

Oracle Private Cloud Appliance

Release Notes for Release 2.4.4



F37609-04
May 2022



Oracle Private Cloud Appliance Release Notes for Release 2.4.4,

F37609-04

Copyright © 2022, Oracle and/or its affiliates.

Contents

Preface

Audience	vii
Related Documentation	vii
Feedback	viii
Conventions	viii
Documentation Accessibility	viii
Access to Oracle Support for Accessibility	viii
Diversity and Inclusion	viii

1 About Oracle Private Cloud Appliance

2 Feature Overview

Changes and Improvements in Release 2.4.4	2-1
Changes and Improvements in Release 2.4.3	2-2
Changes and Improvements in Release 2.4.2	2-3
Earlier Releases	2-5

3 Product Accessibility

Oracle JET User Interface Accessibility Features	3-1
Documentation Accessibility	3-1

4 Configuration Maximums

Virtual Machine Maximums	4-1
Oracle VM Server Maximums	4-2
Server Pool and Cluster Maximums	4-3
Storage Maximums	4-3
Networking Maximums	4-4
Compute Node Maximums	4-5

5 Firmware Qualification

Qualified Firmware for Oracle Private Cloud Appliance Controller Software Release 2.4.4	5-1
---	-----

6 Known Limitations and Workarounds

Oracle Private Cloud Appliance Hardware	6-1
Compute Node Boot Sequence Interrupted by LSI Bios Battery Error	6-1
Reboot From Oracle Linux Prompt May Cause Management Node to Hang	6-1
Oracle ZFS Storage Appliance Firmware Upgrade 8.7.20 Requires A Two-Phased Procedure	6-1
Interruption of iSCSI Connectivity Leads to LUNs Remaining in Standby	6-2
Emulex Fibre Channel HBAs Discover Maximum 128 LUNs	6-2
Fibre Channel LUN Path Discovery Is Disrupted by Other Oracle VM Operations	6-2
Poor Oracle VM Performance During Configuration of Fibre Channel LUNs	6-3
ILOM Firmware Does Not Allow Loopback SSH Access	6-3
Megaraid Firmware Crash Dump Is Not Available	6-4
North-South Traffic Connectivity Fails After Restarting Network	6-4
Some Services Require an Upgrade of Hardware Management Pack	6-4
Compute Node USB Device Unused	6-4
Compute Nodes Containing FC HBA With Maximum FC Paths Paths Ends Up in Dead State During Reprovisioning	6-5
Compute Node FC HBA (Qlogic/Emulex) with FC LUNs Having Path Flapping	6-5
Upgrade Compute Node Fails with Fibre Channel LUNs	6-6
PCA Faultmonitor Check firewall_monitor Fails Due to nc: command not found	6-6
Oracle Private Cloud Appliance Software	6-7
Do Not Install Additional Software on Appliance Components	6-7
Node Manager Does Not Show Node Offline Status	6-7
Compute Node State Changes Despite Active Provisioning Lock	6-7
Compute Nodes Are Available in Oracle VM Server Pool Before Provisioning Completes	6-8
Virtual Machines Remain in Running Status when Host Compute Node Is Reprovisioned	6-8
Ethernet-Based System Management Nodes Have Non-Functional bond0 Network Interface	6-8
Network Performance Is Impacted by VxLAN Encapsulation	6-9
Altering Custom Network VLAN Tag Is Not Supported	6-9
Configuring Uplinks with Breakout Ports Results in Port Group Named 'None'	6-9
DPM Server Pool Policy Interrupts Synchronization of Tenant Group Settings	6-10
Host Network Parameter Validation Is Too Permissive	6-10
Virtual Appliances Cannot Be Imported Over a Host Network	6-11
Customizations for ZFS Storage Appliance in multipath.conf Are Not Supported	6-11
Customer Created LUNs Are Mapped to the Wrong Initiator Group	6-11

Storage Head Failover Disrupts Running Virtual Machines	6-12
Changing Multiple Component Passwords Causes Authentication Failure in Oracle VM Manager	6-12
ILOM Password of Expansion Compute Nodes Is Not Synchronized During Provisioning	6-12
SSH Host Key Mismatch After Management Node Failover	6-13
External Storage Cannot Be Discovered Over Data Center Network	6-13
Mozilla Firefox Cannot Establish Secure Connection with User Interface	6-13
Virtual Machine with High Availability Takes Five Minutes to Restart when Failover Occurs	6-14
CLI Command update appliance Is Deprecated	6-14
Certain CLI Commands Fail in Single-command Mode	6-14
Upgrader Checks Logged in Different Order	6-15
Virtual Machine Loses IP Address Due to DHCP Timeout During High Network Load	6-15
Adding the Virtual Machine Role to the Storage Network Causes Cluster to Lose Heartbeat Networking	6-15
Adding Virtual Machine Role to the Management Network Causes Oracle VM Manager to Lose Contact with the Compute Nodes	6-15
Inadvertant Reboot of Stand-by Management Node During Upgrade Suspends Upgrade	6-16
Loading Incompatible Spine Switch Configuration Causes Storage Network Outage	6-16
Cloud Backup Task Hangs When a ZFSSA Takeover is Performed During Backup	6-17
Export VM to Oracle Cloud Infrastructure Job Shows as Aborted During MN Failover but it is Running in the Background	6-17
Remove Deprecated pca-admin diagnose software Command	6-17
Virtual Machine get Message Failed After 200 Seconds - Observed When kube clusters are Created Concurrently	6-17
Kube Cluster Creation/Deletion Should Not Be in Progress When Management Node Upgrade is Initiated	6-18
o2cb Service Status Reports "Registering O2CB cluster "ocfs2": Failed" State After Compute Node Provisioned	6-18
Compute Node Upgrade Restores Default Repository When Compute Node Was Previously Not Part of Any Tenant Group or Repository	6-19
Check of Local Repository to Ensure Target Compute Node is Empty	6-19
Check No Physical Disks on Target Compute Node Have Repositories	6-21

7 Feedback and Support

Providing Feedback and Reporting Problems	7-1
Contacting Oracle Specialist Support	7-1
Data Collection for Service and Support	7-2

Preface

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

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

The documentation set consists of the following items:

Oracle Private Cloud Appliance Release Notes

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

Oracle Private Cloud Appliance Licensing Information User Manual

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

Oracle Private Cloud Appliance Installation Guide

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

Oracle Private Cloud Appliance Safety and Compliance Guide

The safety and compliance guide is a supplemental guide to the safety aspects of Oracle Private Cloud Appliance.

Oracle Private Cloud Appliance Administrator's Guide

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

Oracle Private Cloud Appliance Quick Start Poster

The quick start poster provides a step-by-step description of the hardware installation and initial software configuration of Oracle Private Cloud Appliance. A printed quick start poster is shipped with each Oracle Private Cloud Appliance base rack, and is intended for data center operators and administrators who are new to the product. The quick start poster is also available in the documentation set as an HTML guide, which contains alternate text for ADA 508 compliance.

Oracle Private Cloud Appliance Expansion Node Setup Poster

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

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

Audience

The Oracle Private Cloud Appliance documentation is written for technicians, authorized service providers, data center operators and system administrators who want to install, configure and maintain a private cloud environment in order to deploy virtual machines for users. It is assumed that readers have experience installing and troubleshooting hardware, are familiar with web and virtualization technologies and have a general understanding of operating systems such as UNIX (including Linux) and Windows.

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

Related Documentation

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



Note:

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

- Oracle Rack Cabinet 1242
https://docs.oracle.com/cd/E85660_01/index.html
- Oracle Server X86 Servers
<https://docs.oracle.com/en/servers/index.html>
- Oracle ZFS Storage Appliance ZS7-2
https://docs.oracle.com/cd/F13758_01/index.html
- Oracle Integrated Lights Out Manager (ILOM)
https://docs.oracle.com/cd/E81115_01/index.html
- Oracle VM
<https://docs.oracle.com/en/virtualization/oracle-vm/index.html>
- Oracle Enterprise Manager Plug-in
<https://docs.oracle.com/en/enterprise-manager/cloud-control/enterprise-manager-cloud-control/13.4/empca/index.html>

Feedback

Provide feedback about this documentation at:

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

Conventions

The following text conventions are used in this document:

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

Documentation Accessibility

For information about Oracle's commitment to accessibility, visit the Oracle Accessibility Program website at <https://www.oracle.com/corporate/accessibility/>.

Access to Oracle Support for Accessibility

Oracle customers that have purchased support have access to electronic support through My Oracle Support. For information, visit <https://www.oracle.com/corporate/accessibility/learning-support.html#support-tab>.

Diversity and Inclusion

Oracle is fully committed to diversity and inclusion. Oracle respects and values having a diverse workforce that increases thought leadership and innovation. As part of our initiative to build a more inclusive culture that positively impacts our employees, customers, and partners, we are working to remove insensitive terms from our products and documentation. We are also mindful of the necessity to maintain compatibility with our customers' existing technologies and the need to ensure continuity of service as Oracle's offerings and industry standards evolve. Because of these technical constraints, our effort to remove insensitive terms is ongoing and will take time and external cooperation.

1

About Oracle Private Cloud Appliance

Oracle Private Cloud Appliance is an Oracle Engineered System designed for virtualization. It 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. Compute resources, network hardware, storage providers, operating systems and applications are engineered to work together but are managed and operated as a single unit.

Installation, configuration, high availability, expansion and upgrading are automated and orchestrated as much as possible. 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. 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).

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

A list of supported operating systems is available in the release notes for the version of Oracle VM in use. All Oracle VM release notes are available at: <https://docs.oracle.com/en/virtualization/oracle-vm/index.html>



Note:

You must comply with all licensing requirements for OS images, other applications, and third-party hardware that you supply.

To view system components, see "Hardware Components" in [Concept, Architecture and Life Cycle of Oracle Private Cloud Appliance](#) in the *Oracle Private Cloud Appliance Administration Guide*.

2

Feature Overview

This chapter provides an overview of the key features of Oracle Private Cloud Appliance.

Changes and Improvements in Release 2.4.4

This section describes functional changes, improvements and bug fixes compared to the previous release.

Oracle Server X9-2 Compute Nodes

This release introduces support for the Oracle Server X9-2 as a compute node for Oracle Private Cloud Appliance.

Software Updates

This release introduces support for Oracle Linux 7 UEK5 and Oracle VM 3.4.7. When you upgrade to Software Controller release 2.4.4, compute nodes and management nodes are reprovisioned with these software updates. For more information about the update procedure, see [Updating Oracle Private Cloud Appliance](#) in *Oracle Private Cloud Appliance Administration Guide*. Release 2.4.4 is currently supported on Ethernet-based racks only.

Oracle Private Cloud Appliance Cloud Native Environment Release 1.2

 **Caution:**

This feature is no longer supported. Kubernetes functions are now available through [Oracle Cloud Native Environment](#).

This release introduces support for the Oracle Private Cloud Appliance Cloud Native Environment Release 1.2 (formerly Kubernetes Engine). The new OVA template (`pca-k8s-1-2-0.ova`) is now the default template for the `create kube-cluster` command.

Ksplice Support

For Ksplice Support on Oracle Private Cloud Appliance 2.4.4 (or later) compute nodes, contact Oracle Support.

Bugs Fixed in Release 2.4.4

The following table lists notable bugs that have been fixed in Oracle Private Cloud Appliance Release 2.4.4.

Table 2-1 List of Fixed Bugs

Bug ID	Description
31811656	"PCA 2.4.3 upgrader fails on storage-network update"
32717711	"pca-admin show a network which doesn't exist in CN node"
32710264	"Changing switch password via dashboard is not persistent on X8-2 systems"
31518408	"Compute node crashes when generating an sos report."
32780994	"PCA upgrade fails to detect running OVMM when JAVA_HOME is declared pointing to incorrect location"
31798407	"PCA 2.4.3 upgrader fails on Cisco steps"
30451985	"Remove obsolete checks from PCA Healthcheck tool"
32356688	"Upgrade cisco switches to add guestshell destroy"
32739839	"Upgrader iscsi precheck should check for unsupported iscsi target"
32047039	"VM Internal Storage networking causes errors during compute node reboot"

Changes and Improvements in Release 2.4.3

This section describes functional changes, improvements and bug fixes compared to the previous release.

Kubernetes as a Service

Caution:

This feature is no longer supported. Kubernetes functions are now available through [Oracle Cloud Native Environment](#).

The Oracle Private Cloud Appliance now supports Kubernetes as a Service. Kubernetes as a Service automates the provisioning of Oracle VM infrastructure and Kubernetes components to provide an integrated solution for Oracle Private Cloud Appliance.

Enhanced Oracle Private Cloud Appliance Fault Monitoring

The new fault monitoring architecture brings together the Oracle Private Cloud Appliance health check, the Oracle VM health check, and the PCA pre-upgrade checks under one service. In addition, you can configure health checking to integrate with ZFS Phone Home service to send reports on a weekly basis to Oracle Service for monitoring, using Oracle's [Automated Service Request \(ASR\)](#) infrastructure.

Oracle Private Cloud Appliance Cloud Backup

The Oracle Private Cloud Appliance Cloud Backup service automates the backup of critical components and configuration data to your customer tenancy in Oracle Cloud Infrastructure (OCI). This feature is designed to recover an Oracle Private Cloud Appliance to a running state after a catastrophic event. It is not designed to backup

virtual machines, guest operating systems, or applications and data hosted on virtual machines.

Support for Oracle VM Exporter Appliance on Oracle Private Cloud Appliance

An Oracle VM Exporter Appliance is a special type of virtual machine used to export another virtual machine from the Oracle VM environment to a tenancy account in the Oracle Cloud Infrastructure.

Storage Network Provides Internal Access to ZFS Storage Appliance

Oracle Private Cloud Appliance administrators can now create custom networks that enable users to access iSCSI and NFS shares on the internal ZFS storage appliance from within their individual Oracle VM guests hosted on the Oracle Private Cloud Appliance. Ensure you have 60 TB of free space on the ZFS storage appliance to use this feature.

Bugs Fixed in Release 2.4.3

The following table lists notable bugs that have been fixed in Oracle Private Cloud Appliance Release 2.4.3.

Table 2-2 List of Fixed Bugs

Bug ID	Description
31683379	"PCA health check failed "Multicast Route Check" from one spine switch "
31133075	"check_storage_space always fails after a given failure scenario"
30594475	"PCA 2.4.2 pca_upgrader fails ERROR (precheck:148) [OVMM Default Network Check (Checking default OVMM networks exist on the management nodes)] Failed " Additional checks have been added to verify the system network configuration.
31176592	"MySQL lock issues on the active management node"
31157042	"Timezone: Incorrect java version in /usr/bin/java for PCA 2.4.3"
31607588	"Pick up the July CPU for JRE 8 on PCA MN" The latest component firmware is incorporated.
30982691	"UI:Dashboard javascript library requires update" The latest UI update is incorporated.
31629904	"Tracking - Pick up OVM Manager build 3.4.6.2684" The latest component firmware is incorporated.
31350870	"Deprecated SSH settings"

Changes and Improvements in Release 2.4.2

This section describes functional changes, improvements and bug fixes compared to the previous release.

Oracle Private Cloud Appliance Release 2.4.2 Software Supports Both Platform Architectures

Controller Software Release 2.4.2 supports both Ethernet-based and InfiniBand-based Oracle Private Cloud Appliances, which includes these enhancements:

- Oracle Private Cloud ApplianceThe **Command Line Interface** supports both architectures and displays only the applicable commands based on the system in question.
- The **Upgrader** has been improved to include more automated checks and works on both architectures by determining the rack type.
- A version of the **Oracle Server X8-2 expansion node** is available for each architecture.

Optional Fibre Channel Cards Available

Oracle Server X8-2 compute nodes can be ordered with optional physical FC cards with dual 32Gbit HBAs. This option is available for both architectures. Note that Fibre Channel HBAs can be installed in existing compute nodes. One or two HBAs are supported per Oracle Server X8-2 compute node. If adding Fibre Channel HBAs to Oracle Server X7-2 or older compute nodes, do not use Fibre Channel connections from the Infiniband hardware to those compute nodes.

Oracle JET User Interface

The Oracle JavaScript Extension Toolkit (JET) user interface leverages popular open-source technologies, and features a rich set of modern UI components with built-in accessibility and mobile support.

Oracle Private Cloud Appliance Health Check Utility

The Health Check utility is built on the framework of the Oracle Private Cloud Appliance Upgrader, and is included in the Upgrader package. It detects the appliance network architecture and runs the sets of health checks defined for the system in question.

Bugs Fixed in Release 2.4.2

The following table lists bugs that have been fixed in Oracle Private Cloud Appliance Release 2.4.2.

Table 2-3 List of Fixed Bugs

Bug ID	Description
26679231	"Error running chgnetid "ImportError: cannot import name disable_warnings"" An issue preventing the changing of network IDs has been fixed.
29681802	"Pickup AK8.8.6 to prevent discovery issue with SI plugin" An issue was fixed that prevented the discovery of Oracle Private Cloud Appliance from Oracle Enterprise Manager, which results in the failures on both the internal and externally attached ZFS storage appliances.
26784732	"PCA 2.3.1 upgrade fails due to bond0 not configured" Additional checks have been added to verify the system configuration on management nodes and compute nodes before upgrade.
29437665	"The PCA 2.3.4 upgrade process breaks ASR if it is installed." A fix is provided that maintains the ASR installation through the upgrade process.
22720391	"Support listing backups in CLI"

Table 2-3 (Cont.) List of Fixed Bugs

Bug ID	Description
29585636	"Unable to Use Uplink to Communicate With Server Out of Rack After Restarting Network" The Cisco switches shipped with Oracle Private Cloud Appliance Release 2.4.2 (Ethernet-based systems) include the latest NX-OS I7(7) firmware, which fixes this connectivity issue. For more information about this bug fix (CSCvj65578) and others, see the Cisco 9000 Series NX-OS Release Notes, Release 7.0(3)I7(7): https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/7-x/release/notes/70377_nxos_rn.html

Earlier Releases

For an overview of features, functional changes and bug fixes in earlier releases of Oracle Private Cloud Appliance, please refer to the release notes for the version in question. You can access the documentation libraries for all releases from the [Documentation Overview](#) page.

3

Product Accessibility

Oracle is committed to making its products, services and supporting documentation accessible and usable to the disabled community. This chapter contains information about the status of Oracle Private Cloud Appliance in terms of compliance with the Americans with Disabilities Action (ADA) requirements.

Oracle JET User Interface Accessibility Features

The Oracle JET User Interface is compliant with the Americans with Disabilities Action (ADA) requirements. For detailed accessibility information, refer to [Oracle JET and Accessibility](#).

Documentation Accessibility

The Oracle Private Cloud Appliance documentation is available in the following formats:

- HTML format is provided and is accessible.
- PDF format is provided but is not accessible.
- Printed versions of specific installation instructions are shipped with the product. Accessible equivalents in HTML format are provided in the online documentation library.

Documentation HTML Access Keys

To use the documentation without using a mouse, you can use HTML access keys. Enter the HTML access key for your browser, plus the access key letter. For example, using Mozilla Firefox, press Alt+Shift+n to go to the next page. See your web browser documentation to find the key combination to use for HTML access keys. The following table lists the tasks you can perform using access keys.

Task	Access Key
Go to the next page	n
Go to the previous page	p
Go to the document home page	h
Go up a level in the document	u
Activate the Contents tab	c
Activate the Search tab	s

In addition to the HTML access keys, the following keyboard shortcuts are available:

Task	Shortcut Key
Toggle hide and show the sidebar	Ctrl + Left Arrow
Toggle hide and show page header	Ctrl + Up Arrow

Documentation Accessibility Issues

The following are the known accessibility issues with Oracle Private Cloud Appliance documentation:

- HTML page heading levels might not start at h1 (**Bug:** 26717728)
- Tables are missing rowheaders (**Bug:** 24300058)
- docs.oracle.com Help drawer HTML heading tags have structure violation (**Bug:** 26560104)

4

Configuration Maximums

This chapter contains the configuration maximums for Oracle Private Cloud Appliance. The limits presented in the following tables represent tested, recommended limits, and are fully supported by Oracle.

NOT_SUPPORTED:

Please also respect all other limitations with regard to the functionality divided between Oracle VM and Oracle Private Cloud Appliance. For example, refer to the "Functional Network Limitations" section in [Monitoring and Managing Oracle Private Cloud Appliance](#) in the *Oracle Private Cloud Appliance Administration Guide*.

At all times, respect the warnings and cautions throughout the documentation, most notably those at the beginning of these chapters in the *Oracle Private Cloud Appliance Administration Guide*:

- [Monitoring and Managing Oracle Private Cloud Appliance](#)
- [Managing the Oracle VM Virtual Infrastructure](#)

Virtual Machine Maximums

Table 4-1 Virtual Machine Maximums

Item	Maximum	Notes
Virtual CPUs	48 per Oracle Server X8-2 64 per Oracle Server X9-2	Equals the maximum number of CPU threads for a single CPU socket.
Virtual RAM (32-bit guest)	63GB	
Virtual RAM (64-bit guest)	1000GB	The maximum is for a HVM guest. If more memory is assigned, performance could be adversely affected, and live migration is no longer possible. A PVHVM guest can use the installed amount of physical RAM minus a safe margin of 32GB for use by dom0. Windows guests with PV drivers are limited to 256GB virtual RAM.
Virtual NICs (paravirtualized guest)	31	
Virtual NICs (hardware virtualized guest)	8	

Table 4-1 (Cont.) Virtual Machine Maximums

Item	Maximum	Notes
Virtual Disks (paravirtualized guest)	PVHVM: 107	
Disks: virtual and other (hardware virtualized guest)	4	In the case of hardware virtualized guests, "disks" refers to virtual disks as well as IDE and SCSI hard drives, and cd-rom drives.

 **Note:**

These maximums do not account for CPUs allocated to any other domains, such as the control domain. The hypervisor can allocate subsets of the overall CPU, memory, and I/O resources of a server to any given logical domain.

Virtual Machine maximums assume a single virtual machine using all available resources.

Oracle VM Server Maximums

Table 4-2 Oracle VM Server Maximums

Item	Maximum	Notes
CPUs	48 per Oracle Server X8-2 64 per Oracle Server X9-2	The Oracle Server X8-2 contains two 24-core CPUs with hyperthreading enabled. The Oracle Server X9-2 contains two 32-core CPUs with hyperthreading enabled.
RAM	1.5TB per Oracle Server X8-2 2T Oracle Server X9-2	This is the maximum installed amount of RAM in each server.
Virtual Machines	depends on RAM size of compute node and VM	As described in Server Pool and Cluster Maximums , the total amount of disk space and RAM in use by virtual machines must not exceed the actual available storage space and physical memory. This restriction applies at the level of a single compute node as well as a server pool or tenant group.

 **Note:**

A limited amount of RAM is required by the hypervisor.

Server Pool and Cluster Maximums

Table 4-3 Server Pool and Cluster Maximums

Item	Maximum	Notes
Oracle VM Servers in a clustered server pool	default: 25 custom: 24	In a base rack configuration, the maximum number of installed compute nodes is 25. All compute nodes may be a member of the default server pool. However, the default server pool cannot be empty, so any server pool associated with a custom tenant group can contain up to 24 compute nodes.
Clusters per Oracle Private Cloud Appliance	8	This maximum includes the default server pool.
Virtual disk space	depends on repository size	The total amount of disk space in use by virtual machines in the same server pool must be lower than the size of the storage repositories presented to the server pool. If this limitation is exceeded, file systems may become disconnected from the virtual machines, and access to the virtual machines may be lost.
RAM usage	depends on compute node RAM	The total amount of RAM used by virtual machines in the same server pool must be lower than the amount of physical RAM installed in the compute nodes that belong to the server pool. If this limitation is exceeded, virtual machines will fail to start. Similarly, you cannot increase the RAM for a virtual machine beyond the available free RAM.

Storage Maximums

Table 4-4 Storage Maximums

Item	Maximum	Notes
iSCSI LUNs and paths per Compute Node	1000 LUNs 2000 paths	The limit is the total number of paths. In this regard, each identical LUN connected as a single multipath device must be considered as a separate path. Examples of supported configurations are: 1000 LUNs with 2 paths each, 500 LUNs with 4 paths each, or 250 LUNs with 8 paths each.
OCFS2 volume size	64TB	
Files per OCFS2 volume	30,000	
Virtual Disk size	10TB	The maximum is determined by the hardware capacity.

Table 4-4 (Cont.) Storage Maximums

Item	Maximum	Notes
Virtual IDE drives per device	4	
LUNs and paths per Oracle Server X8-2 compute nodes with fibre channel cards	1000 LUNs 4000 paths	Examples of supported configurations are: 500 LUNs with 8 paths each, or 1000 LUNs with 4 paths each.

Networking Maximums

Table 4-5 Networking Maximums

Item	Maximum	Notes
NICs per Bond	2	Network bond configurations must never be modified by the customer.
Custom networks per rack	internal: 14 external: 8 total combined max: 16	Oracle Private Cloud Appliance accepts a maximum of 14 internal networks and 8 external networks (8 across all 8 available Cisco Ports). This maximum includes the default internal network and the default external network.
Custom networks per compute node	internal: 5 external: 9	This maximum includes the default internal network and the default external network.
Storage networks	16	This is the maximum number of storage networks allowed to connect to the internal ZFS storage appliance. This maximum is separate from the internal/external network maximums.
VLANs per Server Pool	256	If the number of VLANs is larger than 256, the boot times of the servers are severely impacted. The maximum applies to any given tenant group or compute node.

Table 4-5 (Cont.) Networking Maximums

Item	Maximum	Notes
Packet Size (MTU) by interface type	Bond	Do not modify the maximum transfer unit (MTU) of a network interface, standard port or bond, except as documented explicitly in the Oracle Private Cloud Appliance Administration Guide.
	• active: 9054	
	• inactive: 1500	
	Ethernet	
	• active: 9054	
	• inactive: 1500	
	tun: 9054 bytes	
vxlan: 9000 bytes		
mgmt: 9000 bytes		
storage: 9000 bytes		

Compute Node Maximums

Table 4-6 Compute Node Maximums

PDU	CPU/Memory	Maximum Compute Nodes
24kVA	32 Core processor 1TB DRAM	22
22kVA	32 Core processor 1TB DRAM	20
15kVA	32 Core processor 1TB DRAM	10
24kVA	32 Core processor 2TB DRAM	17
22kVA	32 Core processor 2TB DRAM	15
15kVA	32 Core processor 2TB DRAM	6

5

Firmware Qualification

This chapter provides information about Oracle Private Cloud Appliance component firmware versions that have been qualified, through Oracle internal platform testing, for use with a specific release of the Oracle Private Cloud Appliance Controller Software.

▲ Caution:

Interim Firmware Patches

Oracle periodically releases firmware patches for many products, for example to address security vulnerabilities. It may occur that an important firmware patch is released for a component of Oracle Private Cloud Appliance outside of the normal Controller Software release schedule. When this occurs, the patches go through the same testing as all other appliance firmware, but they are not added to the qualified firmware list or the installation `.iso` for the affected Controller Software release.

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

Qualified Firmware for Oracle Private Cloud Appliance Controller Software Release 2.4.4

The following tables list the firmware versions that have been qualified by Oracle for Release 2.4.4. All required firmware packages are provided on the Controller Software ISO image.

Table 5-1 Qualified Firmware for Release 2.4.4

Component	Version	Patch
System Software	Oracle Private Cloud Appliance 2.4.4	
	<ul style="list-style-type: none">Oracle VM Manager 3.4.7Oracle VM Server 3.4.7Oracle Linux 7	
Kernel	<ul style="list-style-type: none">Management node: kernel-uek-4.14.35-2047.503.1.1.el7uek.x86_64.rpmCompute node: kernel-uek-4.14.35-2047.505.4.2.el7uek.x86_64.rpm	
Management/Compute Nodes		
Oracle Server X8-2 management node/compute node	SW 1.2.3	<ul style="list-style-type: none">ILOM 4.0.4.51.aBIOS 51.02.13.00

Table 5-1 (Cont.) Qualified Firmware for Release 2.4.4

Component	Version	Patch
Oracle Server X9-2 compute node	SW 1.0.1	
	<ul style="list-style-type: none"> • ILOM 5.0.2.20 .a • BIOS 61.02.04.00 	
RAID controller	X8-2: 4.740.00-8440	
	X9-2: 7.5.0.1152	
CX5 100GbE Adapter Dual Port QSFP28	X8-2: 16.28.1002	
	X9-2: 16.29.1414	
1.2TB 10K RPM Drive	X8-2: ORA9 (or above)	
240GB Intel M.2	X9-2: XC311132	
240GB Micron M.2	X9-2: D3MU001	
3.84TB Intel SFF	X9-2: 2CV1RC30	
3.84TB Samsung SFF	X9-2: MPK93R5Q	
Storage Appliance		
Oracle ZFS Storage Appliance ZS7-2 controller server	OS8.8.28	
14TB Drive	A07G (or above)	
Switches		
Cisco Nexus 9348GC-FXP Switch	7.0.3.I7.9 (NXOS and EPLD)	
Cisco Nexus 9336C-FX2 Switch	7.0.3.I7.9 (NXOS and EPLD)	

6

Known Limitations and Workarounds

This chapter provides information about the known limitations and workarounds for Oracle Private Cloud Appliance (PCA).

Oracle Private Cloud Appliance Hardware

This section describes hardware-related limitations and workarounds.

Compute Node Boot Sequence Interrupted by LSI Bios Battery Error

When a compute node is powered off for an extended period of time, a week or longer, the LSI BIOS may stop because of a battery error, waiting for the user to press a key in order to continue.

Workaround: Wait for approximately 10 minutes to confirm that the compute node is stuck in boot. Use the Reprovision button in the Oracle Private Cloud Appliance Dashboard to reboot the server and restart the provisioning process.

Bug 16985965

Reboot From Oracle Linux Prompt May Cause Management Node to Hang

When the reboot command is issued from the Oracle Linux command line on a management node, the operating system could hang during boot. Recovery requires manual intervention through the server ILOM.

Workaround: When the management node hangs during (re-)boot, log in to the ILOM and run these two commands in succession: `stop -f /SYS` and `start /SYS`. The management node should reboot normally.

Bug 28871758

Oracle ZFS Storage Appliance Firmware Upgrade 8.7.20 Requires A Two-Phased Procedure

Oracle PCA racks shipped prior to Release 2.3.4 have all been factory-installed with an older version of the Operating Software (AK-NAS) on the controllers of the ZFS Storage Appliance. A new version has been qualified for use with Oracle PCA Release 2.3.4, but a direct upgrade is not possible. An intermediate upgrade to version 8.7.14 is required.

Workaround: Upgrade the firmware of storage heads twice: first to version 8.7.14, then to version 8.7.20. Both required firmware versions are provided as part of the Oracle PCA Release 2.3.4 controller software. For upgrade instructions, refer to "Upgrading the Operating Software on the Oracle ZFS Storage Appliance" in [Updating Oracle Private Cloud Appliance](#) in the [Oracle Private Cloud Appliance Administration Guide](#).

Bug 28913616

Interruption of iSCSI Connectivity Leads to LUNs Remaining in Standby

If network connectivity between compute nodes and their LUNs is disrupted, it may occur that one or more compute nodes mark one or more iSCSI LUNs as being in *standby* state. The system cannot automatically recover from this state without operations requiring downtime, such as rebooting VMs or even rebooting compute nodes. The standby LUNs are caused by the specific methods that the Linux kernel and the ZFS Storage Appliance use to handle failover of LUN paths.

Workaround: This issue was resolved in the ZFS Storage Appliance firmware version AK 8.7.6. Customers who have run into issues with missing LUN paths and standby LUNs, should update the ZFS Storage Appliance firmware to version AK 8.7.6 or later before upgrading Oracle Private Cloud Appliance.

Bug 24522087

Emulex Fibre Channel HBAs Discover Maximum 128 LUNs

When using optional Broadcom/Emulex Fibre Channel expansion cards in Oracle Server X8-2 and Oracle Server X9-2 compute nodes, and your FC configuration results in more than 128 LUNs between the compute nodes and the FC storage hardware, it may occur that only 128 LUNs are discovered. This is typically caused by a driver parameter for Emulex HBAs.

Workaround: Update the Emulex lpfc driver settings by performing the steps below on each affected compute node.

1. On the compute node containing the Emulex card, modify the file `/etc/default/grub`. At the end of the `GRUB_CMDLINE_LINUX` parameter, append the `scsi_mod` and `lpfc` module options shown.

```
GRUB_CMDLINE_LINUX="crashkernel=auto rd.lvm.lv=vglvroot rd.lvm.lv=vglvswap \
rd.lvm.lv=vglvusrrhgb quiet numa=off transparent_hugepage=never \
scsi_mod.max_luns=4096
                                scsi_mod.max_report_luns=4096
                                lpfc.lpfc_max_luns=4096"
```

2. Rebuild the grub configuration with the new parameters.

```
# grub2-mkconfig -o /boot/efi/EFI/redhat/grub.cfg
```

3. Reboot the compute node.

Bug 30461433, 33114489

Fibre Channel LUN Path Discovery Is Disrupted by Other Oracle VM Operations

During the setup of Fibre Channel storage, when the zones on the FC switch have been created, the LUNs become visible to the connected compute nodes. Discovery operations are started automatically, and all discovered LUNs are added to the multipath configuration on the compute nodes. If the storage configuration contains a large number of LUNs, the multipath configuration may take a long time to complete.

As long as the multipath configuration has not finished, the system is under high load, and concurrent Oracle VM operations may prevent some of the FC LUN paths from being added to multipath.

Workaround: It is preferred to avoid Oracle VM operations during FC LUN discovery. Especially all operations related to compute node provisioning and tenant group configuration are disruptive, because they include refreshing the storage layer. When LUNs become visible to the compute nodes, they are detected almost immediately. In contrast, the multipath configuration stage is time-consuming and resource-intensive.

Use the `lsscsi` command to determine the number of detected LUN paths. The command output is equal to the number of LUN paths plus the system disk. Next, verify that all paths have been added to multipath. The multipath configuration is complete once the `multipath -ll` command output is equal to the output of the `lsscsi` command minus 1 (for the system disk).

```
# lsscsi | wc -l
251
# multipath -ll | grep "active ready running" | wc -l
250
```

When you have established that the multipath configuration is complete, all Oracle VM operations can be resumed.

Bug 30461555

Poor Oracle VM Performance During Configuration of Fibre Channel LUNs

Discovering Fibre Channel LUNs is a time-consuming and resource-intensive operation. As a result, Oracle VM jobs take an unusually long time to complete. Therefore, it is advisable to complete the FC storage configuration and make sure that the configuration is stable before initiating new Oracle VM operations.

Workaround: Schedule Fibre Channel storage setup and configuration changes at a time when no other Oracle VM operations are required. Verify that all FC configuration jobs have been completed, as explained in [Fibre Channel LUN Path Discovery Is Disrupted by Other Oracle VM Operations](#). When the FC configuration is finished, all Oracle VM operations can be resumed.

Bug 30461478

ILOM Firmware Does Not Allow Loopback SSH Access

In Oracle Integrated Lights Out Manager (ILOM) firmware releases newer than 3.2.4, the service processor configuration contains a field, called `allowed_services`, to control which services are permitted on an interface. By default, SSH is not permitted on the loopback interface. However, Oracle Enterprise Manager uses this mechanism to register Oracle Private Cloud Appliance management nodes. Therefore, SSH must be enabled manually if the ILOM version is newer than 3.2.4.

Workaround: On management nodes running an ILOM version more recent than 3.2.4, make sure that SSH is included in the `allowed_services` field of the network configuration. Log into the ILOM CLI through the `NETMGT` Ethernet port and enter the following commands:

```
-> cd /SP/network/interconnect
-> set hostmanaged=false
```

```
-> set allowed_services=fault-transport,ipmi,snmp,ssh
-> set hostmanaged=true
```

Bug 26953763

Megaraid Firmware Crash Dump Is Not Available

ILOM console logs may contain many messages similar to this:

```
[ 1756.232496] megaraid_sas 0000:50:00.0: Firmware crash dump is not available
[ 1763.578890] megaraid_sas 0000:50:00.0: Firmware crash dump is not available
[ 2773.220852] megaraid_sas 0000:50:00.0: Firmware crash dump is not available
```

These are notifications, not errors or warnings. The crash dump feature in the megaraid controller firmware is not enabled, as it is not required in Oracle Private Cloud Appliance.

Workaround: This behavior is not a bug. No workaround is required.

Bug 30274703

North-South Traffic Connectivity Fails After Restarting Network

This issue may occur if you have not up upgraded the Cisco Switch firmware to version NX-OS I7(7) or later. See "Upgrading the Cisco Switch Firmware" in [Updating Oracle Private Cloud Appliance](#) in the [Oracle Private Cloud Appliance Administration Guide](#).

Bug 29585636

Some Services Require an Upgrade of Hardware Management Pack

Certain secondary services running on Oracle Private Cloud Appliance, such as Oracle Auto Service Request or the Oracle Enterprise Manager Agent, depend on a specific or minimum version of the Oracle Hardware Management Pack. By design, the Controller Software upgrade does not include the installation of a new Oracle Hardware Management Pack or server ILOM version included in the ISO image. This may leave the Hardware Management Pack in a degraded state and not fully compatible with the ILOM version running on the servers.

Workaround: When upgrading the Oracle Private Cloud Appliance Controller Software, make sure that all component firmware matches the qualified versions for the installed Controller Software release. To ensure correct operation of services depending on the Oracle Hardware Management Pack, make sure that the relevant `oracle-hmp*.rpm` packages are upgraded to the versions delivered in the Controller Software ISO.

Bug 30123062

Compute Node USB Device Unused

USB devices on Oracle Private Cloud Appliance compute nodes are unused. If a compute node USB device name changes after an upgrade or reboot, there is no functional impact to the compute node or the Oracle Private Cloud Appliance.

Bug 32998531

Compute Nodes Containing FC HBA With Maximum FC Paths Paths Ends Up in Dead State During Reprovisioning

The cause of this failure is an FC LUN path flapping issue that is typically seen with large number of FC LUNs (approximately greater than 100 LUNs) presented to compute node during reprovisioning. If you encounter this follow the workaround.

Workaround:

1. Shutdown compute node and make sure it is powered off.
2. Log in to the external storage array and remove the compute node's FC initiator from the initiator group (the initiator group that was used to present luns to this CN).
This is temporary change. We will add this back later.
3. Restart the compute node.
4. Use the `multipath -F` command to verify that no FC LUNs are present (iSCSI LUNs can be present).
5. Reprovision the compute node.
6. (Emulex only) Log in to the compute node and re-apply the grub customizaiton for the Emulex driver, see [Emulex Fibre Channel HBAs Discover Maximum 128 LUNs](#).
7. Log in to the external storage and re-add the compute node's FC initiator into the initiator group.
8. Log in to the Oracle VM Manager UI and add the compute node as an admin server to the Unmanaged FibreChannel Storage Array.
9. Refresh the Unmanaged FibreChannel Storage Array.
10. If all FC LUNs are not discovered, you may need to apply the workaround for [Compute Node FC HBA \(Qlogic/Emulex\) with FC LUNs Having Path Flapping](#).

Bug 33134228

Compute Node FC HBA (Qlogic/Emulex) with FC LUNs Having Path Flapping

You may encounter path flapping when hundreds of FC LUNs are presented to a compute node in following scenarios:

- After a compute node reprovision
- After a compute node upgrade
- After exposing a compute node with 100's of new LUNs (either by LUNs creation on storage array or fabric rezoning)

If path flapping is occuring on your system, you will see the following errors on your compute node:

- `tailf /var/log/devmon.log` has floods of event messages similar to:

```
AGENT_NOTIFY EVENT : Jul 29 19:56:38 {STORAGE} [CHANGE_DM_SD] (dm-961)
3600144f0d987aa07000061027d9c48c6-10:0:0:1917
(failed:0x2100000e1e1b95c0:3600144f0d987aa07000061027d9c48c6)
AGENT_NOTIFY EVENT : Jul 29 19:56:38 {STORAGE} [CHANGE_DM_SD] (dm-988)
```

```

3600144f0d987aa07000061027db248e1-10:0:0:1971
(failed:0x2100000e1e1b95c0:3600144f0d987aa07000061027db248e1)
AGENT_NOTIFY EVENT : Jul 29 19:56:38 {STORAGE} [CHANGE_DM_SD] (dm-988)
3600144f0d987aa07000061027db248e1-10:0:0:1971
(active:0x2100000e1e1b95c0:3600144f0d987aa07000061027db248e1)
AGENT_NOTIFY EVENT : Jul 29 19:56:39 {STORAGE} [CHANGE_DM_SD] (dm-961)
3600144f0d987aa07000061027d9c48c6-10:0:0:1917
(active:0x2100000e1e1b95c0:3600144f0d987aa07000061027d9c48c6)

```

This issue is resolved when `/var/log/devmon.log` stops logging new `CHANGE_DM_SD` messages.

- `systemd-udevd` process consumes 100% CPU in `top` command.
This is resolved when `systemd-udevd` no longer consumes a large percentage of the CPU.
- `multipath -ll` command may not show all the LUNs. It may show a fraction of the LUNs expected.
- `multipath -ll | grep "active ready running" | wc -l` command may not count all the LUNs. It may show a fraction of the LUNs expected.

Workaround: Follow this procedure to resolve path flapping.

1. Log in to the compute node as root and execute the `systemctl restart multipathd` command.
2. Continue to execute the above detection commands until all 4 outputs are resolved, and you see the correct amount of FC LUNs/paths.
3. If any of the monitoring scenarios does not resolve after 3-4 minutes, repeat step 1.

Bug 33171816

Upgrade Compute Node Fails with Fibre Channel LUNs

If your compute node contains an Emulex or QLogic Fiber Channel HBA, the compute node upgrade procedure may fail because of the Fibre Channel lun path flapping problem. Follow this workaround to avoid the issue: [Oracle Support Document 2794501.1 \(\[PCA 2.4.4\] Upgrader Tool - Upgrade Compute Node Fails with FC Luns More Than Limits\)](#).

PCA Faultmonitor Check `firewall_monitor` Fails Due to `nc: command not found`

If your compute node fails the `Faultmonitor firewall_monitor` check with this log error:

```

[2021-08-03 16:30:15 605830] ERROR (ovmfaultmonitor_utils:487) invalid
literal for int() with base 10: '-bash: nc: command not found'
Traceback (most recent call last):
  File
"/usr/lib/python2.7/site-packages/ovca/monitor/faultmonitor/ovmfaultmonitor/ov
mfaultmonitor_utils.py", line 458, in firewall_monitor
    cmd_outputs[server][port] = int(output.strip())
ValueError: invalid literal for int() with base 10: '-bash: nc: command not
found.

```

you are encountering a port error which creates a false report and pushes it to Phonehome, if Phonehome is enabled. The `firewall_monitor` verifies whether the required ports for Oracle VM Manager and the compute node are opened or not. To manually fix this error apply the workaround documented at: [Oracle Support Document 2797364.1 \(\[PCA 2.4.4\] Faultmonitor Check firewall_monitor Fails due to "nc: command not found"\)](#).

Bug 33186553

Oracle Private Cloud Appliance Software

This section describes software-related limitations and workarounds.

Do Not Install Additional Software on Appliance Components

Oracle Private Cloud Appliance is delivered as an appliance: a complete and controlled system composed of selected hardware and software components. If you install additional software packages on the pre-configured appliance components, be it a compute node, management node or storage component, you introduce new variables that potentially disrupt the operation of the appliance as a whole. Unless otherwise instructed, Oracle advises against the installation or upgrade of additional packages, either from a third party or from Oracle's own software channels like the Oracle Linux YUM repositories.

Workaround: Do not install additional software on any internal Oracle Private Cloud Appliance system components. If your internal processes require certain additional tools, contact your Oracle representative to discuss these requirements.

Node Manager Does Not Show Node Offline Status

The role of the Node Manager database is to track the various states a compute node goes through during provisioning. After successful provisioning the database continues to list a node as *running*, even if it is shut down. For nodes that are fully operational, the server status is tracked by Oracle VM Manager. However, the Oracle Private Cloud Appliance Dashboard displays status information from the Node Manager. This may lead to inconsistent information between the Dashboard and Oracle VM Manager, but it is not considered a bug.

Workaround: To verify the status of operational compute nodes, use the Oracle VM Manager user interface.

Bug 17456373

Compute Node State Changes Despite Active Provisioning Lock

The purpose of a lock of the type `provisioning` or `all_provisioning` is to prevent all compute nodes from starting or continuing a provisioning process. However, when you attempt to reprovision a running compute node from the Oracle Private Cloud Appliance CLI while an active lock is in place, the compute node state changes to "reprovision_only" and it is marked as "DEAD". Provisioning of the compute node continues as normal when the provisioning lock is deactivated.

Bug 22151616

Compute Nodes Are Available in Oracle VM Server Pool Before Provisioning Completes

Compute node provisioning can take up to several hours to complete. However, those nodes are added to the Oracle VM server pool early on in the process, but they are not placed in maintenance mode. In theory the discovered servers are available for use in Oracle VM Manager, but you must not attempt to alter their configuration in any way before the Oracle Private Cloud Appliance Dashboard indicates that provisioning has completed.

Workaround: Wait for compute node provisioning to finish. Do not modify the compute nodes or server pool in any way in Oracle VM Manager.

Bug 22159111

Virtual Machines Remain in Running Status when Host Compute Node Is Re provisioned

Using the Oracle Private Cloud Appliance CLI it is possible to force the reprovisioning of a compute node even if it is hosting running virtual machines. The compute node is not placed in maintenance mode. Consequently, the active virtual machines are not shut down or migrated to another compute node. Instead these VMs remain in *running* status and Oracle VM Manager reports their host compute node as "N/A".

Caution:

Reprovisioning a compute node that hosts virtual machines is considered bad practice. Good practice is to migrate all virtual machines away from the compute node before starting a reprovisioning operation or software update.

Workaround: In this particular condition the VMs can no longer be migrated. They must be killed and restarted. After a successful restart they return to normal operation on a different host compute node in accordance with start policy defined for the server pool.

Bug 22018046

Ethernet-Based System Management Nodes Have Non-Functional `bond0` Network Interface

When the driver for network interface bonding is loaded, the system automatically generates a default `bond0` interface. However, this interface is not activated or used in the management nodes of an Oracle Private Cloud Appliance with the Ethernet-based network architecture.

Workaround: The `bond0` interface is not configured in any usable way and can be ignored on Ethernet-based systems. On InfiniBand-based systems, the `bond0` interface is functional and configured.

Bug 29559810

Network Performance Is Impacted by VxLAN Encapsulation

The design of the all-Ethernet network fabric in Oracle Private Cloud Appliance relies heavily on VxLAN encapsulation and decapsulation. This extra protocol layer requires additional CPU cycles and consequently reduces network performance compared to regular tagged or untagged traffic. In particular the connectivity to and from VMs can be affected. To compensate for the CPU load of VxLAN processing, the MTU (Maximum Transmission Unit) on VM networks can be increased to 9000 bytes, which is the setting across the standard appliance networks. However, the network paths should be analyzed carefully to make sure that the larger MTU setting is supported between the end points: if an intermediate network device only supports an MTU of 1500 bytes, then the fragmentation of the 9000 byte packets will result in a bigger performance penalty.

Workaround: If the required network performance cannot be obtained with a default MTU of 1500 bytes for regular VM traffic, you should consider increasing the MTU to 9000 bytes; on the VM network and inside the VM itself.

Bug 29664090

Altering Custom Network VLAN Tag Is Not Supported

When you create a custom network, it is technically possible – though not supported – to alter the VLAN tag in Oracle VM Manager. However, when you attempt to add a compute node, the system creates the network interface on the server but fails to enable the modified VLAN configuration. At this point the custom network is stuck in a failed state: neither the network nor the interfaces can be deleted, and the VLAN configuration can no longer be changed back to the original tag.

Workaround: Do not modify appliance-level networking in Oracle VM Manager. There are no documented workarounds and any recovery operation is likely to require significant downtime of the Oracle Private Cloud Appliance environment.

Bug 23250544

Configuring Uplinks with Breakout Ports Results in Port Group Named 'None'

When you split uplink ports for custom network configuration by means of a breakout cable, and subsequently start configuring the port pairs through the Oracle Private Cloud Appliance CLI, all four breakout ports are stored in the configuration database at the same time. This means that when you add the first two of four breakout ports to a port group, the remaining two breakout ports on the same cable are automatically added to another port group named "None", which remains disabled. When you add the second pair of breakout ports to a port group, "None" is replaced with the port group name of your choice, and the port group is enabled. The sequence of commands in the example shows how the configuration changes step by step:

```
PCA> create uplink-port-group custom_ext_1 '1:1 1:2' 10g-4x
Status: Success
```

```
PCA> list uplink-port-group
Port_Group_Name  Ports  Mode  Speed  Breakout_Mode  Enabled  State
-----
```



```

default_5_1      5:1 5:2   LAG   10g   10g-4x   True   (up) *
Not all ports are up
default_5_2      5:3 5:4   LAG   10g   10g-4x   False  down
custom_ext_1     1:1 1:2   LAG   10g   10g-4x   True   up
None             1:3 1:4   LAG   10g   10g-4x   False  up

```

```

-----
4 rows displayed
Status: Success

```

```

PCA> create uplink-port-group custom_ext_2 '1:3 1:4' 10g-4x
Status: Success

```

```

PCA> list uplink-port-group
Port_Group_Name  Ports      Mode      Speed      Breakout_Mode  Enabled  State
-----
default_5_1      5:1 5:2   LAG       10g       10g-4x       True    (up) *
Not all ports are up
default_5_2      5:3 5:4   LAG       10g       10g-4x       False   down
custom_ext_1     1:1 1:2   LAG       10g       10g-4x       True    up
custom_ext_2     1:3 1:4   LAG       10g       10g-4x       True    up

```

```

-----
4 rows displayed
Status: Success

```

Workaround: This behavior is by design, because it is a requirement that all four breakout ports must be added to the network configuration at the same time. When a port group is named "None", and it consists of two ports in a 4-way breakout cable, which are otherwise (temporarily) unconfigured, this can be ignored.

Bug 30426198

DPM Server Pool Policy Interrupts Synchronization of Tenant Group Settings

Tenant groups in Oracle Private Cloud Appliance are based on Oracle VM server pools, with additional configuration for network and storage across the servers included in the tenant group. When a compute node is added to a tenant group, its network and storage configuration is synchronized with the other servers already in the tenant group. This process takes several minutes, and could therefore be interrupted if a distributed power management (DPM) policy is active for the Oracle VM server pool. The DPM policy may force the new compute node to shut down because it contains no running virtual machines, while the tenant group configuration process on the compute node is still in progress. The incomplete configuration causes operational issues at the level of the compute node or even the tenant group.

Workaround: If server pool policies are a requirement, it is suggested to turn them off temporarily when modifying tenant groups or during the installation and configuration of expansion compute nodes.

Bug 30478940

Host Network Parameter Validation Is Too Permissive

When you define a host network, it is possible to enter invalid or contradictory values for the Prefix, Netmask and Route_Destination parameters. For example, when you enter a prefix with "0" as the first octet, the system attempts to configure IP addresses on compute node Ethernet interfaces starting with 0. Also, when the netmask part of

the route destination you enter is invalid, the network is still created, even though an exception occurs. When such a poorly configured network is in an invalid state, it cannot be reconfigured or deleted with standard commands.

Workaround: Double-check your CLI command parameters before pressing Enter. If an invalid network configuration is applied, use the `--force` option to delete the network.

Bug 25729227

Virtual Appliances Cannot Be Imported Over a Host Network

A host network provides connectivity between compute nodes and hosts external to the appliance. It is implemented to connect external storage to the environment. If you attempt to import a virtual appliance, also known as assemblies in previous releases of Oracle VM and Oracle Private Cloud Appliance, from a location on the host network, it is likely to fail, because Oracle VM Manager instructs the compute nodes to use the active management node as a proxy for the import operation.

Workaround: Make sure that the virtual appliance resides in a location accessible from the active management node.

Bug 25801215

Customizations for ZFS Storage Appliance in `multipath.conf` Are Not Supported

The ZFS stanza in `multipath.conf` is controlled by the Oracle Private Cloud Appliance software. The internal ZFS Storage Appliance is a critical component of the appliance and the multipath configuration is tailored to the internal requirements. You should never modify the ZFS parameters in `multipath.conf`, because it could adversely affect the appliance performance and functionality.

Even if customizations were applied for (external) ZFS storage, they are overwritten when the Oracle Private Cloud Appliance Controller Software is updated. A backup of the file is saved prior to the update. Customizations in other stanzas of `multipath.conf`, for storage devices from other vendors, are preserved during upgrades.

Bug 25821423

Customer Created LUNs Are Mapped to the Wrong Initiator Group

When adding LUNs on the Oracle Private Cloud Appliance internal ZFS Storage Appliance you must add them under the "OVM" target group. Only this default target group is supported; there can be no additional target groups. However, on the initiator side you should not use the default configuration, otherwise all LUNs are mapped to the "**All Initiators**" group, and accessible for all nodes in the system. Such a configuration may cause several problems within the appliance.

Additional, custom LUNs on the internal storage **must** instead be mapped to one or more custom initiator groups. This ensures that the LUNs are mapped to the intended initiators, and are not remapped by the appliance software to the default "**All Initiators**" group.

Workaround: When creating additional, custom LUNs on the internal ZFS Storage Appliance, always use the default target group, but make sure the LUNs are mapped to one or more custom initiator groups.

Bugs 22309236 and 18155778

Storage Head Failover Disrupts Running Virtual Machines

When a failover occurs between the storage heads of a ZFS Storage Appliance, virtual machine operation could be disrupted by temporary loss of disk access. Depending on the guest operating system, and on the configuration of the guest and Oracle VM, a VM could hang, power off or reboot. This behavior is caused by an iSCSI configuration parameter that does not allow sufficient recovery time for the storage failover to complete.

Workaround: Increase the value of `node.session.timeo.replacement_timeout` in the file `/etc/iscsi/iscsid.conf`. For details, refer to the support note with [Doc ID 2189806.1](#).

Bug 24439070

Changing Multiple Component Passwords Causes Authentication Failure in Oracle VM Manager

When several different passwords are set for different appliance components using the Oracle Private Cloud Appliance Dashboard, you could be locked out of Oracle VM Manager, or communication between Oracle VM Manager and other components could fail, as a result of authentication failures. The problem is caused by a partially failed password update, whereby a component has accepted the new password while another component continues to use the old password to connect.

The risk of authentication issues is considerably higher when Oracle VM Manager and its directly related components Oracle WebLogic Server and Oracle MySQL database are involved. A password change for these components requires the `ovmm` service to restart. If another password change occurs within a matter of a few minutes, the operation to update Oracle VM Manager accordingly could fail because the `ovmm` service was not active. An authentication failure will prevent the `ovmm` service from restarting.

Workaround: If you set different passwords for appliance components using the Oracle Private Cloud Appliance Dashboard, change them one by one with a 10 minute interval. If the `ovmm` service is stopped as a result of a password change, wait for it to restart before making further changes. If the `ovmm` service fails to restart due to authentication issues, it may be necessary to replace the file `/nfs/shared_storage/wsl/servers/AdminServer/security/boot.properties` with the previous version of the file (`boot.properties.old`).

Bug 26007398

ILOM Password of Expansion Compute Nodes Is Not Synchronized During Provisioning

After the rack components have been configured with a custom password, any compute node ILOM of a newly installed expansion compute node does not automatically take over the password set by the user in the Wallet. The compute node provisions correctly, and the Wallet maintains access to its ILOM even though it uses

the factory-default password. However, it is good practice to make sure that custom passwords are correctly synchronized across all components.

Workaround: Set or update the compute node ILOM password using the Oracle Private Cloud Appliance Dashboard or CLI. This sets the new password both in the Wallet and the compute node ILOM.

Bug 26143197

SSH Host Key Mismatch After Management Node Failover

When logging in to the active management node using SSH, you typically use the virtual IP address shared between both management nodes. However, since they are separate physical hosts, they have a different host key. If the host key is stored in the SSH client, and a failover to the secondary management node occurs, the next attempt to create an SSH connection through the virtual IP address results in a host key verification failure.

Workaround: Do not store the host key in the SSH client. If the key has been stored, remove it from the client's file system; typically inside the user directory in `.ssh/known_hosts`.

Bug 22915408

External Storage Cannot Be Discovered Over Data Center Network

The default compute node configuration does not allow connectivity to additional storage resources in the data center network. Compute nodes are connected to the data center subnet to enable public connectivity for the virtual machines they host, but the compute nodes' network interfaces have no IP address in that subnet. Consequently, SAN or file server discovery will fail.

Bug 17508885

Mozilla Firefox Cannot Establish Secure Connection with User Interface

Both the Oracle Private Cloud Appliance Dashboard and the Oracle VM Manager user interface run on an architecture based on Oracle WebLogic Server, Oracle Application Development Framework (ADF) and Oracle JDK 6. The cryptographic protocols supported on this architecture are SSLv3 and TLSv1.0. Mozilla Firefox version 38.2.0 or later no longer supports SSLv3 connections with a self-signed certificate. As a result, an error message might appear when you try to open the user interface login page.

Workaround: Override the default Mozilla Firefox security protocol as follows:

1. In the Mozilla Firefox address bar, type `about:config` to access the browser configuration.
2. Acknowledge the warning about changing advanced settings by clicking *I'll be careful, I promise!*.
3. In the list of advanced settings, use the Search bar to filter the entries and look for the settings to be modified.
4. Double-click the following entries and then enter the new value to change the configuration preferences:
 - `security.tls.version.fallback-limit: 1`
 - `security.ssl3.dhe_rsa_aes_128_sha: false`

- `security.ssl3.dhe_rsa_aes_256_sha: false`
5. If necessary, also modify the configuration preference `security.tls.insecure_fallback_hosts` and enter the affected hosts as a comma-separated list, either as domain names or as IP addresses.
 6. Close the Mozilla Firefox advanced configuration tab. The pages affected by the secure connection failure should now load normally.

Bug 21622475 and 21803485

Virtual Machine with High Availability Takes Five Minutes to Restart when Failover Occurs

The compute nodes in an Oracle Private Cloud Appliance are all placed in a single clustered server pool during provisioning. A clustered server pool is created as part of the provisioning process. One of the configuration parameters is the cluster time-out: the time a server is allowed to be unavailable before failover events are triggered. To avoid false positives, and thus unwanted failovers, the Oracle Private Cloud Appliance server pool time-out is set to 300 seconds. As a consequence, a virtual machine configured with high availability (HA VM) can be unavailable for 5 minutes when its host fails. After the cluster time-out has passed, the HA VM is automatically restarted on another compute node in the server pool.

This behavior is as designed; it is not a bug. The server pool cluster configuration causes the delay in restarting VMs after a failover has occurred.

CLI Command `update appliance` Is Deprecated

The Oracle Private Cloud Appliance command line interface contains the `update appliance` command, which is used in releases prior to 2.3.4 to unpack a Controller Software image and update the appliance with a new software stack. This functionality is now part of the Upgrader tool, so the CLI command is deprecated and will be removed in the next release.

Workaround: Future updates and upgrades will be executed through the Oracle Private Cloud Appliance Upgrader.

Bug 29913246

Certain CLI Commands Fail in Single-command Mode

The Oracle Private Cloud Appliance command line interface can be used in an interactive mode, using a closed shell environment, or in a single-command mode. When using the single-command mode, commands and arguments are entered at the Oracle Linux command prompt as a single line. If such a single command contains special characters, such as quotation marks, they may be stripped out and interpreted incorrectly.

Workaround: Use the CLI in interactive mode to avoid special characters being stripped out of command arguments. If you must use single-command mode, use single and double quotation marks around the arguments where required, so that only the outer quotation marks are stripped out. For example, change this command from:

```
# pca-admin create uplink-port-group myPortGroup '2:1 2:2' 10g-4x
```

to

```
# pca-admin create uplink-port-group myPortGroup "'2:1 2:2'" 10g-4x
```

Do not use doubles of the same quotation marks.

Bug 30421250

Upgrader Checks Logged in Different Order

Due to a change in how the Oracle Private Cloud Appliance Upgrader test are run, the output of the checks could be presented in a different order each time the tests are run.

This behavior is not a bug. There is no workaround required.

Bug 30078487

Virtual Machine Loses IP Address Due to DHCP Timeout During High Network Load

When an Oracle Private Cloud Appliance is configured to the maximum limits and a high load is running, a situation may occur where general DHCP/IP bandwidth limits are exceeded. In this case the DHCP client eventually reaches a timeout, and as a result the virtual machine IP address is lost, then reset to 0.0.0.0. This is normal behavior when the system is operating at full bandwidth capacity.

Workaround: When adequate bandwidth is available, recover from the situation by issuing the `dhclient` command from the virtual machine to request a new IP address.

Bug 30143723

Adding the Virtual Machine Role to the Storage Network Causes Cluster to Lose Heartbeat Networking

Attempting to add the Virtual Machine role to the storage network in Oracle VM Manger on an Oracle Private Cloud Appliance can cause your cluster to lose heartbeat networking, which will impact running Virtual Machines and their workloads. This operation is not supported on Oracle Private Cloud Appliance.

Workaround: Do not add the VM role to the `storage-int` network.

Bug 30936974

Adding Virtual Machine Role to the Management Network Causes Oracle VM Manager to Lose Contact with the Compute Nodes

Attempting to add the Virtual Machine role to the management network in Oracle VM Manger on an Oracle Private Cloud Appliance causes you to lose connectivity with your compute nodes. The compute nodes are still up, however your manager can not communicate with the compute nodes, which leaves your rack in a degraded state. This operation is not supported on Oracle Private Cloud Appliance.

Workaround: Do not add the VM role to the `mgmt-int` network.

Bug 30937049

Inadvertant Reboot of Stand-by Management Node During Upgrade Suspends Upgrade

When upgrading to Oracle Private Cloud Appliance Controller Software release 2.4.3 from either release 2.3.4 or 2.4.x releases you are required to upgrade the original stand-by management node first. Part of that upgrade is a reboot of this node which happens automatically during the upgrade process. After this reboot the original stand-by management node becomes the new active node. The next step is to upgrade the original active management node. However, if instead, you inadvertently reboot the original stand-by node again (the node that is now the new active) you will be unable to proceed with the upgrade because this will cause Oracle Private Cloud Appliance services on the new active node to fail.

Workaround: Reboot the original active node. This restarts the Oracle Private Cloud Appliance services on the new active node and you can proceed with upgrading the original active node.

Bug 30968544

Loading Incompatible Spine Switch Configuration Causes Storage Network Outage

When upgrading to Oracle Private Cloud Appliance Controller Software release 2.4.3 on an ethernet-based system do not attempt to make any manual changes to the spine switch configurations prior to the completion of the storage network upgrade. Doing so could cause the management nodes to lose access to the storage network. The management nodes may also get rebooted.

Additionally, once an upgrade to Controller Software release 2.4.3 is complete on an ethernet-based system, do not attempt to reload a spine switch backup from a prior software release. This could cause the management nodes to lose access to the storage network. The management nodes may also get rebooted. For example, you may see these error messages:

```
192.0.2.1 is unreachable
[root@ovcamn05r1 data]# ping 192.0.2.1
PING 192.0.2.1 (192.0.2.1) 56(84) bytes of data.
```

```
Mount points under shared storage are gone.
[root@ovcamn05r1 ~]# ls /nfs/shared_storage/
logs NO_STORAGE_MOUNTED
```

```
No master management node any more. o2cb service is offline. Both management
nodes are slave now.
[root@ovcamn05r1 ~]# pca-check-master
o2cb service is offline.
NODE: 192.0.2.2 MASTER: False
```

Workaround: Manually roll back the changes made on the spine switch configurations, then reboot both management nodes.

Bug 31407007

Cloud Backup Task Hangs When a ZFSSA Takeover is Performed During Backup

When the connection to the ZFS storage appliance is interrupted, the Oracle Cloud Infrastructure process will terminate the operation and mark it failed in the task database. In some cases, such as a management node reboot, there is no mechanism to update the state.

Workaround: When the task is unable to change state, delete the task from the task database, delete the `oci_backup` lock file, and institute a new backup operation. See "Cloud Backup" in [Monitoring and Managing Oracle Private Cloud Appliance](#) in the [Oracle Private Cloud Appliance Administration Guide](#).

Bug 31028898

Export VM to Oracle Cloud Infrastructure Job Shows as Aborted During MN Failover but it is Running in the Background

If there is an Export VM to Oracle Cloud Infrastructure job running when an active management node reboots or crashes, that job status changes to *Aborted* on Oracle VM Manager. In some cases, the export job will continue on the Exporter Appliance, despite the Abort message.

Workaround: Restart the Export VM to Oracle Cloud Infrastructure job. If the job is still running in the background, a pop up message shows *An export operation is already in progress for VM*. If the export job was aborted gracefully with the management node failover, then the export job is restarted.

Bug 31687516

Remove Deprecated `pca-admin diagnose software` Command

As of the Oracle Private Cloud Appliance Software Controller version 2.4.3 release, the `pca-admin diagnose software` command is no longer functional.

Workaround: Use the diagnostic functions now available through a separate health check tool. See "Health Monitoring" in [Monitoring and Managing Oracle Private Cloud Appliance](#) in the [Oracle Private Cloud Appliance Administration Guide](#) for more information.

Bug 31705580

Virtual Machine `get` Message Failed After 200 Seconds - Observed When `kube clusters` are Created Concurrently

When using the Oracle Private Cloud Appliance Cloud Native Environment release 1.2 OVA to create `kube clusters`, if you attempt to start multiple clusters at the same time, some clusters may fail with the following message:

```
Error_Code          VM_ERROR_004
Error_Message       Error (VM_ERROR_004): Virtual machine
autonas-cc3-master-3 get message failed after 200 seconds:
com.oracle.linux.Keepalived.master-addr,com.oracle.linux.k8s.error,com.oracle.
linux.k8s.script-result,com.oracle.linux.Keepalived.error.
```


Workaround: Stop the kube cluster that has failed, then restart that kube cluster.

Bug 32799556

Kube Cluster Creation/Deletion Should Not Be in Progress When Management Node Upgrade is Initiated

When upgrading management nodes from Software Controller release 2.4.3 to Software Controller release 2.4.4, do not initiate any kube cluster start or stop operations. As part of the upgrade procedure, a management node failover occurs. This failover can cause a kube cluster to go into a degraded state, if the kube cluster was attempting to start or stop at the time of the upgrade.

Workaround: Stop the kube cluster that has failed, then restart that kube cluster. These operations will clean up and recreate any VMs that were corrupted.

Bug 32880993

o2cb Service Status Reports "Registering O2CB cluster "ocfs2": Failed" State After Compute Node Provisioned

After compute nodes are provisioned, during the upgrade from Oracle Private Cloud Appliance release 2.4.3 to 2.4.4, you may encounter error messages with the o2cb service. When queried, the service is in the active state, but some clusters may show a failed state, as seen in the example below.

```
[root@ovcacn08r1 ~]# service o2cb status
Redirecting to /bin/systemctl status o2cb.service
• o2cb.service - Load o2cb Modules
   Loaded: loaded (/usr/lib/systemd/system/o2cb.service; enabled; vendor preset: disabled)
   Active: active (exited) since Thu 2021-04-22 09:00:51 UTC; 21h ago
   Main PID: 2407 (code=exited, status=0/SUCCESS)
   CGroup: /system.slice/o2cb.service

Apr 22 09:00:51 ovcacn08r1 o2cb.init[2407]: Loading stack plugin "o2cb": OK
Apr 22 09:00:51 ovcacn08r1 o2cb.init[2407]: Loading filesystem "ocfs2_dlmfs": OK
Apr 22 09:00:51 ovcacn08r1 o2cb.init[2407]: Creating directory '/dlm': OK
Apr 22 09:00:51 ovcacn08r1 o2cb.init[2407]: Mounting ocfs2_dlmfs filesystem at /dlm: OK
Apr 22 09:00:51 ovcacn08r1 o2cb.init[2407]: Setting cluster stack "o2cb": OK
Apr 22 09:00:51 ovcacn08r1 o2cb.init[2407]: Registering O2CB cluster "ocfs2": Failed
Apr 22 09:00:51 ovcacn08r1 o2cb.init[2407]: o2cb: Unknown cluster 'ocfs2'
Apr 22 09:00:51 ovcacn08r1 o2cb.init[2407]: Unregistering O2CB cluster "ocfs2": Failed
Apr 22 09:00:51 ovcacn08r1 o2cb.init[2407]: o2cb: Cluster 'ocfs2' is not active
Apr 22 09:00:51 ovcacn08r1 systemd[1]: Started Load o2cb Modules.
[root@ovcacn08r1 ~]#
```

Workaround: The failed messages are incorrectly reporting the status of the clusters, the clusters are functioning properly. It is safe to ignore these error messages. To clear the false messages, restart the o2cb service and check the status.

```
[root@ovcacn10r1 ~]# service o2cb restart
Redirecting to /bin/systemctl restart o2cb.service
[root@ovcacn10r1 ~]# service o2cb status
```

Bug 32667300

Compute Node Upgrade Restores Default Repository When Compute Node Was Previously Not Part of Any Tenant Group or Repository

When upgrading compute nodes from Software Controller release 2.4.3 to Software Controller release 2.4.4, if you have a compute node that is a part of the default tenant group but has no assigned repositories, the upgrade process restores the default repository to that compute node.

Workaround: If you wish to keep a compute node with no assigned repositories, and the upgrade process assigns the default repository to that compute node, simply unrepresent the repository from that compute node after upgrade.

Bug 32847571

Check of Local Repository to Ensure Target Compute Node is Empty

When upgrading compute nodes from Software Controller release 2.4.3 to Software Controller release 2.4.4, if the local repository for a compute node being upgraded has any ISO, VM Files, VM Templates, Virtual Appliances or Virtual Disks present, the upgrade precheck fails. This is expected behavior to ensure the data inside the local repository is retained before the upgrade occurs, which erases that data. If your compute node upgrade pre-checks fail, move all objects located in the compute node local repository to another repository, then retry the upgrade. If there is no need to retain any ISOs, VM Files, VM Templates, Virtual Appliances or Virtual Disks, delete them in order to make the local repository empty.

Workaround: Move items to another repository and retry the upgrade.

1. Log in to the Oracle VM Manager Web UI for the compute node you are upgrading.
2. Move each file type as described below:

Table 6-1 Moving Items Out of the Local Repository

Item	Steps
ISO	<ol style="list-style-type: none"> a. Clone the ISOs to other repositories. b. Delete the ISO files from the local repository.
VM Template	<ol style="list-style-type: none"> a. Move the template to another repository using clone customiser.
Virtual Appliances	<ol style="list-style-type: none"> a. Create VMs using each of the virtual appliances. b. Create virtual appliances from those VMs just create using "Export to Virtual Appliance" and point them to other repositories. c. Delete the virtual appliances (created in step 1) from the local repository.

Table 6-1 (Cont.) Moving Items Out of the Local Repository

Item	Steps
Virtual Disks	<ul style="list-style-type: none"> • If the VMs using these Virtual disks are in the local repository, migrate the corresponding VMs along with the Virtual disks (residing in local repository) to some other repository, using clone customiser. • If the VMs using these Virtual disks, are <i>not</i> in the local repository, (for example, few or all Virtual disks of some VMs reside in the local repository), follow these steps: <ol style="list-style-type: none"> a. Stop the VMs using those virtual disks. b. Clone the virtual disks with clone target as some other repository. This clone target repository should be presented to the compute node on which the VMs are hosted. c. Delete the actual virtual disks from the VMs. d. Attach the cloned virtual disks to their corresponding VMs. e. Start the VMs.
VM Files	<ol style="list-style-type: none"> a. Migrate the corresponding VMs to some other repository.

3. Run the `pca_upgrader` in verify mode to confirm the pre-checks pass.

If the pre-checks pass, run the upgrade.

```
[root@ovcamn05r1 ~]# pca_upgrader -V -t compute -c ovcacnXXr1
PCA Rack Type: PCA X8_BASE.
Please refer to log file /nfs/shared_storage/pca_upgrader/log/
pca_upgrader_<timestamp>.log for more details.
```

Beginning PCA Compute Node Pre-Upgrade Checks...

```
Check target Compute Node exists 1/8
Check the provisioning lock is not set 2/8
Check OVCA release on Management Nodes 3/8
Check Compute Node's Tenant matches Server Pool 4/8
Check target Compute Node has no local networks VNICs 5/8
Check target Compute Node has no VMs 6/8
Check local repository of target Compute Node is empty 7/8
Check no physical disks on target Compute Node have repositories 8/8
```

PCA Compute Node Pre-Upgrade Checks completed after 0 minutes

```
-----
PCA Compute Node Pre-Upgrade Checks Passed
-----
Check target Compute Node exists Passed
Check the provisioning lock is not set Passed
Check OVCA release on Management Nodes Passed
Check Compute Node's Tenant matches Server Pool Passed
Check target Compute Node has no local networks VNICs Passed
Check target Compute Node has no VMs Passed
Check local repository of target Compute Node is empty Passed
Check no physical disks on target Compute Node have repositories Passed
-----
Overall Status Passed
-----
```

```
PCA Compute Node Pre-Upgrade Checks Passed
Please refer to log file /nfs/shared_storage/pca_upgrader/log/
pca_upgrader_<timestamp>.log for more details.
```

```
[root@ovcamn05r1 ~]#
```

4. After you successfully perform the upgrade, restore the files you just backed up to the local repository on the newly-upgraded compute node (ovcacnXXr1-localfsrepo). You can use the table above to restore the items, or find detailed instructions in [Repositories Tab](#) in the [Oracle VM Manager User's Guide](#).

Bug 33093080

Check No Physical Disks on Target Compute Node Have Repositories

When upgrading compute nodes from Software Controller release 2.4.3 to Software Controller release 2.4.4, if there are repositories present on any Physical Disks (iSCSI/FC) and those Physical Disks (iSCSI/FC) are only presented to the compute node which is being upgraded, the precheck will fail.

Workaround: Release the ownership of the repository from the physical disk.



Note:

Check all physical disks that are *only* presented to the compute node being upgraded for repositories. You must perform this procedure for each repository that is present on each of these physical disks.

Pre-Upgrade Steps

1. Log in to the Oracle VM Manager Web UI.
2. In the Servers and VMs tab, select the appropriate server pool and validate that the compute node is part of that server pool.
3. From the Repositories tab, select the repository and note the physical disk over which the repository lies.
4. From the Repositories tab, select the repository, then edit the concerned repository and check Release Ownership.
5. From Repository tab, click Show All Repositories, then select the repository and delete it.

This only deletes the repository from Oracle VM Manager and not the actual filesystem on the physical disk.

Retry the Compute Node Upgrade

1. Run the `pca_upgrader` in verify mode to confirm the pre-checks pass.

```
[root@ovcamn05r1 ~]# pca_upgrader -V -t compute -c ovcacnXXr1
PCA Rack Type: PCA X8_BASE.
Please refer to log file /nfs/shared_storage/pca_upgrader/log/
pca_upgrader_<timestamp>.log for more details.
```

```
Beginning PCA Compute Node Pre-Upgrade Checks...
```

```
Check target Compute Node exists
```

```
1/8
```

```

Check the provisioning lock is not set                2/8
Check OVCA release on Management Nodes              3/8
Check Compute Node's Tenant matches Server Pool     4/8
Check target Compute Node has no local networks VNICs 5/8
Check target Compute Node has no VMs                6/8
Check local repository of target Compute Node is empty 7/8
Check no physical disks on target Compute Node have repositories 8/8

```

PCA Compute Node Pre-Upgrade Checks completed after 0 minutes

```

-----
PCA Compute Node Pre-Upgrade Checks                Passed
-----

```

```

Check target Compute Node exists                    Passed
Check the provisioning lock is not set              Passed
Check OVCA release on Management Nodes             Passed
Check Compute Node's Tenant matches Server Pool    Passed
Check target Compute Node has no local networks VNICs Passed
Check target Compute Node has no VMs               Passed
Check local repository of target Compute Node is empty Passed
Check no physical disks on target Compute Node have repositories Passed
-----

```

```

Overall Status                                     Passed
-----

```

```

PCA Compute Node Pre-Upgrade Checks                Passed

```

```

Please refer to log file /nfs/shared_storage/pca_upgrader/log/
pca_upgrader_<timestamp>.log for more details.

```

```
[root@ovcamn05r1 ~]#
```

2. If the pre-checks pass, run the upgrade.

```
[root@ovcamn05r1 ~]# pca_upgrader -U -t compute -c ovcacnXXr1
```

Post Upgrade Steps to Restore the Repository

1. In the Storage tab, click on the SAN Server which hold the physical disks and refresh the physical disk (which held the repository before the upgrade).
2. In the Storage tab, select Shared File System/Local File System for the corresponding file system for the physical disk on which you had the repository, then click the refresh button.
3. In the Repository tab, click Show All Repositories, then confirm the repository (which was deleted earlier in this procedure) is restored.
4. From Repository tab, click Show All Repositories, then edit the repository that was deleted pre-upgrade. Click on take ownership and select the same server pool it was associated with prior to the upgrade.
5. Select the repository and click Refresh Selected Repository.

Bug 33093068

7

Feedback and Support

This chapter contains information about how to provide feedback and contact support for the Oracle Private Cloud Appliance (PCA) product.



Note:

Customers who use Oracle Auto Service Request (ASR) do not need to submit a service request through My Oracle Support for the specific hardware failure scenarios covered by ASR. For details, refer to "Oracle Auto Service Request (ASR)" in [Servicing Oracle Private Cloud Appliance Components](#) in the [Oracle Private Cloud Appliance Administration Guide](#).



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.

Providing Feedback and Reporting Problems

If you need to report an issue and have an Oracle Premier Support Agreement, you should open a case with Oracle Support at [My Oracle Support](#).

If you are reporting an issue, please provide the following information where applicable:

- Description of the problem, including the situation where the problem occurs, and its impact on your operation.
- Machine type, operating system release, browser type and version, locale and product release, including any patches you have applied, and other software that might be affecting the problem.
- Detailed steps on the method you have used, to reproduce the problem.
- Any error logs or core dumps.

Contacting Oracle Specialist Support

If you have an Oracle Customer Support Identifier (CSI), first try to resolve your issue by using [My Oracle Support](#). Your Oracle Premier Support CSI does not cover customization support, third-party software support, or third-party hardware support.

If you cannot resolve your issue, open a case with the Oracle specialist support team for technical assistance on break/fix production issues. The responding support engineer will need the following information to get started:

- Your Oracle Customer Support Identifier.
- The product you are calling about.
- A brief description of the problem you would like assistance with.
- Any logs or support data you have, see [Feedback and Support](#) for details.

If your CSI is unknown, find the correct Service Center for your country (<https://www.oracle.com/support/contact.html>), then contact Oracle Services to open a non-technical service request (SR) to get your CSI sorted. Once you have your CSI, you can proceed to open your case through My Oracle Support.

Data Collection for Service and Support

When submitting a Service Request (SR), please include an archive file with the relevant log files and debugging information as listed in this section. This information can be used by Oracle Support to analyze and diagnose system issues. The support data files can be uploaded for further analysis by Oracle Support.

Tip:

Collecting support files involves logging in to the command line on components in your Oracle Private Cloud Appliance rack and copying files to a storage location external to the appliance environment, in the data center network. This can only be achieved from a system with access to both the internal appliance management network and the data center network. You can set up a physical or virtual system with those connections, or use the active management node.

The most convenient way to collect the necessary files, is to mount the target storage location on the system using `nfs`, and copy the files using `scp` with the appropriate login credentials and file path. The command syntax should be similar to this example:

```
# mkdir /mnt/mynfsshare
# mount -t nfs storage-host-ip:/path-to-share /mnt/mynfsshare
# scp root@component-ip:/path-to-file /mnt/mynfsshare/pca-support-data/
```

Collecting Oracle Private Cloud Appliance Support Data

NOT_SUPPORTED:

For more accurate diagnosis of physical server issues, Oracle Support Services require a system memory dump. For this purpose, `kdump` must be installed and configured on the component under investigation. The entire procedure is described in the support note with [Doc ID 2142488.1](#). By default, `kdump` is installed on all Oracle Private Cloud Appliance compute nodes and configured to write the system memory dump to the ZFS storage appliance at this location: `192.168.4.100:/export/nfs_repository1/`.

Caution:

For diagnostic data collection, Oracle Support Services recommend that the OSWatcher tool be run for an extended period of time. For general information about the use of OSWatcher, please consult the support note with [Doc ID 580513.1](#). OSWatcher is installed by default on all Oracle Private Cloud Appliance compute nodes.

For diagnostic purposes, Oracle Support Services use a tool called `pca-diag` that automatically collects vital troubleshooting information from your Oracle Private Cloud Appliance environment. This tool is part of the Oracle Private Cloud Appliance controller software installed on both management nodes and on all compute nodes. Its capabilities are described in "Oracle Private Cloud Appliance Diagnostics Tool" in "Software Components" in [Concept, Architecture and Life Cycle of Oracle Private Cloud Appliance](#) in the [Oracle Private Cloud Appliance Administration Guide](#).

To collect support data from your system, proceed as follows:

1. Log in to **the active management node** as root.
2. Run `pca-diag` with the appropriate command line arguments.

Note:

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 the most complete set of diagnostic data, run the command with both arguments: `pca-diag ilom vmpinfo`.

- `pca-diag ilom`

Use this command to detect and diagnose potential component hardware and software problems.

```
[root@ovcamn05r1 ~]# pca-diag ilom
```

```
Oracle Private Cloud Appliance diagnostics tool
```

```
Gathering Linux information...
```



```
Gathering system messages...
Gathering PCA related files...
Gathering OS version information...
Gathering host specific information...
Gathering PCI information...
Gathering SCSI and partition data...
Gathering OS process data...
Gathering network setup information...
Gathering installed packages data...
Gathering disk information...
Gathering ILOM Service Processor data... this may take a while
Generating diagnostics tarball and removing temp directory
```

```
=====
=====
Diagnostics completed. The collected data is available in:
/tmp/pcadiag_ovcamn05r1_<ID>_<date>_<time>.tar.bz2
=====
=====
```

- `pca-diag vmpinfo`

Use this command to detect and diagnose potential problems in the Oracle VM environment.

 **Note:**

To collect diagnostic information for a subset of the Oracle VM Servers in the environment, run the command with an additional `servers` parameter, as shown in the example below.

```
[root@ovcamn05r1 ~]# pca-diag vmpinfo servers='ovcacn07r1,ovcacn08r1'

Oracle Private Cloud Appliance diagnostics tool
Gathering Linux information...
Gathering system messages...
Gathering PCA related files...
Gathering OS version information...
Gathering host specific information...
Gathering PCI information...
Gathering SCSI and partition data...
Gathering OS process data...
Gathering network setup information...
Gathering installed packages data...
Gathering disk information...
Gathering FRU data and console history. Use ilom option for complete
ILOM data.
```

When the `vmpinfo3` script is called as a sub-process from `pca-diag`, the console output continues as follows:

```
Running vmpinfo tool...
Starting data collection
Gathering files from servers: ovcacn07r1,ovcacn08r1 This process may
take some time.
Gathering OVM Model Dump files
Gathering sosreports from servers
The following server(s) will get info collected: [ovcacn07r1,ovcacn08r1]
```

```
Gathering sosreport from ovcacn07r1
Gathering sosreport from ovcacn08r1
Data collection complete
Gathering OVM Manager Logs
Clean up metrics
Copying model files
Copying DB backup log files
Invoking manager sosreport
```

When all files have been collected, the data is compressed into two tarballs. One is from the `pca-diag` tool, while `vmpinfo3` writes a separate tarball with its own specific data.

```
Compressing VMPinfo3 <date>-<time>.
=====
Please send /tmp/vmpinfo3-<version>-<date>-<time>.tar.gz to Oracle OVM support
=====

Generating diagnostics tarball and removing temp directory
=====
Diagnostics completed. The collected data is available in:
/tmp/pcadiag_ovcamn05r1_<ID>_<date>_<time>.tar.bz2
=====
```

3. If necessary, run `pca-diag`, with or without the `ilom` argument, on some or all compute nodes as well.
4. To allow better analysis of physical server issues, for example hanging, crashing or rebooting, also include the system memory dump file (`vmcore`). Refer to the *Tip* at the start of this section for a convenient way to collect the files.

The location of the file is: `<kdump-partition-mount-point>/kdump/<compute-node-ip>-<date>-<time>/vmcore`. The partition and mount point are defined during `kdump` configuration. By default, `kdump` writes to `192.168.4.100:/export/nfs_repository1/`.

For details, please consult the support note with [Doc ID 2142488.1](#).

5. When required, collect the OSWatcher logs from the compute nodes. The default location is `/var/log/oswatcher/archive/`.
6. Copy all diagnostic files to a location external to the appliance environment.

Uploading Support Data Files

For support data up to 2 GB, upload the file as part of the Service Request (SR) process in My Oracle Support (MOS).

- If you are still in the process of logging the SR, upload the support data in the **Upload Files/Attachments** step of the SR.
- If you have already logged the SR and need to upload files afterwards, proceed as follows:
 1. Log into MOS and open the **Dashboard** or **Service Request** tab.
 2. In the **Service Request** region, click the SR you want to update.
 3. In the **Update** section, select **Add Attachment**.
 4. In the pop-up window, select the file for upload, include any notes, and click **Attach File**.

If uploading the support data with the SR is not an option, or for support data files over 2 GB in size, use the FTPS file upload service from Oracle support at `transport.oracle.com`. Oracle Support might request that you upload using a different mechanism.

1. Using an FTPS client, for example FileZilla or WinSCP, access the My Oracle Support File Upload Service `transport.oracle.com` in *passive* mode.
2. Log in with your Oracle Single Sign-On user name and password.
3. Select the support data file to upload.
4. Select a destination for the file.

Use the directory path provided by Oracle Support.

Typically, the directory path is constructed as follows: `"/upload/issue/<sr_number>/"`.

The use of a SR number ensures that the file is correctly associated with the service request. Write down the full path to the file and the SR number for future reference in communications with Oracle Support.

5. Upload the file.

When the upload is complete, a confirmation message is displayed.

 **Note:**

If you prefer to use a command-line client, for example cURL, you typically enter a single command to connect, authenticate and complete the upload. A cURL command will look similar to this example: `curl -T <path_to_file> -u "<user>" ftps://transport.oracle.com/upload/issue/<sr_number>/.
For security reasons, it is recommended that you leave the password out of the command and be prompted to enter the password.`

For detailed information about uploading files to Oracle Support, refer to the support note with [Doc ID 1547088.2](#).