

Oracle Private Cloud Appliance

Release Notes for Release 2.4.4.6



G43606-01
September 2025



Oracle Private Cloud Appliance Release Notes for Release 2.4.4.6,
G43606-01

Copyright © 2025, Oracle and/or its affiliates.

Contents

Preface

Audience	ii
Related Documentation	ii
Feedback	iii
Conventions	iii
Documentation Accessibility	iii
Access to Oracle Support for Accessibility	iii
Diversity and Inclusion	iii

1 About Oracle Private Cloud Appliance

2 Feature Overview

Changes and Improvements in Release 2.4.4.6	1
Changes and Improvements in Release 2.4.4.5	2
Changes and Improvements in Release 2.4.4.4	4
Changes and Improvements in Release 2.4.4.3	5
Changes and Improvements in Release 2.4.4.2	6
Changes and Improvements in Release 2.4.4.1	7
Changes and Improvements in Release 2.4.4	9
Changes and Improvements in Release 2.4.3	10
Earlier Releases	11

3 Configuration Maximums

Ethernet-Based System Configuration Maximums	1
--	---

4 Firmware Qualification

Qualified Firmware for Oracle Private Cloud Appliance Controller Software Release 2.4.4.6	1
---	---

5 Known Limitations and Workarounds

Oracle Private Cloud Appliance Hardware	1
Cisco Firmware Configuration Change	1
Compute Node Boot Sequence Interrupted by LSI Bios Battery Error	1
Reboot From Oracle Linux Prompt May Cause Management Node to Hang	1
Oracle ZFS Storage Appliance More Aggressively Fails Slow Disks	2
Emulex Fibre Channel HBAs Discover Maximum 128 LUNs	2
Fibre Channel LUN Path Discovery Is Disrupted by Other Oracle VM Operations	2
Poor Oracle VM Performance During Configuration of Fibre Channel LUNs	3
ILOM Firmware Does Not Allow Loopback SSH Access	3
incorrect opcode Messages in the Console Log	4
Megaraid Firmware Crash Dump Is Not Available	4
Some Services Require an Upgrade of Hardware Management Pack	4
Compute Nodes Containing Emulex HBA Card With Maximum FC Paths Reboots With Errors in Oracle VM Manager UI	5
Compute Nodes Containing FC HBA with Maximum FC Paths in Dead State After Reprovisioning	6
Compute Node FC HBA (QLogic/Emulex) with FC LUNs Having Path Flapping	7
Upgrade Compute Node Fails with Fibre Channel LUNs	8
PCA Faultmonitor Check firewall_monitor Fails Due to nc:command not found	8
Certain TZ Configuration is Failing on Cisco Switches	8
NFS Shares on Internal ZFS Will Fail After ZFS Firmware Update If vnuc Owned by SN02	8
Repository Size is Not Reflecting Properly in OVMM GUI	9
Oracle Private Cloud Appliance Software	9
Upgrade to UEK6 Kernel Fails to Boot with 1000 FC LUNs Attached	9
Fault Reports Are Not Removed Until Another Fault Is Reported	10
Do Not Install Additional Software on Appliance Components	10
Upgrader UI Installation Step Is Skipped	10
Deleting a Storage Network Is Allowed When the Network Is Assigned to NFS or iSCSI Storage	10
Node Manager Does Not Show Node Offline Status	11
Compute Node State Changes Despite Active Provisioning Lock	11
Compute Nodes Are Available in Oracle VM Server Pool Before Provisioning Completes	11
Virtual Machines Remain in Running Status when Host Compute Node Is Reprovisioned	11
Ethernet-Based System Management Nodes Have Non-Functional bond0 Network Interface	12
Network Performance Is Impacted by VxLAN Encapsulation	12
Altering Custom Network VLAN Tag Is Not Supported	12
Configuring Uplinks with Breakout Ports Results in Port Group Named 'None'	13
DPM Server Pool Policy Interrupts Synchronization of Tenant Group Settings	14
Host Network Parameter Validation Is Too Permissive	14

Virtual Appliances Cannot Be Imported Over a Host Network	14
Customizations for ZFS Storage Appliance in multipath.conf Are Not Supported	14
Customer Created LUNs Are Mapped to the Wrong Initiator Group	15
Storage Head Failover Disrupts Running Virtual Machines	15
Changing Multiple Component Passwords Causes Authentication Failure in Oracle VM Manager	15
ILOM Password of Expansion Compute Nodes Is Not Synchronized During Provisioning	16
SSH Host Key Mismatch After Management Node Failover	16
External Storage Cannot Be Discovered Over Data Center Network	16
Mozilla Firefox Cannot Establish Secure Connection with User Interface	16
Virtual Machine with High Availability Takes Five Minutes to Restart when Failover Occurs	17
CLI Command update appliance Is Deprecated	17
Certain CLI Commands Fail in Single-command Mode	18
Upgrader Checks Logged in Different Order	18
Virtual Machine Loses IP Address Due to DHCP Timeout During High Network Load	18
Adding the Virtual Machine Role to the Storage Network Causes Cluster to Lose Heartbeat Networking	18
Adding Virtual Machine Role to the Management Network Causes Oracle VM Manager to Lose Contact with the Compute Nodes	19
Inadvertant Reboot of Stand-by Management Node During Upgrade Suspends Upgrade	19
Cloud Backup Task Hangs When a ZFSSA Takeover is Performed During Backup	19
Export VM to Oracle Cloud Infrastructure Job Shows as Aborted During MN Failover but it is Running in the Background	19
Remove Deprecated pca-admin diagnose software Command	20
Virtual Machine get Message Failed After 200 Seconds - Observed When kube clusters are Created Concurrently	20
Kube Cluster Creation/Deletion Should Not Be in Progress When Management Node Upgrade is Initiated	20
o2cb Service Status Reports "Registering O2CB cluster "ocfs2": Failed" State After Compute Node Provisioned	21
Compute Node Upgrade Restores Default Repository When Compute Node Was Previously Not Part of Any Tenant Group or Repository	21
Check of Local Repository to Ensure Target Compute Node is Empty	22
Check No Physical Disks on Target Compute Node Have Repositories	24
Backup of Config Can Fill Filesystem and Cause Numerous Problems	26
Two Different Release Packages of Same Tool Present in the ISO	26
Intermittent Error When Deleting the Uplink Port Group	26

6 Feedback and Support

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:

- [Oracle Private Cloud Appliance](#)
- [Oracle Private Cloud Appliance Release 2.x](#)
- [Oracle Private Cloud Appliance Release 2.4.x](#)

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 this release of Oracle Private Cloud Appliance. The release notes also provide the link on My Oracle Support to the upgrade guide for the release.

Oracle Private Cloud Appliance Upgrade Guide

The upgrade guide provides specific instructions for how to upgrade an Oracle Private Cloud Appliance to a new release by using the Upgrader Tool.

Newer Oracle Private Cloud Appliance upgrade information is on [My Oracle Support](#). See the release notes for the release for specific links.

Upgrade instructions for Release 2.4.4.1 are in the Oracle Private Cloud Appliance Upgrade Guide for Release 2.4.4.1.

Prior to Oracle Private Cloud Appliance Release 2.4.4.1, upgrade instructions are in a chapter in the Oracle Private Cloud Appliance Administrator's Guide for the release.

Oracle Private Cloud Appliance Patching Guide

The patching guide describes how to patch a Private Cloud Appliance using dedicated channels on the Unbreakable Linux Network (ULN). Patching provides security updates, kernel changes, and bug fixes between larger ISO releases, for both management and compute nodes.

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 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 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 UNIX systems, 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 are included with Oracle Private Cloud Appliance depending on configuration. Following are references to documentation for common additional components:

- Oracle Rack Cabinet 1242
<https://docs.oracle.com/en/servers/options/rack-cabinet-1242/index.html>
- Oracle Server X86 Servers
<https://docs.oracle.com/en/servers/index.html>
- Oracle ZFS Storage Appliance ZS7-2
<https://docs.oracle.com/en/storage/zfs-storage/zfs-appliance/os8-8-x/>
- Oracle Integrated Lights Out Manager (ILOM) 4.0.x
https://docs.oracle.com/cd/E81115_01/index.html
- Oracle Integrated Lights Out Manager (ILOM) 5.0
<https://docs.oracle.com/en/servers/management/ilom/index.html>
- Oracle VM
<https://docs.oracle.com/en/virtualization/oracle-vm/index.html>

- Oracle VM Manager Release 3.4.7
[Oracle VM : What's New with Release 3.4.7 - Latest Information \(Doc ID 2832974.1\)](#)
- Oracle VM Manager/Oracle Linux Server/Kernel versions
[\[PCA 2.x\] OVMM/OVS/Kernel Versions Matrix \(Doc ID 2325981.1\)](#)
- Oracle Enterprise Manager Plug-in
<https://docs.oracle.com/en/enterprise-manager/cloud-control/enterprise-manager-cloud-control/13.5/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

New versions of the Oracle Private Cloud Appliance Controller Software support earlier configurations of the Ethernet-based hardware platform. 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*.

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, two 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 operating system images, other applications, and third-party hardware that you supply.

2

Feature Overview

This chapter provides an overview of the new features and other changes in different releases of Oracle Private Cloud Appliance.

Changes and Improvements in Release 2.4.4.6

This section describes functional changes, improvements, and bug fixes in the 2.4.4.6 release.

Upgrading to Oracle Private Cloud Appliance Release 2.4.4.6

For specific instructions for upgrading to Oracle Private Cloud Appliance Release 2.4.4.6, see the [My Oracle Support](#) document [\[PCA 2.4.4.6\] Upgrade Guide \(Doc ID 3105091.1\)](#).

You can upgrade directly to Oracle Private Cloud Appliance Release 2.4.4.6 from Release 2.4.4.5. For earlier releases, you must first upgrade to Release 2.4.4.5, then complete the upgrade from Release 2.4.4.5 to Release 2.4.4.6. See the Upgrade Guide referenced in the preceding paragraph for prerequisites. See also [Additional Firmware Notes](#) in [Qualified Firmware for Oracle Private Cloud Appliance Controller Software Release 2.4.4.6](#).

To review the general upgrade procedure and command output see [Oracle Private Cloud Appliance Upgrade Guide for Release 2.4.4.1](#).

Upgraded Oracle VM Manager

Oracle VM Manager is upgraded to Release 3.4.7.7 with July 2025 CPUs.

For information about enhancements in Oracle VM Manager Release 3.4.7.7, see [Oracle VM : What's New with Release 3.4.7 - Latest Information \(Doc ID 2832974.1\)](#).

Upgraded Oracle ZFS Storage Appliance Software

The ZFSSA firmware must be updated to OS8.8.83 to apply deferred updates.

Oracle Private Cloud Appliance Release 2.4.4.6 includes the following version of the Oracle ZFS Storage Appliance software:

- Patch 37777352: ORACLE ZFS STORAGE APPLIANCE SOFTWARE 2013.1.8.83 (OS8.8.83) OS 8.8.83 2013.06.05.8.83,1-1.1

This version allows for deferred updates. The ZFSSA firmware must be updated to this version to apply deferred updates.

See the [My Oracle Support](#) document [\[PCA 2.4.4.6\] Upgrade Guide \(Doc ID 3105091.1\)](#) for details about how to upgrade the ZFS Storage Appliance.

Bugs Fixed in Release 2.4.4.6

The following table lists notable bugs that have been fixed in Oracle Private Cloud Appliance Release 2.4.4.6.

Table 2-1 List of Fixed Bugs

Bug ID	Description
36050330	PCA 2.4.4.2 RPM lsof-4.87-6.el7.x86_64 not installed on CNs
29913246	Consider removing the obsolete update appliance command
34732503	grub2-pc-2.02-0.87.0.21.el7_9.9.x86_64.rpm is not installing fully.
34906041	pca_upgrader ui_install step is skipped due to .standby_upgrade_complete check
38376055	Linux kernel update needed
38376343	sqlite update needed
38376150	libxml2 update needed
38376185	python-setuptools update needed
38376126	glibc update needed
38299187	MySQL crashed during upgrade
38275208	Pickup Latest timezone tzdata rpm for MN's and CN's
38316382	pca_upgrader failed with phonehome missing parameters
37873062	PCA Upgrade does not maintain enable_phonehome configuration
34263370	PCA Daily backup needs include all OVMMModelExport xml files
36925686	Cisco Config is not exactly same after fw upgrade
31620556	osimage.ini out-of-date references to no longer used kickstart files.
31715624	Failed PCA oci-backup-restore operation could still create the destination restore share in ZS
32649461	CSP Header issues in ovm console
34139051	creating cloud target with already in use cloud bucket should fail with proper error message
34980329	Need Precheck to Verify Shared Storage doesn't have Ownership of nobody:nobody
35154323	[PCA X8 2.4.4.x] Change MN Shared mount check to be less restrict of NFS version used
34767928	pca-cli allowing to delete storage-network when it is assigned to nfs/iscsi storage
35924931	PCA 2.4.4.2 post upgrade pca_healthcheck intermittent failures

Changes and Improvements in Release 2.4.4.5

This section describes functional changes, improvements, and bug fixes in the 2.4.4.5 release.

Upgrading to Oracle Private Cloud Appliance Release 2.4.4.5

For specific instructions for upgrading to Oracle Private Cloud Appliance Release 2.4.4.5, see the [My Oracle Support](#) document [\[PCA 2.4.4.5\] Upgrade Guide \(Doc ID 3081129.1\)](#).

You can upgrade directly to Oracle Private Cloud Appliance Release 2.4.4.5 from any of the following releases. See the Upgrade Guide referenced in the preceding paragraph for prerequisites. See also [Additional Firmware Notes](#) in [Qualified Firmware for Oracle Private Cloud Appliance Controller Software Release 2.4.4.6](#).

- Private Cloud Appliance Release 2.4.4.4
- Private Cloud Appliance Release 2.4.4.3

- Private Cloud Appliance Release 2.4.4.2
- Private Cloud Appliance Release 2.4.4.1
- Private Cloud Appliance Release 2.4.4

To review the general upgrade procedure and command output see [Oracle Private Cloud Appliance Upgrade Guide for Release 2.4.4.1](#).

Upgraded Oracle VM Manager

Oracle VM Manager is upgraded to Release 3.4.7.6 with January 2025 CPUs.

For information about enhancements in Oracle VM Manager Release 3.4.7.6, see [Oracle VM : What's New with Release 3.4.7 - Latest Information \(Doc ID 2832974.1\)](#).

Upgraded Oracle ZFS Storage Appliance Software

The ZFSSA firmware must be updated to OS8.8.75 to apply deferred updates.

① Note

While running OS8.8.75, you must do the following to work around bug 37294624: You must perform a manual check any time a ZFS head becomes the active head (AKCS_OWNER).

After the deferred updates are applied for OS8.8.75, immediately upgrade to OS8.8.75 IDR 5746.1, which includes the fix for bug 37294624.

Oracle Private Cloud Appliance Release 2.4.4.5 includes the following two versions of Oracle ZFS Storage Appliance software:

- OS8.8.75, 2013.06.05.8.75,1-1.1

This version allows for deferred updates. The ZFSSA firmware must be updated to this version to apply deferred updates.

This version also includes the firmware update for Bear Cove Disk Drives. For more information see [\[PCA 2.x\] Out of Cycle Drive Firmware Upgrade for Bear Cove Disk Drives \(Doc ID 3053539.1\)](#)

- OS8.8.75+IDR 5746.1, 2013.06.05.8.75,1-2.75.5746.1

This version includes the fix for bug 37294624 - FTP access issues due to race in nas_cache. After the deferred updates are applied for OS8.8.75, immediately upgrade to OS8.8.75 IDR 5746.1 to get the fix for bug 37294624.

See the [My Oracle Support](#) document [\[PCA 2.4.4.5\] Upgrade Guide \(Doc ID 3081129.1\)](#) for details about how to upgrade the ZFS Storage Appliance.

Oracle ZFS Storage Appliance software OS8.8.75 supports 22 TB drives.

Bugs Fixed in Release 2.4.4.5

The following table lists notable bugs that have been fixed in Oracle Private Cloud Appliance Release 2.4.4.5.

Table 2-2 List of Fixed Bugs

Bug ID	Description
37727858	Pick up Latest timezone tzdata rpm for MNs and CNs
36975883	2.4.4.5: man-db rpm is not installed on 2.4.4.3->2.4.4.5 upgrade

Changes and Improvements in Release 2.4.4.4

This section describes functional changes, improvements, and bug fixes in the 2.4.4.4 release.

Upgrading to Oracle Private Cloud Appliance Release 2.4.4.4

For specific instructions for upgrading to Oracle Private Cloud Appliance Release 2.4.4.4, see the [My Oracle Support](#) document [\[PCA 2.4.4.4\] Upgrade Guide \(Doc ID 3047453.1\)](#).

You can upgrade directly to Oracle Private Cloud Appliance Release 2.4.4.4 from any of the following releases. See the Upgrade Guide referenced in the preceding paragraph for prerequisites.

- Private Cloud Appliance Release 2.4.4.3
- Private Cloud Appliance Release 2.4.4.2
- Private Cloud Appliance Release 2.4.4.1
- Private Cloud Appliance Release 2.4.4

To review the general upgrade procedure and command output see [Oracle Private Cloud Appliance Upgrade Guide for Release 2.4.4.1](#).

Configuring ULN Mirror Servers

The procedure to set up channels for Private Cloud Appliance Release 2.4.3.x and Private Cloud Appliance Release 2.4.4.x on a ULN mirror server is changed for mirror servers that are running Oracle Linux 8 or Oracle Linux 9.

For instructions, see [Configuring a Mirror Server for Oracle Linux 8 or Oracle Linux 9](#) in *Oracle Private Cloud Appliance Patching Guide for Release 2.4.3.x and Release 2.4.4.x*.

The procedure to configure an Oracle Linux 7 mirror server has not changed.

Upgraded Oracle VM Manager

Oracle VM Manager is upgraded to Release 3.4.7.5 with July 2024 CPUs.

For information about enhancements in Oracle VM Manager Release 3.4.7.5, see [Oracle VM : What's New with Release 3.4.7 - Latest Information \(Doc ID 2832974.1\)](#).

Bugs Fixed in Release 2.4.4.4

The following table lists notable bugs that have been fixed in Oracle Private Cloud Appliance Release 2.4.4.4.

Table 2-3 List of Fixed Bugs

Bug ID	Description
36973470	man-db-2.6.3-11.el7.x86_64 is not installed on 2.4.4.3->2.4.4.4 upgrade

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

Bug ID	Description
36925852	Faultmonitor failing to add details to text and html reports
36493643	Platinum Certified Monitor didn't report wrong version of ZFS on first head
36400073	Reprovision button missing on the Rack view page of Dashboard UI for Dead nodes
36184188	Command "/usr/sbin/pca-admin create uln-repo mirror http://yum-mirror-server/yum" fails
35984246	Cisco Switch vpc role needs auto-recovery enabled
35781596	missing man executable, man-db-2.6.3-11.el7.x86_64 is not installed on management node
35665992	Fault Monitor doesn't clear reports after set number of days unless new ones are created
36611922	[OVM-3.4.7] PCA 2.4.4 Xen Destroys Virtual Machine When it Resumes After Live Migrate
33622799	VIRTUAL FUNCTION IS NOT ASSIGNED TO DOMU SOMETIMES IN XEN OL7

Changes and Improvements in Release 2.4.4.3

This section describes functional changes, improvements, and bug fixes in the 2.4.4.3 release.

Upgrading to Oracle Private Cloud Appliance Release 2.4.4.3

For specific instructions for upgrading to Oracle Private Cloud Appliance Release 2.4.4.3, see the [My Oracle Support](#) document [\[PCA 2.4.4.3\] Upgrade Guide \(Doc ID 2993669.1\)](#).

You can upgrade from Private Cloud Appliance Release 2.4.4, Private Cloud Appliance Release 2.4.4.1, or Private Cloud Appliance Release 2.4.4.2.

To review the general upgrade procedure and command output see [Oracle Private Cloud Appliance Upgrade Guide for Release 2.4.4.1](#).

ULN Patching Provides Updates Between Full Releases

In addition to upgrading a Private Cloud Appliance to a new software release using the Upgrader Tool, you can patch a Private Cloud Appliance using dedicated channels on the Unbreakable Linux Network (ULN). Patching provides security updates, kernel changes, and bug fixes between larger ISO releases, for both management and compute nodes. See [Oracle Private Cloud Appliance Patching Guide for Release 2.4.3.x and Release 2.4.4.x](#).

Upgraded Kernel

The kernel that is installed on the management nodes and compute nodes is updated from UEK5 to UEK6.

Upgraded Oracle VM Manager

Oracle VM Manager is upgraded to Release 3.4.7.4 with October 2023 CPUs.

For information about enhancements in Oracle VM Manager Release 3.4.7.4, see [Oracle VM : What's New with Release 3.4.7 - Latest Information \(Doc ID 2832974.1\)](#).

Bugs Fixed in Release 2.4.4.3

The following table lists notable bugs that have been fixed in Oracle Private Cloud Appliance Release 2.4.4.3.

Table 2-4 List of Fixed Bugs

Bug ID	Description
34152575	Win2019 Cluster(WSFC) Issue on PCA 2.4.4 CNs
35600633	Management node ILOM lockout after spMn-root password revert failed post upgrade to ILOM 5.1.1.21 r150524
35209865	[X8-2] Healthchecker should pick up luns mapped to "All"
35574512	To Pickup Latest "tzdata (2023)" rpm for all MN's and CN's
34946563	pca-vm-tool is present and throws a traceback on execution
35840166	spanning-tree bpdudfilter enable' for uplinks on spine switches failed during MN upgrade
35560535	BVT : / mounted options are different
35103017	pca-faultmonitor fails for storage_path_flapping_monitor
34837709	PCA X8-2 version 2.4.4.1 is loosed network redundancy capability

Changes and Improvements in Release 2.4.4.2

This section describes functional changes, improvements, and bug fixes in the 2.4.4.2 release.

Upgrading to Oracle Private Cloud Appliance Release 2.4.4.2

For specific instructions for upgrading to Oracle Private Cloud Appliance Release 2.4.4.2, see the [My Oracle Support](#) document [\[PCA 2.4.4.2\] Upgrade Guide \(Doc ID 2914968.1\)](#).

You can upgrade from either Private Cloud Appliance Release 2.4.4 or Private Cloud Appliance Release 2.4.4.1.

To review the general upgrade procedure and command output see [Oracle Private Cloud Appliance Upgrade Guide for Release 2.4.4.1](#).

ULN Patching Provides Updates Between Full Releases

In addition to upgrading a Private Cloud Appliance to a new software release using the Upgrader Tool, you can patch a Private Cloud Appliance using dedicated channels on the Unbreakable Linux Network (ULN). Patching provides security updates, kernel changes, and bug fixes between larger ISO releases, for both management and compute nodes. See [Oracle Private Cloud Appliance Patching Guide for Release 2.4.3.x and Release 2.4.4.x](#).

Upgraded Oracle VM Manager

Oracle VM Manager is upgraded to Release 3.4.7.2 with October 2022 CPUs.

For information about enhancements in Oracle VM Manager Release 3.4.7.2, see [Oracle VM : What's New with Release 3.4.7 - Latest Information \(Doc ID 2832974.1\)](#).

pca-diag Enhancements

The pca-diag diagnostic tool is enhanced in the following ways:

- The `vmppinfo` command collects a single tarball of logs and configuration data from the Oracle VM Manager. The `ocvdiag` tarball is merged into the single tarball.
- The `vmppinfo` command collects logs and `sosreport` support data both management nodes and from each Oracle VM Server or compute node.
- Collects data for Cisco Switches.
- Collects ILOM data for each management and compute node.
- Collects ULN configuration files and yum transaction history.

For more information, see the "Oracle Private Cloud Appliance Diagnostics Tool" software component in [Concept, Architecture and Life Cycle of Oracle Private Cloud Appliance](#) in the Oracle Private Cloud Appliance Administration Guide.

Bugs Fixed in Release 2.4.4.2

The following table lists notable bugs that have been fixed in Oracle Private Cloud Appliance Release 2.4.4.2.

Table 2-5 List of Fixed Bugs

Bug ID	Description
34790192	pca-diag vmppinfo giving warning messages for packages as "Header V3 RSA/SHA256 Signature, key ID ec551f03: NOKEY"
34747228	/boot not having enough space on CNs for upgrade to b218
34726510	Upgrade fails on post upgrade tasks during guestshell destroy
34619179	The repository_mount_monitor Reports Warnings on Local Repositories
34597193	[X8-2] Sort the outputs in pca-admin show network
34595672	[X8-2] Remove pre upgrade check and healthchecks if MN uptime is more than 100 days
34581221	Improve Cisco switch firmware upgrade procedure
34563724	Filesystem in Solaris VM Corrupts Frequently after Testing VM HA by Running "shutdown -h now" on Compute Node
34524559	[X8-2] pca-faultmonitor network_monitor warnings on blue rack invalid check of usb0 ports on server
34505885	[X8-2] Update Password for zfs-root throws an error log "Task failed. Host 192.168.4.2 marked as error in error DB"
34500610	[X8-2] pca_upgrader: upgrade scripts error is not logged in case of abrupt failure
34470591	[X8-2] Add 'bpdufilter enable' by default to uplinks in Cisco configs
34297873	Upgrading Spine Switch Lost External/Uplink Connectivity to PCA
31601967	request to block pca-admin command from non root user
31246371	pca-diag should collect data from all sub-systems without requiring ssh
29551849	add new target "list all-networks" to list all internal and external networks

Changes and Improvements in Release 2.4.4.1

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

Ethernet-based Rack Update Only

The 2.4.4.1 software update is designed for the Ethernet-based Oracle Private Cloud Appliance racks. Your Ethernet-based rack must be at version 2.4.4 before you can upgrade to version 2.4.4.1. For upgrade information, see [Upgrading Oracle Private Cloud Appliance](#).

Oracle Private Cloud Appliance Upgrader Changes

For release 2.4.4.1, health checks and additional pre-checks have been added to the upgrader code to ensure the system is healthy and ready for upgrade. For more information, see [Oracle Private Cloud Appliance Upgrade Guide for Release 2.4.4.1](#).

Oracle Private Cloud Appliance Software Installation Changes

The methodology used for upgrading compute nodes has changed for release 2.4.4.1. In the prior release, the `pca_upgrader` tool was used to upgrade compute nodes. For this release, compute node upgrades are performed using the Oracle Private Cloud Appliance `update compute-node` command, which provides more flexibility in scheduling compute node upgrades. For more information, see "Upgrading the Virtualization Platform" in the [Oracle Private Cloud Appliance Upgrade Guide for Release 2.4.4.1](#).

New CLI Functionality for the Phone Home Service

Starting with release 2.4.4.1, you can enable or disable the Phone Home service on the management nodes using the `set system-property` command in the Oracle Private Cloud Appliance CLI.

To set the Phone Home option, log in to the Oracle Private Cloud Appliance CLI on each management node and set the service to the desired option.

```
# ssh root@10.100.1.101
root@10.100.1.101's password:
root@ovcamn05r1 ~]# pca-admin
PCA> set system-property phonehome [enable | disable]
```

Bugs Fixed in Release 2.4.4.1

The following table lists notable bugs that have been fixed in Oracle Private Cloud Appliance Release 2.4.4.1.

Table 2-6 List of Fixed Bugs

Bug ID	Description
34318768	"[X8-2] During boot of CN, RAID Card failure caused provisioning to overwrite FC LUN Repo"
34099145	"kdump service not operational on compute nodes updated to kernel version: 4.14.35-2047.511.5.6.el7uek.x86_64 "
33935672	"PCA 2.4.4 <code>pca_healthcheck</code> fails on valid X9 BIOS version 61040500 "
33925660	"[X8-2] PCA - Provisioning status give false impression of CN state"
32702599	"PCA: The <code>pca_healthchecker</code> Gets <code>java.lang.IllegalArgumentException</code> If Compute Node IPMI Not Enabled "
33731454	"[X8-2] Add <code>tcpdump</code> rpm to CN repos"
33699788	"[X8-2] delete <code>oci-backup/target</code> command without an argument fails with index out of range error"

Table 2-6 (Cont.) List of Fixed Bugs

Bug ID	Description
33238764	"[X8-2] PCA Healthcheck - Root disk need not always be on /dev/sda - Remove assumption "
31721530	"[X8-2] pca_upgrader precheck should be a super set of checks that would fail if the healthcheck fails "

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

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 [Upgrading Oracle Private Cloud Appliance](#). Release 2.4.4 is currently supported on Ethernet-based racks only.

Oracle 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 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 Release 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 2-7 List of Fixed Bugs

Bug ID	Description
31811656	"PCA 2.4.3 upgrader fails on storage-network update "

Table 2-7 (Cont.) List of Fixed Bugs

Bug ID	Description
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

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 Oracle Private Cloud Appliance 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.

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 bugs that have been fixed in Oracle Private Cloud Appliance Release 2.4.3.

Table 2-8 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 "

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

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 "Functional Network Limitations" in [Monitoring and Managing Oracle Private Cloud Appliance](#) in the [Oracle Private Cloud Appliance Administration Guide for Release 2.4.4](#).

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](#)

Ethernet-Based System Configuration Maximums

Table 3-1 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)	63 GB	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 RAM (64-bit guest)	1000 GB	
Virtual NICs (paravirtualized guest)	31	
Virtual NICs (hardware virtualized guest)	8	

Table 3-1 (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.

Table 3-2 Oracle VM Server Maximums

Item	Maximum	Notes
CPUs	48 per Oracle Server X8-2	The Oracle Server X8-2 contains two 24-core CPUs with hyperthreading enabled.
	64 per Oracle Server X9-2	The Oracle Server X9-2 contains two 32-core CPUs with hyperthreading enabled.
RAM	1.5 TB per Oracle Server X8-2	This is the maximum installed amount of RAM in each server.
	2 T per Oracle Server X9-2	
Virtual Machines	Depends on RAM size of compute node and VM	As described in "Server Pool and Cluster Maximums" in Configuration Maximums in <i>Oracle Private Cloud Appliance Release Notes</i> , 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.

Table 3-3 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.

Table 3-4 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	64 TB	
Files per OCFS2 volume	30,000	
Virtual Disk size	10 TB	The maximum is determined by the hardware capacity.
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.

Table 3-5 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 maximum: 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.
Packet Size (MTU) by interface type	Bond • active: 9054 • inactive: 1500 Ethernet • active: 9054 • inactive: 1500 tun: 9054 bytes vxlan: 9000 bytes mgmt: 9000 bytes storage: 9000 bytes	Do not modify the maximum transfer unit (MTU) of a network interface, standard port or bond, except as documented explicitly in the <i>Oracle Private Cloud Appliance Administrator's Guide</i> .

Table 3-6 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 2 TB DRAM	17
22kVA	32 Core processor 2 TB DRAM	15
15kVA	32 Core processor 2 TB DRAM	6

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.

Interim Firmware Patches

Oracle periodically releases firmware patches for many products, for example to address security vulnerabilities. Sometimes 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 2.4.4.6

[Table 4-1](#), "Qualified Firmware Release 2.4.4.6 for Ethernet-based Systems," lists the firmware versions that have been qualified by Oracle for Private Cloud Appliance Release 2.4.4.6. All required firmware packages are provided on the Controller Software ISO image.

Before upgrading to Release 2.4.4.6, be sure to read the following firmware notes.

Additional Firmware Notes

The upgrade to Private Cloud Appliance Release 2.4.4.6 includes the following changes to Cisco Firmware 10.4(5):

- Cisco Firmware 10.3(4a) included the following change: The default configuration for resource limit for VRF under the VDC configuration was increased by 1 to accommodate a new feature in Cisco Firmware 10.3(4a). This change had no functional impact, and this is now reverted in the Cisco Firmware 10.4(5)M update. See Cisco bug "CSCwh68545 Default resource limit change for VRFs" for more information.

The upgrade to Private Cloud Appliance Release 2.4.4.4 includes the following changes to Cisco Firmware 10.3(4a):

- The Cisco spine and leaf switch files are modified to set auto-recovery as suggested by Cisco to fix 35984246. These leaf and spine files are updated during upgrade when `pca-misc-tasks` runs. If the switch is reloaded, run `pca-misc-tasks` to add this change to the leaf and spine switches. For more information, contact Oracle Support and reference [\[PCA 2.x\] - After a Power Outage a PCA Operating on only One Leaf Switch Can Not Locate its Peer vPC \(Doc ID 2987504.1\)](#).
- The update to Cisco Firmware 10.3(4a) includes the following change: The default configuration for resource limit for VRF under the VDC configuration was increased by 1 to accommodate a new feature in Cisco Firmware 10.3(4a). This change has no functional

impact and will be reverted in the next Cisco Firmware update. See Cisco bug "CSCwh68545 Default resource limit change for VRFs" for more information.

The upgrade to Private Cloud Appliance Release 2.4.4.3 includes the following changes:

- The Cisco spine switch files are modified to set "system dot1q-tunnel transit vlan 3040-3088" as suggested by Cisco to fix bug 34837709. The spine switch files are updated during upgrade when `pca-misc-tasks` runs. If the switch is reloaded from a backup that does not include this change, run `pca-misc-tasks` to add this change to the spine switches. For more information, contact Oracle Support and reference "[PCA 2.x] Port Based Network Redundancy Check (Port Down, Cable Pull, Admin Shut etc.) / Port Based Failover Scenario Leads to Connectivity Loss on PCA Spine Switches (Doc ID 2920424.1)."
- The update to Cisco 10.2.5M includes the following changes:
 - The Spine Filesystem Check is removed from the `pca-admin diagnose spine-switch` output, and the "Leaf Filesystem Check" is removed from the `pca-admin diagnose leaf-switch` output.
 - When you use EM and EMVT to discover the Private Cloud Appliance X8-2, the Cisco switches are not discovered, which causes appliance discovery to fail. To prevent appliance discovery from failing, deselect all Cisco switches before you begin appliance discovery.

The upgrade to Private Cloud Appliance Release 2.4.4.2 includes the following changes:

- The Cisco spine switches are modified to set "spanning-tree bpdupfilter enable" for all uplinks except for the default_5_2 port-group and its associated interfaces (eth1/5/3 and eth1/5/4), which are incomplete by default. The spine switch configurations are updated during upgrade when `pca-misc-tasks` runs. If the switch is reloaded from a backup that does not include this change, run `pca-misc-tasks` to add this change to the spine switches. When a new uplink port-group is created, the "spanning-tree bpdupfilter enable" will be set by default in the port-channel interface configuration.

The upgrade to Private Cloud Appliance Release 2.4.4.1 includes the following changes:

- The Cisco leaf/spine switch backup configuration from "show running-config" is incomplete for interfaces forcefully added to "channel-group 1". The correct setting is shown in the default leaf/spine switch base configurations as "channel-group 1 force mode on".

`pca-admin reload` of a leaf/spine switch from a backup file will modify the "channel-group 1" to append the "force mode on" to that setting.

Manual reload of a leaf/spine switch from a backup file that contains the incomplete "channel-group 1" setting will cause a vPC peer-link failure. All manual reloads of a leaf/spine switch must have the "channel-group 1 force mode on" setting corrected before reloading the switch.

- The upgrade to Cisco firmware 9.3(8) and later versions no longer allows a non-default MTU setting on a vPC peer-link port-channel.

The Cisco leaf/spine switch configuration files have the `mtu` setting removed from the "interface port-channel1" description. Switches that are upgraded originally from the NXOS 7.0.3.17.9 firmware may carry the residual `mtu` setting forward when upgrading to later Cisco NXOS versions. The "interface port-channel1" `mtu` setting is not removed during switch upgrade, but is ignored after the upgrade to NXOS 9.3(8) and later versions.

`pca-admin reload` of a leaf/spine switch from a backup file containing the "interface port-channel1" `mtu` setting will remove the `mtu` setting automatically during the `pca-admin reload` of the switch.

Manual reload of a leaf/spine switch from a backup file that contains the "interface port-channel1" mtu setting will cause a vPC peer-link failure. All manual reloads of a leaf/spine switch must have the disallowed mtu setting removed from the "interface port-channel1" description before reloading the switch.

Table 4-1 Qualified Firmware Release 2.4.4.6 for Ethernet-based Systems

Component	Firmware Version	Patch
System Software	- Oracle Private Cloud Appliance 2.4.4.6 - Oracle VM Manager 3.4.7.7 - Oracle VM Server 3.4.7	
Compute/Management Nodes		
Kernel management node/compute node	UEK6 kernel-uek-5.4.17-2136.346.6.el7uek.x86_64.rpm	
Oracle Server X8-2 management node/compute node	SW 3.8.0 - ILOM 5.1.520 r163206 - BIOS 51.15.01.00 - Patch 37913068: Oracle Server X8-2 SW 3.8.0 - FIRMWARE PACK CX-5 Adapter 16.35.4030 (Oracle_Dual_Port_100_Gb_RoCE_Adapter) RAID SAS9361-16i (Vail) : 4.740.00-8477	37913068
Oracle Server X9-2 compute node	SW 3.8.0 - ILOM 5.1.4.25 r160118 - BIOS 61.13.02.00 - Patch 37712172: Oracle Server X9-2 SW 3.8.0 - FIRMWARE PACK CX-5 Adapter 16.35.4030 (Oracle_Dual_Port_CX5_100_Gb_OCP_Adapter)	37712172
Storage		
Oracle ZFS Storage Appliance ZS7-2 Controller Server	Patch 37777352: ORACLE ZFS STORAGE APPLIANCE SOFTWARE 2013.1.8.83 (OS8.8.83) 2013.06.05.8.83,1-1.1 For more information, see Upgraded Oracle ZFS Storage Appliance Software .	37777352
1.2 TB 10 K RPM Drive	ORA6 (or above)	
Switches		
Cisco Nexus 9348GC-FXP Switch	Cisco Firmware 10.4(5)M	
Cisco Nexus 9336C-FX2 Switch	Cisco Firmware 10.4(5)M	

5

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.

Cisco Firmware Configuration Change

The default configuration for resource limit for VRF under the VDC configuration was changed in Cisco Firmware 10.3(4a).

Previous configuration:

```
limit-resource vrf minimum 2 maximum 4096
```

Cisco Firmware 10.3(4a) configuration:

```
limit-resource vrf minimum 2 maximum 4097
```

This change has no functional impact. No action is required. See Cisco bug "CSCwh68545 Default resource limit change for VRFs" for more information.

Bug 36925686

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 More Aggressively Fails Slow Disks

Oracle ZFS Storage Appliance IDR 8.8.44 5185.1 has a fault management architecture that more aggressively fails slower disks (FMA DISK-8000-VP). Disk failures can be seen because the slow-disk telemetry system-wide variable is set lower.

If you encounter this issue, the following command will show `ireport.io.scsi.cmd.disk.dev.slow.read` with DISK-8000-VP and the HDD disk location.

> maintenance problems show

For more information, see the Oracle Support article [Oracle ZFS Storage Appliance: Handling DISK-8000-VP 'fault.io.disk.slow_rw' \(Doc ID 2906318.1\)](#).

Workaround:

If you determine you have a single UNAVAIL disk or multiple disks that are faulted and in a DEGRADED state, engage Oracle Support to investigate and correct the issue.

Emulex Fibre Channel HBAs Discover Maximum 128 LUNs

When using optional Broadcom/Emulex Fibre Channel expansion cards in Oracle Server X8-2 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=vg/lvroot rd.lvm.lv=vg/lvswap \
rd.lvm.lv=vg/lvusr rhgb 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 named `allowed_services` that controls 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

incorrect opcode Messages in the Console Log

Any installed packages that use the `mstflint` command with a device (-d flag) format using the PCI ID will generate the `mst_ioctl 1177: incorrect opcode = 8008d10` error message. Messages similar to the following appear in the console log:

```
Sep 26 09:50:12 ovcacn10r1 kernel: [ 218.707917] MST:: : print_opcode 549:
MST_PARAMS=8028d001
Sep 26 09:50:12 ovcacn10r1 kernel: [ 218.707919] MST:: : print_opcode 551:
PCICONF_READ4=800cd101
Sep 26 09:50:12 ovcacn10r1 kernel: [ 218.707920] MST:: : print_opcode 552:
PCICONF_WRITE4=400cd102
```

This issue is caused by an error in the PCI memory mapping associated with the InfiniBand ConnectX device. The messages can be safely ignored, the reported error has no impact on PCA functionality.

Workaround: Using `mstflint`, access the device from the PCI configuration interface, instead of the PCI ID.

```
[root@ovcamn06r1 ~]# mstflint -d /proc/bus/pci/13/00.0 q
Image type: FS2
FW Version: 2.11.1280
Device ID: 4099
HW Access Key: Disabled
Description: Node Port1 Port2 Sysimage
GUIDs: 0010e0000159ed0c 0010e0000159ed0d 0010e0000159ed0e 0010e0000159ed0f
MACs: 0010e059ed0d 0010e059ed0e
VSD:
PSID: ORC1090120019
```

Bug 29623624

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

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 Nodes Containing Emulex HBA Card With Maximum FC Paths Reboots With Errors in Oracle VM Manager UI

If a compute node contains an Emulex FC HBA and is configured with 500 LUNs/4000 paths, or 1000 LUNs/4000 paths, you might see the following errors upon reboot of that compute node.

Rack1-Repository errors:

```
Description: OVMEVT_00A000D_000 Presented repository: Rack1-Repository,
mount: ovcacn31r1_/OVS/Repositories/0004fb00000300009f334f0aad38872b, no
longer found on server: ovcacn31r1.
Please unpresent/present the repository on this server
(fsMountAbsPath: /OVS/Repositories/0004fb00000300009f334f0aad38872b,
fsMountSharePath: , fsMountName: 0004fb00000500003150bc24d6f7c2d5
```

```
OVMEVT_00A002D_002 Repository: [RepositoryDbImpl]
0004fb00000300009f334f0aad38872b (Rack1-Repository), is unmounted but in Dom0
DB
```

Compute Node error:

```
Description: OVMEVT_003500D_003 Active data was not found. Cluster service is
probably not running.
```

```
[root@ovcacn31r1 ~]# service o2cb status
Driver for "configfs": Loaded
Filesystem "configfs": Mounted
Stack glue driver: Loaded
Stack plugin "o2cb": Loaded
Driver for "ocfs2_dlmfs": Loaded
Filesystem "ocfs2_dlmfs": Mounted
Checking O2CB cluster "59b95c6b5c6bc782": Offline
Debug file system at /sys/kernel/debug: mounted
```

Workaround: Clear the errors for the compute node and the Rack1-Repository as follows.

1. For the compute node follow these directions:

[Attempting to Present a repository fails with "Cannot present the Repository to server: <hostname> . Cluster is currently down on the server" Oracle Support Document 2041602.1.](#)

2. For the Rack1-Repository, acknowledge the critical error, then refresh the repository.

Bug 33124747

Compute Nodes Containing FC HBA with Maximum FC Paths in Dead State After Reprovisioning

If you are reprovisioning a compute node that contains a Fibre Channel HBA with one of the following configurations, reprovisioning fails and leaves the compute node in a dead state.

- 500 FC LUNs/4000 FC paths
- 1000 FC LUNs/4000 FC paths

To avoid this issue, follow the directions below to reprovision these types of compute nodes.

Note

Compute nodes with FC LUNs less than or equal to 128 FC LUNs with 2 paths each succeeds in reprovisioning without this workaround.

Workaround:

1. Log in to the external storage and remove the compute node's FC initiator from the initiator group (the initiator group that was used to create the max FC paths).
2. Log in to the compute node and run the `multipath -F` command to flush out the FC LUNs that are no longer available. `multipath -ll` will now only show 3 default LUNs.

```
[root@ovcacn32r1 ~]# multipath -F

Jul 21 17:23:12 | 3600144f0d0d725c7000060f5ecb30004: map in use
Jul 21 17:23:18 | 3600062b20200c6002889e3a010d81476: map in use
Jul 21 17:23:22 | 3600144f0d0d725c7000060f5ecb10003: map in use
[root@ovcacn32r1 ~]# multipath -ll
3600144f0d0d725c7000060f5ecb30004 dm-502 SUN,ZFS Storage 7370
size=3.0T features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
`-+- policy='round-robin 0' prio=50 status=active
   `- 11:0:0:3   sdbks 71:1664   active ready running
3600062b20200c6002889e3a010d81476 dm-0 AVAGO,MR9361-16i
size=1.0T features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=1 status=active
   `- 8:2:1:0    sdb  8:16     active ready running
3600144f0d0d725c7000060f5ecb10003 dm-501 SUN,ZFS Storage 7370
size=12G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
`-+- policy='round-robin 0' prio=50 status=active
   `- 11:0:0:1   sdbkr 71:1648   active ready running
```

3. Reprovision the compute node.
4. (Emulex only) Log in to the compute node and re-apply the grub customization for the Emulex driver, see [Emulex Fibre Channel HBAs Discover Maximum 128 LUNs](#).
5. Log in to the external storage and re-add the compute node's FC initiator into the initiator group.
6. Log in to the Oracle VM Manager UI and add the compute node as an admin server to the Unmanaged FibreChannel Storage Array. Refresh the Unmanaged FibreChannel Storage Array. Max FC paths should be restored.

Bug 33134228

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

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

- After a compute node reprovision
- After a compute node upgrade
- After exposing a compute node with hundreds of new LUNs (either by LUN creation on the storage array or by fabric rezoning)

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

- The `tailf /var/log/devmon.log` command shows many event messages similar to the following:

```
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.

- The `systemd-udevd` process consumes 100% CPU in `top` command.
This is resolved when `systemd-udevd` no longer consumes a large percentage of the CPU.
- The `multipath -ll` command does not show all the LUNs. The command might show a fraction of the LUNs expected.
- The `multipath -ll | grep "active ready running" | wc -l` command might not count all the LUNs. The command might 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 might fail because of a Fibre Channel LUN path flapping problem. Use the following workaround to avoid the issue.

Workaround: [\[PCA 2.4.4\] Upgrade Compute Node with Fibre Channel Luns may Fail due to FC Path Flapping \(Doc ID 2794501.1\)](#).

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

If your compute node fails the Faultmonitor `firewall_monitor` check and displays the following log error, you are encountering a port error which creates a false report and pushes it to Phone Home, if Phone Home is enabled. The `firewall_monitor` verifies whether the required ports for Oracle VM Manager and the compute node are opened or not.

```
[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.'
```

Workaround: 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"\)](#).

Certain TZ Configuration is Failing on Cisco Switches

Starting in 2016, tzdata implemented numeric timezone abbreviations like "+03" for new timezones. The Cisco switches only support timezones abbreviations with alphabetic abbreviations like "ASTT". Attempting to change the timezone on a Cisco switch could cause the following error in the `ovca.log` file:

```
[2022-05-30 15:08:13 10455] ERROR (cisco:145) Configuration failed
partially:Clock timezone set:: Timezone name should contain alphabets only
```

Workaround: Do not change the timezones on the Cisco switches. Cisco switches will always report the time in UTC.

Bug 34223027

NFS Shares on Internal ZFS Will Fail After ZFS Firmware Update If vnuc Owned by SN02

Starting with ZFSSA AK version 8.8.30, it is now a requirement that the address used to mount shares from a pool must be formally owned by the same head which formally owns the pool, as shown by Configuration -> Cluster.

Software version 2.4.4.1 introduces a new pre-check which flags any storage network interfaces that have owner = ovcasn02r1 so the customer can manually correct the owner to ovcasn01r1 before proceeding with the upgrade. If you see the following error, proceed to the workaround below.

```
[2022-05-24 18:57:22 33554] ERROR (precheck:154) [ZFSSA Storage Network Interfaces Check (Ensure ovcasn01r1 is the owner of all customer-created storage network interfaces)] Failed
```

The check failed: Detected customer-created storage network interface(s) owned by ovcasn02r1: net/vnic10, net/vnic11, net/vnic12, net/vnic7, net/vnic8, net/vnic9

Workaround: See [\[PCA 2.4.4.1\] Pre-check "ZFSSA Storage Network Interfaces Check" Fails \(Doc ID 2876150.1\)](#).

Bug 34192251

Repository Size is Not Reflecting Properly in OVMM GUI

The Oracle VM Manager GUI can report an incorrect repository size, which may cause VMs to hang because the repository is actually full. Use another method to check the repository size, like the compute node df output or the OVM CLI.

Workaround: Check the repository size using OVM CLI.

```
OVMM> show repository name=NFS-ExtZFSSA-Repository
Command: show repository name=NFS-ExtZFSSA-Repository
Status: Success
Time: 2021-10-19 10:48:15,294 UTC
Data:
  File System = 14a6cf21-a170-41aa-9a09-7b768aaabc6f [nfs on
192.168.40.242:/export/NFS-Ext-Repo]
  Manager UUID = 0004fb000001000087ae02edfd0534dc
  File System Free (GiB) = 998.48
  File System Total (GiB) = 1018.84
  File System Used (GiB) = 20.37
  Used % = 2.0
  Apparent Size (GiB) = 25.0
  Capacity % = 2.5
```

Bug 33455258

Oracle Private Cloud Appliance Software

This section describes software-related limitations and workarounds.

Upgrade to UEK6 Kernel Fails to Boot with 1000 FC LUNs Attached

Compute node upgrade to the UEK6 kernel fails during boot with 1000 FC LUNs attached.

The following workaround is applied to /etc/systemd/system.conf during upgrade of the CNs to Private Cloud Appliance 2.4.4.3 and to newly provisioned CNs. You do not need to perform this workaround yourself.

Workaround: Increase systemd timeouts. In /etc/systemd/system.conf, update both DefaultTimeoutStartSec and DefaultTimeoutStopSec as follows:

```
DefaultTimeoutStartSec=900s
DefaultTimeoutStopSec=900s
```

Bug 35780853

Fault Reports Are Not Removed Until Another Fault Is Reported

Fault Monitor reports that are older than `report_dir_cleanup_days` are not removed until a new fault is generated. Fault report cleanup is performed only when a fault is reported.

Bug 35665992

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.

Upgrader UI Installation Step Is Skipped

This issue causes the `pca_upgrader` to skip the `ui_install` step. The following message appears in the log file: "The Oracle VM upgrade has already been completed on the other manager. Skipping OVCA UI installation on *management-node-name*."

If you previously performed an upgrade and did not perform the preventive step described in the following workaround, you might need to manually install the UI.

Workaround: Before you begin the upgrade, delete the following file to prevent skipping UI installation:

```
/nfs/shared_storage/pca_upgrader/pxe_upgrade/.standby_upgrade_complete
```

To manually install the UI if you did not delete the `.standby_upgrade_complete` file prior to performing the upgrade, see [\[PCA 2.x\] Accessing Dashboard Gets HTTP 404 Error After Upgrading Management Nodes \(Doc ID 2491230.1\)](#).

Bug: 34906041

Deleting a Storage Network Is Allowed When the Network Is Assigned to NFS or iSCSI Storage

You are able to delete a `storage_network` network while that network is assigned to `nfs-storage` or `iscsi-storage`.

Workaround: Before you delete a `storage_network` network, use `show nfs-storage` and `show iscsi-storage` to verify that the `storage_network` is not assigned.

Bug 34767928

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 Reprovisioned

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/wls1/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 Manager 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 Manager 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

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 for Release 2.4.4](#).

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 Release 2.4.4](#) 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 Manger Web UI for the compute node you are upgrading
2. Move each file type as described below:

Table 5-1 Moving Items Out of the Local Repository

Item	Steps
ISO	a. Clone the ISOs to other repositories. b. Delete the ISO files from the local repository.
VM Template	a. Move the template to another repository using clone customiser.
Virtual Appliances	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 5-1 (Cont.) Moving Items Out of the Local Repository

Item	Steps
Virtual Disks	- 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: 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	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 ~]#

```

- 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

- Log in to the Oracle VM Manager Web UI.
- In the Servers and VMs tab, select the appropriate server pool and validate that the compute node is part of that server pool.
- From the Repositories tab, select the repository and note the physical disk over which the repository lies.
- From the Repositories tab, select the repository, then edit the concerned repository and check Release Ownership.
- 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

- 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

Backup of Config Can Fill Filesystem and Cause Numerous Problems

Over time, backups of the Private Cloud Appliance configuration information can accumulate at `/nfs/shared_storage/backups` and fill the filesystem. Periodically you must remove old backups to ensure the filesystem does not run out of room.

Workaround: Remove backups using the following procedure.

Removing Backups

1. Log in to the active management node as root user.
2. Move any custom scripts or data located in `/nfs/shared_storage/backups` to a different location. During this cleanup procedure everything in that location will be deleted, *except the backup files for the selected retention period*.
`/nfs/shared_storage/backups` must *only* contain backup tarballs and or uncompressed backup directories.
3. Remove old backups using this command:

```
[root@ovcamn05r1 ~]# find /nfs/shared_storage/backups -maxdepth 1 -mtime
+<retention-period-in-days> -exec rm -rf {} \;
```

For example, to delete all older backups, inclusive of any uncompressed backup directories, older than 30 days, type:

```
[root@ovcamn05r1 ~]# find /nfs/shared_storage/backups -maxdepth 1 -mtime
+30 -exec rm -rf {} \;
```

Bug 33947155

Two Different Release Packages of Same Tool Present in the ISO

There is a chance that an ISO file could contain multiple rpm files for `impitool`. There is no action to be taken, the upgrader tool will install the correct version.

Bug 34375901

Intermittent Error When Deleting the Uplink Port Group

When you create an `uplink-port-group`, then restore the spine switch configuration from a backup the `uplink-port-group` is removed from the switch configuration. However, the `uplink-port-group` still appears in the Oracle Private Cloud Appliance interface. An attempt to delete the `uplink-port-group` may fail.

```
PCA> delete uplink-port-group group-name
*****
WARNING !!! THIS IS A DESTRUCTIVE OPERATION.
*****
Are you sure [y/N]:y
Status: Failure
Error Message: Error (UPLINK_PORT_GROUP_012): Failed to delete the uplink
port group: group-name.
```

Workaround: Repeat the delete command a second time.

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

Bug 34379557

6

Feedback and Support

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

If you cannot find an answer to your question by using information available on [My Oracle Support](#), you can open a service request for help from Oracle Support Services.

Oracle Auto Service Request

If Oracle Auto Service Request (ASR) is enabled on the Private Cloud Appliance, a service request will be created and sent to Oracle Support automatically for some failures. For details, see "Oracle Auto Service Request (ASR)" in [Servicing Oracle Private Cloud Appliance Components](#) in the *Oracle Private Cloud Appliance Administration Guide for Release 2.4.4*.

Important

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 incurs a Time and Materials charge.

Reporting Problems

If the issue is not automatically reported by ASR, open a service request at [My Oracle Support](#) to request assistance from Oracle Support. You need an Oracle Premier Support Agreement and your Oracle Customer Support Identifier (CSI).

If you do not know your CSI, 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.

Your Oracle Premier Support CSI does not cover customization support, third-party software support, or third-party hardware support.

When you open a service request, 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, patches that you have applied, and other software that might be affecting the problem.
- Details of steps that you have taken to reproduce the problem.
- Applicable logs and other support data.

Much of this information can be provided by creating a support bundle and attaching it to the service request as described in "Data Collection for Service and Support" in [Monitoring and Managing Oracle Private Cloud Appliance](#) in the *Oracle Private Cloud Appliance Administration Guide for Release 2.4.4*.