

Oracle Private Cloud Appliance

Release Notes for Release 3.0.2



F71311-02
February 2023



Oracle Private Cloud Appliance Release Notes for Release 3.0.2,

F71311-02

Copyright © 2022, 2023, Oracle and/or its affiliates.

Contents

Preface

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

1 Accessibility and Oracle Private Cloud Appliance

Oracle JET User Interface Accessibility Features	1-1
Oracle Server X9-2 Accessibility Features	1-1
Oracle Server X9-2 Hardware Accessibility	1-1
Oracle Integrated Lights Out Manager Manager Accessibility	1-2
Oracle Hardware Management Pack Accessibility	1-2
BIOS Accessibility	1-3

2 Feature Updates

Latest Features	2-1
Features Released in Software Version 3.0.1-b741265 (November 2022)	2-3
Features Released in Software Version 3.0.1-b697160 (August 2022)	2-4

3 Service Limits

Tenancy Resource Configuration Limits	3-1
System Load and Concurrency Limits	3-2

4 Known Issues and Workarounds

Platform Issues	4-1
Compute Node Provisioning Takes a Long Time	4-1
Not Authorized to Reconfigure Appliance Network Environment	4-1

Grafana Service Statistics Remain at Zero	4-2
Terraform Provisioning Requires Fully Qualified Domain Name for Region	4-2
Synchronizing Hardware Data Causes Provisioning Node to Appear Ready to Provision	4-2
Rack Elevation for Storage Controller Not Displayed	4-2
Free-Form Tags Used for Extended Functionality	4-3
Imported Images Not Synchronized to High-Performance Pool	4-4
API Server Failure After Management Node Reboot	4-4
CLI Command Returns Status 500 Due To MySQL Connection Error	4-4
Microservice Pods Unresponsive After Storage Controller Failover	4-5
Administrators in Authorization Group Other Than SuperAdmin Must Use Service CLI to Change Password	4-5
Password Policy for Monitoring Framework Differs from Other Infrastructure Components	4-6
Service Web UI and Grafana Unavailable when HAProxy Is Down	4-6
Lock File Issue Occurs when Changing Compute Node Passwords	4-7
Compute Node Hangs at Dracut Prompt after System Power Cycle	4-7
No Error Reported for Unavailable Spine Switch	4-7
ZFS Storage Appliance Controller Stuck in Failsafe Shell After Power Cycle	4-8
First Tenancy Creation Fails and Successive Attempts Report It Already Exists	4-8
Concurrent Compute Node Provisioning Operations Fail Due to Storage Configuration Timeout	4-8
Switch Fault Reported After Password Change	4-9
Kubernetes Health Checker Cannot Find Endpoints	4-9
Leaf Switch Stuck in Password Change	4-10
Data Switch Fails to Boot Due to Active Console Connection	4-10
User Interface Issues	4-11
Moving Resources Between Compartments Is Not Supported in the Compute Web UI	4-11
No Available Compute Web UI Operation to Update Instance Pool	4-11
Saving Resource Properties Without Modifications Briefly Changes Status to Provisioning	4-11
NFS Export Squash ID Not Displayed	4-12
Scrollbars Not Visible in Browser	4-12
Authorization Failure When Retrieving Compartment Data	4-12
Object List Is Not Updated Automatically	4-13
File Storage Mount Target Link Not Available	4-13
UDP Ports Not Displayed In Security List Rules Table	4-13
Not All Resources Shown in Drop-Down List	4-13
Volume Group Can Be Created Without Name	4-14
File Systems and Mount Targets Not Displayed	4-14
Optional ICMP Security Rule Parameters Cannot Be Removed	4-14
Compartment Selector Not Available When Creating DHCP Options	4-15
Custom Search Domain Error Not Rolled Back When Operation Is Canceled	4-15

DHCP Options Error Message for Custom Search Domain Is Misleading	4-16
Unclear Error when Logging in to Service Web UI with Insufficient Privileges	4-16
Cloning Snapshot to New File System Not Supported from Compute Web UI	4-17
No Details Displayed for File System Cloned from Snapshot	4-17
Network Environment Information Is Not Refreshed Automatically After Admin Network Is Enabled	4-17
Unable to Display Details of Instance Backup	4-18
IP Address List on VNIC Detail Page Not Updated	4-18
No Error Displayed When Attempting to Reserve Public IP Address While All Public IPs In Use	4-18
Incorrect Error Message Displayed When Deleting Exadata Network	4-18
When Creating Volume Group from Backup Its Volumes Are Not Displayed	4-19
Networking Issues	4-19
DNS Zone Scope Cannot Be Set	4-19
To Update a DNS Record the Command Must Include Existing Protected Records	4-19
Oracle Linux 8 Instance Host Name Resolution Fails	4-19
Create Route Table Fails With Confusing Error Message	4-20
VCN Creation Uses Deprecated Parameter	4-20
File Storage Traffic Blocked By Security Rules	4-20
Stateful and Stateless Security Rules Cannot Be Combined	4-22
VCN With Single Subnet of Same Size Not Supported	4-22
Routing Failure With Public IPs Configured as CIDR During System Initialization	4-23
Admin Network Cannot Be Used for Service Web UI Access	4-23
Network Configuration Fails During Initial Installation Procedure	4-23
External Certificates Not Allowed	4-24
Low Uplink MTU Setting Prevents Browser Interface Loading	4-24
VCNs Stuck in Provisioning After Concurrent Creation Operations	4-24
Admin Network Configuration Blocks Access to Services	4-24
After Upgrade to Release 3.0.2 the Metadata of Existing Instances Contains No FQDN	4-25
DNS Entries on Oracle Linux 8 Instances Incorrect After Upgrade to Release 3.0.2	4-25
Compute Service Issues	4-25
No Consistent Device Paths for Connecting to Block Volumes	4-25
Instance Pools Cannot Be Terminated While Starting or Scaling	4-26
Network Interface on Windows Does Not Accept MTU Setting from DHCP Server	4-26
Oracle Solaris Instance in Maintenance Mode After Restoring from Backup	4-26
Instance Disk Activity Not Shown in Compute Node Metrics	4-27
Attached Block Volumes Not Visible Inside Oracle Solaris Instance	4-27
Host Name Not Set In Successfully Launched Windows Instance	4-27
Oracle Solaris Instance Stuck in UEFI Interactive Shell	4-28
Instance Backups Can Get Stuck in an EXPORTING or IMPORTING State	4-28
Instances Experience Extended Downtime After Compute Node Reboot	4-29

Migrated Instance Fails to Start Due to Incomplete Cleanup of Guest Configuration on Target Compute Node	4-29
Instance Not Started After Fault Domain Change	4-29
Instance Start or Migration Fails Due to Libvirt Process Write Error	4-30
Instance Migration Stuck in MOVING State	4-30
Migration Task Stays Active Even After the VM Migration Completed	4-31
Storage Services Issues	4-31
Creating Image from Instance Takes a Long Time	4-31
Large Object Transfers Fail After ZFS Controller Failover	4-31
Use Multipart Upload for Objects Larger than 100MiB	4-31
File System Export Temporarily Inaccessible After Large Export Options Update	4-32
Block Volume Stuck in Detaching State	4-32
Detaching Volume Using Terraform Fails Due To Timeout	4-32
Scheduled Volume Backups Do Not Appear in Backup List	4-33
OCI CLI Might Not Return the Correct Value for Object Storage Namespace	4-33
Creating File System Export Fails Due To Timeout	4-34
Concurrent File System Creation Operations Cause Exception During Flush	4-34
File System Access Lost When Another Export for Subset IP Range Is Deleted	4-34
File System Export UID/GID Cannot Be Modified	4-35
ZFS Pool Usage Decreases After Instance Migration Due to Disk Allocation Change	4-35
Serviceability Issues	4-35
Minimum Upgrader Package Version for Upgrade to Release 3.0.2	4-35
Order of Upgrading Components Has Changed	4-36
DR Configurations Are Not Automatically Refreshed for Terminated Instances	4-36
Rebooting a Management Node while the Cluster State is Unhealthy Causes Platform Integrity Issues	4-37
ULN Mirror Is Not a Required Parameter for Compute Node Patching	4-37
Patch Command Times Out for Network Controller	4-38
Upgrade Commands Fail when One Storage Controller Is Unavailable	4-38
Management Node Upgrade Causes Loki Outage	4-38
Instances with a Shared Block Volume Cannot Be Part of Different Disaster Recovery Configurations	4-39
Time-out Occurs when Generating Support Bundle	4-39
DR Operations Intermittently Fail	4-39
Update of Disaster Recovery Replication Target Fails	4-40

Preface

This publication is part of the customer documentation set for Oracle Private Cloud Appliance Release 3.0.2. Note that the documentation follows the release numbering scheme of the appliance software, not the hardware on which it is installed. All Oracle Private Cloud Appliance product documentation is available at <https://docs.oracle.com/en/engineered-systems/private-cloud-appliance/index.html>.

Oracle Private Cloud Appliance Release 3.x is a flexible general purpose Infrastructure as a Service solution, engineered for optimal performance and compatibility with Oracle Cloud Infrastructure. It allows customers to consume the core cloud services from the safety of their own network, behind their own firewall.

Audience

This documentation is intended for owners, administrators and operators of Oracle Private Cloud Appliance. It provides architectural and technical background information about the engineered system components and services, as well as instructions for installation, administration, monitoring and usage.

Oracle Private Cloud Appliance has two strictly separated operating areas, known as enclaves. The Compute Enclave offers a practically identical experience to Oracle Cloud Infrastructure: It allows users to build, configure and manage cloud workloads using compute instances and their associated cloud resources. The Service Enclave is where privileged administrators configure and manage the appliance infrastructure that provides the foundation for the cloud environment. The target audiences of these enclaves are distinct groups of users and administrators. Each enclave also provides its own separate interfaces.

It is assumed that readers have experience with system administration, network and storage configuration, and are familiar with virtualization technologies. Depending on the types of workloads deployed on the system, it is advisable to have a general understanding of container orchestration, and UNIX and Microsoft Windows operating systems.

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.

Convention	Meaning
<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, code in examples, text that appears on the screen, or text that you enter.
\$ prompt	The dollar sign (\$) prompt indicates a command run as a non-root user.
# prompt	The pound sign (#) prompt indicates a command run as the <code>root</code> user.

Documentation Accessibility

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

For information about the accessibility of the Oracle Help Center, see the Oracle Accessibility Conformance Report at <https://www.oracle.com/corporate/accessibility/templates/t2-11535.html>.

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

Accessibility and Oracle Private Cloud Appliance

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

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

For information about the accessibility of the Oracle Help Center, see the Oracle Accessibility Conformance Report at <https://www.oracle.com/corporate/accessibility/templates/t2-11535.html>.

Oracle JET User Interface Accessibility Features

The Oracle Private Cloud Appliance user interface is built with Oracle JavaScript Extension Toolkit (JET) which is compliant with the Americans with Disabilities Action (ADA) requirements. For detailed accessibility information about JET, refer to [Oracle JET and Accessibility](#).

Oracle Server X9-2 Accessibility Features

Oracle strives to make its products, services, and supporting documentation usable and accessible to the disabled community. To that end, products, services, and documentation include features that make the product accessible to users of assistive technology.

The accessibility features of Oracle Server X9-2 are detailed within the following product components:

- Oracle Server X9-2 hardware
- Oracle Integrated Lights Out Manager (ILOM)
- Oracle Hardware Management Pack
- BIOS

Oracle Server X9-2 Hardware Accessibility

Oracle Server X9-2 hardware has color-coded labels, component touch points, and status indicators (LEDs) that provide information about the system. These labels, touch points, and indicators can be inaccessible features for sight-impaired users. The product's HTML documentation provides context and descriptive text available to assistive technologies to aid in interpreting status and understanding the system.

You can also use the built-in Oracle Integrated Lights Out Manager (ILOM) to obtain information about the system. Oracle ILOM provides a browser-based user interface (UI) and a command-line interface (CLI) that support assistive technologies for real-time viewing of

system status, indicator interpretation, and system configuration. For details, see [Oracle Integrated Lights Out Manager Manager Accessibility](#).

Oracle Integrated Lights Out Manager Manager Accessibility

You can use the Oracle Integrated Lights Out Manager (ILOM) UI to monitor and manage the server hardware. The Oracle ILOMUI does not require a special accessibility mode; rather, its accessibility features are always available. The UI was developed using standard HTML and JavaScript and its features conform to accessibility guidelines.

To navigate a UI page and select items or enter commands, use standard keyboard inputs, such as the Tab key to go to a selection, or the up and down arrow keys to scroll through the page. You can use standard keyboard combinations to make menu selections.

For example, using the Oracle ILOM Open Problems UI page, you can identify faulted memory modules (DIMMs) or processors (CPUs) that would otherwise be identified by a lighted LED indicator on the motherboard. Likewise, you can use the Oracle ILOM UI to monitor the hardware power states that are also indicated by flashing LED indicators on the hardware.

The Oracle ILOM CLI is an alternative and equivalent way to access the Oracle ILOM UI features and functionality. Because the operating systems that run on the Oracle server hardware support assistive technologies to read the content of the screen, you can use the CLI as an equivalent means to access the color-based, mouse-based, and other visual-based utilities that are part of the UI. For example, you can use a keyboard to enter CLI commands to identify faulted hardware components, check system status, and monitor system health.

You can use the Oracle ILOM Remote Console Plus application to access both a text-based serial console and a graphics-based video console that enable you to remotely redirect host server system keyboard, video, mouse, and storage devices. Note, however, that the Oracle ILOM Java Remote Console Plus does not support scaling of the video frame within the Java application. You need to use assistive technology to enlarge or reduce the content in the Java Remote Console Plus display.

As an alternative method to using the BIOS Setup Utility to configure BIOS settings, Oracle ILOM provides a set of configurable properties that can help you manage the BIOS configuration parameters on an Oracle x86 server. Using Oracle ILOM, you can do the following:

- Back up a copy of the BIOS configuration parameters to an XML file using the Oracle ILOM UI.
- Edit the XML file using a standard XML editor. The BIOS XML tags correlate directly to the BIOS screen labels.
- Restore the XML file of the backed up or edited configuration parameters to BIOS.

The UI and CLI methods for using Oracle ILOM are described in the accessible HTML documentation for Oracle ILOM at <https://www.oracle.com/goto/ilom/docs>.

Oracle Hardware Management Pack Accessibility

Oracle Hardware Management Pack software is a set of CLI tools. Oracle Hardware Management Pack software does not include product-specific accessibility features.

Using a keyboard, you can run the CLI tools as text commands from the operating system of a supported Oracle server. All output is text-based.

Additionally, most Oracle Hardware Management Pack tools support command output to a text log file or XML file, which can be used for text-to-speech conversion. Accessible man pages are available that describe the Hardware Management Pack tools on the system on which those tools are installed.

You can install and uninstall Oracle Hardware Management Pack by using text commands entered from the CLI. Assistive technology products such as screen readers, digital speech synthesizers, or magnifiers can be used to read the content of the screen.

Refer to the assistive technology product documentation for information about operating system and command-line interface support.

The CLI tools for using the software are described in the accessible HTML documentation for Hardware Management Pack at <https://www.oracle.com/goto/ohmp/docs>.

BIOS Accessibility

When viewing BIOS output from a terminal using the serial console redirection feature, some terminals do not support function key input. However, BIOS supports the mapping of function keys to Control key sequences when serial redirection is enabled. Descriptions of the function key to Control key sequence mappings are provided in the product documentation, typically within the server Service Manual. You can navigate the BIOS Setup Utility by using either a mouse or keyboard commands.

As an alternative method of configuring BIOS settings using the BIOS Setup Utility screens, Oracle ILOM provides a set of configurable properties that can help you manage the BIOS configuration parameters on an Oracle x86 server. For more information, see [Oracle Integrated Lights Out Manager Manager Accessibility](#).

2

Feature Updates

This section contains a list of the new features and changes to features that have been added to the Oracle Private Cloud Appliance software since its initial release. You can obtain the latest features and bug fixes by applying patches to your system. For more information, see the [Oracle Private Cloud Appliance Patching Guide](#).

Latest Features

Platform Images

Platform images are available to all compartments in all tenancies without being imported to any compartment by users.

The following platform images are delivered with this Private Cloud Appliance release:

Oracle Linux 8	uln-pca-Oracle-Linux-8-2022.08.29_0.oci
Oracle Linux 7.9	uln-pca-Oracle-Linux-7.9-2022.08.29_0.oci
Oracle Solaris 11.4	uln-pca-Oracle-Solaris-11.4.35-2021.09.20_0.oci

New platform images are delivered through Private Cloud Appliance installation, upgrade, and patch.

Important:

The Service Enclave administrator must import platform images after Private Cloud Appliance installation and should import platform images after every upgrade and patch in case new images were delivered. See "Providing Platform Images" in [Hardware Administration](#) in the *Oracle Private Cloud Appliance Administrator Guide*.

The following Windows versions have been tested and are supported on this release of Oracle Private Cloud Appliance as custom images:

- Microsoft Windows Server 2022
- Microsoft Windows Server 2019
- Microsoft Windows Server 2016
- Microsoft Windows Server 2012 R2, with Window2k12 Datacenter R2 64 Bit
- Microsoft Windows Server 2012, with Window2k12 Datacenter R2 64 Bit

Instance Backup and Restore

Oracle Private Cloud Appliance provides API commands that enable you to back up instances. The commands are flexible to suit a variety of use cases, including:

- Back up instances and any attached block volumes.
- Store the backups on another server for safekeeping.
- Restore a faulty instance and any attached block volumes.
- Use the backup to create matching instances.
- Use the backup and restore feature to migrate instances to another tenancy, or to another appliance.

 **Note:**

The maximum recommended object size supported is 10TB of total data and the maximum recommended object part size in a multipart upload is 5 GB.

For details see *Instance Backup and Restore* in the [Oracle Private Cloud Appliance Concepts Guide](#) and *Backing Up and Restoring an Instance* in the [Oracle Private Cloud Appliance User Guide](#).

Instance Shape Update

When you update an instance, you can change the shape. You can change from any shape to any other shape. If the flexible shape is specified, you can change the shape configuration. For more information, see "Updating an Instance" in [Compute Instance Deployment](#) in the Oracle Private Cloud Appliance User Guide.

Enhanced Compute Instance Availability

If a compute node is lost due to a failure, a new reboot migration process is invoked. Its purpose is to evacuate compute instances to other compute nodes. Fault domain preference is strictly enforced with instance migration. If a compute instance cannot be migrated to another compute node in the same fault domain due to insufficient capacity, the instance is stopped and must be restarted manually.

File System Clones

You can use the OCI CLI to create file system clones. A clone is a new file system that is created from a snapshot of an existing file system. Snapshots preserve the state of the data of a file system at a particular point in time. If you take snapshots of a file system at regular intervals, you can create clones of the file system as it existed at multiple points in its lifetime.

Cloned file systems are managed in the same way that any other file system is managed. See the [File System Storage](#) chapter in the Oracle Private Cloud Appliance User Guide.

Tags for Specifying Certain Property Values

Private Cloud Appliance provides defined tags that enable you to set values for some properties. Applying these tags is the only way to set these particular properties.

The following defined tags are in the OraclePCA tag namespace.



Note:

Do not create your own tags in the OraclePCA tag namespace.

Resource, Operation	Tag Name	Values
Block volume, create and update	logBias	LATENCY, THROUGHPUT
	secondaryCache	ALL, METADATA, NONE
File system, create	databaseRecordSize	512, 1024, 2048, 4096, 8192, 16384, 32768, 65536, 131072, 262144, 524288, 1048576

You must use the OCI CLI to set these tags. See the OCI CLI procedures in "Working with Resource Tags" in [Resource Tag Management](#) in the Oracle Private Cloud Appliance User Guide.

For examples, see "Creating a Block Volume" in [Block Volume Storage](#) in the Oracle Private Cloud Appliance User Guide and "Creating a File System" in [File System Storage](#) in the Oracle Private Cloud Appliance User Guide.

Capacity Monitoring

Administrators have direct access to the current consumption of key physical resources: CPU, memory and storage space. For more information, see "Monitoring System Capacity" in the chapter [Status and Health Monitoring](#) of the Oracle Private Cloud Appliance Administrator Guide.

Full Administration Network Segregation

In an environment with elevated security requirements, you can optionally segregate administrative appliance access from the data traffic. The administration network physically separates configuration and management traffic from the operational activity on the data network. In this configuration, only the administration network provides access to the Service Enclave, which includes the monitoring, metrics collection and alerting services, the API service, and all component management interfaces.

Service Request Diagnostic Data

If the Private Cloud Appliance is registered for Oracle Auto Service Request (ASR), certain hardware failures cause a service request and diagnostic data to be automatically sent to Oracle support. The collection of diagnostic data is also called a support bundle. A Service Enclave administrator can also create and send a service request and supporting diagnostic data separate from ASR. For more information about ASR and support bundles, see [Status and Health Monitoring](#) in the *Oracle Private Cloud Appliance Administrator Guide*.

Features Released in Software Version 3.0.1-b741265 (November 2022)

Flexible Compute Shapes

A flexible compute shape lets you customize the number of OCPUs and the amount of memory when launching your instance. This flexibility lets you create instances that meet

your workload requirements, while optimizing performance and using resources efficiently. For details see [Standard Shapes](#) in the [Oracle Private Cloud Appliance Concepts Guide](#).

GUI Support for Viewing CPU and Memory Metrics

As of this release, you can view Memory and CPU metrics at a fault domain level using the Service Enclave GUI. For details, see [Viewing CPU and Memory Usage by Fault Domain](#) in the [Oracle Private Cloud Appliance Administrator Guide](#).

Features Released in Software Version 3.0.1-b697160 (August 2022)

Compute Instance Availability

When compute instances go down because of a compute node reboot or failure, the system takes measures to recover the compute instances automatically. For details, see [Compute Instance Availability](#) in the [Oracle Private Cloud Appliance Concepts Guide](#).

Optimized NUMA Alignment

Algorithm optimizations are in place to ensure that the hypervisor assigns compute instances on physical resources (CPU and memory) with best possible alignment to compute node NUMA architecture. For details, see [Physical Resource Allocation](#) in the [Oracle Private Cloud Appliance Concepts Guide](#).

View CPU and Memory Metrics at the Fault Domain Level

Memory and CPU usage metrics are available at the compute nodes level already. Each node belongs to a fault domain. New functionality provides the option to view these metrics at a fault domain level. For details, see [Fault Domain Observability](#) in the [Oracle Private Cloud Appliance Concepts Guide](#), and [Viewing CPU and Memory Usage by Fault Domain](#) in the [Oracle Private Cloud Appliance Administrator Guide](#).

Secondary Private IP Addresses

After an instance is launched, you can attach secondary private IP addresses to the primary VNIC or to any secondary VNICs. These secondary private IP addresses are especially useful when running multiple services or endpoints on a single instance, or for instance failover scenarios.

For more information, see "About Secondary Private IPs" under "IP Addressing" in the [Virtual Networking Overview](#) section of the [Oracle Private Cloud Appliance Concepts Guide](#).

For procedures, see "Assigning a Secondary Private IP Address" in the [Networking](#) chapter of the [Oracle Private Cloud Appliance User Guide](#).

3

Service Limits

This chapter contains the service limits for Oracle Private Cloud Appliance. The limits presented here have been tested and are fully supported by Oracle.

The minimum appliance configuration contains three compute nodes and one high-capacity disk shelf with 100TB of usable disk space. Both compute and storage capacity can be expanded by adding compute nodes and disk shelves.

Tenancy Resource Configuration Limits

This section lists the resource limits that are dependent on the appliance architecture. Oracle Private Cloud Appliance supports up to 8 tenancies; these are default limits per tenancy. The numbers provided here apply to any Private Cloud Appliance installation, regardless of its hardware configuration.

Service	Resource Type	Limit
IAM Service	Users	100
IAM Service	Groups	100
IAM Service	Users per group	100
IAM Service	Groups per user	50
IAM Service	Compartments	50
IAM Service	Policies	100
IAM Service	Policy statements	50 per policy
IAM Service	Identity providers	3
IAM Service	Group mappings	100 per identity provider
Networking Service	VCNs	10
Networking Service	Subnets	20 per VCN
Networking Service	Dynamic routing gateways	8 total across all tenancies
Networking Service	Internet gateways	1 per VCN
Networking Service	Local peering gateways	5 per VCN
Networking Service	NAT gateways	1 per VCN
Networking Service	Service gateways	1 per VCN
Networking Service	Reserved public IPs	1/16th of customer-defined block
Networking Service	Ephemeral public IPs	2 per compute instance
Networking Service	DHCP options	30 per VCN
Networking Service	Route tables	20 per VCN
Networking Service	Route rules	50 per route table
Networking Service	Network security groups	100 per VCN

Service	Resource Type	Limit
Networking Service	VNICs in network security group	As many VNICs as are in the VCN. A VNIC can belong to max. 5 network security groups
Networking Service	Security rules	50 per network security group
Networking Service	Security lists	20 per VCN 5 per subnet
Networking Service	Ingress rules	30 per security list
Networking Service	Egress rules	30 per security list
Compute Service	Custom images	100
Block Storage Service	Aggregated size of block volumes	100TB (with default storage capacity)
Block Storage Service	Block volume backups	100 (with default storage capacity)
File Storage Service	File systems	100
File Storage Service	Mount targets	100
File Storage Service	File system size	3.3PB
Object Storage Service	Buckets	10000

System Load and Concurrency Limits

This section shows how many concurrent operations of a given type Oracle Private Cloud Appliance can manage at any given time. The limits presented in the table apply across the entire system and all tenancies. For each of these limits it is assumed that no other operations of any kind are running at the same time. When a limit is exceeded, an error with code 409 or 429 is displayed.

Resource Type	Operation	Concurrency Limit
compute instance	back up or restore an instance	10
compute instance	launch/terminate instance	15
compute instance	reset/stop/start instance	15
compute instance	update fault domain (live migration)	10
compute image	create image from instance	10
compute image	import image	10
block volume	create/delete volume	10
block volume	attach/detach boot volume	15
block volume	attach/detach data volume	15
block volume	resize volume	15
file system	create/delete file system	10
mount target	create/delete mount target	10
VCN	create/delete VCN	10

Resource Type	Operation	Concurrency Limit
VCN gateway	create/delete gateway (all types)	10
subnet	create/delete subnet	10
route table	create/delete route table	10
security list	create/delete security list	10
network security group	create/delete network security group	10
VNIC	attach/detach VNIC	15
public IP	create/delete public IP	10
private IP	create/delete private IP	10
all networking resources	update network resource	10

 **Note:**

In addition, there is a system limit on the number of concurrent user sessions:

- Compute Web UI: 15 tenancy users (5 sessions per management node)
- Service Web UI: 9 administrators (3 sessions per management node)

An authentication error is displayed when the limit is reached. An inactive user session times out after 1 hour.

4

Known Issues and Workarounds

This chapter provides information about known issues and workarounds for Oracle Private Cloud Appliance. They are presented in separate sections per category, thus allowing you to navigate more easily.

Platform Issues

This section describes known issues and workarounds related to the appliance platform layer.

Compute Node Provisioning Takes a Long Time

The provisioning of a new compute node typically takes only a few minutes. However, there are several factors that may adversely affect the duration of the process. For example, the management nodes may be under a high load or the platform services involved in the provisioning may be busy or migrating between hosts. Also, if you started provisioning several compute nodes in quick succession, note that these processes are not executed in parallel but one after the other.

Workaround: Unless an error is displayed, you should assume that the compute node provisioning process is still ongoing and will eventually complete. At that point, the compute node provisioning state changes to *Provisioned*.

Bug: 33519372

Version: 3.0.1

Not Authorized to Reconfigure Appliance Network Environment

If you attempt to change the network environment parameters for the rack's external connectivity when you have just completed the initial system setup, your commands are rejected because you are not authorized to make those changes. This is caused by a security feature: the permissions for initial system setup are restricted to only those specific setup operations. Even if you are an administrator with unrestricted access to the Service Enclave, you must disconnect after initial system setup and log back in again to activate all permissions associated with your account.

Workaround: This behavior is expected and was designed to help protect against unauthorized access. In case you need to modify the appliance external network configuration right after the initial system setup, log out and log back in to make sure that your session is launched with the required privileges.

Bug: 33535069

Version: 3.0.1

Grafana Service Statistics Remain at Zero

The Grafana Service Monitoring folder contains a dashboard named Service Level, which displays statistical information about requests received by the fundamental appliance services. These numbers can remain at zero even though there is activity pertaining to the services monitored through this dashboard.

Workaround: No workaround is currently available.

Bug: 33535885

Version: 3.0.1

Terraform Provisioning Requires Fully Qualified Domain Name for Region

If you use the Oracle Cloud Infrastructure Terraform provider to automate infrastructure provisioning on Oracle Private Cloud Appliance, you must specify the fully qualified domain name of the appliance in the region variable for the Terraform provider.

Synchronizing Hardware Data Causes Provisioning Node to Appear Ready to Provision

Both the Service Web UI and the Service CLI provide a command to synchronize the information about hardware components with the actual status as currently registered by the internal hardware management services. However, you should not need to synchronize hardware status under normal circumstances, because status changes are detected and communicated automatically.

Furthermore, if a compute node provisioning operation is in progress when you synchronize hardware data, its Provisioning State could be reverted to *Ready to Provision*. This information is incorrect, and is caused by the hardware synchronization occurring too soon after the provisioning command. In this situation, attempting to provision the compute node again is likely to cause problems.

Workaround: If you have started provisioning a compute node, and its provisioning state reads *Provisioning*, wait at least another five minutes to see if it changes to *Provisioned*. If it takes excessively long for the compute node to be listed as *Provisioned*, run the Sync Hardware Data command.

If the compute node still does not change to *Provisioned*, retry provisioning the compute node.

Bug: 33575736

Version: 3.0.1

Rack Elevation for Storage Controller Not Displayed

In the Service Web UI, the Rack Units list shows all hardware components with basic status information. One of the data fields is *Rack Elevation*, the rack unit number

where the component in question is installed. For one of the controllers of the ZFS storage appliance, `pcasn02`, the rack elevation is shown as *Not Available*.

Workaround: There is no workaround. The underlying hardware administration services currently do not populate this particular data field. The two controllers occupy 2 rack units each and are installed in RU 1-4.

Bug: 33609276

Version: 3.0.1

Fix available: Please apply the latest patches to your system.

Free-Form Tags Used for Extended Functionality

You can use the following free-form tags to extend the functionality of Oracle Private Cloud Appliance.

Note:

Do not use these tag names for other purposes.

- `PCA_no_lm`

Use this tag to instruct the Compute service not to live migrate an instance. The value can be either True or False.

By default, an instance can be live migrated, such as when you need to evacuate all running instances from a compute node. Live migration can be a problem for some instances. For example, live migration is not supported for instances in a Microsoft Windows cluster. To prevent an instance from being live migrated, set this tag to True on the instance.

Specify this tag in the Tagging section of the Create Instance or Edit *instance_name* dialog, in the `oci compute instance launch` or `oci compute instance update` command, or using the API.

The following is an example option for the `oci compute instance launch` command:

```
--freeform-tags '{"PCA_no_lm": "True"}
```

Setting this tag to True on an instance will not prevent the instance from being moved when you change the fault domain. Changing the fault domain is not a live migration. When you change the fault domain of an instance, the instance is stopped, moved, and restarted.

- `PCA_blocksize`

Use this tag to instruct the ZFS storage appliance to create a new volume with a specific block size.

The default block size is 8192 bytes. To specify a different block size, specify the `PCA_blocksize` tag in the Tagging section of the Create Block Volume dialog, in the `oci bv volume create` command, or using the API. Supported values are a power of 2 between 512 and 1M bytes, specified as a string and fully expanded.

The following is an example option for the `oci bv volume create` command:

```
--freeform-tags '{"PCA_blocksize": "65536"}'
```

The block size cannot be modified once the volume has been created.

Use of these tags counts against your tag limit.

Version: 3.0.1

Imported Images Not Synchronized to High-Performance Pool

In an Oracle Private Cloud Appliance with default storage configuration, when you import compute images, they are stored on the ZFS Storage Appliance in an `images` LUN inside the standard ZFS pool. If the storage configuration is later extended with a high-performance disk shelf, an additional high-performance ZFS pool is configured on the ZFS Storage Appliance. Because there is no replication between the storage pools, the images from the original pool are not automatically made available in the new high-performance pool. The images have to be imported manually.

Workaround: When adding high-performance storage shelves to the appliance configuration, import the required compute images again to ensure they are loaded into the newly created ZFS pool.

Bug: 33660897

Version: 3.0.1

API Server Failure After Management Node Reboot

When one of the three management nodes is rebooted, it may occur that the API server does not respond to any requests, even though it can still be reached through the other two management nodes in the cluster. This is likely caused by an ownership issue with the virtual IP shared between the management nodes, or by the DNS server not responding quickly enough to route traffic to the service pods on the available management nodes. After the rebooted management node has rejoined the cluster, it may still take several minutes before the API server returns to its normal operating state and accepts requests again.

Workaround: When a single management node reboots, all the services are eventually restored to their normal operating condition, although their pods may be distributed differently across the management node cluster. If your UI, CLI or API operations fail after a management node reboot, wait 5 to 10 minutes and try again.

Bug: 33191011

Version: 3.0.1

CLI Command Returns Status 500 Due To MySQL Connection Error

When a command is issued from the OCI CLI and accepted by the API server, it starts a series of internal operations involving the microservice pods and the MySQL database, among other components. It may occur that the pod instructed to execute an operation is unable to connect to the MySQL database before the timeout is reached. This exception is reported back to the API server, which in turn reports that the request could not be fulfilled due to an unexpected condition (HTTP status code 500). It is normal for this type of exception to result in a generic server error code. More detailed information may be stored in logs.

Workaround: If a generic status 500 error code is returned after you issued a CLI command, try to execute the command again. If the error was the result of an intermittent connection problem, the command is likely to succeed upon retry.

Bug: n/a

Version: 3.0.1

Microservice Pods Unresponsive After Storage Controller Failover

When a failover or failback occurs between the controllers of the ZFS Storage Appliance, NFS shares mounted on the management nodes or Kubernetes microservices pods can become unresponsive. The higher the number of storage resources and I/O load, the greater the risk that the mounted NFS shares become unresponsive. Services will be impacted when their pods run on the management nodes with unresponsive shares.

Do not attempt to unmount and re-mount any NFS file systems, as this generates stale file handles and makes recovery significantly more difficult. The correct recovery process is to identify which management nodes have unresponsive shares, restart the network connection used for the NFS traffic, and monitor the service pods to make sure they are restored either automatically or manually.

Workaround: The appliance administrators are not permitted to execute the operations required to diagnose and resolve this issue. Please contact Oracle Support to request assistance.

Bug: 33495030

Version: 3.0.1

Administrators in Authorization Group Other Than SuperAdmin Must Use Service CLI to Change Password

Due to high security restrictions, administrators who are not a member of the *SuperAdmin* authorization group are unable to change their account password in the Service Web UI. An authorization error is displayed when an administrator from a non-SuperAdmin authorization group attempts to access their own profile.

Workaround: Log in to the Service CLI, find your user id in the user preferences, and change your password as follows:

```
PCA-ADMIN> show UserPreference
Data:
  Id = 1c74b2a5-c1ce-4433-99da-cb17aab4c090
  Type = UserPreference
[...]
  UserId = id:5b6c1bfa-453c-4682-e692-6f0c91b53d21 type:User name:dcadmin
```

```
PCA-ADMIN> changePassword id=<user_id> password=<new_password>
confirmPassword=<new_password>
```

Bug: 33749967

Version: 3.0.1

Password Policy for Monitoring Framework Differs from Other Infrastructure Components

The password rules for appliance infrastructure components are: minimum length of 8 characters, containing at least 1 lower case letter (a-z), upper case letter (A-Z), digit (0-9), and symbol (@\$!#%*&). It is possible to run a Service CLI command to apply a password that matches these rules but violates the password policy of the appliance's cloud native monitoring and alerting framework, internally referred to as *Sauron*.

When you run the `updateSauronCredentials` command and the monitoring framework rejects the new password, the error message explains the password rules for appliance infrastructure components, not the monitoring framework. In some cases the Service CLI even reports that the password update command was successful although it was not accepted by the monitoring framework.

Workaround: Make sure that the password you set for the appliance monitoring framework complies with its specific password policy: length of 12-20 characters, containing at least 1 lower case letter (a-z), upper case letter (A-Z), digit (0-9), and special character (-_+=.).

Bug: 34816137

Version: 3.0.2

Service Web UI and Grafana Unavailable when HAProxy Is Down

HAProxy is the load balancer used by the Private Cloud Appliance platform layer for all access to and from the microservices. When the load balancer and proxy services are down, the Service Web UI and Grafana monitoring interface are unavailable. When you attempt to log in, you receive an error message: "*Server Did Not Respond*".

Workaround: Log in to one of the management nodes. Check the status of the HAProxy cluster resource, and restart if necessary.

```
# ssh pcamn01
# pcs status
Cluster name: mncluster
Stack: corosync
[...]
Full list of resources:

scsi_fencing (stonith:fence_scsi): Stopped (disabled)
Resource Group: mgmt-rg
  vip-mgmt-int (ocf::heartbeat:IPaddr2): Started pcamn03
  vip-mgmt-host (ocf::heartbeat:IPaddr2): Started pcamn03
  vip-mgmt-ilom (ocf::heartbeat:IPaddr2): Started pcamn03
  vip-mgmt-lb (ocf::heartbeat:IPaddr2): Started pcamn03
  vip-mgmt-ext (ocf::heartbeat:IPaddr2): Started pcamn03
  llapi (systemd:llapi): Started pcamn03
  haproxy (ocf::heartbeat:haproxy): Stopped (disabled)
  pca-node-state (systemd:pca_node_state): Started pcamn03
  dhcp (ocf::heartbeat:dhcpd): Started pcamn03
  hw-monitor (systemd:hw_monitor): Started pcamn03
```


To start HAProxy, use the `pcs resource` command as shown in the example below. Verify that the cluster resource status has changed from "Stopped (disabled)" to "Started".

```
# pcs resource enable haproxy
# pcs status
[...]
Resource Group: mgmt-rg
    haproxy          (ocf::heartbeat:haproxy):      Started pcamm03
```

Bug: 34485377

Version: 3.0.2

Lock File Issue Occurs when Changing Compute Node Passwords

When a command is issued to modify the password for a compute node or ILOM, the system sets a temporary lock on the relevant database to ensure that password changes are applied in a reliable and consistent manner. If the database lock cannot be obtained or released on the first attempt, the system makes several further attempts to complete the request. Under normal operating circumstances, it is expected that the password is eventually successfully changed. However, the command output may contain error messages such as "Failed to create DB lockfile" or "Failed to remove DB lock", even if the final result is "Password successfully changed".

Workaround: The error messages are inaccurate and can be ignored as long as the password operations complete as expected. No workaround is required.

Bug: 34065740

Version: 3.0.2

Compute Node Hangs at Dracut Prompt after System Power Cycle

When an appliance or some of its components need to be powered off, for example to perform maintenance, there is always a minimal risk that a step in the complex reboot sequence is not completed successfully. When a compute node reboots after a system power cycle, it can hang at the `dracut` prompt because the boot framework fails to build the required `initramfs/initrd` image. As a result, primary GPT partition errors are reported for the root file system.

Workaround: Log on to the compute node ILOM. Verify that the server has failed to boot, and is in the `dracut` recovery shell. To allow the compute node to return to normal operation, reset it from the ILOM using the `reset /System` command.

Bug: 34096073

Version: 3.0.2

No Error Reported for Unavailable Spine Switch

When a spine switch goes offline due to loss of power or a fatal error, the system gives no indication of the issue in the Service Enclave UI/CLI or Grafana. This behavior is the result of the switch client not properly handling exceptions and continuing to report the default "healthy" status.

Workaround: There is currently no workaround to make the system generate an error that alerts the administrator of a spine switch issue.

Bug: 34696315

Version: 3.0.2

ZFS Storage Appliance Controller Stuck in Failsafe Shell After Power Cycle

The two controllers of the Oracle ZFS Storage Appliance operate in an active-active cluster configuration. When one controller is taken offline, for example when its firmware is upgraded or when maintenance is required, the other controller takes ownership of all storage resources to provide continuation of service. During this process, several locks must be applied and released. When the rebooted controller rejoins the cluster to take back ownership of its assigned storage resources, the cluster synchronization will fail if the necessary locks are not released correctly. In this situation, the rebooted controller could become stuck in the failsafe shell, waiting for the peer controller to release certain locks. This is likely the result of a takeover operation that was not completed entirely, leaving the cluster in an indeterminate state.

Workaround: There is currently no workaround. If the storage controller cluster ends up in this condition, contact Oracle for assistance.

Bug: 34700405

Version: 3.0.2

First Tenancy Creation Fails and Successive Attempts Report It Already Exists

After completing the initial installation process, you create the first tenancy in your cloud environment. If the microservice pod, responsible for executing the tenancy creation command, crashes between the start and completion of the relevant operation, the creation process fails. However, when you try to create the tenancy again, the system returns an exception indicating the tenancy already exists. This is because part of the tenancy configuration was stored before the pod crashed on the first attempt.

Workaround: This problem is typically caused by bad timing. When you complete the initial installation process, the microservices pods need to update their region and realm data and restart with the new values. This takes approximately 3 minutes. Wait until the pods have been restarted before creating the first tenancy on your system.

Bug: 34743353

Version: 3.0.2

Concurrent Compute Node Provisioning Operations Fail Due to Storage Configuration Timeout

When the Private Cloud Appliance has just been installed, or when a set of expansion compute nodes have been added, the system does not prevent you from provisioning all new compute nodes at once. Note, however, that for each provisioned node the storage initiators and targets must be configured on the ZFS Storage Appliance. If there are too many configuration update requests for the storage appliance to process,

they will time out. As a result, all compute node provisioning operations will fail and be rolled back to the unprovisioned state.

Workaround: To avoid ZFS Storage Appliance configuration timeouts, provision compute nodes sequentially one by one, or in groups of no more than 3.

Bug: 34739702

Version: 3.0.2

Switch Fault Reported After Password Change

When you execute the command to update the password of the switches, there is a brief delay before the change is applied to the component. If a health check runs during this delay, it detects an authentication error and changes the component run state to "fail". This, in turn, results in an active fault, which appears in the Service Enclave UI and CLI and can trigger an alert. Once the password update is completed successfully and the health check passes, the fault is cleared.

Workaround: Keep monitoring the state of the switch. If the fault was caused by the password change then it will be cleared automatically once the new password is accepted.

```
PCA-ADMIN> list fault
  id
name                               status  severity
--
----
6418d404-7702-41e8-de2a-83dcc69d08a8
RackUnitRunStateFaultStatusFault (pcaswsp01)  Active  Critical
90611c79-cf3f-4012-94c4-db709c97e390
RackUnitRunStateFaultStatusFault (pcaswsp02)  Active  Critical

PCA-ADMIN> show fault id=90611c79-cf3f-4012-94c4-db709c97e390
Data:
  Type = Fault
  Category = Status
  Severity = Critical
  Status = Active
  Associated Attribute = runStateFault
  Last Update Time = 2022-11-17 21:30:50,251 UTC
  Cause = RackUnit pcaswsp02 attribute 'runStateFault' = 'CRITICAL'.
[...]
```

Bug: 34850833

Version: 3.0.2

Kubernetes Health Checker Cannot Find Endpoints

After upgrade to Release 3.0.2 the Kubernetes health checker can end up in an unhealthy state and unable to connect to the necessary endpoints. The issue is caused by the *prometheus container* inside the Prometheus Kubernetes pod running out of memory ("*OOMKilled*") and restarting over and over.

```
# kubectl describe pod prometheus-k8s-0 -n monitoring
[...]
```

```
prometheus:
  Container ID:
  cri-o://5f1ef5176f4c1124085662556e94008e46c9090d35a2894e8187f9b16df14fbe
```

```
Image:          253.255.0.31:5000/sauron/prometheus-mirror:2.20.0-1
Image ID:
253.255.0.31:5000/sauron/prometheus-mirror@sha256:9a2552ae5a53ad11b83cdd91463f
d83ae8826620f31ee8c3eefc6ef99c1da0ab
Port:          9090/TCP
Host Port:     0/TCP
[...]
State:        Waiting
Reason:       CrashLoopBackOff
Last State:   Terminated
Reason:       OOMKilled
```

The container memory limit has been increased to prevent the problem, but these settings only take effect after the Prometheus Kubernetes pod is restarted. Based on the latest Helm chart, the pod will restart with the appropriate container memory limit applied.

```
# kubectl delete pod prometheus-k8s-0 -n monitoring
pod "prometheus-k8s-0" deleted
```

Workaround: The appliance administrators are not permitted to execute the operations required to diagnose and resolve this issue. Please contact Oracle Support to request assistance.

Bug: 34858888

Version: 3.0.2

Leaf Switch Stuck in Password Change

When changing the password of the leaf switches, the process could hang during the password update on the switch itself. The switch properties and logs indicate this by showing the switch provisioningState parameter as "Update_Switch_Password". Because the password change on the switch may have a considerable lag, the operation may still be completed successfully within 10-30 minutes after the command was executed. If the situation is not resolved automatically within that time frame, the switch can be considered stuck in this state and out of sync with the platform.

Workaround: If the switch is no longer in sync with the appliance platform, the password must be applied manually in the switch operating software.

```
# config
(config)# username admin password <password>
(config)# end
# copy running-config startup-config
```

Bug: 34819799

Version: 3.0.2

Data Switch Fails to Boot Due to Active Console Connection

If a Cisco Nexus 9336C-FX2 Switch has an active local console session, for example when a terminal server is connected, the switch could randomly hang during reboot. It is assumed that the interruption of the boot sequence is caused by a ghost session on the console port. This behavior has not been observed when no local console connection is used.

Workaround: Do not connect any cables to the console ports of the data switches. There is no need for a local console connection in a Private Cloud Appliance installation.

Bug: 32965120

Version: 3.0.2

User Interface Issues

This section describes known issues and workarounds related to the graphical user interface.

Moving Resources Between Compartments Is Not Supported in the Compute Web UI

The Compute Web UI does not provide any function to move a resource from one compartment to another. Operations to change the compartment where a cloud resource resides, can only be performed through the CLI. However, note that not all resource types support compartment changes. For example, none of the network resources can be moved.

Workaround: If you need to move a resource from its current compartment to another compartment, use the CLI. After successfully executing the CLI command, you can see the resulting changes in the Compute Web UI.

Bug: 33038606

Version: 3.0.1

No Available Compute Web UI Operation to Update Instance Pool

The Compute Web UI does not provide the functionality to update the properties of an existing instance pool. There is no user interface implementation of the `UpdateInstancePool` and `UpdateInstancePoolDetails` API resources.

Workaround: Update the instance pool through the CLI. Use this command:

```
oci compute-management instance-pool update [OPTIONS]
```

Bug: 33393214

Version: 3.0.1

Saving Resource Properties Without Modifications Briefly Changes Status to Provisioning

If you open the Edit dialog box in the Compute Web UI to modify the properties of a resource, and you click Save Changes without actually modifying any of the properties, the status of the resource does change to *Provisioning* for a few seconds. This is the normal response of the UI to a user clicking the Save button. It has no adverse effect.

Workaround: To prevent the resource status from changing to *Provisioning* if you have not made any changes to it, click the Cancel button in the dialog box instead.

Bug: 33445209

Version: 3.0.1

NFS Export Squash ID Not Displayed

In the Compute Web UI, the detail page of an NFS export does not display the squash ID in the NFS export options. The squash ID is required for anonymous access to the NFS export, but you can only retrieve it by editing the export options.

Workaround: To obtain an NFS export squash ID, go to the NFS Export detail page, scroll down to the NFS Export Options, and click Edit Options. Alternatively, look up the export options through the CLI.

Bug: 33480572

Version: 3.0.1

Scrollbars Not Visible in Browser

The browser-based interfaces of Private Cloud Appliance are built with Oracle JavaScript Extension Toolkit(JET) and follow Oracle's corporate design guidelines. Scrollbars are meant to remain hidden as long as you are not actively using the part of the screen where content does not fit within the space provided – for example: large tables, long drop-down lists, and so on. Not all browsers or browser versions display scrollbars in the intended way. For example, Google Chrome typically hides the scrollbars as intended while Mozilla Firefox does not hide them at all.

Workaround: The behavior of the scrollbars is by design. It applies to both the Compute Web UI and Service Web UI. In areas where content runs beyond the allocated screen area, scrollbars appear automatically where appropriate when the cursor is placed over the content in question.

Bug: 33489195

Version: 3.0.1

Authorization Failure When Retrieving Compartment Data

The Identity and Access Management service allows you to control users' access permissions to resources in a fine-grained way through policies. Those policies determine which operations a group of users is authorized to perform on resources of a particular type or residing in a particular compartment. In certain situations, the Compute Web UI is unable to hide all the resources that a user has no access to. Consequently, a user operation may result in a request for data the account is not authorized to access.

While using the Compute Web UI you may run into authorization failures in case your operation triggers an attempt to retrieve data that you have no permission for. In this situation an error appears in your browser, indicating that the application has stopped working due to account permissions. The compartment tree, in particular, is prone to this type of failure because it can display compartments that you are not allowed to access.

Workaround: When the error is caused by a compartment tree access issue, it is likely that the intended page is displayed when you click Try Again. Otherwise, contact your tenancy administrator to request additional permissions to access the required data.

Bug: 33497526, 33520207

Version: 3.0.1

Object List Is Not Updated Automatically

In the Compute Web UI you display the objects stored in a bucket by browsing through its directory structure. The list or table of objects is not automatically refreshed at regular intervals, so any object changes will only become visible when you refresh the page manually. There is no available function for the UI to poll the status of a bucket.

Workaround: To display the current list of objects in the Compute Web UI, refresh the page manually. This behavior is not specific to the object storage service; it may occur in other areas of the UI as well. If a resource list is not updated automatically at regular intervals, you should refresh it manually.

Bug: 33519215

Version: 3.0.1

File Storage Mount Target Link Not Available

In the File Storage area of the Compute Web UI, you can find a mount target URL as follows: display the list of mount targets, select the mount target you are interested in, click the export to display its detail page, and locate the Mount Target field. However, it may occur that the mount target is shown as Not Available even though it does exist.

Workaround: The UI can not always reliably retrieve resource details from the file system storage service. Refresh the page or navigate away and back in order to force the UI to retrieve the details again.

Bug: 33571007

Version: 3.0.1

UDP Ports Not Displayed In Security List Rules Table

Ingress and egress rules belonging to a security list are displayed in a table in the Resources section of the Security List detail page. When you create rules related to UDP ports, the port numbers are not displayed in the Port Range columns. The UDP settings are not lost; you can view them in the Edit window.

Workaround: The Edit Security List window does display all relevant settings, including the UDP ports. In the Actions menu (on the right edge of the row) select Edit as you would to modify an ingress or egress rule.

Bug: 33575269

Version: 3.0.1

Not All Resources Shown in Drop-Down List

When you need to use a drop-down list to select a resource, you may notice that not all items are shown if the list is very long. As you scroll through the list, more items are loaded, yet you may still be unable to find the item you are looking for. The reason for this behavior is that UI components are designed to respond quickly rather than slowing down the user due to long load times. You are encouraged to filter a long list by typing part of a resource name in the text field, instead of scrolling through a complete alphabetical list. This is characteristic of

Oracle JavaScript Extension Toolkit, so it affects both the Compute Web UI and the Service Web UI

Workaround: Scrolling is not the preferred way to search for an item in a long drop-down list. Instead, start typing the name of the resource you are looking for, and the available list items will be reduced to those matching what you type.

Bug: 33583708

Version: 3.0.1

Volume Group Can Be Created Without Name

When you create a volume group in the Block Storage area of the Compute Web UI, you are not required to enter a name. If you leave the name field blank, the volume group appears in the list as *Unnamed Item*. However, if you do not provide a name when creating a volume group in the CLI, a name is automatically assigned based on the time of creation.

Workaround: This is not a code bug: the name is technically not a required parameter. To avoid having volume groups with meaningless names, make sure you provide an appropriate name at the time of creation, in the Compute Web UI as well as the CLI. If you accidentally created the volume group without specifying a name, you can edit the volume group afterwards and add the name of your choice.

Bug: 33608462

Version: 3.0.1

File Systems and Mount Targets Not Displayed

When users have access to the resources in a particular compartment, but have no permission to view the content of the root compartment of the tenancy, the Compute Web UI might not display the resources that a user is allowed to list. Instead, an authorization error is displayed. For the file system service specifically, file systems or mount targets in a particular compartment are not displayed, even if the user has full access to that compartment and the resources it contains.

This behavior is caused by the way the API request is made from the UI, using the OCID of the root compartment. In contrast, the CLI requires that you specify the OCID of the compartment that effectively contains the requested resources, so it is not affected by the same authorization issue as the UI.

Workaround: The tenancy administrator should make sure that users of the file system service have read access to the root compartment. Users who cannot list the file systems and mount targets they are authorized to use, should ask their tenancy administrator to verify their account permissions and make the necessary adjustments.

Bug: 33666365

Version: 3.0.1

Optional ICMP Security Rule Parameters Cannot Be Removed

When you add an ingress or egress security rule to the security list of a VCN, you can specifically select the ICMP protocol. The Compute Web UI indicates that selecting a

Parameter Type and *Parameter Code* from the respective lists is optional. This is incorrect, because the *Parameter Type* is mandatory for ICMP rules.

If you specified both *Type* and *Code* in your ICMP rule, it is possible to remove the *Parameter Code*. Edit the security rule, place your cursor in the *Code* text field, and delete its content. This is how drop-down lists work in the UI; there is no "empty" option to select.

Workaround: When working with ICMP security rules, always specify the *Parameter Type*. To remove an optional parameter selected from a drop-down list, select and delete the content of the text field.

Bug: 33631794

Version: 3.0.1

Compartment Selector Not Available When Creating DHCP Options

When you create or modify DHCP options for a VCN through the Compute Web UI, there is no way to add the DHCP options to another compartment. Because the compartment selector is not available in the create and edit windows, the DHCP options are implicitly stored in the same compartment as the VCN itself. However, it is supported to store DHCP options in another compartment. If you wish to do so, please use the CLI.

Workaround: If you want DHCP options to be stored in a different compartment than the VCN they apply to, create the DHCP options through the CLI, or use the CLI to move them to the desired compartment.

Bug: 33722013

Version: 3.0.1

Fix available: Please apply the latest patches to your system.

Custom Search Domain Error Not Rolled Back When Operation Is Canceled

The Networking service allows you to control certain instance boot configuration parameters by setting DHCP options at the level of the VCN and subnet. One of the DHCP options you can control is the search domain, which is appended automatically to the instance FQDN in a DNS-enabled VCN and subnet.

The search domain must be specified in the format `example.tld` – where TLD stands for top-level domain. If you attempt to save an invalid search domain, two error messages are displayed: one appears immediately under the Search Domain field, and the other at the bottom of the DHCP Options window, indicating that "Some fields are incomplete or invalid". When you cancel the creation or modification of the DHCP options and subsequently reopen the DHCP Options window, the error messages and the incorrect value may still be displayed, even though the settings were not applied.

Workaround: If you close and reopen the DHCP Options window after a failed attempt to save new settings, and the same error messages and invalid value still appear, you may ignore the error. Refreshing the browser window may clear the error message. Enter a custom search domain in the required format: `example.tld`

Bug: 33734400

Version: 3.0.1

DHCP Options Error Message for Custom Search Domain Is Misleading

The Networking service allows you to control certain instance boot configuration parameters by setting DHCP options at the level of the VCN and subnet. One of the DHCP options you can control is the search domain, which is appended automatically to the instance FQDN in a DNS-enabled VCN and subnet.

The search domain must be specified in the format `example.tld` – where TLD stands for top-level domain. However, the Compute Web UI does not validate this parameter; this is done by the Networking service when you save the DHCP options. The Compute Web UI checks that the value contains no spaces. If it does, an error message appears under the Search Domain field: *"Must be in the format of example.tld"*. This is technically inaccurate as it merely indicates the value contains a space.

Workaround: Enter a custom search domain in the required format: `example.tld`. Spaces are not allowed in domain names. If the error message in question appears, correct the value you entered in the Search Domain field and try to save the DHCP options again.

Bug: 33753758

Version: 3.0.1

Unclear Error when Logging in to Service Web UI with Insufficient Privileges

In the Service Enclave, an administrator who is a member of the *SuperAdmin* authorization group can create other administrator accounts and determine to which authorization group those accounts belong. It is possible to set access restrictions that prevent certain administrators from logging in to the Service Web UI, even though they can still access the Service CLI.

If this situation occurs, the Service Web UI returns an error message that does not clearly describe the authorization problem – for example:

```
{"message":"AUTH_000008: An error occurred getting system config state.",
"errorCode":"AUTH_SERVICE_GET_CONFIG_STATE_ERROR",
"cause":["Caused by:
com.oracle.pca.ui.common.server.exception.PcaConsoleException:
SERVICE_000002: Calling underlying service resulted in an error:
Failed to get PcaSystem object from admin service [Bad
Request]."],"csrfToken":null,"idps":null,"serviceError":null}
```

Workaround: If you receive an error similar to this example, and you should be able to access the Service Web UI, ask an appliance administrator with the appropriate permissions to correct the access restrictions for your administrator account.

Bug: 34522989

Version: 3.0.2

Cloning Snapshot to New File System Not Supported from Compute Web UI

It is possible to clone a snapshot of a file system and use that clone as a new file system. However, this functionality is not provided by the Compute Web UI.

Workaround: To clone a file system, use the OCI CLI. Assuming you already have a file system snapshot, this is the command syntax to clone it to a new file system:

```
oci fs file-system create
--availability-domain <availability_domain_name>
--compartment-id <compartment_OCID>
--display-name <fs_display_name>
--source-snapshot-id <snapshot_OCID>
```

Bug: 34566069

Version: 3.0.2

No Details Displayed for File System Cloned from Snapshot

When you create a file system, all relevant information about the file system is displayed in the Compute Web UI detail page. It contains practically the same information as the output from the OCI CLI command `oci fs file-system get`. However, when you clone a file system snapshot through the OCI CLI, to use as a new file system, the detail page of the clone file system does not provide the same data. Relevant missing fields include Source Snapshot, Parent File System, Clone Root, Hydration, etc.

Workaround: Use the OCI CLI instead if you need those details. The CLI displays all the data fields for a clone-based file system.

Bug: 34566735

Version: 3.0.2

Network Environment Information Is Not Refreshed Automatically After Admin Network Is Enabled

In the Network Environment area of the Service Web UI, when you edit the appliance network configuration to enable the Admin Network for segregated administrative connectivity, the Network Environment Information page is not refreshed when you save your configuration changes.

Workaround: To view the Admin Network parameters you just applied, you need to manually refresh the Network Environment Information page in your browser.

Bug: 34769734

Version: 3.0.2

Unable to Display Details of Instance Backup

When you create a backup of an instance, it is stored in an object storage bucket. When the backup operation has completed, and you try to view the details of the backup object, the system may return an HTTP 404 error.

Workaround: Manually reload the browser page. The backup details should be displayed. Additional error messages might appear in a pop-up but those can be ignored.

Bug: 34777856

Version: 3.0.2

IP Address List on VNIC Detail Page Not Updated

In the Compute Web UI you can manage the IP addresses assigned to a compute instance from the detail pages of the instance's attached VNICs. In the IP Addresses table in the Resources section of the VNIC detail page you can add and remove private and public IPs. However, the Reserve Public IP pop-up window is not always rendered correctly, and changes applied there are not displayed in the VNIC IP Address table.

Workaround: No workaround is available.

Bug: 34797160

Version: 3.0.2

No Error Displayed When Attempting to Reserve Public IP Address While All Public IPs In Use

In the Compute Web UI, when you try to reserve a public IP address when no more public IPs are available from the pool, no error message is displayed. The Reserve Public IP window hangs, but provides no feedback about the operation you are attempting. Since no IPs are available to reserve, the operation does not complete.

Workaround: Close the Reserve Public IP window or reload the browser interface page. An administrator needs to expand the public IP pool for the Private Cloud Appliance environment before you can reserve another public IP address.

Bug: 34832457

Version: 3.0.2

Incorrect Error Message Displayed When Deleting Exadata Network

When you delete an Exadata network using the Service Web UI, an error message may report that the operation has failed, even though the network was successfully deleted.

Workaround: Do not rely on the error message. Verify that the Exadata network has been deleted and no longer appears in the Exadata Networks list.

Bug: 34680351

Version: 3.0.2

When Creating Volume Group from Backup Its Volumes Are Not Displayed

You can create a backup of a set of volumes contained in a volume group, and subsequently create a new volume group from that volume group backup. The new volume group then contains the same volumes that were part of the backup. However, despite the operation being completed successfully, the individual volumes that are part of the new volume group are not displayed in the volume group detail page in the Compute Web UI.

Workaround: Use the OCI CLI to confirm that the volumes are present, and to perform any required operations on the volumes.

Bug: 33997759

Version: 3.0.2

Networking Issues

This section describes known issues and workarounds related to all aspects of the appliance networking, meaning the system's internal connectivity, the external uplinks to the data center, and the virtual networking for users' compute instances.

DNS Zone Scope Cannot Be Set

When creating or updating a DNS zone, scope cannot be set. In command line output, the value of the `scope` property is `null`.

Bug: 32998565

Version: 3.0.1

To Update a DNS Record the Command Must Include Existing Protected Records

When updating a DNS record, it is expected that you include all existing protected records in the update command even if your update does not affect those. This requirement is intended to prevent the existing protected records from being inadvertently deleted. However, the checks are so restrictive with regard to SOA records that certain updates are difficult to achieve.

Workaround: It is possible to update existing records by either providing the SOA record as part of the command, or by setting the domain to not include the SOA domain. In practice, most record updates occur at a higher level and are not affected by these restrictions.

Bug: 33089111

Version: 3.0.1

Fix available: Please apply the latest patches to your system.

Oracle Linux 8 Instance Host Name Resolution Fails

When you try to interact with an Oracle Linux 8 compute instance using only its host name, the DNS server cannot fulfill the request and name resolution fails. When you use the

instance's fully qualified domain name (FQDN), it is resolved as expected. This problem is caused by an incomplete configuration of the DNS domain at the level of the VCN: there is no NS or SOA record, which are both minimum requirements for a working DNS zone.

Workaround: To access Oracle Linux 8 instances, always use the FQDN. Oracle Linux 7 instances are not affected by this issue: DNS requests using only their host name are processed correctly.

Bug: 34734562

Version: 3.0.1

Create Route Table Fails With Confusing Error Message

When you create a route table, but make a mistake in the route rule parameters, the API server may return an error message that is misleading. That specific message reads: *"Route table target should be one of LPG, NAT gateway, Internet gateway, DRG attachment or Service gateway."* In that list of possible targets, DRG attachment is not correct. The dynamic routing gateway itself should be specified as a target, not its DRG attachment.

Workaround: Ignore the error message in question. When configuring route rules to send traffic through a dynamic routing gateway, specify the DRG as the target.

Bug: 33570320

Version: 3.0.1

VCN Creation Uses Deprecated Parameter

When creating a VCN, you typically specify the CIDR range it covers. In the Compute Web UI, you simply enter this in the applicable field. However, the CLI provides two command parameters: `--cidr-block`, which is now deprecated, and `--cidr-blocks`, which is a new parameter that is meant to replace the deprecated one. When using the OCI CLI with Private Cloud Appliance you must use `--cidr-block`. The new parameter is not supported by the API server.

Workaround: Ignore any warning messages about the deprecated parameter. Use the `--cidr-block` parameter when specifying the CIDR range used by a VCN.

Bug: 33620672

Version: 3.0.1

File Storage Traffic Blocked By Security Rules

To allow users to mount file systems on their instances, security rules must be configured in addition to those in the default security list, in order to allow the necessary network traffic between mount targets and instances. Configuring file storage ports and protocols in Oracle Private Cloud Appliance is further complicated by the underlay network architecture, which can block file storage traffic unexpectedly unless the source and destination of security rules are set up in a very specific way.

Scenario A – If the mount target and instances using the file system service reside in the same subnet, create a security list and attach it to the subnet in addition to the default security list. The new security list must contain the following stateful rules:

```

+++ Ingress Rules ++++++

Source          Protocol    Source Ports    Destination Ports
-----
<subnet CIDR>  TCP        All             111, 389, 445, 4045,
                                     2048-2050, 20048
<subnet CIDR>  UDP        All             111, 289, 445, 2048,
                                     4045, 20048

+++ Egress Rules ++++++

Destination     Protocol    Source Ports    Destination Ports
-----
<subnet CIDR>  TCP        111, 389, 445, 4045,
                                     2048-2050, 20048    All
<subnet CIDR>  TCP        All             111, 389, 445, 4045,
                                     2048-2050, 20048
<subnet CIDR>  UDP        111, 389, 445,
                                     4045, 20048        All
<subnet CIDR>  UDP        All             111, 389, 445,
                                     4045, 20048

```

Scenario B – If the mount target and instances using the file system service reside in different subnets, create a new security list for each subnet, and attach them to the respective subnet in addition to the default security list.

The new security list for the subnet containing the mount target must contain the following stateful rules:

```

+++ Ingress Rules ++++++

Source          Protocol    Source Ports    Destination Ports
-----
<instances subnet CIDR>  TCP        All             111, 389, 445, 4045,
                                     2048-2050, 20048
<instances subnet CIDR>  UDP        All             111, 289, 445, 2048,
                                     4045, 20048

+++ Egress Rules ++++++

Destination     Protocol    Source Ports    Destination Ports
-----
<instances subnet CIDR>  TCP        111, 389, 445, 4045,
                                     2048-2050, 20048    All
<instances subnet CIDR>  UDP        111, 389, 445,
                                     4045, 20048        All

```

The new security list for the subnet containing the instances using the file system service must contain the following stateful rules:

```

+++ Ingress Rules ++++++

Source          Protocol    Source Ports    Destination Ports
-----
<mount target subnet CIDR>  TCP        111, 389, 445, 4045,
                                     2048-2050, 20048    All
<mount target subnet CIDR>  UDP        111, 289, 445, 2048,
                                     4045, 20048        All

+++ Egress Rules ++++++

```

Destination Ports	Protocol	Source Ports	Destination
-----	-----	-----	
<mount target subnet CIDR> 445, 4045, 20048	TCP	All	111, 389, 2048-2050,
<mount target subnet CIDR>	UDP	All	111, 389, 445, 4045, 20048

Workaround: Follow the guidelines provided here to configure ingress and egress rules that enable file system service traffic. If the unmodified default security list is already attached, the proposed egress rules do not need to be added, because there already is a default stateful security rule that allows all egress traffic (destination: 0.0.0.0/0, protocol: all).

Bug: 33680750

Version: 3.0.1

Stateful and Stateless Security Rules Cannot Be Combined

The appliance allows you to configure a combination of stateful and stateless security rules in your tenancy. The access control lists generated from those security rules are correct, but may cause a wrong interpretation in the virtual underlay network. As a result, certain traffic may be blocked or allowed inadvertently. Therefore, it is recommended to use either stateful or stateless security rules.

Workaround: This behavior is expected; it is not considered a bug. Whenever possible, create security rules that are either all stateful or all stateless.

Note:

If you have a specific need, you can have stateful and stateless rules combined, but if you use stateless rules they must be symmetrical, meaning you cannot have a stateless egress rule, and a stateful ingress rule for the same flow.

Bug: 33744232

Version: 3.0.1

VCN With Single Subnet of Same Size Not Supported

When you create a VCN, you assign it a CIDR range with a maximum size of "/16" – for example: 10.100.0.0/16, or 172.16.64.0/18. The Networking service expects that you create subnets within a VCN, but rather than subdividing it into several smaller subnets you might prefer to use a single large subnet. However, it is a requirement for a subnet to be smaller than the VCN it belongs to.

Workaround: If you intend to use one large subnet within your VCN, make sure that the VCN is set up with a CIDR that is larger than the IP address range you need for the subnet.

Bug: 33758108

Version: 3.0.1

Fix available: Please apply the latest patches to your system.

Routing Failure With Public IPs Configured as CIDR During System Initialization

When you complete the initial setup procedure on the appliance (see "Complete the Initial Setup" in the chapter [Configuring Oracle Private Cloud Appliance](#) of the Oracle Private Cloud Appliance Installation Guide), one of the final steps is to define the data center IP addresses that will be assigned as public IPs to your cloud resources. If you selected BGP-based dynamic routing, the public IPs may not be advertised correctly when defined as one or more CIDRs, and thus may not be reachable from outside the appliance.

Workaround: To ensure that your cloud resources' public IPs can be reached from outside the appliance, specify all IP addresses individually with a /32 netmask. For example, instead of entering 192.168.100.0/24, submit a comma-separated list: 192.168.100.1/32,192.168.100.2/32,192.168.100.3/32,192.168.100.4/32, and so on.

Bug: 33765256

Version: 3.0.1

Fix available: Please apply the latest patches to your system.

Admin Network Cannot Be Used for Service Web UI Access

The purpose of the (optional) Administration network is to provide system administrators separate access to the Service Web UI. The current implementation of the Administration network is incomplete and cannot provide the correct access.

Workaround: None available. At this point, *do not* configure the Admin Network during initial configuration.

Bug: 34087174, 34038203

Version: 3.0.1

Network Configuration Fails During Initial Installation Procedure

After physical installation of the appliance rack, the system must be initialized and integrated into your data center environment before it is ready for use. This procedure is documented in the chapter titled "[Configuring Oracle Private Cloud Appliance](#)" of the Oracle Private Cloud Appliance Installation Guide. If the network configuration part of this procedure fails – for example due to issues with message transport or service pods, or errors returned by the switches – there are locks in place that need to be rolled back manually before the operation can be retried.

Workaround: None available. Please contact Oracle for assistance.

If possible, confirm the state of the network configuration from the Service CLI.

```
PCA-ADMIN> show
networkConfig
```

```
Data:  
[...]  
  Network Config Lifecycle State = FAILED
```

Bug: 34788596

Version: 3.0.2

External Certificates Not Allowed

At this time, Oracle Private Cloud Appliance does not allow the use of external CA-signed certificates.

Workaround: Please contact Oracle support for a workaround.

Bug: 33025681

Version: 3.0.2

Low Uplink MTU Setting Prevents Browser Interface Loading

When you set the maximum transmission unit (MTU) for the uplink ports to an unusually low size, the network performance between the Private Cloud Appliance and the data center environment will be degraded, and connectivity can be disrupted. For example, while a segregated administration network might require a lower MTU, a symptom of an excessively small packet size is that the Service Web UI cannot be loaded. However, the system does not enforce a particular operating range, so no errors are returned when you configure an unusual MTU.

Workaround: The default uplink MTU is 9216 bytes. If your data center setup requires a different setting, please use a common packet size, such as 1500 bytes.

Bug: 34841495

Version: 3.0.2

VCNs Stuck in Provisioning After Concurrent Creation Operations

When multiple VCNs are created concurrently, some may become stuck in provisioning state, meaning they never become available. This is caused by a rare race condition where, for a short amount of time, an object life cycle state can be updated by different threads.

Workaround: After concurrently creating a series of VCNs, if one of the VCN objects ends up in "PROVISIONINGFAIL" state, use the interface of your choice (UI, CLI, Terraform) to delete the stuck resource and create it again.

Bug: 34841815

Version: 3.0.2

Admin Network Configuration Blocks Access to Services

In a Private Cloud Appliance environment with administration network enabled, all administrative traffic is strictly segregated from the data network. As a consequence, peer-to-peer connectivity between compute instances and the appliance's administrative services is not allowed.

Workaround: To access administration functionality, please use a host external to your Private Cloud Appliance.

Bug: 34849776

Version: 3.0.2

After Upgrade to Release 3.0.2 the Metadata of Existing Instances Contains No FQDN

For host name resolution of compute instances the DNS configuration relies on certain parameters stored in the instance metadata. Improvements have been made in the software of Release 3.0.2 to make name resolution more reliable. However, instances created before the software upgrade do not take advantage of the metadata changes: the relevant fields are not updated automatically with the fully qualified domain name (FQDN).

Workaround: There is currently no workaround for existing instances. The metadata of instances created after the software upgrade contains the correct values.

Bug: 34859081

Version: 3.0.2

DNS Entries on Oracle Linux 8 Instances Incorrect After Upgrade to Release 3.0.2

After the appliance software is upgrade to Release 3.0.2, the name resolution settings in the compute instance operating system are not automatically updated. Up-to-date network parameters are obtained when the instance's DHCP leases are renewed. Until then, due to the way Oracle Linux 8 responds to DNS server messages, it can fail to resolve short host names although queries with FQDNs are successful. Oracle Linux 7 instances are not affected by this issue.

Workaround: Restart the DHCP client service (`dhclient`) on the command line of your Oracle Linux 8 instances. Rebooting the instance also resolves the issue.

Bug: 34918899

Version: 3.0.2

Compute Service Issues

This section describes known issues and workarounds related to the compute service.

No Consistent Device Paths for Connecting to Block Volumes

When you attach a block volume to an instance, it is not possible to specify a device path that remains consistent between instance reboots. It means that for the `attach-paravirtualized-volume` CLI command the optional `--device` parameter does not work. Because the device name might be different after the instance is rebooted, this affects tasks you perform on the volume, such as partitioning, creating and mounting file systems, and so on.

Workaround: No workaround is available.

Bug: 32561299

Version: 3.0.1

Instance Pools Cannot Be Terminated While Starting or Scaling

While the instances in a pool are being started, and while a scaling operation is in progress to increase or decrease the number of instances in the pool, it is not possible to terminate the instance pool. Individual instances, in contrast, can be terminated at any time.

Workaround: To terminate an instance pool, wait until all instances have started or scaling operations have been completed. Then you can successfully terminate the instance pool as a whole.

Bug: 33038853

Version: 3.0.1

Network Interface on Windows Does Not Accept MTU Setting from DHCP Server

When an instance is launched, it requests an IP address through DHCP. The response from the DHCP server includes the instruction to set the VNIC maximum transmission unit (MTU) to 9000 bytes. However, Windows instances boot with an MTU of 1500 bytes instead, which may adversely affect network performance.

Workaround: When the instance has been assigned its initial IP address by the DHCP server, change the interface MTU manually to the appropriate value, which is typically 9000 bytes for an instance's primary VNIC. This new value is persistent across network disconnections and DHCP lease renewals.

Alternatively, if the Windows image contains `cloudbase-init` with the `MTUPlugin`, it is possible to set the interface MTU from DHCP. To enable this function, execute the following steps:

1. Edit the file `C:\Program Files\Cloudbase Solutions\Cloudbase-Init\conf\cloudbase-init.conf`. Add these lines:

```
mtu_use_dhcp_config=true  
plugins=cloudbaseinit.plugins.common.mtu.MTUPlugin
```
2. Enter the command `Restart-Service cloudbase-init`.
3. Confirm that the MTU setting has changed. Use this command: `netsh interface ipv4 show subinterfaces`.

Bug: 33541806

Version: 3.0.1

Oracle Solaris Instance in Maintenance Mode After Restoring from Backup

It is supported to create a new instance from a backup of the boot volume of an existing instance. The existing instance may be running or stopped. However, if you

use a boot volume backup of an instance based on the Oracle Solaris image provided with Private Cloud Appliance, the new instance created from that backup boots in maintenance mode. The Oracle Solaris console displays this message: *"Enter user name for system maintenance (control-d to bypass)."*

Workaround: When the new Oracle Solaris instance created from the block volume backup has come up in maintenance mode, reboot the instance from the Compute Web UI or the CLI. After this reboot, the instance is expected to return to a normal running state and be reachable through its network interfaces.

Bug: 33581118

Version: 3.0.1

Instance Disk Activity Not Shown in Compute Node Metrics

The virtual disks attached to compute instances are presented to the guest through the hypervisor on the host compute node. Consequently, disk I/O from the instances should be detected at the level of the physical host, and reflected in the compute node disk statistics in Grafana. Unfortunately, the activity on the virtual disks is not aggregated into the compute node disk metrics.

Workaround: To monitor instance disk I/O and aggregated load on each compute node, rather than analyzing compute node metrics, use the individual VM statistics presented through Grafana.

Bug: 33551814

Version: 3.0.1

Attached Block Volumes Not Visible Inside Oracle Solaris Instance

When you attach additional block volumes to a running Oracle Solaris compute instance, they do not become visible automatically to the operating system. Even after manually rescanning the disks, the newly attached block volumes remain invisible. The issue is caused by the hypervisor not sending the correct event trigger to re-enumerate the guest LUNs.

Workaround: When you attach additional block volumes to an Oracle Solaris compute instance, reboot the instance to make sure that the new virtual disks or LUNs are detected.

Bug: 33581238

Version: 3.0.1

Host Name Not Set In Successfully Launched Windows Instance

When you work in a VCN and subnet where DNS is enabled, and you launch an instance, it is expected that its host name matches either the instance display name or the optional host name you provided. However, when you launch a Windows instance, it may occur that the host name is not set correctly according to the launch command parameters. In this situation, the instance's fully qualified domain name (FQDN) does resolve as expected, meaning there is no degraded functionality. Only the host name setting within the instance itself is incorrect; the VCN's DNS configuration works as expected.

Workaround: If your instance host name does not match the specified instance launch parameters, you can manually change the host name within the instance. There is no functional impact.

Alternatively, if the Windows image contains `cloudbase-init` with the `SetHostNamePlugin`, it is possible to set the instance host name (*computer name*) based on the instance FQDN (*hostname-label*). To enable this function, execute the following steps:

1. Edit the file `C:\Program Files\Cloudbase Solutions\Cloudbase-Init\conf\cloudbase-init.conf`. Make sure it contains lines with these settings:

```
plugins=cloudbaseinit.plugins.common.sethostname.SetHostNamePlugin  
allow_reboot=true
```

2. Enter the command `Restart-Service cloudbase-init`.
3. Confirm that the instance host name has changed.

Bug: 33736674

Version: 3.0.1

Oracle Solaris Instance Stuck in UEFI Interactive Shell

It has been known to occur that Oracle Solaris 11.4 compute instances, deployed from the image delivered through the management node web server, get stuck in the UEFI interactive shell and fail to boot. If the instance does not complete its boot sequence, users are not able to log in. The issue is likely caused by corruption of the original `.oci` image file during the import into the tenancy.

Workaround: If your Oracle Solaris 11.4 instance hangs during UEFI boot and remains unavailable, proceed as follows:

1. Terminate the instance that fails to boot.
2. Delete the imported Oracle Solaris 11.4 image.
3. Import the Oracle Solaris 11.4 image again from the management node web server.
4. Launch an instance from the newly imported image and verify that you can log in after it has fully booted.

Bug: 33736100

Version: 3.0.1

Instance Backups Can Get Stuck in an EXPORTING or IMPORTING State

In rare cases, when an instance is exporting to create a backup, or a backup is being imported, and the system experiences a failure of one of the components, the exported or imported backup gets stuck in an EXPORTING or IMPORTING state.

Workaround:

1. Delete the instance backup.
2. Wait 5 minutes or more to ensure that all internal services are running.
3. Perform the instance export or import operation again.

See *Backing Up and Restoring an Instance* in [Compute Instance Deployment](#).

Bug: 34699012

Version: 3.0.1

Instances Experience Extended Downtime After Compute Node Reboot

When a compute node reboots while instances are still running, the system recovers the impacted compute instances automatically. When the compute node returns to normal operation, the hypervisor attempts to restart the instances on the same host compute node.

A restart command is only issued when an active connection has been established between the compute node agent and the hypervisor. If the hypervisor connection is down, the instance restart attempt fails, and a new attempt is made after 5 minutes. Up to 5 restart attempts are made with a 5 minute interval.

Workaround: After 5 unsuccessful restart attempts, or approximately 25 minutes after the rebooted compute node has returned to normal operation, the instances that remain unavailable must be restarted manually.

Bug: 34343810

Version: 3.0.1

Migrated Instance Fails to Start Due to Incomplete Cleanup of Guest Configuration on Target Compute Node

When a compute instance is live-migrated from one compute node to another, it goes through a series of internal operations and state changes to ensure that it comes up in good working condition on its new host, and that the hypervisor configuration on both compute nodes correctly reflects the changes resulting from the migration activity. Similarly, when a compute node's data network connection is interrupted, a cold migration process starts where compute instances also go through a series of internal operations and state changes.

When instances are involved in multiple successive migration operations, it is possible that the compute service is unable to clean up all state changes and hypervisor configurations quickly enough, even though all migration activity appears successful to the user. If you attempt to start a compute instance after cold migration where not all cleanup operations were completed, the start command will fail with an internal server error indicating that the *"Requested operation is not valid: Setting different SELinux label on /var/lib/libvirt/qemu/nvram/<filename>_VARS.fd which is already in use"*.

Workaround: This scenario is highly unlikely to occur in production environments. However, avoid actions that may trigger multiple instance migrations within minutes of each other. If you do run into this situation where instances are unrecoverable after migration, please contact Oracle for assistance.

Bug: 34640667

Version: 3.0.2

Instance Not Started After Fault Domain Change

When you change the fault domain of a compute instance, the system stops it, cold-migrates it to a compute node in the selected target fault domain, and restarts the instance on the new host. This process includes a number of internal operations to ensure that the instance can

return to its normal running state on the target compute node. If one of these internal operations fails, the instance could remain stopped.

The risk of running into issues with fault domain changes increases with the complexity of the operations. For example, moving multiple instances concurrently to another fault domain, especially if they have shared block volumes and are migrated to different compute nodes in the target fault domain, requires many timing-sensitive configuration changes at the storage level. If the underlying iSCSI connections are not available on a migrated compute instance's new host, the hypervisor cannot bring up the instance.

Workaround: After changing the fault domain, if a compute instance remains stopped, try to start it manually. If the instance failed to come up due to a timing issue as described above, the manual start command is likely to bring the instance back to its normal running state.

Bug: 34550107

Version: 3.0.2

Instance Start or Migration Fails Due to Libvirt Process Write Error

For compute instances with shared disks in read/write mode, the hypervisor manages the guest's SCSI-based disks through SCSI persistent reservation. All related requests are handled by a helper process (`qemu-pr-helper`) associated with the guest domain. The ID of this process is stored in a Libvirt control group (`cgroup`). When a compute instance is started, it may occur that the hypervisor is unable to find the ID of the `qemu-pr-helper` process associated with the instance. This results in a write error similar to this example:

```
Unable to write to '/sys/fs/cgroup/cpu,cpuacct/machine.slice/machine-qemuXYZ.scope/tasks': No such process
```

Workaround: The issue is intermittent. When you repeat the same operation, it is likely to succeed.

Bug: 34376870

Version: 3.0.2

Instance Migration Stuck in MOVING State

When migrating VMs using the Service Web UI it is possible that a migration can get stuck in the MOVING lifecycle state and you will be unable to continue further migrations.

This error can occur when administrative activities, such as live migrations, are running during a patching or upgrading process, or administrative activities are started before patching or upgrading processes have fully completed.

Workaround: Contact Oracle Support to resolve this issue.

Bug: 33911138

Version: , 3.0.1, 3.0.2

Migration Task Stays Active Even After the VM Migration Completed

When migrating VMs using the Service Web UI it is possible that a migration will complete successfully, but the migration job will show the Run State as `active`, rather than `completed`.

In this case, there is no action to take because the migration completed successfully, only the Run State of the job ID is incorrect.

Workaround: No action required.

Bug: 33974395

Version: , 3.0.1, 3.0.2

Storage Services Issues

This section describes known issues and workarounds related to the functionality of the internal ZFS storage appliance and the different storage services: block volume storage, object storage and file system storage.

Creating Image from Instance Takes a Long Time

When you create a new compute image from an instance, its boot volume goes through a series of copy and conversion operations. In addition, the virtual disk copy is non-sparse, which means the full disk size is copied bit-for-bit. As a result, image creation time increases considerably with the size of the base instance's boot volume.

Workaround: Wait for the image creation job to complete. Check the work request status in the Compute Web UI, or use the work request id to check its status in the CLI.

Bug: 33392755

Version: 3.0.1

Large Object Transfers Fail After ZFS Controller Failover

If a ZFS controller failover or failback occurs while a large file is uploaded to or downloaded from an object storage bucket, the connection may be aborted, causing the data transfer to fail. Multipart uploads are affected in the same way. The issue occurs when you use a version of the OCI CLI that does not provide the retry function in case of a brief storage connection timeout. The retry functionality is available as of version 3.0.

Workaround: For a more reliable transfer of large objects and multipart uploads, use OCI CLI version 3.0 or newer.

Bug: 33472317

Version: 3.0.1

Use Multipart Upload for Objects Larger than 100MiB

Uploading very large files to object storage is susceptible to connection and performance issues. For maximum reliability of file transfers to object storage, use multipart uploads.

Workaround: Transfer files larger than 100MiB to object storage using multipart uploads. This behavior is expected; it is not considered a bug.

Bug: 33617535

Version: n/a

File System Export Temporarily Inaccessible After Large Export Options Update

When you update a file system export to add a large number of 'source'-type export options, the command returns a service error that suggests the export no longer exists ("code": "NotFound"). In actual fact, the export becomes inaccessible until the configuration update has completed. If you try to access the export or display its stored information, a similar error is displayed. This behavior is caused by the method used to update file system export options: the existing configuration is deleted and replaced with a new one containing the requested changes. It is only noticeable in the rare use case when dozens of export options are added at the same time.

Workaround: Wait for the update to complete and the file system export to become available again. The CLI command `oci fs export get --export-id <fs_export_ocid>` should return the information for the export in question.

Bug: 33741386

Version: 3.0.1

Block Volume Stuck in Detaching State

Block volumes can be attached to several different compute instances, and can even have multiple attachments to the same instance. When simultaneous volume detach operations of the same volume occur, as is done with automation tools, the processes may interfere with each other. For example, different work requests may try to update resources on the ZFS storage appliance simultaneously, resulting in stale data in a work request, or in resource update conflicts on the appliance. When block volume detach operations fail in this manner, the block volume attachments in question may become stuck in *detaching* state, even though the block volumes have been detached from the instances at this stage.

Workaround: If you have instances with block volumes stuck in *detaching* state, the volumes have been detached, but further manual cleanup is required. The *detaching* state cannot be cleared, but the affected instances can be stopped and the block volumes can be deleted if that is the end goal.

Bug: 33750513

Version: 3.0.1

Fix available: Please apply the latest patches to your system.

Detaching Volume Using Terraform Fails Due To Timeout

When you use Terraform to detach a volume from an instance, the operation may fail with an error message indicating the volume attachment was not destroyed and the volume remains in attached state. This can occur when the storage service does not send confirmation that the volume was detached, before Terraform stops polling the

state of the volume attachment. The volume may be detached successfully after Terraform has reported an error.

Workaround: Re-apply the Terraform configuration. If the errors were the result of a timeout, then the second run will be successful.

Bug: 34732321

Version: 3.0.2

Scheduled Volume Backups Do Not Appear in Backup List

When you configure a block volume backup policy, a backup snapshot is created at regular intervals. However, when you list all the block volume backups in the compartment, these automated backups do not appear in the list. This happens by design: retrieving the list of backups for all volumes from the ZFS storage appliance is relatively time-consuming, so the synchronization only occurs on a per-volume basis and when an explicit manual request is made.

Workaround: To access the list of volume backups created through a backup policy, use the list command as shown below with `--volume-id <bv_ocid>`, to specifically display the backups of the volume in question. This synchronizes the entries to the list of volume backups, meaning they will all be displayed the next time you list all block volume backups in the compartment.

```
oci bv backup list --volume-id <bv_ocid> --compartment-id <compt_ocid>
```

Bug: 33785277

Version: 3.0.1

OCI CLI Might Not Return the Correct Value for Object Storage Namespace

Many commands that take the namespace value provide a default value, which is obtained using the CLI command `oci os ns get`. The data returned by that command might not be correct, causing the commands that use this value as the namespace to fail.

Workaround: Do not rely on the default value for object storage namespace. Enter the namespace value explicitly on the command line whenever the parameter is required.

To get the correct object storage namespace value, use the Web UI. Click your user menu in the top right corner of the Compute Web UI, and then click Tenancy. The namespace value is listed under Object Storage Settings.

If you need to obtain the namespace value through the OCI CLI, explicitly add the "iaas" endpoint to the command.

```
oci os ns get --endpoint https://iaas.<my pca>.example.com
{
  "data": "<myobjstor>"
}
```

For developers it is important to note that this behavior is different from Oracle Cloud Infrastructure.

Bug: 34133183

Version: 3.0.1

Creating File System Export Fails Due To Timeout

At a time when many file system operations are executed in parallel, timing becomes a critical factor and could lead to an occasional failure. More specifically, the creation of a file system export could time out because the file system is unavailable. The error returned in that case is: "*Internal Server Error: No such filesystem to create the export on*".

Workaround: Because this error is caused by a resource locking and timeout issue, it is expected that the operation will succeed when you try to execute it again. This error only occurs in rare cases.

Bug: 34778669

Version: 3.0.2

Concurrent File System Creation Operations Cause Exception During Flush

When the first file system resource is created within a new tenancy, a specific database entry is created. If another concurrent file system creation operation is in progress, it may occur that both jobs attempt to create the same database entry. This results in a conflict and returns an error similar to this: "*Session transaction has been rolled back due to a previous exception during flush.*".

Workaround: Because this error is caused by a timing issue, it is expected that the operation that failed and was rolled back, will succeed when you try to execute it again.

Bug: 34774136

Version: 3.0.2

File System Access Lost When Another Export for Subset IP Range Is Deleted

A virtual cloud network (VCN) can contain only one file system mount target. All file systems made available to instances connected to the VCN must have exports defined within its mount target. File system exports can provide access to different file systems from overlapping subnets or IP address ranges. For example: *filesys01* can be made available to IP range 10.25.4.0/23 and *filesys02* to IP range 10.25.5.0/24. The latter IP range is a subset of the former. Due to the way the mount IP address is assigned, when you delete the export for *filesys02*, access to *filesys01* is removed for the superset IP range as well.

Workaround: If your file system exports have overlapping source IP address ranges, and deleting one export causes access issues with another export similar to the example above, then it is recommended to delete the affected exports and create them again within the VCN mount target.

Bug: 33601987

Version: 3.0.2

File System Export UID/GID Cannot Be Modified

When creating a file system export you can add extra NFS export options, such as access privileges for source IP addresses and identity squashing. Once you have set a user/group identity (UID/GID) squash value in the NFS export options, you can no longer modify that value. When you attempt to set a different ID, an error is returned: "Uid and Gid are not consistent with FS AnonId: <currentUID>"

Workaround: If you need to change the UID/GID mapping, delete the NFS export options and recreate them with the desired values. If you are using the OCI CLI, you must delete the entire file system export (not just the options) and recreate the export, specifying the desired values with the `--export-options` parameter.

Bug: 34877118

Version: 3.0.2

ZFS Pool Usage Decreases After Instance Migration Due to Disk Allocation Change

The virtual disks or block volumes attached to compute instances are iSCSI LUNs configured in a ZFS pool on the ZFS Storage Appliance. The disk space available in these LUNs is fully allocated at the time of creation. However, when compute instances are migrated to another compute node, the LUNs associated with their block volumes are changed to sparse disk space allocation. This means not the configured size but only the disk space effectively in use is allocated, which is apparent from the decrease in ZFS pool usage.

```
PCA-ADMIN> show ZfsPool name=PCA_POOL
Data:
  Id = e898b147-7cf0-4bd0-8b54-e32ec83d04cb
  Type = ZfsPool
  Pool Status = Online
  Free Pool = 44879343128576
  Total Pool = 70506183131136
  Pool Usage Percent = 0.3634693989163486
  Name = PCA_POOL
  Work State = Normal
```

Workaround: There is no workaround available. The disk space allocation method is controlled through the Oracle ZFS Storage Appliance operating software.

Bug: 34922432

Version: 3.0.2

Serviceability Issues

This section describes known issues and workarounds related to service, support, upgrade and data protection features.

Minimum Upgrader Package Version for Upgrade to Release 3.0.2

In the upgrade path from appliance software Release 3.0.1 to Release 3.0.2 there are host upgrade commands that could fail and need to be executed a second time to complete the

operation. These issues can be avoided if a recent version of the upgrader is used to execute the appliance software upgrade procedure. The minimum package version that should be installed on the system before the upgrade to Release 3.0.2 is `pca-upgrader-3.0.101.17.g813e14f-1.el7`.

Workaround: From the command line of one of the management nodes, verify that your upgrader package is the version listed or a newer version. If not, please contact Oracle Support for assistance before attempting to upgrade to Release 3.0.2.

```
# rpm -qa | grep 'pca-upgrader-3.0.'  
pca-upgrader-3.0.101.17.g813e14f-1.el7
```

Bug: 34077896

Version: 3.0.2

Order of Upgrading Components Has Changed

When updating the platform, **you must update the compute nodes first**. Failing to update the compute nodes in this order can cause the upgrade to fail and disrupt the system.

Workaround: Complete platform upgrades in this order:

1. Compute Nodes
2. Management Nodes
3. Management Node Operating System
4. MySQL Cluster Database
5. Secret Service
6. Component Firmware
7. Kubernetes Cluster
8. Microservices

Bug: 34358305

Version: 3.0.1

DR Configurations Are Not Automatically Refreshed for Terminated Instances

If you terminate an instance that is part of a DR configuration, then a switchover or failover operation will fail due to the terminated instance. The correct procedure is to remove the instance from the DR configuration first, and then terminate the instance. If you forget to remove the instance first, you must refresh the DR configuration manually so that the entry for the terminated instance is removed. Keeping the DR configurations in sync with the state of their associated resources is critical in protecting against data loss.

Workaround: This behavior is expected. Either remove the instance from the DR configuration before terminating, or refresh the DR configuration if you terminated the instance without removing it first.

Bug: 33265549

Version: 3.0.1

Rebooting a Management Node while the Cluster State is Unhealthy Causes Platform Integrity Issues

Rebooting the management nodes is a delicate procedure because it requires many internal interdependent operations to be executed in a controlled manner, with accurate timing and often in a specific order. If a management node fails to reboot correctly and rejoin the cluster, it can lead to a destabilization of the appliance platform and infrastructure services.

Symptoms include: microservice pods in *CrashLoopBackOff* state, data conflicts between MySQL cluster nodes, repeated restarts of the NDB cluster daemon process, and so on.

Workaround: Before rebooting a management node, always verify that the MySQL cluster is in a healthy state. From the management node command line, run the command shown in the example below. If your output does not look similar and indicates a cluster issue, you should power-cycle the affected management node through its ILOM using the `restart / System` command.

As a precaution, if you need to reboot all the management nodes – for example in a full management cluster upgrade scenario –, observe an interval of at least 10 minutes between two management node reboot operations.

```
# ndb_mgm -e show
Connected to Management Server at: pcamn01:1186
Cluster Configuration
-----
[ndbd(NDB)]      3 node(s)
id=17   @253.255.0.33 (mysql-8.0.25 ndb-8.0.25, Nodegroup: 0)
id=18   @253.255.0.34 (mysql-8.0.25 ndb-8.0.25, Nodegroup: 0, *)
id=19   @253.255.0.35 (mysql-8.0.25 ndb-8.0.25, Nodegroup: 0)

[ndb_mgmd(MGM)] 3 node(s)
id=1     @253.255.0.33 (mysql-8.0.25 ndb-8.0.25)
id=2     @253.255.0.34 (mysql-8.0.25 ndb-8.0.25)
id=3     @253.255.0.35 (mysql-8.0.25 ndb-8.0.25)

[mysqld(API)]   18 node(s)
id=33   @253.255.0.33 (mysql-8.0.25 ndb-8.0.25)
id=34   @253.255.0.33 (mysql-8.0.25 ndb-8.0.25)
[...]
```

Bug: 34484128

Version: 3.0.2

ULN Mirror Is Not a Required Parameter for Compute Node Patching

In the current implementation of the patching functionality, the ULN field is required for all patch requests. The administrator uses this field to provide the URL to the ULN mirror that is set up inside the data center network. However, compute nodes are patched in a slightly different way, in the sense that patches are applied from an secondary, internal ULN mirror on the shared storage of the management nodes. As a result, the ULN URL is technically not required to patch a compute node, but the patching code does consider it a mandatory parameter, so it must be entered.

Workaround: When patching a compute node, include the URL to the data center ULN mirror as a parameter in your patch request. Regardless of the URL provided, the secondary ULN mirror accessible from the management nodes is used to perform the patching.

Bug: 33730639

Version: 3.0.1

Patch Command Times Out for Network Controller

When patching the platform, the process may fail due to a time-out while updating the network controller. If this is the case, logs will contain entries like "ERROR [pcanwctl upgrade Failed]".

Workaround: Execute the same patch command again. The operation should succeed.

Bug: 33963876

Version: 3.0.1

Upgrade Commands Fail when One Storage Controller Is Unavailable

The ZFS Storage Appliance has two controllers operating in an HA cluster, meaning it continues to operate when one of the controllers goes down. However, with one controller unavailable, upgrade-related operations will fail due to a connection error in the RabbitMQ internal message bus: "*Error in RabbitMQ service: No response received after 90 seconds*". Even viewing the upgrade job history is not possible, because the upgrade service is unable to send a response.

Workaround: Make sure that both storage controllers are up and running. Then, rerun the required upgrade commands.

Bug: 34507825

Version: 3.0.2

Management Node Upgrade Causes Loki Outage

During the upgrade of the management node cluster, each management node is upgraded and rebooted individually. Generally speaking, there is continuation of service because the cluster provides high availability. However, a brief outage of Loki is to be expected during upgrades of the management node cluster. This is considered normal behavior for a containerized microservice, running within pods distributed across multiple server nodes, which are rebooted one by one.

Workaround: If Loki log aggregation is interrupted during a management node upgrade, wait until the upgrade operations are completed and verify that Loki returns to normal operation. Automatic full recovery is expected as part of the Kubernetes container lifecycle.

Bug: 34452451

Version: 3.0.2

Instances with a Shared Block Volume Cannot Be Part of Different Disaster Recovery Configurations

Multiple instances can have a block volume attached that is shared between them. If you add those instances to a disaster recovery (DR) configuration, their attached volumes are moved to a dedicated ZFS storage project. However, if the instances belong to different DR configurations, each one with its own separate ZFS storage project, the system cannot move any shared block volume as this always results in an invalid DR configuration. Therefore, the Disaster Recovery service does not support adding compute instances with shared block volumes to different DR configurations.

Workaround: Consider including instances with a shared block volume in the same DR configuration, or attaching different block volumes to each instance instead of a shared volume.

Bug: 34566745

Version: 3.0.2

Time-out Occurs when Generating Support Bundle

When you request assistance from Oracle Support, it is usually required to upload a support bundle with your request. A support bundle is generated from a management node using a command similar to this example:

```
# support-bundles -m time_slice --all -s 2022-01-01T00:00:00.000Z -e  
2022-01-02T00:00:00.000Z
```

If there is a very large number of log entries to be collected for the specified time slice, the process could time out with API exception and an error message that says *"unable to execute command"*. In actual fact, the data collection will continue in the background, but the error is caused by a time-out of the websocket connection to the Kubernetes pod running the data collection process.

Workaround: If you encounter this time-out issue when collecting data for a support bundle, try specifying a shorter time slice to reduce the amount of data collected. If the process completes within 30 minutes the error should not occur.

Bug: 33749450

Version: 3.0.2

DR Operations Intermittently Fail

During certain conditions of heavy load, Site Guard users performing DR operations on the Private Cloud Appliance 3.0 can encounter out-of-session errors when Site Guard EM scripts attempt to perform DR operations using the PCA DR REST API.

This condition occurs when the system is overloaded with requests.

Workaround: Retry the operation.

Bug: 33934952

Version: 3.0.1, 3.0.2

Update of Disaster Recovery Replication Target Fails

When you update the configuration of the disaster recovery service to modify the replication target, there are several changes that need to be applied to the ZFS Storage Appliances and their network connectivity. Some settings cannot be changed until the existing connection has timed out. As a result, the disaster recovery service may need several attempts to apply the new settings to the ZFS Storage Appliance.

Workaround: Retrying the configuration update is built into the disaster recovery code. No administrator action is required. It is expected that the configuration is eventually changed successfully.

Alternatively, if the replication target update failure is permanent, it can be necessary to destroy and recreate the disaster recovery setup. This involves deleting all DR configuration from both systems, deleting the disaster recovery service on both systems, recreating the disaster recovery service on both racks with the new IPs, and recreating all DR configurations.

Bug: 34773761

Version: 3.0.2