Oracle Private Cloud Appliance
Release Notes

F74805-18
May 2024

Oracle Private Cloud Appliance Release Notes,

F74805-18

Contents

Preface

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

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 3.0.2-b1001356 (December 2023) 2-3
Features Released in Software Version 3.0.2-b925538 (August 2023) 2-5
Features Released in Software Version 3.0.2-b892153 (July 2023) 2-6
Features Released in Software Version 3.0.2-b852928 (May 2023) 2-8
Features Released in Software Version 3.0.2-b819070 (March 2023) 2-9
Features Released in Software Version 3.0.2-b799577 (February 2023) 2-10
Features Released in Software Version 3.0.2-b776803 (December 2022) 2-10
Features Released in Software Version 3.0.1-b741265 (November 2022) 2-12
Features Released in Software Version 3.0.1-b697160 (August 2022) 2-13

3 Service Limits

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

Guest Operating System Matrix 3-4

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
- Error Changing Hardware Component Password 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 4-3
- Free-Form Tags Used for Extended Functionality 4-3
- Do Not Use Reserved Tag Namespace 4-4
- Imported Images Not Synchronized to High-Performance Pool 4-4
- API Server Failure After Management Node Reboot 4-5
- CLI Command Returns Status 500 Due To MySQL Connection Error 4-5
- Administrators in Authorization Group Other Than SuperAdmin Must Use Service CLI to Change Password 4-5
- Service Web UI and Grafana Unavailable when HAProxy Is Down 4-6
- Lock File Issue Occurs when Changing Compute Node Passwords 4-6
- 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 4-7
- Concurrent Compute Node Provisioning Operations Fail Due to Storage Configuration Timeout 4-8
- Data Switch Fails to Boot Due to Active Console Connection 4-8
- Federated Login Failure after Appliance Upgrade 4-8
- Ensure No Storage Buckets Are Present Before Deleting a Compartment or Tenancy 4-9
- Syntax Issue when Generating Certificate Signing Request 4-9
- Listing Upgrade Jobs Fails with RabbitMQ Error 4-10
- Availability Domain Name Change in Version 3.0.2-b1001356 4-10
- No Packages Available to Patch MySQL Cluster Database 4-10
- Uppercase Letters Are Not Supported in Domain Names 4-12

User Interface Issues 4-12
- Moving Resources Between Compartments Is Not Supported in the Compute Web UI 4-12
- Saving Resource Properties Without Modifications Briefly Changes Status to Provisioning 4-12
- NFS Export Squash ID Not Displayed 4-12
- Scrollbars Not Visible in Browser 4-13
- Authorization Failure When Retrieving Compartment Data 4-13
- Object List Is Not Updated Automatically 4-13
- Not All Resources Shown in Drop-Down List 4-14

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 4-15
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
No Details Displayed for File System Cloned from Snapshot 4-17
Error Message when Exadata Network Is Deleted Successfully 4-17
Unable to Display Details of Instance Backup 4-17
IP Address List on VNIC Detail Page Not Updated 4-17
No Error Displayed When Attempting to Reserve Public IP Address While All Public IPs In Use 4-18
When Modifying Listener of Load Balancer, Existing Configuration Parameters Are Ignored 4-18
Create Backend to Add Security List/Configure Automatically Using IP Address Does Not Display Egress/Ingress Lists 4-18
Tag Area Above Resource Tables Does Not Expand 4-19
Service API Reference Displays Wrong Path for Disaster Recovery Commands 4-19
Details of an Instance Import Are Displayed as an Instance Export 4-20
When Canceling an Instance Import the Operation Does Start 4-20
Networking Issues 4-20
Possible Impact to BGP Links After Upgrading to 3.0.2-b1081555 or Higher 4-20
DNS Zone Scope Cannot Be Set 4-21
To Update a DNS Record the Command Must Include Existing Protected Records 4-21
Create Route Table Fails With Confusing Error Message 4-21
VCN Creation Uses Deprecated Parameter 4-21
File Storage Traffic Blocked By Security Rules 4-22
Stateful and Stateless Security Rules Cannot Be Combined 4-23
Routing Failure With Public IPs Configured as CIDR During System Initialization 4-24
Admin Network Cannot Be Used for Service Web UI Access 4-24
Network Configuration Fails During Initial Installation Procedure 4-24
External Certificates Not Allowed 4-25
DNS Entries on Oracle Linux 8 Instances Incorrect After Upgrade to Release 3.0.2 4-25
Network Load Balancer Does Not Report Detailed Backend Health Status 4-25
Route Table Stuck in Provisioning State Failure 4-25
Updating Route Table Using Terraform Fails Because DRG Is Not Attached 4-26
Failure Executing Terraform Destroy Due to Route Table in Provisioning State 4-26
When Configuring BGP Authentication the Password Is a Required Parameter 4-27
Uplink VRRP Mesh Configuration Sets Second Switch IP Incorrectly 4-27
Real Application Cluster (RAC) Environment Loses Access to Oracle Exadata Database Instances During Appliance Upgrade 4-28
Compute Service Issues 4-28

Possible VM Impact When OS Images are Deleted 4-28
E5.Flex Instance Shape Is Not Supported on the X9-2 Hardware Platform 4-28
Displaced Instances Not Returned to Their Selected Fault Domains 4-29
Terraform Cannot Be Used for Instance Update 4-29
No Consistent Device Paths for Connecting to Block Volumes 4-30
Instance Pools Cannot Be Terminated While Starting or Scaling 4-30
TypeError Returned when Attaching an Instance to an Instance Pool 4-30
Network Interface on Windows Does Not Accept MTU Setting from DHCP Server 4-31
Oracle Solaris Instance in Maintenance Mode After Restoring from Backup 4-31
Instance Disk Activity Not Shown in Compute Node Metrics 4-32
Attached Block Volumes Not Visible Inside Oracle Solaris Instance 4-32
Host Name Not Set In Successfully Launched Windows Instance 4-32
Oracle Solaris Instance Stuck in UEFI Interactive Shell 4-33
Instance Backups Can Get Stuck in an EXPORTING or IMPORTING State 4-33
Instance Not Started After Fault Domain Change 4-33
Instance Migration Stuck in MOVING State 4-34
OCI CLI Commands Fail When Run From a Compute Instance 4-34
Cannot Install OCI CLI on Oracle Linux 9 Instance 4-35
Instance Launch Fails at 80 Percent Complete with Libvirt Error 4-36
Instance Principal Unavailable Until Next Certificate Renewal Check 4-36
Unable to Delete Tag Due to Instance Principal Error 4-36
When Instance Is Shut Down from OS, Soft Stop Results in Conflict 4-37
Storage Services Issues 4-37
Updating Terraform Changes File Storage Export Path 4-37
Creating Image from Instance Takes a Long Time 4-38
Large Object Transfers Fail After ZFS Controller Failover 4-38
Use Multipart Upload for Objects Larger than 100MiB 4-38
File System Export Temporarily Inaccessible After Large Export Options Update 4-38
Block Volume Stuck in Detaching State 4-39
Detaching Volume Using Terraform Fails Due To Timeout 4-39
Creating File System Export Fails Due To Timeout 4-39
File System Access Lost When Another Export for Subset IP Range Is Deleted 4-40
File System Export UID/GID Cannot Be Modified 4-40
Block Volume Performance Level Not Preserved During Cloning 4-40
Internal Backups for Instance Cloning Not Displayed 4-41
Limit for Volume Backups Not Enforced 4-41
NFS Service Interruption During ZFS Storage Appliance Firmware Upgrade or Patching 4-41
Container Engine Issues 4-41
Container Engine for Kubernetes Requires Switch Firmware Upgrade on Systems with Administration Network 4-42
Tag Filters Not Available for Kubernetes Node Pools and Nodes 4-42
Kubernetes Node Tags Not Available in Compute Web UI 4-42

Node Doctor Script Not Available in Worker Nodes 4-42
Unable to Delete Kubernetes Cluster in Failed State 4-43
Kubernetes Cluster Creation Failure Due to Load Balancer Limit Returning Unclear Error 4-43
Intermittent Failures when Using Terraform to Create Kubernetes Cluster 4-43
API Reference on Appliance Not Up-to-Date for OKE Service 4-44
Serviceability Issues 4-44
Order of Upgrading Components Has Changed 4-44
DR Configurations Are Not Automatically Refreshed for Terminated Instances 4-44
Rebooting a Management Node while the Cluster State is Unhealthy Causes Platform Integrity Issues 4-45
ULN Mirror Is Not a Required Parameter for Compute Node Patching 4-45
Patch Command Times Out for Network Controller 4-46
Upgrade Commands Fail when One Storage Controller Is Unavailable 4-46
Instances with a Shared Block Volume Cannot Be Part of Different Disaster Recovery Configurations 4-46
Time-out Occurs when Generating Support Bundle 4-47
DR Operations Intermittently Fail 4-47
MN01 Host Upgrade Fails When it is the Last Management Node to Upgrade 4-47
Failure Draining Node when Patching or Upgrading the Kubernetes Cluster 4-48
Oracle Auto Service Request Disabled after Upgrade 4-48

Preface

This publication is part of the customer documentation set for Oracle Private Cloud Appliance Release 3.0. 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.
italic	Italic type indicates book titles, emphasis, or placeholder variables for which you supply particular values.

Convention	Meaning
`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 `root` user.

Documentation Accessibility

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

Access to Oracle Support for Accessibility

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

Diversity and Inclusion

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

1

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.

To check which components require upgrading or patching with a released software build, see My Oracle Support. Refer to the note with Doc ID 2907892.1. For information and recommendations about released builds, refer to the note with Doc ID 2906831.1.

Latest Features

> **Caution:**
>
> Prior to patching or upgrading to the latest release, ensure that all compute nodes are in the provisioned state.

Platform Images

New platform images are made available for Compute Enclave users through Private Cloud Appliance installation, upgrade, and patching.

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

Oracle Linux 9	`uln-pca-Oracle-Linux-9-2023.09.26_0.oci`
Oracle Linux 8	`uln-pca-Oracle-Linux-8-2023.09.26_0.oci`
Oracle Linux 7.9	`uln-pca-Oracle-Linux-7.9-2023.09.26_0.oci`
Oracle Solaris 11.4	`uln-pca-Oracle-Solaris-11-2023.10.16_0.oci`

Container Engine for Kubernetes	`uln-pca-Oracle-Linux8-OKE-1.26.6-20240210.oci`
	`uln-pca-Oracle-Linux8-OKE-1.27.7-20240209.oci`
	`uln-pca-Oracle-Linux8-OKE-1.28.3-20240210.oci`

Container Engine for Kubernetes

Oracle Private Cloud Appliance Container Engine for Kubernetes (OKE) is a scalable, highly available service that can be used to deploy any containerized application to the cloud. OKE uses Cluster API Provider (CAPI) and Cluster API Provider for Oracle Cloud Infrastructure (CAPOCI) to orchestrate the cluster on the Private Cloud Appliance. OKE uses Kubernetes, the open-source system for automating deployment, scaling, and management of containerized applications across clusters of hosts. Kubernetes groups the containers that make up an application into logical units called pods for easy management.

See the Oracle Private Cloud Appliance Container Engine for Kubernetes Guide for information about how to configure the network, create a Kubernetes cluster, create a node

pool, expose containerized applications outside the appliance, and provide persistent storage for containerized applications. See the OKE Monitoring folder in Grafana for OKE dashboards.

Instance Principals

An instance principal is a compute instance that is authorized to perform actions on service resources. Applications running on an instance principal can call services and manage resources similar to the way Private Cloud Appliance users call services to manage resources. The instance is a principal actor just as a user is a principal actor. When you use instance principals, you do not need to configure user credentials or a configuration file on the instance to run applications that need to manage service resources.

To grant authorizations to an instance principal, include the instance as a member of a dynamic group. A dynamic group provides authorizations to instances just as a user group provides authorizations to users.

See "Configuring Instances for Calling Services" and "Creating and Managing Dynamic Groups" in the Identity and Access Management chapter of the *Oracle Private Cloud Appliance User Guide*.

Upgrade History and Enhancements

The upgrade history presents information from all upgrade and patch jobs in a categorized way, providing insight into which version upgrades have been performed, which jobs have been run for each of those upgrades, and from which source (ISO upgrade or ULN patch). Details include build versions, component versions before and after, job completion, success or failure, time stamps, and duration.

Oracle Integrated Lights Out Manager (ILOM) has been included in the upgrade or patch workflow of compute node and management node hosts, reducing the overall process duration and the number of reboots. The upgrade plan has also been refined and covers the Oracle Cloud Infrastructure images provided with Private Cloud Appliance.

An appliance software prerequisite version check is performed during the upgrade or patch preparation phase. The Upgrader service proceeds only if the running version passes this check, otherwise you must install the minimum required version first, before proceeding with the intended target version.

Upgrading or patching to this version of the appliance software also adds the *region registry*, which contains resources required for the operation of the Oracle Container Engine for Kubernetes (OKE).

All changes are reflected in the Oracle Private Cloud Appliance Upgrade Guide and Oracle Private Cloud Appliance Patching Guide.

ULN Mirror for Appliance Patching on Oracle Linux 8

The ULN mirror in the data center, which the appliance uses to retrieve new packages to patch components to the latest version, can now be configured on an Oracle Linux 8 server. For detailed information, refer to the chapter "Configure Your Environment for Patching" in the Oracle Private Cloud Appliance Patching Guide.

Backup Space Management

The Backup and Restore Service has been enhanced further to optimize storage space consumption on the ZFS Storage Appliance. When purging backups older than the retention period, the service also removes the large temporary files of previous MySQL database backups that are no longer required.

DRGv1+ Support

DRGv1+ provides VRF/VLAN-backed isolation for network traffic when using a DRG.

Features Released in Software Version 3.0.2-b1001356 (December 2023)

X10 Rack Configuration

The Private Cloud Appliance X10 rack configuration ships from the factory with appliance software version 3.0.2-b1001356 or newer installed. This rack configuration is characterized by the use of Oracle Server X10 compute nodes (2U). Compared to the X9 rack configuration, the component order and cabling are slightly different. The storage and network infrastructure components are identical.

In terms of compute capacity, a 2U compute node is comparable to two 1U compute nodes. To deploy compute instances on the Oracle Server X10 compute nodes, you must select the *VM.PCAStandard.E5.Flex* shape, which offers adjustable CPUs, memory, and network bandwidth.

For more information about the X10 rack configuration, refer to the chapter Hardware Overview in the "Oracle Private Cloud Appliance Concepts Guide".

Platform Images

New platform images are made available for Compute Enclave users through Private Cloud Appliance installation, upgrade, and patching.

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

Oracle Linux 9	`uln-pca-Oracle-Linux-9-2023.08.31_0.oci`
Oracle Linux 8	`uln-pca-Oracle-Linux-8-2023.08.31_0.oci`
Oracle Linux 7.9	`uln-pca-Oracle-Linux-7.9-2023.08.31_0.oci`
Oracle Solaris 11.4	`uln-pca-Oracle-Solaris-11-2023.09.20_0.oci`

Network Load Balancer

A network load balancer provides automated layer 4 traffic distribution from one public or private entry point where incoming requests are received, to a set of backend servers in the virtual cloud network (VCN) where the requests are processed. For efficient resource management, you can attach a compute instance pool as a backend set.

Network load balancers and previously implemented (layer 7) application load balancers can coexist in your environment, and have shared resource configuration limits. The key difference is that the network load balancer operates at OSI network layer 4 and manages TCP traffic. It provides better performance but lacks the layer 7 routing intelligence. Distribution to the backend servers is controlled by a 5-tuple, 3-tuple, or 2-tuple hash policy. However, for architectural reasons these are mapped internally to a source IP hash, which ensures that a client's requests are all directed to the same backend server.

For information about layer 4 load balancing, see the chapter Network Load Balancing Overview of the *Oracle Private Cloud Appliance Concepts Guide*. Instructions to configure network load balancers can be found in the chapter Network Load Balancers of the *Oracle Private Cloud Appliance User Guide*

File System Service Improvements

File System Quota

You can set a space quota on a file system when you create the file system and when you update the file system. The quota includes the data in the file system and all snapshots created under the file system. You cannot set a quota smaller than the current usage of the file system.
For more information, see "Creating a File System" and "Updating a File System" in the *Oracle Private Cloud Appliance User Guide*.

File System High Performance Backing Store

By default, the backing store of a file system instance is the default pool of the attached ZFS Storage Appliance. You can specify that you want to use a high performance pool for the backing store. See `poolName` in "Creating a File System" in the *Oracle Private Cloud Appliance User Guide*.

Block Storage Volume Performance Option

By default, block volumes have Balanced performance. When you create block storage, you can optionally enable High performance. For a comparison of Balanced Performance and High Performance, see "Block Volume Performance Options" in the *Oracle Private Cloud Appliance Concepts Guide*.

Instance Pool Updates

Instance Pool Instance Attach and Detach

You can attach an existing instance to an instance pool or detach an instance that is attached to an instance pool.
When you detach an instance from a pool, you have the following choices:

- Regarding the detached instance, you can choose to terminate the instance or keep the instance as a standalone instance.
- Regarding the instance pool, you can choose to leave the pool as a smaller pool or create a new instance in the pool, using the instance configuration parameters for the pool.

For more information, see "Updating an Instance Pool" in the *Oracle Private Cloud Appliance User Guide*.

Instance Pool Soft Stop and Soft Reboot

When you use the Compute Web UI to stop or reboot an instance pool, by default soft stop or soft reboot is selected. A dialog enables you to stop or reboot all the instances in the pool immediately.
When you use the OCI CLI, you can specify `softstop` or `softreset`.
For more information, see "Stopping and Starting Instances in an Instance Pool" in the *Oracle Private Cloud Appliance User Guide*.

Compute Service Improvements

Instance Serial Console

To troubleshoot an instance that is not running, you can connect to the instance serial console as an alternative to using the instance VNC console. For more information, see "Remotely Troubleshooting an Instance by Using a Console Connection" in the *Oracle Private Cloud Appliance User Guide*.

Instance Configuration from an Existing Instance
In addition to creating an instance configuration by entering values in the Compute Web UI or in a file, you can create an instance configuration by using the configuration information from an existing compute instance. See "Working with Instance Configurations" in the *Oracle Private Cloud Appliance User Guide*.

More Bandwidth for Flexible Shape Instance Configurations
Maximum bandwidth for the VM.PCAStandard1.Flex shape is updated to more closely match the maximum bandwidth for fixed shapes. For 1-24 OCPUs, the maximum bandwidth is 24.6 Gbps. For 25-32 OCPUs, the maximum bandwidth is 1 Gbps per OCPU.
For the VM.PCAStandard.E5.Flex shape, maximum bandwidth for 1-24 OCPUs is 24.6 Gbps. Maximum bandwidth for 25-40 OCPUs is 1 Gbps per OCPU. Maximum bandwidth for 41-96 OCPUs is 40.0 Gbps.

IMDS Version 2 Endpoints
The Instance Metadata Service is available in two versions: version 1 and version 2. To increase the security of metadata requests, upgrade applications to use the IMDS version 2 endpoints.
New options are available in instance create and instance update to disable recognition of IMDSv1 endpoints. For more information, see "Retrieving Instance Metadata from Within the Instance" in the *Oracle Private Cloud Appliance User Guide*.

Features Released in Software Version 3.0.2-b925538 (August 2023)

Platform Images

New platform images are made available for Compute Enclave users through Private Cloud Appliance installation, upgrade, and patching.

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

Oracle Linux 9	`uln-pca-Oracle-Linux-9-2023.05.24_0.oci`
Oracle Linux 8	`uln-pca-Oracle-Linux-8-2023.05.24_0.oci`
Oracle Linux 7.9	`uln-pca-Oracle-Linux-7.9-2023.05.24_0.oci`
Oracle Solaris 11.4	`uln-pca-Oracle-Solaris-11-2023.04.18_0.oci`

New Kubernetes Version 1.25

In this release, the Kubernetes cluster hosted on the management nodes is upgraded to version 1.25, which implies that the environment goes through 5 full upgrade cycles. The Upgrader manages the entire process for you, but note that it takes at least 4 hours on a minimum appliance configuration, and up to 18 hours on a fully populated rack.

Streamlined Upgrade Process

Based on the upgrade plan, upgrade and patch operations for all components except firmware follow a prescribed order. All steps to prepare the upgrade environment must be completed before any upgrade or patch command can be run.

The upgrade plan logic has been improved to ensure that unnecessary component reboots are avoided, This shortens the overall upgrade or patch duration and minimizes the risk of failures, degraded performance, or downtime. In addition, the estimated durations listed in the upgrade

plan are calculated more accurately, and the reboot and upgrade requirement indicators are more reliable.

Systems running appliance software version 3.0.1 need to be upgraded twice to reach the latest release. An intermediate upgrade to a version between 3.0.2-b776803 and 3.0.2-b892153 is required.

All changes are reflected in the Oracle Private Cloud Appliance Upgrade Guide and Oracle Private Cloud Appliance Patching Guide.

New Uplink Configuration Options

The data network uplinks (ports 1-4) from the appliance to the data center network can now be configured as a static routing connection without the use of vPC/MLAG for link aggregation. Data uplinks in active/active mode use the ECMP protocol; active/passive uplinks use VRRP.

The optional administration network uplink (port 5) now provides similar routing options as the data network. The administration network can be configured to use BGP-based dynamic routing. Static routing can now be configured without requiring vPC/MLAG for link aggregation. In active/active mode the uplink uses the ECMP protocol; an active/passive uplink uses VRRP.

See Configuring Oracle Private Cloud Appliance in the Oracle Private Cloud Appliance Installation Guide

DNS Mapping for Disaster Recovery

Custom CA certificates can now be used on systems that have been configured for disaster recovery. To allow the replication IP addresses to be resolved, pointer records must be added to the data center DNS configuration. These PTR records must map the ZFS Storage Appliance host names to the replication interface IPs of the remote system in the DR configuration. For more information and instructions, refer to the chapter Disaster Recovery in the Oracle Private Cloud Appliance Administrator Guide.

New Generation SSD

The Oracle Server X9-2 management and compute nodes use a pair of 240GB M.2 SATA hard drives as boot devices. As the current model is being phased out, a new generation has been qualified for Private Cloud Appliance. The new SSDs are functionally the same as the earlier model.

Features Released in Software Version 3.0.2-b892153 (July 2023)

Upgrade Enhancements

Both the upgrade and patching processes are now based on an *upgrade plan*, which is the result of a metadata comparison between the new version to be installed and the version currently running on the system. The upgrade plan ensures that components are only patched or upgraded if the latest version is newer. Overwriting an installed component with the same version can be forced, if necessary.

Every upgrade and patch command has a *verify-only* option. This can be used to test in advance for system health issues that would prevent the operation from completing successfully.

Patching or upgrading the operating system on the management nodes no longer requires the administrator to manually reassign the cluster primary role to another node in between

operations. The code detects which node holds the primary role and automatically modifies the cluster configuration in the background.

User-friendly commands to retrieve IP addresses of nodes and their ILOMs have been added to the Service CLI.

The full management node cluster upgrade automatically runs the `upgradeOCIImages` command at the end. This updates the Oracle Cloud Infrastructure compute images across all tenancies.

All changes are reflected in the Oracle Private Cloud Appliance Upgrade Guide and Oracle Private Cloud Appliance Patching Guide.

Your Own CA Trust Chain

In the Oracle Private Cloud Appliance architecture, you can provide your own CA certificates which allows you to use your CA trust chain to access the rack's external interfaces.

> **Note:**
>
> OpenSSH clients must be at least version openssh-clients-7.4p1 or later.

For instructions, see "Accessing External Interfaces with Your Certificate Authority Trust Chain" in the *Oracle Private Cloud Appliance Administrator Guide*.

Oracle Defined Volume Backup Policies

For scheduled (policy-based) backups of block volumes and boot volumes, you can select an Oracle defined policy to use as an alternative to defining your own policy. For descriptions of the Oracle defined backup policies, see "Volume Backups and Clones" in the Block Volume Storage Overview chapter in the *Oracle Private Cloud Appliance Concepts Guide*. For information about how to use these policies, see "Managing Backup Policies" in the Block Volume Storage chapter of the *Oracle Private Cloud Appliance User Guide*.

Platform Images

New platform images are made available for Compute Enclave users through Private Cloud Appliance installation, upgrade, and patching.

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-11-2023.04.18_0.oci`

Support for Oracle Exadata Multi-cluster VMs

Private Cloud Appliance now provides support for Oracle Exadata multi-cluster VMs. Note that when you are connecting more than one Oracle Exadata system to the Private Cloud Appliance rack, you must ensure that the configured IP address ranges do not overlap.

System Backup Service Updates

System backups that were created by running either a daily scheduled `brs` backup or a manual (`backup-now`) `brs` backup are retained for no more than 14 days.

After you upgrade the Private Cloud Appliance to this release and run either type of `brs` backup, all `brs` backups that were previously created and are older than 14 days are deleted.

ZFSSA manual snapshots are not deleted if they were not created by using any `brs` job, and their snapshot name is not in the following form (the `brs` snapshot naming convention):

projectname*/*filesystemname_timestamp

For more information, see "Backup and Restore" in the Appliance Administration Overview chapter of the *Oracle Private Cloud Appliance Concepts Guide* and the Backup and Restore chapter in the *Oracle Private Cloud Appliance Administrator Guide*.

Features Released in Software Version 3.0.2-b852928 (May 2023)

Load Balancer as a Service (Layer 7)

The Load Balancing service provides automated traffic distribution from one public or private entry point where incoming requests are received, to a set of backend servers in the virtual cloud network (VCN) where the requests are processed.

To efficiently manage compute resources associated with a load balancer, you can attach a compute instance pool to a load balancer backend set. Doing this adds each instance in the pool as a backend server in the backend set.

A load balancer manages TCP or HTTP traffic based on typical distribution policies like round-robin, least connections, or IP hash.

For optimal utilization of the backend resources, request routing on the listener side can be further refined by using multiple virtual host names and path route rules. Load balancers also provide configuration options for SSL traffic handling and session persistence.

For information about load balancing, see the Load Balancing Overview chapter of the *Oracle Private Cloud Appliance Concepts Guide*. Instructions to configure load balancers can be found in the chapter Load Balancer as a Service of the *Oracle Private Cloud Appliance User Guide*

Instance Pool Autoscaling

An instance pool defines a set of compute instances that is managed as a group. Autoscaling a pool enables you to use resources more effectively by stopping or removing instances when demand is lower and starting or adding instances when demand is higher.

For more information, see "Using Schedule-Based Autoscaling" in the section titled "Working with Instance Pools" in the *Oracle Private Cloud Appliance User Guide*.

You can use instance pool autoscaling along with load balancing by attaching an instance pool that has an autoscaling configuration to a load balancer backend set. See "Managing Instance Pool Load Balancer Attachments" in the section titled "Working with Instance Pools" in the *Oracle Private Cloud Appliance User Guide*.

High Availability Configuration and Fault Domain Enforcement Default

The Service Enclave has new commands that give administrators more control over how to implement instance high availability.

- Instance high availability: When enabled, instances are automatically reboot migrated off of an unreachable compute node. The default is enabled.

- Instance fault domain resolution: When enabled, instances that are running in a fault domain that is not the fault domain that is specified in their instance configuration (their selected fault domain) are automatically migrated back to their selected fault domain when resources become available in that fault domain. The default is enabled.

  Instances can become displaced (running in a fault domain that is not their selected fault domain) during compute node evacuation or failure. You can list all currently displaced instances.

- Instance restart: When enabled, instances that were stopped by the Compute service (not by an administrator) are automatically restarted in their selected fault domain when resources become available in that fault domain. The default is enabled.

  Instances can be stopped by the Compute service during compute node evacuation or compute node failure when no fault domain has resources to accommodate the instances, or when strict fault domain enforcement is enabled and no other compute node in the selected fault domain can accommodate the instances. You can list all instances that are currently stopped by the Compute service.

- Strict fault domain enforcement: When enabled, instances that cannot be accommodated in the current fault domain during compute node evacuation or failure will be stopped by the Compute service. If the force option is not used for compute node evacuation, instances will be still running in the current compute node and the compute node evacuation will fail.

  The default is disabled: Instances will be migrated to a different fault domain if possible.

For more information, see "Migrating Instances from a Compute Node" and "Configuring the Compute Service for High Availability" in the Hardware Administration chapter of the *Oracle Private Cloud Appliance Administrator Guide*.

Routing Options in Virtual Networking

The virtual networking configuration in the Compute Enclave provides additional routing options:

- setting a private IP address as a route rule target
- associating a route table with a dynamic routing gateway (DRG) attachment.

With this enhancement, existing data center infrastructure and applications can be integrated into the network communication between Private Cloud Appliance compute instances, or between instances and other network services external to the appliance.

Features Released in Software Version 3.0.2-b819070 (March 2023)

Fault Domain Enforcement for Compute Node Evacuation

When evacuating a compute node, you can specify the behavior you want if some instances cannot be accommodated in other compute nodes in the same fault domain.

- If strict enforcement is disabled, instances that cannot be accommodated in the current fault domain will be migrated to other fault domains if possible.
- If strict enforcement is enabled, instances that cannot be accommodated in the current fault domain will be left running in the current compute node.

See "Migrating Instances from a Compute Node" in the chapter Hardware Administration of the Oracle Private Cloud Appliance Administrator Guide for more information.

New Upgrade Guide

The instructions to upgrade an appliance were included in the Oracle Private Cloud Appliance Administrator Guide. That content is now moved to a separate Oracle Private Cloud Appliance Upgrade Guide.

Features Released in Software Version 3.0.2-b799577 (February 2023)

Fault Descriptions Improved

Hardware-related fault messages appearing in command output, logs and monitoring data now describe the observed issue more accurately. When a fault message is related to resource utilization, for example when a threshold is exceeded, the description explicitly mentions it. This prevents utilization warnings from being mistaken for actual hardware problems and makes troubleshooting easier.

Mounting File Systems Across Appliances

In a deployment comprising multiple Private Cloud Appliance systems it is now possible to mount an NFS export from the ZFS Storage Appliance of one appliance on a compute instance hosted on another appliance. The VCNs, mount target and export options must be appropriately configured to grant instances access to the remote file system.

Features Released in Software Version 3.0.2-b776803 (December 2022)

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*.

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 Compute 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 Monitoring System Capacity 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 Monitoring System Capacity 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 chapter 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.

Compute Node Physical Resources

A part of each compute node's CPU and RAM capacity is reserved for system use. The table shows available physical resources by compute node model.

Compute Node Model	Total Physical Resources	Available Compute Capacity
Oracle Server X9-2	CPU: 64 cores (2x32) RAM: 1024 GB (16x64)	CPU: 60 cores RAM: 960 GB
Oracle Server X10	CPU: 192 cores (2x96) RAM: 2304 GB (24x96)	CPU: 184 cores RAM: 2224 GB

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, unless indicated otherwise. 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	40 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

Service	Resource Type	Limit
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
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
Networking Service	DNS zones	1000 across all tenancies (in addition to any internal zones)
Networking Service	DNS records	25000 per zone
Compute Service	Custom images	100
Block Storage Service	Aggregated size of block volumes	100TB (with default storage capacity)
Block Storage Service	Block volume backups	100000 across all tenancies
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
(Network) Load Balancing Service	Load balancers (Network LB and LBaaS combined)	32 total across all tenancies, 20 in a single VCN
(Network) Load Balancing Service	IP address	1 per load balancer
(Network) Load Balancing Service	Network security groups	5 per load balancer
(Network) Load Balancing Service	Listeners	16 per load balancer
(Network) Load Balancing Service	Backend sets	4 per load balancer
(Network) Load Balancing Service	Backend servers	512 per load balancer and per backend set
Kubernetes Container Engine	Clusters	10 per tenancy
Kubernetes Container Engine	Worker nodes	128 per cluster (across all pools)
Kubernetes Container Engine	Pods	110 per node (Kubernetes default)

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
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
Kubernetes cluster	create/update/delete cluster	10
Kubernetes node pool	create/update/delete node pool	5
Kubernetes node	create/update/delete node	15

> **Note:**
>
> In addition, there is a system limit on the number of concurrent user sessions:
>
> - Compute Web UI: 10 tenancy users
> - Service Web UI: 6 administrators
>
> An authentication error is displayed when the limit is reached. An inactive user session times out after 1 hour.

Guest Operating System Matrix

Oracle Cloud Infrastructure compute images of Oracle Linux and Oracle Solaris are provided as part of the appliance software, and new image versions are added through upgrades and patches. Updates of the Oracle-provided images are listed by software version in the Feature Updates chapter.

Oracle Private Cloud Appliance supports other guest operating systems, which you can add to your appliance environment as custom images. Several guest operating systems are part of Oracle testing and are known to work in Private Cloud Appliance compute instances. The table below provides an overview.

Guest Operating System	Oracle-Provided Image	Custom Image	Oracle-Tested
Oracle Linux 9.x	Y	Y	Y
Oracle Linux 8.x	Y	Y	Y
Oracle Linux 7.x	Y	Y	Y
Oracle Solaris 11.x	Y	Y	Y
Red Hat Enterprise Linux 9.2		Y	
Red Hat Enterprise Linux 8.x		Y	
Red Hat Enterprise Linux 7.x		Y	
CentOS Linux 8.x		Y	
CentOS Linux 7.x		Y	
SUSE Linux Enterprise Server 15 (latest)		Y	
SUSE Linux Enterprise Server 12 SP4		Y	
Ubuntu 20.04 and later		Y	
Ubuntu 18.04 and later		Y	
AlmaLinux OS 9.2		Y	Y
Kali Linux		Y	Y
Microsoft Windows Server 2022		Y	Y
Microsoft Windows Server 2019		Y	Y
Microsoft Windows Server 2016		Y	Y
Microsoft Windows Server 2012 R2		Y	Y
Microsoft Windows Server 2012		Y	Y

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

Error Changing Hardware Component Password

The hardware layer of the Oracle Private Cloud Appliance architecture consists of various types of components with different operating and management software. As standalone

products their password policies can vary, but the appliance software enforces a stricter rule set. If an error is returned when you try to change a component password, ensure that your new password complies with the Private Cloud Appliance policy for hardware components.

For more information about password maintenance across the entire appliance environment, refer to the Oracle Private Cloud Appliance Security Guide.

Workaround: For hardware components, use the Service CLI to set a password that conforms to the following rules:

- consists of at least 8 characters
  - with a maximum length of 20 characters for compute nodes, management nodes, and switches
  - with a maximum length of 16 characters for ILOMs and the ZFS Storage Appliance
- contains at least one lowercase letter (a-z)
- contains at least one uppercase letter (A-Z)
- contains at least one digit (0-9)
- contains at least one symbol (@$!#%*&)

Bug: 35828215

Version: 3.0.2

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

Do Not Use Reserved Tag Namespace

Oracle Private Cloud Appliance uses a reserved tag namespace named `OraclePCA` to enable additional functionality. For example, the File Storage Service supports defined tags to set file system quota or database record size. There is no protection mechanism in place to prevent users from using that same namespace for other purposes. However, the reserved tag namespace must not be used for any other tags than those defined by the system.

Workaround: Do not use the `OraclePCA` tag namespace to create and use your own defined tags. Create your tags in a different tag namespace.

Bug: 35976195

Version: 3.0.2

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

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

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
     l1api            (systemd:l1api):                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 pcamn03
```

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

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

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

Federated Login Failure after Appliance Upgrade

Identity federation allows users to log in to Private Cloud Appliance with their existing company user name and password. After an upgrade of the appliance software, the trust relationship between the identity provider and Private Cloud Appliance might be broken, causing all federated logins to fail. During the upgrade the Private Cloud Appliance X.509 external server certificate could be updated for internal service changes. In this case, the certificate on the identity provider side no longer matches.

Workaround: If the identity provider allows it, update its service provider certificate.

1. Retrieve the appliance SAML metadata XML file from `https://iaas.<domain>/saml/<TenancyId>` and save it to a local file.
2. Open the local file with a text editor and find the `<X509Certificate>` element.

```
<SPSSODescriptor>
    <KeyDescriptor use="signing">
        <KeyInfo>
            <X509Data>
                <X509Certificate>
                    <COPY CERTIFICATE CONTENT FROM HERE>
                </X509Certificate>
            </X509Data>
        </KeyInfo>
```

```
    </KeyDescriptor>
</KeyDescriptor>
```

3. Copy the certificate content and save it to a new `*.pem` file, structured as follows:

```
-----BEGIN CERTIFICATE-----
<PASTE CERTIFICATE CONTENT HERE>
-----END CERTIFICATE-----
```

4. Update the identity provider with this new service provider certificate for your Private Cloud Appliance.

If the identity provider offers no easy way to update the certificate, we recommend that you delete the service provider and reconfigure identity federation. For more information, refer to the section "Federating with Microsoft Active Directory" in the Oracle Private Cloud Appliance Administrator Guide.

Bug: 35688600

Version: 3.0.2

Ensure No Storage Buckets Are Present Before Deleting a Compartment or Tenancy

When a command is issued to delete a compartment or tenancy, the appliance software cannot reliably confirm that no object storage buckets exist, because it has no service account with access to all buckets present on the ZFS Storage Appliance. As a result, access to certain object storage buckets could be lost when their compartment is deleted.

Workaround: Before deleting a compartment or tenancy, verify that no object storage buckets are present in that compartment or tenancy.

Bug: 35811594

Version: 3.0.2

Syntax Issue when Generating Certificate Signing Request

When you plan to implement a new custom certificate, you need to provide a certificate signing request (CSR) to the Certificate Authority (CA) that will create your CA certificate. You generate the CSR from the Service CLI using the `generateCustomerCsr` command, typically with optional parameters to add details about your organization.

To include multiple Organization Units (OU) in the CSR, you enter them as a comma-separated list. For example: `generateCustomerCsr organization="My Company" organizationUnit=Division-1,Division-2`. However, the CLI always interprets the comma as a separator in this situation, and no escape character can be used to alter the behavior. This means you cannot enter a name containing a comma, as it will be interpreted as a set of two comma-separated strings.

Workaround: Do not enter strings containing commas when adding optional parameters to the `generateCustomerCsr` command. Use a comma only as a separator.

Bug: 35946278

Version: 3.0.2

Listing Upgrade Jobs Fails with RabbitMQ Error

When you run the Service CLI command `getUpgradeJobs`, the following error might be returned:

```
PCA-ADMIN> getUpgradeJobs
Status: Failure
Error Msg: PCA_GENERAL_000012: Error in RabbitMQ service: null
```

Workaround: The issue is temporary. Please retry the command at a later time.

Bug: 35999461

Version: 3.0.2

Availability Domain Name Change in Version 3.0.2-b1001356

In software version 3.0.2-b1001356 (December 2023), Private Cloud Appliance's single availability domain has been renamed from "`ad1`" to "`AD-1`". This change was required for compatibility with Oracle Cloud Infrastructure. The availability domain is a mandatory parameter in a small set of commands, and an optional parameter in several other commands.

The `--availability-domain` parameter is required with the following commands:

```
oci bv boot-volume create
oci bv boot-volume list
oci bv volume create
oci bv volume-group create
oci compute instance launch
oci compute boot-volume-attachment list
oci fs file-system create
oci fs file-system list
oci fs mount-target create
oci fs mount-target list
oci fs export-set list
oci iam fault-domain list
```

Workaround: Ensure that the correct value is used to identify the availability domain in your commands, depending on the version of the appliance software your system is running. If you are using scripts or any form of automation that includes the `--availability-domain` parameter, ensure that your code is updated when you upgrade or patch the appliance with version 3.0.2-b1001356 or newer.

Bug: 36094977

Version: 3.0.2

No Packages Available to Patch MySQL Cluster Database

With the release of appliance software version 3.0.2-b1001356, new MySQL RPM packages were added to the ULN channel *PCA 3.0.2 MN*. However, a package signing issue prevents the ULN mirror from downloading them, which means the MySQL cluster database on your system cannot be patched to the latest available version.

When patching the system, you will see no error message or abnormal behavior related to the missing MySQL packages. Follow the workaround to obtain the new packages. Once these have been downloaded to the ULN mirror, you can patch the MySQL cluster database.

> **Note:**
>
> For new ULN mirror installations, the steps to enable updates of MySQL packages have been included in the Oracle Private Cloud Appliance Patching Guide under "Configure Your Environment for Patching".

To determine if a system is affected by this issue, check the ULN mirror for the presence of MySQL packages in the yum directory referenced by the `pca302_x86_64_mn` soft link. If the search returns no results, the ULN mirror was unable to download the MySQL packages. The default location of the yum setup directory is `/var/www/html/yum`, which is used in the following example:

```
# ls -al /var/www/html/yum/pca302_x86_64_mn/getPackage/ | grep mysql
-rw-r--r--. 1 root root  85169400 Dec 19 03:19 mysql-cluster-commercial-
client-8.0.33-1.1.el7.x86_64.rpm
-rw-r--r--. 1 root root   4751220 Dec 19 03:19 mysql-cluster-commercial-client-
plugins-8.0.33-1.1.el7.x86_64.rpm
-rw-r--r--. 1 root root    689392 Dec 19 03:19 mysql-cluster-commercial-
common-8.0.33-1.1.el7.x86_64.rpm
-rw-r--r--. 1 root root  12417692 Dec 19 03:19 mysql-cluster-commercial-data-
node-8.0.33-1.1.el7.x86_64.rpm
-rw-r--r--. 1 root root   2229080 Dec 19 03:19 mysql-cluster-commercial-icu-data-
files-8.0.33-1.1.el7.x86_64.rpm
-rw-r--r--. 1 root root   2236184 Dec 19 03:19 mysql-cluster-commercial-
libs-8.0.33-1.1.el7.x86_64.rpm
-rw-r--r--. 1 root root   1279012 Dec 19 03:19 mysql-cluster-commercial-libs-
compat-8.0.33-1.1.el7.x86_64.rpm
-rw-r--r--. 1 root root   3478680 Dec 19 03:19 mysql-cluster-commercial-management-
server-8.0.33-1.1.el7.x86_64.rpm
-rw-r--r--. 1 root root 364433848 Dec 19 03:19 mysql-cluster-commercial-
server-8.0.33-1.1.el7.x86_64.rpm
-rw-r--r--. 1 root root   2428848 Dec 19 03:19 mysql-connector-j-
commercial-8.0.33-1.1.el7.noarch.rpm
-rw-r--r--. 1 root root   4570200 Dec 19 03:19 mysql-connector-odbc-
commercial-8.0.33-1.1.el7.x86_64.rpm
```

Workaround: When you import the appropriate GPG key on your ULN mirror, it can download the updated MySQL packages. Proceed as follows:

1. Log in to the ULN mirror server.
2. Download the MySQL GPG key from https://repo.mysql.com/RPM-GPG-KEY-mysql-2022.
3. Import the GPG key.

   ```
   # rpm --import RPM-GPG-KEY-mysql-2022
   ```

4. Update the ULN mirror.

   ```
   # /usr/bin/uln-yum-mirror
   ```

   If the key was imported successfully, the new MySQL packages are downloaded to the ULN mirror.

5. For confirmation, verify the signature using one of the new packages.

   ```
   # rpm --checksig mysql-cluster-commercial-management-server-8.0.33-1.1.el7.x86_64.rpm
   mysql-cluster-commercial-management-server-8.0.33-1.1.el7.x86_64.rpm: rsa sha1 (md5)
   pgp md5 OK
   ```

Bug: 36123758

Version: 3.0.2

Uppercase Letters Are Not Supported in Domain Names

Uppercase letters aren't supported in domain names. The domain name for your system is used as the base domain for the internal network, and by Oracle Private Cloud Appliance public facing services. This attribute has a maximum length of 190 characters. Acceptable characters are "a"→"z", "0"→"9", "-"

Bug: 36484125

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

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

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 ensure 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

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

Error Message when Exadata Network Is Deleted Successfully

When you delete an Exadata network using the Service Web UI, an error message might appear that says the network could not be deleted. Check the Exadata Network table for confirmation. The delete operation might succeed despite the error message.

Workaround: If the Exadata Networks table indicates that a network was deleted successfully, you can safely ignore the error message that indicates the opposite.

Bug: 35644436

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

When Modifying Listener of Load Balancer, Existing Configuration Parameters Are Ignored

When you use the Compute Web UI to modify a listener configuration of a load balancer, several parameters are not preserved in the Edit Listener window. Although you are changing the configuration of an existing listener, certain parameters are replaced with default values, forcing you to enter and verify all the parameters again.

As an example, let's assume your load balancer has a listener configured for HTTPS at port 8443, with an SSL certificate and cipher suite set up. When you edit this listener through the UI, the protocol appears as HTTP and the SSL configuration is not shown. Next, when you select the already configured HTTPS protocol again, the custom port, TLS version and cipher suite are reverted to their default values instead of the currently configured parameters.

Parameters that are not preserved in the Edit Listener window: protocol other than HTTP, custom port number, SSL settings (TLS version and cipher suite). Other parameters are not reverted: backend set, host names, path route set, idle timeout.

Workaround: If you modify a listener configuration through the Compute Web UI, always ensure that you have selected the right protocol first. Then set all listener parameters as if you were creating it from scratch, without moving back and forth between the policy options in the UI. Alternatively, modify the listener configuration using the OCI CLI.

Bug: 36032894

Version: 3.0.2

Create Backend to Add Security List/Configure Automatically Using IP Address Does Not Display Egress/Ingress Lists

If you use the following method in the Compute Web UI to create backend servers to a backend set of a load balancer, the egress and ingress security lists/rules are not displayed:

Select IP Address under the Backend Servers section, enter the backend server IP address, and then select Configure Automatically under the Security Rules section.

Workaround: Create the backend servers using one of these other methods:

- Select Computed Instances under the Backend Servers section when creating the backend, and select security rules Configure Automatically.
- Select IP Address under the Backend Servers section when creating the backend, and select security rules Configure Manually. Then, manually add the ingress and egress security list rules to the security lists attached to the corresponding subnet.

Bug: 35570701

Version: 3.0.2

Tag Area Above Resource Tables Does Not Expand

Resource pages in the UI display a table listing resources of a particular type. Just above the table you find control options to auto-reload the page, refresh the resource table, and filter the entries displayed. When filtering by tag, you can select multiple tags for your filter. However, the space available to show the filter tags does not expand, which causes elements of the UI to overlap as more tags are added.

Workaround: There is no workaround. Only the first few tags used for filtering are displayed while additional tags are hidden behind other UI elements.

Bug: 35975108

Version: 3.0.2

Service API Reference Displays Wrong Path for Disaster Recovery Commands

Developers can access the Service API Reference in a browser by appending *api-reference* to the base URL of the Service Web UI. For example: `https://adminconsole.myprivatecloud.example.com/`***`api-reference`***.

For commands related to disaster recovery configuration, the reference pages are split between the `ComputeInstance` and `DrConfig` nodes. This presentation is incorrect: all disaster recovery commands belong to the `DrConfig` object. As a result, the paths shown on the pages in the `ComputeInstance` section are wrong. Those commands will result in an HTTP 404 - not found error.

Workaround: While the request and response information is correct in the pages of the `ComputeInstance` section, you must replace that part of the object path with `DrConfig`. The following commands are affected:

Documented Path	Correct Path
`/admin/wsapi/rest/3.0.1/ComputeInstance/drAddComputeInstance`	`/admin/wsapi/rest/3.0.1/`***`DrConfig`***`/drAddComputeInstance`
`/admin/wsapi/rest/3.0.1/ComputeInstance/drGetComputeInstance`	`/admin/wsapi/rest/3.0.1/`***`DrConfig`***`/drGetComputeInstance`
`/admin/wsapi/rest/3.0.1/ComputeInstance/drGetComputeInstances`	`/admin/wsapi/rest/3.0.1/`***`DrConfig`***`/drGetComputeInstances`

Documented Path	Correct Path
`/admin/wsapi/rest/3.0.1/ComputeInstance/drRemoveComputeInstance`	`/admin/wsapi/rest/3.0.1/DrConfig/drRemoveComputeInstance`

Bug: 36267982

Version: 3.0.2

Details of an Instance Import Are Displayed as an Instance Export

When you import an instance from an object storage backup, it appears in the Instance Imports table of the Compute service. You click the instance import to display its detail page. However, the instance import detail page has the layout and data fields of an instance *export* detail page, including an Import button.

Workaround: Do not click the Import button; the instance has already been imported. The "export" text labels are wrong and can be ignored. The OCIDs and properties displayed on the page are correct for the instance import you selected.

Bug: 36310038

Version: 3.0.2

When Canceling an Instance Import the Operation Does Start

When you export an existing instance to object storage, it appears in the Instance Exports table of the Compute service. You can directly import an existing instance export using the Import button on the Instance Export detail page. A pop-up window appears where you either confirm or cancel the import operation. However, even if you click Cancel the import operations is started.

Workaround: Considering the import operation is started regardless of the option you select, either try to cancel the import in progress or delete it after it has finished.

Bug: 36311280

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.

Possible Impact to BGP Links After Upgrading to 3.0.2-b1081555 or Higher

When running BGP in a Mesh configuration, you may experience a situation where BGP links show an IDLE state or never connected, when upgrading to 3.0.2-b1010555 or later. If you are using BGP in a Mesh configuration and are currently running a release prior to 3.0.2-b1010555 then contact Oracle support, who can assist in updating and correcting your uplink configuration prior to upgrade. If an upgrade to 3.0.2-b1010555 or later has already been performed and you see BGP link state as IDLE then contact Oracle support who can also assist, post upgrade.

Bug: 36525352

Version: 3.0.2

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.

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,    All
                               2048-2050, 20048
<subnet CIDR>     TCP          All                     111, 389, 445, 4045,
                                                       2048-2050, 20048
<subnet CIDR>     UDP          111, 389, 445,          All
                               4045, 20048
<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,    All
                                           2048-2050, 20048
```

```
<instances subnet CIDR>       UDP          111, 389, 445,          All
                                           4045, 20048
```

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,    All
                                           2048-2050, 20048
<mount target subnet CIDR>    UDP          111, 289, 445, 2048,    All
                                           4045, 20048

+++ Egress Rules ++++++++++++++++++++

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

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

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

Network Load Balancer Does Not Report Detailed Backend Health Status

Users of Oracle Cloud Infrastructure might be familiar with the detailed health statuses it provides for backend servers of network load balancers. In case a backend server is not entirely healthy, the health check status provides an indication of the problem, for example: connection failure, time-out, regex mismatch, I/O error, invalid status code. Due to the specific load balancer implementation in Oracle Private Cloud Appliance, the Network Load Balancer service can only report whether a backend server is healthy (`OK`) or unhealthy (`CRITICAL`).

Workaround: There is no workaround. Backend health checks cannot provide extra status information.

Bug: 35993214

Version: 3.0.2

Route Table Stuck in Provisioning State Failure

When updating a route table that is associated as an attachment to a Dynamic Routing Gateway (DRG) to have a Local Peering Gateway (LPG) as a target, this known issue can leave the route table stuck in the provisioning state:

```
{
  "timestamp": "2023-06-28T15:30:58.635+0000",
  "rid":
"7FCCBAEBA62848878983FDA3098EE4DB/330fc100-b86f-4137-a1f9-2437a512b8e8/7b003c9
```

```
7-11c4-4e4a-8c9a-11861532db0d,
  "process": 1,
  "ocid": null,
  "levelname": "ERROR"
  "src_lineno": 481
  "src_pathname": "/usr/lib/python3.6/site-packages/pcanwctl/framework.py",
  "message::
  "Exception on function call: update_route_table, error: (404,
NotAuthorizedOrNotFound', 'No Subnet was found'), start exception rollback",
  "tag": "pca-nwctl.log"
}
```

Workaround: Delete and recreate the route table to avoid the error in the update routine.

Bug: 35547644

Version: 3.0.2

Updating Route Table Using Terraform Fails Because DRG Is Not Attached

When deploying network resources with Terraform, it may occur that a route table cannot be updated because an expected Dynamic Routing Gateway (DRG) attachment appears not to exist. Although the DRG is attached to the VCN, the operation is not fully completed when the command is issued to update the route table. The quick succession of commands through Terraform can reveal this timing issue, but it is highly unlikely to occur as a result of human user actions.

Workaround: Assuming the route table update failure is the result of a timing issue, repeating the route table update command is expected to succeed. Reapply the Terraform configuration or update the route table manually.

Bug: 36297777

Version: 3.0.2

Failure Executing Terraform Destroy Due to Route Table in Provisioning State

When you run a *terraform destroy* operation, it might fail because a route table object is still in 'provisioning' state instead of 'available'. This typically occurs when many updates are made to a route table in a short amount of time, resulting in commands taking longer to complete than expected by the Terraform provider. Strictly speaking, this is not a bug but rather a timing issue.

Workaround: Assuming the failure is the result of a timing issue, no route table or other resource is permanently stuck in 'provisioning' state. Repeating the terraform destroy command is expected to successfully remove the remaining objects. If necessary, increase the wait times for specific resources in your Terraform settings.

Bug: 36352218

Version: 3.0.2

When Configuring BGP Authentication the Password Is a Required Parameter

When the appliance uplinks to the data center network are configured for dynamic routing, two Autonomous Systems – meaning the spine switches on the appliance side, and the ToR switches on the data center side – are set up as BGP (Border Gateway Protocol) peers. The sessions between the BGP peers can be protected with password-based authentication. BGP authentication can be enabled for the data network as well as the optional separate administration network.

You can set the BGP password using the `setDay0DynamicRoutingParameters` command in the Service CLI. Two command parameters must be provided for each network.

- data network: `bgpAuthentication=True` and `bgpPassword=`***`<mypassword>`***
- administration network: `adminBgpAuthentication=True` and `adminBgpPassword=`***`<mypassword>`***

However, the CLI command is accepted if you set BGP authentication to "true" without providing a password. This has no adverse effects, but BGP authentication remains disabled.

Workaround: When you enable BGP authentication on the data network, and the administration network if present, make sure you also specify the BGP password as part of the command parameters.

Bug: 35737959

Version: 3.0.2

Uplink VRRP Mesh Configuration Sets Second Switch IP Incorrectly

When you try to configure the appliance data and administration network uplinks in mesh topology with VRRP (Virtual Router Redundancy Protocol), the command results in a CLI error. The problem occurs when the spine switches' second IP address is configured: the switch interprets the parameters as overlapping network settings and rejects them.

The following example shows an administration network with a 4-port mesh uplink topology. The same behavior applies to data network uplinks.

```
PCA-ADMIN> edit networkConfig enableAdminNetwork=True adminportcount=4
admintopology=MESH adminportspeed=10
adminspine1Ip=10.1.1.97,10.1.1.98 adminspine2Ip=10.1.1.101,10.1.1.102
adminSpineVip=10.1.1.105 [...]

PCA-ADMIN> show networkconfig
Data:
  [...]
Error:
UpdateFirstBootHandler: {'http_status_code': 500, 'code': 'InternalServerError',
'message': 'SwitchCliError on 100.96.2.20:
overlapping network for ipv4 address: 10.1.1.98/28 on po46, 10.1.1.97/28 already
configured on po45\\n for cmd [...]
```

Workaround: There is no workaround. This specific uplink configuration cannot be applied at this time.

Bug: 36063880

Version: 3.0.2

Real Application Cluster (RAC) Environment Loses Access to Oracle Exadata Database Instances During Appliance Upgrade

Oracle Real Application Cluster (RAC) environments are typically deployed with application servers running as compute instances on the Private Cloud Appliance, and database clusters on a directly connected Oracle Exadata system. When the appliance software is upgraded or patched, the RAC listener services can experience failures, making the database clusters inaccessible on the Exadata network.

If the RAC listener service is down, the Exadata database node is inaccessible for applications making new connections. However, existing sessions are not interrupted as long as another database node remains online.

Workaround: In most cases the Oracle Clusterware Agent recovers the listener service automatically after the outage caused by the appliance software upgrade. If not, a manual restart of the listener service is required to restore database access on a node. Use the `lsnrctl` utility from the Grid home as the Grid user.

Bug: 36446341

Version: 3.0.2

Compute Service Issues

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

Possible VM Impact When OS Images are Deleted

Under certain operating conditions, if an OS image is deleted while boot volumes and VM instances based on this image are still present in the system, there can be an impact to all VMs' boot devices based on the deleted image. The symptoms can include but are not limited to the following:

- Input/output error messages in a running VM's logs or console and possible VM application failures
- VMs fail to boot with a `no bootable device` message on the VMs' console
- Re-attaching a boot volume to a stopped VM might fail

Workaround: To avoid this situation, do not delete an OS image unless all the VM instances, all the boot volumes, their backups and their clones originating from this image have been properly terminated first.

Bug: 36489907

Version: 3.0.2

E5.Flex Instance Shape Is Not Supported on the X9-2 Hardware Platform

Compute instance shapes are tied to the architecture of the underlying compute nodes. The VM.PCAStandard.E5.Flex shape was added specifically to create instances on Oracle Server X10 compute nodes. It is the only shape supported on the X10 rack configuration. On a Private Cloud Appliance X9-2, all other shapes – including flex shapes – are supported.

Workaround: Select a suitable shape for your Private Cloud Appliance compute node architecture. If the compute nodes in your appliance are Oracle Server X10, always select the VM.PCAStandard.E5.Flex shape. Systems with Oracle Server X9-2 compute nodes support all shapes except VM.PCAStandard.E5.Flex. If you need a flexible shape, select the VM.PCAStandard1.Flex shape instead.

Bug: 35549831

Version: 3.0.2

Displaced Instances Not Returned to Their Selected Fault Domains

A displaced instance is an instance that is running in a fault domain that is not the fault domain that is specified in the configuration for that instance. An instance can become displaced during compute node evacuation or failure.

When Auto Recovery is enabled, a displaced instance is automatically returned to the fault domain that is specified in its configuration when resources become available in that fault domain. Auto Recovery is enabled by default.

Workaround:

If your Private Cloud Appliance is running Software Version 3.0.2-b852928 or Software Version 3.0.2-b892153, or if you upgrade to either of these releases, disable Auto Recovery from the Service CLI:

```
PCA-ADMIN> disableAutoResolveDisplacedInstance
```

If your Private Cloud Appliance is running a release that is newer than Software Version 3.0.2-b892153, you can enable Auto Recovery.

See "Migrating Instances from a Compute Node" and "Configuring the Compute Service for High Availability" in the Hardware Administration chapter of the *Oracle Private Cloud Appliance Administrator Guide* for more information about these commands.

If your Private Cloud Appliance is affected by this bug and an instance is displaced, stop and restart the instance to return the instance to its selected fault domain. See "Stopping, Starting, and Resetting an Instance" in the Compute Instance Deployment chapter of the *Oracle Private Cloud Appliance User Guide*.

Bug: 35601960, 35703270

Version: 3.0.2

Terraform Cannot Be Used for Instance Update

Starting with the May 2023 release of the Oracle Private Cloud Appliance software, the Oracle Cloud Infrastructure Terraform provider cannot be used to update an instance on Oracle Private Cloud Appliance. Only the instance update operation is affected by this issue.

Instance update fails when done using Terraform because the `is_live_migration_preferred` property does not exist for Terraform. Because the property is unknown, when the property is seen, Terraform treats the property value as `false`, which is not a supported value.

Workaround: Use the Compute Web UI or the OCI CLI to perform instance update.

Bug: 35421618

Version: 3.0.2

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

TypeError Returned when Attaching an Instance to an Instance Pool

When you attach an existing compute instance to an instance pool, you can include parameters with the OCI CLI command so it reports when the instance reaches the intended ("active") lifecycle state. However, a bug in the OCI CLI could lead to the following error:

```
# oci compute-management instance-pool-instance attach \
--instance-id ocid1.instance....unique_ID --instance-pool-id
ocid1.instancePool....unique_ID \
--wait-for-state ACTIVE --wait-interval-seconds 120 --max-wait-seconds 1200
Action completed. Waiting until the resource has entered state: ('ACTIVE',)
Encountered error while waiting for resource to enter the specified state. Outputting
last known resource state
{
  "data": {
    "availability-domain": "AD-1",
    "compartment-id": "ocid1.tenancy....unique_ID",
    "display-name": "Standard1.4",
    "fault-domain": "FAULT-DOMAIN-3",
    "id": "ocid1.instance....unique_ID",
    "instance-configuration-id": null,
    "instance-pool-id": "ocid1.instancePool....unique_ID",
    "lifecycle-state": "ATTACHING",
    "load-balancer-backends": [],
    "region": "mypca.mydomain.com",
    "shape": "VM.PCAStandard1.Flex",
    "state": "RUNNING",
    "time-created": "2023-10-28T03:22:45+00:00"
  },
  "opc-work-request-id": "ocid1.workrequest....unique_ID"
```

```
}
TypeError: get_instance_pool_instance() missing 1 required positional argument:
'instance_id'
```

Workaround: The command option `--wait-for-state` is unreliable at this time. As an alternative you can use the command `list-instance-pool-instances` to check the state of the instances in the pool.

Bug: 35956140

Version: 3.0.2

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 issues 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

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 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

OCI CLI Commands Fail When Run From a Compute Instance

Compute instances based on Oracle Linux images provided since early 2023 are likely to have a firewall configuration that prevents the OCI CLI from connecting to the Private Cloud Appliance identity service. In Oracle Cloud Infrastructure the identity service must now be accessed through a public IP address (or FQDN), while Oracle Private Cloud Appliance provides access through an internal IP address. The Oracle Cloud Infrastructure images are configured by default to block all connections to this internal IP address.

The issue has been observed with these images:

- uln-pca-oracle-linux-7-9-2023-08-31-0-oci
- uln-pca-oracle-linux-8-2023-08-31-0-oci
- all Oracle Linux 9 images with a 2023 availability date

Workaround: If you intend to use the OCI CLI from a compute instance in your Private Cloud Appliance environment, verify its access to the identity service. If connections are refused, check the instance firewall configuration and enable access to the identity service.

1. Test the instance connection to the identity service. For example, use telnet or netcat.

```
# curl -v telnet://identity.mydomain.us.oracle.com:443
* connect to 169.254.169.254 port 443 failed: Connection refused

-- OR --
# nc -vz identity.mydomain.us.oracle.com 443
Ncat: Connection refused.
```

2. Confirm that the firewall output chain contains a rule named *BareMetalInstanceServices*.

```
# iptables -L OUTPUT --line-numbers
Chain OUTPUT (policy ACCEPT)
num  target                     prot   opt   source           destination
1    BareMetalInstanceServices  all    --    anywhere         169.254.0.0/16
```

3. Disable the bare metal instance rules in the firewall configuration.
   a. Rename the file that defines these firewall rules (`/etc/firewalld/direct.xml`).
   b. Restart the firewalld service.

   Detailed instructions are provided in the note with Doc ID 2983004.1.

Bug: 35234468

Version: 3.0.2

Cannot Install OCI CLI on Oracle Linux 9 Instance

To run the OCI CLI on an Oracle Linux 9 compute instance, the package `python39-oci-cli` and its dependencies are required. These are provided through the *Oracle Linux 9 OCI Included Packages* (`ol9_oci_included`) repository, but this repository cannot be accessed outside Oracle Cloud Infrastructure.

An Oracle Linux 9 compute instance on Oracle Private Cloud Appliance must instead retrieve the required packages from the public Oracle Linux 9 repositories – specifically: *Oracle Linux 9 Development Packages* (`ol9_developer`) and *Oracle Linux 9 Application Stream Packages* (`ol9_appstream`). These repositories are not enabled by default in the provided Oracle Linux 9 image.

Workaround: Enable the `ol9_developer` and `ol9_appstream` public yum repositories to install `python39-oci-cli`.

```
$ sudo yum --disablerepo="*" --enablerepo="ol9_developer ol9_appstream" install python39-
oci-cli -y
Dependencies resolved.
================================================================================================
=========================
 Package                                Architecture       Version
Repository              Size
================================================================================================
=========================
Installing:
 python39-oci-cli                       noarch             3.40.2-1.el9
ol9_developer           39 M
Upgrading:
 python39-oci-sdk                       x86_64             2.126.2-1.el9
ol9_developer           74 M
Installing dependencies:
 python3-arrow                          noarch             1.1.0-2.el9
ol9_developer          153 k
 python3-importlib-metadata             noarch             4.12.0-2.el9
ol9_developer           75 k
 python3-jmespath                       noarch             0.10.0-4.el9
ol9_developer           78 k
 python3-prompt-toolkit                 noarch             3.0.38-4.el9
ol9_appstream          1.0 M
 python3-terminaltables                 noarch             3.1.10-8.0.1.el9
ol9_developer           60 k
 python3-wcwidth                        noarch             0.2.5-8.el9
ol9_appstream           65 k
 python3-zipp                           noarch             0.5.1-1.el9
ol9_developer           24 k

Transaction Summary
================================================================================================
========================
```

```
Install  8 Packages
Upgrade  1 Package
[...]
Complete!
```

Bug: 35855058

Version: 3.0.2

Instance Launch Fails at 80 Percent Complete with Libvirt Error

When an instance is launched, Libvirt processes a set of instructions to bring the requested virtual machine to a running state. During this process, the connection with the agent sending the requests might be interrupted, which means the responses are not received. As a result, the instance is terminated. The Compute Service logs an internal server error similar to this example:

```
INFO (errors:66) Libvirt Error: code: 38, domain: 7, message: Cannot recv data: Input/
output error
```

Workaround: This is an intermittent issue. If an instance fails to launch, retry the operation.

Bug: 36100146

Version: 3.0.2

Instance Principal Unavailable Until Next Certificate Renewal Check

An instance principal is a compute instance that is authorized to perform actions on service resources. Before allowing these operations, the Identity and Access Management Service (IAM) validates the instance principal security token: a TLS certificate that expires after 30 days.

The system checks for expired certificates every 24 hours and renews them if necessary. However, an instance principal might lose its authorization after an outage, system maintenance, or upgrade activity. In that case, it cannot obtain an updated certificate until the next renewal check, which could be up to 24 hours later.

Similarly, after upgrading from a release that does not support instance principals to a release that does support instance principals, compute instances might have to wait up to 24 hours to receive their TLS certificates.

Workaround: If you need to have this certificate installed or renewed immediately, contact Oracle for assistance.

Bug: 36165739

Version: 3.0.2

Unable to Delete Tag Due to Instance Principal Error

When cleaning up a set of resources including compute instances with defined tags, it might occur that the compute service is unable to remove the tag from an instance. As a result, the tag key definition cannot be deleted: the associated work request fails, and typically returns the error "`Error message from compute: create instance principal`". It indicates that an instance principal certificate was regenerated for an instance that belongs to a compartment that no longer exists.

This situation can occur if the tag key definition belongs to a different compartment than the tagged instance, and the cleanup operations are performed in this order: first the tagged instance, then the compartment it belongs to, and then the tag key definition. When attempting to delete the tag key definition, the tagged instances have already been terminated, yet the instance principal certificate can be regenerated during tag deletion. At this point the error is logged.

Workaround: Terminated instances are purged from the database after 24 hours. Once they have been purged, the tag key definition can be deleted. Alternatively, the terminated instances can be removed manually from the database. Contact Oracle for assistance.

Bug: 36348781

Version: 3.0.2

When Instance Is Shut Down from OS, Soft Stop Results in Conflict

To avoid data corruption in applications that take a long time to stop, the recommendation is to shut down a compute instance from within its operating system, before issuing the soft stop command for a graceful shutdown from the Compute Enclave. However, due to a problem with the instance action logic, when an instance has been shut down from the OS, the soft stop command returns a conflict (error 409) because the instance is no longer in a running state.

Workaround: To shut down a compute instance with lowest possible risk of data corruption, use the OS shutdown command first. Then, issue either the *force stop* command from the Compute Web UI or the instance action `STOP` from the OCI CLI.

Bug: 36299430

Version: 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.

Updating Terraform Changes File Storage Export Path

When you use Terraform to create a file system export, you must specify `AUTOSELECT` for the value of `path` in the `oci_file_storage_export` definition.

You must also include the `lifecycle` stanza to ignore any updates to the path. If you do not ignore updates to the path, the path is automatically deleted and re-created when you update the Terraform, even if you do not explicitly update this path. Updating this path can interrupt clients that have an active mount via the export.

Workaround: Set the path and include the `lifecycle` stanza as shown in the following example:

```
resource "oci_file_storage_export" "pcauserExport" {
  export_set_id  = local.Okit_MT_1702774958525ExportSet_id
  file_system_id = local.Okit_FS_1702774481898_id
  path           = "AUTOSELECT"
  lifecycle {
    ignore_changes = [
      path,
    ]
```

```
    }
}
```

Bug: 36116003

Version: 3.0.2

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: 35256335

Version: 3.0.2

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

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

Block Volume Performance Level Not Preserved During Cloning

The block volumes provisioned on the ZFS Storage Appliance are located in either the standard or high-performance pool. The performance level is reflected in the properties of each block volume as *volume performance units* (VPU) per GB. However, when cloning a volume group or volume group backup, the performance level of all new block volumes produced by the clone operation is set to 0. CLI output will show the parameter `"vpus-per-gb": 0` in the properties of the block volume clone.

Workaround: There is no workaround available. The block volume clones are placed in the correct storage pool, meaning their performance level is as intended.

Bug: 35333587

Version: 3.0.2

Internal Backups for Instance Cloning Not Displayed

When you clone a compute instance, an internal backup of the boot and block volumes is created. In appliance software versions up to 3.0.2-b852928 those internal backups are visible to users. While not recommended, the backups could technically be used to create additional instances. Existing internal backups are not deleted during appliance upgrade or patching. However, in newer software versions the internal backups are no longer exposed.

Workaround: Do not create clones or new compute instances from the existing internal volume (group) backups. To remove old backups of storage volumes, ensure that all other backups and clones of the original source volume are terminated first.

Bug: 35406033

Version: 3.0.2

Limit for Volume Backups Not Enforced

The "Service Limits" chapter in the Oracle Private Cloud Appliance Release Notes specifies a limit of 100 volume backups per tenancy for a system with default storage capacity. This limit is not enforced: you can continue to create volume backups beyond the documented maximum.

Workaround: In theory, the maximum number of volume backups is limited by available storage on the ZFS Storage Appliance. The system is expected to handle thousands of volume backups across all tenancies. However, we recommend that an administrator monitors storage space consumption proactively if users create many volume backups.

Bug: 35509673

Version: 3.0.2

NFS Service Interruption During ZFS Storage Appliance Firmware Upgrade or Patching

When the firmware of the appliance's ZFS Storage Appliance is upgraded or patched, compute instances could encounter an interruption of NFS connectivity. The service outage occurs when failover/failback is performed between the storage appliance controllers, and it could take over 2 minutes to reestablish the NFS service. There could be multiple factors contributing to the delay: the NFS server's 90 second grace period to allow NFSv4 clients to recover locking state after an outage, the NFS protocol attempting to reconnect to the same TCP port, and the NFS client's kernel version.

Workaround: To reduce the outage time of NFS connectivity, it is recommended to use the mount options described in the note with Doc ID 359515.1. While the document describes optimizations for Oracle RAC and Oracle Clusterware, the mount options also improve NFS performance and stability in a Private Cloud Appliance environment.

Bug: 36348165

Version: 3.0.2

Container Engine Issues

This section describes known issues and workarounds related to the Oracle Container Engine for Kubernetes.

Container Engine for Kubernetes Requires Switch Firmware Upgrade on Systems with Administration Network

If your Private Cloud Appliance is configured with a separate administration network, the appliance and data center networking need reconfiguration to enable the traffic flows required by the Oracle Container Engine for Kubernetes (OKE). In addition, the reconfiguration of the network is dependent on functionality included in a new version of the switch software.

Workaround: Upgrade or patch the software of the switches in your appliance. Reconfigure the network. You can find details and instructions in the following documentation sections:

- "Upgrading the Switch Software" in the Oracle Private Cloud Appliance Upgrade Guide
- "Patching the Switch Software" in the Oracle Private Cloud Appliance Patching Guide
- "Securing the Network" in the Oracle Private Cloud Appliance Security Guide

  This section includes a port matrix for systems with a separate administration network. Use it to configure routing and firewall rules, so the required traffic is enabled in a secure way.

Bug: 36073167

Version: 3.0.2

Tag Filters Not Available for Kubernetes Node Pools and Nodes

Unlike Oracle Cloud Infrastructure, Private Cloud Appliance currently does not provide the functionality to use Tag Filters for tables listing Kubernetes node pools and nodes. Tag filtering is available for Kubernetes clusters.

Workaround: There is no workaround. The UI does not provide the tag filters in question.

Bug: 36091835

Version: 3.0.2

Kubernetes Node Tags Not Available in Compute Web UI

The Compute Web UI does not allow users to apply defined or freeform tags to all nodes in a Kubernetes node pool. However, tags can be applied to one node at a time from the UI. Tagging all nodes in a node pool at once must be done using the OCI CLI.

Workaround: To apply tags to all nodes in a node pool, use the OCI CLI command options `--node-defined-tags` and `--node-freeform-tags`.

Bug: 36156349

Version: 3.0.2

Node Doctor Script Not Available in Worker Nodes

In Oracle Cloud Infrastructure, the Oracle Container Engine for Kubernetes (OKE) provides a troubleshooting utility called *Node Doctor*. Its purpose is to help resolve problems with worker nodes that are not in *Active* state. You will see references to worker node troubleshooting and the Node Doctor script in the Compute Web UI. However, the functionality is not available in Private Cloud Appliance. Even if you install the script on your worker nodes, its operations will fail because of missing or unexpected environment data.

Workaround: There is no workaround. The Node Doctor script is not available on worker nodes in Private Cloud Appliance.

Bug: 35807245

Version: 3.0.2

Unable to Delete Kubernetes Cluster in Failed State

To deploy a Kubernetes cluster, the Oracle Container Engine for Kubernetes (OKE) uses various types of cloud resources that can also be managed through other infrastructure services, such as compute instances and load balancers. However, Kubernetes cluster resources must be manipulated only through the OKE Service, to avoid inconsistencies. If the network load balancer of a Kubernetes cluster is deleted outside the control of the OKE Service, that cluster ends up in a failed state and you will no longer be able to delete it.

Workaround: This is a known issue with the Cluster API Provider. If a cluster is in failed state and its network load balancer is no longer present, it must be cleaned up manually. Contact Oracle for assistance.

Bug: 36193835

Version: 3.0.2

Kubernetes Cluster Creation Failure Due to Load Balancer Limit Returning Unclear Error

When the maximum number of load balancers has already been deployed in your tenancy or appliance environment, a new Kubernetes cluster cannot be created. The cluster creation attempt results in a failure, but the error message returned does not state that the limit was reached. A more generic *cluster reconciliation error* message is returned instead.

```
# oci ce work-request-error list --compartment-id ocid1.tenancy...unique_id
--work-request-id ocid1.workrequest...unique_id
  "data": [
    {
      "code": "GetWorkRequestGeneric",
      "message": "OCICluster reconciliation failed: ReconcileError",
      "timestamp": "2024-02-24T17:24:48.615203+00:00"
    }
  ]
```

Workaround: To confirm that the failure is caused by the load balancer limit, check the number of load balancers deployed in the appliance environment. If you have insufficient access rights, ask an administrator who has the necessary privileges.

Bug: 36335225

Version: 3.0.2

Intermittent Failures when Using Terraform to Create Kubernetes Cluster

When creating Kubernetes clusters from Terraform, there might be intermittent failures, which return a generic error message : "`Failed to create cluster due to an Unknown error`". These failures are known to occur when the Load Balancing service pods are not in sync. Testing shows that this issue is most likely to affect the first cluster creation attempt in the appliance environment.

Workaround: When running into this type of failure, users should delete the Kubernetes cluster that failed to deploy, and retry the operation to create the cluster. Particularly when the first cluster creation on the system fails, subsequent create operations tend to be successful.

Bug: 36379853

Version: 3.0.2

API Reference on Appliance Not Up-to-Date for OKE Service

Every Private Cloud Appliance provides online API reference pages, conveniently accessible from your browser. For the Compute Enclave, these pages are located at https://console.*mypca.mycompany.com*/api-reference. Unfortunately, the API reference in the appliance software version providing the initial release of the Oracle Container Engine for Kubernetes (OKE) has not been updated with the OKE Service APIs.

Workaround: Contact Oracle for assistance and open a service request.

Bug: 35710716

Version: 3.0.2

Serviceability Issues

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

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

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

MN01 Host Upgrade Fails When it is the Last Management Node to Upgrade

Upgrades and patches to the management nodes are performed in a sequential order. When `MN01` falls last in that order, the management node upgrade or patch operation fails. To avoid this issue, ensure that the Management Node Virtual IP address is assigned to `MN02` before you start any management node upgrade or patching operations.

Workaround: Assign the Management Note Virtual IP address to `MN02` before you upgrade or patch.

```
# pcs resource move mgmt-rg pcamn02
```

Bug: 35554754

Version: 3.0.2

Failure Draining Node when Patching or Upgrading the Kubernetes Cluster

To avoid that microservice pods go into an inappropriate state, each Kubernetes node is drained before being upgraded to the next available version. The Upgrader allows all pods to be evicted gracefully before proceeding with the node. However, if a pod is stuck or is not evicted in time, the upgrade or patch process stops.

Workaround: If a Kubernetes node cannot be drained because a pod is not evicted, you must manually evict the pod that causes the failure.

1. Log on to the Kubernetes node using ssh, and run the following command, using the appropriate host name:

   ```
   # kubectl drain pcamn00 --ignore-daemonsets --delete-local-data
   ```

   Wait for the draining to complete. The command output should indicate: `node/pcamn00 drained`.

2. If the drain command fails, the output indicates which pod is causing the failure. Either run the drain command again and add the `--force` option, or use the delete command.

   ```
   # kubectl delete pod pod-name --force
   ```

3. Rerun the Kubernetes upgrade or patch command. The Upgrader continues from where the process was interrupted.

Bug: 35677796

Version: 3.0.2

Oracle Auto Service Request Disabled after Upgrade

When a Private Cloud Appliance has been registered for Oracle Auto Service Request (ASR), and the service is enabled on the appliance, the ASR service may become disabled after an upgrade of the appliance software. The issue has been observed when upgrading to version 3.0.2-b925538.

Workaround: After the appliance software upgrade, verify the ASR configuration. If the ASR service is disabled, manually enable it again. See "Using Auto Service Requests" in the Status and Health Monitoring chapter of the Oracle Private Cloud Appliance Administrator Guide.

Bug: 35704133

Version: 3.0.2