interactive Mode

The Oracle PCA 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 are able to enter commands specific to the management of the Oracle PCA. 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 3.1 An example of interactive mode usage of the CLI

```
# pca-admin
Welcome to PCA! Release: 2.1.1
PCA>
```

```
PCA> list management-node
Management_Node  IP_Address  Provisioning_Status  ILOM_MAC            Provisioning_State  Master
---------------  ----------  -------------------  --------            ------------------  ------
ovcamn05r1       192.168.4.3 RUNNING              00:10:e0:65:30:5f   running             Yes
ovcamn06r1       192.168.4.4 RUNNING              00:10:e0:65:30:65   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.

3.1.1.1 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, 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 3.2 Examples showing tab-completion

```
PCA> <tab>
EOF          backup        create        diagnose      get           list          quit
reprovision  set           show          stop          add           configure     delete
exit         help          q             remove        rerun         shell         start
update

PCA> list <tab>
compute-node     lock              network          network-port      opus-port
storage-network  tenant-group      wwpn-info        config-error      management-node
network-card     ofm-network       server-profile   task              update-task

PCA> list m<tab>anagement-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.

3.1.1.2 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
Mon Jun 27 09:49:50 UTC 2016
PCA> !uptime > /tmp/uptime-today
PCA> !rm /tmp/uptime-today
```

3.1.2 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 3.1.3.1, "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
------------  ----------    ----------           --------          ----------
ovcacn08r1    192.168.4.9   RUNNING              00:10:e0:65:2f:b7  running
ovcacn28r1    192.168.4.10  RUNNING              00:10:e0:62:31:81  running
ovcacn10r1    192.168.4.7   RUNNING              00:10:e0:65:2f:cf  running
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

#
```

3.1.3 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.

3.1.3.1 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:e0:65:2f:46": {
        "name": "ovcacn10r1",
        "ilom_state": "running",
        "ip": "192.168.4.7",
        "state": "RUNNING",
        "networks": "mgmt_public_eth, extnet, vm_private, vm_public_vlan",
        "ilom_mac": "00:10:e0:65:2f:4b"
    },
    "00:10:e0:65:2f:ee": {
```

```
        "name": "ovcacn08r1",
        "ilom_state": "running",
        "ip": "192.168.4.5",
        "state": "RUNNING",
        "networks": "intnet, mgmt_public_eth, extnet, vm_public_vlan, extnet6, vm_private",
        "ilom_mac": "00:10:e0:65:2f:f3"
    },
    "00:10:e0:62:98:de": {
        "name": "ovcacn09r1",
        "ilom_state": "running",
        "ip": "192.168.4.10",
        "state": "RUNNING",
        "networks": "mgmt_public_eth, extnet, vm_private, vm_public_vlan, intnet",
        "ilom_mac": "00:10:e0:62:98:e3"
    },
    "00:10:e0:65:2f:8e": {
        "name": "ovcacn07r1",
        "ilom_state": "running",
        "ip": "192.168.4.8",
        "state": "RUNNING",
        "networks": "intnet, mgmt_public_eth, extnet, vm_public_vlan, pvi_network, extnet6, vm_private",
        "ilom_mac": "00:10:e0:65:2f:93"
    }
}
```

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.

3.1.3.2 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 the `--sorted-by` option is case sensitive, but you can use tab-completion to easily obtain the options for different column names. See Section 3.1.1.1, "Tab Completion" for more information.

3.1.3.3 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 case insensitivity, 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 the `--filter-column` option is case sensitive, but you can use tab-completion to easily obtain the options for different column names. See Section 3.1.1.1, "Tab Completion" for more information.

3.1.4 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      configure  delete    get    list     reprovision  set    start   update
backup   create     diagnose  help   remove   rerun             show   stop

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 PCA, 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.
  network              List netbundle networks.
```

```
  network-card          List network card.
  network-port          List network port.
  ofm-network           List OFM networks.
  opus-port             List OPUS port.
  server-profile        List server profile.
  storage-network       List storage networks.
  task                  List task.
  tenant-group          List tenant-group.
  update-task           List update task.
  wwpn-info             List wwpn info.

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 ovcacn11r1

  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', 'config-error', 'management-node', 'lock',
'network-port', 'server-profile', 'tenant-group', 'storage-network', 'opus-port', 'task',
'compute-node', 'ofm-network', 'wwpn-info', 'network-card', 'network'].
```

3.2 CLI Commands

This section describes all of the documented commands available via the CLI.

3.2.1 add network

Connects a server node to an existing network. To create a new custom network, see Section 3.2.7, "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
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.3 Connecting a Compute Node to a Custom Network

```
PCA> add network MyNetwork ovcacn09r1
Status: Success
```

3.2.2 add network-to-tenant-group

Associates a custom network with an existing tenant group. To create a new tenant group, see Section 3.2.8, "create tenant-group". To create a new custom network, see Section 3.2.7, "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.

Options

The following table shows the available options for this command.

Option	Description
`--json`	Return the output of the command in JSON format

Option	Description
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.4 Associating a Custom Network with a Tenant Group

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

3.2.3 add server

Adds a server to an existing tenant group. To create a new tenant group, see Section 3.2.8, "create tenant-group".

Syntax

```
add server 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 servers to, and `node` is the name of the server node that should be added to the selected tenant group.

Description

Use the `add server` 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
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.5 Adding a Compute Node to a Tenant Group

```
PCA> add server ovcacn09r1 myTenantGroup
Status: Success
```

3.2.4 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 PCA 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, 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 3.41, "Show Task" for more information.

Options

The following table shows the available options for this command.

Option	Description
--json	Return the output of the command in JSON format
--less	Return the output of the command one screen at a time for easy viewing, as with the less command on the Linux command line. This option allows both forward and backward navigation through the command output.
--more	Return the output of the command one screen at a time for easy viewing, as with the more command on the Linux command line. This option allows forward navigation only.
--tee=OUTPUTFILENAME	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.6 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
```

```
-------         ------  -------- ----------          ---------
341e7bc74f339c  RUNNING None     06-27-2016 09:59:36  backup
---------------
1 row displayed

Status: Success
```

3.2.5 configure vhbas

Configures vHBAs on compute nodes.

Syntax

configure vhbas { 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 vhbas 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
ALL | node	Configure vHBAs for all compute nodes or for one or more specific compute nodes.
--json	Return the output of the command in JSON format
--less	Return the output of the command one screen at a time for easy viewing, as with the less command on the Linux command line. This option allows both forward and backward navigation through the command output.
--more	Return the output of the command one screen at a time for easy viewing, as with the more command on the Linux command line. This option allows forward navigation only.
--tee=OUTPUTFILENAME	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.7 Configuring the vHBAs for Specific Compute Nodes

```
PCA> configure vhbas ovcacn11r1 ovcacn14r1
Compute_Node        Status
------------        ------
ovcacn14r1          Succeeded
ovcacn11r1          Succeeded
----------------
2 rows displayed
Status: Success
```

3.2.6 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 | database | install | manufacturing | 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
`all_provisioning`	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.
`database`	Impose a lock on the databases during the management node update process. The lock is released after the update.
`install`	Placeholder lock type. Currently not used.
`manufacturing`	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 `ovca` service does not start.
`provisioning`	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.
`service`	Placeholder lock type. Behavior is identical to manufacturing lock.
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.

Option	Description
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.8 Imposing a Provisioning Lock

```
PCA> create lock provisioning
Status: Success
```

3.2.7 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 'ports' | host_network 'ports' 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 Fabric Interconnect ports where the network terminates must also be specified as `'ports'`.

If the network type is `host_network`, then additional arguments are expected. The subnet agruments 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
{ `rack_internal_network` | `external_network` | `host_network` }	The type of custom network to create. The options are: • a network internal to the rack

Option	Description
	• a network with external connectivity
	• a network with external connectivity, accessible for physical hosts
`external_network 'ports'`	To create a custom network with external connectivity, you must specify the ports on the Fabric Interconnect as well. Ports are identified by their I/O module number and port number, separated by a colon. Put the port identifiers between quotes as a space-separated list, for example: `'4:2 5:2'`.
`host_network 'ports' prefix netmask [route-destination gateway]`	To create a custom host network, you must specify the ports on the Fabric Interconnect as with an external network. Ports are identified by their I/O module number and port number, separated by a colon. Put the port identifiers between quotes as a space-separated list, for example: `'4:2 5:2'`. 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.
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.9 Creating an Internal Custom Network

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

Example 3.10 Creating a Custom Network with External Connectivity

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

3.2.8 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 servers using the `add server` 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
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.11 Creating a Tenant Group

```
PCA> create tenant-group myTenantGroup

Status: Success
```

3.2.9 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 3.31, "List All Configuration Errors" for more information.

Options

The following table shows the available options for this command.

Option	Description
`--confirm`	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
`--force`	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.
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.12 Removing a Configuration Error

```
PCA> delete config-error 87
************************************************************
 WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
************************************************************
Are you sure [y/N]:y
Status: Success
```

3.2.10 delete lock

Removes a lock that was previously imposed on certain appliance functionality.

Syntax

delete lock { all_provisioning | database | install | manufacturing | 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
{ `all_provisioning` | `database` | `install` | `manufacturing` | `provisioning` | `service` }	The type of lock to be removed. For a description of lock types, see Section 3.2.6, "create lock".

Option	Description
`--confirm`	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
`--force`	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.
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.13 Unlocking Provisioning

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

3.2.11 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 3.2.17, "remove network".

Default Oracle PCA networks are protected and any attempt to delete them will fail.

Options

The following table shows the available options for this command.

Option	Description
--confirm	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
--force	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.
--json	Return the output of the command in JSON format
--less	Return the output of the command one screen at a time for easy viewing, as with the less command on the Linux command line. This option allows both forward and backward navigation through the command output.
--more	Return the output of the command one screen at a time for easy viewing, as with the more command on the Linux command line. This option allows forward navigation only.
--tee=OUTPUTFILENAME	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.14 Deleting a Custom Network

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

Example 3.15 Attempting to Delete a Default Network

```
PCA> delete network vm_private
Status: Failure
Error Message: Error (NETWORK_003): Exception while deleting network: vm_private.
['INVALID_NAME_002: Invalid Network name: vm_private. Name is reserved.']
```

3.2.12 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
`--confirm`	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
`--force`	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.
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.16 Removing a Task

```
PCA> delete task 341e7bc74f339c
***************************************************************
 WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
***************************************************************
Are you sure [y/N]:y
Status: Success
```

3.2.13 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 and the virtual IP address becomes available for re-use by another 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
`--confirm`	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
`--force`	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.
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.17 Deleting a Tenant Group

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

Status: Success
```

3.2.14 diagnose

Performs various diagnostic checks against the Oracle Private Cloud Appliance for support purposes.

Syntax

diagnose { ilom | software } [--json] [--less] [--more] [--tee=OUTPUTFILENAME]

Description

Use the `diagnose` command to initiate a diagnostic check of various components that make up Oracle PCA.

Options

The following table shows the available options for this command.

Option	Description
`{ ilom	software }`	Specify the type of diagnostic check that should be performed as the command target. The `ilom` diagnostic checks that the ILOM for each component is accessible on the management network. The `software` diagnostic triggers the Oracle PCA software acceptance tests.

Option	Description
--json	Return the output of the command in JSON format
--less	Return the output of the command one screen at a time for easy viewing, as with the less command on the Linux command line. This option allows both forward and backward navigation through the command output.
--more	Return the output of the command one screen at a time for easy viewing, as with the more command on the Linux command line. This option allows forward navigation only.
--tee=OUTPUTFILENAME	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.18 Running the ILOM Diagnostic

```
PCA> diagnose ilom
Checking ILOM health............please wait..

IP_Address      Status          Health_Details
----------      ------          --------------
192.168.4.129   Not Connected
192.168.4.128   Not Connected
192.168.4.127   Not Connected
192.168.4.126   Not Connected
192.168.4.125   Not Connected
192.168.4.124   Not Connected
192.168.4.123   Not Connected
192.168.4.122   Not Connected
192.168.4.121   Not Connected
192.168.4.120   Not Connected
192.168.4.101   OK
192.168.4.102   OK
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
192.168.4.106   OK
192.168.4.109   OK
192.168.4.108   OK
192.168.4.112   OK
192.168.4.113   Not Connected
192.168.4.110   OK
192.168.4.111   OK
192.168.4.116   Not Connected
192.168.4.117   Not Connected
192.168.4.114   Not Connected
192.168.4.115   Not Connected
192.168.4.118   Not Connected
192.168.4.119   Not Connected
----------------
27 rows displayed

Status: Success
```

Example 3.19 Running the Software Diagnostic

```
PCA> diagnose software
PCA Software Acceptance Test runner utility
```

```
Test -    701 - OpenSSL CVE-2014-0160 Heartbleed bug Acceptance [PASSED]
Test -    785 - PCA package Acceptance [PASSED]
Test -   1083 - Mgmt node xsigo network interface Acceptance [PASSED]
Test -    787 - Shared Storage Acceptance [PASSED]
Test -    973 - Simple connectivity Acceptance [PASSED]
Test -   1078 - Test for ovs-agent service on CNs Acceptance [PASSED]
Test -   1079 - Test for shares mounted on CNs Acceptance [PASSED]
Test -   1080 - ovs-log check Acceptance [PASSED]
Test -    788 - PCA services Acceptance [PASSED]
Test -    789 - PCA config file Acceptance [PASSED]
Test -   1300 - All compute nodes running Acceptance [PASSED]
Test -   1318 - Check support packages in PCA image Acceptance [PASSED]
Test -    928 - Repositories defined in OVM manager Acceptance [PASSED]
Test -   1107 - Compute node xsigo network interface Acceptance [PASSED]
Test -   1316 - PCA version Acceptance [PASSED]
Test -   1117 - Network interfaces check Acceptance [PASSED]
Test -    824 - OVM manager settings Acceptance [PASSED]
Test -    927 - OVM server model Acceptance [PASSED]
Test -    925 - PCA log Acceptance [PASSED]
Test -    926 - Networks defined in OVM manager for CNs Acceptance [PASSED]
Test -    822 - Compute node network interface Acceptance [PASSED]
Status: Success
```

3.2.15 get log

Retrieves the log files from the selected components and saves them to a temporary directory.

> **Note**
>
> Currently the Oracle Fabric Interconnect F1-15 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 the temporary directory on the management node where you run the command. 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
--json	Return the output of the command in JSON format
--less	Return the output of the command one screen at a time for easy viewing, as with the less command on the Linux command line. This option allows both forward and backward navigation through the command output.
--more	Return the output of the command one screen at a time for easy viewing, as with the more command on the Linux command line. This option allows forward navigation only.

Option	Description
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.20 Collecting the Log Files from the Fabric Interconnects

Note that the CLI uses 'xsigo' as the internal alias for an Oracle Fabric Interconnect F1-15.

```
PCA> get log xsigo
Log files to be found: /tmp/xsigo-log*
Status: Success
```

3.2.16 list

The `list` command can be used to list the different components and tasks within the Oracle PCA. 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` { `compute-node` | `config-error` | `lock` | `management-node` | `network` | `network-card` | `network-port` | `ofm-network` | `opus-port` | `server-profile` | `storage-network` | `task` | `tenant-group` | `update-task` | `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 3.1.3.2, "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 3.1.3.3, "Filtering" for more information.

The following table describes each possible target of the `list` command.

Command Target	Information Displayed
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.
network	Displays all networks configured in the environment.
network-card	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 Fabric Interconnects.

Command Target	Information Displayed
ofm-network	Displays network configuration, read directly from the Oracle Fabric Manager software on the Fabric Interconnects.
opus-port	Displays connection information about every port of every Oracle Switch ES1-24 in the Oracle PCA environment.
server-profile	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	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.
update-task	Displays a list of all software update tasks that have been started on the appliance.
wwpn-info	Displays a list of all World Wide Port Names (WWPNs) for all ports participating in the Oracle PCA 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 PCA. 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
{ `compute-node` | `config-error` | `lock` | `management-node` | `network` | `network-card` | `network-port` | `ofm-network` | `opus-port` | `server-profile` | `storage-network` | `task` |	The command target to list information for.

Option	Description
`tenant-group` | `update-task` | `wwpn-info` }	
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.
[--sorted-by `SORTEDBY`]	Sort the table by the values within a particular column in the table, specified by replacing `SORTEDBY` with the name of the column that should be used to perform the sort.
[--sorted-order `SORTEDORDER`]	Used to specify the sort order, which can either be `ASC` for an ascending sort, or `DES` for a descending sort. You must use the `--sorted-by` option in conjunction with this option.
[--filter-column `FILTERCOLUMN`]	Filter the table for a value within a particular column in the table, specified by replacing `FILTERCOLUMN` with the name of the column that should be used to perform the sort. You must use the `--filter` option in conjunction with this option.
[--filter `FILTER`]	The filter that should be applied to values within the column specified by the `--filter-column` option.

Note

When the target of the `list` command is either `management-node` or `compute-node`, the Node_State column is used to show the software completeness of a node through the provisioning process. It does not indicate the system status of a node.

Examples

Example 3.21 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:e0:65:30:5f  running             Yes
ovcamn06r1       192.168.4.4 RUNNING              00:10:e0:65:30:65  running             None
----------------
2 rows displayed

Status: Success
```

Example 3.22 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 3.23 List All Tenant Groups

```
PCA> list tenant-group

Name                 Default    State
----                 -------    -----
Rack1_ServerPool     True       ready
myTenantGroup        False      ready
----------------
2 rows displayed

Status: Success
```

Example 3.24 List Appliance Networks

```
PCA> list network

Network_Name                     Trunkmode    Description
----------                       ---------    -----------
vm_public_vlan                   True         2014-07-12 02:11:27
mgmt_public_eth                  True         2014-07-12 02:09:42
mgmt_pvi                         False        2014-07-12 02:08:38
vm_private                       False        2014-07-12 02:10:35
discovered-network               False
----------------
5 rows displayed

Status: Success
```

Example 3.25 List the Network Ports Configured on the Fabric Interconnects

```
PCA> list network-port

Port      Director     Type                 State     Networks
----      --------     ----                 -----     --------
3:2       ovcasw15r1   sanFc8GbPort         down      None
3:1       ovcasw15r1   sanFc8GbPort         down      None
4:1       ovcasw15r1   nwEthernet10GbPort   up        mgmt_public_eth, vm_public_vlan
4:4       ovcasw15r1   nwEthernet10GbPort   down      None
4:3       ovcasw15r1   nwEthernet10GbPort   down      None
4:2       ovcasw15r1   nwEthernet10GbPort   down      None
5:1       ovcasw15r1   nwEthernet10GbPort   up        mgmt_public_eth, vm_public_vlan
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
12:2      ovcasw15r1   sanFc8GbPort         down      None
12:1      ovcasw15r1   sanFc8GbPort         down      None
3:2       ovcasw22r1   sanFc8GbPort         down      None
3:1       ovcasw22r1   sanFc8GbPort         down      None
4:1       ovcasw22r1   nwEthernet10GbPort   up        mgmt_public_eth, vm_public_vlan
```

```
4:4        ovcasw22r1       nwEthernet10GbPort       down        None
4:3        ovcasw22r1       nwEthernet10GbPort       down        None
4:2        ovcasw22r1       nwEthernet10GbPort       down        None
5:1        ovcasw22r1       nwEthernet10GbPort       up          mgmt_public_eth, vm_public_vlan
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
10:2       ovcasw22r1       nwEthernet10GbPort       up          None
10:1       ovcasw22r1       nwEthernet10GbPort       down        None
11:4       ovcasw22r1       nwEthernet10GbPort       up          None
11:3       ovcasw22r1       nwEthernet10GbPort       down        None
11:2       ovcasw22r1       nwEthernet10GbPort       down        None
11:1       ovcasw22r1       nwEthernet10GbPort       down        None
12:2       ovcasw22r1       sanFc8GbPort             down        None
12:1       ovcasw22r1       sanFc8GbPort             down        None
-----------------
40 rows displayed

Status: Success
```

Example 3.26 List the I/O Modules installed in the Fabric Interconnects

```
PCA> list network-card

Slot    Director       Type                                       State       Number_Of_Ports
----    --------       ----                                       -----       ---------------
4       ovcasw22r1     nwEthernet4Port10GbCardEthIb               up          4
10      ovcasw22r1     nwEthernet4Port10GbCardEthIb               up          4
3       ovcasw22r1     sanFc2Port8GbLrCardEthIb                   up          2
12      ovcasw22r1     sanFc2Port8GbLrCardEthIb                   up          2
5       ovcasw22r1     nwEthernet4Port10GbCardEthIb               up          4
11      ovcasw22r1     nwEthernet4Port10GbCardEthIb               up          4
12      ovcasw15r1     sanFc2Port8GbLrCardEthIb                   up          2
11      ovcasw15r1     nwEthernet4Port10GbCardEthIb               up          4
5       ovcasw15r1     nwEthernet4Port10GbCardEthIb               up          4
10      ovcasw15r1     nwEthernet4Port10GbCardEthIb               up          4
3       ovcasw15r1     sanFc2Port8GbLrCardEthIb                   up          2
4       ovcasw15r1     nwEthernet4Port10GbCardEthIb               up          4
-----------------
12 rows displayed

Status: Success
```

Example 3.27 List ports on each Oracle Switch ES1-24 using a Filter

Note that the CLI uses the internal alias for an Oracle Switch ES1-24, referencing an 'opus' switch or port. In this example the command displays all internal Ethernet connections to the Net Mgt port on the destination component. A wildcard is used in the `--filter` option.

```
PCA> list opus-port --filter-column=Dest_Port --filter=*Mgt

Dest        Dest_Port  Hostname     Key        OPUS        RACK   RU    Src_Port   Type
----        ---------  --------     ---        ----        ----   --    --------   ----
22          Net Mgt    ovcasw22r1   OPUS-1-9   OPUS-1      1      22    9          switch
19          Net Mgt    ovcasw19r1   OPUS-1-10  OPUS-1      1      19    10         switch
PDU-A       Net Mgt    ovcapoPDU-Ar1 OPUS-1-20 OPUS-1      1      0     20         power
15          Net Mgt    ovcasw15r1   OPUS-2-9   OPUS-2      1      15    9          switch
20          Net Mgt    ovcasw20r1   OPUS-2-10  OPUS-2      1      20    10         switch
PDU-B       Net Mgt    ovcapoPDU-Br1 OPUS-2-20 OPUS-2      1      0     20         power
16          Net Mgt    ovcasw16r2   OPUS-3-10  OPUS-3      2      16    10         switch
PDU-A       Net Mgt    ovcapoPDU-Ar2 OPUS-3-20 OPUS-3      2      0     20         power
20          Net Mgt    ovcasw20r2   OPUS-4-10  OPUS-4      2      20    10         switch
PDU-B       Net Mgt    ovcapoPDU-Br2 OPUS-4-20 OPUS-4      2      0     20         power
16          Net Mgt    ovcasw16r3   OPUS-5-10  OPUS-5      3      16    10         switch
```

```
PDU-A      Net Mgt   ovcapoPDU-Ar3   OPUS-5-20  OPUS-5    3    0    20       power
20         Net Mgt   ovcasw20r3      OPUS-6-10  OPUS-6    3    20   10       switch
PDU-B      Net Mgt   ovcapoPDU-Br3   OPUS-6-20  OPUS-6    3    0    20       power
----------------
14 rows displayed

Status: Success
```

Example 3.28 List All Tasks

```
PCA> list task
Task_ID          Status   Progress Start_Time             Task_Name
-------          ------   -------- ----------             ---------
341e7bc74f339c   SUCCESS       100 06-27-2016 09:59:36    backup
341e73748f5182   SUCCESS       100 06-27-2016 09:00:01    backup
341e0edfa4891e   SUCCESS       100 06-26-2016 21:00:01    backup
341daa4a2d3624   SUCCESS       100 06-26-2016 09:00:01    backup
341d45b5424c16   SUCCESS       100 06-25-2016 21:00:01    backup
341ce11fc6c39c   SUCCESS       100 06-25-2016 09:00:01    update_download_image
341c7c8afcc86a   SUCCESS       100 06-24-2016 21:00:02    backup
----------------
7 rows displayed

Status: Success
```

Example 3.29 List Storage Networks for External Fibre Channel Storage

```
PCA> list storage-network

Network_Name             Description
------------             -----------
Cloud_D                  Default Storage Cloud ru15 port2 - Do not delete or modify
Cloud_A                  Default Storage Cloud ru22 port1 - Do not delete or modify
Cloud_C                  Default Storage Cloud ru15 port1 - Do not delete or modify
Cloud_B                  Default Storage Cloud ru22 port2 - Do not delete or modify
----------------
4 rows displayed

Status: Success
```

Example 3.30 List WWPNs To Configure External Fibre Channel Storage

```
PCA> list wwpn-info

WWPN                    vHBA    Cloud_Name   Server        Type  Alias
----                    ----    ----------   ------        ----  -----
50:01:39:70:00:69:F1:06 vhba01  Cloud_A      ovcacn08r1    CN    ovcacn08r1-Cloud_A
50:01:39:70:00:69:F1:04 vhba01  Cloud_A      ovcacn09r1    CN    ovcacn09r1-Cloud_A
50:01:39:70:00:69:F1:08 vhba01  Cloud_A      ovcacn10r1    CN    ovcacn10r1-Cloud_A
50:01:39:70:00:69:F1:0C vhba01  Cloud_A      ovcacn32r1    CN    ovcacn32r1-Cloud_A
50:01:39:70:00:69:F1:0A vhba01  Cloud_A      ovcacn07r1    CN    ovcacn07r1-Cloud_A
50:01:39:70:00:69:F1:0E vhba01  Cloud_A      ovcacn31r1    CN    ovcacn31r1-Cloud_A
50:01:39:70:00:69:F1:07 vhba02  Cloud_B      ovcacn08r1    CN    ovcacn08r1-Cloud_B
50:01:39:70:00:69:F1:05 vhba02  Cloud_B      ovcacn09r1    CN    ovcacn09r1-Cloud_B
50:01:39:70:00:69:F1:09 vhba02  Cloud_B      ovcacn10r1    CN    ovcacn10r1-Cloud_B
50:01:39:70:00:69:F1:0B vhba02  Cloud_B      ovcacn07r1    CN    ovcacn07r1-Cloud_B
50:01:39:70:00:69:F1:0F vhba02  Cloud_B      ovcacn31r1    CN    ovcacn31r1-Cloud_B
50:01:39:70:00:69:F1:0D vhba02  Cloud_B      ovcacn32r1    CN    ovcacn32r1-Cloud_B
50:01:39:70:00:6A:11:0E vhba03  Cloud_C      ovcacn31r1    CN    ovcacn31r1-Cloud_C
50:01:39:70:00:6A:11:0A vhba03  Cloud_C      ovcacn07r1    CN    ovcacn07r1-Cloud_C
50:01:39:70:00:6A:11:0C vhba03  Cloud_C      ovcacn32r1    CN    ovcacn32r1-Cloud_C
50:01:39:70:00:6A:11:08 vhba03  Cloud_C      ovcacn10r1    CN    ovcacn10r1-Cloud_C
50:01:39:70:00:6A:11:04 vhba03  Cloud_C      ovcacn09r1    CN    ovcacn09r1-Cloud_C
50:01:39:70:00:6A:11:06 vhba03  Cloud_C      ovcacn08r1    CN    ovcacn08r1-Cloud_C
50:01:39:70:00:6A:11:0D vhba04  Cloud_D      ovcacn32r1    CN    ovcacn32r1-Cloud_D
50:01:39:70:00:6A:11:0F vhba04  Cloud_D      ovcacn31r1    CN    ovcacn31r1-Cloud_D
50:01:39:70:00:6A:11:0B vhba04  Cloud_D      ovcacn07r1    CN    ovcacn07r1-Cloud_D
```

```
50:01:39:70:00:6A:11:09    vhba04    Cloud_D          ovcacn10r1    CN    ovcacn10r1-Cloud_D
50:01:39:70:00:6A:11:05    vhba04    Cloud_D          ovcacn09r1    CN    ovcacn09r1-Cloud_D
50:01:39:70:00:6A:11:07    vhba04    Cloud_D          ovcacn08r1    CN    ovcacn08r1-Cloud_D
----------------
24 rows displayed

Status: Success
```

Example 3.31 List All Configuration Errors

```
PCA> list config-error

ID        Module                      Host              Timestamp
--        ------                      ----              ---------
87        Management node password    192.168.4.4       Thu Aug 21 02:45:42 2015
84        MySQL management password   192.168.4.216     Thu Aug 21 02:44:54 2015
53        Compute node vHBA           192.168.4.7       Thu Aug 21 01:29:18 2015
---------------
3 rows displayed

Status: Success
```

3.2.17 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
`--confirm`	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
`--force`	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.
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.

Option	Description
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.32 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
```

3.2.18 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
`--confirm`	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
`--force`	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.
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.

Option	Description
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.33 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
```

3.2.19 remove server

Removes a server from an existing tenant group.

Syntax

`remove server 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 servers from, and `node` is the name of the server node that should be removed from the selected tenant group.

Description

Use the `remove server` command to remove the required server 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 server, 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
`--confirm`	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
`--force`	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.
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux

Option	Description
	command line. This option allows both forward and backward navigation through the command output.
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.34 Removing a Compute Node from a Tenant Group

```
PCA> remove server ovcacn09r1 myTenantGroup
************************************************************
 WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
************************************************************
Are you sure [y/N]:y

Status: Success
```

3.2.20 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 3.40, "Show System Properties" for more information.

Options

The following table shows the available options for this command.

Option	Description
compute-node	The command target to perform the reprovision operation against.
--save-local-repo	Skip the HMP step in the provisioning process in order to save the local storage repository.
--json	Return the output of the command in JSON format.
--less	Return the output of the command one screen at a time for easy viewing, as with the less command on the Linux command line. This option allows both forward and backward navigation through the command output.
--more	Return the output of the command one screen at a time for easy viewing, as with the more command on the Linux command line. This option allows forward navigation only.
--force	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.
--tee=OUTPUTFILENAME	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.35 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
Status: Success
```

3.2.21 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 3.31, "List All Configuration Errors" for more information.

You can use the show task command to view the status of the task after you have triggered it to run. See Example 3.41, "Show Task" for more information.

Options

The following table shows the available options for this command.

Option	Description
config-task	The command target to perform the rerun operation against.
--json	Return the output of the command in JSON format
--less	Return the output of the command one screen at a time for easy viewing, as with the less command on the Linux command line. This option allows both forward and backward navigation through the command output.
--more	Return the output of the command one screen at a time for easy viewing, as with the more command on the Linux command line. This option allows forward navigation only.
--tee=OUTPUTFILENAME	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.36 Re-run a configuration task

```
PCA> rerun config-task 84
Status: Success
```

3.2.22 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 3.40, "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
ftp_proxy	Set the value for the IP address of an FTP Proxy
http_proxy	Set the value for the IP address of an HTTP Proxy
https_proxy	Set the value for the IP address of an HTTPS Proxy
log_count	Set the value for the number of log files that should be retained through log rotation
log_file	Set the value for the location of a particular log file. This property can be defined separately for the following log files: backup, cli, diagnosis, monitor, ovca, snmp, and syncservice.
log_level	Set the value for the log level output. Accepted log levels are: CRITICAL, DEBUG, ERROR, INFO, NOTSET, VERBOSE, WARN, and WARNING. This property can be defined separately for the following log files: backup, cli, diagnosis, monitor, ovca, snmp, and syncservice.
log_size	Set the value for the maximum log size before a log is rotated
timezone	Set the time zone for the location of the Oracle PCA.
--json	Return the output of the command in JSON format
--less	Return the output of the command one screen at a time for easy viewing, as with the less command on the Linux command line. This option allows both forward and backward navigation through the command output.
--more	Return the output of the command one screen at a time for easy viewing, as with the more command on the Linux command line. This option allows forward navigation only.
--tee=OUTPUTFILENAME	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.37 Changing the location of the sync service log file

```
PCA> set system-property log_file syncservice /mnt/filer/logs/syncservice.log
Status: Success
```

Note

Log configuration through the CLI is described in more detail in Section 7.2, "Setting the Oracle Private Cloud Appliance Logging Parameters".

Example 3.38 Configuring and unconfiguring an HTTP proxy

```
PCA> set system-property http_proxy http://proxy.example.com:8080
Status: Success
```

```
PCA> set system-property http_proxy ''
Status: Success
```

> **Note**
>
> Proxy configuration through the CLI is described in more detail in Section 7.1,
> "Adding Proxy Settings for Oracle Private Cloud Appliance Updates".

Example 3.39 Configuring the Oracle PCA Time Zone

```
PCA> set system-property timezone US/Eastern
Status: Success
```

3.2.23 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	Storage Network/Cloud Name
compute-node	Compute Node Name
network	Network Name
rack-layout	Rack Architecture or Type
server-profile	Server Name
storage-network	Storage Network/Cloud Name
system-properties	(none)
task	Task ID
tenant-group	Tenant Group Name
version	(none)
vhba-info	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
`cloud-wwpn` | `compute-node` | `network` | `rack-layout` | `server-profile` | `storage-network` | `system-properties` | `task` | `tenant-group` | `version` | `vhba-info`	The command target to show information for.
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.40 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                US/Eastern
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       DEBUG
```

```
Snmp.Log_File          /nfs/shared_storage/logs/ovca_snmptrapd.log
Snmp.Log_Level         DEBUG
-------------------------------------

Status: Success
```

Example 3.41 Show Task

```
PCA> show task 341e7bc74f339c

-------------------------------------
Task_Name              backup
Status                 RUNNING
Progress               70
Start_Time             06-27-2016 09:59:36
End_Time               None
Pid                    1503341
Result                 None
-------------------------------------

Status: Success
```

Example 3.42 Show Rack Layout

```
PCA> show rack-layout x5-2_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]
25  ovcasw22r1    infrastructure  switch      ovn-switch2   [25, 24, 23, 22]
24  ovcasw22r1    infrastructure  switch      ovn-switch2   [25, 24, 23, 22]
23  ovcasw22r1    infrastructure  switch      ovn-switch2   [25, 24, 23, 22]
22  ovcasw22r1    infrastructure  switch      ovn-switch2   [25, 24, 23, 22]
21  ovcasw21br1   infrastructure  switch      opus-switch2  [21]
21  ovcasw21ar1   infrastructure  switch      opus-switch1  [21]
20  ovcasw20r1    infrastructure  switch      ib-switch2    [20]
19  ovcasw19r1    infrastructure  switch      ib-switch1    [19]
18  ovcasw15r1    infrastructure  switch      ovn-switch1   [18, 17, 16, 15]
17  ovcasw15r1    infrastructure  switch      ovn-switch1   [18, 17, 16, 15]
16  ovcasw15r1    infrastructure  switch      ovn-switch1   [18, 17, 16, 15]
15  ovcasw15r1    infrastructure  switch      ovn-switch1   [18, 17, 16, 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   N / A          infrastructure   storage      disk-shelf   [4, 3]
3   N / A          infrastructure   storage      disk-shelf   [4, 3]
2   ovcasn02r1     infrastructure   storage      zfs-head2    [2]
1   ovcasn01r1     infrastructure   storage      zfs-head1    [1]
0   ovcapduBr1     infrastructure   pdu          pdu2         [0]
0   ovcapduAr1     infrastructure   pdu          pdu1         [0]
----------------
45 rows displayed

Status: Success
```

Example 3.43 Show the Ports and VNICs in the mgmt_public_eth Network

```
PCA> show network mgmt_public_eth

----------------------------------------
Cloud_Name         mgmt_public_eth
Trunkmode          True
Description        2015-07-12 02:09:42
Ports              ovcasw22r1:5:1, ovcasw22r1:4:1, ovcasw15r1:5:1, ovcasw15r1:4:1
Vnics              ovcacn27r1-eth5, ovcacn13r1-eth5, ovcacn14r1-eth5, ovcacn09r1-eth5,
                   ovcacn36r1-eth5, ovcamn06r1-eth5, ovcacn29r1-eth5, ovcacn07r1-eth5,
                   ovcacn12r1-eth5, ovcacn30r1-eth5, ovcacn42r1-eth5, ovcacn40r1-eth5,
                   ovcacn10r1-eth5, ovcacn28r1-eth5, ovcacn41r1-eth5, ovcamn05r1-eth5,
                   ovcacn38r1-eth5, ovcacn34r1-eth5, ovcacn37r1-eth5, ovcacn11r1-eth5,
                   ovcacn26r1-eth5, ovcacn39r1-eth5, ovcacn35r1-eth5
Status             None
Network_Type       external_network
Compute_Nodes      ovcacn27r1, ovcacn13r1, ovcacn14r1, ovcacn09r1, ovcacn36r1,
                   ovcacn06r1, ovcacn29r1, ovcacn07r1, ovcacn12r1, ovcacn30r1,
                   ovcacn42r1, ovcacn40r1, ovcacn10r1, ovcacn28r1, ovcacn05r1,
                   ovcacn34r1, ovcacn11r1, ovcacn26r1, ovcacn35r1
Prefix             192.168.100
Netmask            None
Route Destination  None
Route Gateway      None
----------------------------------------

Status: Success
```

Example 3.44 Show Details of a Storage Network

```
PCA> show storage-network Cloud_A

----------------------------------------
Cloud_Name    Cloud_A
Description   Default Storage Cloud ru22 port1 - Do not delete or modify
Ports         ovcasw22r1:12:1, ovcasw22r1:3:1
vHBAs         ovcacn11r1-vhba01, ovcacn29r1-vhba01, ovcacn12r1-vhba01,
              ovcacn09r1-vhba01, ovcacn14r1-vhba01, ovcacn28r1-vhba01,
              ovcacn08r1-vhba01, ovcacn30r1-vhba01, ovcacn26r1-vhba01,
              ovcacn13r1-vhba01, ovcacn27r1-vhba01, ovcacn31r1-vhba01,
              ovcacn10r1-vhba01, ovcamn06r1-vhba01, ovcamn05r1-vhba01
----------------------------------------

Status: Success
```

Example 3.45 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    192.168.140.131
Networks            ['myPublicNetwork']
Pool Filesystem ID  3600144f0d29d4c86000057162ecc0001
--------------------------------------

Status: Success
```

Example 3.46 Show Details of a Custom Network

```
PCA> show network myHostNetwork

--------------------------------------
Network_Name        myHostNetwork
Trunkmode           True
Description         User defined network
Ports               ['11:4']
vNICs               ovcacn02r2-eth8, ovcacn01r2-eth8, ovcacn42r1-eth8
Status              ready
Network_Type        host_network
Compute_Nodes       ovcacn42r1, ovcacn01r2, ovcacn02r2
Prefix              10.10.10
Netmask             255.255.255.0
Route Destination   10.10.20.0/24
Route Gateway       10.10.10.250
--------------------------------------

Status: Success
```

Example 3.47 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 3.48 Show the vHBA configuration for a Compute Node

```
PCA> show vhba-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 3.49 Show Oracle PCA Version Information

```
PCA> show version
```

```
----------------------------------------
Version            2.2.1
Build              384
Date               2016-04-19
----------------------------------------

Status: Success
```

3.2.24 start

Starts up a rack component.

Note

Feature disabled in this release.

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
`compute-node CN` | `management-node MN`	Start either a compute node or a management node. Replace *CN* or *MN* respectively with the host name of the server to be started.
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.50 Starting a Compute Node

```
PCA> start compute-node ovcacn11r1
Status: Success
```

3.2.25 stop

Shuts down a rack component or aborts a running task.

> **Note**
>
> Stopping a rack component is disabled in this release.

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
`compute-node CN` | `management-node MN`	Shut down either a compute node or a management node. Replace *CN* or *MN* respectively with the host name of the server to be shut down.
`task id` | `update-task id`	Aborts a running task. Use the `update-task` 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. > **Caution** >ocupa > Stopping an update task is a risky operation and should be used with extreme caution.
`--confirm`	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.

Option	Description
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.51 Aborting a Task

```
PCA> stop task 341d45b5424c16
*************************************************************
 WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*************************************************************
Are you sure [y/N]:y

Status: Success
```

3.2.26 update appliance

Updates the Oracle PCA software stack by downloading and installing a new ISO image.

Syntax

`update appliance { get_image url | install_image } [ --confirm ] [ --json ] [ --less ] [ --more ] [ --tee=OUTPUTFILENAME ]`

where `url` is the download location for the zipped ISO file containing the software update.

Description

Use the `update appliance` command to first download a new version of the Oracle PCA software stack, and then launch the software update process. 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
`get_image url`	Start the task to download the zipped ISO file containing the software update from the specified url.
`install_image`	Start the task to update the appliance software stack with the previously downloaded image file.
`--confirm`	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.

Option	Description
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.52 Downloading an Oracle PCA Software Update

```
PCA> update appliance get_image http://myserver.org/images/ovca-2.1.1-b999.iso.zip
************************************************************
 WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
************************************************************
Are you sure [y/N]:y

The update job has been submitted. Use "show task <task id>" to monitor the progress.

Task_ID        Status  Progress Start_Time           Task_Name
-------        ------  -------- ----------           ---------
333dcc8b617f74  RUNNING None     09-15-2015 11:40:53  update_download_image
--------------
1 row displayed

Status: Success
```

3.2.27 update password

Modifies the password for one or more components within the Oracle Private Cloud Appliance.

Syntax

update password { mgmt-root | mysql-ovs | nm-ilom-admin | nm-ilom-operator | nm-root | opus-root | ovm-admin | spCn-root | spMn-root | spZfs-root | system-root | wls-weblogic | xsigo-admin | xsigo-recovery | xsigo-root | zfs-root } [PCA-password target-password] [--confirm] [--json] [--less] [--more] [--tee=OUTPUTFILENAME]

where `PCA-password` is the current password of the Oracle PCA 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 PCA password and the new target component password with the command. If not, you are prompted for the current password of the Oracle PCA 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
`mgmt-root`	Sets a new password for the `root` user on the management nodes.
`mysql-ovs`	Sets a new password for the `ovs` user in the MySQL database.

Option	Description
`nm-ilom-admin`	Sets a new password for the `ilom-admin` user in the InfiniBand switches' ILOMs.
`nm-ilom-operator`	Sets a new password for the `ilom-operator` user in the InfiniBand switches' ILOMs.
`nm-root`	Sets a new password for the `root` user in the InfiniBand switches' ILOMs.
`opus-root`	Sets a new password for the `root` user on the Ethernet switches.
`ovm-admin`	Sets a new password for the `admin` user in Oracle VM Manager.
`spCn-root`	Sets a new password for the `root` user in the compute node ILOMs.
`spMn-root`	Sets a new password for the `root` user in the management node ILOMs.
`spZfs-root`	Sets a new password for the `root` user on the ZFS storage appliance as well as its ILOM.
`system-root`	Sets a new password for the `root` user on all compute nodes.
`wls-weblogic`	Sets a new password for the `weblogic` user in WebLogic Server.
`xsigo-admin`	Sets a new password for the `admin` user on the Fabric Interconnects.
`xsigo-recovery`	Sets a new password for the `rcli` user on the Fabric Interconnects.
`xsigo-root`	Sets a new password for the `root` user on the Fabric Interconnects.
`zfs-root`	Sets a new password for the `root` user on the ZFS storage appliance as well as its ILOM.
`--confirm`	Confirm flag for destructive command. Use this flag to disable the confirmation prompt when you run this command.
`--json`	Return the output of the command in JSON format
`--less`	Return the output of the command one screen at a time for easy viewing, as with the `less` command on the Linux command line. This option allows both forward and backward navigation through the command output.
`--more`	Return the output of the command one screen at a time for easy viewing, as with the `more` command on the Linux command line. This option allows forward navigation only.
`--tee=OUTPUTFILENAME`	When returning the output of the command, also write it to the specified output file.

Examples

Example 3.53 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
```

```
 WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
Are you sure [y/N]:y
Current PCA Password:
```

Chapter 4 Managing the Oracle VM Virtual Infrastructure

Table of Contents

Warning

Access to the Oracle VM Manager web user interface and command line interface is provided without restrictions. The configuration of Oracle Private Cloud Appliance (PCA) components within Oracle VM Manager is automatic and handled by the Oracle PCA 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.

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 PCA.

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, use of this on your Oracle PCA should only be attempted by advanced users with a thorough understanding of Oracle VM and the limitations of its usage within an Oracle PCA context.

The information provided in here, is a description of the Oracle VM Manager Web UI within the context of the Oracle PCA. Where particular actions within the Oracle VM Manager Web UI are referenced, a link to the appropriate section within the Oracle VM User's Guide is provided. The complete Oracle VM User's Guide is available at this URL: http://docs.oracle.com/cd/E35328_01/E35332/html/index.html.

Note

When consulting the Oracle VM documentation directly, keep in mind the limitations imposed by using it within Oracle PCA.

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 active node within the cluster. If the 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.

4.1 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, and those that may be attempted with due care, are presented as guidelines and limitations that should be followed when working within Oracle VM Manager.

The following actions must not be performed, except if Oracle gives specific instructions to do so.

Do Not:

- attempt to discover or modify or remove servers or their configuration;

- attempt to add or modify or remove server pools or their configuration;

- attempt to move servers out of the existing server pool;

- attempt to add or modify or remove server processor compatibility groups;

- attempt to modify or remove the existing repository named **Rack1-repository**, or the local disk repositories;

- attempt to delete or modify any of the preconfigured networks;

- attempt to connect virtual machines to the appliance management network;

- attempt to add or configure VLAN Groups, except to add VLAN tags to the default VLAN Group configuration;

- attempt to modify or delete any existing Storage elements that are already configured within Oracle VM;

- attempt to configure global settings in the **Tools and Resources** tab, such as the NTP or YUM Update configuration.

While unlikely to cause a problem with your appliance, the following actions are not recommended:

Not Recommended:

- add or configure additional networks within Oracle VM Manager;

- use the Oracle VM Manager Command Line Interface or Utilities.

If you ignore this advice, the Oracle PCA 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 PCA.

4.2 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.

4.3 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 PCA Dashboard. See Section 2.3, "Hardware View" for more information.

The **Statistics** subtab available on the Health tab provides statistical information, including real-time graphs, for CPU and memory usage. These statistics can be viewed at a global scale to determine overall usage, 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.

For detailed information on using the Health tab, please refer to the section entitled Health Tab in the Oracle VM User's Guide.

4.4 Creating and Managing Virtual Machines

The **Servers and VMs** tab is used to create and manage your virtual machines. By default, compute nodes are listed as belonging to a single server pool called **Rack1_ServerPool**. The configuration of this server pool must not be altered. There is no need to discover servers, as compute nodes are automatically provisioned and discovered within an Oracle PCA. Editing the configuration of the server pool, servers and processor compatibility groups is not supported. The primary purpose of this tab within the Oracle PCA 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

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 4.5, "Managing Virtual Machine Resources" for more information on importing these resources into the Oracle PCA repository.

Before you create a virtual machine that requires network connectivity, or a paravirtualized machine which requires network connectivity to perform the operating system install, you must generate virtual

network interfaces using the VNIC Manager. See Section 4.6, "Configuring Network Resources for Virtual Machines" for information on using the VNIC Manager.

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 User's Guide:

Managing Virtual Machines

- Create a virtual machine

  You can create a virtual machine following the instructions provided in the section entitled Creating a 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 PCA 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 Viewing Virtual Machine Information and Events.

- Edit a virtual machine

  You can edit virtual machine parameters as described in the section entitled Editing a Virtual Machine.

  You can also convert a virtual machine to use paravirtualization. This involves making changes to the virtual machine itself, and then editing its Oracle VM Domain Type. Refer to the section entitled Converting to Paravirtualized Guests or Installing Paravirtualized Drivers for information on converting a virtual machine to use paravirtualization.

- Start a virtual machine

  Further information is provided in the section entitled Starting Virtual Machines.

- Stop a virtual machine

  Further information is provided in the section entitled Stopping (Shutting Down) Virtual Machines.

- Kill a virtual machine

  Further information is provided in the section entitled Killing Virtual Machines.

- Restart a virtual machine

  Further information is provided in the section entitled Restarting Virtual Machines.

- Suspend a virtual machine

  Further information is provided in the section entitled Suspending Virtual Machines.

- Resume a virtual machine

  Further information is provided in the section entitled Resuming a Virtual Machine.

- Move a virtual machine between repositories

  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. Refer to the section entitled Moving Virtual Machines Between Repositories for more information.

- Move a virtual machine from one server to another

  Further information is provided in the section entitled Moving Virtual Machines Between Oracle VM Servers.

- Move a virtual machine to or from the Unassigned Virtual Machines folder

  Further information is provided in the section entitled Moving Virtual Machines To/From Unassigned Virtual Machines Folder.

- Migrate a virtual machine

  Because there is only a single server pool available in an Oracle PCA base rack, migration of virtual machines can only be achieved between servers and between a server and the Unassigned Virtual Machines folder. 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 Groups.

  Information on migrating virtual machines is provided in the section entitled Migrating Virtual Machines.

- 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 Anti-Affinity Groups.

- Delete a virtual machine

  Further information is provided in the section entitled Deleting Virtual Machines.

- 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 Sending Messages to Virtual Machines for more information.

- Connect to a virtual machine console

  Further information is provided in the section entitled Connecting to a Virtual Machine.

Monitoring Compute Node Information and Events

- View Oracle VM Server information and events

  For more information on the options available to you, refer to the section entitled Viewing Oracle VM Server Information and Events.

Figure 4.1 A view of the Servers and VMs tab
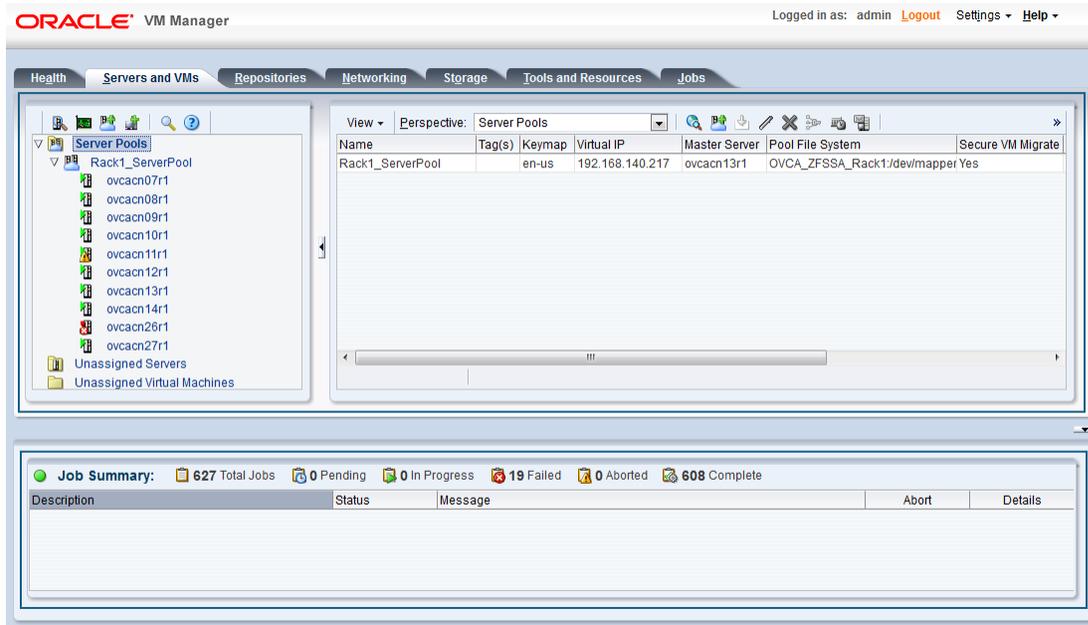
4.5 Managing Virtual Machine Resources

The **Repositories** tab provides a view of the Oracle PCA repository. By default, a shared repository is configured on the 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.

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, assemblies 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 User's Guide:

- Manage Virtual Machine Templates

  - Import a template

  - Create a template

  - Edit a template

  - Delete a template

  - Clone a template

  - Move a template

  - Manage a template clone customizer

  All documentation for these actions can be found in the section entitled Virtual Machine Templates.

- Manage Virtual Machine Assemblies

  - Import an assembly

  - Create a template from an assembly

  - Edit an assembly

  - Delete an assembly

  - Refresh an assembly

  All documentation for these actions can be found in the section entitled Assemblies.

- Manage Virtual Machine ISO Image Files

  - Import an ISO

  - Edit an ISO

  - Delete an ISO

  - Clone an ISO

  All documentation for these actions can be found in the section entitled ISO Files (CD/DVD Images).

- Manage Virtual Disks

  - Create a virtual disk

  - Import a virtual disk

  - Edit a virtual disk

  - Delete a virtual disk

- Clone a virtual disk

  All documentation for these actions can be found in the section entitled Virtual Disks.

- View Virtual Machine Configuration Entries

  For more information, refer to the section entitled Virtual Machine Configuration Files.

While it is possible to create additional repositories using alternate storage, this is beyond the scope of a normal Oracle PCA configuration and is not recommended.

4.6 Configuring Network Resources for Virtual Machines

The **Networking** tab is used to manage networks within the Oracle VM environment running on the Oracle PCA.

> **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 PCA software layer.

The 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 VLAN 1 from the VLAN Group named `vm_public_vlan_vgrp`, meaning no additional configuration is required.

  If you prefer to use VLANs for your VM networking, configure the VLAN tags of your choice as follows:

  1. Add the tags or IDs to the VLAN Group `vm_public_vlan_vgrp`.

     The process for editing VLAN Groups is described in the Oracle VM User's Guide in the section entitled Editing a VLAN Group.

  2. Create a new VLAN-only network with the *VM role* for each VLAN tag you added to the VLAN Group. Each new network should be configured like the `vm_public_vlan` network, but with a different VLAN segment.

     The process for creating VLAN-only networks is described in the Oracle VM User's Guide in the section entitled Creating a Network.

3. 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. It uses VLAN 1 from the VLAN Group named `vm_private_vgrp`. To use the VLANs of your choice configure them as follows:

  1. Add the tags or IDs to the VLAN Group `vm_private_vgrp`.

      The process for editing VLAN Groups is described in the Oracle VM User's Guide in the section entitled Editing a VLAN Group.

  2. Create a new VLAN-only network with the *VM role* for each VLAN tag you added to the VLAN Group. Each new network should be configured like the `vm_private` network, but with a different VLAN segment.

      The process for creating VLAN-only networks is described in the Oracle VM User's Guide in the section entitled Creating a 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 PCA 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 PCA network configuration, see Section 1.2.4, "Network Infrastructure".

> **Caution**
>
> Do not create a new network in Oracle VM using the `eth0` network ports of the servers in the server pool. Those ports are assigned an IP address by the DHCP server on the internal appliance management network (`192.168.4.0`). If virtual machines are connected to this network, they are likely to cause IP conflicts and security issues.

Do not edit or delete any of the networks listed here. Doing so may cause your appliance to malfunction. In an Oracle PCA context, use the Networking tab to configure and manage Virtual NICs and VLANs for use by your virtual machines.

The creation and management of Virtual NICs is important for the facilitation of networking within your virtual machines. Use the **Virtual NICs** subtab to create batches of Virtual NICs that virtual machines are

able to use. The process for creating Virtual NICs is described in the Oracle VM User's Guide in the section entitled Managing VNICs.

Figure 4.2 A view of the Networking tab

4.7 Viewing and Managing Storage Resources

The storage resources underlying the built-in Oracle PCA 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 exported through NFS and is listed under the File 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.

It is possible to attach additional storage for use within Oracle VM, if you intend to create an additional repository to store virtual machine resources. The options to extend the storage capacity of an Oracle PCA are explained in detail in the *Oracle Private Cloud Appliance Installation Guide*: refer to the chapter entitled Extending Oracle Private Cloud Appliance - External 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 User's Guide*. Refer to the section entitled Managing Storage.

You are also fully capable of using other networked storage, available on the public network, within your own Virtual Machines. The limitations described here apply only to internal ZFS storage used by the Oracle VM environment running within your Oracle PCA.

Figure 4.3 A view of the Storage tab
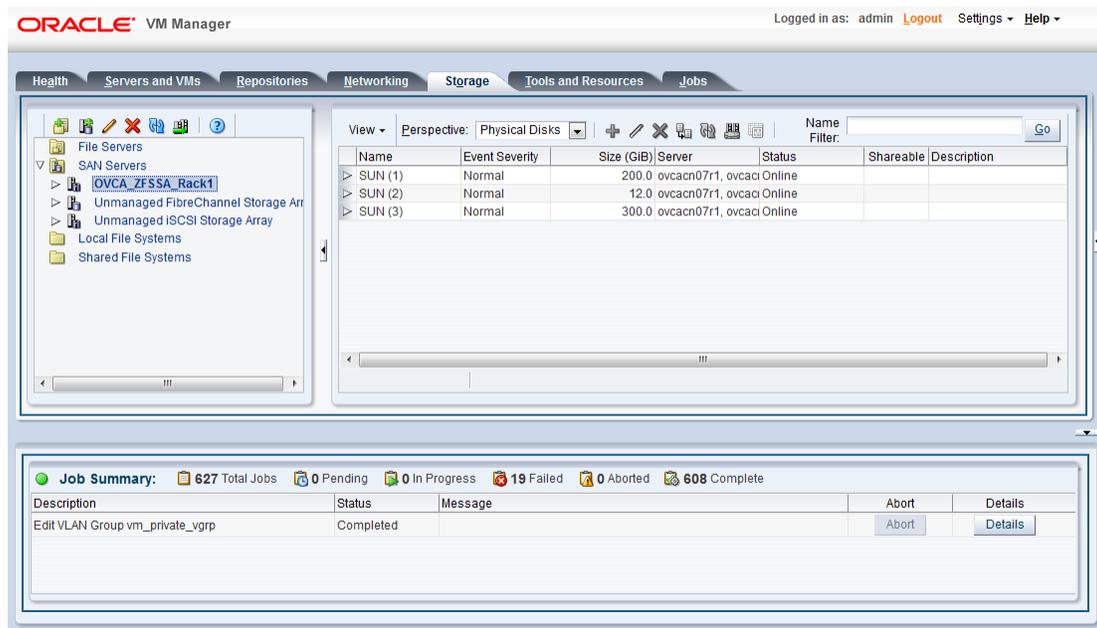
4.8 Tagging of Resources in Oracle VM Manager

The **Tools 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 PCA, 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.

4.9 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 User's Guide in the section entitled Working with the Jobs Framework.

Since the Recurring Jobs, described in the Oracle VM User's Guide, are all automated and handled directly by the Oracle PCA, you must not edit any of the settings for recurring jobs.

Chapter 5 Automating Oracle Private Cloud Appliance Service Requests

Table of Contents

Oracle Private Cloud Appliance (PCA) is qualified for Oracle Auto Service Request (ASR). ASR is 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.

The information in this chapter is based on the support note with Doc ID 2032791.1. For the latest information about installing Oracle Auto Service Request (ASR) on Oracle PCA, refer to My Oracle Support or contact your Oracle representative.

> ⚠ **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: http://www.oracle.com/technetwork/systems/asr/documentation/index.html.

5.1 Understanding Oracle Auto Service Request (ASR)

ASR is designed to automatically open service requests when specific Oracle PCA hardware faults occur. To enable this feature, the Oracle PCA 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: http://www.oracle.com/technetwork/systems/asr/overview/index.html.

- Oracle Auto Service Request user documentation: http://docs.oracle.com/cd/E37710_01/index.htm.

5.2 ASR Prerequisites

Before you install ASR, make sure that the prerequisites in this section are met.

Verifying ASR Prerequisites

1. Make sure that you have a valid My Oracle Support account.

    If necessary, create an account at https://support.oracle.com.

2. Ensure that the following are set up correctly in My Oracle Support:

    - technical contact person at the customer site who is responsible for Oracle PCA

    - valid shipping address at the customer site where the Oracle PCA is located, so that parts are delivered to the site where they must be installed

3. 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 PCA. 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/.

4. Verify connectivity to the Internet using HTTPS.

    For example, try `curl` to test whether you can access https://support.oracle.com.

5.3 Installing ASR Manager

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 PCA, 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.

Downloading the Current ASR Packages

1. Verify whether ASR Manager is already installed. If any version older than 5.0 is installed, uninstall it first.

   For detailed information, refer to the document *ASR Manager 5.x Installation, Backup and Upgrade* (Doc ID 1392042.1).

2. Download ASR Manager Release 5.0 or newer from the ASR download page.

   Select the download for the Linux x86-64 platform.

   > **Note**
   >
   > The version tested at the time of writing is 5.2.1. The latest version is 5.3.

3. Unpack the `*.zip` file in a directory available to both management nodes. For example: `/nfs/shared_storage/ASRM/pkg`.

```
[root@ovcamn05r1 ~]# cd /nfs/shared_storage/ASRM/pkg
[root@ovcamn05r1 pkg]# unzip p21056276_521_Linux-x86-64.zip
Archive:  p21056276_521_Linux-x86-64.zip
  inflating: asrmanager-5.0.2-20141215170108.rpm
  inflating: readme.html
  inflating: license_agreement.html
```

Installing ASR Components on the Management Nodes

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 ~]# pca-check-master
NODE: 10.100.1.101  MASTER: True
```

2. Go to the shared directory where you stored the ASR package.

```
# cd /nfs/shared_storage/ASRM/pkg
```

3. Install the ASR Manager package.

```
# rpm -i asrmanager-5.2.1-20150501144949.rpm
Copyright [2008,2014], Oracle and/or its affiliates. All rights reserved.
License and Terms of Use for this software are described at https://support.oracle.com/
(see Legal Notices and Terms of Use).

ASR Manager (pid 345193) is RUNNING.
Installing ASR Manager bundles...
Successfully installed ASR Manager bundles.
asrm: unrecognized service
Successfully added ASR Manager (asrm) service.

******************************************************
To allow a non-root user to manage ASR Manager service:
```

```
Add the following line to /etc/sudoers file
'<userName> ALL=(root) NOPASSWD:/opt/asrmanager/bin/asr start,/opt/asrmanager/bin/asr stop,
/opt/asrmanager/bin/asr status,/opt/asrmanager/bin/asr restart'
******************************************************

The ASR Manager application is installed in '/opt/asrmanager'. Log files are located in '/var/opt/asrmanage

ASR Administration command is now available at /opt/asrmanager/bin/asr.
ASR Manager is stopped.
ASR Manager (pid 347799) is RUNNING.
Checking ASR Manager status ....Installation of asrmanager was successful.
```

> **Warning**
>
> For Oracle PCA the auto-update feature of ASR must not be used.

At the end of the installation, the ASR Manager service (`asrm`) is started automatically.

4. Disable auto-update.

```
# /opt/asrmanager/bin/asr disable_autoupdate
```

5. Stop the ASR Manager service and prevent it from starting automatically.

```
# service asrm status
ASR Manager (pid 357177) is RUNNING.
# service asrm stop
ASR Manager is stopped.
# chkconfig asrm off
```

6. Configure the `init` service to log when the ASR Manager service starts and stops.

   a. Open the file `/etc/init.d/asrm` for editing.

   b. In the start and stop sections, add the echo commands as shown.

   ```
   start)
     /opt/asrmanager/bin/asrm start
     echo "START: `hostname`  `date "+%D %T"`" >> /nfs/shared_storage/ASRM/ASRHA.log
     ;;

   stop)
     /opt/asrmanager/bin/asrm stop
     echo "STOP:  `hostname`  `date "+%D %T"`" >> /nfs/shared_storage/ASRM/ASRHA.log
     ;;
   ```

   c. Save and close the file `/etc/init.d/asrm`.

7. Preserve the data in the original ASR directories.

```
# mv /var/opt/asrmanager /var/opt/orig_asrmanager
```

8. Repeat this procedure on the other management node.

```
# ssh root@10.100.1.102
root@10.100.1.102's password:
[root@ovcamn06r1 ~]# pca-check-master
NODE: 10.100.1.102  MASTER: False
```

> **Caution**
>
> The secondary management node does not need to take over the master role
> during the ASR installation.

The ASR Manager service must remain stopped on both management nodes.

Configuring ASR Manager Failover

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. Create the following directories and log file on the shared storage:

   ```
   # mkdir /nfs/shared_storage/ASRM
   # mkdir /nfs/shared_storage/ASRM/VarOptAsrmanager
   # touch /nfs/shared_storage/ASRM/ASRHA.log
   ```

3. Copy the original ASR data from the master management node local file system to the shared storage.

   ```
   # cp -r /var/opt/orig_asrmanager/* /nfs/shared_storage/ASRM/VarOptAsrmanager
   ```

4. Create a symlink to the directory on the shared storage.

   ```
   # ln -s /nfs/shared_storage/ASRM/VarOptAsrmanager /var/opt/asrmanager
   ```

5. Add the `asrm` service to the Oracle PCA system configuration.

   a. Open the file `/var/lib/ovca/ovca-system.conf` for editing.

   b. Insert a new line to add `asrm` to the foundational services.

   ```
   [mgmt_init]
   [...]
   foundational_services: ovmm_mysql
                          ovmm
                          tinyproxy
                          xms
                          dhcpd
                          asrm
   [...]
   ```

   c. Save and close the file `/var/lib/ovca/ovca-system.conf`.

6. Log into the other management node and execute steps 4 and 5.

   a. Create a symlink to the directory on the shared storage.

   b. Add the `asrm` service to the Oracle PCA system configuration.

Starting ASR Manager

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. Start the ASR Manager service.

> **Note**
>
> You can monitor the process by tailing the log file `/nfs/shared_storage/ASRM/ASRHA.log`.

```
# service asrm start
```

3. Register the ASR Manager.

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 PCA ASRM as stand-alone. This means it communicates directly with the Oracle backend systems, which is the standard registration method.

For details and instructions, refer to the section entitled "Register the ASR Manager" in the *Oracle Auto Service Request ASR Manager User's Guide*.

The basic registration command is: `/opt/asrmanager/bin/asr register`.

A relay ASRM could be used if an established ASRM is already in use by other systems in their facility. Since the other systems **cannot** be routed through the Oracle PCA ASR Manager, you must route the Oracle PCA ASR Manager through the existing one.

The instructions for setting up the ASR Manager in relay mode are in the ASR documentation. In short you must register the Oracle PCA ASR Manager to the web address generated by the relay ASR Manager, which must already be configured for relaying the data. (Follow the instructions in the documentation for your particular ASR version, as it may be different with each version of ASR.) To get the destination address, enter the following command at the ASR prompt on the destination ASR Manager:

```
# /opt/asrmanager/bin/asr show_http_receiver
        HTTP Receiver configuration:
        HTTP Receiver Status: Enabled
        Host Name: Relay_Node.Company.com
        HTTP Port: 7777
        HTTPS/SSL configuration is not enabled.
```

The basic ASR Manager relay command is: `/opt/asrmanager/bin/asr register -e http://Relay_Node.Company.com:7777/asr`.

4. Test ASR Manager failover by rebooting the current master management node.

> **Note**
>
> You can monitor the process by tailing the log file `/nfs/shared_storage/ASRM/ASRHA.log`.

```
[root@ovcamn05r1 ]# reboot
[root@ovcamn06r1 ]# tail -f /nfs/shared_storage/ASRM/ASRHA.log
STOP:  ovcamn05r1  07/20/15 14:35:20
START: ovcamn05r1  07/20/15 14:35:20
STOP:  ovcamn05r1  07/20/15 14:37:37
START: ovcamn06r1  07/20/15 14:39:15
```

5.4 Configuring TinyProxy

For Oracle ASR to work, both Oracle PCA management nodes must be running TinyProxy version 1.8.3 or later. Verify the installed version and upgrade if necessary. Configure TinyProxy according to the procedure in this section.

> **Warning**
>
> Strictly apply the configuration changes presented in this section. Do not modify TinyProxy in any other way. Certain modifications to TinyProxy can cause the provisioning functionality to break.

Configuring TinyProxy for ASR

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 ~]# pca-check-master
   NODE: 10.100.1.101  MASTER: True
   ```

2. Stop the `tinyproxy` service and verify the version installed on the system.

   ```
   # service tinyproxy stop
   Stopping tinyproxy: [ OK ]
   # rpm -q tinyproxy
   tinyproxy-1.8.2-1.el6.x86_64
   ```

3. If the installed TinyProxy is older than version 1.8.3, upgrade it to the version referenced in the support note with Doc ID 2032791.1.

   Download the tinyproxy rpm to a temporary directory on the management node and run the upgrade command.

   ```
   # cd /tmp
   # ls -al tinyproxy*
   -rw-r--r-- 1 root root 61224 Apr 4 13:44 tinyproxy-1.8.3-1.el6.x86_64.rpm
   # rpm -Uvh tinyproxy-1.8.3-1.el6.x86_64.rpm
   warning: tinyproxy-1.8.3-1.el6.x86_64.rpm: Header V3 RSA/SHA256 Signature,
   key ID ec551f03: NOKEY
   Preparing...                ########################################### [100%]
      1:tinyproxy              ########################################### [100%]
   ```

4. Make the necessary changes to the TinyProxy configuration to enable ASR traffic through the management node.

   a. Open the file `/etc/tinyproxy/tinyproxy.conf` for editing.

   b. Navigate to the section that defines on which interfaces tinyproxy listens. It starts with "`# Listen:`".

      Comment out the line that binds tinyproxy to the interface in the Oracle VM management network by adding a hash character (#) in front of it.

      ```
      # Listen: If you have multiple interfaces this allows you to bind to
      # only one. If this is commented out, tinyproxy will bind to all
      ```

```
# interfaces present.
#
#Listen 192.168.140.4
```

c. Scroll down to the section Customization of authorization controls. It starts with "`# Allow:`".

Add two lines to allow traffic from the internal appliance management network and Oracle VM management network.

```
# Allow: Customization of authorization controls. If there are any
# access control keywords then the default action is to DENY. Otherwise,
# the default action is ALLOW.
#
# The order of the controls are important. All incoming connections are
# tested against the controls based on order.
#
#Allow 127.0.0.1
Allow 192.168.4.0/24
Allow 192.168.140.0/24
```

d. Save and close the file `/etc/tinyproxy/tinyproxy.conf`.

5. Restart the `tinyproxy` service and verify the installed version.

```
# service tinyproxy start
Starting tinyproxy: [ OK ]
# tinyproxy -version
tinyproxy 1.8.3
```

6. Repeat this procedure on the other management node.

5.5 Activating ASR Assets

The Oracle Private Cloud Appliance components that are qualified as ASR assets are:

- Server compute nodes

- NM2-36P Sun Datacenter InfiniBand Expansion Switches

- ZFS Storage Appliance

> **Note**
>
> There is no ASR support for the Oracle Switch ES1-24 switches and Oracle Fabric Interconnect F1-15s in Oracle Private Cloud Appliance Release 2.2. The two management nodes also must not be activated.

This section provides the instructions to activate each type of Oracle Private Cloud Appliance ASR asset.

5.5.1 Activating Compute Nodes

The ASR activation mechanism 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. Follow the procedure below to activate the compute nodes as ASR assets.

Activating Compute Nodes as ASR Assets

1. Using SSH, log into the compute node ILOM as `root`.

2. Go to the alert rules directory and display rule set 2.

> **Caution**
>
> As of Release 2.2.1 the ILOM rule set 1 is used by the Oracle PCA controller software. The first available rule set is typically set 2.

```
-> cd /SP/alertmgmt/rules
/SP/alertmgmt/rules

-> show 2
  /SP/alertmgmt/rules/2
      Targets:
      Properties:
          type = snmptrap
          level = disable
          destination = 0.0.0.0
          destination_port = 0
          community_or_username = public
          snmp_version = 1
          testrule = (Cannot show property)
->
```

If the rule set is available (destination = 0.0.0.0), configure it for ASR. Otherwise, look for the next available rule set.

3. Define a new rule set with the following properties:

```
-> cd 2
-> set level=minor destination=192.168.4.216 destination_port=162 snmp_version=2c
-> show
  /SP/alertmgmt/rules/2
      Targets:
      Properties:
          type = snmptrap
          level = minor
          destination = 192.168.4.216
          destination_port = 162
          community_or_username = public
          snmp_version = 2c
          testrule = (Cannot show property)
      Commands:
          cd
          set
          show
->
```

The *destination* property is a virtual IP in the appliance management network, shared by both management nodes, but active only on the master.

Use a space character to separate rule set properties.

4. Activate the compute node as an ASR asset in ASR Manager.

   a. Using SSH, log into the master management node.

   b. Activate the asset by adding its ILOM IP to the ASR configuration with the following command:

   ```
   # /opt/SUNWswasr/bin/asr activate_asset –i <asset_ilom_ip>
   ```

   > **Caution**
   >
   > Verify that you are activating the asset with the IP address of its ILOM, and not the IP assigned to another network interface.

5. Repeat this procedure for each compute node you wish to activate as an ASR asset.

5.5.2 Activating InfiniBand Switches

The ASR activation mechanism requires operations in two separate locations. First the switches' 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. Follow the procedure below to activate the NM2-36P Sun Datacenter InfiniBand Expansion Switches as ASR assets.

Activating InfiniBand Switches as ASR Assets

1. Using SSH, log into the NM2-36P Sun Datacenter InfiniBand Expansion Switch ILOM as `root`.

2. Launch the Service Processor shell.

```
[root@ilom-ovcasw19r1 ~]# spsh
->
```

3. Go to the alert rules directory and display rule set 2.

⚠️ **Caution**

As of Release 2.2.1 the ILOM rule set 1 is used by the Oracle PCA controller software. The first available rule set is typically set 2.

```
-> cd /SP/alertmgmt/rules
/SP/alertmgmt/rules

-> show 2
  /SP/alertmgmt/rules/2
      Targets:
      Properties:
          type = snmptrap
          level = disable
          destination = 0.0.0.0
          destination_port = 0
          community_or_username = public
          snmp_version = 1
          testrule = (Cannot show property)
      Commands:
          cd
          set
          show
->
```

4. Define a new rule set with the following properties:

```
-> cd 2
-> set level=minor destination=192.168.4.216 destination_port=162 snmp_version=2c
Set 'level' to 'minor'
Set 'destination' to '192.168.4.216'
Set 'destination_port' to '162'
Set 'snmp_version' to '2c'

-> show
  /SP/alertmgmt/rules/2
      Targets:
      Properties:
          type = snmptrap
          level = minor
          destination = 192.168.4.216
          destination_port = 162
          community_or_username = public
```

```
            snmp_version = 2c
            testrule = (Cannot show property)
        Commands:
            cd
            set
            show
->
```

The *destination* property is a virtual IP in the appliance management network, shared by both management nodes, but active only on the master.

Rule set properties are space-separated.

5. Make sure that SNMP version 2c is enabled on the switch.

```
-> cd /SP/services/snmp
/SP/services/snmp
-> show
  /SP/services/snmp
      Targets:
          communities
          mibs
          users
      Properties:
          engineid = (none)
          port = 161
          servicestate = enabled
          sets = disabled
          v1 = disabled
          v2c = disabled
          v3 = enabled
      Commands:
          cd
          set
          show
-> set v2c=enabled
```

6. Activate the NM2-36P Sun Datacenter InfiniBand Expansion Switch as an ASR asset in ASR Manager.

   a. Using SSH, log into the master management node.

   b. Activate the asset by adding its ILOM IP to the ASR configuration with the following command:

   ```
   # /opt/SUNWswasr/bin/asr activate_asset -i <asset_ilom_ip>
   ```

7. Repeat this procedure for the second InfiniBand switch in order to activate it as an ASR asset.

5.5.3 Activating ASR on the ZFS Storage Appliance

The ZFS Storage Appliance differs from the other ASR assets because it 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. See Section 5.4, "Configuring TinyProxy".

For detailed information about enabling ASR on the ZFS Storage Appliance, refer to the ASR documentation and the support note with Doc ID 1285455.1.

You must set up the ZFS Storage Appliance to relay its ASR data through the ASR Manager on the master management node, using the following proxy settings:

• **inet addr:** 192.168.4.216

- **broadcast:** 192.168.4.255

- **mask:** 255.255.255.0

- **port:** 8888

Enabling ASR from the CLI on the Active Storage Controller

1. Using SSH and an account with superuser privileges, log into the master management node, and then log into the active storage controller from there.

    > **Note**
    >
    > The data center IP address used in this procedure is an example.
    >
    > The two storage controllers of the ZFS Storage Appliance use the IP addresses 192.168.4.1 and 192.168.4.2.

    ```
    # ssh root@10.100.1.101
    root@10.100.1.101's password:
    [root@ovcamn05r1 ~]# ssh root@192.168.4.1

    ovcasn01r1:> ls

    Properties:
            showcode = false
            showstack = false
            exitcoverage = false
            showmessage = true
            asserterrs = false

     Children:
            analytics => Manage appliance analytics
            configuration => Perform configuration actions
            maintenance => Perform maintenance actions
            raw => Make raw XML-RPC calls
            shares => Manage shares
            status => View appliance status
    ```

2. Work your way down the configuration tree to the SCRK service.

    ```
    ovcasn01r1:> configuration
    ovcasn01r1:configuration> services
    ovcasn01r1:configuration services> scrk
    ovcasn01r1:configuration services scrk> ls

    Properties:
                            <status> = disabled
                              soa_id =
                        soa_password =
                            proxy_on = false
                          proxy_host =
                          proxy_user =
                      proxy_password =
    ```

3. Use the set command to make the following eight changes.

    > **Note**
    >
    > Make sure you have an active user account on My Oracle Support, because you need the user name and password to set the configuration parameters.

    ```
    ovcasn01r1:configuration services scrk> set soa_id=<registered-MOS-user>
    ```

```
ovcasn01r1:configuration services scrk> set soa_password=<MOS-user-password>
ovcasn01r1:configuration services scrk> set proxy_on=true
ovcasn01r1:configuration services scrk> set proxy_host=192.168.4.216:8888
ovcasn01r1:configuration services scrk> set proxy_user=root
ovcasn01r1:configuration services scrk> set proxy_password=Welcome1
ovcasn01r1:configuration services scrk> set updatecheck_on=false
ovcasn01r1:configuration services scrk> set time_updatecheck=7d
```

4. Commit the configuration changes and then verify that the new settings are correct.

```
ovcasn01r1:configuration services scrk> commit

ovcasn01r1:configuration services scrk> show
Properties:
                       <status> = online
                         soa_id = first.last@oracle.com
                   soa_password = ********
                       proxy_on = true
                     proxy_host = 192.168.4.216:8888
                     proxy_user = root
                 proxy_password = ********
                  updatecheck_on = false
               time_updatecheck = 7d
```

Chapter 6 Servicing Oracle Private Cloud Appliance Components

Table of Contents

This chapter contains an overview of replaceable components in your Oracle Private Cloud Appliance (PCA), and provides servicing instructions for customer-replaceable units.

6.1 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.

6.1.1 Rack Components

The following table lists the replaceable components of the Oracle PCA rack.

> **Note**
>
> For the current list of replacement parts and their manufacturing part numbers, refer to the Oracle PCA 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
Sun Rack II 1242:		
Jumper Cable C13-C14, 2m	FRU	Yes
Jumper Cable C19-C20, 1m	FRU	Yes
10Gbps QSFP to QSFP Cable, 3m	FRU	Yes
Ethernet Cable, Category 5/5E, RJ45 to RJ45, 10ft, Blue	FRU	Yes
Ethernet Cable, Category 5/5E, RJ45 to RJ45, 1m, Grey	FRU	Yes
Ethernet Cable, Category 5/5E, RJ45 to RJ45, 7ft, Black	FRU	Yes
Ethernet Cable, Category 5/5E, RJ45 to RJ45, 7ft, Green	FRU	Yes
Ethernet Cable, Category 5/5E, RJ45 to RJ45, 7ft, Yellow	FRU	Yes
Ethernet Cable, Category 6A, RJ45 to RJ45, 10ft, Black	FRU	Yes
Ethernet Cable, Category 6A, RJ45 to RJ45, 10ft, Blue	FRU	Yes
1U/2U Screw-Mount Slide Rail Kit	FRU	
1U/2U Cable Management Arm (Snap-in)	FRU	
10Gbps SFP+ TwinX Cable, 1m	FRU	Yes
1.25Gb/Sec Copper SFP Transceiver	FRU	
10Gbps QSFP to QSFP Cable, 1m	FRU	Yes
Power Distribution Units (PDUs):		

Component Description	FRU/CRU	Hot-Swap
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 6.3, "Servicing the Oracle Private Cloud Appliance Rack System".

6.1.2 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 PCA 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 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
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 6.4, "Servicing an Oracle Server X6-2".

6.1.3 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 PCA 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 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
1U/2U Remote Battery Assembly	CRU	No
Cable Kit	FRU	No

For Oracle Server X5-2 component servicing instructions, see Section 6.5, "Servicing an Oracle Server X5-2".

6.1.4 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 PCA 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 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 6.6, "Servicing a Sun Server X4-2".

6.1.5 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 PCA 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 Sun Server X3-2 Components

Component Description	FRU/CRU	Hot-Swap
System Board Assembly	FRU	No

Component Description	FRU/CRU	Hot-Swap
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 6.7, "Servicing a Sun Server X3-2".

6.1.6 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 PCA 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 Oracle ZFS Storage Appliance ZS3-ES Components

Component Description	FRU/CRU	Hot-Swap
Oracle ZFS Storage Appliance ZS3-ES Storage Head:		
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
4GB USB 2.0 Flash Drive	FRU	No

Component Description	FRU/CRU	Hot-Swap
2-Slot PCI Express Riser Assembly	FRU	No
1-Slot PCI Express Riser Assembly	FRU	No
Power Distribution Board	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
1U/2U Remote Battery Assembly	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
Oracle Storage DE2-24P Disk Shelf:		
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	
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 Oracle ZFS Storage Appliance ZS3-ES component servicing instructions, see Section 6.8, "Servicing the Oracle ZFS Storage Appliance ZS3-ES".

6.1.7 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 PCA 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 ZFS Storage Appliance 7320 Components

Component Description	FRU/CRU	Hot-Swap
Sun ZFS 7320 Storage Head:		

Component Description	FRU/CRU	Hot-Swap
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	
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
Oracle Storage DE2-24P Disk Shelf:		
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 6.9, "Servicing the Sun ZFS Storage Appliance 7320".

6.1.8 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 PCA 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 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 6.10, "Servicing an Oracle Switch ES1-24".

6.1.9 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 PCA 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 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 6.11, "Servicing an NM2-36P Sun Datacenter InfiniBand Expansion Switch".

6.1.10 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 PCA 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 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 × 8 Gigabit Fibre Channel I/O Module	FRU	Yes
F1-15 Chassis without Power Supply, Fan, Fabric Board, Front Panel	FRU	No

For Oracle Fabric Interconnect F1-15 component servicing instructions, see Section 6.12, "Servicing an Oracle Fabric Interconnect F1-15".

6.2 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 (E71902-02)*

- Follow the safety guidelines described in the *Oracle Private Cloud Appliance Installation Guide (E71900-02)*:

  - 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 PCA 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)

6.3 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 6.2, "Preparing Oracle Private Cloud Appliance for Service".

6.3.1 Powering Down Oracle Private Cloud Appliance (When Required)

Some service procedures may require you to power down the Oracle PCA.   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 User's Guide.

- Stopping Virtual Machines

- Stopping Oracle VM Servers

▋ •

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

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. Press the Power button on both Fabric Interconnects.

5. Toggle the rack Power switches to the Off position.

> **Note**
>
> The Ethernet and InfiniBand 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

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 and InfiniBand switches are powered on with the PDUs.

3. Press the Power button on each Fabric Interconnect.

   Wait until the Status LED stops blinking and is solid green, indicating the system has finished booting.

4. 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.

5. Press the Power button on the management nodes.

The management node that completes booting first assumes the master role.

6. When the management nodes are up, press the Power button on the compute nodes.

   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.

6.3.2 Service Procedures for Rack System Components

For parts that are not hot-swappable, power down the Oracle PCA before starting the service procedure. Generally speaking, hot-swappable components can be serviced without specific additional steps.

Table 6.11 Service Instructions for Rack System Components

Replaceable Part(s)	Hot-Swap	Instructions
Power cables		
Ethernet cables		
InfiniBand cables		
SFP transceivers (Oracle-qualified service technician only)		
Cable management arms (CMAs) (Oracle-qualified service technician only)		For removal and installation of a cable management arm, refer to the Oracle Server X5-2 Installation Guide (part no. E48314). • "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 X5-2 Service Manual (part no. E48320). • "Remove the Server From the Rack" • "Reinstall the Server Into the Rack" For slide rail installation instructions, refer to the section Attach Tool-less Slide-Rail Assemblies in the Oracle Server X5-2 Installation Guide (part no. E48314). To remove the slide rails, reverse the installation steps.

6.4 Servicing an Oracle Server X6-2

This section provides instructions to service replaceable components (CRUs/FRUs) in an Oracle Server X6-2 management node or compute node. Before starting any service procedure, read and follow the guidelines in Section 6.2, "Preparing Oracle Private Cloud Appliance for Service".

6.4.1 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:

> **Note**
>
> The management nodes are not placed in maintenance node 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 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. For details, refer to the section "Placing an Oracle VM Server into Maintenance Mode" in the Oracle VM User's Guide.

1. Log in to the Oracle VM Manager Web UI.

   For details, refer to the section "Section 4.2, "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. 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 X6-2 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.

6.4.2 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 PCA. 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 6.12 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

6.5 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 6.2, "Preparing Oracle Private Cloud Appliance for Service".

6.5.1 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 node 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. For details, refer to the section "Placing an Oracle VM Server into Maintenance Mode" in the Oracle VM User's Guide.

1. Log in to the Oracle VM Manager Web UI.

   For details, refer to the section "Section 4.2, "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://$manager\text{-}vip$:7002/ovm/console.

      Replace $manager\text{-}vip$ 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.

6.5.2 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 PCA. 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 6.13 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

6.6 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 6.2, "Preparing Oracle Private Cloud Appliance for Service".

6.6.1 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 node 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 User's Guide.

1. Log in to the Oracle VM Manager Web UI.

    For details, refer to the section "Section 4.2, "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.

6.6.2 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 PCA. 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 6.14 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
PCI Express cards (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E36975_01/html/E38045/z40000f91037409.html#scrolltoc

Replaceable Part(s)	Hot-Swap	URL
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

6.7 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 6.2, "Preparing Oracle Private Cloud Appliance for Service".

6.7.1 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 node 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 User's Guide.

1. Log in to the Oracle VM Manager Web UI.

   For details, refer to the section "Section 4.2, "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.

6.7.2 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 PCA. 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 6.15 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

6.8 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 6.2, "Preparing Oracle Private Cloud Appliance for Service".

6.8.1 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 PCA 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.

6.8.2 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 PCA. 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 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 (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E37831_01/html/E48559/ z40000f91037409.html#scrolltoc
Internal USB flash drive (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E37831_01/html/E48559/ z4000a6d1442801.html#scrolltoc
Battery	No	http://docs.oracle.com/cd/E37831_01/html/E48559/ z40003f01423753.html#scrolltoc
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

6.9 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 6.2, "Preparing Oracle Private Cloud Appliance for Service".

6.9.1 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.

6.9.2 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 PCA. 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 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__hardwa
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/E28317_01/html/E38247/ maintenance__hardware__details__7x20.html#maintenance__hardwa
Storage head power supplies	Yes	http://docs.oracle.com/cd/E28317_01/html/E38247/ maintenance__hardware__details__7x20.html#maintenance__hardwa
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__hardwa
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__hardwa
Battery	No	http://docs.oracle.com/cd/E28317_01/html/E38247/ maintenance__hardware__details__7x20.html#maintenance__hardwa
System indicator boards (Oracle-qualified service technician only)	Yes	http://docs.oracle.com/cd/E26765_01/html/E26399/ maintenance__hardware__details__7x20.html#maintenance__hardwa

Replaceable Part(s)	Hot-Swap	URL
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

6.10 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 6.2, "Preparing Oracle Private Cloud Appliance for Service".

6.10.1 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.

6.10.2 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 PCA. 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 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

6.11 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 6.2, "Preparing Oracle Private Cloud Appliance for Service".

6.11.1 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.

6.11.2 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 PCA. 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 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

6.12 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 6.2, "Preparing Oracle Private Cloud Appliance for Service".

6.12.1 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.

6.12.2 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 PCA. 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 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 (Oracle-qualified service technician only)	Yes	http://docs.oracle.com/cd/E38500_01/html/E50997/z40004411020496.html#scrolltoc
Front panel assembly, including system control processor (Oracle-qualified service technician only)	No	http://docs.oracle.com/cd/E38500_01/html/E50997/z40004411020323.html#scrolltoc http://docs.oracle.com/cd/E38500_01/html/E50997/z400037d1022426.html#scrolltoc

Chapter 7 Troubleshooting

Table of Contents

This chapter describes how to resolve a number of common problem scenarios.

7.1 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 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 PCA password.

2. Use the CLI, as described in Chapter 3, *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. If your proxy server expects a user name and password, this can be specified in the following format:

   ```
   http://username:password@IP:PORT
   ```

   • 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.

7.2 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 PCA. 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 PCA password.

2. Use the CLI, as described in Chapter 3, *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 LEVEL
```

     Where *LEVEL* 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.

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.

7.3 Configuring Data Center Switches for VLAN Traffic

The Oracle PCA 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. 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 PCA 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 PCA 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 PCA.

7.4 Enabling Support Tools on Systems Provisioned with Software Release 1.0.2

Oracle PCA controller software Release 1.1.1 and later installs the support tools on the management nodes and compute nodes. However, compute nodes provisioned and discovered in an earlier release do not have the packages for these support tools. After you update the appliance software to the new release, any compute node you provision or reprovision will be installed with the Oracle PCA support tools. On running compute nodes that were provisioned with controller software Release 1.0.2, and already host virtual machines, you can install the support tools manually by following the instructions in this section.

Installing Support Tools Manually on a Compute Node

1. Make sure that the appliance software has been updated successfully to Release 1.1.3.

2. Using SSH and an account with superuser privileges, log into the compute node where you want to install the support tools.

   You can connect to a compute node from a workstation connected to the appliance management network. Alternatively, log into a management node first, and then connect to the compute node through either the appliance management network (`192.168.4.x`) or the Oracle VM management network (`192.168.140.x`). The IP address of a discovered compute node can be found in the *Servers and VMs* tab of Oracle VM Manager.

   > **Note**
   >
   > The management node 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 ~]# ssh root@192.168.4.6
   root@192.168.4.6's password:
   [root@ovcacn09r1 ~]#
   ```

3. On the compute node, mount the shared partition on the storage appliance where the software install media is unpacked.

   ```
   [root@ovcacn09r1 ~]# mkdir /tmp/shared_storage
   [root@ovcacn09r1 ~]# mount 192.168.40.1:/export/MGMT_ROOT /tmp/shared_storage
   ```

4. Locate and install the `ovca-support` package.

   ```
   [root@ovcacn09r1 ~]# ls -l /tmp/shared_storage/ovm_image/Server/ovca-support*
   -r--r--r--+ 1 root root 162726 Mar 23 16:14 /tmp/shared_storage/ovm_image/Server/ \
   ovca-support-1.1.1-106.el5.noarch.rpm
   [root@ovcacn09r1 ~]# rpm -ivh /tmp/shared_storage/ovm_image/Server/ovca-support*
   warning: /tmp/shared_storage/ovm_image/Server/ovca-support-1.1.1-106.el5.noarch.rpm:
       Header V3 DSA signature: NOKEY, key ID 1e5e0159
   ```

```
    Preparing...                          ######################################## [100%]
       1:ovca-support                     ######################################## [100%]
[root@ovcacn09r1 ~]#
```

5. Verify that the package has been installed correctly, for example by checking the software version as follows:

```
[root@ovcacn09r1 ~]# ovca-diag version
1.1.3_3_2014-05-19
[root@ovcacn09r1 ~]#
```

6. Unmount the shared partition and remove the temporary directory you created.

```
[root@ovcacn09r1 ~]# umount /tmp/shared_storage
[root@ovcacn09r1 ~]# rmdir /tmp/shared_storage
[root@ovcacn09r1 ~]#
```

7. Log out of the compute node.

8. Repeat this procedure on all other compute nodes installed in your Oracle PCA environment.

7.5 Verifying and Re-applying Oracle VM Manager Tuning after Software Update

During a Software Update from Release 1.0.2 to Release 1.1.3 it may occur that certain Oracle VM Manager tuning settings are not applied properly and default settings are used instead. After updating the Oracle PCA software stack, you must verify these tuning settings, and re-apply them if necessary. Run the following procedure:

Verifying Oracle VM Manager Tuning Settings

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. Verify that you are logged in to the **master** management node.

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

If the command returns MASTER: False, log in to the other management node and run the same command.

3. Log in to the Oracle VM shell as the admin user.

```
# /usr/bin/ovm_shell.sh -u admin
Password:
OVM Shell: 3.2.<version_id> Interactive Mode
---
>>>
```

4. At the Oracle VM shell prompt, enter the following command:

```
>>> OvmClient.getOvmManager().getFoundryContext().getModelManager().getMaxCacheSize()
150000
```

```
>>>
```

To exit Oracle VM shell, press `Ctrl+D`.

If the value returned is **not** `300000`, proceed with the next step.

5. From the Oracle Linux command line on the master management node, apply the required Oracle VM Manager tuning settings by running the following Oracle VM shell script as the `admin` user:

```
# /usr/bin/ovm_shell.sh -u admin -i /var/lib/ovca/ovm_scripts/ovmm_tuning.py
Password:
live events max age: 24 hours
archive events max age: 72 hours
max cache size: 150000 objects
live jobs max age: 168 hours
archive jobs max age: 14 hours
live jobs max age (after): 24 hours
archive jobs max age (after): 168 hours
live events max age (after): 3 hours
archive events max age (after): 6 hours
max cache size (after): 300000 objects
```

6. When the tuning script completes successfully, log out of the master management node.

7.6 Upgrading Existing Compute Node Configuration from Release 1.0.2

When you run a Software Update on an existing Oracle PCA installation with controller software Release 1.0.2, you must execute an additional upgrade procedure on provisioned compute nodes. The procedure is designed to reconfigure the compute node storage connections and the Oracle VM Agent, allowing the compute node to continue to function reliably as a member of the clustered server pool after the software update. By following the instructions in this section, you also make sure that previously deployed virtual machines remain fully functional when the appliance update to Release 1.1.3 is complete.

> **Warning**
>
> First update any Release 1.0.2 system to appliance software Release 1.1.3. Execute the procedure in this section on each compute node provisioned with software Release 1.0.2, *after* the Release 1.1.3 software update on the management nodes has completed successfully. Only resume normal operations or proceed with further updates when these steps have been completed.

> **Note**
>
> A large part of this procedure must be executed from Oracle VM Manager. An overview of its usage in the context of Oracle PCA, and links to detailed instructions in the Oracle VM documentation, are provided in this Administrator's Guide, in Chapter 4, *Managing the Oracle VM Virtual Infrastructure*.

Upgrading the Compute Node Configuration from Release 1.0.2 to Release 1.1.3

1. Make sure that the appliance software has been updated successfully to Release 1.1.3.

   You can verify this by logging into the master management node and entering the following command:

```
# ovca-diag version
1.1.3_<build>_<date>
```

2. Log in to Oracle VM Manager.

For details, see Section 4.2, "Logging in to the Oracle VM Manager Web UI".

3. Identify the master server in the server pool.

   a. In the **Servers and VMs** tab, select the server pool in the navigation pane.

   b. In the management pane, select the **Info** perspective from the list in the toolbar.

      A configuration overview of the server pool is displayed.

   c. Locate the *Master Server* setting and write down which server is the pool master.

   > **Warning**
   >
   > The server pool master server must be the last Oracle VM Server to go through this upgrade procedure.

4. 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 User's Guide section entitled Migrating Virtual Machines.

5. Place the compute node in maintenance mode.

   Information on maintenance mode is provided in the Oracle VM User's Guide section entitled Placing an Oracle VM Server into 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. 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.

6. Install the `ovca-support` package by executing the steps described in Section 7.4, "Enabling Support Tools on Systems Provisioned with Software Release 1.0.2". However, do not log out of the compute node yet.

7. From the Oracle Linux command line on the compute node, run this script:

   ```
   # /usr/sbin/cn_upg_111.sh
   ```

8. When the upgrade script completes successfully, log out of the compute node and return to the Oracle VM Manager UI.

9. Run the Oracle VM Server update on the compute node in question.

   Information on the Oracle VM Server update functionality is provided in the Oracle VM User's Guide section entitled Updating and Upgrading Oracle VM Servers.

   a. In the **Servers and VMs** tab, select the server pool in the navigation pane.

   b. In the management pane, select the **Servers** perspective from the list in the toolbar.

      A table is displayed with all servers in the server pool.

c. Select the appropriate compute node in the management pane table and click **Update Server** in the management pane toolbar.

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.

10. 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 **Maintenance Mode** check box. Click OK.

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

11. If the compute node is in *WARNING* status, unpresent the *Rack1-Repository* from this node, and remove the node from the *Rack1_ServerPool*.

    a. In the **Repositories** tab, select the *Rack1-Repository* in the management pane. Click **Present/ Unpresent Repository** in the management pane toolbar.

    The **Present Repository to Servers** dialog box is displayed.

    b. Select the **Servers** option at the top of the dialog box. Move the server in question out of the *Present to Server* pane. Click OK.

    The repository is unpresented from the selected Oracle VM Server.

    c. In the **Servers and VMs** tab, select *Rack1_ServerPool* and click **Edit Server Pool**.

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

    d. In the **Servers** tab of the **Edit Server Pool** dialog box, move the server in question out of the *Selected Servers* pane. Click OK.

    The Oracle VM Server is removed from the server pool.

12. Repeat this procedure for each compute node that requires upgrade from Release 1.0.2 to Release 1.1.3.

> **Warning**
>
> The server pool master server must be the last Oracle VM Server to go through this upgrade procedure.

13. When all compute nodes have been upgraded and rebooted successfully, acknowledge all events for the Oracle VM Servers and the *Rack1-Repository*.

    Information on acknowledging events is provided in the Oracle VM User's Guide section entitled Acknowledging Events/Errors.

14. Remove the stale LUN mappings left behind from the previous storage configuration.

    a. Go to the **Storage** tab of the Oracle VM Manager UI.

    b. In the navigation pane, select the SAN Server named *OVCA_ZFSSA_Rack1*.

    c. In the management pane, select the LUN or LUNs in *WARNING* status and click **Delete Physical Disk** in the toolbar.

15. Re-add any node that was removed from the *Rack1_ServerPool*. Make sure the *Rack1-Repository* is presented to all servers in the pool.

If no errors have occurred during this procedure, the software update to Release 1.1.3 is now complete.

7.7 Enabling Fibre Channel Connectivity on a Provisioned Appliance

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 PCA 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 PCA CLI in interactive mode.

   ```
   # pca-admin
   Welcome to PCA! Release: 2.1.1
   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 vhbas 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 vhbas 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*.

7.8 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 changed 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:

```
# /u01/app/oracle/ovm-manager-3/bin/createBackup.sh -n ManualBackup1
```

2. In the Oracle PCA 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.

```
# 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
```

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.

7.9 Enabling Compute Node IPMI in Oracle VM Manager

Oracle PCA controller software Release 1.1.1 and later automatically configures the Intelligent Platform Management Interface (IPMI) for all compute nodes in Oracle VM Manager. Compute nodes provisioned and discovered in an earlier release do not have IPMI enabled in Oracle VM Manager. If compute nodes are shut down and have no IPMI configured, they must be powered on by pressing the power button on the front panel of the server. To be able to control the power status of compute nodes remotely, you may add their IPMI configuration manually by following the instructions in this section.

> **Note**
>
> The compute node's IPMI is the built-in Oracle Integrated Lights Out Manager (ILOM). Its default user name is root ; the default password is Welcome1 . Look up the ILOM IP in the Hardware View of the Oracle PCA Dashboard: roll the mouse over the compute node in question and note the IP address in the pop-up window.

Configuring IPMI for Compute Nodes

1. Log in to the Oracle VM Manager UI.

   For details, see Section 4.2, "Logging in to the Oracle VM Manager Web UI".

2. Go to the **Servers and VMs** tab.

3. In the navigation pane, select the compute node you wish to control through the IPMI.

   In the toolbar, click `Edit Server`.

4. In the Edit Server dialog box, go to the `IPMI` tab.

5. In the IPMI tab, make the following changes:

   • Select the *Enable Server IPMI* and *Change IPMI Password* check boxes.

   • Enter the administrative *User Name* and *Password* in the respective fields.

   • Enter the *IP Address* for the IPMI of this compute node.

6. Click OK. The compute node configuration in Oracle VM Manager is updated.

7.10 Enabling SNMP Server Monitoring

For troubleshooting or hardware monitoring, it may be useful to enable SNMP on the servers in your Oracle PCA. 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 PCA 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 PCA ISO image, which is either already installed on your system or downloaded and unpacked on the shared storage during the update process:

   • 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

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 PCA 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 PCA environment.

7.11 Using a Custom CA Certificate for SSL Encryption

By default, Oracle PCA and Oracle VM Manager use a self-signed SSL certificate. 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 PCA Dashboard and the Oracle VM Manager web interface run on Oracle WebLogic Server. Oracle WebLogic Server provides the functionality to update the digital certificate and keystore. To add your own trusted CA certificate and keystore, see the procedures set out in the Oracle WebLogic documentation:

- Configuring Keystores: http://docs.oracle.com/cd/E17904_01/apirefs.1111/e13952/taskhelp/security/ConfigureKeystoresAndSSL.html

- Configuring Identity and Trust: http://docs.oracle.com/cd/E23943_01/web.1111/e13707/identity_trust.htm

> **Caution**
>
> Each Oracle PCA management node runs its own instance of Oracle WebLogic Server. You must apply the same changes separately to the master and standby management node.

Accessing the Oracle WebLogic Server Administration Console

1. In your browser, enter the address `https://IP:7002/console` .

   In the address, `IP` refers to the IP address that you have configured for each of the management nodes during installation.

2. Log in with the user name `weblogic` and the global password.

   If you have changed the WebLogic Server password, enter the password that you configured for WebLogic Server when you reset this password.

7.12 A Compute Node Fails to Complete Provisioning

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 reprovision 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 PCA 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
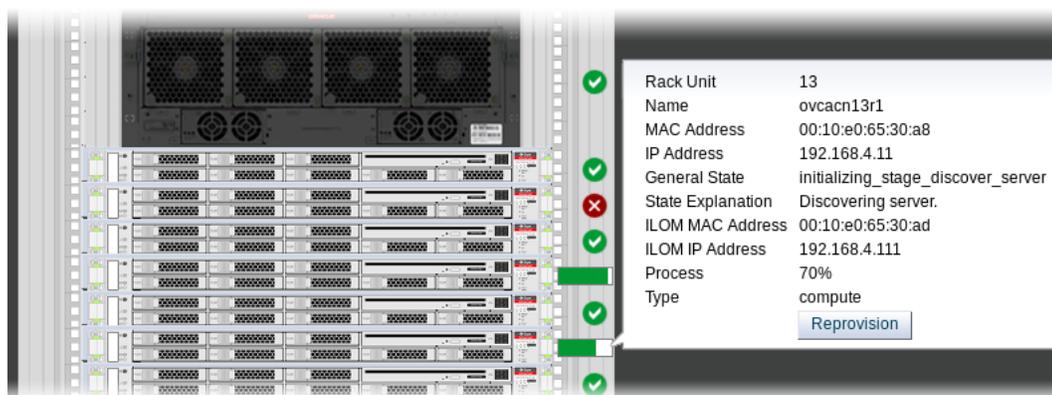
4. If the compute node provisioning is incomplete and the server is in error status or stuck in an intermittent state for several hours, click the **Reprovision** button in the pop-up window.

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 4.2, "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 User's Guide. Refer to the sections entitled Jobs Tab and Working with the Jobs Framework.

4. Remove the compute node from the Oracle VM server pool.

   Refer to the Oracle VM User's Guide and follow the instructions in the section entitled Removing an Oracle VM Server from a Server Pool.

5. Delete the compute node from Oracle VM Manager.

Refer to the Oracle VM User's Guide and follow the instructions in the section entitled Deleting Oracle VM Servers from Oracle VM Manager.

When the failing compute node has been removed from the Oracle VM configuration, return to the Oracle PCA Dashboard, to reprovision it. If the compute node is powered off and reprovisioning cannot be started, power on the server manually.

7.13 Oracle VM Server Pool Is Offline After Network Services Restart

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 PCA 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 PCA 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 PCA CLI. See Section 2.7.2, "Configuring Tenant Groups".

As a result, Oracle VM Manager and Oracle PCA 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 PCA 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 PCA 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 PCA 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 PCA 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 PCA are in sync again.

Index

A

B

C

D

E

F

H